

The Value Proposition of Distributed Satellite Systems for Space Science Missions

by

Benjamin Andrew Corbin

B.S., Aerospace Engineering, University of Central Florida, 2008
S.M., Aeronautics and Astronautics, Massachusetts Institute of Technology, 2011
S.M., Earth and Planetary Sciences, Massachusetts Institute of Technology, 2011

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy

at the Massachusetts Institute of Technology
September 2015

© 2015 Benjamin Andrew Corbin. All Rights Reserved

The author hereby grants to MIT permission to reproduce and to distribute publicly paper and electronic copies of this thesis document in whole or in part in any medium now known or hereafter created.

Signature of Author.....

Benjamin Andrew Corbin
Department of Aeronautics and Astronautics
May 1st, 2015

Certified by.....

Prof. Jeffrey Hoffman
Professor of the Practice of Aerospace Engineering
Thesis Committee Chair

Certified by.....

Prof. Sara Seager
Professor, Earth and Planetary Sciences
Thesis Committee Member

Certified by.....

Dr. Adam Ross
Research Scientist, Engineering Systems; Lead Research Scientist, SEArI
Thesis Committee Member

Accepted by.....

Prof. Paulo C. Lozano
Associate Professor, Aeronautics and Astronautics
Chair, Graduate Program Committee

The Value Proposition of Distributed Satellite Systems for Space Science Missions

by
Benjamin Andrew Corbin

Submitted to the Department of Aeronautics and Astronautics on August 20th, 2015 in Partial Fulfillment of the Requirements for the Degree of Doctor of Philosophy

Abstract

The resources available for planetary science missions are finite and subject to some uncertainty. Despite decreasing costs of spacecraft components and launch services, the cost of space science missions is increasing, causing some missions to be canceled or delayed, and fewer science groups have the opportunity to achieve their goals due to budget limits. New methods in systems engineering have been developed to evaluate flexible systems and their sustained lifecycle value, but these methods are not yet employed by space agencies in the early stages of a mission's design. Previous studies of distributed satellite systems (DSS) showed that they are rarely competitive with monolithic systems; however, comparatively little research has focused on how DSS can be used to achieve new, fundamental space science goals that simply cannot be achieved with monolithic systems.

The Responsive Systems Comparison (RSC) method combines Multi-Attribute Tradespace Exploration with Epoch-Era Analysis to examine benefits, costs, and flexible options in complex systems over the mission lifecycle. Modifications to the RSC method as it exists in previously published literature were made in order to more accurately characterize how value is derived from space science missions. A tiered structure in multi-attribute utility theory allows attributes of complex systems to be mentally compartmentalized by stakeholders and more explicitly shows synergy between complementary science goals. New metrics help rank designs by the value derived over their entire mission lifecycle and show more accurate cumulative value distributions.

A complete list of the emergent capabilities of DSS was defined through the examination of the potential benefits of DSS as well as other science campaigns that leverage multiple assets to achieve their scientific goals. Three distinct categories consisting of seven total unique capabilities related to scientific data sampling and collection were identified and defined. The three broad categories are fundamentally unique, analytically unique, and operationally unique capabilities. This work uses RSC to examine four case studies of DSS missions that achieve new space science goals by leveraging these emergent capabilities. ExoplanetSat leverages *shared sampling* to conduct observations of necessary frequency and length to detect transiting exoplanets. HOBOCOP leverages *simultaneous sampling* and *stacked sampling* to study the Sun in far greater detail than any previous mission. ÆGIR leverages *census sampling* and *self-sampling* to catalog asteroids for future ISRU and mining operations. GANGMIR leverages *staged sampling* with *sacrifice sampling* and *stacked sampling* to answer fundamental questions related to the future human exploration of Mars.

In all four case studies, RSC showed how scientific value was gained that would be impossible or unsatisfactory with monolithic systems. Information gained in these studies helped stakeholders more accurately understand the risks and opportunities that arise as a result of the added flexibility in these missions. The wide scope of these case studies demonstrates how RSC can be applied to any science mission, especially one with goals that are more easily achieved with (or impossible to achieve without) DSS. Each study serves as a blueprint for how to conduct a Pre-Phase A study using these methods.

Thesis Supervisor: Jeffrey Hoffman

Title: Professor of the Practice of Aerospace Engineering

Acknowledgements

No research is conducted in a vacuum. Every researcher stands on the shoulders of giants who paved the way for them to push the boundary of the unknown, but we climb up onto those shoulders with the help of many on the long journey to the top of those giants. This dissertation could not have been completed without the help of so many people who have been a part of my life since the beginning and those who have given me encouragement, wisdom, or help along the way.

First and foremost, I would like to thank my parents, Andy and Therese, for encouraging me and guiding me along the path I chose of my own volition. Without your solid parenting and countless hours of dedication, whether changing my diapers, cooking homemade meals, driving me to various extracurricular activities, or allowing me to go on dangerous adventures, I would not be the man I am today. I would like to thank my older brother, Brian, for growing up with me, serving as an example of what to do in life (and sometimes what *not* to do), and for all the years of encouragement and wisdom when I had no one else to talk to about unusual problems. I would like to thank my younger sister, Claire, for reminding me that no matter how small you are you can still beat up things bigger than you, and for having the courage to blaze your own trail despite contrary advice from people who aren't in your shoes.

I would like to thank Dr. Jeffrey Hoffman for many years of guidance and encouragement, and for giving me a chance to earn a higher education at the best institute in the world. Your career serves as a model of what I hope mine will be, and working with you has reminded me that it is possible to achieve my dreams because someone I know already has. That reminder has been necessary during difficult periods throughout my graduate education.

I would like to thank Dr. Sara Seager for taking a gamble on me when I was a student down on his luck with little sense of direction, and supporting me as a mentor and a friend throughout this long process. Although the path towards my degree was never a straight one, the opportunities to learn from a diverse portfolio of projects has strengthened me as a student and as a researcher.

I would like to thank Dr. Adam Ross for his contributions to the field of systems engineering upon which this work is heavily based. Without your technical guidance and suggestions, I would have never come up with a feasible thesis topic nor would I have the understanding to finish it. Yours are the shoulders upon which I am adding another block of knowledge, so another may stand on top of it.

Other professors have gone out of their way to help me with many aspects of my work at MIT. I particularly would like to thank Paulo Lozano for his guidance and support of the MIT Rocket Team and my modeling of electric propulsion systems in space missions; David Miller for his classes in systems engineering, and all the help he has provided not just me but countless other students in the Space Systems Lab over the years; and Rick Binzel for giving me a crash course in asteroids and planetary defense that propelled my ideas into the international spotlight. They have helped me on projects and research and took time out of their schedule despite no additional contributions to their research from me.

I have also worked with other scientists and engineering instructors on many facets of this research. I would like to thank Dr. Mahboubeh Asgari-Targhi, Dr. Sabrina Corpino, Dr. Alessandro Golkar, and Dr.

Alessandra Babuscia for their contributions to my knowledge in different fields outside of my expertise. I would also like to thank Sabrina and Alessandro for allowing me to travel and work with them, and as a result I have had to pleasure of visiting five additional countries throughout the course of this research.

I would not have begun working for Sara if it were not for my previous positions as a research assistant as a master's student and an undergraduate. I would like to thank Dr. John Clarke, Dr. Carol Carveth, Dr. Chris Sweeney, and Dr. Tim Cooke at Boston University for having the inexhaustible patience to work with me for four years. I would like to thank Dr. Eric Petersen and Dr. Yanga Fernandez for giving me a chance to do research as an undergraduate at the University of Central Florida in the fields of gas dynamics and solar system astronomy.

I especially would like to thank many fellow graduate students for their help along the way, especially Niraj Inamdar for his help on numerous projects including our trip to Planetary Resources and difficult MATLAB programming challenges for ExoplanetSat; Mary Knapp for her work with the MATLAB/STK interface; Gwendolyn Gettliffe for her help with support vector machines; Kevin Stout for his help with thermal modeling of spacecraft; Fabio Nichele for his help on developing science value propositions and modeling ground penetrators; Lorenzo Feruglio for his help developing a Mars EDL simulation; Ignasi Lluch for his help with satellite communications and his research in federated satellites; Matt Fitzgerald, Capt Paul La Tour, Michael Curry, Michael Schaffner, and other graduate and undergraduate students in the SEArI laboratory for their help and guidance with systems engineering methodologies, small spacecraft tradespace analysis, and programs necessary for communicating the results of the case studies; Dr. Renyu Hu for his inputs on Mars atmospheric science goals; and Dr. Farah Alibay for helping me with my proposal.

I would like to thank several members of the support staff in machine shops that have guided me in my design work and taken the time out of their schedules to teach me how to design and build parts correctly when my classroom experienced failed me, particularly Abdulbaset Benwali at UCF, Todd Billings at MIT, and Heitor Mourato and Bob Fazio at Boston University.

Other notable instructors over the years who have positively influenced me and helped me get to this point, from kindergarten through university, include Mrs. Fran, Mr. Terry Webster, Mr. Wendell Griffith, Mrs. De Cook, Dr. Lisa Hagen, Mrs. Charla Cotton, Mrs. Mary Jane Robertson, Dr. John Calhoun, and Dr. Daniel Britt. These instructors in particular have had a profound impact on my life, my behavior, and how I engage the universe I inhabit and the people I meet. They have sparked my intellectual curiosity and I intend to repay them by keeping the flame alive for as long as I live and sparking the flame in others along the way.

I would also like to thank many student organizations for allowing me to do more with my time as a student, including AIAA, SEDS, SPS, and both the MIT and UCF Rocket Teams. I particularly want to thank Ryan McLinko, Joshua Nelson, Todd Romberger, Alex LePage, Matthew Stephens, and Christian Valledor for being leaders in those organizations. I would also like to thank the leadership, founders, and members of the Space Frontier Foundation, the International Space University, the National Space Society, the Space Generation Advisory Council, and the Mars Society for allowing me to work on projects and meet people that further my interests in space exploration. I will always be a part of the Space Mafia and you will always be with me in my journey exploring space.

Additionally, I would like to thank many friends who have been there for me even if they haven't directly impacted my career or research. I would like to thank John Alexander, not only for being a surrogate older brother to me in college and a true friend for over a decade, but also for getting me out of bed at 8 a.m. every morning to get me into a routine. Without your radio show, I would not have been as productive as I was during my time at UCF. I also want to thank John as well as Matt Spikes for showing me how to push the boundaries, how to be funny, and for helping me through particularly tough times. Thank you both for keeping me grounded and reminding me there's a home to come back to when this is all over. There are countless other people who have been positive influences in my life, and to each and every one of you I am thankful for the time we've spent together on this planet, no matter how brief.

In the realm of inspiration, I would like to thank all the members of Iron Maiden, DragonForce, Dream Theater, Amon Amarth, Manowar, Megadeth, System of a Down, Rage Against The Machine, and Judas Priest among countless other artists for their continued efforts to provide me with inspirational music to fight through the difficult and lonely times. I would like to thank Ron Howard and George Lucas for their work in filmmaking that partially inspired me to pursue a career in space exploration at a young age.

Last but certainly not least, I would like to thank my friend Sean Flynn. He agreed to edit this monstrosity of a thesis in exchange for only a single 12-oz beer and a mention in the acknowledgements for his contributions. Thank you for convincing me to buy a motorcycle and ride the Viking Trail with you. When I turn this in, we'll go back to Cape Breton and have a beer at the Red Shoe Pub.

Biographical Note

Ben Corbin was born in Fort Walton Beach, Florida, the son of Andy and Therese Corbin. He has two siblings, an older brother, Brian, and a younger sister, Claire. Growing up in a town sandwiched between three active Air Force bases taught Ben to look to the skies and yearn to go higher and faster.

Ben graduated from Okaloosa-Walton Community College (now Northwest Florida State College) Collegiate High School in 2004 with his high school diploma and associate degree, where he was the captain of the high school county championship-winning academic team, a member of the state championship college Brain Bowl team, and the Prom King. Scientists have identified no other realms, planes of existence, or dimensions where this stereotype-defying combination exists.

Ben is an Eagle Scout with over 60 merit badges and was active in Venture Crew throughout high school. The 5th Ranger Training Battalion taught him mountaineering, and he hiked one of the most strenuous treks available at Philmont Scout Ranch. His passion for outdoor adventure is boundless, as is his love of travel. He has been to 13 countries and hopes to quadruple that number in the next five to ten years.

Ben earned his boater's license in 1999, his driver's license in 2002, his scuba certification in 2003, his pilot's license in 2006, his skydiving license in 2011, his motorcycle license in 2012, and his suborbital scientist specialist's license in 2013. He plans to earn his B-license in skydiving in the next year and instrument flight rating when he has a real job. His dream house will have a garage and a basement so he can work on his motorcycles and build a sweet man cave for hosting sports and gaming parties.

It wasn't until Ben started at MIT that he was finally able to play ice hockey. He has played on the Aero/Astro intramural teams at every level available, mostly as a goalie. He served as the backup goalie for the B-league team (the highest IM league at the time) and was called upon to play in the league's championship game. He notched a shutout as Aero/Astro won yet another championship.

Ben earned master's degrees in Aeronautics and Astronautics and Earth and Planetary Sciences from MIT in 2011. No, there are not supposed to be any commas in that statement. His original goal was to make his doctoral dissertation shorter than his double master's thesis. He failed miserably and is so sorry to anyone who has to read this monstrosity.

Ben has played many musical instruments over the years, including piano, baritone, hand bells, trombone, flugelbone, trumpet, bass guitar, electric guitar, acoustic guitar, and drums. He enjoys a rich variety of music, including classic heavy metal, death metal, melodic death metal, technical progressive death metal, folk metal, power metal, speed metal, thrash metal, Viking metal, groove metal, and G-Funk-era hip hop.

Ben has been a member of the Students for the Exploration and Development of Space, the American Institute of Aeronautics and Astronautics, and several student rocket teams during his academic career. He is the outgoing EHS Safety Representative for the MIT Rocket Team, which has had zero injuries in his six years of service thanks to strict adherence to his oft-quoted motto: "Safety 3rd." He has served as the Party Liaison for the Space Frontier Foundation at the NewSpace conference for several years running, hosting wild parties disguised as "networking receptions" at a professional conference.

Table of Contents

ABSTRACT	3
ACKNOWLEDGEMENTS	4
BIOGRAPHICAL NOTE	7
TABLE OF CONTENTS	8
LIST OF FIGURES	17
LIST OF TABLES	23
LIST OF ACRONYMS	25
1. INTRODUCTION	28
1.1 Motivation and Background.....	28
1.2 Problem Statement and Thesis Objectives	31
1.3 Research Approach.....	31
1.4 Thesis Scope.....	33
1.5 Thesis Structure.....	33
2. LITERATURE REVIEW	35
2.1 Systems Engineering Methodologies	35
2.1.1 <i>Traditional Systems Engineering Methods</i>	35
2.1.1.1 Historical System Engineering Methods	35
2.1.1.2 NASA’s System Engineering Methods	36
2.1.1.3 The NASA Project Lifecycle.....	36
2.1.2 <i>Tradespace Exploration</i>	37
2.1.2.1 Development of MATE.....	38
2.1.2.2 Improvement of MATE and the Development of EEA and RSC.....	39
2.1.2.3 Other Tradespace Exploration Methods	40
2.1 Distributed Satellite Systems.....	41
2.1.1 <i>Satellite Constellations</i>	41
2.1.1.1 Constellations Forms	41
2.1.1.2 Constellation Performance Objectives	41
2.1.1.3 Constellations in Science Missions	42
2.1.2 <i>Fractionated Satellite Systems</i>	42
2.1.3 <i>Federated Satellite Systems</i>	43
2.1.4 <i>Design Optimization in Tradespace Exploration</i>	44

2.2	Satellite Communications.....	45
2.2.1	<i>Communications Subsystems Technologies</i>	45
2.2.1.1	Deployable Antennas	45
2.2.1.2	Inflatable Antennas.....	46
2.2.1.3	Laser/Optical Communications	46
2.2.2	<i>The Deep Space Network</i>	47
2.3	Science Mission Value Modeling.....	47
2.3.1	<i>Mission Selection Processes</i>	47
2.3.2	<i>Value-Centric Design Methodologies</i>	48
2.3.2.1	Economics-based VCDMs	48
2.3.2.2	Utility-based VCDMs.....	49
2.3.2.3	Alternative VCDMs	52
2.3.3	<i>Value in Time</i>	53
2.3.3.1	Epoch-Era Analysis.....	53
2.3.3.2	Decision Networks	57
2.3.4	<i>Value Elicitation and Structuring</i>	57
2.3.5	<i>Other Value Modeling Tools</i>	58
2.3.5.1	Value at Risk	58
2.3.5.2	Value Synergy and Redundancy.....	59
2.3.5.3	Alternate Aggregation Methods	60
2.4	Engineering Modeling.....	61
2.4.1	<i>Performance Modeling</i>	61
2.4.1.1	Coverage Modeling	61
2.4.1.2	Orbit Modeling.....	61
2.4.1.3	Propulsion and Trajectory Modeling	62
2.4.2	<i>Mass Modeling</i>	62
2.4.2.1	Instrument Models.....	62
2.4.2.2	Common Small Satellite Components.....	63
2.4.3	<i>Cost Modeling</i>	63
2.4.3.1	Simple Cost Models	63
2.4.3.2	Cost Estimating Relationships.....	64
2.4.3.3	Launch and Operations Costs	64
2.4.3.4	Direct Pricing	64
2.4.3.5	Cost Improvement	65
2.5	Research Gap Summary	65

3. THE EMERGENT CAPABILITIES OF DISTRIBUTED SATELLITE SYSTEMS.....67

3.1	Potential Benefits of DSS.....	67
3.2	Defining the Emergent Capabilities of DSS.....	70
3.2.1	<i>Fundamentally Unique Emergent Capabilities</i>	70
3.2.1.1	Shared Sampling	70
3.2.1.2	Simultaneous Sampling.....	71
3.2.1.3	Self-Sampling.....	72
3.2.2	<i>Analytically Unique Emergent Capabilities</i>	74
3.2.2.1	Census Sampling	74

3.2.2.2	Stacked Sampling	75
3.2.3	<i>Operationally Unique Emergent Capabilities</i>	76
3.2.3.1	Staged Sampling	76
3.2.3.2	Sacrifice Sampling	78
3.3	Chapter Summary	79

4. THE RESPONSIVE SYSTEMS COMPARISON METHOD FOR DISTRIBUTED SATELLITE SYSTEMS.....80

4.1	Overview	80
4.1.1	<i>General Framework Description</i>	81
4.1.2	<i>Summary of User Roles</i>	83
4.1.3	<i>Chapter Structure</i>	84
4.2	Phase 1: Value Modeling	84
4.2.1	<i>Contextualizing Phase 1 for Distributed Satellite Systems</i>	85
4.2.2	<i>Value Modeling Process 1: Value-Driving Context Definition</i>	87
4.2.3	<i>Value Modeling Process 2: Value-Driven Design Formulation</i>	88
4.2.3.1	VCDMs Revisited	88
4.2.3.2	Multi-Tiered MAUT Structure	90
4.2.3.3	Value in Multiple Measurements	92
4.2.3.4	Eliciting Single-Attribute Utility Curves	93
4.2.3.5	Eliciting Multi-Attribute Weighting Coefficients	95
4.2.3.6	Mental and Constructed Value Models	97
4.2.4	<i>Value Modeling Process 3: Epoch Characterization</i>	99
4.2.4.1	Epoch Shifts Due to Exogenous Mission Factors	100
4.2.4.2	Epoch Shifts Due to Endogenous Mission Factors	102
4.2.5	<i>Key Assumptions and Limitations</i>	103
4.3	Phase 2: Performance Modeling	105
4.3.1	<i>Contextualizing Phase 2 for Distributed Satellite Systems</i>	105
4.3.2	<i>Performance Modeling Process 1: Identifying Architecture and Design Variables</i>	107
4.3.3	<i>Performance Modeling Process 2: Design Modeling</i>	108
4.3.4	<i>Performance Modeling Process 3: Cost Modeling</i>	108
4.3.5	<i>Key Assumptions and Limitations</i>	110
4.4	Phase 3: Tradespace Exploration	111
4.4.1	<i>Contextualizing Phase 3 for Distributed Satellite Systems</i>	111
4.4.2	<i>Tradespace Exploration Processes</i>	112
4.4.3	<i>Key Assumptions and Limitations</i>	113
4.5	Phase 4: Epoch-Era Analysis	113
4.5.1	<i>Contextualizing Phase 4 for Distributed Satellite Systems</i>	113
4.5.2	<i>EEA Process 1: Single-Epoch Analyses</i>	114
4.5.3	<i>EEA Process 2: Multi-Epoch Analysis</i>	114
4.5.4	<i>EEA Process 3: Single-Era Analyses</i>	114
4.5.5	<i>EEA Process 4: Multi-Era Analysis</i>	114

4.5.6	<i>Cumulative Value Distribution and Value at Opportunity</i>	115
4.5.6.1	Value at Opportunity in Science Missions	115
4.5.6.2	Value at Opportunity in Other Distributed Systems	117
4.5.7	<i>Value Delivery over the Mission Lifecycle</i>	118
4.5.7.1	Continuous Value Delivery	118
4.5.7.2	Discrete Value Delivery	121
4.5.7.3	Combined Value Delivery	124
4.5.7.4	Value Delivery Summary	127
4.5.8	<i>Value Model Epoch Shifts in Era Analysis</i>	127
4.5.8.1	Retroactive Era Analysis	128
4.5.8.2	Added Value Era Analysis	129
4.5.8.3	Comparison to Previous Work	131
4.5.9	<i>Key Assumptions and Limitations</i>	132
4.5.9.1	General Assumptions and Limitations of EEA	132
4.5.9.2	Assumptions and Limitations of Proposed Modifications to EEA	132
4.6	Method Synthesis and Summary	133
4.6.1	<i>Revisiting Known Weaknesses</i>	133
4.6.2	<i>RSC Summary and Flowcharts</i>	134
4.7	Introduction to Case Studies.....	136
5.	CASE STUDY: EXOPLANETSAT	138
5.1	Case Study Introduction and Motivation.....	138
5.1.1	<i>Primary Mission Goals and Stakeholders</i>	138
5.1.2	<i>Mission Concept of Operations</i>	139
5.1.2.1	Single Satellite Model	139
5.1.2.2	Single Satellite Operations	141
5.1.2.3	Full Mission Operations	141
5.1.3	<i>Team X Study and Results</i>	142
5.1.4	<i>Motivation</i>	144
5.2	Case-Specific Literature Review	144
5.2.1	<i>Exoplanet Detection</i>	144
5.2.1.1	The Transit Method	144
5.2.1.2	Exoplanet Detection Requirements	145
5.2.2	<i>Similar Science Campaigns and Missions</i>	146
5.2.2.1	Previous and Current Missions.....	146
5.2.2.2	Future and Proposed Missions.....	147
5.2.3	<i>Performance Modeling in ExoplanetSat</i>	148
5.3	The RSC Model for the ExoplanetSat Mission	148
5.3.1	<i>ExoplanetSat RSC Phase 1: Value Modeling</i>	149
5.3.1.1	Value Modeling Process 1: Value-Driving Context Definition.....	149
5.3.1.2	Value Modeling Process 2: Value-Driven Design Formulation	150
5.3.1.3	Value Modeling Process 3: Epoch Characterization	154
5.3.2	<i>ExoplanetSat RSC Phase 2: Performance Modeling</i>	156
5.3.2.1	Performance Modeling Process 1: Identifying Architecture and Design Variables	156
5.3.2.2	Performance Modeling Process 2: Design Modeling	157
5.3.2.3	Performance Modeling Process 3: Cost Modeling	168

5.3.2.4	Performance Model Limitations	171
5.3.3	<i>ExoplanetSat RSC Phase 3: Tradespace Exploration</i>	172
5.3.4	<i>ExoplanetSat RSC Phase 4: Epoch-Era Analysis</i>	175
5.3.4.1	EEA Process 1: Single-Epoch Analyses	176
5.3.4.2	EEA Process 2: Multi-Epoch Analysis	177
5.3.4.3	EEA Process 3: Single Era Analyses	180
5.3.4.4	EEA Process 4: Multi-Era Analysis	183
5.4	Discussion and Conclusions	186
5.4.1	<i>Discussion of RSC Method Results</i>	186
5.4.1.1	Stakeholder Value Modeling	186
5.4.1.1	Optimal Constellation Design	186
5.4.1.2	Value Degradation and Constellation Replenishment	187
5.4.2	<i>Recommended Designs Solutions</i>	188
5.4.3	<i>Suggested Improvements for ExoplanetSat</i>	188
5.4.3.1	Detector Field of View	189
5.4.3.2	Detector Sensitivity	189
5.4.3.3	Solar Baffle	189
5.4.3.4	Operations Cost	189
5.4.4	<i>Comparisons to Monolithic Systems</i>	190
5.4.5	<i>Future Work</i>	190
5.5	Chapter Summary	191
6.	CASE STUDY: HOBOCOP	194
6.1	Case Study Introduction and Motivation	194
6.1.1	<i>Primary Mission Goals and Stakeholders</i>	195
6.1.1.1	Primary Science Goals	195
6.1.1.2	Primary Stakeholders	196
6.1.1.3	Secondary Mission Goals and Stakeholders	196
6.1.2	<i>Mission Concept of Operations</i>	196
6.1.2.1	Single Satellite Model	196
6.1.2.2	Single Satellite Concept of Operations	197
6.1.2.3	Full Mission Operations	201
6.2	Case-Specific Literature Review	202
6.2.1	<i>The Sun and Solar Phenomena</i>	202
6.2.1.1	The Solar Magnetic Field	202
6.2.1.2	The Solar Wind and Interplanetary Magnetic Field	203
6.2.1.3	The Heliospheric Current Sheet	205
6.2.2	<i>Similar Science Campaigns and Missions</i>	205
6.2.2.1	Campaigns to Study the Sun	205
6.2.2.2	Campaigns to Study Earth's Magnetic Field with DSS	207
6.2.3	<i>Performance Modeling in HOBOCOP</i>	208
6.2.3.1	Electric Propulsion Systems	208
6.2.3.2	Planetary Flybys	208
6.2.3.3	Radiation Environment	208
6.2.3.4	Thermal Environment	209
6.2.3.5	Instrument Modeling	209
6.3	The RSC Model for the HOBOCOP Mission	211

6.3.1	<i>HOBOCOP RSC Phase 1: Value Modeling</i>	211
6.3.1.1	Value Modeling Process 1: Value-Driving Context Definition.....	211
6.3.1.2	Value Modeling Process 2: Value-Driven Design Formulation.....	212
6.3.1.3	Value Modeling Process 3: Epoch Characterization.....	219
6.3.2	<i>HOBOCOP RSC Phase 2: Performance Modeling</i>	222
6.3.2.1	Performance Modeling Process 1: Identifying Architecture and Design Variables.....	222
6.3.2.2	Performance Modeling Process 2: Design Modeling.....	224
6.3.2.3	Performance Modeling Process 3: Cost Modeling.....	229
6.3.2.4	Performance Model Limitations.....	230
6.3.3	<i>HOBOCOP RSC Phase 3: Tradespace Exploration</i>	231
6.3.4	<i>HOBOCOP RSC Phase 4: Epoch-Era Analysis</i>	240
6.3.4.1	EEA Process 1: Single Epoch Analyses.....	240
6.3.4.2	EEA Process 2: Multi-Epoch Analysis.....	244
6.3.4.3	EEA Process 3: Single-Era Analyses.....	244
6.3.4.4	EEA Process 4: Multi-Era Analysis.....	245
6.4	Discussion and Conclusions.....	252
6.4.1	<i>Discussion of RSC Method Results</i>	252
6.4.1.1	The Value of Heterogeneity.....	252
6.4.1.2	Decision-Making with Cumulative Value Distribution.....	253
6.4.1.3	The Value of Communications Path Flexibility.....	253
6.4.2	<i>Recommended Design Solutions</i>	254
6.4.3	<i>Comparison to Monolithic Missions</i>	255
6.4.4	<i>Future Work</i>	255
6.5	Chapter Summary.....	256
7.	CASE STUDY: ÆGIR	259
7.1	Case Study Introduction and Motivation.....	259
7.1.1	<i>Primary Mission Goals and Stakeholders</i>	260
7.1.1.1	Primary Science Goals.....	262
7.1.1.2	Primary Stakeholders.....	263
7.1.1.3	Secondary Mission Goals and Stakeholders.....	263
7.1.2	<i>Mission Concept of Operations (Single Satellite)</i>	263
7.2	Case-Specific Literature Review.....	264
7.2.1	<i>Asteroid Classes and Compositions</i>	264
7.2.2	<i>Measuring Water on Asteroids</i>	265
7.2.2.1	Neutron Spectroscopy.....	265
7.2.2.2	Near-Infrared Spectroscopy.....	266
7.2.2.3	Radar Observations.....	266
7.2.3	<i>Similar Science Campaigns and Missions</i>	267
7.2.3.1	Asteroid Classification Taxonomies.....	267
7.2.3.2	Asteroid Ground Observations.....	268
7.2.3.3	Asteroid Visit Missions.....	268
7.2.3.4	Estimated Number of Asteroids in the Solar System.....	269
7.2.4	<i>Performance Modeling in ÆGIR</i>	270
7.2.4.1	The Traveling Asteroid Prospector Problem.....	270
7.2.4.2	Other Asteroid Prospectors.....	271

7.3	The Value Model for the ÆGIR Mission	271
7.3.1	<i>Value Modeling Process 1: Value-Driving Context Definition</i>	271
7.3.2	<i>Value Modeling Process 2: Value-Driven Design Formulation</i>	272
7.3.2.1	Top-Level Goal Attribute Weights.....	274
7.3.2.2	Instrument-Level Attributes	276
7.3.2.3	Lifetime Attributes	277
7.3.2.4	Statistical Attributes Among Spectral Types.....	277
7.3.2.5	Value Model Synthesis.....	279
7.3.3	<i>Value Modeling Process 3: Epoch Characterization</i>	281
7.3.3.1	Alternative Perceptions of Utility in Time	281
7.3.3.2	Alternative Stakeholder Perceptions of Statistical Attributes.....	282
7.3.3.1	Alternative Goal-Satisfying Instruments	285
7.3.3.2	Other Value Model Epoch Shifts.....	285
7.4	Discussion and Conclusions	286
7.4.1	<i>Discussion of RSC Value Modeling Results</i>	287
7.4.1.1	Challenges of Value Modeling with Multiple Targets	287
7.4.1.2	Payload and Bus Design Synergies	288
7.4.1.3	Adapting to New Technologies	288
7.4.1.4	Developing Operational Strategies Based on Value Models.....	288
7.4.1.5	Alternative Elicitation Techniques	289
7.4.1.6	Implications for Performance Modeling.....	289
7.4.2	<i>Comparison to Monolithic Missions</i>	289
7.4.3	<i>Future Work</i>	290
7.5	Chapter Summary	290
8.	CASE STUDY: GANGMIR.....	292
8.1	Case Study Introduction and Motivation.....	292
8.1.1	<i>Primary Mission Goals and Stakeholders</i>	293
8.1.1.1	Primary Science Goals	293
8.1.1.2	Primary Stakeholders	294
8.1.1.3	Secondary Mission Goals and Stakeholders.....	294
8.1.2	<i>Mission Concept of Operations</i>	294
8.1.2.1	Penetrator Operations	296
8.1.2.2	Small Satellite Operations	297
8.1.2.3	Combined Operations.....	297
8.2	Case-Specific Literature Review	298
8.2.1	<i>Mars Science Goals</i>	298
8.2.1.1	Mars Atmospheric Science.....	298
8.2.1.2	Mars Ground Science	298
8.2.2	<i>Similar Science Campaigns and Missions</i>	299
8.2.2.1	Previous Successful Mars Campaigns	300
8.2.2.2	Landing Site Selection Processes	301
8.2.2.3	Distributed Deployment on Mars	302
8.2.3	<i>Performance Modeling in GANGMIR</i>	303
8.2.3.1	Penetrators in Space Science Missions.....	303
8.2.3.2	Small Satellite Instrument Modeling	306
8.3	The RSC Model for the GANGMIR Mission	307

8.3.1	<i>GANGMIR RSC Phase 1: Value Modeling</i>	307
8.3.1.1	Value Modeling Process 1: Value-Driving Context Definition.....	307
8.3.1.2	Value Modeling Process 2: Value-Driven Design Formulation.....	309
8.3.1.3	Value Modeling Process 3: Epoch Characterization.....	316
8.3.2	<i>GANGMIR RSC Phase 2: Performance Modeling</i>	319
8.3.2.1	Performance Modeling Process 1: Identifying Architecture and Design Variables.....	319
8.3.2.2	Performance Modeling Process 2: Design Modeling.....	322
8.3.2.3	Performance Modeling Process 3: Cost Modeling.....	328
8.3.2.4	Performance Module Limitations and Scope.....	329
8.3.3	<i>GANGMIR RSC Phase 3: Tradespace Exploration</i>	330
8.3.3.1	Homogeneous Asset Comparison.....	330
8.3.3.1	Heterogeneous Asset Comparison.....	340
8.3.4	<i>GANGMIR RSC Phase 4: Epoch-Era Analysis</i>	340
8.3.4.1	EEA Process 1: Single Epoch Analyses.....	340
8.3.4.2	EEA Process 2: Multi-Epoch Analysis.....	340
8.3.4.3	EEA Process 3: Single Era Analyses.....	340
8.3.4.4	EEA Process 4: Multi-Era Analysis.....	340
8.4	Discussion and Conclusions.....	346
8.4.1	<i>Discussion of RSC Method Results</i>	346
8.4.1.1	Compromises Among Conflicting Small Satellite Design Variables.....	346
8.4.1.2	Substitution Effects in Instrument Attributes.....	347
8.4.1.3	Quantifying the Added Value of Stacked Sampling with Heterogeneous Assets.....	348
8.4.1.4	Quantifying the Added Value of Staged Sampling with Expendable Assets.....	348
8.4.2	<i>Recommended Design Solutions</i>	349
8.4.3	<i>Comparison to Monolithic Missions</i>	349
8.4.4	<i>Future Work</i>	350
8.5	Chapter Summary.....	351
9.	DISCUSSION AND CONCLUSIONS	352
9.1	Research Questions Revisited.....	352
9.1.1	<i>Research Question #1: Emergent Capabilities of DSS</i>	352
9.1.2	<i>Research Question #2: Analytical Techniques and Metrics</i>	354
9.1.3	<i>Combining Research Questions #1 and #2</i>	355
9.1.3.1	Quantifying Certain Benefits for Space Science Missions Leveraging DSS.....	355
9.1.3.2	Quantifying Uncertain Benefits for Responsive Systems.....	355
9.1.3.3	Comparisons to Current Methods for Evaluating Space Science Missions.....	356
9.1.4	<i>Linking Emergent Capabilities to Specific Analytical Techniques</i>	357
9.1.5	<i>Summary of Case-Specific Insights</i>	358
9.1.5.1	Case Study #1: ExoplanetSat.....	358
9.1.5.2	Case Study #2: HOBOCOP.....	359
9.1.5.3	Case Study #3: ÆGIR.....	361
9.1.5.4	Case Study #4: GANGMIR.....	362
9.2	Potential Impact on Systems Engineering.....	363
9.2.1	<i>Consequences of the Proposed MAUT Hierarchy</i>	363
9.2.2	<i>The Nature of Performance Modeling</i>	363
9.2.3	<i>Modified RSC Method in Other Applications</i>	364

9.2.4	<i>Recommendations for Concurrent Design Studios</i>	365
9.3	Potential Impact on Space Science Missions	366
9.3.1	<i>Mission Concept Development</i>	366
9.3.2	<i>Articulating Science Returns</i>	367
9.3.3	<i>Complementary and Competing Science Goals</i>	367
9.4	Additional Research Observations and Findings	368
9.4.1	<i>The Truth of Phenomena</i>	368
9.4.2	<i>The Nature of Modeling Uncertainty in RSC</i>	371
9.4.3	<i>The Difficulties of Stakeholder Value Modeling</i>	372
9.4.3.1	Method Acting As a Scientist.....	373
9.4.3.2	Stakeholder Interview Difficulties.....	373
9.4.3.3	Stakeholder Risk Attitudes	374
9.4.3.4	Stakeholder Availability Issues	375
9.4.3.5	Stakeholder Minimum Satisfaction Levels and Attribute Independence.....	375
9.4.3.6	Quantifying Scientific Benefits Versus Ranking Design Alternatives	377
9.4.4	<i>Adding Instruments or Assets to a DSS</i>	378
9.5	Future Work and Open Research Questions	378
9.5.1	<i>Future Work on Developing Methods</i>	379
9.5.2	<i>Future Work on Case Studies</i>	379
9.5.3	<i>Performance Modeling Research Questions</i>	379
9.5.4	<i>Value Modeling Research Questions</i>	380
9.6	Summary of Unique Thesis Contributions	380
9.6.1	<i>Definitions of Emergent Capabilities</i>	380
9.6.2	<i>Case Study Data Sets</i>	380
9.6.3	<i>Epoch-Era Analysis Value Modeling</i>	381
9.6.4	<i>Constellations for Astrophysics</i>	381
9.6.5	<i>Interplanetary Federated Communications Network</i>	381
9.7	Final Thoughts.....	381
	REFERENCES	382
	A. APPENDIX – EXOPLANETSAT	403
	B. APPENDIX – HOBOCOP	408
	C. APPENDIX – GANGMIR	414

List of Figures

Figure 1-1: Traditional (left) versus more recent (right) distributions of resources during a product development lifecycle (reprinted with permission from [6])	30
Figure 1-2: The four step research approach used in this work.	32
Figure 2-1: NASA Project Lifecycle (reprinted under Fair Use from [12])	37
Figure 2-2: Examples of monotonic single attribute utility curves. Left to right: linear, diminishing, increasing, and threshold returns.	50
Figure 2-3: Partitioning short run and long run into epochs and eras (reprinted with permission from [39])	54
Figure 2-4: Example epochs in an era for a serviceable satellite system (reprinted with permission from [185])	54
Figure 2-5: Graphical representation of the relationship between value, redundancy, and synergy with multiple measurements.	59
Figure 4-1: Major categories of processes in MATE and RSC (reprinted with permission from [25] and [58])	81
Figure 4-2: Major phases of the RSC method as described in this work.	82
Figure 4-3: Roles in an RSC case study for a DSS science mission.	84
Figure 4-4: Standard hierarchy of MAUT to calculate total MAU of a design.	90
Figure 4-5: Modified hierarchy of MAUT to calculate total MAU when multiple science goals and instruments are present.	90
Figure 4-6: ASSESS software examples with certainty equivalence with constant probability (left) and lottery equivalence/probability (right) for eliciting a single attribute utility curve.	95
Figure 4-7: ASSESS software demonstrating a choice that will help assign weights to the attributes.	96
Figure 4-8: Plots of MAU versus SAU for k -weightings that show synergy (left) and redundancy (left)	96
Figure 4-9: Potential MAUT hierarchy with one constructed model that describes stakeholder satisfaction.	99
Figure 4-10: Potential MAUT hierarchy where a constructed model or attributes is a part of what describes stakeholder satisfaction.	99
Figure 4-11: Examples of attribute curves for the primary mission (left), a mission extension (center), and possible ways to interpret the combination of the two (right).	101
Figure 4-12: Additional modification to MAUT hierarchy that more accurately expresses synergy between a set of top-level goals that does not express additional synergy with the remaining top-level goals.	104
Figure 4-13: Block diagram of the system performance and value models	105
Figure 4-14: Cost models of producing multiple units using (left) constant discount cost and (right) learning curve discount cost.	109
Figure 4-15: Comparison of (left) financial VaR and (right) systems VaO.....	115
Figure 4-16: Comparison of the probability distributions of value for designs with different numbers of assets and probabilities of failure given the same probability of discovering opportunity.....	116
Figure 4-17: Example experienced or static utility during the mission lifecycle for five different designs.	118
Figure 4-18: (Left) Constant attribute curve definition for coverage time. (Right) Cumulative value delivery model with a constant attribute curve definition.	119
Figure 4-19: (Left) Convex attribute curve definition for coverage time. (Right) Cumulative value delivery model with a convex attribute curve definition.	120
Figure 4-20: (Left) Concave attribute curve definition for coverage time. (Right) Cumulative value delivery model with a concave attribute curve definition.....	120
Figure 4-21: Linear returns on an impulsive value attribute. (Top left) Utility function. (Top right) Utility of four designs in time. (Bottom left) Impulsive value delivery. (Bottom right) Cumulative value delivery.	122
Figure 4-22: Diminishing returns on an impulsive value attribute. (Top left) Utility function. (Top right) Utility of four designs in time. (Bottom left) Impulsive value delivery. (Bottom right) Cum. value delivery.....	122

Figure 4-23: Increasing returns on an impulsive value attribute. (Top left) Utility function. (Top right) Utility of four designs in time. (Bottom left) Impulsive value delivery. (Bottom right) Cumulative value delivery.	123
Figure 4-24: Model of how value is delivered with both short-term and long-term measurements.	124
Figure 4-25: Less naive time value model that accounts for both discrete and long-term mission value and shows how value builds as more assets are deployed.	125
Figure 4-26: Complete model of time value delivery with delays in time-dependent value returns.	126
Figure 4-27: Example era with epoch shifts resulting in value model changes on an attribute.	128
Figure 4-28: Epoch shift showing how expectations can change. Although the value delivered in the first epoch is not diminished by the epoch shift, the expectations and highest levels of satisfaction have been raised significantly as a result of the epoch shift.	130
Figure 4-29: Affine scaling of the value model in the event of a mission extension with no changes in the components of the value model. The slope of the line decreases under the extension as the value is scaled down but does not change otherwise.	130
Figure 4-30: Affine scaling of the value model in a mission extension with changes in the components of the value model. The slope of the function can be different on one side of the epoch shift compared to the other if the components change.	131
Figure 4-31: Summary of RSC Phase 1: Value Modeling	134
Figure 4-32: Summary of RSC Phase 2: Performance Modeling	135
Figure 4-33: Summary of RSC Phase 4: Epoch-Era Analysis	135
Figure 5-1: Exploded view of the interior components of a single ExoplanetSat satellite (reprinted with permission from [273]).	140
Figure 5-2: Focal plane showing the detectors on ExoplanetSat (reprinted with permission from [271]. Image credit: Matt Smith).....	140
Figure 5-3: Concept of operations for single Exoplanetsat asset (reprinted with permission from [273]).	141
Figure 5-4: Three-year coverage of several stars by a constellation in a terminator, Sun-synchronous orbit (adapted from a model originally developed by Mary Knapp).	143
Figure 5-5: Example light curve from an exoplanet transit (reprinted with permission from [271]. Image credit: John Johnson).	145
Figure 5-6: Single-attribute utility for the expected number of exoplanets detected by ExoplanetSat.	151
Figure 5-7: Heuristic functions for probability of transit detection for observation time per orbit and number of shared satellite sample times (left) and observational lifetime (right).	153
Figure 5-8: Evolution of the test stars on the celestial sphere as viewed from the celestial North Pole for (1) initial low resolution scan during Team X study, (2) low resolution 4-direction scan, (3) medium resolution scan, and (4) high resolution scan for complete coverage mapping.	159
Figure 5-9: Coverage timeline for a single satellite at 28 degrees inclination viewing stars along the celestial prime meridian from the equator to the North Pole.	160
Figure 5-10: Illustration of how orbital precession prevents a single satellite from observing a target year-round, and how shared sampling can be leveraging to	161
Figure 5-11: Example coverage timeline of a particular test star (RA = 0, Dec = 68) by a 6-plane Walker constellation of ExoplanetSat assets.	161
Figure 5-12: Examples of how the path length within Earth's umbra and the length of orbit night changes with orbital inclination and altitude during (left) the Summer Solstice and (right) either seasonal equinox. (Model code originally written by Niraj Inamdar)	162
Figure 5-13: Graph showing how changes in orbital altitude and inclination affect the length of orbit night when the line of nodes is perpendicular to the Earth-Sun line during the Summer Solstice. (Model code originally written by Niraj Inamdar)	163
Figure 5-14: Coverage of a single star (Dec 60, RA 231) by one unique constellation design (3 planes, 52 degrees inclination, 500 km altitude) separated out by minimum length of coverage per orbit by a single satellite. (Left) Coverage from 4 to 13 minutes per orbit. (Right) Coverage from 14 to 23 minutes per orbit.	165

Figure 5-15: (Left) Coverage map of a particular constellation over the northern celestial hemisphere. Black dots are covered the entire year, blue dots are covered at least 90% of the year, red dots are covered for less than that. (Right) SVM modeling that defines the area covered and counts the number of actual target stars within the covered region. In these figures, the Vernal Equinox (RA = 0 degrees) is to the right, and during the Summer Solstice, the Sun is at the top (RA = 90 degrees).....	166
Figure 5-16: Coverage maps for constellations that only differ in inclination. All constellations are 6 planes at 500 km altitude observing targets for a minimum of 20 minutes per orbit. From left to right, orbital inclinations are 38, 44, and 52 degrees.	167
Figure 5-17: ExoplanetSat tradespace comparing total cost to Expected Number of exoplanets detected, which is the primary attribute for stakeholder satisfaction.	173
Figure 5-18: ExoplanetSat tradespace comparing total cost to probability of detecting at least 1 exoplanet. This is an alternative way for the stakeholder to view how value can be gained from the mission.	173
Figure 5-19: ExoplanetSat tradespace comparing total cost to single attribute utility.....	174
Figure 5-20: ExoplanetSat tradespace comparing total cost to single attribute utility. From this, we see that only designs that have no redundancy and observe for 20 minutes per orbit are on the Pareto front up until the point where the number of possible targets is less than the number of observable targets.....	174
Figure 5-21: ExoplanetSat tradespace comparing total cost to single attribute utility. Constellations that are “unsaturated” meaning there are more targets that can be observed than there are satellites times targets per satellite, tend to not be on the Pareto front.	175
Figure 5-22: Tradespaces showing variations in the Single-Attribute Utility in contexts ranging from R^*/a equals 1/215 (left), 1/150 (center), and 1/100 (right). As the ratio gets higher, utility tends to rise for all designs, though the Pareto front does not change significantly on the lower end of the curve. Even in these contexts, only unsaturated constellations with maximum observation time and no redundancy are on the Pareto front.	178
Figure 5-23: Tradespaces showing variations in Single-Attribute Utility in contexts ranging from number of targets per satellite equals 2 (left), 3 (center), and 4 (right).	178
Figure 5-24: Tradespaces showing variations in the Expected Number of exoplanets detected (top) and Single-Attribute Utility (bottom) by using the JPL Cost Model (left), the Custom Cost Model #1 (center), and the Custom Cost Model #2 (right).....	179
Figure 5-25: Multi-epoch analysis of (left) nine designs on the Pareto front examining (center) the FPN of each design in all 81 epochs, and (right) a closer look at Designs 1 and 9 and how they leave the Pareto front.	179
Figure 5-26: (Left) Overhead and (right) angled back view of an 8-plane constellation illustrating how losing a plane perpendicular to Earth-Sun line (red) at the Summer Solstice lowers the coverable area of the sky to that of a 4-plane constellation (green).	181
Figure 5-27: Example era with both failures and upgrades to demonstrate how static utility is not always the best metric for a system's value delivery.	182
Figure 5-28: Drop in EN (top) and SAU (bottom) versus operational lifetime for an unsaturated constellation of 15 satellites given satellite failure rates of 2.5% (left), 5% (center), and 10% (right) per year of operation.	184
Figure 5-29: Drop in EN (top) and SAU (bottom) versus number of satellites in an unsaturated constellation given satellite failure rates of 2.5% (left), 5% (center), and 10% (right) per year of operation.	185
Figure 6-1: Preliminary concept sketch of a HOBOSAT primary spacecraft bus. (Left) Close-approaching HOBOSAT with sun shield. (Right) Polar-orbiting HOBOSAT with VFM.	198
Figure 6-2: Preliminary concept sketch of a HOBOSAT primary spacecraft bus and expendable propulsion module.	198
Figure 6-3: An individual HOBOSAT satellite concept of operations (for close-approaching HOBOSATs).	199
Figure 6-4: An individual HOBOSAT satellite concept of operations (for polar orbiting HOBOSATs).....	199

Figure 6-5: Concept of Operations of a constellation of close-approaching HOBOS in ¼-year orbits with COPs. Royal blue is Earth’s orbit. Red curves are HOBOS’ burn trajectories. Yellow orbits are HOBOS’ operational orbits. Green curves are the COP burn trajectories. Cyan represents possible COP orbits. ..200	200
Figure 6-6: Concept of Operations of polar-orbiting HOBOS in ½-year period polar orbits with no COPs. Red lines are the inclination-changing burns. Yellow lines are coasting orbits. Green lines are operational orbits.200	200
Figure 6-7: WSA-ENLIL Cone Model of the Solar Wind. On July 14th, 2012, a CME hit Earth, which is shown here (reprinted under Fair Use. Image credit: NASA Goddard Space Weather Laboratory). ..204	204
Figure 6-8: Artist renditions of the heliospheric current sheet (reprinted under Fair Use from [322]).....204	204
Figure 6-9: Single Attribute Utility functions for the Vector Field Magnetometer instrument.213	213
Figure 6-10: Single Attribute Utility functions of HOBOCOP Science Goal #1213	213
Figure 6-11: Single Attribute Utility functions of HOBOCOP Science Goal #2216	216
Figure 6-12: Single Attribute Utility functions of HOBOCOP Science Goal #3218	218
Figure 6-13: Single Attribute Utility function of HOBOCOP Mission Lifetime.....219	219
Figure 6-14: (Left) Primary MAUT structure similar to what is shown in Figure 4-5. (Right) Alternative MAUT structure that better represents an enhanced relative synergy between Science Goals #2 and #3.219	219
Figure 6-15: Comparison of two different weighting of attributes and their effect on MAU of Science Goal #1. (Top) PSS weightings and how utility increases as number of viewing angles increasing from left to right. (Bottom) Alternative weightings showing higher weight on inclination angle and number of viewing angles.....221	221
Figure 6-16: Illustration of multiple satellites within the same orbital petal with different inclinations.224	224
Figure 6-17: Sketch of the one-dimensional thermal model to ensure survivable temperature of a HOBOS satellite with a heat shield.226	226
Figure 6-18: MATLAB model of the heliospheric current sheet at solar minimum (left) and near solar maximum (right).....229	229
Figure 6-19: Tradespace of HOBOCOP with only homogeneous assets, color-coded by payload option. (Top left) MAU of Goal #1. (Top right) MAU of Goal #2. (Bottom left) MAU of Goal #3. (Bottom Right) Total Mission MAU.....233	233
Figure 6-20: Tradespace of HOBOCOP with both homogeneous and heterogeneous assets, color-coded by payload option. (Top left) MAU of Goal #1. (Top right) MAU of Goal #2. (Bottom left) MAU of Goal #3. (Bottom Right) Total Mission MAU.234	234
Figure 6-21: Tradespace of HOBOCOP with both homogeneous and heterogeneous assets, color-coded by number of petals. (Top left) MAU of Goal #1. (Top right) MAU of Goal #2. (Bottom left) MAU of Goal #3. (Bottom Right) Total Mission MAU.235	235
Figure 6-22: Attributes and utility levels for Science Goal #1.236	236
Figure 6-23: Attributes and utility levels for Science Goal #2.237	237
Figure 6-24: Attributes and utility levels for Science Goal #3.238	238
Figure 6-25: Comparison of tradespaces showing (left) total mass versus (right) complex mass.240	240
Figure 6-26: Epoch comparisons of Mission MAU. (Top left) primary mission context. (Top right) pessimistic technology epoch. (Bottom left) alternate stakeholder value model changing weights on Goal #1. (Bottom right) alternate stakeholder value model and pessimistic technology epoch.241	241
Figure 6-27: Same as previous figure, but with X-Band technology instead of Ka-band in each epoch.....241	241
Figure 6-28: Epoch comparison of Goal #1. (Top left) original stakeholder value model, Ka-band (Top right) Original stakeholder value model, X-band. (Bottom left) Alternate stakeholder value model, Ka-band. (Bottom right) Alternate stakeholder value model, X-band.243	243
Figure 6-29: Multi-era analysis results for Design 112110.246	246
Figure 6-30: Multi-era analysis results for Design 144387.247	247
Figure 6-31: Instantaneous distribution of HOBOS satellites in an extreme case with many satellites with multiple petal configurations (2, 3, 4, and 5 are shown) and potential locations for COPs.249	249

Figure 6-32: Cumulative value distribution comparison of a single design that downlinks data (left) once per year, (center) twice per year, (right) three times per year.....	251
Figure 7-1: Proposed timeline of the exploration and pilot projects of the ÆGIR mission (timeline created by Jared Atkinson).....	261
Figure 7-2: Bus-DeMeo taxonomy of asteroids (reprinted with permission from [435]).....	267
Figure 7-3: Comparison of number of the number of discovered asteroids versus the number of asteroids that are assumed to exist (originally from [395], modified in [396] and reprinted with permission).....	270
Figure 7-4: MAU hierarchy for ÆGIR, showing all levels for Science Goal #1. The parts of the tree connected by blue lines are easily integrated into a standard MAU hierarchy. The parts of the tree connected by red lines are not, and an alternative value-ranking method must be developed.....	272
Figure 7-5: 3D utility functions showing the Total MAU when (left) Goal #1 MAU = 0, and (right) Goal #1 MAU = 1.....	275
Figure 7-6: Lifetime attribute utility function for ÆGIR Phase 1, labeled with Phase 2 milestones.....	277
Figure 7-7: Alternative special MAU hierarchy for degenerate case where no differences are perceived among the science goals for the spectral type adjusted and satisfaction weights. This structure inverts the top two levels compared to Figure 7-4.....	280
Figure 7-8: Timeline comparison of missions that trade frequency of asteroid visits with instrument capability.....	281
Figure 7-9: Alternative lifetime utility function with discounted utility over time.....	282
Figure 7-10: Alternative timeline comparing three designs with a different set of satisfaction weights.....	283
Figure 7-11: Comparison of varying levels of measurement uncertainty and varying sample means with the number of samples per spectral type and the relative size of the 95% confidence interval.....	284
Figure 8-1: GANGMIR mission assets concept design. (Top left) Individual penetrator. (Bottom left) Individual small satellite. (Right) Combined set with two small satellites and one penetrator.....	295
Figure 8-2: Concept of operations for a penetrator.....	295
Figure 8-3: Combined operations in GANGMIR with small satellites and ground penetrators.....	296
Figure 8-4: Mass fraction of water on Mars as determined by gamma ray spectrometer on Mars Odyssey (Image credit: NASA/JPL/Los Alamos National Laboratory [479]).....	299
Figure 8-5: Cross-section sketch of the Deep Space 2 micropenetrator (reprinted under fair use from [502]).....	304
Figure 8-6: Cross-section of a planetary penetrator (reprinted under fair use from [506]).....	305
Figure 8-7: Single Attribute Utility functions for Science Goal #1 of the GANGMIR mission.....	311
Figure 8-8: Hierarchy of the attributes of GANGMIR Science Goal #1: Observe Global Temperature Field.....	311
Figure 8-9: Map of desired landing sites for a Mars penetrator mission (data taken from [514]) and previous successful landing campaigns.....	312
Figure 8-10: Single Attribute Utility functions of Science Goal #2 of the GANGMIR mission.....	314
Figure 8-11: Vensim diagram showing the interactions of potential GANGMIR variables.....	320
Figure 8-12: Code map of the system integrator in GANGMIR.....	328
Figure 8-13: Tradespace for Science Goals 1-4 from left to right, showing the gained from (top) small satellites, (middle) ground penetrators, and (bottom) combined utility. Color coding: red = 1 or 2 assets; yellow = 3 or 4; green = 5 or 6; cyan = 7 or 8; blue = 9 or 10; black = 11 or 12; magenta = 13 or 14.....	332
Figure 8-14: Tradespace for Science Goals 5-8 from left to right, showing the gained from the ground instruments (top), orbiting instruments (middle), and combined utility (bottom). Color coding: red = 1 or 2 assets; yellow = 3 or 4; green = 5 or 6; cyan = 7 or 8; blue = 9 or 10; black = 11 or 12; magenta = 13 or 14.....	333
Figure 8-15: Tradespaces of instrument-level MAUs for goals related to the small satellite radiometer showing variations in (top) the number of wavelength bands and (bottom) aperture diameter.....	334
Figure 8-16: Tradespaces of instrument-level MAUs for goals related to the small satellite camera showing variations in (left) aperture diameter and (right) focal ratio.....	335
Figure 8-17: Tradespaces of total mission MAU showing how utility varies with (left) camera aperture diameter, (left center) camera F#, (right center) number of radiometer bands, and (right) radiometer aperture diameter.....	335

Figure 8-18: Tradespaces of instrument-level MAUs for goals related to the penetrators' (top) weather package options and (bottom) spectrometer package options. The right-most plots show total mission MAU	336
Figure 8-19: Tradespace of total mission MAU for designs with either smaller satellites or penetrators but not both.	338
Figure 8-20: Tradespace of total mission MAU for designs with 1:1 ratio of small satellites to penetrators.	339
Figure 8-21: Tradespace of total mission MAU for designs with 2:1 ratio of small satellites as penetrators.	339
Figure 8-22: Cumulative value distribution curves for assets that leverage staged sampling to achieve opportunities that arise from endogenous epoch shifts.	342
Figure 8-23: Additional cumulative value distribution comparisons.	343
Figure 8-24: Example cumulative probability distribution functions for differing number of penetrators. Value at risk (VaR) is shown in red, value at opportunity (VaO) is shown in green.	343
Figure 8-25: Comparison of (left) the ratio of the average and expected value and (right) value at opportunity (VaO) versus the number of penetrators.	344
Figure 8-26: Updated tradespace incorporating shifts due to average value calculated from EEA.	345
Figure 8-27: Change in the expected value for penetrators.	345
Figure 8-28: Value probability distribution for a rover. VaR is shown in red (1%), VaO is in green (23.7%)	350
Figure 9-1: Spectrum of science knowledge for a science goal.	370
Figure 9-2: Comparison between the reported minimum acceptable levels of attributes for (left) independent attributes, and (right) dependent attributes or incorrectly reported acceptable minima.	376
Figure 9-3: Comparison of different interpretations of "zero value" when examining systems with multiple discrete measurements.	377
Figure A-1: Comparison of the Expected Number vs. mission lifetime for 15 (top left), 20 (top right), 25 (bottom left), and 30 (bottom right) satellites per constellation with 2.5% failure rate.	403
Figure A-2: Comparison of the Expected Number vs. mission lifetime for 15 (top left), 20 (top right), 25 (bottom left), and 30 (bottom right) satellites per constellation with 5% failure rate.	404
Figure A-3: Comparison of the Expected Number vs. mission lifetime for 15 (top left), 20 (top right), 25 (bottom left), and 30 (bottom right) satellites per constellation with 10% failure rate.	405
Figure A-4: Tradespaces showing variations in the Expected Number of Exoplanets Discovered (top) and Single-Attribute Utility (bottom) in contexts ranging from R^*/a equals 1/215 (left), 1/150 (center), and 1/100 (right). As the ratio gets higher, utility tends to rise for all designs, though the Pareto front does not change significantly on the lower end of the curve.	406
Figure A-5: Tradespaces showing variations in the Expected Number of Exoplanets Discovered (top) and Single-Attribute Utility (bottom) in contexts ranging from Number of Targets Per Satellite equals 2 (left), 3 (center), and 4 (right).	407
Figure B-1: Enlarged tradespace of HOBOCOP.	409
Figure B-2: Era analysis results for a particular design (designation 153081) with one contact per year.	410
Figure B-3: Era analysis results for a particular design (designation 153081) with two contacts per year.	411
Figure B-4: Era analysis results for a particular design (designation 153081) with three contacts per year.	412
Figure B-5: Era analysis results for a particular design (designation 153081) with four contacts per year.	413
Figure C-1: SAU functions for Science Goal #3 GANGMIR (Observe Global Aerosol Composition).....	417
Figure C-2: SAU functions for Science Goal #4 in GANGMIR (Measure Global Surface Pressure).....	417
Figure C-3: SAU functions for Science Goal #5 GANGMIR (Observe Local/Regional Weather).....	418
Figure C-4: SAU functions for Science Goal #6 in GANGMIR (Observe Dust and Aerosol Activity).....	418
Figure C-5: SAU functions for Science Goal #7 in GANGMIR (Measure Presence of Ground Ice).....	419
Figure C-6: SAU functions for Science Goal #8 in GANGMIR (Obtain Atmospheric EDL Profiles).	419

List of Tables

Table 2-1: Estimated cost improvement schedule for spacecraft recurring costs [227].....	65
Table 3-1: List of the potential benefits of DSS (items in bold are in the “Big Four”)	68
Table 4-1: Summary of the properties of typical space science missions and how well popular VCDMs can be used to assess their value	89
Table 4-2: Possible outcomes of a science mission and operational situations with multiple assets.....	103
Table 4-3: Summary of the properties of typical space science missions and how well variations of MAUT can be used to assess their value	134
Table 4-4: Summary of the important differences among the four case studies.	136
Table 5-1: Comparison of current exoplanet detection missions and requirements to find Earth-sized planets.....	147
Table 5-2: Comparison of future exoplanet detection missions and requirements to find Earth-sized planets.	148
Table 5-3: Context variables and their values for ExoplanetSat.....	155
Table 5-4: Design variables and their discrete values for ExoplanetSat.....	157
Table 5-5: Important parameters for the STK coverage module of ExoplanetSat.....	158
Table 5-6: Team X’s ExoplanetSat cost model line items with multiplying factors for additional units.	169
Table 5-7: Fixed costs in Team X’s ExoplanetSat cost model.....	169
Table 5-8: Costs for hardware and ground support operations, which do not scale the same way as other cost elements listed in the previous tables.	169
Table 5-9: Summary of total costs per satellite in the ExoplanetSat constellation.	170
Table 5-10: Custom Cost Models #1 and #2 line items with constant multiplying factors for more units.	170
Table 5-11: Custom Cost Models #1 and #2 for payload and ground support operations.	170
Table 6-1: Primary Science Goals in HOBOCOP	196
Table 6-2: A sample of the campaigns launched or planned to study solar phenomena.....	205
Table 6-3: Available CubeSat magnetometer options.	209
Table 6-4: Available CubeSat energetic particle analyzer options.	210
Table 6-5: Single Attribute Utility weights for the Vector Field Magnetometer instrument.	212
Table 6-6: Single Attribute Utility weights for HOBOCOP Science Goal #1.....	214
Table 6-7: Single Attribute Utility weights for HOBOCOP Science Goal #2.....	215
Table 6-8: Single Attribute Utility weights for HOBOCOP Science Goal #3.....	217
Table 6-9: Science Goal MAU weights for HOBOCOP	217
Table 6-10: Alternative science goal MAU weight for HOBOCOP with split synergy.	219
Table 6-11: Comparison of parameters used in both the Optimistic context and the SMAD-Based context for spacecraft mass estimations.....	220
Table 6-12: Comparison of differing stakeholder opinion on attribute weights of Science Goal #1.	221
Table 6-13: Context variables and their values for HOBOCOP.....	222
Table 6-14: Payload Packages available for HOBOSatellites.....	223
Table 6-15: Architectural and design variables and their discrete values for HOBOCOP	223
Table 6-16: Known components of the HOBOSatellite, suppliers, and masses.....	225
Table 6-17: Cost Parameters for HOBOCOP Cost Estimate.....	230
Table 6-18: Comparison of the fraction of feasible designs in different technology epochs.....	243
Table 6-19: Comparison of the difference in cumulative value between satellites that communicate only with Earth and satellites that can communicate multiple times per year through relay satellites.....	250
Table 7-1: Timeline of events during Phase 1 of ÆGIR (schedule created by Jared Atkinson).....	261
Table 7-2: Timeline of events during Phase 2 of ÆGIR (schedule created by Jared Atkinson).....	262
Table 7-3: Primary Science Goals in ÆGIR.....	262

Table 7-4: Potential stakeholders and their relative importance in the ÆGIR mission (modified from an original table created by Jared Atkinson).	263
Table 7-5: Potential concept architectures, instruments, and instrument attributes for the Exploration Phase of ÆGIR (created in conjunction with Jared Atkinson).	273
Table 7-6: Goal-level attribute weights elicited from the initial stakeholder interview using ASSESS.	274
Table 7-7: Results from hypothetical “science broker” interview selling combinations of goals.	275
Table 7-8: Goal-level attribute weights scaled from the results of the “science broker” interview.	275
Table 7-9: Desired spectral types of asteroids to study and their relative importance as they were elicited from stakeholder interviews.	278
Table 7-10: Relative satisfaction for number of subsequent visits per spectral class.	279
Table 8-1: Top Ten science goals of the MEPAG and P-SAG.....	294
Table 8-2: Notable successful Mars missions.....	300
Table 8-3: Instruments designed for the LunarNet penetrator mission (data from).....	305
Table 8-4: Instruments design for the METNET mission (data from [507]).....	305
Table 8-5: MCS spectral channel bands and functions (data from [512]) arranged by band center.	306
Table 8-6: Summary of small satellite camera data for use in GANGMIR.	307
Table 8-7: Complete list of science goals being considered in GANGMIR.	308
Table 8-8: Single Attribute Utility weights for GANGMIR Science Goal #1.....	311
Table 8-9: Landing sites identified by the students by the “Ground” group at Politecnico di Torino.....	312
Table 8-10: Complete set of value-driving attributes of all science goals in GANGMIR.....	316
Table 8-11: Summary of possible epochs in GANGMIR that could occur through staged sampling.....	318
Table 8-12: Design variables and their discrete values for GANGMIR.....	321
Table 8-13: Instrument options for the penetrator subsurface package.	322
Table 8-14: Instrument options for the penetrator surface package.....	323
Table 8-15: Cost Parameters for GANGMIR Cost Estimate.....	328
Table 9-1: Summary of the properties of typical space science missions and how well some VCDMs can be used to assess their value. (This table is a combination of Table 4-1 and Table 4-3).	357
Table B-1: List of natural constants.....	408
Table B-2: List of parameter assumptions in the thermal submodule.....	408
Table B-3: List of satellite component parameters and assumptions.....	408
Table B-4: List of propulsion system component parameters and assumptions.	409
Table B-5: List of assumed instrument data rates.....	409
Table C-1: List of natural constants.....	414
Table C-2: List of parameters of payloads used to build parametric payload models.	414
Table C-3: List of parameters for payload performance.....	414
Table C-4: List of satellite components and assumptions.....	415
Table C-5: List of mothership components and assumptions.	415
Table C-6: List of mothership arrival parameters.....	415
Table C-7: List of penetrator components and assumptions.....	416

List of Acronyms

ÆGIR	Asteroid Exploration and Geohydrology for In-situ Resources
AHP	Analytic Hierarchy Process
ARL	Applied Research Laboratory
ASTERIA	Arcsecond Space Telescope Enabling Research in Astrophysics
ATH	Above-the-Horizon
AU	Astronomical Unit (1 AU = 149,597,871 km)
BRITE	Bright Target Explorer
BOL	Beginning- of-Life
CBA	Cost-Benefit Analysis
CER	Cost Estimating Relationship
CHEOPS	Characterizing Exoplanets Satellite
CME	Coronal Mass Ejection
COCOMO	Constructive Cost Model
CoRoT	Convection, Rotation and planetary Transits
CPT	Cumulative Prospect Theory
CPU	Central Processing Unit
DARPA	Defense Advanced Research Projects Agency
DSN	Deep Space Network
DSS	Distributed Satellite Systems
ΔV	Delta-V, or change in velocity
DVM	Design Value Matrix
EDL	Entry, Descent, and Landing
EEA	Epoch-Era Analysis
EIRP	Equivalent Isotropically Radiated Power
EOL	End-of-Life
EXPLORE	Exosolar Planet Occultation Research
FCC	Federal Communications Commission
F6	Future, Fast, Flexible, Fractionated Free-Flying spacecraft
fNPT	Fuzzy Normalized Pareto Trace
FOV	Field-of-View
FPA	Fast Particle Analyzer
FPN	Fuzzy Pareto Number
FSS	Federated Satellite Systems (not to be confused with Fractionated Satellite Systems)
FWHM	Full-Width Half Maximum
GA	Genetic Algorithms
GANGMIR	Groups of Atmospheric Nanosatellites and Ground Penetrators for Mars In-situ Research
GCR	Galactic Cosmic Rays
GEO	Geosynchronous Earth Orbit
GINA	Generalized Information Network Analysis
GRACE	Gravity Recovery and Climate Experiment
GRAIL	Gravity Recovery and Interior Laboratory
GT-FAST	Georgia Tech F6 Architecture Synthesis Tool
HOBOCOP	Heliocentric-Orbiting Baseline-Optimized Cosmos Observation Paradigm
HPOP	High-Precision Orbital Propagator
I_{sp}	Specific Impulse
ISRU	In-Situ Resource Utilization

ISS	International Space Station
ITAR	International Traffic of Arms Regulations
ITU	International Telecommunications Union
IVTea	Interactive Visual Tradespace Exploration Application
JPL	Jet Propulsion Laboratory
LADEE	Lunar Atmosphere and Dust Environment Explorer
LEO	Low Earth Orbit
MATE	Multi-Attribute Tradespace Exploration
MATE-CON	Multi-Attribute Tradespace Exploration with Concurrent Engineering
MAU	Multi-Attribute Utility
MAUT	Multi-Attribute Utility Theory
MAVEN	Mars Atmosphere and Volatile Evolution
MCS	Mars Climate Sounder
MEPAG	Mars Exploration Program Analysis Group
MER	Mars Exploration Rover
MDO	Multidisciplinary Design Optimization
MIST	MATE Interview Software Tool
MIT	Massachusetts Institute of Technology
MLI	Multi-Layer Insulation
MMDOSA	Multi-objective, Multidisciplinary Design Optimization System Architecting
MRO	Mars Reconnaissance Orbiter
MSL	Mars Science Laboratory
NASA	National Aeronautics and Space Administration
NBI	Normal-Boundary Intersection
NEO	Near-Earth Object
NEN	Near-Earth Network
NICM	NASA Instrument Cost Model
NPT	Normalized Pareto Trace
NPV	Net-Present Value
PHI	Polarimetric and Helioseismic Imager
PHO	Potentially Hazardous Object
PI	Principal Investigator
PLATO	Planetary Transits and Oscillations of stars
P-SAG	Precursor Strategy Analysis Group
PSO	Particle Swarm Optimization
PSS	Primary Science Stakeholder
QFD	Quality Function Deployment
RA	Right Ascension
RAAN	Right Ascension of the Ascending Node
RSC	Responsive Systems Comparison
RTG	Radioisotope Thermoelectric Generator
RV	Radial Velocity
SA	Simulated Annealing
SAU	Single Attribute Utility
SBAG	Small Bodies Assessment Group
SBD	Set-Based Design
SEArI	Systems Engineering Advancement Research Initiative
SEET	Space Environment and Effects Tool
SMAD	Space Mission Analysis and Design
SME	Space Mission Engineering
SMS	Secondary Mission Stakeholders

SNR	Signal-to-Noise Ratio
SOHO	Solar and Heliophysics Observatory
SPE	Solar Particle Event
SPENVIS	Space Environment, Effects, and Education System
SSL	Space Systems Laboratory
SSPARC	Space Systems Policy and Architecture Research Consortium
SSS	Secondary Science Stakeholders
SVM	Support Vector Machine
STEREO	Solar Terrestrial Relations Observatory
STK	Systems Toolkit (formerly Satellite Toolkit)
TDN	Time-Expanded Decision Network
THEMIS	Time History of Events and Macroscale Interactions during Substorms (Chapter 6)
THEMIS	Thermal Emission Imaging System (Chapter 8)
TRL	Technology Readiness Level
TSE	Tradespace Exploration
TOPSIS	Technique for Order Preference by Similarity to Ideal Solution
TPF	Terrestrial Planet Finder
USCM8	Unmanned Space vehicle Cost Model v.8
VaO	Value at Opportunity
VaR	Value at Risk
VASC	Valuation Approach for Strategic Changeability
VCDM	Value-Centric Design Methodology
VF	Value Functions
VFM	Vector Field Magnetograph
VWFO	Value-Weighted Filtered Outdegree
YORP	Yarkovsky–O'Keefe–Radzievskii–Paddack effect

1. Introduction

*“We make the pact, we took our first breath. **Fearless, this is our quest**
That leads me on... man the star-gates! Star gates! Go!
We are the dark and the light! We have the power of time!
We are the serpents intertwined! The galaxy will unite!
This is the beginning! It’s all starting now!”*

DETHKLOK, “THE GALAXY”

Most doctoral dissertations are highly specialized, narrow works that push the boundary of a specific field by building on the work of those who came immediately before them. Standing at the top of a tower centuries in the making, they add a single brick to the collective knowledge of humanity so that their tower can stand taller above the unenlightened ground from which mankind has risen.

This is not like those dissertations. *This work is intended to serve as a how-to framework* to anyone wishing to conduct a concept design or early phase study of a space science mission, in addition to standing alone as a contribution to the field of systems engineering worthy of the award of a doctoral degree. This work is intentionally broad to cover many audiences, but detailed enough for a reader with any background in space science, space mission engineering, or systems engineering to utilize. Rather than standing atop the tower of any one field, this dissertation serves as a bridge among many academic towers, offering a multidisciplinary approach to analyzing space science missions in dynamic and evolving contexts.

Specifically, this work focuses on space science missions that leverage the *emergent capabilities* of distributed satellite systems in order to achieve fundamentally new space science goals. This dissertation asks “What interesting scientific questions can be answered *because* multiple satellites are working in concert?” and “What low-hanging scientific fruit is waiting to be picked because it is only possible to do so with multiple satellites?”

1.1 Motivation and Background

This work stands at the intersection of several changing paradigms in spaceflight as well as challenges that have been present since the beginning of the space age. First, the budgets of national space agencies are limited, and more must be accomplished with fewer resources in order to demonstrate that taxpayer dollars are being spent wisely. No space agency has ever had an unlimited budget, and every agency must prioritize its goals and achieve as much as possible with the resources available.

Thanks to the digital revolution, prices of individual satellite components have decreased significantly over the past thirty years, and especially in the last ten thanks to an explosion in commercial space companies developing technologies for small satellites. Additionally, the cost of access to space has decreased significantly due to competition in the launch market, thanks to commercial launch vehicles competing for contracts. These two paradigm shifts have sparked a revolution in small satellites. The price point for developing, launching, and operating a satellite has fallen enough for some universities to

be able to afford their own satellite educational curricula. As more numerous and more efficient resources, instruments, and technologies are developed, the operational value of small satellites will continue to grow.

Despite all this, the cost of major, government-funded planetary science missions has *risen* [1]. Fewer missions are sharing the same pool of resources as more complex missions are being proposed. If those missions go over budget, they siphon resources from other missions in their early phases. Some of these resource-starved missions are delayed; others are simply canceled [2]–[4]. These missions may have had the potential to answer novel science questions, but were cut due to the expanding needs of missions that had already been committed.

Another changing paradigm in space missions is the rise in multidisciplinary design optimization techniques and other systems engineering methodologies to help increase the knowledge gained early on in the mission lifecycle. In addition to being able to optimize many variables across many engineering disciplines, more concept architectures can be studied to see which ones deliver the most benefit at cost.

Over the past 15 years in particular, much effort has gone into understanding the value of changeable systems – systems that can operate in a variety of contexts, are capable of responding to situations beyond their original design plan, are not simply passively robust to exogenous circumstances, and can take advantage of opportunities that arise during the mission lifecycle. Methods and frameworks to evaluate such systems and rank them against alternative design solutions have been developed to analyze the unarticulated value in these systems over their lifecycles.

These methods are also intended to increase the knowledge gained early in the mission development process. A modest increase in the cost incurred during the concept development phase can significantly increase the knowledge gained, delay the cost committed, and extend management leverage available in the design process (as illustrated in Figure 1-1). Knowing more about a system long before a single part is fabricated can prevent cost overruns later in the lifecycle.

Unfortunately, this research has not been formally incorporated into space agencies' decision-making processes early in mission selection phases. Concurrent engineering studios tend to gravitate towards optimizing point designs or conducting a small series of *trade studies* to provide a handful of design choices; they tend not to explore a *tradespace* of possible designs by examining the effects of changing many design variables simultaneously. Additionally, trade studies do not necessarily fundamentally articulate what drives *value* in space science missions, whereas new methods begin by understanding and articulating what is truly important for mission success.

A third changing paradigm in space missions is the use of distributed satellite systems (DSS). Distributed satellites, for the purposes of this work, are defined as any satellite system that consists of multiple space-based assets working on the same mission, whether they are homogeneous or heterogeneous, whether they are spatially distant or locally clustered, and whether they are operating in the same environment. Fractionated satellites are a subset of DSS, whereas federated satellite systems may be considered a subset or a superset, depending on the definition and purpose within a “system of systems” [5].

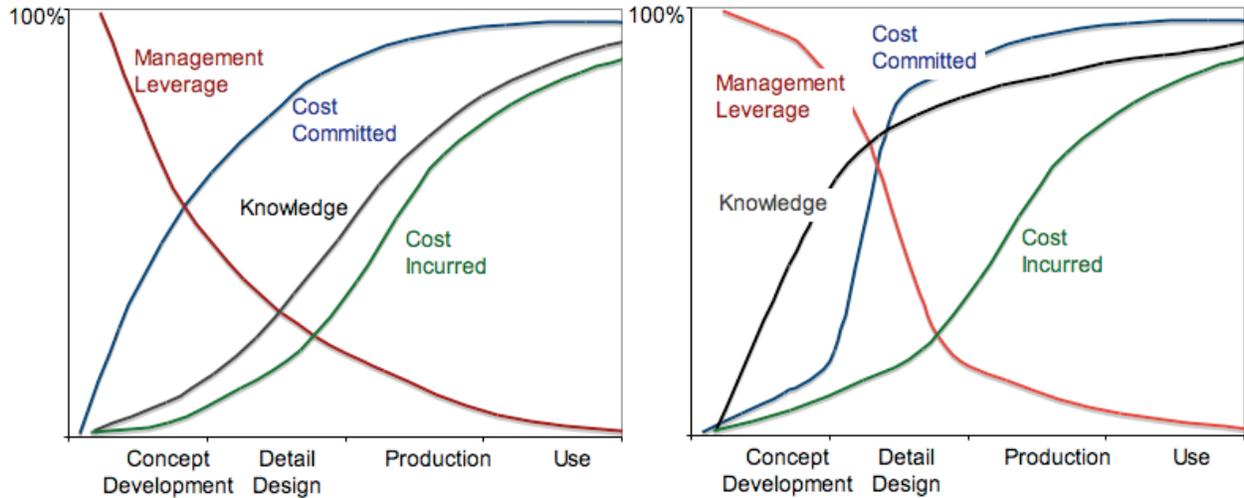


Figure 1-1: Traditional (left) versus more recent (right) distributions of resources during a product development lifecycle (reprinted with permission from [6])

There is great potential to accomplish more with DSS than with monolithic, or single-satellite, systems. Systems that utilize multiple assets have been shown to be less prone to system failures caused by the failure of individual parts, more easily replaced and upgraded, and more likely to be able to react to changing circumstances in their performance expectations. These lifecycle properties of distributed systems have been the subject of much research over the past twenty years.

Those added lifecycle properties come at a price. Having more independently operating components requires more redundant subsystems. Specific payload components can be shared or split, but so many other satellite subsystems must be reproduced that the mass penalty (and therefore cost penalty) is likely to outweigh the added benefit. Previous research has shown that the value proposition of fractionated and distributed satellites is poor compared to monolithic systems, and that the added benefits from the emergent lifecycle benefits were difficult to communicate to decision makers and tended to result in added value only after long mission lifetimes. Benefits like upgradability may be even less useful in space science missions beyond low-Earth orbit because of the added costs associated with dedicated launches to travel to the mission's operational area with no other systems to spread costs.

However, these previous studies should be reexamined from two new viewpoints. First, these studies used either traditional or relatively immature evaluation methods to make the comparison between monolithic and distributed systems. New developments over the past five years have shown great promise in better articulating the lifecycle value of lifecycle properties of DSS, especially in space science missions.

Second, these studies were making direct comparisons between monolithic and distributed systems in operations wherein monolithic systems have traditionally dominated, and for good reasons. The mission objectives were typically well-suited for and could be accomplished by a single satellite. If a monolithic satellite is designed well enough, and the stakeholder expectations change little between the start of the detailed design process and the end of the mission lifecycle, there is no way the added value gained through lifecycle properties of any DSS could compete with the cost of the monolithic satellite.

Rather than simply retread the ground that these studies have already covered with the updated methods as described in the first viewpoint, this research seeks to address an alternative approach alluded to in the second viewpoint. Instead of attempting to shoehorn a DSS into a mission in which monolithic systems have performed well, this work seeks to flip the perspective: what scientific objectives can distributed satellite systems accomplish that monolithic satellite systems cannot?

Comparatively little research has been conducted on what science questions can be answered by leveraging the *emergent capabilities* of DSS, or the value of answering such questions. In particular, this work hypothesizes that there may be new scientific vistas that have not yet been explored because they can only be explored with DSS, and the value of doing so can be better articulated with the methods described herein.

1.2 Problem Statement and Thesis Objectives

The purpose of this dissertation is to formulate the value proposition for distributed satellites by understanding and leveraging the unique, emergent capabilities of distributed satellite architectures. This will be accomplished by developing various metrics that describe system performance, quantifying the inherent value of a system to scientific stakeholders, and leveraging new techniques of tradespace exploration. In particular, we seek answers to the following research questions:

1. **What emergent capabilities can be demonstrated to achieve science objectives with distributed satellite systems that cannot be achieved with monolithic satellite systems?**
2. **Which analytical techniques and metrics should be applied or modified in the development of distributed satellite systems for science missions to more accurately characterize the overall benefit over the entire mission lifecycle?**

Put another way, we ask, “What can multiple satellites working together do that simply cannot be done with a single satellite, no matter how advanced its instrument package?” and “In what new ways can we predict the return on investment for a distributed system by examining both the baseline mission scenario and the changing contexts within which the system may be forced to operate?”

1.3 Research Approach

To answer these research questions, a four step research methodology shown in Figure 1-2 is followed: (1) knowledge capture and synthesis, (2) capabilities definitions, (3) method development, and (4) case studies and applications. Each step is not a discrete chapter in a serial process, because knowledge is continually captured that shifts the direction and scope of later steps, and results from early iterations of the case studies feeds back into the method development. Additionally, some emergent capabilities are identified before a rigorous definitions process is conducted that change how the missions in the case studies leverage these capabilities.

First, a thorough literature review is conducted to capture knowledge in a variety of fields, including the development of complex systems, space science, and value modeling. This review examines the methods and practices currently in use and in development to handle a number of subjects that arise in the development of space science missions. Scientific literature related to each case study is also examined to help understand the meaning and impact that the science goals that could be achieved through these missions would have on its field and related fields.

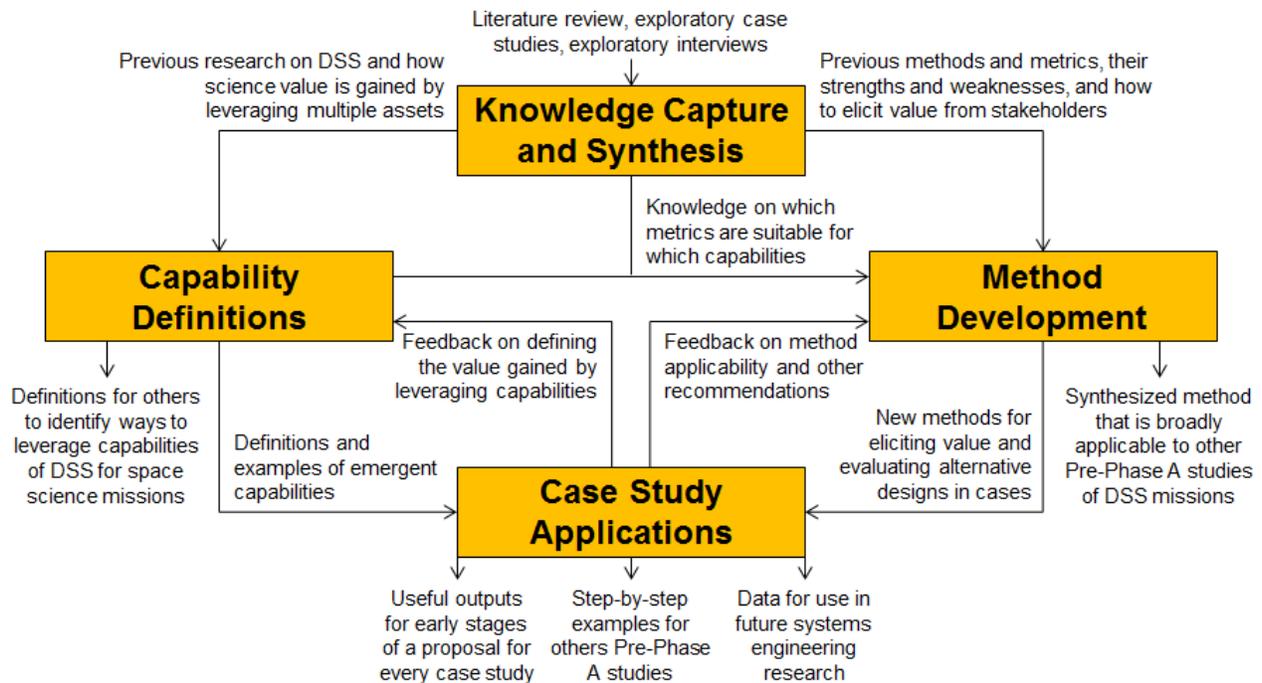


Figure 1-2: The four step research approach used in this work.

The second step is to define the emergent capabilities of distributed satellite systems. A cursory examination of several systems that use multiple assets revealed a subset of these capabilities. This subset helped frame potential case studies and why leveraging those capabilities would deliver benefits to the scientific community. A more rigorous examination that combines sets of potential benefits of distributed satellites as they are published in the literature reveals a complete set of capabilities that encompasses all of the known ways multiple systems can be used in exploration. This step alone answers the first research question to a fairly satisfactory degree.

Next, the literature related to methods in value modeling, tradespace exploration, and expressing uncertain benefits is synthesized and examined. Gaps in the existing methods are examined, and modifications and additions are made to help fill these gaps.

Finally, several case studies explore aspects of potential future distributed satellite systems using the developed methods that focus on assessing unique science capabilities and lifecycle benefits of these architectures across uncertain futures. These case studies all involve multiple satellites observing multiple targets, but cover a variety of mission profiles, satellite locations, scientific targets, and instrument packages.

Most importantly, the case studies that are selected and developed in this dissertation are all concrete examples of missions that achieve well-established science goals whose results will have major impacts across many fields. These goals either *cannot* be accomplished by a monolithic system, or a monolithic alternative is technically or financially infeasible with technology available today.

The variety and breadth of these case studies is intentional; not only is it useful for empirically proving the answers to the questions posed above, but the breadth may enable readers from different science or engineering backgrounds to connect with some case studies more than others. Each case study is meant to

serve as a standalone blueprint for how an agency or team would conduct a concept study to explore mission architectures and select valuable, feasible designs.

The MATLAB code used in each of these case studies is publicly available on the appendix DVD so that any student or mission manager can look into them in more detail. It is especially important that these case studies be modifiable and upgradeable so that new models can be inserted to give higher-fidelity information about all aspects of the systems. Not only are the assumptions and coefficients to certain functions constantly changing as technology progresses, thereby changing the final results of the value proposition, but new instruments or performance metrics can be included and new value models can be added based on different stakeholders' inputs so they can see their own value proposition of the mission.

1.4 Thesis Scope

The fields of remote sensing and Earth observation have extensive previous and currently ongoing research. The value of multiple satellites in low Earth orbit pointing downward has been demonstrated in published literature and by numerous successful practical applications in operation, such as communications, military applications, weather monitoring, and ionospheric modeling. Rather than tread on ground being covered by that research, this work focuses solely on planetary science missions that do not have goals related to Earth science and operate anywhere from low Earth orbit (while looking upward instead of downward) and beyond.

The purpose of the case studies is to build many designs from the top down, not a few from the ground up. For this reason, no detailed design work will be conducted, and parametric models will be used where appropriate.

This thesis primarily focuses on how particular science goals can be achieved using a variety of technologies. There are also a number of emergent opportunities that could arise as a result of these missions. Where applicable, alternative objectives will be discussed, specifically alternate science goals that can be answered using data sets from the missions, but they will not be explored in detail beyond their qualitative impact.

The methods described herein are meant to be used *prescriptively* in order to better understand the options that are available. The goal is to understand the relationships among design variables, performance attributes, value models, and changing contexts in a mission. Specific design recommendations can be made as a result of these studies, but these decisions are necessarily the responsibility of the designer. Best efforts were made to ensure accuracy with validated modeling approaches; however, due to the breadth of this work, time constraints, and limitations on the availability of information, there is inherently some uncertainty associated with the cost, value, and performance models that cannot feasibly be eliminated, nor does it need to be in the early stages of design that these case studies intend to explore. Where appropriate, the limits of these models are recognized and described.

1.5 Thesis Structure

The remainder of this thesis is structured as follows:

Chapter 2 presents a literature review of previous work in systems engineering methodologies, value models, distributed satellites, and other work related to the core hypothesis of the thesis and all case

studies. Because of the broad scope of the case studies being used to test the hypothesis, literature that is relevant only to specific case studies will be presented in the appropriate chapters. This will not only help separate the theoretical systems engineering and value proposition formulation from the applied scenarios and specific science background, but it will also keep literature relevant to the performance modeling that is specific to each case study contained within that case study's chapter.

Chapter 3 defines and discusses the emergent capabilities of distributed satellite systems. Examples of how these capabilities exist in other terrestrial systems, how value is derived from them, and how these capabilities can be leveraged to achieve science goals that are unachievable with a single satellite are presented and discussed.

Chapter 4 presents a detailed breakdown of the primary processes and the mathematical formulation of the Responsive Systems Comparison (RSC) method and its subcomponents. Additional elements that have been developed or modified for purposes of this work that apply to one or more case studies will also be presented here.

Chapters 5 through 8 present each of the four case studies used to explore the value proposition of distributed satellites for space science missions. The formulation of the performance models will be discussed in detail, since each case study requires unique contributions to be developed before the value propositions could be explored in detail. These chapters can be read in any order, although the simplest case that could be used as an introduction to the applied methods is the first case study.

Chapter 5 presents ExoplanetSat, a mission that uses a constellation of small satellites in low Earth orbit to detect transiting exoplanets. This mission uses homogeneous satellites to achieve one science goal.

Chapter 6 presents HOBOCOP, a mission that uses a network of small satellites to study several phenomena related to the Sun. This mission uses homogeneous or heterogeneous satellites to achieve three science goals while also showing the added benefit of a federated satellite system providing communication services for space science missions.

Chapter 7 presents ÆGIR, a mission that uses a fleet of satellites to explore asteroids and characterize their water content, distribution, and composition. This case uses homogeneous and fractionated satellites to achieve several science goals related to asteroid mining and resource utilization. Due to computational limitations, only the value models will be explored in this chapter.

Chapter 8 presents GANGMIR, a mission that uses ground penetrators and a constellation of satellites to study Mars. This mission uses heterogeneous assets in different operating environments to achieve multiple science goals that must be achieved before human missions to Mars can be designed.

Chapter 9 discusses the commonalities among the case studies and some of the more important themes, as well as ways scientists or engineers could use or modify the methods in order to explore their own mission concepts in more detail before making high-level architectural decisions. The chapter ties the discussion and results back to the original problem statement and thesis objectives while reviewing the contributions to many fields that were made as a result of this work.

2. Literature Review

*“In a barren land where the dead will walk and the dragons fly
Blood across the sand, good intentions crumble as the demons rise
Welcome to my own hell, a place I know well, there’s nothing like
Hope you came prepared, because this is where angels come to die.”*

JT MACHINIMA, “PREPARE TO DIE”

Chapter 2 presents a literature review of the development of the engineering methodologies that will be applied in later chapters, previous and ongoing research in distributed satellite systems, and other topics relevant to all case studies. Unlike most dissertations, which are designed to be extremely narrow in scope to advance concepts within a single field, this dissertation is intentionally broad because it bridges major gaps between science and engineering while utilizing elements of economics, value modeling, and to some extent game theory and statistical forecasting. The scientific literature that is specific to the individual case studies will be confined to their appropriate chapters.

2.1 Systems Engineering Methodologies

Difficulties in the design of complex systems and interdisciplinary designs led to the development of the field of systems engineering. This section addresses the development of some traditional methods as well as newer methods that are at the core of the methods used in the case studies. Later sections of this chapter will delve deeper into how the methodologies work.

2.1.1 Traditional Systems Engineering Methods

The field of systems engineering emerged to address design challenges too complex for a single engineering discipline to encompass. It designates “the management of the design and development of technological systems,” and can incorporate strategies balancing business, economics, and politics in addition to engineering disciplines [7].

2.1.1.1 Historical System Engineering Methods

Bell Labs started developing systems engineering techniques in the 1930s and use of the discipline spread to everything from military products, operations research, computers and software, communications, and rocket design [8].

Manufacturing of commercial products also greatly improved thanks to systems engineering. From 1948 to 1975, Toyota developed the Toyota Production system, a socio-technical system to deliver products on time [9]. In 1972, Mitsubishi developed and used Quality Function Deployment (QFD) and the House of Quality. These principles led to superior product designs compared to American competitors [10]. QFD is a “method to transform user demands into design quality, to deploy the functions forming quality, and to deploy methods for achieving the design quality into subsystems and component parts, and ultimately to specific elements of the manufacturing process.” [11]

During the 1970s, the focus of systems engineering moved towards software development as software became more complex. Software developers starting using the waterfall model of product development,

which was used in hardware design. This model starts the design process with requirements, and then flows into design, implementation, verification, and finally maintenance. However, this model has no feedback loop, and efforts to improve upon the feedback helped create the Vee and Spiral models of systems engineering.

Concurrent engineering is a more recent product development method where spacecraft subsystems are developed simultaneously in parallel. Multiple designers are connected over a network so that information is shared in real time, usually in one room that is similar to a “War Room” for engineering design. When one designer changes or adds a variable to a document, the effects automatically cascade to the other systems and can be modeled in other areas by their respective experts.

2.1.1.2 NASA’s System Engineering Methods

NASA uses its own Systems Engineering Handbook with its “systems engineering engine,” last updated in 2010. NASA’s design process uses Vee models and Spiral models simultaneously at different resolutions of the development process guided by continuously operating technical management processes [12]. While this handbook goes into great detail describing technical management processes definitions, it is very light on the sections describing the concept generation phases of mission design. There are very few references to anything written past the year 2000, and no reference deals with new techniques in tradespace exploration. Additionally, there is usually very little money budgeted for detailed tradespace exploration for a given mission [13].

Essentially, while NASA’s own systems engineering handbook may be appropriate for project management and product realization final point designs, it is not effective at bridging the gap between the science goals of the organizations and how those flow down to requirements for an individual design. It is a handbook written for engineers, not scientists and high-level decision makers.

Since NASA engineers rapidly spiral towards point solutions, the number of design options being considered for a mission is prematurely reduced. Engineers converge on locally optimal design solutions without the tools to find a global optimum [14]. Even advanced studies, such as the ones conducted at JPL’s Team X concurrent design studio, focus more on optimizing a single point solution rather than exploring a multitude of options.

Team X works very closely with project scientists and key decision makers to quickly develop a concurrent engineering model for spacecraft, but the team focuses heavily on requirements definitions (like most engineering teams do) rather than exploring the value proposition of the mission as it relates to the scientists’ needs. The performance and cost models may be high fidelity, but the stakeholder value models are not. While many individual design trades are performed, they are not all performed in parallel and compared in one tradespace, and the number of design options is generally limited.

2.1.1.3 The NASA Project Lifecycle

Product engineering typically goes through four phases, known as CDIO – Conceive, Design, Implement, and Operate. NASA programs are split into similar phases, Phase A through Phase F. Phase A is concept and technology development. Phase B is preliminary design and technology completion. Phase C is final design and fabrication. Phase D is system assembly, integration and test, and launch. Phase E is operations and sustainment. Phase F is closeout or retirement. These phases and the many mission reviews along the way are summarized in Figure 2-1.

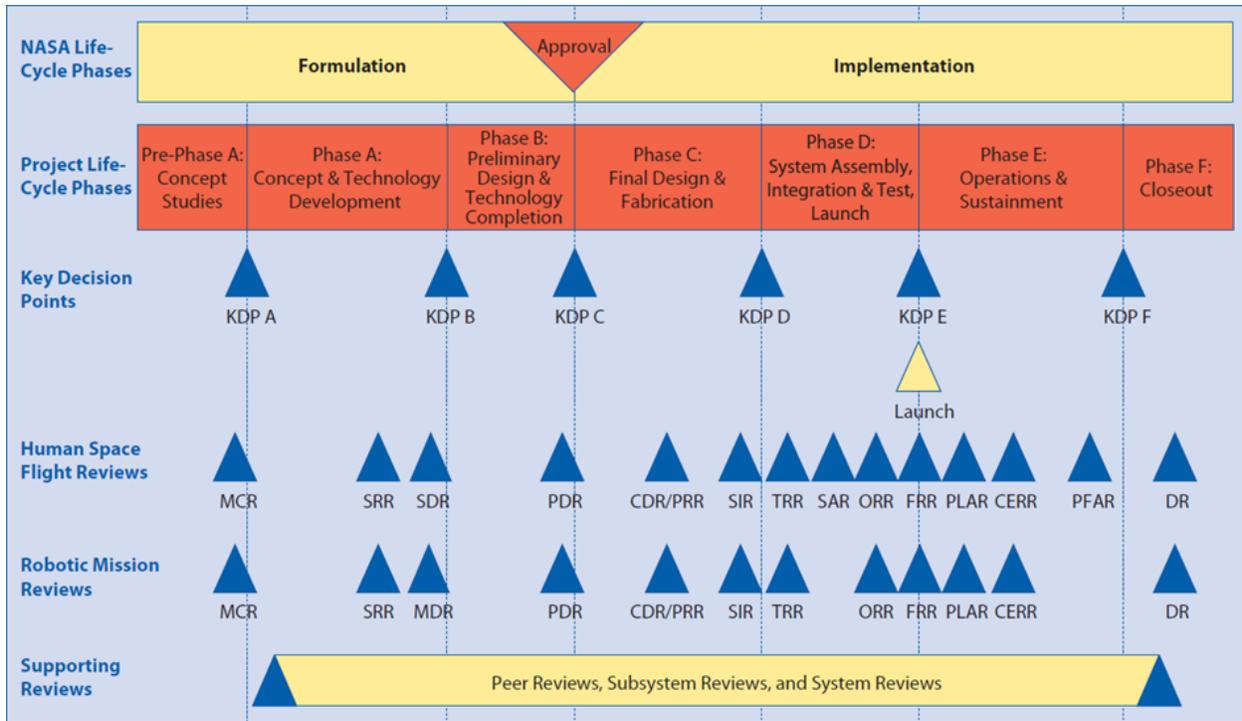


Figure 2-1: NASA Project Lifecycle (reprinted under Fair Use from [12])

Before Phase A, however, is Pre-Phase A, the concept studies phase. This is the phase where ideas for missions are compared at their basic level with trade studies. Trade studies compare a few possible architectures, designs, or technologies across a range of metrics to identify the best option. However, these studies are usually done by vastly different groups from a variety of backgrounds, often with little formal guidance. In fact, the NASA Systems Engineering Handbook devotes less than a full page to the description of Pre-Phase A activities, so it is difficult to replicate the quality of Pre-Phase A activities across many missions and concepts. As a result, very few designs are considered or evaluated, so the full range of design options is rarely explored.

New methods have been developed over the past fifteen years to solve this problem and are discussed in detail in this chapter. These practices are just now starting to take root in real design situations and have not found a way into NASA’s hierarchy. **Pre-Phase A is the phase of a space science mission’s lifecycle on which this dissertation focuses entirely.**

2.1.2 Tradespace Exploration

A tradespace is the set of all possible designs that could be considered to satisfy a customer’s needs or complete a mission, and it can be thought of as a combination of multiple trade studies wrapped into one package. While an individual trade study may examine how system performance or cost varies when one variable is changed among a few options, a tradespace examines all the possible combinations of trades that have been considered in the design space. Tradespace exploration (TSE) is the act of examining as many designs as possible in the tradespace to develop an understanding of which variable or combinations of variables create the largest changes in the value proposition, both in the value derived from the product and the cost. This section addresses the historical development of some tradespace exploration methodologies and identifies other research relevant in the field.

2.1.2.1 Development of MATE

NASA isn't the only agency struggling with the conceptual design phase of complex engineering programs. Military programs also lack the tools needed to understand complex systems early in the design process and how their value delivery can change over the course of the mission lifecycle. Prof. Daniel Hastings was the Chief Scientist of the Air Force from 1997-1999, and during his tenure he became frustrated with the need to make decisions based on too little information [15]. When he returned to the MIT Space Systems Laboratory (SSL), he created the Space Systems Policy and Architecture Research Consortium (SSPARC) to address the problem. This led to the development of Multi-Attribute Tradespace Exploration (MATE), Epoch-Era Analysis (EEA), and the Responsive Systems Comparison (RSC) methods, which are key techniques that will be used in the case studies for this thesis.

The precursor to MATE is the Generalized Information Network Analysis Methodology (GINA). GINA uses parametric models to evaluate a large number of designs by using information metrics as consistent evaluation criteria across the design space, which potentially allows for the comparison of different architectures on the same basis [16], [17]. This methodology was partially made possible by increased desktop computing speeds. GINA compares user-centric attributes to greatly simplify tradeoffs between architectures. GINA was combined with multidisciplinary design optimization (MDO) methods to fill in information gaps, which created Multi-objective, Multidisciplinary Design Optimization System Architecting (MMDOSA) [18]–[20]. MMDOSA applies simulated annealing techniques to find Pareto fronts in multi-objective tradespaces rather than individual, locally optimized point designs so that designers can understand the full tradespace better.

Multi-Attribute Utility Theory (MAUT) extends economic analysis as a metric for consumer choice from single attributes to multiple attributes [21], [22] (see Section 2.3.2.2). When MAUT was combined with the parametric model and design space evaluation aspects of GINA, it became MATE. The advantage of MATE is that it is more generalizable. Examples of early tradespace studies using GINA, MMDOSA, and MATE include: the Terrestrial Planet Finder (TPF), a distributed satellite cluster that uses the principles of nullifying interferometry to directly image Earth-sized planets; many iterations of the terrestrial observer swarm (TOS); and an orbital transfer vehicle (Space Tug) [19], [23], [24].

Concurrently, other research focused on different areas of MATE. There were three doctoral dissertations and eight masters theses based on MATE, which extended it to other areas. While Diller's and Ross's masters' theses outlined the motivation and process of MATE [14], [25], other students explored ways to apply and improve MATE. Weigel [1] explored how policy trade-offs can be quantified by proposing real options as a means to minimize designs that were sensitive to changes in launch policy or budget cuts [26]. Walton created project portfolios to manage uncertainty using MATE [27]–[29]. Derleth applied MATE to the spiral development of a small diameter bomb [30], and Spaulding applied MATE to a space-based radar system [31]. Seshasai developed the Multi-Attribute Interview Software Tool (MIST) to facilitate interviews to elicit the value of attributes [32], [33]

However, there was an increased motivation to understand how big an impact certain system lifecycle properties (also known as “ilities”) made on the value of the system over its lifecycle. In his doctoral dissertation, Ross explored the relationship between flexibility, adaptability, robustness, and scalability in aerospace systems, how they relate to unarticulated value, and how they could be quantified, evolving MATE from its “static” form into its “dynamic” form [34]–[36]. At the same time, McManus and others

devised ways of understanding and mitigating uncertainty and a framework for incorporating “ilities” into tradespace studies [37], [38].

2.1.2.2 Improvement of MATE and the Development of EEA and RSC

In 2006, Dr. Ross and former INCOSE president Dr. Donna Rhodes formed the Systems Engineering Advancement Research Initiative (SEARI) and continued to develop MATE. Around this time, they developed a new method called Epoch-Era Analysis (EEA), which leverages techniques from economic analysis to evaluate how systems can sustain value in changing contexts [39]. In EEA, the changing contexts over the mission lifecycle are considered to see if systems can sustain value over time. This led to the evaluation of passive robustness using a Pareto trace metric across multiple epochs, scenario planning, and value-weighted filtered outdegree (VWFO) as a way of quantifying flexibility [40]–[44].

Research combining EEA and MATE continued to build upon these results. Richards et al. used EEA to evaluate survivability in dynamic MATE by proposing new metrics for survivability [45], [46]. Fitzgerald and Ross developed ways to show how EEA can be used to compare sustained lifecycle value, mitigate contextual uncertainties, and assess uncertain benefits [47]–[49]. Rader used Monte-Carlo simulation with EEA to compare architectures in uncertain environments and applied this to a tool that populates subsystems in a satellite constellation tradespace [50], [51]. Ross et al. identified roles for multi-stakeholder negotiations within interactive tradespaces and helped formalize the tradespace exploration process [52], [53]. Beesemyer et al. examined several satellite constellations and how they could have been better designed to respond to changes in context across the system lifecycle [54].

Ross et al. summarized many of the value-centric design methodologies (VCDMs) and gave a guide on how to choose between Net Present Value (NPV), Cost-Benefit Analysis (CBA), and MAUT given certain assumptions in the tradespace exploration process [52]. While MAUT is an appropriate method to use when considering preferences of a single stakeholder, MAUT is insufficient when multiple stakeholder preferences cannot be aggregated because of the “nonexistence of a social welfare function” [55]. Such a function could be developed for political reasons in multiple-stakeholder negotiations between scientists to show how some missions can satisfy multiple stakeholders.

Several masters’ and doctoral theses examined MATE and EEA, and the marriage of the two methods is called the Responsive Systems Comparison (RSC) method. Schofield applied a rudimentary version of RSC to a Coast Guard cutter program [56]. Mekdeci examined survivability and pliability to achieve viable systems of systems [57]. Schaffner mathematically expressed how to conduct multi-era analysis and used it to examine many possible futures of a naval operations portfolio [58], [59]. Piña applied EEA to the selection of a distributed power generation system for homeowners [60], and Viscito applied EEA to selecting a home to purchase [44]. Schaffner, Wu, and Ricci researched affordable systems and how EEA could be used to understand how systems can become unaffordable over their lifecycle [61]–[63]. Ricci also examined evolvability options [64] and explored stakeholder value models with interactive visualization [65]. Fitzgerald examined different methods for multi-stakeholder negotiations with MATE and EEA, an area that could lead to a solution to the problem of MAUT not having a “social welfare function” when multiple stakeholders are present [66].

The Engineered Resilient Systems program (ERS) now sees TSE as a core capability for future mission planning, design, and selection within the Department of Defense (DoD) [67]. This strong endorsement

validates much of the research SEARI has conducted as it applies to complex programs, and this thesis adds to that validation by showing that it works with DSS in space science missions.

2.1.2.3 Other Tradespace Exploration Methods

Outside of SEARI, other researchers examined the value proposition of systems with flexibility and how this added value is expressed to decision makers during the tradespace exploration process. Saleh applied flexibility to space telescope servicing missions and argued that “[it is] only by accounting for flexibility that the true value of on-orbit servicing can be evaluated” [68]–[71]. Nilchiani developed a comprehensive six-element framework for measuring the value of flexibility and applied it to an orbital transportation network [72]–[74]. Joppin, Long, and Richards investigated the value proposition of on-orbit satellite servicing, satellite upgrades, and serviceability of satellites [75]–[77]

De Weck et al. showed how flexibility could be used to minimize risk when deploying a satellite constellation [78]. They argued that adding the ability to change orbits costs more in the design phase, but can mitigate risks if demand falls short of expectations, so customers could still get service globally without launching any more satellites than necessary. Further research in this area led to the development of Time-Expanded Decision Networks (TDNs) that can help apply the costs of changing designs at certain points in the lifecycle to determine how the cost of changeability in a system can affect the decisions made in the design phase [79]. Another group used similar network techniques to research how value flow mapping can be used to inform stakeholder analysis [80].

De Weck et al. also investigated relationships among many “ilities” [81]. Bailey examined a tradespace model for robotic lunar exploration [82]. Alibay evaluated spatially and temporally distributed multi-vehicle architectures for planetary exploration [83], [84]. Selva developed a rule-based system for architecting and evaluating Earth observation satellite systems [85], [86] in addition to examining how synergy between scientific data sets could add to the value of the scientific products achieved in a mission [87]. Golkar described a framework for systems architecting under ambiguity of stakeholder objectives, an important consideration when facing stakeholder needs and opinions that can change with known unknowns and unknown unknowns [88].

Other groups have also focused on tradespace exploration using Set-Based Design (SBD). This framework is built on the foundation of decision theory [89]. Ward laid the framework for these and applied a theory of quantitative interference to the design of a mechanical compiler [90]. Sobek et al. applied SBD to the Toyota Production System as a case study [91]. Malak et al. modeled fluid power systems using predictive tradeoff models [92]. Malak and Paredis have also used both parameterized Pareto sets [93] and Support Vector Machines [94] to model design concepts.

Visualization in tradespaces has also been researched significantly in SEARI as well as other institutions. Stump and Carlsen wrote their masters’ theses on multi-criteria design optimization using multi-dimensional data visualization [95] and visual steering commands during TSE [96]. This group also created the Applied Research Laboratory (ARL) Trade Space Visualizer [97] and applied it to satellite constellation design [98] and other design by shopping problems and steering-controlled visualization problems [99], [100]. Jordan et al. specifically explored optimal continuous-thrust trajectories with visual TSE [101].

2.1 Distributed Satellite Systems

A common theme of many of the projects involved in the development of MATE, EEA, and RSC is the study of distributed satellite systems for a variety of missions and stakeholders. These are not the only projects to involve distributed satellites; multiple satellites working in concert have been in use for decades. This section addresses different types of distributed systems investigated in the literature, including satellite constellations, fractionated satellites, and federated satellites.

2.1.1 Satellite Constellations

Satellite constellations have many applications, most notably in communications, and were the first type of distributed satellite system ever launched. The idea of a satellite communication network was first proposed by Arthur C. Clarke in 1945 as a way to relay signals from one side of the world to the other [102]. Since then, there have been many technically successful (though not always commercially successful) satellite constellations built for communications, navigation, and Earth science. In contrast, relatively little work has been done on constellations with astrophysical objectives.

2.1.1.1 Constellations Forms

One form of satellite constellation is a Walker constellation [103], [104], wherein the entire constellation is expressed with three numbers: orbital plane number, satellite in plane number, and mean anomaly spacing. For instance, 4/2/1 represents a total of 8 satellites in 4 planes, with 2 satellites in each plane; the satellites in each successive plane are spaced out by 90 degrees in mean anomaly (so if the first plane's satellites are at mean anomalies of 0 and 180 degrees, the second plane's first two satellites are at 90 and 270 degrees, and the third plane repeats the first plane, so this third number could be thought of as the "number of planes between repeats mean anomalies"). For a 6/6/3 Walker constellation, there are 36 satellites in 6 planes with 6 satellites per plane, and each plane's satellites are staggered by 20 degrees in mean anomaly (in each plane, each satellite is spaced by 60 degrees, and it takes 3 planes to go all the way around, so the first satellite of three adjacent planes are at a mean anomalies of 0, 20, and 40 degrees, while the next three planes would repeat that same configuration).

Not all constellations require satellites to be in similar orbits. The XM/Sirius satellite radio constellation uses a combination of satellites in both geosynchronous and highly elliptical orbits. Geosynchronous orbits such as the XM satellites have a period of one sidereal day and remain relatively fixed over a small band of longitudes during their orbits. The highly elliptical orbits that the Sirius satellites use are known as "tundra orbits" [105]. These orbits also have a period of one sidereal day but spend much of that orbit in the slow-moving high end of the orbit. Its inclination is near 63.4 degrees. Satellites in this type of orbit spend as much as 16 hours per day in the northern hemisphere, making them excellent for communications at high latitudes. A Molniya orbit (молния, Russian for "lightning") is similar but has an orbit of one half of a sidereal day.

2.1.1.2 Constellation Performance Objectives

The performance of a constellation depends on a number of factors, including the selected orbits, instrument field of view (FOV), and number of spacecraft in the constellation. However, determining how "good" or "valuable" a constellation is depends on the constellation's intended purpose. The definition of "coverage" and how it is used to evaluate a constellation can take on many forms. For communications, coverage may only be important for a certain region (for instance, Sirius/XM only covers North America), or global coverage may be a requirement. Whether a constellation is judged on its average,

minimum, or maximum revisit time depends on how critical coverage is during certain time periods. Among the many metrics that may be important include percent coverage, maximum coverage gap, time average gap, and mean response time.

According to James Wertz, “Earth coverage is not a Gaussian parameter and statistical data can give very misleading results” [106]. Although statistical tools can be used to evaluate constellation performance, they can be misleading. Wertz provides examples of how removing satellites from a constellation can inadvertently *increase* the statistical performance compared to a constellation that, when viewed with other techniques, was objectively better. Typically, there are no analytic solutions to constellation performance, and numerical calculations and simulations must be conducted to properly understand the performance and capabilities of a constellation.

There has been much research on constellation optimization for ground coverage, space-based situational awareness, and above the horizon (ATH) coverage [107]–[109]. More recently, multi-altitude ATH coverage optimization techniques have been more thoroughly developed by a group at the University of Texas [110]–[113]. However, these techniques are designed for targeting other satellites for short periods of time, not sections of the sky for long periods of time.

2.1.1.3 Constellations in Science Missions

The world’s first satellite constellation for astrophysics research is the BRiGht Target Explorer (BRITE) nanosatellite constellation developed by Canada and Austria [114]. The mission of this constellation is to study variations in brightness of the brightest stars in the sky using differential photometry. Two satellites will fly in the same orbit with different filters and observe a region of interest for 100 days or longer. The first two satellites were launched in February 2013 while another four were launched between November 2013 and August 2014 [115]. These satellites are in Sun-synchronous orbits [116], [117].

While the mission of the BRITE Constellation is somewhat similar to the first case study in this thesis, the orbital coverage constraints and photometric precision requirements are far less stringent in comparison (see Section 5.2.1). Areas of the sky that will be observed will only be targeted for a season at a time, not for an entire year. Additionally, lunar pointing constraints do not interfere with BRITE’s mission. While coverage techniques for observing parts of the sky from LEO have been developed, no techniques have been used to show how to cover any part of the sky for an entire year. Techniques like these are valuable for finding transiting exoplanets, but also have other uses in astrophysical observations, such as high-precision gamma-ray burst source monitoring [118].

2.1.2 Fractionated Satellite Systems

The cost and mass of satellite components has dropped thanks to advances in the personal computer industry, and as a result, the mass required to be launched into orbit has dropped along with total launch costs. However, satellites are still vulnerable to single-point failures, and there are no feasible or cost-efficient methods to repair satellites on orbit. This problem has led to the idea of fractionated satellite systems – satellites that are composed of smaller, heterogeneous elements that operate in close range of each other.

If only one component on one free-flying element of a fractionated system suffers from a single-point failure, only that element would need to be replaced. This is a much cheaper investment alternative compared to launching an entire new, complete system in order to continue mission operations and derive

value from the system. Additionally, elements of a fractionated system can be easily replaced with upgraded components that can extend the operational mission lifetime and increase product quality of the system, both of which greatly increase the value of the system, for a relatively small investment.

In June 2007, DARPA proposed System F6 to explore how fractionated satellite systems could change how business is done in space. Compared to monolithic systems, fractionated systems are inherently more responsive, meaning they can react to various forms of uncertainty, as demonstrated by Brown and Eremenko with the application of value-centric design methodologies [5], [119]–[121]. To handle the rich tradespace of fractionated architectures, several groups have developed tools to rapidly design and analyze fractionated spacecraft. A group at Georgia Tech created the Georgia Tech F6 Architecture Synthesis Tool (GT-FAST) [122], a group of Lockheed Martin and Aurora Flight Sciences engineers working collaboratively developed a systems value modeling (SVM) tool [123], a group at Boeing developed a competing tool called RAFTIMATE [124], and a group at MIT developed the Spacecraft Evaluation Tool (SET) [125]–[128], while others compared how effective these tools are [129]

However, after several years and over \$200 million in total investments, System F6 was canceled [130]. Part of the program’s failure stemmed from the lack of a prime contractor to integrate the work of the many companies who were awarded contracts, but also from the lack of a clear business case for heterogeneous spacecraft at the Department of Defense. These spacecraft were simply sharing data with each other in space, not leveraging emergent capabilities to deliver additional value.

2.1.3 Federated Satellite Systems

After System F6 was canceled, there was still interest in continuing distributed systems research, but instead of focusing on fractionated systems, the focus has shifted toward federated architectures [131]–[134]. Federated satellite systems are examples of true “systems of systems” in that satellites can have managerial and operational independence but are still able to work together. Some of the resources that can be traded among satellites include downlink time, data storage, and data processing.

This idea of sharing resources partially stems from the fact that larger satellites tend to not be operating with 100% of their available resources all the time. No satellite is always operating at peak performance, and the unused and available resources during nominal or lower performance periods can be sold to other satellites [135]. The terrestrial analogy to this idea is a cloud computing network, where computers are solving their own problems when it is demanded, but helping other computers solve other problems when the resources are available [136].

There are two major takeaways from the concept of federated satellite systems. The first major takeaway is that larger satellites can have additional revenue streams by acting as “suppliers”, or doing jobs that have been outsourced from other satellites. Monetary risks can be mitigated by supplying these resources to other satellites so that decision makers can hedge their bets to justify possibly riskier decisions.

The second major takeaway is that satellites that will not be suppliers can be intentionally designed and built with lower capability. For instance, smaller storage devices and processors can be used to reduce total mass if data storage and processing functions can be outsourced to other satellites. If a satellite can connect to a federated satellite network, there are more opportunities for ground stations to connect to the satellite, so the satellite would not have as many orbital constraints. These space-based resources could be used to ease design constraints, lower development costs, and decrease total mass required on orbit.

FSS is an open area of research that will be focusing on four topics: stakeholder analysis, systems architecting study, communications technology development, and innovative mission concepts [134]. The first two topics are intended to address aspects of distributed satellites that System F6 did not, which some argue is why the program failed despite so much investment and promise [131].

Communications technology development is especially important to FSS because more satellites will be communicating with each other in space, creating more noise and interference. There are also security concerns with having sensitive data shared over open space. Lluch and Golkar examined coverage optimization for opportunistic intersatellite links as well as resource balancing in federated satellite systems [137], [138]. The final topic, innovative mission concepts, is an area that this thesis partially intends to address, since there is the potential for monetary gain to come from networks of federated satellites aiding in the operations of science missions (see Chapter 6).

While some of the research discussed so far in this chapter has focused on how multiple assets can be distributed spatially and temporally at lower costs to achieve the same goals, none have provided a methodical approach to focus on what can be accomplished as a result of the emergent capabilities of distributed satellites.

2.1.4 Design Optimization in Tradespace Exploration

Tradespace exploration itself is not an optimization process. As the name suggests, it is meant to explore many available options early in the design process so that viable solutions can be weighed against each other. Design optimization has typically been saved for later in the design process, but as computers have become more capable, optimization techniques are being utilized earlier. This is especially true in complex multidisciplinary systems. Multidisciplinary design optimization (MDO) is a relatively new field in systems engineering that can be used concurrently with TSE within performance models.

Ideally, TSE narrows down the list of design variables that are explored so that variables that do not affect the value proposition are held constant. However, there can be thousands of variables in a performance model, and showing a meaningful tradespace is dependent on accurately characterizing what these other variables could be. Optimization can complement TSE by calculating the optimum values for the remaining variables. This is especially important in space missions, where mass and cost are expected to be minimized and objectives and performance are to be maximized.

There are many optimization methods used in mathematics and engineering. One particular area of optimization that has had success in MDO is heuristic methods. Simulated Annealing (SA) is one heuristic method that has been used frequently. SA has a good chance of finding the global optimum because it explores a wide range of values within the bounds of the problem rather than quickly converging on a local optimum that may be dependent on the initial guess like some analytical optimization techniques [139]–[141]. This is an especially important property when local derivatives are unknown or when the data is not continuous. Jilla demonstrated the effectiveness of simulated annealing coupled with MDO in constellation design problems [18]–[20].

Particle Swarm Optimization (PSO) algorithms [142] and Genetic Algorithms (GA) are two other classes of heuristic optimization techniques that have been used in design problems [143]–[145]. However, when data is exclusively discrete, with no continuous variables, the Nelder-Mead Simplex method has been shown to be a fast, effective alternative to these cumbersome alternatives [146], [147]. This method finds

minima by performing a series of pivots within a multi-dimensional matrix. This method was used on an earlier version of a heliocentric satellite network optimization problem that was a precursor to a case study in this thesis [148].

For multi-objective optimization, an effective technique that has been put into practice is Normal-Boundary Intersection (NBI) [149]–[151]. This method is effective for finding the Pareto front for multiple objectives, and is especially well-suited for problems with concave Pareto fronts compared to weighted sum methods. Other techniques to generate Pareto fronts of data sets include the normal constraint method and the adaptive weighted sum method for bi-objective optimization [152], [153]. However, NBI becomes cumbersome when considering more than two objectives, and is not ideal for TSE as a standalone tool because it throws out suboptimal designs that could otherwise be studied to learn more about how the tradespace varies with design variables.

2.2 Satellite Communications

Satellite communications subsystems play a vital role in the tradespace exploration process when considering distributed satellite systems for use in science missions conducted beyond low Earth orbit. Generally, the primary product that satellites deliver is data, unless the satellite is on a sample return mission; and even then, the satellite is in communication with operators on the ground. Over long distances, the maximum amount of data that can be sent back to Earth may limit the total value a system can deliver. Understanding the sensitivity of the value proposition to the communications system model may drive key architectural decisions early on, and demonstrating how miniaturized technologies can lower overall mass is especially important when assets are produced in larger numbers.

2.2.1 Communications Subsystems Technologies

Radio communication has been used on every satellite ever launched, but because radio is such a mature technology, the physical limits have been pushed on the data volume that can be transmitted via radio frequencies. Although there are many different types of antennas to choose from, most deep space missions choose parabolic antennas because they provide the highest gain of any shape relative to its size. These are usually large; the specific size depends on how far away the spacecraft is from Earth and what the data rate requirements are. Gilchrist’s sizing model for communications systems has been used in parametric sizing models for decades [154], although recent work by Lluch and Golkar has provided more modern estimates [138]. Generalized sizing for communications subsystems on small spacecraft is especially complicated because small satellite manufacturers are hesitant to publish or reveal their hardware capabilities.

However, as the size of the antenna grows wider than the faring of the launch vehicle or as satellites get smaller, they cannot feasibly or affordably use rigid parabolic antennas. If this is the case, there are essentially three other options: deployable antennas, inflatable antennas, and optical communication.

2.2.1.1 Deployable Antennas

Satellites that operate far from Earth, such as interplanetary missions that require high data rates, may chose a communications architecture that includes larger, deployable antennas. The diameter of these antennas can be much larger than the maximum diameter allowed by the launch vehicle, but there are risks associated with this choice.

The Galileo satellite (a satellite to study Jupiter, not to be confused with the ESA navigation constellation of the same name from a previous section) carried a large deployable antenna that failed to deploy. Engineers concluded that the delay in the launch caused by the Challenger disaster allowed the lubricants on the deployment mechanisms to dry out, and only 15 of the 18 ribs on the antenna deployed [155]. As a result, data transmission was reduced from 134 kbps to 16 bps via the low-gain antenna. Upgrades to the Deep Space Network (DSN) raised this rate to 160 bps, and advances in data compression techniques raised the effective data rate to 1 kbps [156].

Small, deployable antennas have received some attention in the CubeSat world because they are lightweight and cheap. High-gain, deployable Ka-band antennas are currently under development [157], but the ones that are commercially available only provide a few dB in signal gain. These antennas are useful for communicating with small satellites in low Earth orbit (LEO), but other options are needed for small satellites that are far away from Earth such as ones conducting interplanetary missions.

2.2.1.2 Inflatable Antennas

Inflatable antennas have been tested in space aboard the Space Shuttle. In 1996, an inflatable antenna 14-meters in diameter was deployed and recovered [158], [159]. Although it deployed in a dynamically uncontrolled manner, it did not require external forces (such as a human on an extravehicular activity) to deploy it. The satellite used canisters, springs, and other moving parts to deploy the antenna. While this system suffices for large satellites that require a 14 m antenna, these mechanical devices would be difficult to miniaturize for implementation on small satellites.

Little work on inflatable antennas had been done since then until recently, when an MIT team designed, built, and laboratory tested a 1 m diameter inflatable antenna for CubeSats [160], [161]. The antenna inflated using sublimating powder rather than compressed gas. This antenna requires no moving parts other than a gate to open the backside of the satellite so the antenna can inflate. The structure was constructed with Mylar, weighs less than 1 kg, and fits into an area smaller than a 1U cube (10 cm x 10 cm x 10 cm). This antenna was tested in a vacuum chamber at MIT's Space Propulsion Laboratory to demonstrate how it would inflate and at the Jet Propulsion Laboratory in an anechoic chamber to test its signal gain. Both tests were successful and one configuration achieved a gain of 15 dB at 2.4 GHz.

2.2.1.3 Laser/Optical Communications

The concept of deep space laser communication has existed since the 1980s [162], [163], but had never been demonstrated in space until October 2013 when the Lunar Atmosphere and Dust Environment Explorer (LADEE) transmitted data back to Earth via a laser from Lunar orbit [164], [165].

Subsystems models are essential for modeling the complex dynamics of a spacecraft and how it operates. Plenty of incredibly detailed resources have been published for modeling traditional communications subsystems [166], but there is very little that has been published which unifies all the concepts of laser communications into a single subsystems model. For instance, although Chen elaborates on link design [167], he references two other papers by Biswas and Piazzolla [168] and Moision and Hamkins [169] to fill in the details. While these two papers describe different aspects of laser communication (the link and the modulation) very thoroughly, they are not reconciled at the most basic level.

Additionally, other resources fail to succinctly wrap all the high-level details of a laser communications link model into one package that delivers actionable information to a systems engineer. Summary articles

found in the literature search do not deliver necessary details for calculating link budgets [164] and other works also fail to deliver what the systems engineer needs for modeling mass and cost [170]. Even a master's thesis published by the U.S. Navy glosses over important assumptions and details [171]. Developing a subsystems performance and cost model will give systems engineers the opportunity to understand how choosing an alternate form of communication can affect the overall mission design.

2.2.2 The Deep Space Network

Regardless of the mission, most interplanetary spacecraft rely on the Deep Space Network (DSN), a collection of 35 m and 70 m antennas around the world to receive signals from most of the spacecraft still operating beyond cislunar orbit [172]. However, this network was built in the 1960s and is heavily subscribed; the DSN operates above 95% capacity every month, and it usually operates closer to 99% capacity. If more spacecraft, especially distributed spacecraft all trying to communicate with Earth, are launched into interplanetary space, the extra required time (and cost [173]) on the network not only hurts programs trying to get added to the schedule, but detracts from the overall amount of data the world can receive from these spacecraft.

There are currently no plans to update the network and add more dishes, so this bottleneck could affect *all* science missions in the future unless the network is expanded or space-based resources can be used to enhance our ability to communicate with more satellites. How this is achieved is relevant to the future of how satellites will consider their communication hardware. If ground-based laser receivers are built, future spacecraft can trade smaller hardware and higher data rate for higher pointing stability and more complex transmitters. Additionally, space-based laser receivers in LEO acting as relays have the advantage of less noise from atmospheric absorption and more availability because of the lack of weather constraints, so they can achieve higher data rates and relay that to the ground using radio frequency links.

Regardless of what frequencies are used, there have been many groups working on technologies and designs necessary to have a satellite internet in LEO [174]. Less work has been done on the development of an interplanetary internet [175], although both are essential to operational FSSs. Future science missions using multiple satellites would be able to tap into federated satellite networks to reduce their communications subsystem mass.

2.3 Science Mission Value Modeling

The value proposition of distributed satellite systems for space science missions is not only the title of this dissertation, but quantifying and expressing that value proposition is its main objective. This section gives background on how missions are selected and the options for expressing the value of the mission quantitatively, not only so they may be compared more fairly to monolithic systems, but also to show how new ways to leverage the emergent capabilities of distributed satellite systems can be used to explore new scientific vistas that had not been explored previously.

2.3.1 Mission Selection Processes

Every ten years, the United States Research Council publishes the Planetary Science Decadal Survey for NASA and other government agencies [176]. This survey identifies possible missions and what can be accomplished with interesting scientific objectives and splits them up into Small, Medium, and Large missions. While it is up to committees and high-level decision makers to decide which missions will accomplish the science goals discussed in the latest decadal survey, the survey itself only addresses

missions; it does not seek to resolve unanswered fundamental science questions. With monolithic systems, these two are usually so closely related that it is difficult to perceive the difference. However, when considering distributed satellite systems, the emergent benefits of one mission may help answer questions in another field.

Fundamentally, the value of science missions comes from answering unanswered questions, so missions that can answer high-impact questions in a wide variety of fields are the most valuable. A mission manager could propose a valuable mission that leverages the emergent capabilities of distributed satellite systems to answer science questions that have never been tackled before because these emergent capabilities were previously unknown, and as a result the tools for planning missions with multiple satellites were unavailable.

2.3.2 Value-Centric Design Methodologies

Value-centric design methodologies (VCDMs) attempt to quantify the “value” or “benefit at cost” of a system. Depending on the system and the “product” that is being evaluated, decision makers can choose a particular methodology to measure value and compare similar designs.

2.3.2.1 Economics-based VCDMs

Typically, companies manufacturing commercial products will use Net Present Value (NPV) to determine the value of a product [177]. NPV quantifies the net cash flow and discounts the money spent or gained in the future by a percentage per unit time, usually matching the rate of other potential investments plus inflation. This is perhaps the most straightforward way to measure how successful a product can be, especially when investors care more about the return on investment compared to other investment options over the product lifecycle.

Other approaches may be best-suited by Cost-Benefit Analysis (CBA) [178]. CBA quantifies the net benefits produced by a system relative to the net costs. It is well-suited to weighing positive and negative effects of various outcomes and combining them into a single metric. However, since the costs and benefits are compared as a single monetary metric, the benefits associated with the system must be transformed into monetary units. While this may work well in some settings, e.g. public works projects where the benefits are publicly visible or translate directly to economic impact, it is not suited for situations when the returns are not easily monetized. This is especially true in cases where cost is allocated by a different person or stakeholder set than by whom the benefit is perceived, such as in space science missions funded by governments using citizens’ tax dollars. However, NPV and CBA are most applicable when directly comparing dollars to dollars; it is not easily applicable when the returns on investment cannot be easily monetized.

It is sometimes difficult to quantify the value of information, or how much a stakeholder or decision maker is willing to pay for information before making a decision, that is gained through space missions [179]. Commercial space companies can use dollar-based value models like NPV because they have natural monetization in their business cases. For some missions, there can be a direct dollars-to-dollars comparison between cost and benefit, such as with weather and remote-sensing satellites, which are able to inform many stakeholders about future patterns so they can make decisions on how to react in certain situations.

When exploring beyond Earth, it is very difficult to use economic gain alone to justify a mission because the product from that mission is generally only knowledge, but that knowledge may come with the prospect of economic gain in some cases or risk mitigation in others. Asteroid prospecting and solar monitoring missions are examples of science missions that have both scientific and economic value propositions. When proposing such missions, it is important to quantify, when possible, the economic value proposition using existing techniques for quantifying the value of the information derived from said missions [179]. However, other public works, whether they are ground-based or space-based, tend to not be able to use NPV alone and must make cost-benefit and value of information comparisons using some other method.

2.3.2.2 Utility-based VCDMs

For more general situations when benefits are not necessarily measured in dollars, utility theory enables numerical quantification of benefits or attributes of a system for a given outcome. By definition, *utility* is a stakeholder perception of a product's ability to satisfy preferences or needs. Over the past 50 years, utility theory has been used as a tool for consistent decision making in the face of uncertain outcomes. Due to the importance of utility-based VCDMs in this work, more space is dedicated here to explaining these VCDMs in detail for readers who are not familiar with utility theory or value modeling.

Single Attribute Utility

Von Neumann and Morgenstern created the first operational single-attribute utility (SAU) functions by leveraging game theory in economics and outlining some simple assumptions [21]. These assumptions include:

1. The decision maker knows what he/she likes (existence of a preference and indifference).
2. The decision maker's preferences are transitive (if A is better than B and B is better than C, then A is better than C).
3. If a decision maker is indifferent or equally happy between two options or outcome, he/she is also willing to substitute one for the other in a lottery.
4. The decision maker will always accept a lottery between the best and worst outcome over a sure intermediate outcome if the probabilities are properly adjusted (Archimedean property).

The last assumption holds for any decision that could result in many possible outcomes, even expensive space systems with single-point failures. The von Neumann-Morgenstern utility functions capture the dependence of value of a system to a particular attribute, so as the level of that attribute changes (and therefore the outcome), the perceived value, or utility, of the system changes. These functions assume monotonically increasing value with increasing attribute levels.

Unlike with NPV and CBA, where "more" is always "better" when money accumulates, utility theory assumes there are upper and lower limits to how much value a particular attribute of a system can be perceived to deliver. The "minimum acceptable" level of an attribute gives a utility equal to 0, whereas the "maximum perceptible" level gives a utility equal to 1. For instance, if one were buying a car and an attribute that delivered value was "Top Speed," one would perceive some minimum top speed a car must be able to reach (such as a highway speed limit). However, most consumers rarely drive in excess of 100 mph, so a car buyer would perceive no difference whatsoever between a car with a top speed of 120 mph versus one with 160 mph, so if all other factors were the same but the buyer needed to spend more for the faster car, the buyer would choose the slower car to save money. Examples of these von Neumann-

Morgenstern utility functions and the different ways a stakeholder may perceive the growth in utility as an attribute's level changes are shown in Figure 2-2.

Most products, especially complex systems, deliver value based on many attributes and their associated utilities, not just a single attribute. However, it is difficult for humans to make decisions based on many metrics simultaneously. This is a fundamental limitation of human cognition, and decision makers will tend to gravitate toward simpler metrics. This is one reason why NPV and CBA are more appealing and used extensively in business; there is only one output to consider that benefits most stakeholders in the same measurable way (dollars). The challenge is how to combine multiple value metrics together without making bad assumptions or skewing the interpretation of the data.

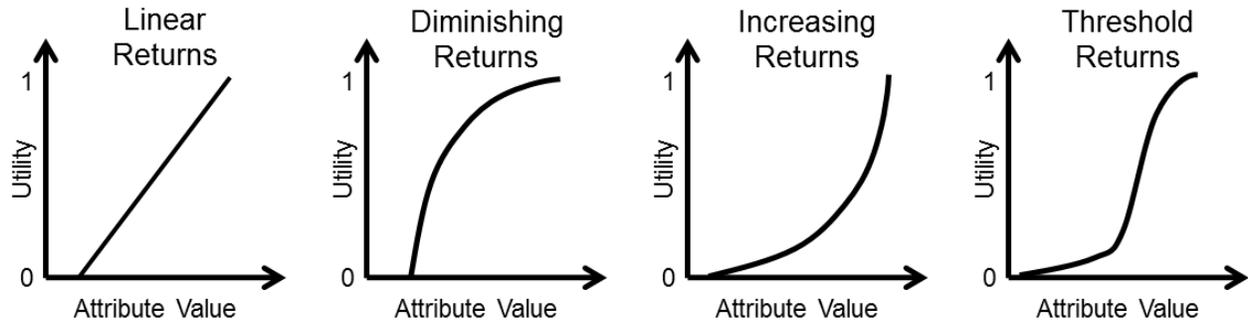


Figure 2-2: Examples of monotonic single attribute utility curves. Left to right: linear, diminishing, increasing, and threshold returns.

Multi-Attribute Utility

Keeney and Raiffa expanded the work of von Neumann and Morgenstern to develop multi-attribute utility theory (MAUT) [22], [180], which aggregates the stakeholder attributes based on the relative benefits perceived by each attribute to rank the options on a single scale. They make two key assumptions about stakeholder decision making with multiple attributes:

1. *Preferential Independence*: that the preference of the decision between (X'_1, X'_2) and (X''_1, X''_2) is independent of the levels of X_3, X_4, \dots, X_N .
2. *Utility Independence*: that the shape of one utility curve $U(X_i)$ is independent of the levels of the other attributes $X_{j \neq i}$

The governing equation of MAUT is

$$K \cdot U(\hat{X}) + 1 = \prod_{i=1}^n (K \cdot k_i \cdot U(X_i) + 1), \quad K = -1 + \prod_{i=1}^n (K \cdot k_i + 1) \quad (2-1)$$

Here, $U(\hat{X})$ is the aggregate utility value of multiple attributes and their respective utilities, $U(X_i)$; k_i is the i th corner point, which is a swing-weighting factor for the i th attribute, X_i ; n is the total number of attributes; and K is a normalization constant. The MAUT function $U(\hat{X})$ scales arbitrarily but is commonly expressed on a dimensionless scale from 0 to 1. This scale is an interval scale, meaning the relative distances between points are meaningful but the absolute distances from an origin and ratios of points are not. For example, a design with a MAU of 0.8 is objectively better than designs with MAUs of

0.79 and 0.75, and the relative distances 0.01 and 0.05 carry meaning about how much more the first design meets stakeholder expectations than the other two.

However, as a consequence of the arbitrary scaling, there is no absolute zero point, so comparisons based on ratios cannot be made. For example, a design with a MAU of 0.8 is not “exactly twice as good” as a design with a MAU of 0.4, though the distance between the two points provides useful information as described in the previous paragraph. This limits MAUT’s ability to be used to compare different designs using ratio comparisons, although addition and subtraction comparisons are valid.

One key property of the multiplicative formulation of MAUT is that designs that are weak in one attribute are penalized more than designs that meet all attributes relatively well. This prevents designs with major deficiencies in any attributes from being ranked higher. Weighted sums and other additive aggregation methods (discussed in Section 2.3.5.3) would unfairly hide designs that are well rounded compared to ones that excel in most attributes but fail in others, especially when the stakeholder is unwilling to risk low satisfaction on any attribute in particular (though MAUT can also accommodate cases where the stakeholder perceives high satisfaction when only a few attributes are highly satisfied)

MAUT is the VCDM that is typically chosen due to its ability to rank systems based on their aggregate benefits relative to predetermined attributes. These attributes are based on stakeholder preferences and are characterized by the MAUT metric with attached uncertainty so that probabilities can be incorporated into scenarios.

Value, as interpreted by MAUT, is “the aggregation of benefits relative to the monetary cost of obtaining those benefits” [55]. MAUT is used in MATE and MATE-CON because it is considered the most appropriate existing VCDM for space science missions (this will be discussed further in Chapter 4).

MAUT can be used as a consensus building tool between decision makers where consequences can be easily visualized and understood. However, a consequence of MAUT’s formulation is that the preferences of multiple stakeholders cannot be aggregated into a single MAU because there is no absolute, universal scale for utility. For example, two stakeholders working on the exact same mission may have completely different perceptions on the minimum and maximum attribute levels, attribute functions, and attribute weights. One stakeholder’s definitions may mean that a particular design scores 0.1 while the other scores the same design as 0.9; even if the attribute weights are the same, the first stakeholder clearly has higher minimum acceptable levels, higher maximum perceivable levels, or both. The differences can be magnified when science missions address fundamentally new objectives; a MAU score of 0.1 for such a mission may be revolutionary according to one value model, whereas a well-studied objective may simply experience “more of the same” with anything less than an MAU score of 0.9.

Multiple stakeholder preferences can be aggregated with a social welfare function, but it would make substantial assumptions that rank the importance of the stakeholders and how utility is perceived across individuals. An active area of research involves stakeholder negotiations using MAUT, and the cases studies in this thesis may also be useful for that research as well [66].

Key Assumptions and Limitations

MAUT makes several key assumptions that add limitations in its implementation. One limitation of MAUT is the selection of attributes to be ranked. Attributes are the validated set of criteria that reflect the

“perceived goodness” of an alternative to a decision maker. Keeney and Raiffa describe a set of attributes as having five properties: complete, operational, decomposable, non-redundant, and minimal [22]. If the chosen attributes are not perceived to be independent of each other (meaning the level of one influences the level of another), or if the attributes are otherwise poorly chosen or difficult to decompose, the stakeholder interview process can be much more cumbersome and the final output may poorly represent how the stakeholders feel about the designs, leading to an inaccurate representation of the value of the system. Choosing good attributes up front is a key to ensuring the product of this VCDM is worthwhile. If necessary, dependent attributes can be redefined as a single, coupled attribute to encapsulate those dependencies into a higher level of abstraction.

The output from MAUT is a score on a dimensionless scale that is more difficult to communicate to decision makers than the output from other VCDMs with easily understood metrics like dollars. As a result, communication of what the ranking means to decision makers can be particularly difficult, especially since MAUT does not *directly* quantify the value of the system; rather, it quantifies the perceived utility of the system. This limitation is especially apparent when considering system lifecycle properties in the design space; a flexible system usually costs more to implement and gives no direct advantages if external variables do not change, so there is no incentive to be flexible. This is why MAUT is usually used as a screening process to show the “good” alternatives to a decision maker in terms of each design’s attribute scores; it is a decision-making tool, not a decision maker.

Finally, MAUT only captures a static snapshot of the stakeholder preferences and perception of value and it assumes that, over time, the preferences of the stakeholders do not change. Time is not explicitly addressed in MAUT, and different stakeholder preferences under different conditions or evolving preferences in time are not considered. This limitation is especially severe for spacecraft development because of the long time scales that are usually required and how quickly the context of the environments (e.g., political funding situation and financial market growth rate) can change.

These last two limitations are especially detrimental when exploring architectures with multiple satellites. Multiple satellites provide inherent desirable system lifecycle properties such as redundancy, flexibility, and reliability, but these properties alone do not contribute directly to the value proposition of the system in the context of the stakeholder’s current preferences or the static environment in which the system is being considered to operate. As a result, the added expense required to achieve these lifecycle properties results in no easily measured gain in static value.

2.3.2.3 Alternative VCDMs

In addition to NPV and CBA, some other alternative VCDMs include Cumulative Prospect Theory (CPT), Value Functions (VFs), Analytic Hierarchy Process (AHP), and Technique for Order Preference by Similarity to Ideal Solution (TOPSIS). However, none of these are as appropriate as MATE for exploring space science missions. CPT is similar to MATE but uses descriptive rather than quantitative stakeholder preferences and requires monetized value similar to NPV and CBA [181]. VFs are additive in nature and assume *mutual additive independence*, which is a much stronger assumption than MAUT’s *mutual utility independence* [22]. This means MAUT penalizes the effects of one attribute’s shortcomings by multiplying the whole by a small number, whereas VFs may hide some attributes’ shortcomings if other attribute values are large.

AHP and TOPSIS rely on ad-hoc assignment of weights for attributes, whereas MAUT uses stakeholder preferences to map these weights [182]–[184]. AHP uses pairwise comparisons of alternatives and the pairwise eigenvector method to obtain attribute weights and provides a cardinal ranking similar to NPV and CBA without requiring benefits to be monetized. However, its methods for deriving value can be very complex and quickly lose transparency, making visualization difficult, and this method becomes unrealistically time-consuming as the number of alternatives grows. TOPSIS has a nearly identical output as MAUT (ordinally ranked alternatives) but also suffers from mutual additive independence like VFs, NPV, and CBA.

2.3.3 Value in Time

Previous research has shown that what value DSS may lose in extra costs due to redundant subsystems, it can make up over the course of its mission lifecycle [5], [128]. Several tools have been developed to help analyze value over the entire system lifecycle that show the added value of changeability and other emergent system lifecycle properties within DSS better than VCDMs alone.

2.3.3.1 Epoch-Era Analysis

Epoch-Era Analysis (EEA) is a structured way to consider the different possible scenarios when certain contextual variables that were once considered unchangeable parameters can change. It was developed to revisit the assumptions in MATE, forcing designers and decision makers to consider contextual and needs-based uncertainty in both the short- and long-term up front rather than as an afterthought later in the design process. Usually, such contextual uncertainties are considered in the sensitivity analysis phase of mission design; at this point, time and funding may be limited so these processes may be performed only at a high level or skipped entirely. Once again, knowledge generation at the front-end of the design is a key motivating factor for using EEA.

Introduction to EEA

Much as MAUT has *stakeholder expectations* and MATE has *design variables*, EEA uses *context variables*, or variables that the designer does not have direct control over, to enumerate all the possible scenarios for the environment in which the system will perform. A time period wherein all context variables and stakeholder expectations are constant is called an *epoch*. The entire mission lifecycle, which may be built up by one or more epochs, is called an *era*. While static MATE without EEA only considered the baseline epoch and assumed it lasts for the entire mission lifecycle, MATE with EEA assumes that events can occur that change the context.

In addition to changing context variables, each epoch will have its own set of stakeholder satisfaction criteria and expectations so that changes in value delivery from one epoch to the next can be considered. If either a context variable or an expectation changes, it results in an *epoch shift*. This is especially important given that stakeholder preferences may change based on changing conditions in which the product will be operating. This quality of EEA wipes away the primary limitations of MAUT and MATE as static snapshots and enables them to operate dynamically.

The set of all epochs is the combination of all contexts and sets of expectations. Once the epochs are generated, they can be combined in sequence to build an era. Since there is some logical progression for how some contexts may come about, rules need to be in place to prevent illogical context jumps from one epoch to the next. A mission lifecycle era can be partitioned into epochs, and how the expectations and utility delivered over time can change is shown in Figure 2-3.

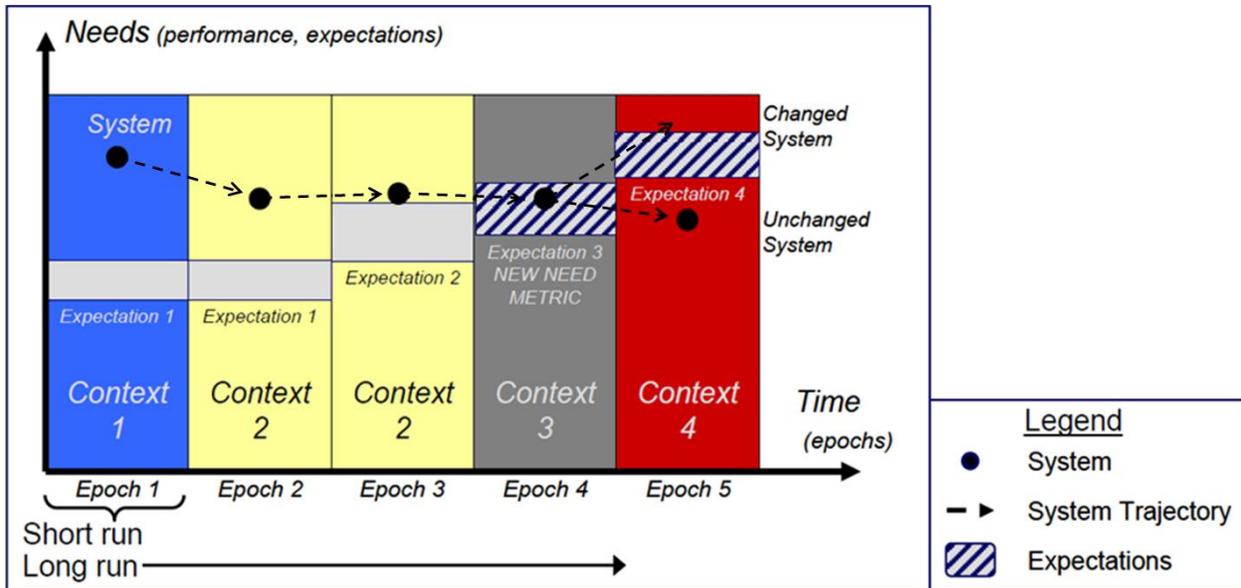


Figure 2-3: Partitioning short run and long run into epochs and eras (reprinted with permission from [39])

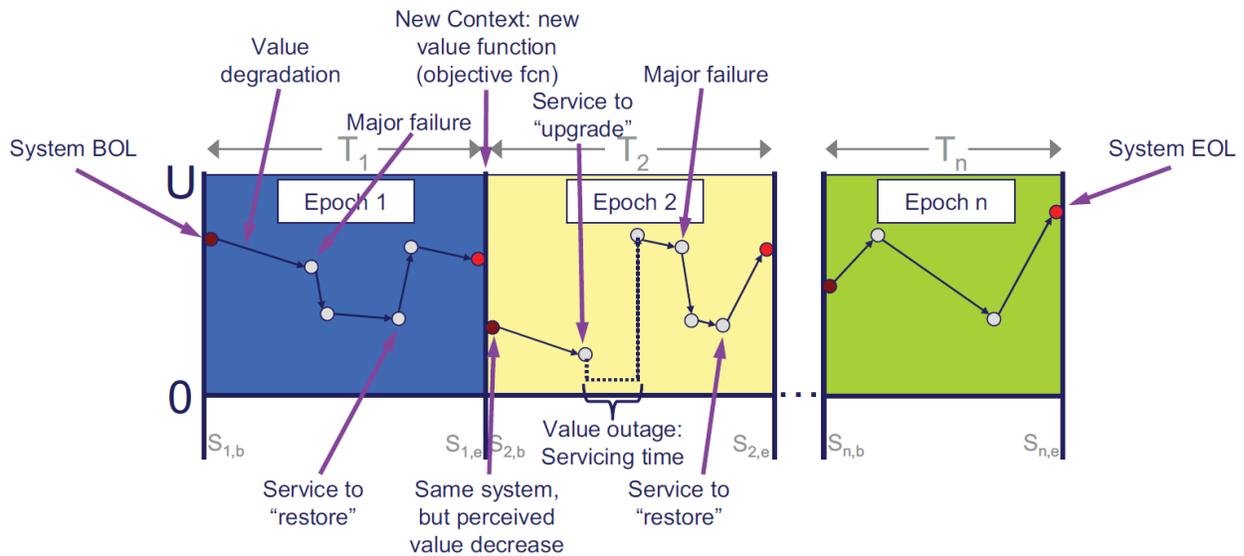


Figure 2-4: Example epochs in an era for a serviceable satellite system (reprinted with permission from [185])

Perturbations that can affect the mission are divided into two categories: *disturbances* and *shifts*. Disturbances affect the design but do not change the epoch (context or expectations). Shifts affect the epoch (context or expectations) but do not change the design. An example of how different disturbances and shifts affect a design over the course of its lifecycle is shown in Figure 2-4.

EEA can give information to the designers about which systems are not just passively robust but also how to incorporate changeability into the mission lifecycle. For instance, if an event occurs and the context changes and some subset of designs have the option to make a change, this change and the cost to make the change can be considered in the process. This way, lifecycle costs with possible changes or planned upgrades can also be analyzed in more detail to accurately capture the value proposition over time.

Additionally, MATE with EEA has the potential to be viewed from multiple levels. While most mission managers will only need to consider mission-specific parameters and variables, higher-level decision makers need to consider the context in which a certain mission affects an entire program or agency. When multiple levels are considered, strategic decisions can be made in the planning stages of one mission that could affect other missions decades from when the first mission launches. Distributed satellites have the ability to carry resources that can be used in federated networks, and future missions can leverage those resources to save costs in the long run. EEA is a framework that can be used to better express the cascading long-term effects better than a typical trade study.

Activities of EEA

There are four primary sets of activities within EEA: Single Epoch Analyses, Multi-Epoch Analysis, Single-Era Analyses, and Multi-Era Analysis.

Single Epoch Analyses are similar to traditional TSE, but the only difference is that now there are up to $n \times m$ tradespaces to analyze, where n is the number of unique contexts and m is the number of sets of stakeholder expectations. Common activities include searching for value driving variables and analyzing the Fuzzy Pareto Number (FPN) [186].

Multi-Epoch Analysis considers how designs perform over many contexts. One useful tool in this activity is Normalized Pareto Trace (NPT), which reflects the percentage of all epochs for which a given design is on the Pareto front [41], [44]. Additionally, not all epoch need to be viewed equally, and weights can be applied depending on the probability of an epoch occurring or other metrics to skew the NPT to reflect a more accurate vision of all epochs.

Schaffner separates “*Era Construction*” as a separate event in EEA between epoch and era analyses, so although this step lays the groundwork for era analyses, it does not generate useful data on its own [58]. Constructing eras can be done manually (especially for single-era analysis) or by Monte Carlo simulation or other randomized methods (required for multi-era analysis). Some contextual changes would also end the mission early, so the length of the system lifecycle itself is also an uncertainty that is inherently considered with sound design of rules. Each era of each design has some net utility as a result of the system lifecycle that can be approximated as the sum of the value delivered across all epochs in the era.

Single-Era Analyses examine individual eras of designs to see how the utility of a design changes over the course of some possible mission lifecycle. Designs that are on the Pareto front in one epoch may be unacceptable in others, whereas designs that were far enough away from the Pareto front to be hidden in both single-epoch and multi-epoch analyses but have changeability options built in may rise in ranking. Single-era analyses will also show which epochs are particularly challenging for all designs.

Multi-Era Analysis examines many eras of many designs. These eras are constructed automatically with Monte Carlo simulation. Schaffer provided the first detailed explanation and application of multi-era analysis in his master’s thesis. In this process, there are four required inputs:

- A set of design vectors \mathbf{X} , where each X is a unique configuration of design variable levels.
- A set of epoch vectors \mathbf{E} , where each E is a unique configuration of contexts and expectations.
- A set of attribute vectors \mathbf{Y} , where each Y is the evaluated performance model output for each element in the cross product $X \times E$.

- A set of value vectors \mathbf{V} , where each V is a value model output for each attribute vector Y in \mathbf{Y}

There are three optional inputs, at least one of which must be present:

- A set of eras \mathbf{R} , where each R is an n -tuple of frames. A frame r is a 2-tuple of an epoch and its duration: $\langle E, t \rangle$. The number of frames in an era n determines the length of each R
- A set of perturbations \mathbf{P} , where each P is an operator on one or more of the following: design variable levels, performance model, context variables, and/or the value model.
- A set of change options $\mathbf{\Delta}$, where each Δ is a path enabler providing edges on a tradespace network that can link changeable designs and track the costs associated with change options.

And one more optional input:

- A set of tactics \mathbf{T} , which are epoch-level transition strategies defined in VASC (e.g. maximize efficiency, survive) if era-level strategies are being studied [49].

Schaffer identifies seven activities that can be performed within multi-era analysis: identify desired variant, select designs and epoch-level metrics of interest, create era-level strategies, identify design-strategy pairs of interest, create eras, design-strategy-era evaluation, and results analysis [58]. He also develops several meta-metrics that can be applied, including expedience, variability, greatest instant change, and range. Strategies include leveraging changeability options to minimize lifecycle cost, maximize lifecycle efficiency, and survive, and these influence the set of tactics \mathbf{T} .

Key Assumptions and Limitations

One of the key limitations to how effective EEA can be, specifically in era and multi-era analysis, is that knowledge about the future is required to do the scenario planning. In order for a particular context to be considered in multi-era analysis, designers must first consider what variables could change and what the probability of that change occurring is. There is room for error when estimating both the probability that a context will change and the impact that change will have on the result. However, if an event's probability of occurring is well understood, EEA is an effective way of applying statistical analysis to the lifecycle in order to characterize the expected value better. Without such probability data, all epochs are treated equally. While this is sufficient for many cases, designs where evolution and flexibility are highly desirable may be sensitive to changes in perceived probabilities.

An obvious limitation of multi-era analysis is that there are *infinite* possible eras; there is absolutely no way to sample even one billionth of a percent of all possible futures. The idea is that hopefully by enumerating many possible scenarios, the likely ones will emerge to show the value added by changeable designs.

Combining MATE and EEA can be a computationally expensive process. Each design needs to be considered in each epoch, so the total number of design-epochs that needs to be analyzed is the full factorial of all the possible design variables and context variables. For instance, the simple Trucklandia study from SEAr's Tradespace Exploration Working Group explored 31,104 possible designs and 96 possible contexts, resulting in 2,985,984 possible epochs [187]. Once these are enumerated, era generation can finally begin, which can be an even more computationally intensive process depending on how many eras are generated for each design. There are many ways to reduce the required computational

time required, including subsampling the full factorial design-epoch space with a design of experiments, creating surrogate subsystems models that perform faster on large tradespaces, and performing EEA on the MATE designs that meet some maximum Pareto-fuzziness number.

Although the results from EEA combined with MATE may more accurately describe the possible system lifecycle value, it still suffers from the same communication problem as MAUT in that the output is not easy to understand from an outsider's perspective. While MAUT outputs a single MAU, EEA outputs timelines and many possible scenarios that are far less likely to be fully understood and mentally digestible by a decision maker.

Value cannot be interpreted simply as the integration of utility over time because utility does not represent value directly; it only represents the *perceived ability* of a system to deliver value. Thus, a single metric like “integrated utility”, “utility-time”, or “value gained” cannot be calculated simply as the average area under the utility curve in all calculated eras.

Additionally, there is no guarantee that an amazing revelation will come from EEA that will show beyond a reasonable doubt and for a justifiable increase in cost that one design that was considered unattractive in static MATE will suddenly be more attractive with the addition of EEA. This is exactly how good metrics and better visualizations can help, both of which are active areas of research.

2.3.3.2 Decision Networks

Another way to analyze future scenarios to assess the value of changeability or ramped upgradeability to reduce risk is decision networks. Silver and de Weck proposed the concept of Time-Expanded Decision Networks (TDNs) [79] as a way to make major decisions in large-scale technical systems independently of EEA. TDNs were applied to launch vehicle development and the deployment of a commercial communications constellation.

TDNs leverage Monte Carlo simulation in ways that are similar to EEA. However, this methodology works best with few design options with highly coupled and well-defined change options. It does not scale well with large design spaces, especially ones that are diverse with few change options, and operates as an optimization technique more than a TSE technique. Overall, EEA is a more mature, more applicable method to assess value over time.

2.3.4 Value Elicitation and Structuring

“Mental models” as defined by cognitive science are explanations of someone's thought process or the synthesis of their perceptions on the world [188], [189]. Humans do not have the processing capacity to constantly imagine the entire world all at once, and the selected concepts and relationships that are perceived represent their mental model. Mental models can also be used to describe how a person perceives value from a given product.

Eliciting people's thoughts, opinions, and values is central to many social sciences, but it is a critical and often misunderstood part of engineering. One of the problems is that the person someone is trying to elicit value from cannot or has not articulated that value. Fischhoff identifies three paradigms for eliciting value based on three assumptions on how well the stakeholder can articulate that value: “People know what they want about all possible questions (to some degree of precision),” “People have stable but incoherent perspectives, causing divergent responses to formally equivalent forms,” and “People lack articulated

values on [a] specific topic (but have pertinent basic values).” [190] Identifying which paradigm a stakeholder being interviewed to elicit a value model is mentally operating in can influence the strategy the interviewer uses to elicit value.

In contrast to the field of cognitive science, engineering design is typically rooted in rigorous mathematics, but when subjective statements of value are mixed with mathematical methods, validating a design can be problematic. Pedersen et al. proposed the Validation Square to validate the performance of a design when subjective value statements are present as an alternative to “formal, rigorous, and quantitative” validation methods [191].

When creating a value model, a designer is building a *constructed*, mathematical representation of a stakeholder’s *mental* value model. The designer will be able to create a better model or understand the inherent limitations in any constructed value model if he or she understands which of Fischhoff’s mental model paradigms the stakeholder is thinking within; the more ambiguity there is, the more potential there is for value to be unarticulated that must be extracted through careful stakeholder questioning.

One of the goals of stakeholder interviews is to digest a person’s mental model into a list of attributes that fully describe their thought process and perception of value. Recent research shows how interactive visualization and value model trading among many models can help identify the mathematical model that best represents a stakeholder’s opinions [65], [192].

In space systems, there is often too much emphasis on calculating the cost and not enough emphasis on calculating the value, whether that is a revenue stream in the case of a business venture or an expression of utility in military and science campaigns [193]. Value-driven design is an effort to shift emphasis away from focusing so heavily on costs and instead focus on value [194]. Using value in the design and acquisition process of complex systems can increase satisfaction, reduce cost overruns, and ensure that designers are putting customer validation up front rather than treating it as an afterthought [195].

2.3.5 Other Value Modeling Tools

In addition to the VCDM’s discussed in Section 2.3.2, there are other tools and concepts that have been previously developed and applied extensively in other fields that can help express the value proposition of distributed satellite systems for space science missions.

2.3.5.1 Value at Risk

Value at Risk (VaR) is a risk metric used widely in the financial industry to express how much an investment could be worth in the future [196], usually shown as a cumulative probability distribution to express the probability that the investment will be worth less than or equal to some amount. VaR is typically applied to stocks or bonds, but is also applicable to investments in other ventures, including engineering projects. This distribution is dependent on a number of probabilistic factors, such as market changes, cost overruns, or performance degradation.

Since VaR is typically associated with dollars-to-dollars comparisons and not CBA or MAUT analyses, it has not been applied to space science missions. Space missions tend to focus on risk calculation or reliability analysis to determine the probability of failure; as a result, they tend not to focus on the probability of success beyond what is expected over the course of the mission. This is unfortunate, because not only do space missions face the risk of failure due to myriad reasons, they also are inherently

searching for something unknown, and therefore if there is something to be discovered by a science mission, the value proposition can be greatly enhanced by considering how much more success could be achieved by one mission compared to others. This is especially important for distributed satellite systems because their inherent flexibility allows them to react to new information, which may provide quantifiable value in the lifecycle of science missions. Converting this concept from a dollars comparison to a MAUT comparison would allow it to be used in MATE and RSC.

2.3.5.2 Value Synergy and Redundancy

One important aspect of the value proposition of distributed satellites is how science goals are addressed in combination. With multiple data sets simultaneously being sent to scientists, it is important to understand the value of how different data sets can be used to answer larger scientific questions when combined. Previous authors have identified instances where multiple measurements have been combined or could be combined to give added value. These combinations tend to either provide more detail than a comparable monolithic data set [197] or can be used to help answer far greater, fundamental science questions [198].

Synergy is defined as the interaction of elements that, when combined, produce a total that is greater than the sum of its parts. Redundancy is the opposite, where the combination is less than the sum of its parts. Selva explored the concept of value synergy and redundancy with multiple instruments on a single spacecraft, but generalized his methods for understanding synergy and redundancy among sets of instruments using synergy matrices [85]–[87]. A graphical representation of how the value of two instruments can exhibit both redundant and synergistic values is shown in Figure 2-5.

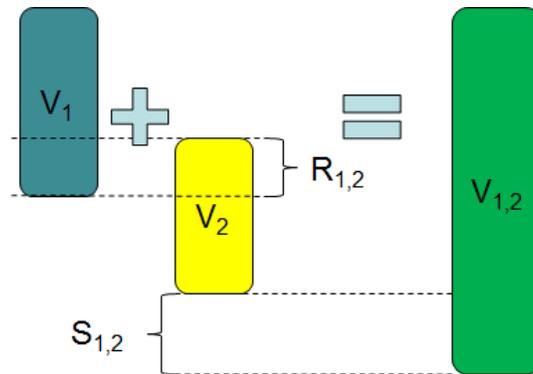


Figure 2-5: Graphical representation of the relationship between value, redundancy, and synergy with multiple measurements.

Mathematically, this is represented as:

$$V_1 + V_2 - R_{1,2} + S_{1,2} = V_{1,2} \quad (2-2)$$

where V_1 and V_2 are the values gained by assets 1 and 2 working individually, $R_{1,2}$ is the redundancy causing a value loss between the two instruments working together, $S_{1,2}$ is the synergy causing a value gain between the two instruments working together, and $V_{1,2}$ is the total combined value. In a more general form, with n assets acting as sources of value, the total value can be represented as

$$V_T = \sum_{i=1}^n V_i + \sum_{\substack{i,j=1 \\ i \neq j, j > i}}^n (S_{i,j} - R_{i,j}) \quad (2-3)$$

Understanding exactly how the value of these measurements stack and overlap usually requires extensive knowledge in the subject field. This is why stakeholder interviews are so critical to the process of creating the value proposition; proper stakeholder interviews are essential to understanding how the stakeholder perceives a number of different options.

2.3.5.3 Alternate Aggregation Methods

Keeney and Raiffa's aggregation method in MAUT is the not only option for combining multiple attributes into a single score. Several other methods exist in engineering design to aggregate multiple attributes that will be discussed here [199].

Quality Function Deployment (QFD) organizes customer requirements, relative importance weights, and potential design variables into the House of Quality in a way that allows experts to identify correlations, couplings, and benchmarks for products [10]. This tool is especially useful in manufacturing processes that are repeated and improved upon (such as in the auto industry, where this method was developed). A simplified version of QFD is used in the early stages of MATE and MATE-CON to identify value-driving design variables for an initial iteration of the performance model.

Pugh's Method for concept selection is useful in trade studies to compare qualities of a few designs to one design that is designated as the datum [200]. Pluses and minuses for each attribute relative to the datum are summed and the best design is the one that has the highest score. A modified version of this puts weights on these criteria to better reflect the relative importance among the attributes so that not all plusses and minuses are equal [201].

Lexicographic selection methods rank the attributes to identify the most valuable attribute and select the set of designs from the design space that best satisfy that attribute [202], [203]. Then, the remaining designs are measured against the second most important attribute, and the ones that satisfy the second attribute best are selected. The design space is whittled down attribute by attribute until one design is left, which should be the "best" design to choose.

However, these methods all suffer from inherent deficiencies in their selection process. Hazelrigg identified ten favorable properties of a design method [204]. Those properties are:

1. The method should provide a rank ordering of candidate designs.
2. The method should not impose preferences on the designer.
3. The method should permit the comparison of design alternatives under conditions of uncertainty and with risky outcomes.
4. The method should be independent of the discipline of engineering.
5. If the method recommends alternative A when compared to the set of alternative options $S = \{B, C, D, \dots\}$, then it should also recommend A when compared to any reduced set of options such as $S_R = \{C, D, \dots\}$.
6. The method should make the same recommendation regardless of the order in which the design alternatives are considered.

7. The method should not impose constraints on the design or the design process.
8. The method should be such that the addition of a new design alternative should not make the existing alternatives less favorable.
9. The method should be such that information is always beneficial.
10. The method should not have internal inconsistencies.

Of the aggregation techniques discussed here, only MAUT satisfies all of these properties. Hazelrigg [205] and Ross and Hastings [13] have shown that these alternative aggregation techniques may not only suggest a design that is not the best among a handful of options, they may even select *exactly the worst* choice for a stakeholder.

2.4 Engineering Modeling

The level of fidelity in engineering models can range from basic concept to full detailed design. Some aspects of a mission can be estimated fairly accurately with basic physics equations, while others require full computational fluid dynamics analysis. How accurate a model needs to be is dependent on a number of factors that must be examined on a case-by-case basis. Where possible, assumptions that are made in modeling will be spelled out explicitly in each case study. This section introduces some of the techniques that are relevant to modeling the performance, mass, and cost of space systems.

2.4.1 Performance Modeling

Performance in TSE is usually measured as an attribute. Stakeholders care about the attributes; therefore they should be calculated with reasonable precision and fidelity. It is infeasible to generate complete, bottom-up models of every system design, so top-down models are generally what TSE uses.

There are a number of tools for modeling space systems. One of the most detailed and comprehensive guides is Space Mission Analysis and Design (SMAD) [206] and its more recent update Space Mission Engineering (SME): The New SMAD [207]. These resources will be used as the primary references for much of the performance modeling put forth in this dissertation including optics, communications, thermodynamics, propulsion, orbital mechanics, and power.

Other references that are specific to individual case studies will be discussed where appropriate in the literature respective chapters. The remainder of this section covers more complex aspects of performance modeling like coverage, orbits, and maneuvers.

2.4.1.1 Coverage Modeling

In case studies where any definition of “coverage” is a metric for success or an attribute, Systems Tool Kit (STK) (formerly Satellite Tool Kit) was used [208]. STK has been used in many studies and is a tool that is available to students in the MIT SSL.

2.4.1.2 Orbit Modeling

Modeling spacecraft trajectories, orbits, and thrusts can be done in a variety of ways. The simplest methods that are usually acceptable at the TSE level of fidelity and are fast enough to run on standard laptop computers are basic physics equations. For simple, short-term orbits, standard Keplerian laws are useful [209], [210]. However, when considering long-term missions, the assumption that the central body is perfectly spherical is not valid. Nodal precession due to the oblateness of a planet causes significant

changes in the orbital trajectory over the period of weeks depending on the orbital inclination that cannot be ignored.

For long-term orbit models around oblate central bodies, the STK high precision orbit propagator (HPOP) gives the highest accuracy, but also takes much longer to run. An acceptable compromise is to use the J4 Propagator. This model has been shown to be accurate enough over periods of years and fast enough for TSE to justify its use in previous case studies. STK has a number of built-in propagators for several planets, moons, and the Sun that are appropriate in a variety of situations.

2.4.1.3 Propulsion and Trajectory Modeling

Most TSE studies use chemical propulsion, and thus, the assumption that any thrust maneuver is impulsive is valid. For these maneuvers, it is appropriate to use the Tsiolkovsky rocket equation to calculate the ΔV and fuel burned for a given thrust and Keplerian mechanics to propagate the trajectory of the spacecraft after the thrust [211].

Thrust maneuvers that use electric propulsion take far more time, and as a result, neither the Tsiolkovsky equation nor Keplerian mechanics are valid for estimating the fuel burn and the expected trajectory [212]. In these cases, STK's Astrogator can be used to propagate a finite burn accurately.

However, STK does not handle atmospheric entry, descent, and landing (EDL) well. Although STK has several well-established 3D models of Earth's atmosphere, it only has 1D models for other planets built in. Additionally, STK does not accurately calculate drag on objects as they transition from hypersonic to subsonic velocities. There exists software to accurately calculate EDL, but this is incredibly restricted because of its use in ICBM trajectory calculations and is thus unavailable for use in the case studies presented in this work. For EDL in TSE, a time-integrated script can be written that incorporates Keplerian orbits, a 3D model of the atmosphere, and ballistic physics equations for objects transitioning from hypersonic to subsonic velocities [213]–[217].

2.4.2 Mass Modeling

In addition to performance, SMAD also gives tools to provide estimates for spacecraft subsystems mass and total mass. Where possible, mass models can be validated against similar components, instruments, or entire spacecraft to at least ensure sanity in mass models if there is no way to truly validate the model. Some mass estimates, such as specific component mass, are known accurately, while others, such as structural mass, are gross estimates based on a set fraction of the total mass.

2.4.2.1 Instrument Models

A simple method for estimating the mass of an instrument based off an existing model is to scale it with a mass-defining dominant design variable. For optical instruments, aperture diameter is usually the design variable that dominates, so the ratio $R = A_i/A_o$, where A_i is the aperture diameter of the instrument whose mass is being estimated and A_o is the aperture diameter of a known instrument, can be used to estimate the mass of the new instrument. SMAD argues that instrument volume scales as the cube of aperture diameter, and that mass and power scale with volume:

$$M_i \sim KR^3M_o \quad P_i \sim KR^3P_o \quad (2-4)$$

where $K = 1$ for large payloads and $K = 2$ for apertures less than 0.5 meters. This assumption is rather naïve; power scales with the electronics associated with the payload, not the aperture diameter, and given that a telescope is mostly hollow volume, mass is more likely to scale by a power between a square and a cube rather than exactly as a cube.

When more than one instrument is available to compare to, a parametric estimate will likely be more accurate. There are a number of example instruments and publications that can be used for this and will be referenced appropriately in their respective case study chapters.

2.4.2.2 Common Small Satellite Components

Typical components like reaction control wheels and solar panels are used in many small satellites. Where possible, performance and mass data can be taken directly from the manufacturer. Manufacturers whose components were used for mass estimates include GomSpace, Clyde Space, Sinclair Interplanetary, and Maryland Aerospace Inc.

2.4.3 Cost Modeling

In TSE and early Pre-Phase A studies, cost estimation is not meant to be as exhaustive as it is during later phases when a single design has been chosen. What is important is that the same cost estimating tools are used to compare all designs in a tradespace the same way. There are a number of ways to estimate costs that will be explored in this section.

2.4.3.1 Simple Cost Models

A simple but effective cost model for estimating the cost of space missions during the early mission phases is to simply calculate the total mass and multiply by a cost coefficient that converts mass into dollars as shown here.

$$C = c_T M_T \quad (2-5)$$

where C is the cost, M_T is the total mass, and c_T is the cost coefficient. This simple estimate works relatively well, but mostly because it is typically used in a static, monolithic satellite systems paradigm. This model does not take into account a variety of factors that would unfairly penalize changeable and/or distributed satellite systems, nor does it characterize the error in the estimate. The next easiest step is to separate the spacecraft dry mass and fuel mass, because the cost of fuel is much cheaper per unit mass compared to the cost of the spacecraft itself. Separating the fuel mass gives:

$$C = c_D M_D + c_F M_F \quad (2-6)$$

where M_D is the dry mass, M_F is the fuel mass, and c_D and c_F are the cost coefficients for the dry mass and fuel mass. Typically, $c_D/c_F \sim 10$, so if two missions that were identical except the second one carried 10% more fuel mass, the cost increase would only be 1% compared to the first, while a third mission that carried 10% more dry mass would cost 10% more than the first. This is the cost model used in the Space Tug tradespace data set [24].

Other adjustments are due to the fact that the payloads are typically much more expensive than other components. There are many complexity metrics associated with different payloads that act as cost

multipliers depending on a number of factors, including technology readiness level (TRL) [218]–[220]. These multipliers transform the equation by separating payload mass and multiplying by that factor:

$$C = c_S M_S + c_P M_P + c_F M_F \quad (2-7)$$

where M_S is the dry mass of the spacecraft without the payload, M_P is the mass of the payload (so $M_S + M_P = M_D$), and the ratio c_P/c_S is a complexity scaling factor.

2.4.3.2 Cost Estimating Relationships

For more detailed cost estimates based on the budget breakdowns of previous space missions of a similar scale and scope, cost estimating relationships (CERs) can be used. An aggregate of several CERs is called a parametric cost model. Space Mission Engineering [207] provides examples of several publicly available cost models, including the Unmanned Space Vehicle Cost Model version 8 (USCM8) [221], the NASA Instrument Cost Model (NICM) [222], [223], and the Cost Construction Model (COCOMO) [224] for estimating software development costs.

Unfortunately, these cost models are not available to the public in their entirety. The full version of the NICM was not available for use without registration; when reached for contact, the designated license distributor did not respond. The Aerospace Corporation’s Small Satellite Cost Model is even more appropriate for this work, but when reached for contact, access was denied because it has to be made available to every student at the Institute, and every student must have a US passport. The full version of USCM8 is restricted by International Traffic in Arms Regulations (ITAR). Team X also has a proprietary cost model that is not for distribution outside of JPL, but results from their studies have been analyzed and reverse engineered for exploring the tradespace of the case study that they also explored. When possible, the example tables in SME will be used for spacecraft development and software costs [225].

Another factor that makes CERs difficult to use is that the relationships are only valid within the bounds of the missions in the database. The same estimates that are used on large satellites and payloads do not scale down to small satellite equivalents because there is no data on which to base an analogous relationship. For this reason, the CERs in SME are not applicable to all case studies in this thesis.

2.4.3.3 Launch and Operations Costs

Launch and operations costs are usually calculated separately from the design, fabrication, and integration, testing, and assembly phases of space missions. However, typical launch cost estimates assume that the vehicle will be the primary payload. If this is not the case, launch costs can be obtained directly from the launch service providers. Spaceflight Services was contacted directly to get estimates on CubeSat launch prices and discounted with bulk purchases.

Operations costs for small satellite missions tend to be more aligned with operations costs for larger missions, so CERs can be used. In cases where operations costs are entirely calculable, such as when the DSN or other networks publish their cost estimating tools, those can be used as well [172], [226].

2.4.3.4 Direct Pricing

When available, a direct quote or price listing from the manufacturer will be used. Parts that will have these will be listed individually but include common CubeSat components like standard solar panels, reaction control wheels, CPUs, and antennas.

2.4.3.5 Cost Improvement

Cost improvement, also referred to as cost progress, experience curves, or learning curves, is the reduction in price of recurring cost per unit as workers produce more identical or nearly identical units. This relationship is defined mathematically as:

$$C(Q) = T_1 \cdot Q^b \quad b = 1 + \ln(S)/\ln(2) \quad (2-8)$$

where $C(Q)$ is the cost of the Q th unit, T_1 is the cost of the first unit (not including non-recurring costs), and S is the slope of the cost improvement ($S < 1$).

Cost improvement theory is well-established in manufacturing industries, but its use and predictability is very immature in spaceflight missions since so few missions use many redundant copies. Fox et al. point out that in the USCM8 spacecraft database, there are 60 lots from 52 programs, and over half use only one unit; only two programs use more than 20 units [227]. As a result, estimating the costs for additional units is still a heuristic metric that can be wildly uncertain. Fox et al. also examined other historical data and gave rules of thumb for the cost improvement slope for spacecraft shown in Table 2-1.

Table 2-1: Estimated cost improvement schedule for spacecraft recurring costs [227].

Total Program Quantity	Cost Improvement Slope S (%)
1-10	95
11-50	90
>50	85

An example of a real-world, mass-produced small satellite is Planet Labs' Dove constellation [228]. When contacted, Planet Labs would not divulge cost or cost improvement information on their satellites (nor their subsystems mass or anything beyond what is available in FCC and ITU filings).

However, the cost improvement theory formula only works for recurring costs; if non-recurring and recurring costs are not separated by the cost model, another heuristic estimate is needed. Experts estimate that additional units for small numbers of satellites (< 10) may be approximately 10-15% of the cost of the first units' recurring plus non-recurring costs [229].

Understanding cost improvement is absolutely critical to the value proposition of DSS since lower unit prices on a design will shift design points to the left on the tradespace (cost axis), but the focus of this work is mostly on how designs leveraging DSS move up and down on the tradespace (utility axis). Cost improvement will be treated as a context variable where appropriate with alternative slope values.

2.5 Research Gap Summary

In summary, the following gaps in published literature have been identified that relate to the overarching goals of this research:

- The methods NASA uses in Pre-Phase A studies do not include the latest developments in tradespace exploration techniques pioneered over the past decade.
- There has been little research (outside of space interferometry) on the emergent capabilities of distributed satellites and how these capabilities add to the value proposition in science missions. More case studies using MATE and EEA can be used to help express the value proposition.

- While there has been research that shows the comparison of distributed satellite systems to monolithic systems in the baseline scenario context of a science mission, no research has yet shown how the comparison changes in multiple contexts with contextual uncertainty, nor has much research been done on parameter spaces where monolithic systems are inherently incapable of operating feasibly.
- No satellite constellation in LEO has yet been required to view targets on the celestial sphere for more than a year; consequently, no year-round coverage methods have been developed.
- Research in FSS has not yet addressed space science missions and how the emergent capabilities of federated satellites can enable ground-breaking science missions, nor how science missions can benefit from future federated networks in space to leverage resources.
- Communications technologies are advancing while our methods of exploring tradespace options with them remain in the past. Better subsystems models are required so that designers can weigh the options between different subsystem architecture choices.
- Multiple location sampling can increase scientific understanding of larger environments.
- Value modeling in EEA is an area of research that is relatively immature and would benefit from more studies on how it can be applied in multiple contexts.
- While it is difficult to put a price on scientific information, the economic impact and benefit of information can be considered when weighing other options for scientific priority.

3. The Emergent Capabilities of Distributed Satellite Systems

*“Into the fires of forever we will fly through the heavens
With the power of the universe **we stand strong together**
Through the force of power, it will soon reach the hour
For victory we ride, Fury of the Storm!”*

DRAGONFORCE, “FURY OF THE STORM”

While many previous authors have explored lifecycle properties of systems – including distributed satellite systems – in detail, little work has been done to focus on the *emergent capabilities* of DSS and how they can be leveraged to achieve new science goals that are impossible, infeasible, or otherwise unattractive to stakeholders within a monolithic satellite systems paradigm. This chapter seeks to answer the first research question presented in Section 1.2 by examining the known *potential benefits* of distributed satellite systems and then combining subsets of those benefits to form potentially valuable *emergent capabilities*. Emergent capabilities are what can be accomplished when multiple assets combine their individual capabilities. For example, a camera can take 2D images, but two cameras together can take 3D images, and many cameras working together can build 3D models of the object being studied.

Section 3.1 defines and examines the potential benefits of DSS as they have been described in previous literature. Section 3.2 examines how those benefits can be combined to form an emergent capability. Within each subsection defining an emergent capability, examples of science missions that have previously leveraged that capability will be provided along with a corresponding case study in Chapters 5 through 8. These cases will examine how that capability can be leveraged in a space science mission.

3.1 Potential Benefits of DSS

In his Ph.D. dissertation, Cyrus Jilla identified 18 properties of DSS that are potentially beneficial to the lifecycle value proposition [20]. The complete list of the potential benefits of DSS identified by Jilla is shown in Table 3-1 and comes from Appendix A of his dissertation.

Other sources of literature were examined to identify other potential benefits or emergent characteristics of DSS, fractionated satellite systems, and FSS, but none provided any additional information that was not already encompassed by Jilla’s work or that could be defined as some combination of the properties listed in Table 3-1.

Jilla then allocated these potential benefits into multiple categories to show how these benefits, when combined in subsets, created additional system lifecycle properties that are potentially valuable over the course of the mission lifecycle. He identified 13 categories and listed the required potential benefits to be combined to form that benefit. Some of his examples include:

Table 3-1: List of the potential benefits of DSS (items in bold are in the “Big Four”)

Potential Benefits of DSS Identified by Jilla [20]
Decentralization of Resources
Smaller, Simpler Satellites
Enables ‘Mass Production’ of Satellites
Promotes Modular Designs
Spatial Distribution of Payloads
Reduced Range to Target
Enables Separate, Sparse Apertures
Multiple Viewing Angles
Reduced Computational Workload
More Satellites
Task Division
Increases SNR
Clutter Rejection
Redundancy
Path Diversity
Allows for Ramped Deployment and Upgrades
Minimizes Required Failure Compensation Costs
Multi-Mission Capability

- **Survivability:** Decentralization of resources, spatial distribution, and more satellites.
- **Reduced Time to Initial Operating Capability:** Smaller, simpler satellites, enables “mass production” of satellites, and modular designs.
- **Improved Revisit Time:** Spatial distribution, reduced range to target, more satellites.
- **Upgradability Ease:** Modular designs.
- **Improved Rate:** Multiple viewing angles, reduced computational workload, more satellites, task division, increases SNR, and clutter rejection.

A similar method that Jilla used to identify categories of potential benefits of DSS will be used to identify emergent capabilities. Not all 18 benefits are directly linked to an emergent capability, but they are all listed in Table 3-1 for completeness. Even if an emergent capability that is being leveraged in a DSS mission for space science does not leverage one of these potential benefits, the mission itself could still leverage that potential benefit to improve its value proposition. For example, “Enables ‘Mass Production’ of Satellites” also reduces cost due to learning curve savings when multiple copies of a payload or satellite are manufactured, even though it is not related to any of the emergent capabilities defined in this chapter.

For the reader’s convenience, the potential benefits identified by Jilla that will be linked to any emergent capabilities defined in Section 3.2 are defined here (items in bold are addressed at the beginning of 3.2):

- **Decentralization of Resources:** When payloads are dispersed among multiple spacecraft so that failure of any single spacecraft does not jeopardize the mission.
- **Spatial Distribution of Payloads:** When payloads are placed on different satellites separated by different orbits or locations. The distance between the payloads is an important element to this benefit when considering the emergent capabilities of DSS.

- *Enables Separate, Sparse Apertures*: When smaller payloads can combine their observations to achieve performance levels much greater than comparably priced monolithic payloads, such as with aperture synthesis interferometry.
- *Multiple Viewing Angles*: When multiple satellites can observe the same phenomenon from different positions, either simultaneously or in series.
- *More Satellites*: When more than one total satellite or payload is in operation. These satellites can be spatially distributed with redundant payloads.
- *Task Division*: When multiple satellites split responsibilities so that the workload of an individual satellite decreases as the total number of satellites increases.
- *Increased SNR*: When the signal-to-noise ratio is increased, either by reducing noise or increasing some metric of resolution in individual observations, thereby increasing the integrity and rate of the system
- *Redundancy*: When more than one payload is operating. Whereas traditional satellite systems gained value from redundant systems by improving reliability, availability, and robustness while reducing risk, DSS may also gain an emergent capability by using these redundant systems in concert to increase performance metrics.
- *Path Diversity*: When information has a variety of options to flow back to the users because of intersatellite communications links.
- *Allows for Ramped Deployment and Upgrades*: When spacecraft can be launched over time so that investment can be spread over longer time to reduce risk or easily capitalize on increased opportunity, but also so that information can be gained from the previous phases, allowing the system to react accordingly.
- *Minimization of Failure Compensation Costs*: When failed spacecraft can be cheaply replaced because only small parts of the whole system have failed, reducing the cost to develop and launch new assets.

The remaining definitions can be found in [20]. It is important to note that Jilla's listing of potential benefits is not entirely complete. The potential benefit "path diversity" is defined specifically for the flexibility of the *communications* path among multiple satellites and how data can be linked back to the ground through many possible routes. However, a DSS mission can also operate along a flexible *operational* path; the decision points within a mission can give mission managers the opportunity to change aspects of the target as information is gained during the mission. This can better inform the decision makers on where the mission should go next. For the remainder of this thesis, two new definitions are proposed that will replace Jilla's incomplete definition of "path diversity":

- *Communications Path Diversity*: When information has a variety of options to flow back to the primary stakeholder due to diverse networking options.
- *Operational Path Diversity*: When information gained during the mission lifecycle influences the decisions of future stages of operations due to new revelations.

This distinction is necessary for rigor and for avoiding confusion.

3.2 Defining the Emergent Capabilities of DSS

The *potential benefits* of DSS described in the previous section were then selected in many combinations to identify ways these benefits could result in an *emergent capability*. Additionally, successful scientific endeavors that used more than one asset were examined to determine under what circumstances they were successful, how multiple assets contributed to their success, and which potential benefits were being leveraged to create the emergent capability they employed. From these combinations of benefits and examinations of scientific endeavors, three categories of emergent capabilities and seven individual emergent capabilities have been identified that can be leveraged to achieve unique science goals or dramatically alter the value proposition of a science mission.

Unlike Jilla’s categorization of benefits, where some of the categories are defined by a single potential benefit, no emergent capabilities are defined without several benefits. During the course of identifying emergent capabilities, it was noted that all of the emergent capabilities leveraged four specific potential benefits in concert: **Decentralization of resources, spatial distribution of payloads, more satellites, and redundancy**. These four benefits will be referred to as “The Big Four” for the remainder of this section and are shown in bold in the list of definitions and Table 3-1.

A monolithic system may carry multiple payloads, but rarely are any of those payloads redundant because there is little value in taking the same measurement with an identical payload from the same location at the same time. However, when more satellites are used with redundant payloads spatially distributed over some domain of operations, capabilities begin to emerge. For these reasons, the Big Four are at the core of all the emergent capabilities that have been identified.

The remainder of this section will examine those three categories and the individual capabilities within each category; fundamentally, analytically, and operationally unique emergent capabilities.

3.2.1 Fundamentally Unique Emergent Capabilities

The first category of emergent capabilities refers to operations that a monolithic system, *no matter how capable, optimized, or big it is*, cannot perform. These capabilities are the ones that will most assuredly allow us to capture low-hanging scientific fruit in a monolithic satellite systems paradigm. These emergent capabilities heavily leverage the benefit of “spatial distribution between payloads” of the Big Four to either perform fundamental science that cannot be done any other way or use a technique to perform science in a way that is more feasible than any alternative monolithic system. Three unique emergent capabilities have been identified that fall into this category: shared sampling, simultaneous sampling, and self-sampling.

3.2.1.1 Shared Sampling

When the Big Four are combined with the potential benefit of “task division,” the result is the emergent capability of *shared sampling*. For the purposes of this work, the emergent capability of shared sampling is defined as:

Shared Sampling: When multiple assets trade responsibilities for making the same measurement at different times, particularly when it is impossible for any single asset to make the measurement over the time period required for satisfaction.

One terrestrial example of this is when ground telescopes on different continents observe a target over an uninterrupted period longer than a single telescope could due to time constraints with the solar diurnal cycle (i.e., night time is not long enough and the Sun rises). With telescopes spread around the world working together, time-sensitive campaigns where observation windows are longer than several hours but the entire time window is important can be successful. This is especially crucial when observing close-approaching asteroids whose opportunities for observation are rare and windows of observation are on the order of 24 hours. Additionally, having a network of telescopes working together rather than the bare minimum makes the system more robust because observations cannot be made when atmospheric conditions are prohibitive.

In remote sensing of Earth, this emergent capability should not be confused with “coverage,” for which there are a number of definitions depending on the application and needs of the stakeholders. For example, an agency may be interested in observing a particular region over a long period of time. One option is to have multiple satellites in low Earth orbit that can make passes over the region at different times. Another option is to put a satellite in geosynchronous Earth orbit (GEO) that can observe that region (and possibly *only* that region) for continuous, long periods of time. Granted, the GEO satellite would require high-performance optics that may be *financially* infeasible or be more costly over the lifecycle than the LEO satellites working together, but it is still *technically* feasible and would still deliver stakeholder satisfaction; hence, it is still an option within the tradespace.

However, the same GEO satellite would *not* have the ability to cover more than one target or region at a time, and an agency interested in covering multiple regions on opposite sides of the globe would need to consider using a distributed system to achieve the desired coverage. For that reason, a distributed satellite system leverages the emergent capability of shared sampling when a single satellite, no matter how capable it is in its performance, *cannot* achieve the objective due to one or more constraints.

For example, the Hubble Space Telescope orbits the Earth at an altitude of 559 km with an inclination of 28.5 degrees. Earth obstructs the telescope’s view of any target for approximately half of the orbital period, meaning Hubble cannot conduct individual observations that last more than approximately 50 minutes (slightly more than half of its 97-minute orbital period). Additionally, because of solar exclusion angle constraints, parts of the sky may be unobservable for three to four months of the year, meaning Hubble alone cannot conduct a campaign to observe some targets year-round. Furthermore, Hubble observation time is expensive and fiercely competitive, so it would be rare to use the platform for extended campaigns anyway. Thus, long-period observations with a low probability of success would neither be feasible nor possible with such a valuable and capable monolithic system.

A distributed satellite system, on the other hand, could potentially leverage shared sampling to observe large or distant areas of the sky for long periods of time. One such an application of that need is in the detection of transiting exoplanets, which requires long-term, high-precision photometric observations. A mission like this is explored in detail in Chapter 5.

3.2.1.2 Simultaneous Sampling

When the Big Four are combined with the potential benefit of “multiple viewing angles,” the result is the emergent capability of *simultaneous sampling*. For the purposes of this work, the emergent capability of simultaneous sampling is defined as:

Simultaneous Sampling: When multiple assets conduct a measurement of a common target at the same time from different locations such that the combination of the resultant data sets provides more detailed information that could not have been gathered by a single asset moving between locations and taking measurements at different times.

An obvious example of simultaneous sampling is how the human brain processes data from the spatially separated eyes to give depth perception. A similar, more advanced example is motion capture work using an array of cameras for use in special effects in modern films and video games, most notably with characters like Gollum from Lord of the Rings and Voldemort from the Harry Potter series. This should not be confused with “bullet time” techniques used in The Matrix, where many cameras are distributed spatially but capture shots individually in time as Neo dodges bullets while falling, because a single camera moving at high speed could feasibly create the same effect.

One terrestrial example of simultaneous sampling that does not specifically involve visualization of light waves is seismometry. Networks of seismometers around the globe monitor seismic activity to deliver far more scientific value than just finding the epicenter of an earthquake. Global seismic monitoring helps us understand the internal structure of the Earth, predict catastrophic volcanic and seismic activity, and even monitor for subterranean nuclear weapons testing. This would be impossible if there was only one seismometer working alone; even if the seismometer could move locations rapidly, it still would not be able to provide the data that would satisfy scientific stakeholders.

One example of an active distributed satellite mission that leverages simultaneous sampling is STEREO (Solar Terrestrial RELations Observatory), a mission to obtain stereoscopic images of the Sun using both ground and heliocentric orbital assets. The fact that the system has more than one satellite makes it possible to gather and combine multiple sets of data that provide scientific return greater than the sum of the individual measurements.

Missions involving measurements of rapidly changing, large force fields are ideal candidates for missions that would leverage simultaneous sampling to achieve unprecedented space science goals. Such an application for that is in the study of the solar magnetosphere, which requires many simultaneous observations from spatially separated payloads to build a more complete picture of the heliocentric current than ever before. A mission like this is explored in detail in Chapter 6.

3.2.1.3 Self-Sampling

When the Big Four are combined with the potential benefits of “communications path diversity” and “enables sparse, separate apertures,” the result is the emergent capability of *self-sampling*. For the purposes of this work, the emergent capability of self-sampling is defined as:

Self-Sampling: When multiple assets measure signals generated by each other, or the precise position and velocity of each other, to infer information about a target or phenomenon indirectly rather than measuring the phenomenon directly.

Whereas the previous two emergent capabilities in this category allow scientific vistas to be explored that cannot be explored in *any* way using a monolithic system, this capability allows scientific vistas to be explored in *alternative* ways that may be far more realistic, feasible, affordable, or marketable to

stakeholders. Additionally, there are subtle nuances, similarities, and potential combinations between self-sampling and simultaneous sampling that will be discussed here.

To date, self-sampling has been successfully leveraged in several missions, but with DSS using only two satellites acting as a fractionated satellite system; it is difficult to identify other cases where value scales with the number of satellites. Additionally, missions that leverage self-sampling tend to suffer from complexity issues arising from formation flight.

There are three examples of classes of self-sampling missions that have been highly successful in previous space science missions. The first example is seen in the GRACE and GRAIL missions [230], [231]. Both of these missions used two satellites flying in a precise formation to measure the gravity anomaly of the Earth and the Moon, respectively. In each mission, the two satellites measured perturbations in their counterparts' orbits to infer the gravitational anomaly of the area they were orbiting over at the time the perturbation was measured. From this data, accurate gravity maps could be made. Another way one could theoretically measure the lunar gravity anomaly is to drive a rover with a gravimeter on board across the entire surface; however, this would not only take a very long time, but also a rover mission with even 1% of that lifetime and mileage has never been launched. For this reason, a space-based DSS system was employed to map the gravity anomaly of the Moon. This mission may not scale meaningfully if more satellites or payloads are launched; there may be improvements in the SNR as multiple, independent measurements between every combination of two satellites could be averaged to find and reduce the standard deviation by $1/\sqrt{n}$ for n satellites, but the improved SNR alone may not justify having more than two satellites. For this reason, and because this mission has already proven successful, this type of mission will not be considered as a case study.

The second technique, which has been successful in terrestrial campaigns but has only been studied in space science missions with DSS and not actually applied, is aperture synthesis interferometry (ASI) [31], [232], [233] (though it has been successful in monolithic missions with deployable booms to achieve spatial separation between multiple instruments [234]). This technique combines multiple distributed telescopes to create a much larger effective aperture area than the sum of the individual aperture areas to achieve higher resolution and signal-to-noise ratios. ASI is technically a combination of both self-sampling and simultaneous sampling; in order to successfully perform this technique, not only are all the assets required to conduct their own separate, simultaneous observations, but also the precise relative locations among the satellites must be known. One could argue that the same resolution could be achieved by building a single, massive telescope rather than many smaller ones; however, such a telescope would be laughably infeasible to design and build, whether it was ground-based or space-based. (One could also argue that because of this, this form of self-sampling is not a fundamentally unique emergent capability and belongs with the definitions of analytically unique capabilities; in this case, the aperture synthesis itself is fundamentally unique, but the data is analytically unique). Since this type of mission has been studied in extensive detail in previous MATE studies, it will not be considered as a case study.

The third and simplest example of a mission leveraging self-sampling involves radio occultation studies. With this technique, a radio signal from the ground (or another satellite) is sent to an orbiting satellite, and the attenuation of that signal can be used to identify properties of the atmosphere such as density, chemical composition, and water content [235]–[238]. This type of measurement can also be done using the Sun as the emitter without using a second satellite to study planetary atmospheres. Performing radio

occultation on asteroids to study their interior composition, however, would require at least two satellites working in concert. A mission partially leveraging this technique to enhance its value proposition will be explored in Chapter 7.

3.2.2 Analytically Unique Emergent Capabilities

The second category of emergent capabilities refers to operations that an incredibly capable monolithic system could theoretically perform, but such systems are incredibly unrealistic or expensive. These capabilities allow scientists to gather large volumes of data quickly to better understand populations of many targets with much greater certainty. These emergent capabilities heavily leverage the benefit of “more satellites” of the Big Four in a way that makes gathering information on multiple targets or in multiple environments more feasible than any alternative monolithic system. Two unique emergent capabilities have been identified that fall into this category: Census sampling and stacked sampling.

3.2.2.1 Census Sampling

When the Big Four are combined with the potential benefits of “task division” and “increased SNR,” the result is the emergent capability of *census sampling*. For the purposes of this work, the emergent capability of census sampling is defined as:

Census Sampling: When multiple assets conduct multiple measurements of a subset (or the full set) of a collection of similar targets in order to have greater certainty in the variance of the desired properties or characteristics of the whole target population.

As the name implies, this is similar to a full census or poll of a population. When the target population is large, it is impossible for a single asset, no matter how capable, to sample every target in a timely manner. The time constraint could be human-imposed, e.g., there is a deadline that must be met for satisfaction to be achieved; or the time constraint could be naturally imposed, e.g., the population is changing faster than it can be sampled for the data to be meaningful by the time the mission is completed.

For instance, it would be ludicrous to send one person to every residence in the United States to collect census data. Not only is this an inefficient means to collect data, by the time that person was finished the population statistics would have changed and the results would not represent the current state of the population. Instead, the United States mails out census packages delivered and collected by a network of mail carriers to quickly gather this data.

Similar techniques are used to quickly assess the opinions of the population. During the lead up to elections, polling agencies do not simply have one person calling over a thousand people to answer a question so that campaign managers can understand the strategic moves they need to make. Those agencies employ teams to intelligently sample the population in a much timelier manner.

Previous examples involving space science are rare. Decadal mapping of stars to measure mean motion and other properties is typically done with multiple telescopes, as are campaigns to find asteroids, near-Earth objects, and potentially hazardous objects, although a single, highly capable infrared telescope in a near-Venus orbit would quickly detect more PHOs than any terrestrial network of telescopes could.

However, there is opportunity for future missions to capitalize on the value added from multiple measurements in many locations to create synergies in the value proposition. Castillo-Rogez et al. discuss

many ways multiple measurements, both with homogeneous measurements on heterogeneous targets and with heterogeneous measurements on a single target, can be used to gain more value from the combination of those measurements than would be gained by the sum of the measurements alone [198]. One notable example they put forward involves measurements of Phobos and Deimos:

*“In-situ exploration of Phobos would most probably solve the mystery of the origin of that satellite. However, the **combined exploration** of both Phobos and Deimos would lead to a **far more fundamental understanding** of the early history of the Martian system, the origin of Mars volatiles, and the genetic relationship between Mars and Earth.” (Emphasis added)*

CASTILLO-ROGEZ ET AL., 2012 [198]

In this example, the sum of the data sets from two different bodies can provide unprecedented understanding of far more than just those individual targets. Modeling that added value would be a challenge, but it should not be ignored. Granted, the “census” in this case involves only two targets, not thousands or millions, but this case clearly demonstrates how value can be obtained from multiple targets.

An application of a DSS that leverages census sampling with a much larger population to explore is in the detailed study and characterization of asteroids, of which there are over half a million known targets. A mission like this is explored in detail in Chapter 7.

3.2.2.2 Stacked Sampling

In general, DSSs leverage mass production to lower costs of additional assets in a mission when the assets are homogeneous. However, this assumption is not always true, since the definition of DSS does not strictly forbid heterogeneous assets operating in concert.

When the Big Four are combined with the potential benefits of “more viewing angles” or “increased SNR,” the result is the emergent capability of *stacked sampling*. For the purposes of this work, the emergent capability of stacked sampling is defined as:

Stacked Sampling: When heterogeneous assets are deployed into different environments or locations to make measurements of the same phenomenon using different instruments such that the combination of data sets is more valuable than any individual data set.

Heterogeneous environment sampling is an especially interesting capability to leverage due to the fact that the actual combined value of the heterogeneous measurements is heavily dependent on the type of phenomenon being observed and the stakeholder value model used to assess the value of the combined measurements.

Heterogeneous measurements can have both synergy and redundancy. Measurement redundancy means the value added by a second data set gives much of the same information as the first data set. In this case, the value derived from the second measurement may come primarily through measurement confirmation, from which statistical significance can be more easily achieved and instrument calibrations can be more accurate, but the combined value or utility of the two data sets is inherently less than the sum of its parts. Measurement synergy means the value of the combined data sets is greater than the sum of the individual measurements. In this case, not only may there be measurement confirmation and confidence, there may also be additional information that can be derived from the two different data sets that could not have

been derived from either individual measurement. Usually combined data sets experience both synergy and redundancy simultaneously.

For example, if someone were trying to measure the temperature of a pot of water, they may use either a traditional mercury thermometer to make an in-situ measurement, or they may use a handheld infrared thermometer to make a similar measurement from a distance. Both instruments use fundamentally different physical properties of the boiling water to measure the temperature; the mercury thermometer absorbs some of the kinetic energy of the water molecules, whereas the infrared thermometer measures the emitted thermal radiation of the water. However, there is nothing that is fundamentally gained as a result of combining the data sets because there is so much redundancy. Synergy may result from the increased confidence, but in the end the temperature is all that is known.

In contrast, an example of synergy in the value of heterogeneous measurements involves understanding the history of the water content on Venus. It has taken many missions over four decades to understand how much water Venus had in the past and the mechanisms by which it lost that water [239]–[243]. Many measurements, both in-situ and remote, had to be taken in consideration together to understand this. For example, Venus Express made in-situ measurements of the H₂O/HDO ratio in the upper atmosphere, while sounding rocket experiments remotely measured the D/H ratio above the homopause. In this case, neither individual measurement gives particularly compelling evidence to form a solid conclusion, but taken together the data sets can confirm powerful conclusions.

Stacked sampling should not be confused with census sampling, where homogeneous measurements of multiple targets are made. Instead, this capability focuses on heterogeneous measurements of one target. Additionally, these measurements can sometimes be made over the course of several campaigns and many years, but gathering data in a few years as opposed to decades is more valuable to stakeholders; therefore, time is an important factor in the value proposition. Stacked sampling is mostly valuable when remote and in-situ measurements are combined. A mission partially leveraging stacked sampling to enhance its value proposition by studying the atmosphere of Mars from the ground and from orbit will be explored in Chapter 8.

3.2.3 Operationally Unique Emergent Capabilities

The third category of emergent capabilities applies to mission strategies that would be impossible or irresponsible for a monolithic system to perform and would never be considered by rational decision makers. These capabilities allow scientists to make decisions that can fundamentally alter the destinations or objectives while the mission is in progress or allow greater risks from the beginning with the possibility of greater reward; the operational advantages of these capabilities not only give a mission room for error and failure, they allow opportunity and flexibility to be leveraged as the mission context changes with new data. These emergent capabilities heavily leverage the benefit of “redundancy” of the Big Four in a way that makes both gathering information on known unknowns and discovering unknown unknowns more probable than any alternative monolithic system. Two unique emergent capabilities have been identified that fall into this category: Staged sampling and sacrifice sampling.

3.2.3.1 Staged Sampling

When the Big Four are combined with the potential benefits of “allows for ramped deployment and upgrades,” “mass production,” and “operational path diversity,” the result is the emergent capability of

staged sampling, also known as distributed deployment. For the purposes of this work, the emergent capability of staged sampling is defined as:

Staged Sampling: When additional assets are deployed after knowledge gained by the first asset's or wave of assets' data has been received and analyzed such that the location or orbit of deployment can be re-evaluated and altered to seize opportunity that was previously unknown or provide similar measurements for more confidence in results.

For a terrestrial example of this, consider the process of calibrating the sights on a rifle. First, the rifle is securely mounted to a bench, and a shot is fired. A technician aims for the bullseye, fires, measures the distance between the bullseye and the impact point, and adjusts the sights. The next shot should be closer to where the sights are aligned. With a few iterations of this, the sights will be calibrated correctly. If the technician only had one bullet, there would be no way to know if the sights were properly calibrated once the gun was released from the mount (unless the first shot was a bullseye, and even then the technician would want to fire another shot for confirmation). In this way, information gained from the result of each shot informs the decision maker on the correct course of action for deploying the next asset.

Similar techniques are used for exploration in the energy sector. Companies drill many small exploratory wells to discover oil sources (this phase of exploration also leverages census sampling) [244], [245]. When a source is found, companies will send in the bigger, more expensive drills and related equipment to increase overall oil yield. This saves the company money over long periods of time since relocating and operating the larger equipment is expensive and it would be a waste to drill where the potential yield is unknown. Failure due to uncertainty is taken as a part of the cost of business and is mitigated through cheaper exploration endeavors; unlike with manned spaceflight missions, failure is an option.

In unmanned space science missions, it is often infeasible for one asset to observe a target location and then move to a completely different target location. While an orbiting asset has the ability to observe a wide range of targets during its orbital sweep of the central body, orbital maneuvers (particularly, inclination changes) are expensive, so satellites tend not to vary their orbits drastically and usually carry only enough fuel for maintaining their orbits over the course of their lifecycle. They typically do not intentionally bring spare fuel to seize opportunities that may arise from knowledge gained during their mission lifecycle.

Ground-based assets are far more limited in their target range in comparison to orbital assets, even though the in-situ data they gather may be considered more valuable than orbital data. This is especially apparent with the many successful ground assets that have been deployed on Mars. Since 1971, NASA has successfully deployed three landers and four rovers on the surface of Mars. The three landers do not have the ability to move beyond their landing locations. The Mars Sojourner rover traveled a total of 100 m during its three-month lifespan. The Mars Exploration Rovers (MERs) traveled significantly further; Spirit traveled 7.7 km over six years before its wheels were caught in soft soil, rendering it immobile [246]. Opportunity holds the record for longest distance traveled by a rover, having traveled a total of 41.4 km as of December 17th, 2014 since it landed on Mars on January 25th, 2004 [247]. For reference, a marathon is 42.2 km. The Mars Science Laboratory (MSL) is quickly catching up with this distance, having traveled 10.2 km as of January 15th, 2015 since it landed on August 6th, 2012 [248].

It should be noted that these distances are total, not linear, distances from the landing site. No rover has gone beyond the region it has started in, nor are they designed to travel distances long enough to explore multiple regions or climates (for instance, if a rover landed on Earth in the desert, it would be in the desert throughout its entire lifecycle, and would not reach a temperate zone or jungle environment). If a rover landed in a scientifically uninteresting area, there would be no way to change locations drastically enough to seize the opportunity to achieve more science.

However, in science missions where multiple spacecraft are landing on a planetary surface and the decisions regarding the landing locations are not completely rigid, the flexibility to temporally distribute the landing of those spacecraft can mean that information gained by one unit can be used to inform mission operations for the remaining units that have not been deployed. For instance, a mission looking for specific geological structures may have several target landing sites. If the first units find interesting results, it may be more valuable to send the remaining units to similar locations; if the first units do not find interesting results, it would be more valuable to consider dissimilar landing sites. Distributed satellites have the capability to change paths during later operational phases based on information gained in earlier stages by leveraging their inherent operational path diversity.

Whether a mission manager would choose to use a monolithic system or a distributed system depends on the balance between the cost of adding extra capabilities to move from one location to another and the value that could be gained by investigating more areas. If there is little interest in visiting multiple areas, or those areas are relatively close to each other, and there is little chance for opportunity to arise, a rover may be more valuable than a series of landers or penetrators. However, if the locations of interest are far apart, and the mission is investigating something that is completely unknown or not known well, opportunity may arise that justifies the added expense.

An application of DSS that would leverage staged sampling is a penetrator mission to investigate subsurface chemistry on Mars. A mission like this is explored in detail in Chapter 8.

3.2.3.2 Sacrifice Sampling

As the number of assets in a DSS mission rises, the relative importance of an individual asset goes down and the risk of a single-point mission failure decreases. Traditionally, mission managers tend to gravitate not only toward safer designs that minimize risk, but also toward safer environments. As a result, risky operating environments are typically ignored in feasibility studies due to the risk of mission failure. However, if only a small percentage of the assets need to survive for value to be gained from the mission, or if value can be gained through intentionally destroying assets by placing them in environments where they will not survive, a mission manager may choose to enter this environment with a DSS.

When the Big Four are combined with the potential benefits of “minimization of failure compensation costs” and “mass production,” the result is the emergent capability of *sacrifice sampling*. For the purposes of this work, the emergent capability of sacrifice sampling is defined as:

Sacrifice Sampling: When assets are knowingly deployed to an unsafe operating environment, such that most if not all of the deployed assets will be destroyed, but the scientific returns can justify the development and launch of the entire mission.

Sacrifice sampling enables mission managers to capture scientific value that would otherwise be left out of consideration for exploration in a monolithic systems paradigm. A fictional example of this capability used in a mission comes from the final battle in the Orson Scott Card novel *Ender's Game*. During this mission, thousands of human spaceships were sent to attack the alien home world, penetrate the space defense network, and fire a doomsday payload to destroy the planet [249]. Many ships were outfitted with the superweapon, and in the chaos of the battle only one ship managed to punch through close enough to fire its weapon. One shot was all that was needed, however, and because Ender's attack strategy involved multiple assets operating in a high-risk environment, the mission was a success.

A nonfiction example of leveraging sacrifice sampling in a DSS is the Galileo mission. The Galileo spacecraft was technically a fractionated spacecraft because it consisted of an orbiter and an atmospheric probe. The probe was sent into the atmosphere of Jupiter and survived for 58 minutes, transmitting data back to the orbiter before being crushed by atmospheric pressure and heat. This data has since revolutionized our understanding of outer planet atmospheric structure, composition, and thermal profiles [250], [251]. The scientific findings from this probe could not have been achieved feasibly with a system that was made to survive. (The Galileo orbiter was intentionally deorbited into the atmosphere of Jupiter to prevent possible biological forward contamination of Europa; it was not designed to perform any investigations during this phase.) A concept that is in development that will leverage sacrifice sampling is the Sprite spacecraft [252], [253]. These complete spacecraft weigh less than a quarter and can be used to make many stochastic measurements of atmospheric data.

A DSS mission that leverages sacrifice sampling is very similar to a mission that leverages staged sampling, but one that delivers value impulsively as opposed to one that delivered value by surviving for long periods of time. For this reason, this is the only emergent capability that will not be explicitly explored in any case study; however, the techniques used to assess the impulsive value delivery for staged sampling are identical to what would be applied to a mission that leveraged sacrifice sampling without time-dependent observations. This difference in value delivery will be explored in Chapter 8.

3.3 Chapter Summary

Three categories and seven specific definitions of emergent capabilities of DSS have been derived from the known benefits of DSS. Examples of how each capability has been used previously in both terrestrial and space-based applications and examples of how new scientific vistas could be achieved by leveraging them in future space science missions have been illustrated. Those categories and capabilities are:

- Fundamentally Unique Emergent Capabilities
 - Shared Sampling
 - Simultaneous Sampling
 - Self-Sampling
- Analytically Unique Emergent Capabilities
 - Census Sampling
 - Stacked Sampling
- Operationally Unique Emergent Capabilities
 - Staged Sampling
 - Sacrifice Sampling

4. The Responsive Systems Comparison Method for Distributed Satellite Systems

*“It's a piece of cake to bake a pretty cake. If the way is hazy,
You gotta do the cooking by the book; **you know you can't be lazy!**
Never use a messy recipe, the cake would end up crazy,
If you do the cooking by the book, then you'll have a cake!”*

LAZY TOWN, “COOKING BY THE BOOK”

Chapter 3 answered the first research question presented in Section 1.2 by defining the emergent capabilities of DSS. The remaining chapters are dedicated to answering the second research question. This chapter gives an overview of the primary method used to explore the case studies, the Responsive Systems Comparison (RSC) method, and goes into detail on the justification for its use. Additionally, elements of the method that have been modified in order to address the research in this dissertation will be discussed in detail.

4.1 Overview

The Responsive Systems Comparison method is the marriage of Multi-Attribute Tradespace Exploration (MATE) and Epoch-Era Analysis (EEA) and is the method that has been chosen to study the cases that will be explored in Chapters 5 through 8. Previous work using MATE and EEA for analyzing space science missions that is discussed in Chapter 2 provides some justification for their use in understanding the value proposition of distributed satellite systems, but more on these justifications will be discussed in this chapter. Many modifications and recommendations have been made to tailor the RSC method as it exists in the literature to be more appropriate for evaluating space science missions that leverage the emergent capabilities of distributed satellite systems.

Some of these recommendations, prescriptions, and suggested modifications are being made based on the synthesis of many published techniques and methods and discussions with practitioners of those methods about their weaknesses and how they can be applied specifically to evaluating DSSs in space science missions. The literature review presented in Chapter 2 uncovered many possible methods that could be used that may give similar results to these case studies; however, the methods presented here showed the most promise for surpassing the shortcomings present in previous studies of DSS. Those studies chose to examine DSSs specifically in comparison to monolithic systems and how they may be better suited to fill roles already dominated by monolithic systems; this work instead asks what DSSs can do to capture opportunities that monolithic systems cannot. This chapter will illustrate that monolithic systems-based thinking has been so prevalent in space systems engineering over the past half century that inexpensive ways for planetary science missions to capitalize on DSS opportunities have been ignored or rejected for the wrong reasons. Within a DSS paradigm, scientists and systems engineers can learn to find these opportunities and design cost-effective systems to maximize the scientific returns of these missions.

Other prescriptions, modifications, and general rules of thumb have come from personal experience developing the case studies herein. Some of the lessons learned are the result of completely starting over on case studies, or coming to conclusions that clearly disagreed with stakeholder perceptions and had to be revisited. These recommendations also come from experience gained catching mistakes that would have gone unnoticed otherwise by systems engineers who do not have an intimate knowledge of the scientific literature associated with the goals of a particular mission; the author’s background in both aerospace engineering and planetary science gives him a unique perspective that allowed him to correct errors and understand scientific value that even the most talented systems engineers would not necessarily see. Further details on the specific insights and what caused them can be found in the individual case study chapters, but this chapter contains the general overview and recommendations.

4.1.1 General Framework Description

There are many different ways to define the exact number of processes and steps in RSC, MATE, and EEA. Dr. Adam Ross, the senior researcher at SEArI who was instrumental in the development of MATE, describes MATE as a framework that connects a set of related concepts and activities that can be implemented in a number of ways to better understand the design space and choose designs that maximize value. One particular instance of MATE he defined is a 48-activity process that can be explicitly followed in some cases, though other users may desire a more fluid, less structured approach. MATE splits into three general categories: Needs Identification, Architecture-level Analysis, and Design-level Analysis.

EEA is a framework that can consist of any of the following activities: single-epoch, multi-epoch, single-era, and multi-era analysis. Within each activity there are a number of metrics and techniques that can be applied to learn more about the system being designed. Schaffner defined RSC as a nine-process method with three major phases: Gather-Evaluate-Analyze. The major steps of MATE and RSC are shown as they are published in Ross and Schaffer’s masters’ theses in Figure 4-1.

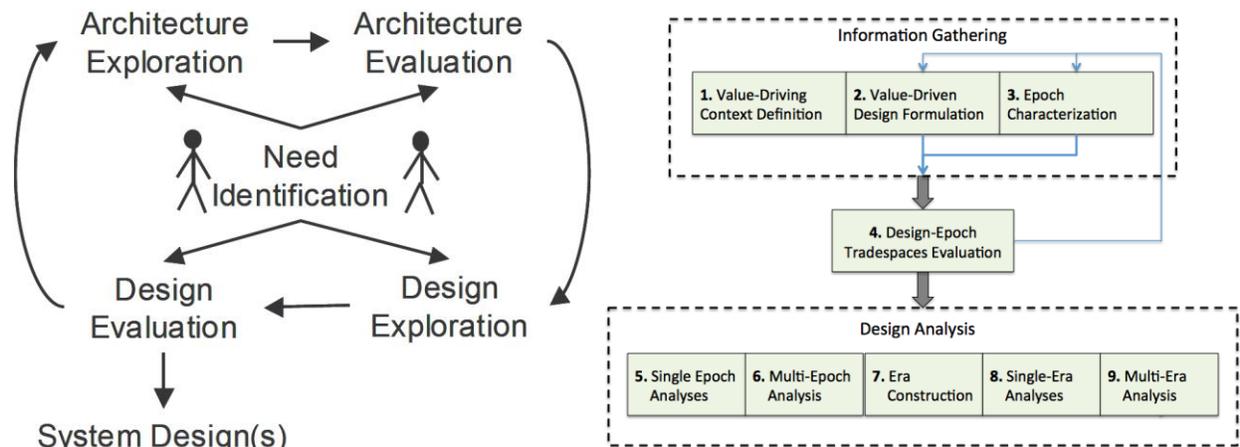


Figure 4-1: Major categories of processes in MATE and RSC (reprinted with permission from [25] and [58])

There is much more within the bibliography and in Chapter 2 about the development of these methods and many of the activities and metrics that can be performed and calculated alongside the primary process, but the central processes and metrics will be explained in detail in this chapter, especially when modifications are made or applications of new techniques are also included.

For purposes of this thesis, the major phases of the RSC method have been conceptually split into four distinct phases: value modeling, performance modeling, tradespace exploration, and EEA. This is shown in Figure 4-2.

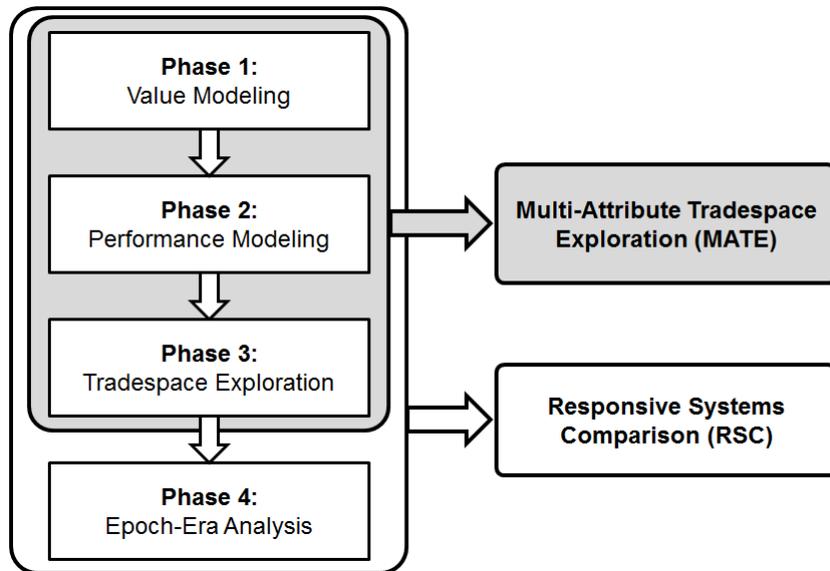


Figure 4-2: Major phases of the RSC method as described in this work.

Value modeling consists of soliciting stakeholder inputs to identify major goals, attribute definitions, and utility curves (the “Information Gathering” blocks shown on the left in Figure 4-1). From these inputs, design and context variables are chosen that appear to have the largest influence on the total value. Performance modeling takes context and design variables and calculates the levels of the attributes to find the utility as the stakeholder would define it from the value model generated in Phase 1. While some would argue that “cost modeling” is separate from “performance modeling,” others would argue that cost is simply an attribute that is important to the value proposition and is calculated just like every other attribute. For this reason, cost and performance modeling will be considered part of the same phase throughout the rest of this dissertation, even though they are intentionally separated in Chapter 2.

Next, tradespace exploration combines the cost estimate and expected utility to plot a snapshot of all designs on a tradespace. Changes in value due to changes in design variable levels can be seen, as can relationships between two or more design variables. Understanding these trades is especially important when a change in a design variable level that should increase performance, according to the designer’s intuition, actually decreases lifecycle value. Finally, EEA looks at changes in the context variables and possible mission timelines to analyze the future contexts the system will operate in over the course of the system lifecycle to better express how systems deliver, retain, or preserve value in the face of uncertainty due to both endogenous and exogenous factors and events.

While theoretically one could perform these four phases in this exact order once, in reality the method delivers more meaningful results if there are iterations between the phases and each phases is revisited several times throughout a case study. This is especially true when stakeholder ambiguity is present, because once stakeholders are presented initial results, they may react by articulating a better mental value model. That new value model could include new elements that had not been considered before, and the performance model would need to be retooled to evaluate those elements. None of the three graphics

in Figure 4-1 and Figure 4-2 adequately capture the feedback loops that are present between the different phases and between the designers and the stakeholders.

In all case studies, value modeling and performance modeling are closely intertwined and were performed concurrently for validation of both models, and in practice each model was iterated a minimum of three times because of several reasons, most notably: (a) key information about the stakeholder's expression of their value model was not discovered until after one or more passes through the performance modeling phase, most notably how new design variables that were not considered beforehand could affect the value proposition in meaningful ways, (b) the performance model was either unnecessarily detailed for calculating an attribute or intermediate variable, meaning it could be simplified or removed from the model to speed up calculations, or it was not detailed enough and more submodules had to be added, or (c) more options for one or more design variables or new design variables were added to the design space to ensure that the limits of the tradespace were being enumerated.

4.1.2 Summary of User Roles

Before going further, it is useful to define the roles that are involved in this method. A single person could theoretically play every single role in an entire case study if he or she had the appropriate scientific background to determine what is necessary for the mission, the systems engineering background to translate those needs into an engineering problem, and the modeling background to simulate the designs, but this is practically never the case in detailed studies.

Like most major product acquisition processes in engineering, one can think of there being two sets of roles in the development of a DSS for a space science mission: the "customer" side and the "developer" side (as illustrated by the two stick figures on the left of Figure 4-1). Many of these terms may be used interchangeably if the study consists of a few individuals who perform several roles, but in complex cases with many actors, it is important for each actor to know their roles.

The "customer" side consists mostly of the stakeholders, who may include scientists, principal investigators, subject matter experts, the "end users" who analyze scientific data the mission would return, and other people or groups who have vested interest or could benefit from the products of the design. These are the people whose needs are being satisfied by the mission or may benefit from the mission despite not directly being a part of it. The "developer" side consists mostly of the group that conducts the engineering, which may include system architects, designers, system modelers, analysts, and engineers. These are the people who take the inputs of the stakeholders and develop concepts and designs that could satisfy the stakeholders' needs.

Perhaps the most critical role in the entire study is the person who elicits the value model from the stakeholders and translates that into the inputs the engineers need to generate designs. This person will be referred to as the "interviewer," though the interviewer could be the same person as the head system architect or the principal investigator (PI). Alternatively, "value elicitor" could be used as a more general term, especially when no stakeholder is available to interview to directly elicit a value model. In these cases, the value model may be based on published recommendations, minimum satisfaction requirements, or other inputs. In this work, the "interviewer" also assumes these duties, and it will be made clear when a value model comes from a direct interview or from other sources.

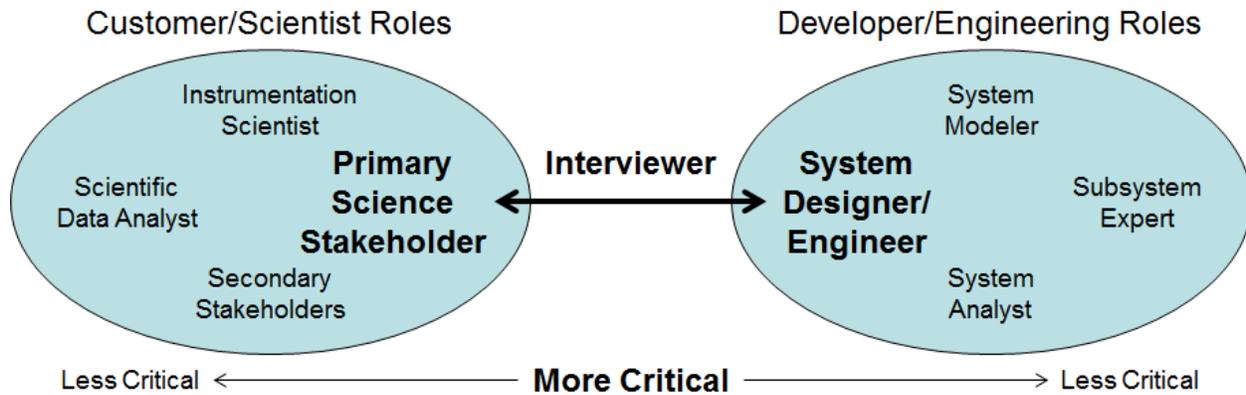


Figure 4-3: Roles in an RSC case study for a DSS science mission.

Elicitation of the value model may be the most crucial part of designing for DSS, especially when a monolithic system may be satisfactory in comparison. In this work, the author designated someone to be a “Primary Scientific Stakeholder” (PSS) for each case study, who would be interviewed to solicit opinions from; other applications may also require multiple levels of stakeholders and inputs (as described in Section 4.2.3.2). These roles are summarized in Figure 4-3. Many of the titles of the less critical roles are used interchangeably in this work.

It is important to note that the term “critical” in this figure is not meant to note a lack of importance to the overall process; it is meant to signify that the role that is expected does not change considerably under a DSS paradigm compared to the monolithic systems paradigm a person with that title may be used to working within. The roles of subsystems modelers and instrumentation scientists in these studies would differ little from their everyday duties and expectations, whereas high-level stakeholders, systems designers, and interviewers all need to understand how leveraging the emergent capabilities of DSS impacts the value of the mission.

4.1.3 Chapter Structure

Each phase shown in Figure 4-2 is described in further detail in the following sections with justification for the use of each component that is relevant to this work and an explanation of what changes or additional metrics will be applied. Each section begins with a brief discussion contextualizing the phase from the perspective of DSS science missions with key insights for what makes it different from the traditional approach with RSC on monolithic systems. The individual processes of each phase are discussed, including how the proposed modifications or alternative strategies or mentalities are relevant to DSS missions. A brief discussion of the key limitations and weaknesses follows at the end of each phase’s description. A summary of the method and its weaknesses are synthesized following the complete description, and a brief introduction to the case studies that compose the remainder of this dissertation concludes this chapter.

4.2 Phase 1: Value Modeling

The first phase of RSC is Value Modeling. For any engineering endeavor, the primary stakeholders, needs, and goals need to be identified before possible solutions can be evaluated. In RSC, the goal of this phase is to more explicitly understand how the stakeholder obtains or perceives value for the system or product they are soliciting or purchasing (note: the stakeholder may not necessarily be the purchasing

agent as well). In the case of space science missions, the scientists who serve as the primary stakeholders communicate what products they want and what aspects or “attributes” of those products are desirable. The end goal of this phase is to understand the primary goals, identify all the relevant attributes, weigh the relative importance among those attributes, construct functions that show how those attributes satisfy the needs of the stakeholders, and identify other variables that could change the stakeholder’s expectations or the performance of the design that are outside of the direct control of the designers.

4.2.1 Contextualizing Phase 1 for Distributed Satellite Systems

Before going further, it is necessary to point out some key differences between scientists and engineers, between monolithic and distributed systems, and how value is gained from space science missions. A designer entering the value modeling phase with unintentional biases can cause the stakeholder that is being interviewed to create additional biases and affect how the perception of value is communicated from the stakeholders to the designers. This breakdown in communication could result in a list of design attributes that a stakeholder will clearly agree represents their value statement and that are easily verified and validated, but doesn’t truly capture value that was not expressed due to those biases.

Although the general public may confuse the two, science and engineering are very different fields that require different mentalities. For centuries, it was possible for scientific progress to be made by an individual scientist with tools they could make on their own. Nowadays, progress is made with complex systems that require scientists and engineers to work collaboratively. Although it is difficult to argue against this progression, one consequence of this is that engineers have dominated the final designs because they are the ones who touch it last, and sometimes the scientists’ true opinions on how a product obtains value are lost along the way.

As a consequence of the current practices in systems engineering, engineers tend to operate based on requirements. While requirements provide healthy boundaries during the later design phases, they can be detrimental to the creative process and to design exploration [12], [254]. Because engineers use requirements as boundaries to set up the problem and work towards a point solution, they tend to think about “what can we feasibly do?” before understanding “what needs to be done?”

Scientists used to thinking within the limits of engineering requirements hardly ever stop to think about what they *really* want. Engineers are so quick to provide a solution derived from requirements that scientists typically don’t get to express what data products would answer scientific questions the best or truly be revolutionary. In theory, requirements are derived from scientists’ needs and act as a proxy for their value proposition, but effective requirements elicitation may not occur in practice, leading to requirements that do not represent the stakeholders’ needs. During one of the case study stakeholder interviews, the scientist actually exclaimed “Wow, I’ve never thought about what I want; no engineer has ever asked me these questions before!” Further, neither scientists nor engineers are typically used to thinking within a DSS paradigm, where goals could be better satisfied with *more* instruments instead of *better* instruments.

Within a monolithic satellite systems paradigm, the trend is to build larger, more complex missions that carry bigger and better instruments or gather data with higher resolution and SNR. This is especially evident in the Mars exploration program as rovers have gotten larger, leaving different members of the science community behind because they did not have enough political decision-making power to

influence the design and operations (e.g. landing sites, not just instruments). Is this trend really what scientists *want*? Is this trend conducive to unlocking the secrets of the universe?

Of course, what scientists *actually* want is often impossible for engineers to deliver. Ideally, their data would be ultra-high spatial resolution with no noise, with an impossibly low frame rate and maximum coverage. A typical engineer who is firmly rooted in monolithic systems thinking might focus on the spatial resolution with little regard for the coverage and temporal resolution when acting as the interviewer to elicit a value model. In a monolithic satellite systems paradigm, the focus could never consider how viewing from multiple angles simultaneously, quick follow-up observations, or multiple operating zones could amplify the value of the data gathered or its ability to confidently answer fundamental science questions.

The monolithic paradigm as seen with the Mars rovers shows that science questions can still be addressed with higher-quality data from more sophisticated instruments, but if a mission only seeks to provide better resolution of phenomena that have already been characterized to satisfy a select group of stakeholders in one scientific field, there is little chance of a valuable scientific breakthrough that will lead to finding the answers to far more fundamental scientific questions – questions that address larger, broader phenomena and the emergent properties of extended targets to provide data that has the potential to provide evidence to answer a wider range of questions.

Castillo-Rogez et al. suggest that measurements of two targets (Phobos and Deimos) could answer fundamental questions about the nature of the origin not only of the two target bodies but also of Earth and Mars [198]; other examples include studying the Sun from multiple angles to better understand the physics of fusion, studying sites around an entire planet to discover historic cataclysmic events, and studying many asteroids to understand solar system dynamics and formation. A wider net can be cast with distributed missions that have a better chance of leading to breakthroughs in multiple fields.

In contrast, highly specialized monolithic systems tend to serve a more limited range of scientific stakeholders; there was considerable debate over where to land the Mars Science Laboratory, and the groups whose goals were not a priority of the agency will not have the opportunity to explore their desired landing site for at least several years. There is a chance that because the coverage of the rover is so small compared to the target body (the entire surface of Mars), major discoveries could be missed. This is why “better data” – better spatial resolution in the monolithic sense – is not always better than “more data” – better temporal resolution or coverage in the distributed sense.

This is why the interview process in MATE is so critical but can be easily skewed – the scientist responding to questions from a designer will answer questions with a biased view if the designer presents them with questions whose framing is biased. In order to help minimize over-fixation on a particular *form* of the potential system, Dr. Ross recommends that the stakeholder imagine there’s a magic black box that can give him/her what he/she wants rather than assuming any form to the solution. This framing can help the stakeholder focus on his/her own values and expectations, rather than indirectly biasing their responses based on expectations of what a particular system solution could potentially provide. In the end, the interviewer needs to know what is important about what comes out of the black box, not what’s inside the black box. The engineers are then free to develop the inner workings of the black box after the stakeholder interview defines the objectives for what must come out of the box.

Outcome-oriented engineering, sometimes called capability-based, effects-based, or function-based (collectively considered as “value-focused thinking” when oriented around stakeholder values), is considered best practice in systems engineering, as opposed to alternative-focused thinking, especially during early phases of the system lifecycle [255]–[258]. The “black box” artifice helps frame the interview into a value-focused one, allowing separation for each role to do what their best at: scientists describing their scientific needs, and engineers developing technical solutions.

This is why the decadal survey should focus purely on fundamental questions rather than missions, destinations, and what can be discovered at those destinations. If instead of asking “What can we accomplish on Mars with this budget?” the committee could ask “How can we answer as many of these unanswered questions given the resources we expect to have over the next decade?” Scientists or engineers proposing missions can then show how fundamental scientific questions can be answered by that mission or how well they can be answered.

More importantly, using the methods described in this chapter, they can show how the stakeholder satisfaction or data products derived from combinations of scientific goals can multiply, not just add, value in a space science mission. By starting from the top-level needs and working downwards, scientists can identify goals with synergistic value. When these synergies are identified, the true impact of the missions can be better communicated. A mission to visit both Phobos and Deimos may end up costing twice as much as a mission to either object, but the combination of those measurements is worth far more than twice the price considering what other questions they can answer [198].

The case studies that have been selected for this dissertation are all examples of missions that answer science goals that have either never been addressed in meaningful ways or that have not been addressed to their full extent. This thesis is meant to be a “How-To” guide for both scientists and engineers. For scientists, it is a guide to think about how they can leverage emergent capabilities of DSS to answer fundamental science questions. For engineers, it is a guide to understand how to better satisfy the scientists’ *actual* needs, not the ones they have solicited through a possibly biased interview process.

4.2.2 Value Modeling Process 1: Value-Driving Context Definition

The first process in value modeling is identifying whose needs or wants are being satisfied with the product or system being delivered. For space science missions, there is usually a Principal Investigator (PI) who is the ultimate decision maker, though there are a variety of other stakeholders or external pressures that can dictate the course of the mission lifecycle. While PIs may not have complete control over budgets, requirements, and subcontractors, they usually have the strongest opinion about how value is defined, interpreted, or delivered by the mission. In more general cases, such as in the context of a commercial satellite manufacturing company selling a concept or design to the scientific community, there may not be a single PI or stakeholder whose needs dictate the course of the mission more than any other individual, and value may be interpreted using a variety of metrics simultaneously depending on the stakeholders, which also include the developers and operators in addition to the scientists. This general case is beyond the scope of this research, however.

For space science missions, there must be some decision maker or scientific authority that can serve as the Primary Science Stakeholder (PSS). In order to maximize the chances that a mission will deliver perceived high value to a particular segment of the science community, the representative decision maker or group of decision makers serving the role of PSS should have expertise in the field being studied by the

space science mission and have expert opinions about how valuable different options for systems would be. They are the ones that should be interviewed to elicit their value proposition. Once a PSS has been identified, the science goals must be determined and defined as explicitly as possible.

Other organizations or individuals may be Secondary Mission Stakeholders (SMS). Other scientists may be able to use the data products that a mission gathers; these are Secondary Science Stakeholders (SSS), who may have additional science goals or opinions on how the data from the mission adds value to their fields. The scope of this dissertation does not fully include these parties, but examples of their needs and influence will be qualitatively addressed where appropriate.

4.2.3 Value Modeling Process 2: Value-Driven Design Formulation

The second process of the value modeling phase is value-driven design formulation, where stakeholder interviews are conducted. This is typically done in several stages over several interviews. The first interview is mostly general discussions about the desires of the PSS, learning more about the science goals and the intricacies and nuances of the goals. By the end, the designers should have some idea of what the attributes should be to begin generating concept architectures, identifying instruments, and proposing design variables. Design-Value Matrices (DVMs) are also helpful for prioritizing design variables and models. Usually at this point the first iteration of the performance model is constructed at a basic level to get some understanding of how the model will go from design variables to attribute values.

It is also advisable to explore the limitations of the attributes. For instance, a stakeholder may suggest that “landing accuracy” is an important attribute. Modeling entry, descent, and landing onto a planetary surface to an accuracy of a few hundred meters is not an easy process. On the other hand, if the stakeholder cares about accuracy only on the order of tens or hundreds of kilometers, a high-fidelity simulation would not be necessary, and the modelers would be wasting resources by creating an EDL model when they should be focusing their modeling efforts on other aspects of the mission.

By the second interview, the PSS and designers should agree on a set of attributes that would be used to define how satisfaction can be achieved. At this point, it is necessary to select a VCDM to continue. MAUT was discussed in detail in Chapter 2, but it has been modified from its standard formulation to be applied to science missions leveraging DSS.

4.2.3.1 VCDMs Revisited

Some of the many properties associated with space science missions and whether or not some of the VCDMs discussed in Chapter 2 can be accurately applied *as they are defined in the literature without modification* are summarized in Table 4-1 (a red X signifies inappropriate use, a green check signifies appropriate use, and an orange circle signifies possible use or that it should be used with caution). The first seven properties are discussed by Ross et al. [55] and the last two have been added here.

Table 4-1: Summary of the properties of typical space science missions and how well popular VCDMs can be used to assess their value

Properties of Typical Space Science Missions	NPV	CBA	TOPSIS	MAUT
Value can be derived from multiple sources	X	✓	✓	✓
Value is NOT discounted cash flow	X	X	✓	✓
Mutual Utility Independence	X	X	X	✓
Stakeholder decision is under uncertainty	X	X	✓	✓
Value is an ordered, non-ratio comparison	X	X	✓	✓
Context may change over time	X	X	X	X
Applicable with multiple stakeholders	✓	✓	X	X
Value is delivered over time	✓	✓	X	X
Value can be delivered synergistically	X	X	O	O

From this list we see that although MAUT does have several weaknesses, it is the most appropriate VCDM to use in space science missions based how many of the properties of space science missions MAUT is appropriate for. This is why it has been used as the basis for tradespace exploration in MATE. This table will be revisited at the end of this chapter to show how modifications to the RSC method can either improve or eliminate these weaknesses when applied to DSSs in space science missions.

The first weakness that can be addressed is “context may change over time”. When a method is used that accounts for the dynamic environment that a spacecraft is developed and operated in that can handle contextual uncertainty, the value of these lifecycle properties can be more explicitly communicated, so this limitation of MAUT can be mitigated. This is why the RSC method combines MAUT with EEA – EEA shows what can happen over the course of the mission lifecycle, so stakeholders can see more than a static snapshot.

The second weakness, “Applicable with multiple stakeholders,” is an active area of research [65], [66], [259], [260]. Early results have shown that through additional negotiations with tradespace visualization, consensus can be achieved in collaborative design problems using MAUT with the addition of MATE and EEA (though this is not the same as aggregation across multiple stakeholders). Whether or not this limitation can or needs to be overcome is dependent on the application and how different the needs and opinions of the stakeholders are.

When the system is expanded in time as it is in EEA, the relationship between “utility” and “value” changes, and the definition of “value” needs to be better suited to what is being delivered throughout the course of the mission lifecycle rather than simply being an aggregate or integration of utility over time. The weakness “Value is delivered over time” is not mitigated by the introduction of EEA to MAUT alone [34], [46]. However, a modified MAUT model that accounts for the data products from a science mission as they are delivered over time, not just as a static output, can make up for the weakness when using MAUT alone. The modifications to MAUT will be addressed in Section 4.5.6.

Finally, MAUT alone does have some power to address synergy and redundancy already, but the framework is typically used to express the stakeholder’s risk attitudes and tolerances explicitly in the general case. By reframing the interview and the line of questioning, the same system can be used to address the synergy between science goals as they are perceived by the PSS. This will be addressed in the next section.

4.2.3.2 Multi-Tiered MAUT Structure

Typical studies use the aggregation of attributes, usually no more than seven, to calculate the Multi-Attribute Utility (MAU) of a particular design. The standard form of this hierarchy is shown in Figure 4-4. In this hierarchy, if any of the individual attributes of a design do not meet the minimum satisfaction criteria, then the entire mission is unacceptable to the stakeholders. This hierarchy is defined for a single set of preferences elicited from one stakeholder (or a preference set that a set of stakeholders has agreed upon). However, more complex mental preference sets may have multiple levels of associated utility satisfaction that need to be taken into account that the standard hierarchy does not account for alone.

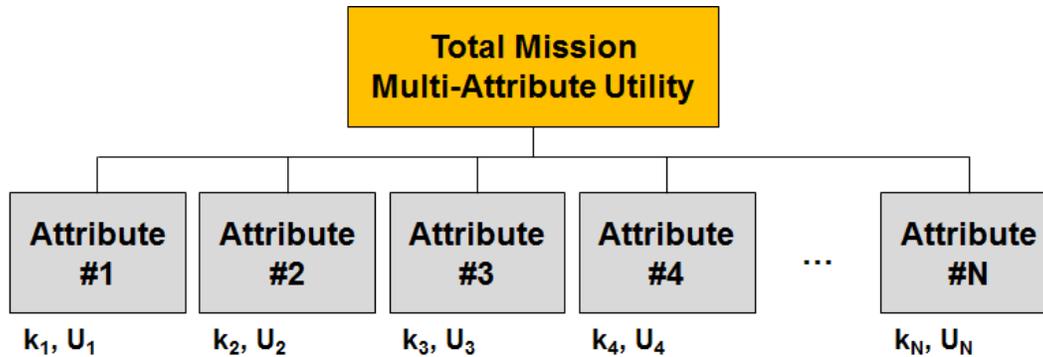


Figure 4-4: Standard hierarchy of MAUT to calculate total MAU of a design.

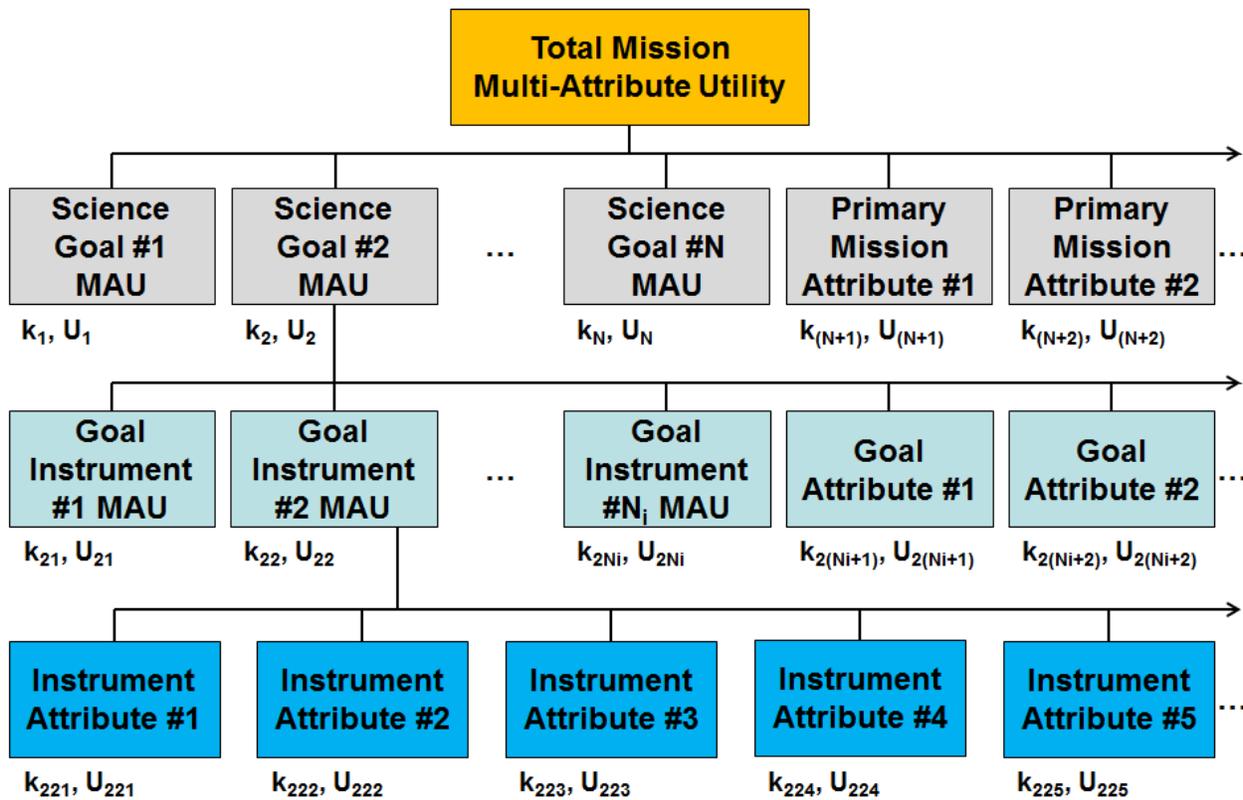


Figure 4-5: Modified hierarchy of MAUT to calculate total MAU when multiple science goals and instruments are present.

Keeney and Raiffa also discussed the hierarchical nature of objectives and how to construct hierarchy trees within MAUT to organize attributes [180]. Attributes at the second level of Figure 4-4 may be

complex objectives with attributes of their own, which require a third level of attribute definitions. Such a hierarchy also removes separation concerns if one person does not have preferences or authority over the entire hierarchy tree. This led to the development of a MAUT hierarchy tree that is representative of space science missions that leverage DSS to achieve multiple science goals. The proposed modifications to the typical hierarchy that will be used to assess mission utility in the case studies that follow this chapter are shown in Figure 4-5.

From the top, this hierarchy takes into account that there may be multiple science goals that a single mission may be trying to address. Although MAUT cannot be used to balance needs between multiple stakeholders with a “social welfare” function, a single stakeholder could have very important opinions about the relative importance among different science goals if more than one science goal is being achieved or considered in the concept design phase.

The top-level k_i weights are assigned to show the relative contribution of each science goal and other primary mission attributes that are important to the overall scientific satisfaction perceived by the stakeholders. If all the weights k_i are equal and $\sum k_i = 1$, then each goal and attribute is equally important and contributes additively toward mission satisfaction. If all the weights k_i are *not* equal but $\sum k_i = 1$ is still true, the higher-weighted goals contribute more towards overall satisfaction than the others. If $\sum k_i < 1$, then the combination of goals provides greater returns to the stakeholders than simply the sum of the parts due to synergy between goals. In contrast, if $\sum k_i > 1$, then the stakeholders begin to experience substitution effects or redundancies.

The second level k_{ij} weights are the attributes associated with the utility of how well an individual science goal is satisfied. This level exists because there may be multiple options to satisfy a science goal or multiple elements that can be joined together to satisfy that goal. For instance, the surface temperature of a planet could be measured by both an in-situ thermometer and an orbiting radiometer; if the instruments are measuring the same effect, there is a high probability that $\sum k_{ij} > 1$. Alternatively, combinations of data sets from different instruments may complement each other to provide synergy, and in these cases, $\sum k_{ij} < 1$.

What is important here, and what makes this method different from the standard implementation of MAUT, is that at both of these levels, a design *could* have $U_i < 0$ or $U_{ij} < 0$ but still have $MAU > 0$, meaning the stakeholder could still be satisfied with the mission despite unacceptable values for certain attributes, depending on the mission and whether or not every attribute is necessary for the stakeholder to achieve satisfaction. For instance, if multiple science goals are being considered, a mission may not be able to satisfy one of the goals at all, but it may still satisfy one or more. Additionally, at the instrument level, there could be multiple instrument options that could satisfy a science goal, some of which are not present in the mission — but having at least one instrument present that can create satisfaction would give the overall mission a satisfactory utility rating.

Decisions about which U_i and U_{ij} can be less than zero yet still produce a positive MAU are dependent on the needs of the stakeholders and the definition of the mission. For instance, if time is a factor and a mission must be completed within 10 years, any designs that would require more time would be unacceptable to the stakeholders. Time could be considered a primary mission attribute in this case, and is mostly relevant to missions that leverage *census sampling* and *stacked sampling*. At the instrument level,

any science goals that can only be satisfied by having more than one $U_{ij} > 0$ would also require careful consideration on a case-by-case basis.

The most important consequence of this structure in MAUT is that measures of synergy and redundancy are a natural byproduct. Selva's methods of synergy matrices introduced in Section 2.3.5.2 show both the synergistic and redundant components, while this implementation in MAUT only shows the sum of the two. However, implementation in MAUT is much simpler in comparison, and only the *sum* of the redundancy and synergy is relevant for ranking design alternatives. (Synergy and redundancy will be discussed in more detail in Section 4.2.3.5.)

A single person *could* have the expertise and authority to define the attributes, weights, and utility functions at all levels of this hierarchy, but in reality, a DSS mission might have multiple stakeholders at these different levels that would have expertise in different aspects of the mission. For instance, a research analyst or instrumentation scientist would have the knowledge and expert opinion necessary to define the attributes, utility functions, and relative weights for a single instrument (k_{ijk} and U_{ijk}). Another project scientist that may be responsible for accomplishing the science goal would trust the expert opinion of the stakeholder at the lower level when defining the other goal attributes and weights (k_{ij} for all j and U_{ij} for $j > N_i$), and the highest level PSS or PI that is responsible for defining the weights among the top-level science goals and primary mission attributes (k_i for all i and U_i for $i > N$) would trust the inputs from the project scientist. The MAU functions at one level can be treated as SAU functions in the parent level with a one-to-one mapping, hence why PIs and project scientists would not define the utilities of attributes that are aggregates defined at a lower level (i.e. U_{ij} for $j < N_i$ and U_i for $i < N$).

4.2.3.3 Value in Multiple Measurements

There are many ways that multiple measurements can add value to a space science mission. Understanding how utility is gained with multiple measurements is crucial to expressing the value proposition of DSS. For strictly redundant measurements, there is little value in multiple measurements beyond the original measurement. This is a clear case of diminishing returns, though there is still value in redundancy when risk is well understood.

For missions that leverage *simultaneous sampling*, there are clearly increasing returns as more assets are added after the first, but there are also threshold returns as the number grows higher. In the case of the STEREO mission, which has two satellites, most of the value comes from having more than one viewpoint in space; however, adding a third satellite may add little value in comparison to the value of adding the second. A scientist being interviewed as a stakeholder may not care about a third viewpoint except as a redundant option, or he or she may be able to use the data from that third viewpoint to strengthen their predictive solar weather models in a way that is more valuable than the cost of a third spacecraft. This is an important distinction that this approach has the power to represent well with proper value elicitation; stakeholders may have preferences on both the “ends” (e.g. redundancy to reduce risk) and the “means” (e.g. value of multiple measurements), and both of them must be accounted for, especially as a way to express the added value per asset in a DSS mission when examining designs over the course of their lifecycle in EEA.

In value modeling, it is important to understand the returns regardless of cost. Whether the mission needs two, five, twenty, one hundred, or one thousand assets, at some point, there will be diminishing returns. At this point, adding more to the system will produce little or no additional value.

Understanding this threshold is critical to properly examine how monolithic systems are inherently limited in their ability to make complex scientific discoveries compared to DSS. When the stakeholder's expectations are aligned with a traditional viewpoint, where one asset is good enough, but that asset needs to be high quality, a monolithic system will most likely be the most cost-effective solution to satisfying the stakeholder's needs — even when using advanced tradespace exploration methodologies to explore concepts that utilize DSS. After all, this model has proven to be successful for many years of space exploration.

If, on the other hand, the stakeholder was interested in a phenomenon that covered a wide area, there may be no limit to the number of assets that could be deployed. For instance, global warming studies benefit from centuries of noisy data (low “measurement resolution”) from measurements made over the entire globe rather than one thermometer with extreme precision (high “measurement resolution”). The stakeholder's perception of value may be based on a resolution of the wide-range phenomenon in both space and time. It is important to consider *measurement, spatial, and temporal resolutions* as part of any value model for space science missions. Failing to account for all three would result in a designer erroneously choosing a monolithic system over a distributed system to study a phenomenon that is better studied by a distributed system.

In contrast, a designer who is too enthusiastic about DSS can make the critical mistake of improper attribute selection. Simply choosing “number of measurements” or “number of assets” as an attribute is not necessarily the correct course of action to model the value of each asset. Instead, a designer should think about what is being accomplished by having multiple assets. If the assets are searching for something, the attribute could be the probability of finding it. If the assets are characterizing a large phenomenon or covering a target planet, various resolution or coverage metrics would be appropriate to use as attributes. “Number of assets” is a great design variable; it is not a great stakeholder attribute.

4.2.3.4 Eliciting Single-Attribute Utility Curves

Once a mission's attributes and the hierarchy of how they contribute to overall mission satisfaction have been defined, SAU curves can be elicited (see Figure 2-2). A software tool called ASSESS™ was used in this research and can help guide the interviewer and PSS to construct attribute definition curves using as few as one intermediate point [261], [262]. First, the lower and upper bounds on an attribute must be given. These are the “minimum acceptable level” and “maximum perceived levels” of the attribute.

(Note: attributes usually take the form of numbers and can be on any scale, including natural or constructed/artificial scales, continuous or discrete scales, and ratio or ordinal scales. Mathematically, these numbers may be called a “value.” To avoid confusion with other definitions of “value” in this work, the word “level” will be used to refer to the number that an attribute of a design takes. For instance, if “maximum speed” were an attribute for a car, the attribute level of one particular car may be 200 km/hr).

The “minimum acceptable” level corresponds to the lowest possible attribute level that the stakeholder would consider accepting. This may correspond to some limit that the stakeholder perceives; for instance, if “detector resolution” were an attribute, the minimum acceptable resolution may be close to or below the

size of the object of study, or, if the system is a continuation of a previous campaign, the minimum acceptable resolution might be some small leap better than that. This level is defined as $SAU \equiv 0$. Any design evaluated with an attribute level worse this point is considered unacceptable and not shown in the tradespace. Typically, designs that are unacceptable are defined as $SAU \equiv -1$, and under the traditional MAUT hierarchy, if *any* $SAU = -1$, then $MAU = -1$. Under the hierarchy described in Section 4.2.3.2, there could be exceptions depending on the stakeholder's preferences.

Eliciting a “maximum acceptable” level for an attribute can be quite a bit more difficult. For the “detector resolution” example, the maximum acceptable level could be as strict as the wavelength being observed. A stakeholder may also have a more reasonable idea of a maximum detector resolution that is relative to the object of study. This level is defined to be $SAU \equiv 1$. Any design evaluated with an attribute level that is perceived to be better than this point is considered to be saturated and is assigned $SAU = 1$ because the stakeholder no longer perceived extra value from improving this attribute.

(Note: “Better” and “worse” in the previous paragraphs could refer to either higher or lower attribute levels, depending on the attribute. For instance, “resolution” typically is perceived as better when the attribute level decreases, whereas “maximum speed” typically is perceived as worse as the attribute level decreases. As long as the function is monotonic, it can be used as an attribute utility function.)

If an attribute like “detector resolution” has a very limited window between the upper and lower bounds, it may still be important to the stakeholders, but there's a chance that the design variables that drive the attribute will make very small differences. If this is the case, it would be better to assume that the attribute is a parameter or a requirement instead (e.g. a binary attribute as acceptable or not). The concept of “requirement” should be avoided unless necessary because requirements inhibit the creative process to explore alternative designs, and may mask the difference in value delivered between design options. If requirements are necessary, make them traceable with requirements traceability matrices.

While the attribute levels of consumer products may have easily identifiable upper bounds, in space science missions there is always a tendency to want to continue to push the limits of what is available. Having a scientific stakeholder that is new to the MAUT process assign an upper limit to an attribute usually ends in one of two ways, based on the author's experience interviewing scientists in the case studies herein. These could be framed as the “dangers” for interviewees that are new to MAU thinking. First: the stakeholder may think far too timidly, resulting in an upper limit that is barely above state-of-the-art compared to what has already been accomplished by monolithic missions. This is an example of anchoring and availability biases. If this happens, the tradespace may not fully explore designs that are feasible, so ground-breaking missions that can deliver high scientific value may not even be examined, and the designs that RSC assigns a high utility to may be nowhere near the actual maximum satisfaction of the stakeholder. Second: the stakeholder thinks too far out of the box, and provides an upper limit far above what is possible given current technology or their budget. It is *much better* to err on the higher side with grandiose visions than it is to be limited by a low upper bound.

Once the upper and lower bounds are identified, the interviewer uses ASSESS to elicit points to generate the SAU curve. ASSESS shows the stakeholder different scenarios depending on which of the four utility assessment methods are chosen: certainty equivalence with constant probability, certainty equivalence with variable probability, probability equivalence, or lottery equivalence/probability (LEP). De Neufville shows that the LEP approach results in less bias when eliciting SAU curves [263]. The stakeholder is

presented with a choice between two probabilistic scenarios that help pinpoint the risk attitudes of the stakeholder and what he or she would be willing to accept. Points on the SAU curve are determined when the stakeholder is indifferent about a choice between two options. Examples of these choices for certainty equivalence (left) and lottery equivalence (right) are shown in Figure 4-6.

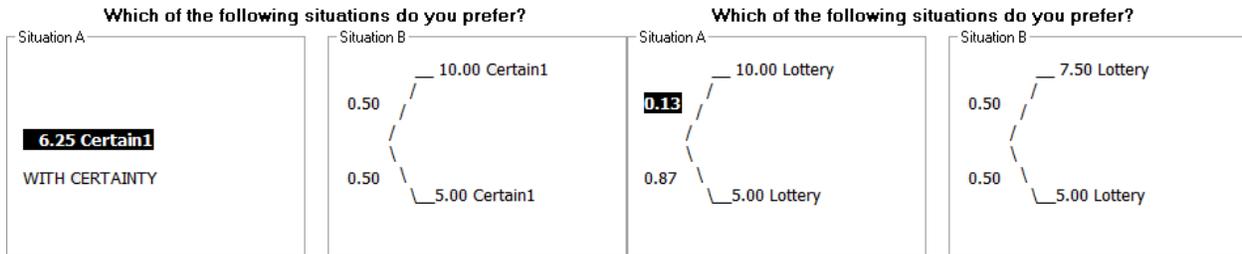


Figure 4-6: ASSESS software examples with certainty equivalence with constant probability (left) and lottery equivalence/probability (right) for eliciting a single attribute utility curve.

In the event that the stakeholder chooses an upper bound that is too high, elicitation of points along the attribute curve using ASSESS should naturally create a utility function showing diminishing returns. For instance, if a stakeholder chooses an upper bound on the attribute “Number of Viewpoints” of 100 when 25 viewpoints would be beyond revolutionary, elicitation with ASSESS may show that 25 viewpoints gives an $SAU = 0.95$, meaning the attribute utility curve has a long, low slope between attribute levels 25 and 100. In contrast, if the stakeholder chose a low upper bound on “Number of Viewpoints” such as 5, the design space may not even explore designs that could add considerable scientific value based on the increase in viewpoints.

This is why multiple iterations on the value model and validation with tradespace data in follow-up interviews are essential before declaring the results of RSC finalized. The interviewer does not need to elicit the perfect value model on the first iteration; it only needs to be good enough to motivate the design variables. A stakeholder may change his or her mind and update the attribute bounds or a point on the curve in the middle of the interview after seeing the final result, which is perfectly fine. It also helps if the interviewer is familiar with the science field to understand what motivates those boundaries.

Elicitation of points on the attribute curve is a proxy for the risk that the stakeholder is willing to take. The reason the utility curve should naturally flatten out if the upper bound is too high is because a stakeholder would be unwilling to risk the possibility of a lower but still highly satisfactory attribute level for the upper bound. Even a risk-tolerant scientist would not be willing to risk much more to achieve the relatively unattainable upper bound (such as 100 viewpoints) if they already have something that still provides incredible satisfaction (such as 25 viewpoints).

4.2.3.5 Eliciting Multi-Attribute Weighting Coefficients

Once the utility curves are assessed, the relative weights can be assigned among attributes at the same level in the MAU structure as shown in Figure 4-5. ASSESS presents the stakeholder with a choice between one of the attributes being at the highest acceptable level and all others at their lowest versus a probability between all attributes at their maximum or at their minimum levels, shown in Figure 4-7. In this example, all lower bounds are 5 and all upper bounds are 10, and the user is being asked to pick between situation A, where all attributes are at the minimum levels except for one, which is at its maximum level, or situation B, where there is a probability of all attributes being at their maximum level

or all attributes being at their minimum level. As the stakeholder chooses between various A and B scenarios, the probability changes and eventually converges on that attribute's appropriate weighting.

This is the point where synergy between measurements can be shown depending where the attributes are on the MAU hierarchy. Two example scenarios regarding the weights of an MAU that is a function of only two SAUs are shown in Figure 4-8. On the left, the weights of both SAUs are $k_i = 0.15$ (such that $\sum k_i < 1$), and on the right, both are $k_i = 0.85$ (such that $\sum k_i > 1$). When the sum of the weights is below one, synergy is represented; on the left, the MAU is higher when $SAU_1 = SAU_2 = 0.3$ than it is when either $SAU_i = 1$ and the other is zero. The mathematical representation of a stakeholder who is either unwilling to risk low SAU on one attribute for a high SAU on another attribute and a stakeholder who sees the added value of combinations of attributes is the same.

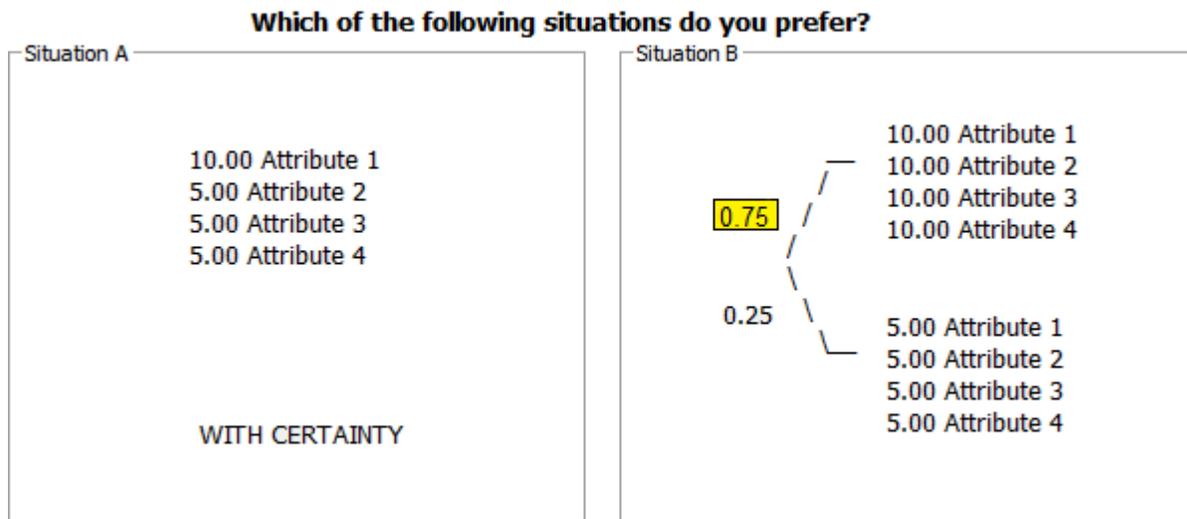


Figure 4-7: ASSESS software demonstrating a choice that will help assign weights to the attributes.

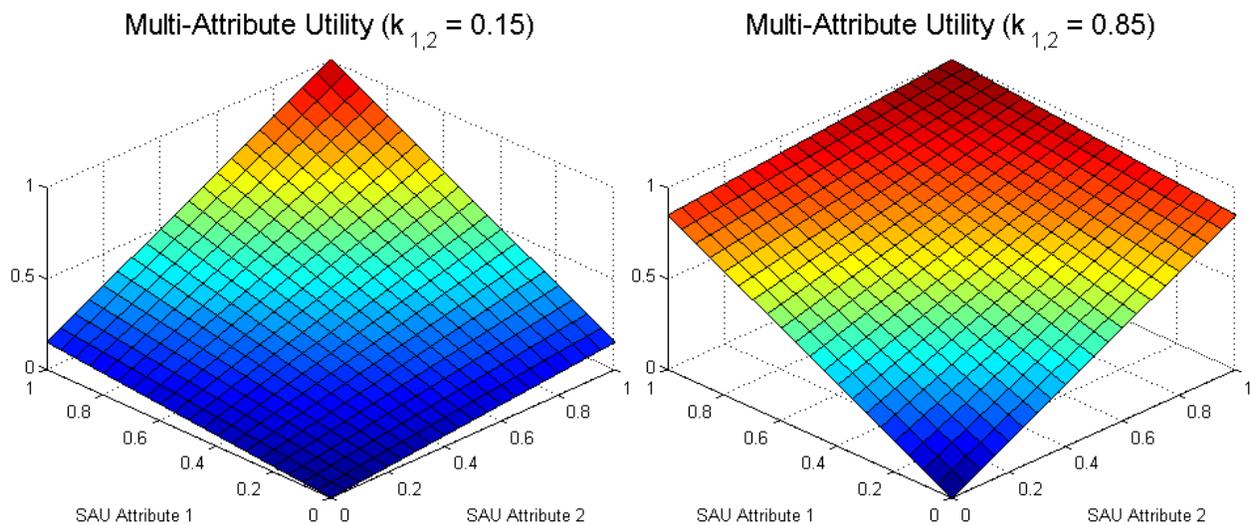


Figure 4-8: Plots of MAU versus SAU for k -weightings that show synergy (left) and redundancy (left)

On the right in Figure 4-8, redundancy, risk-tolerance, and substitution effects are represented. In this case, the stakeholder is mostly satisfied if *either* SAU is near their maximum level. If additional increases

in both SAUs require similar increases in cost, it is more valuable to pursue designs that maximize the value along one SAU than it is to try to achieve both. This would be the case where two instruments could be used to complete a science objective, but only one high-quality instrument can deliver high satisfaction. While this graphic represented an MAU represented by only two SAUs, the interpretation still applies for more SAUs and the weights are elicited in the same way.

In the event that $\sum k_i = 1$, the multi-dimensional shape that describes the multi-attribute utility would be hyperplane (flat), and the slope of the plane along any dimension is equal to the k_i weight of that utility that is represented by that dimension.

In the general case when $k_i \neq k_j$ for all $j \neq i$, if $\sum SAU_{j \neq i} = 0$, and $SAU_i = 1$, then $MAU = k_i$. Graphically, the corners of the surface that represents the multi-attribute utility are at a height equal to k_i . For example, if the k_i weights for a two-attribute system were $k_1 = 0.50$ and $k_2 = 0.15$, the surface would look similar to the graph in the left of Figure 4-8, but the left corner would be raised to 0.5 (additional examples with varying weights and more than two dimensions are shown in Figure 6-15).

4.2.3.6 Mental and Constructed Value Models

In some instances of eliciting value models from stakeholders to describe their satisfaction with complex systems, the stakeholder's satisfaction may not be most appropriately modeled simply as the aggregation of attributes that describe the system. The value model may be based on specific objectives that are calculable with the metrics that may typically be attributes for other systems, but in reality satisfaction is achieved through a particular metric or set of metrics. In some cases, the stakeholder *may not even know* about these metrics that more accurately describe what the complex system is being built to achieve. In business cases where NPV is the primary VCDM, this is usually not a problem; in systems where satisfaction is not measured by an economic metric like dollars, it can be in certain cases.

A classic example for describing the formulation of MAUT and the attribute selection process is through the purchase of a car. Most drivers have some idea of the attributes of the car they want before they buy it. Some drivers want a comfortable car with excellent fuel efficiency for daily commutes. Drivers in warm climates may heavily favor cars with excellent air conditioning systems, whereas drivers in temperate climates or rural areas may prefer cars with four-wheel drive. Some drivers may desire large trunk space for hauling groceries, or higher crash test ratings to protect their children. In a normal stakeholder interview to elicit the attributes and weights, MAUT is an appropriate VCDM to use in its standard form to describe the stakeholder's satisfaction with the product or system because value is based on a *mental* model, or the perceptions and opinions of an individual.

However, in some cases, stakeholders may have a specific objective or metric in mind for the system or product being designed or purchased. In contrast to a *mental* model, there may be a *constructed* value model that describes the stakeholder's satisfaction. This constructed model usually contains elements that an outsider would assume lead to stakeholder satisfaction, which is why eliciting these attributes through an interview assuming MAUT may reveal *most* of the information needed to model the value, but the aggregation of attributes using MAUT is not the exact expression of value a designer should be satisfying.

One automotive example can be found in fictional street racing. In the movie *The Fast and the Furious*, Brian O'Connor (played by Paul Walker) loses a street race to Dominic Toretto (played by Vin Diesel)

and owes him a “ten-second car,” which is a car that is capable of driving ¼-mile down a straight track within ten seconds [264].

In this example, Dominic does *not* specify what the car’s top speed should be, nor does he have thoughts on how satisfied he would be with certain levels of acceleration, torque, or horsepower. He does not specify what make, model, or year this car must be. He does not specify engine fuel type, nor does he opine on transmission type. He does not put requirements on body color, tire lifetime, or safety standards. He does not specify if it needs a spoiler, nitrous oxide injection systems, or forced induction systems. He does not even restrict the concept to an internal combustion engine and thus leaves the interpretation open for Brian to potentially give him a solar-powered hovercraft (if those are legal under the rules of illegal street racing).

Maximum speed, starting line acceleration, and engine torque are potentially important attributes of a race car under a *mental* value model. They are also ingredients to a physics-based set of equations that dictate the *constructed* value model. A simple performance model to calculate this constructed value model would involve equations of motion and an acceleration profile. A higher-fidelity performance model may involve conducting aerodynamics simulations, modeling tire traction, and calculating the optimal time and location to inject nitrous oxide into the fuel-air mixture.

If a constructed value model *does* exist, this does *not* mean that MAUT is no longer applicable for creating the stakeholder’s value model. Attributes such as serviceability, safety, and engine lifetime may also be potentially important to street racers, and those attributes could influence their decision-making process if they have a choice. These attributes could also be considered within the MAUT hierarchy with potentially lower weightings relative to the constructed model, which could now be considered just like any other constructed attribute. Stakeholder inputs are also still necessary to understand how satisfaction increases with the constructed attribute. In this case, the constructed attribute is would be “the number of seconds to drive ¼-mile,” and ten seconds corresponds to the minimum satisfaction level. There were no other inputs on at what level maximum satisfaction is achieved or whether gains in this level are linear, increasing, diminishing, or threshold, but those could be ascertained via a stakeholder interview. An example of a MAUT hierarchy tree where a single constructed model dictates stakeholder satisfaction is shown in Figure 4-9. An example where a constructed model forms only a part of the stakeholder’s satisfaction is shown in Figure 4-10.

In space science missions, satisfaction of some goals may be described more accurately by a mathematical function rather than by an aggregation of attributes. These functions may be rooted in performance calculations such as the ten-second car example, but they are more likely to have roots in statistics, depending on the mission. Identifying these constructed value models or attributes is especially important when considering DSSs because of the statistical variance that is associated with multiple measurements, risk, and opportunity that are inherent when certain emergent capabilities are leveraged to achieve space science goals.

Classical statistical metrics like expected value, confidence interval, distribution descriptors, and variance may be attributes that are important to a stakeholder that a DSS is far more capable of satisfying compared to monolithic systems. When the science goals of a mission go beyond simply probing a place that has never been explored, these metrics become important because they can be used to find the

balance between the raw capability of fewer, larger satellites and the broad capability of many, smaller satellites.

Identifying these metrics is difficult when the stakeholder is completely unaware of them. This is why iterations on the value model after the performance model shows initial results are important, and why it is necessary for the interviewer to be familiar with the scientific literature associated with the field of study. It is especially helpful when the designers notice patterns and can create other outputs that may not have been deemed necessary from the initial interviews. These additional outputs may sway a decision maker, causing them to reevaluate their own perceptions of satisfaction based on the tradeoff between instrument fidelity and measurement abundance.

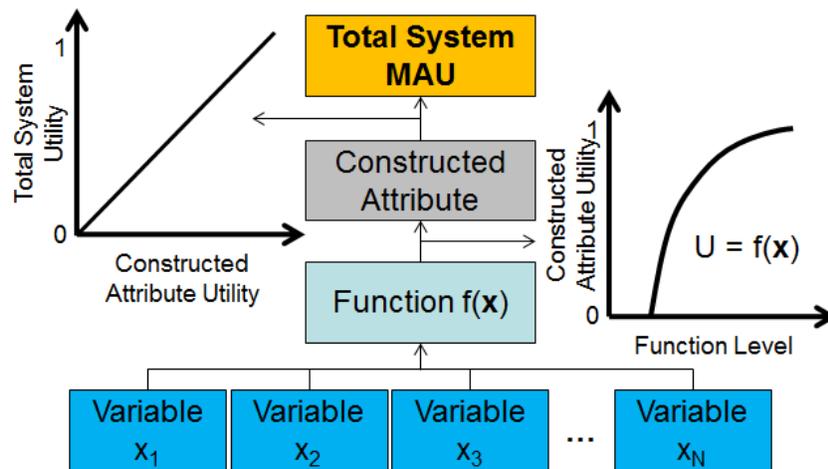


Figure 4-9: Potential MAUT hierarchy with one constructed model that describes stakeholder satisfaction.

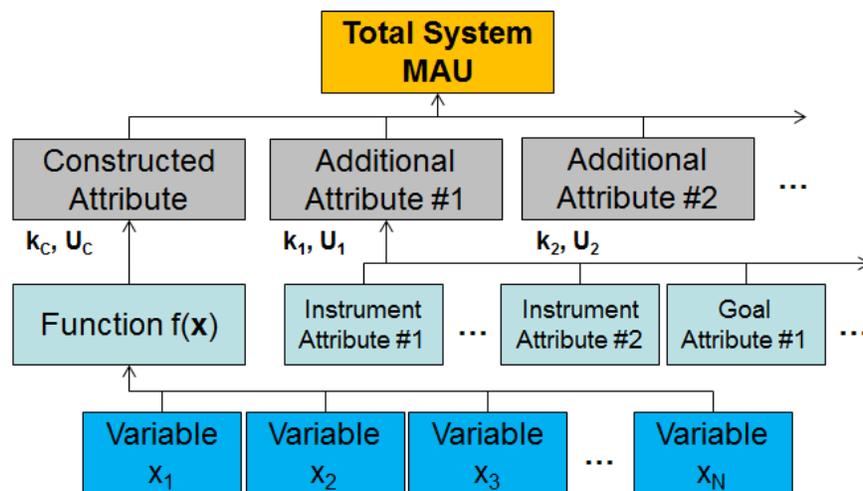


Figure 4-10: Potential MAUT hierarchy where a constructed model or attributes is a part of what describes stakeholder satisfaction.

4.2.4 Value Modeling Process 3: Epoch Characterization

In the final value modeling process, key uncertainties associated with the mission and possible future contexts are identified and contextual variables are parameterized based off of these. Uncertainties and ambiguities in stakeholder needs are also elicited to understand how the stakeholder's expectations can change, either as a result of such context changes or from any factor. These stakeholder uncertainties can

be modeled as separate value models entirely, not only with different weightings on the attributes but also a different set of attributes entirely. The epoch space is the complete set of all possible combinations of unique contexts and stakeholder expectations. A shift from one epoch to another represents the realization of particular levels of the uncertain context and stakeholder needs at a point in time.

Context variables, which are value-driving factors that affect the *performance* model and are beyond the control of the designer, would typically be non-adjustable parameters in a static modeling sense. Examples include federal safety regulations requirements for airplanes, the price of gasoline of cars, cost models and estimated pricing discount schemes, and availability of certain services or technologies.

Epoch shifts are discrete, value driving events that change *either* the value model due to changes in stakeholder needs and expectations *or* any context variable, and these epoch shifts can happen for a variety of reasons. It is impossible for the stakeholder to know and the designer to elicit all possible future value models the stakeholder may have, but knowing what could change their perceptions of the way value is delivered in missions could affect the design that is eventually chosen. Examples of these include competitor products that could launch and make the perception of the chosen design seem worse in comparison, changing or reprioritizing mission goals due to political climate, and changes in the stakeholders themselves (e.g. a decision maker retires before the end of the system lifecycle). Changes in context variables can also result in changes to stakeholder value models if they have some impact on the perceptions of the needs of the stakeholder.

4.2.4.1 Epoch Shifts Due to Exogenous Mission Factors

Epoch shifts due to factors outside the control of a mission can happen for a number of reasons. Researchers have used RSC to examine epoch shifts due to a number of factors including changes in the environment a system operates in, changes in primary objectives, and changes in stakeholders and their value models [54]. Space science missions are less prone to epoch shifts compared to the Navy study in Schaffner's thesis [58], but there are several important shifts to recognize from previous space science missions.

The most easily recognizable and frequent epoch shift in space science missions occurs when a mission reaches the end of its planned lifetime and receives an extension. Mission extensions are especially attractive because they represent ways to gain value with zero acquisition costs; the only costs are for operations, which are usually a fraction of the budget of most planetary science missions. If an analogy to a business model that uses NPV, this is like being handed a free factory, where the only expense is the unit (recurring) cost of the product, essentially creating a no-risk situation for investors (though going back to the science mission, the "product" is scientific data).

One historical example of this epoch shift is the Mars Exploration Rovers (MERs). Both Spirit and Opportunity were originally designed to operate for only 90 Martian days (sols). During these 90 sols, the rovers maintained their expected utility and delivered value according to a timeline that was expected from the design. However, after 90 sols, the rovers continued to operate and deliver value, so the mission was extended. Throughout the course of their mission lifecycle, they have far exceeded their original expectations. A mission being designed with RSC needs to consider how a mission extension could affect the how the stakeholder perceives the value delivered by a mission (though planning for scenarios where the potential operational lifetime is 4,100 percent greater than its planned lifetime may be out of scope for most missions).

Assuming the system operated at its designed utility level for the entire mission, it would have delivered known value up to that point. For example, a design with $MAU = 1$ would have delivered maximum value, so there would be no reason to continue operating *if one assumes the value model has not changed*. However, if a mission is extended, someone has decided that continuing operations must be worthwhile because the mission will continue to deliver value through the course of the mission extension. If the design's $MAU < 1$, extending the mission does not necessarily mean that the lifetime attribute alone can bring about full satisfaction compared to a design with $MAU = 1$ that operated for exactly the planned mission lifetime. The mission lifetime attribute is already at the maximum level, so no additional utility or value can be gained *according to the original value model*.

An epoch shift because of a change in the stakeholder value model clearly happens and must be accounted for. The change in the value model may be as simple as modifying the mission lifetime attribute; if this is the case, all designs are affected equally within an era that invokes a mission extension. However, if at the end of the mission, another stakeholder “acquires” the mission from the first stakeholder, or if the first stakeholder has a different idea about what will bring satisfaction during the mission extension phase, the new value model could undergo significant changes. These changes could include different weights, different utility functions, different minimum and maximum levels, and even entirely different attributes.

This problem is illustrated for a single attribute (mission lifetime) in Figure 4-11. The primary mission may reach its maximum utility after five years in the initial value model, but if a mission survives that long and receives an extension, a new value model needs to be applied to the extension. Exactly how valuable the extension is (center) compared to the primary mission (left) depends on the slopes of the comparing lines of the two on the same timeline (right). The slopes determine the “rate” at which value is delivered during the mission. In practice, one would conduct an additional stakeholder interview after the mission to determine the value of a mission extension and assume a new baseline starting where $t_{0,extension} = t_{end,primary}$ and minimum satisfaction would occur starting at the new t_0 , but for prescriptive methods like RSC, having this information in advance is useful.

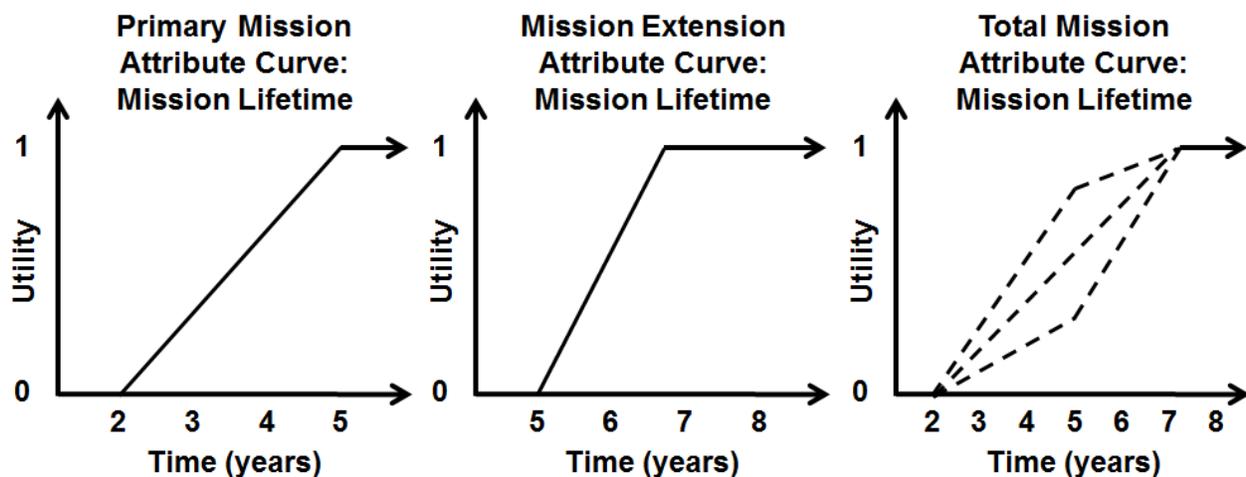


Figure 4-11: Examples of attribute curves for the primary mission (left), a mission extension (center), and possible ways to interpret the combination of the two (right).

Despite the fact that no one would argue that the primary mission delivered less than expected, the utility never goes above 1. If *only* the mission lifetime attribute changes, the attribute curve for the mission would remain flat and scale to a new maximum $SAU = 1$ because of the affine nature of the other attributes in the formulation of MAUT; however, if other attributes change, comparing the cumulative values across this epoch shift can be difficult. This will be discussed in more detail in Section 4.5.8.

These mission extension value models are especially difficult for people who are unfamiliar with MAUT to comprehend at first. It is counterintuitive to think that a mission that has delivered complete satisfaction ($MAU = 1$) the day before the mission extension somehow delivers less utility than that the day after. It is more natural to think “I already have something, and if I go for more, the number should be higher” rather than “I already have something; if I want more, that means I instantly have less.” The reduction in utility is strictly a consequence of how this implementation of MAUT strictly defines 0 and 1 as the minimally acceptable and maximally perceived levels. Having a mission that extends to deliver “double” the value does *not* mean the utility jumps from 1 to 2, but rather that the stakeholder has simply changed his or her mind and redefined what the maximum perceived value of the mission can be.

In monolithic science missions, applying this type of value structure in different contexts makes little sense. However, in dynamic and distributed systems, and especially changeable systems, capturing the changing contexts is important because mission lifetime is not dependent on a single spacecraft. Upgradeable systems need to show that the added value is justified under these future scenarios. It is easier to use NPV to show the value in system replenishment in a communications constellation that is generating revenue, especially if that revenue is growing and demand is increasing towards system capacity [68], [70]. A space science mission typically does not benefit from the same sense of urgency or opportunity, especially if long-duration measurements begin repeating themselves and add no real value to the data sets that have already been gathered; however, this framework helps show those opportunities explicitly and allows direct comparisons in value between designs that cannot seize opportunity and those that can.

Another epoch shift a designer could encounter with a space science mission with an especially long lifecycle is a change in stakeholders. Historical examples include the Voyager probes and the ISS. Many of the original scientists who help define the Voyager missions have passed away, and new scientists have different expectations for what they want from the spacecraft now. ISS is particularly vulnerable to epoch shifts due to the political nature of the multinational effort to keep it operational, and if one nation pulls out of the agreement, a new value model would be needed to best represent how the station could maximize its value delivery with the change options it has. An epoch shift such as this is outside the scope of this work but is mentioned for completeness.

4.2.4.2 Epoch Shifts Due to Endogenous Mission Factors

Epoch shifts can also happen due to scientific information gained during the course of the mission. When new knowledge is gained, a stakeholder may suddenly have a new idea of what measures mission success or a desire to change the course of the mission to pursue data related to this new knowledge. This type of epoch shift is especially important for showing the value of changeability in systems, specifically operational flexibility. In this work, this epoch shift shows how missions that leverage *staged sampling* and *sacrifice sampling* can capitalize on opportunity that arises during the mission lifecycle or mitigate losses as a result of the failure of assets.

A quote from Carl Sagan succinctly summarizes how science is valued: “Extraordinary claims require extraordinary evidence.” To the scientific community, it is more valuable to have an abundance of evidence for a claim compared to too little. Furthermore, it is extremely difficult (if not impossible) to scientifically prove that something *does not* exist: “The absence of evidence is not the evidence of absence.” Missions that leverage multiple assets can not only be distributed to explore a wider range of targets or environments to gain confidence of something’s nonexistence, they can also provide additional evidence for the burden of proof in the event of a major scientific find. If one asset makes a discovery that could fundamentally change the way scientists view a target body, they may have a desire to explore it with another asset; therefore, the value model that governs the remainder of the mission has changed from its initial state.

There are two possible operational situations and two possible scientific outcomes from deploying an asset in a mission, creating four possible ways a scientist could view the success of a mission. These outcomes are summarized in Table 4-2.

Table 4-2: Possible outcomes of a science mission and operational situations with multiple assets.

Asset Situation	Nothing noteworthy discovered	Something noteworthy discovered
No Spare Assets Remain	Ignorance is bliss. The limited data set leaves too much uncertainty to be conclusive with a non-result, but this is not enough reason to justify Congress spending more on a follow-up mission.	There is no opportunity to repeat the measurement and have more statistical confidence. Debate in the community continues for years before a follow-up mission can be designed and launched.
Spare Assets Remain	Thorough confidence in the ability to report conclusive findings without undue skepticism is achieved, and the scientific community now has a solid, disappointing, but valuable non-result.	The mission capitalizes on the opportunity to confirm the discovery with additional observations and data. This data leaves little room for debate over the possibility of this being an incorrect result.

If a mission brings spare assets to deploy to sampling locations, a traditional value modeler may think that the added mass would be wasted, and resources could have been better spent improving the existing assets. However, if there are no spares and any asset fails, there would be less satisfaction than the mission expected compared to one that managed to deploy to all the sampling locations despite asset failure. With no failures and spare assets, the stakeholders have the option to explore the same locations again or identify new locations, which would also create an epoch shift similar to a mission extension since there would be some additional value gained by this additional exploration.

Spare assets also solve the problem of satisfying the needs of multiple stakeholders with clashing priorities in the deployed environments. In exploration missions where a scientific discovery is uncertain, such as in the cases where a mission’s primary goal is to search for evidence of life, being able to explore multiple arenas is useful when there are multiple stakeholders. For instance, there was debate over where exactly to send the Mars Science Laboratory. Scientists from different fields of expertise were split between two locations, and because there was only one rover, multiple groups left dissatisfied with the outcome. Leveraging multiple assets help alleviate this issue and satisfy a wider range of stakeholders.

4.2.5 Key Assumptions and Limitations

Most of the limitations of MAUT and its relation to RSC have been discussed either in Chapter 2 or Section 4.2.3.1. These limitations will be revisited at the end of the chapter. The limitations discussed in

this section deal with modifications or assumptions that have been proposed and the ramifications that affect processes further along in the RSC method.

When developing the MAUT hierarchy for a mission, there may be a primary mission attribute that affects all science goals but relates differently to each one (e.g. different goals have different desired observation lifetimes). If so, these primary mission attributes would need to be brought down one level into each goal as a primary goal attribute. If this happens, it makes the cumulative value calculation process as described in Section 4.5.6 more difficult to implement than it is presented.

The value that can be gained through opportunity is difficult to predict. In space science missions, it might be considered irresponsible to claim that a mission had a finite probability of making a discovery that would change the world, or upend most of the knowledge that has previously been gathered. Estimating the probability that an epoch shift will occur and how much that shift would change the value model depends on what could happen during the mission. Since each mission is unique and the opportunity that may arise is unpredictable (and may not be present at all), determining the cumulative value distribution in later stages is difficult to do. Instead, designs can be compared to how much opportunity they can seize given certain probabilities that affect epoch shifts.

There may be significant synergy between two or more top-level goals, but little or no symmetry between those goals and any other top-level goal. For instance, a mission could have two top-level scientific goals related to atmospheric chemistry and one related to regolith chemistry. In this case, it may be appropriate to combine those synergistic goals as a single top-level goal. It is important for the interviewer to remember that if this is the case, the new top level goal encompasses both goals and should have a higher weight than either goal had before combining them. An example of this additional layer in the MAUT hierarchy is shown in Figure 4-12.

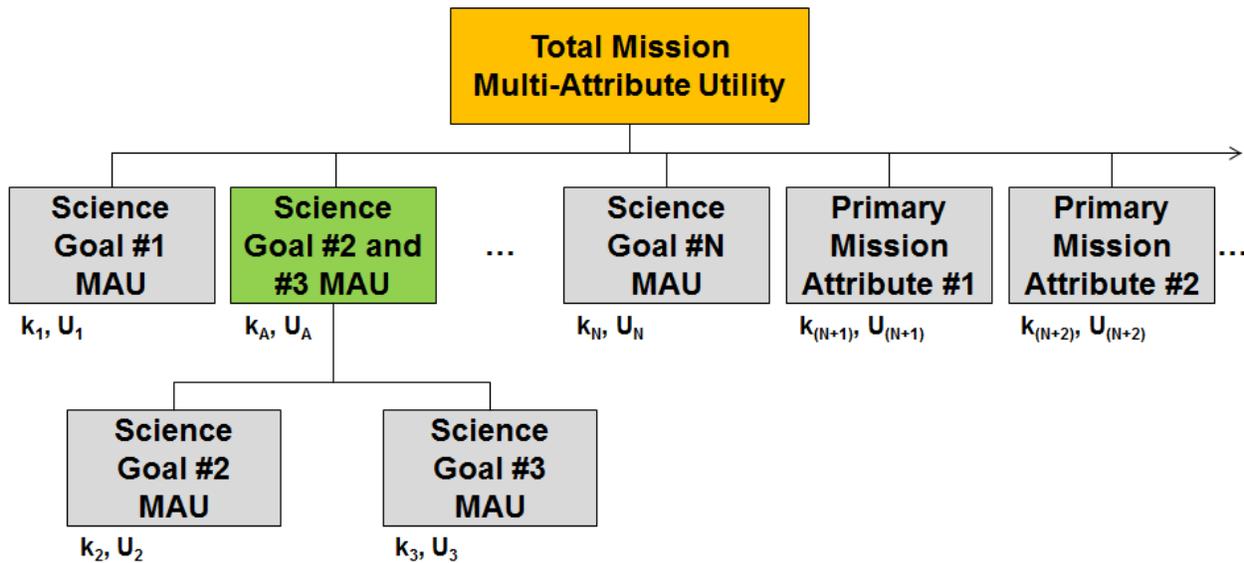


Figure 4-12: Additional modification to MAUT hierarchy that more accurately expresses synergy between a set of top-level goals that does not express additional synergy with the remaining top-level goals.

4.3 Phase 2: Performance Modeling

In this phase, the attribute definitions of the value modeling phase are used to motivate architectural and design variables. The end goal of this phase is the full or partial enumeration and evaluation of the tradespace in each epoch so that overall performance (total system MAU) can be compared to the cost of the design.

After the attributes and their upper and lower bounds have been identified, the next step is to generate concept architectures, identify value-changing design variables, and build a system performance model to enumerate and evaluate as many designs as possible in the tradespace. Designers build mathematical models and simulations that use the context and design variables as inputs.

There also may be intermediate variables and other non-varying parameters that are critical to understanding performance. Parameters that the designer can control but chooses not to should not drive the value of the design as much as the other design variables, but it is important to check that they do not. The performance model will output the attribute levels and cost of each design. The attribute levels are then fed into the value model to output the MAU of every design. A simplified block diagram of the entire performance modeling process and how it works in conjunction with the stakeholder value model derived in the previous phase of RSC is shown in Figure 4-13.

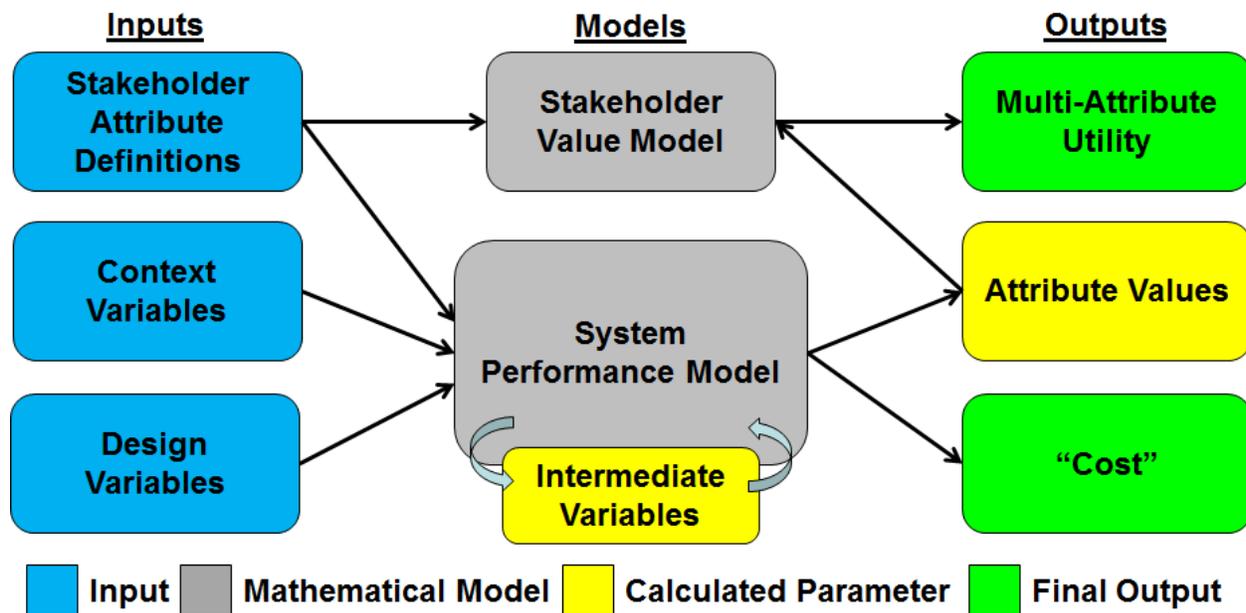


Figure 4-13: Block diagram of the system performance and value models

4.3.1 Contextualizing Phase 2 for Distributed Satellite Systems

Performance modeling is what engineers excel at, and there are only a few changes between modeling the performance of DSSs versus monolithic systems. One of the important considerations is cost reduction. Most space missions use mass alone to estimate cost at the beginning stages. In distributed systems, the mass penalty associated with so many redundant subsystems is partially offset by the decreased recurring costs associated with learning curve improvements. Proprietary cost models are relatively immature at handling DSS, but they will get better with more real-world data from future DSS missions.

Another important factor in DSS is how an attribute like “coverage” is defined. This is a problem Earth observation constellations deal with because there are many ways to define “coverage” depending on the needs of the stakeholder. This is not just a communications issue during the Value Modeling phase or an example of a stakeholder ambiguity in their needs, but it is also an issue in Performance Modeling. Wertz provides examples of how adding satellites to a constellation counterintuitively decreased the performance metrics that are being used to evaluate the constellation, despite the fact that the stakeholder was clearly more satisfied with the system as a result of those additional satellites [106]. This is because the statistical metrics used can be misleading.

To avoid these problems, explicit coverage definitions in the Value Modeling phase help, but the designers need to be sure that they are modeling the correct types of coverage in the first place. For instance, in one case study in this work, the coverage definitions were very strict, well defined, and needed to meet many parameters for satisfaction; in another, the average revisit time calculated over a short period was good enough to show the stakeholders and gain their satisfaction.

A problem that arises with multiple concepts being explored in a single tradespace is how to enumerate the design vector. In these case studies, there are distinctions between “architectural variables” and “design variables.” Loosely defined, architectural variables determine what building blocks go into a concept, whereas design variables describe those building blocks. For instance, if one concept uses cold gas thrusters for attitude control, a design variable in that concept could be “gas mass.” A different concept that opted for reaction control wheels would have no use for that design variable. If one concept uses nuclear power and another uses solar power, a design variable like “solar panel area” would clearly not apply to the nuclear-powered architecture.

This leads to minor difficulties in creating the design vector because each distinct concept or architecture would have its own design vector. The design vector is no longer the simple, complete combination of all design variables and all values they can take. When creating design vectors, rules can be applied to prevent problems when writing the complete design vectors, but it is better to know this up front than find out halfway through the process.

For high-fidelity models where computation time is a factor, it is counterintuitively better to *limit* the tradespace by not enumerating all designs, especially if some of the designs are clearly infeasible. For instance, if “maximum thrust” were a design variable that had an option at a low level, it may be pointless to bother calculating whether it can satisfactorily push a large satellite. In general, the philosophy of MATE and TSE is “more designs are better because we can learn something even from bad designs,” because the definition of a “bad” design may change, or there may be contexts or epochs where “bad” is “good.”

Designers may still use heuristic rules to eliminate designs that are clearly infeasible and would otherwise waste time, especially in tradespaces with many design variables. This is usually recommended when extremely large tradespaces combined with models that cannot be simplified or use interpolated data as a substitute (all the case studies in this work have used some form of simplification to reduce computational complexity or redundancy). There are many enumeration and sampling techniques to evaluate subsets of the tradespace, such as a design of experiments, to search for a convergence of insights into the tradespace that would give the designer an intuitive, intimate knowledge of the system and how design

variables affect the value of the design. These techniques are necessary when it is computationally impractical to evaluate the entire tradespace.

A better way to speed up computation time by limiting the total amount of computation required is to split the tradespace into subsets. For instance, coverage calculations may be the same for a subset of designs, and that subset repeats itself within the full set. If this is the case, it is much faster to run simulations on the subspace and store the data so it can be copied over and over as the subspace is calculated within the main performance model. The performance module that requires the most computation time should be at the deepest level of the performance model loop so that it can be repeated the most. This is only an option if certain attributes are only changed by a subset of design variables but will be applied in the case studies, especially the one in Chapter 5.

Finally, when there are enough assets that can operate in a variety of zones, it becomes increasingly difficult to separate design variables from operational decisions. “Deploy two assets to Zone 1, two assets to Zone 2, and one asset to Zone 3” is the same design as “Deploy five assets to Zone 1” if there are no differences in the required performance to get to those zones, but the value delivered may be very different. If this is the case, the design space can grow exponentially as the combinations of options for a single design grows. The options a decision maker would choose are usually dependent on the context variables, but it is not a good idea to wait until EEA to address the operational variables in case critical modeling components are missing before valid decisions can be made.

4.3.2 Performance Modeling Process 1: Identifying Architecture and Design Variables

In the first performance modeling process, the attribute definitions that were elicited in the second process of value modeling are used to motivate and identify the value-driving architectural and design variables. Typically, this is done numerically using a simplified form of QFD, which is a first-order model in and of itself, but there are a number of heuristic methods that can be used. If there is simply not enough subject expertise to trust the results from QFD, modeling can still begin, but this typically requires even more model iterations as more insight into the system is gained with each improvement to make up for the initial lack of direction in the modeling process.

Performance models for TSE and Phase-A studies do not need to be perfect or high-fidelity; however, they must have *enough* fidelity to accurately show how value-driving design variables affect the value or overall utility. This is why it is important not just to know the attributes but also to know the minimum and maximum levels that are being designed to. Different parametric models may apply to different regimes of the design space. A designer may focus too much effort on developing a high-fidelity model that takes too long to compute an individual design when a low-fidelity model gives good enough information for TSE. While it is usually better to err on the side of higher fidelity, it can reduce the speed at which tradespaces are generated, or siphon efforts from other areas that actually need higher fidelity modeling than the designers realize at first. It is up to the engineers to decide on a case-by-case basis what level is “good enough” at this early stage of the design process.

One concept that is borrowed from previous MATE studies is the inclusion of architectural variables in addition to design and context variables. Architectural variables represent which of the many possible assets or components of a mission are present, whereas design variables give more specific information

about those components. For example, if a mission has the option to use penetrators, the architectural variable options for “penetrator” would be “[no yes]”. If the option is “yes”, then any design variables relating to penetrators would need to be enumerated and calculated; however, if the option is “no,” then all design variables regarding penetrators can be set to zero since they would have no effect on the design performance or value model. This strategy helps limit the size of the tradespace by eliminating degenerate designs from the design vector.

4.3.3 Performance Modeling Process 2: Design Modeling

Design modeling is the process that engineers are most familiar with. This is the process where mathematical models and simulations are built to convert design variables into attribute levels, which then translate into utility through the value model.

Although each case study in this dissertation is unique, there are many commonalities among them. Many use the same small satellite components and mass estimates. A complete list of all assumptions identified in each case study is presented in the Appendix associated with that case study.

Other than the factors that have already been discussed in previous sections, the design modeling process for leveraging DSSs is similar to any other engineering design work. This is the process that engineers should have the least trouble with, especially if the modeling is being done in their field. The multidisciplinary nature makes it difficult for a single person to accomplish all tasks alone, but this process is well-suited for concurrent design engineering.

4.3.4 Performance Modeling Process 3: Cost Modeling

If trusted cost models for space science missions such as those presented in Section 2.4.3.2 are not available, then a low-fidelity cost estimate based on total mass is an appropriate way to start the tradespace exploration process and understand what variables might drive cost. However, total mass alone is not an accurate representation of cost, especially with DSSs and the mass penalties associated with the inherent added complexity and subsystems redundancy that comes with them.

We can improve the mass-based cost model shown in 2.4.3.1 in several ways. Once again, that model is

$$C = c_S M_S + c_P M_P + c_F M_F \quad (4-1)$$

where c_S and M_S are the cost multiplier and mass of the spacecraft without the payload, c_P and M_P are the cost multiplier and mass of the payload, and c_F and M_F are the cost multiplier and mass of the fuel. When distributed satellites consist of many repeated components, they can take advantage of the cost improvement and mass production. However, this simple mass model does not separate out recurring and non-recurring costs; the cost improvement formula from 2.4.3.5 does not apply directly in this case without first estimating what the cost of a second unit relative to the first may be. Successive assets may be priced as a flat fraction of the first unit, or as an exponential learning curve from some fraction of the first unit’s cost (usually ~10-15% in both cases). Both of these possibilities are illustrated in Figure 4-14 (the values have been exaggerated for illustrative purposes).

Payloads typically undergo much more testing and development, so each successive unit costs a higher fraction than the first unit compared to other parts of the spacecraft. This can be account for by having separate cost improvement factors for payloads and other spacecraft hardware.

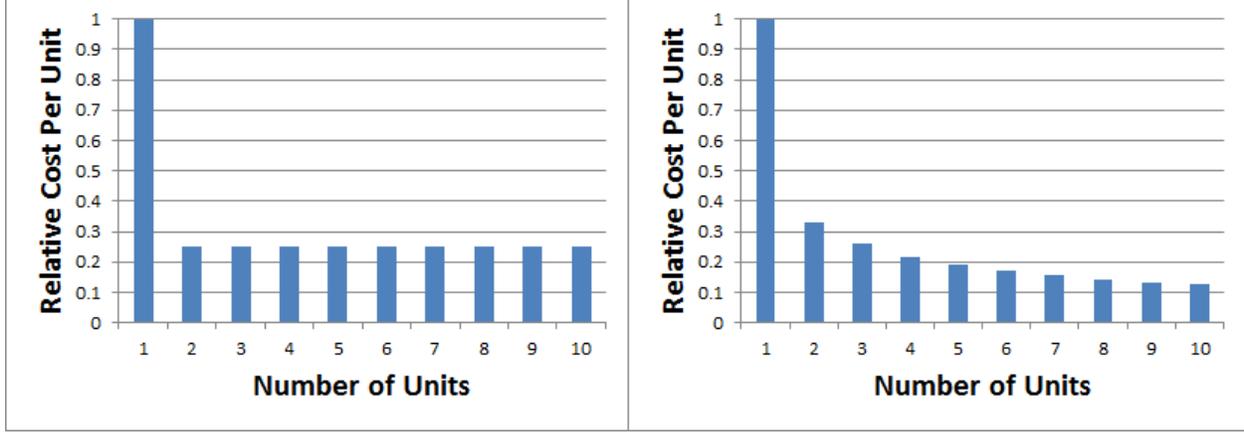


Figure 4-14: Cost models of producing multiple units using (left) constant discount cost and (right) learning curve discount cost.

With these assumptions, a general form of the mass/cost approximation equation that accounts for payload complexity, cost improvement, and heterogeneous assets being launched by a single carrier can be created. For a mission with m distinct assets and n_i copies of the i th distinct asset, the “Complex Cost” of a spacecraft can be estimated as:

$$C_{SC} = c_C M_C + \sum_{i=1}^m \left(c_S M_{S_i} + c_{P_i} M_{P_i} + \sum_{j=2}^{n_i} \left(c_{SD_j} c_S M_{S_i} + c_{SP_j} c_{P_i} M_{P_i} \right) \right) + c_F M_F \quad (4-2)$$

where M_C is the mass of the mothership carrier (if one exists), c_C is the cost/mass ratio of the mothership carrier (which is similar to c_S in the previous equation), M_{S_i} is the mass of the structure of the i th distinct asset, M_{P_i} is the mass of the payload on the i th distinct asset, c_{P_i} is the cost/mass ratio of the payload on the i th distinct asset (this takes into account that some payloads can be more complex and therefore expensive than others but still have the same mass), c_{SD_j} is cost discount on the cost of the j th structure of the i th distinct asset (where $j \neq 1$), which is also dependent on which discounting model is being used, and c_{SP_j} is the cost discount on the cost of the j th payload of the i th distinct asset (where $j \neq 1$). In the event that structural costs scale differently for different types of assets, the cost multiplier of the spacecraft without the payload c_S can be replaced with c_{S_i} .

For example, if a mission carried 8 small satellites, 5 penetrators, and 2 rovers, then $m = 3$, $n_1 = 8$, $n_2 = 5$, and $n_3 = 2$, and the cost coefficients and discount schedules could be completely different for each of them but still be represented by this formula.

Additional modifications can be made if an asset carries multiple payloads, and each payload requires separate cost improvement coefficients. In this case, if the i th distinct asset has l_i payloads, the k th payload would have a mass $M_{P_{ik}}$, a cost coefficient $c_{P_{ik}}$, and a cost improvement coefficient $c_{SP_{jk}}$. This sort of separation would be necessary in an asset that held both a simple, off-the-shelf payload that required extensive testing and a state-of-the-art payload that is expensive to develop but costs little to duplicate and test. Although this separation will not be necessary going forward in this work, the most general representation of this formula is:

$$C_{SC} = c_C M_C + \sum_{i=1}^m \left(c_S M_{Si} + \sum_{k=1}^{l_i} (c_{P_{ik}} M_{P_{ik}}) + \sum_{j=2}^{n_i} \left(c_{SD_j} c_S M_{Si} + \sum_{k=1}^{l_i} (c_{SP_{jk}} c_{P_{ik}} M_{P_{ik}}) \right) \right) + c_F M_F \quad (4-3)$$

In the event that multiple fuel types are being used on a mission, the final term representing the fuel mass could be further adjusted to take this into account, but in this representation M_F is the combined fuel mass for all assets and the mothership carrier.

However, this cost model only accounts for the cost to produce the spacecraft; it does not account for launch vehicle and mission operations costs. Although operations costs tend to be small compared development and launch costs, they are non-negligible. Launch costs are usually estimated by first calculating the vehicle mass and volume, then selecting the cheapest vehicle that is capable of delivering the payload to the desired orbit.

Operations costs can vary depending on the complexity of the mission, if that has not been accounted for in some way in the spacecraft cost model, and lifetime of the mission. A long mission that requires significant use of a communications downlink may have a greater lifecycle cost than a comparable mission that is more independent or can downlink data at faster speeds, thereby reducing the downlink time and costs. This is an important sensitivity in the communications subsystem that does not get examined often with a complete MDO model. Usually, the communications network publishes their cost estimates for the ground segment, which scales with the required downlink time and is quoted by the hour, though there may be startup fees and billable hours for the time for the antennas to point to the target. If the number of downlinks and time length of each downlink is known, the cost to use the network over the mission lifecycle can be estimated with reasonable accuracy.

The total cost of the mission can therefore be estimated as

$$C_T = C_{SC} + C_{LV} + C_{Ops}L \quad (4-4)$$

where C_{LV} is the launch vehicle cost, C_{Ops} is the cost per unit time of operations, and L is the mission lifetime. Mission lifetime can be treated as both an attribute, such as in the case of desired coverage time, and a design variable that drives other intermediate variables, such as in the case of the amount of radiation shielding to use or expendable resources the spacecraft carries (e.g. liquid helium in IR telescopes that naturally escapes over the mission).

4.3.5 Key Assumptions and Limitations

Estimating costs is a major challenge in MATE, as it is with all conceptual design methods and techniques, especially when cost models are not well established. While cost-estimating relationships (CERs) for heritage components with high TRLs have proven to be reliable, cost estimating functions for unproven technologies or new operating modes are difficult to validate or trust. As Walton showed in [28], when uncertainty is not well understood, the assumption that the uncertainty behaves according to the normal distribution is invalid. At the subsystem level, this makes it especially difficult to compare how effective implementing new technologies into specific subsystems can be since these cost models aren't mature. At the mission level, it is difficult to accurately estimate the operations costs of multiple satellites since there are relatively few systems to compare to.

One way to overcome these limitations is to treat some of the cost variables as context variables, or sets of values for cost models as a context variable. For instance, the context variable “Cost Model” could take on values of [Cost Model 1, Cost Model 2]. Cost Model 1 could assume a 10% reduction in fuel mass; a tiered schedule of discount for multiple payloads where units 2-4 are 90% of the cost of the first, units 5-10 are 80%, and all additional units are 20%; a payload mass cost multiplier of 5 compared to the structural cost coefficient; and a constant discount of 15% for each additional spacecraft structure after the first produced. Cost Model 2 could assume a 5% reduction in fuel mass; a constant discount of 20% for each additional payload after the first produced; a payload mass cost multiplier of 6 compared to the structural cost coefficient; and a constant discount of 15% for each additional spacecraft structure after the first produced.

However, if all of these variables and options are treated independently, the context space can quickly expand, and all of these cost variables could easily hide how the tradespace changes with other context variable changes. This could cause some of the designs that would be discovered with Pareto Trace or Value-Weighted Filtered Outdegree (VWFO) to remain hidden from consideration (see Section 4.5.3).

One way to validate tradespace models is to consider where an existing system with known design variables and lifecycle costs fits on the tradespace. If the model closely matches one or more existing architectures, it makes a stronger argument that the model is reliable. However, this requires a mission of a certain type to already exist and be similar enough to what is being designed for it to exist in the tradespace. It will be difficult to validate such models for multiple spacecraft when there are very few systems to compare to.

Because most of the missions in the case studies are going beyond LEO, it is difficult to assume that there would be much opportunity to add additional assets to the system. With ExoplanetSat, this is a very real option; with GANGMIR, sending a few additional penetrator units to Mars would more than likely require an entire new mission, and modeling the opportunities for piggybacking on the next Mars mission opportunity two years later is beyond the scope of this work. Although other DSS studies focused on replenishment, upgradability, and other lifecycle extension decision, this work will only mention where they can be applied but they will not be explored in detail.

4.4 Phase 3: Tradespace Exploration

Phases 1 and 2 are the building phases of RSC, while Phases 3 and 4 are where the real power of RSC can be seen. Phase 3 is tradespace exploration. In MATE and MATE-CON, this is essentially the final phase, and numerous activities can take place during this phase to explore the design space. Traditional optimization typically would choose designs along the Pareto front of the utility and cost dimensions, but TSE allows users to visualize *why* designs are on the Pareto front and recognize when a Pareto design is not as robust or valuable as it may seem by examining all the factors that give a design value.

4.4.1 Contextualizing Phase 3 for Distributed Satellite Systems

In TSE, there is little to contextualize that is specific to DSS that is not already part of the TSE phase. With large numbers of assets, the tradespace may be difficult to parse out and some information may be hidden, but TSE is inherently designed to help.

If a design is implemented with architectural variables that has sets of heterogeneous assets (e.g. a mission could have rovers only, hoppers only, or both rovers and hoppers), it is useful to examine tradespaces of those architectures individually to see how the difference in design drive value. The Pareto fronts for different mixtures of assets may also show how value is driven by the ratio of heterogeneous assets (e.g. all hoppers, 3/1 ratio of hoppers to rovers, 1/1 ratio, and 1/3 ratio, depending on the maximum numbers of each architectural asset).

In missions with multiple science goals, it is definitely useful to examine each goal individually in addition to the tradespace showing total mission MAU. This is especially important to do in missions with synergy. It is often difficult to capture stakeholder preferences on how goals may combined synergistically, and seeing how some options affect the final ranking could cause them to reconsider the science goal attribute weights k_i .

4.4.2 Tradespace Exploration Processes

There are a number of activities that can be performed in TSE. It is helpful to have specialized tools to aid the communications between designers and stakeholders so they both understand large datasets and find details that would be otherwise hidden. Visualization is an important aspect of communication in TSE, which is why SEArI and other groups have invested in applications and software to aid in this process. SEArI's Interactive Visualization for Tradespace Exploration and Analysis (IVTea) software suite allows users to display the tradespace in a number of ways and adjust stakeholder value models in real time.

TSE is also a convenient checkpoint to bring the stakeholders in to compare designs and make sure the value model accurately reflects the stakeholder's needs. If something seems off, the attribute curves, weights, and levels can be adjusted. A value model can be validated by presenting the stakeholder with pairs of designs that are different in their design vector but have the same MAU. If the stakeholder is indifferent between those choices, the value model is in agreement. If there is a clear preference of one design over another, the value model should be readjusted and revalidated.

One important metric for ranking valuable designs (designs that deliver the most utility for their cost) in TSE is Fuzzy Pareto Number (FPN). A design that is Pareto efficient dominates in both cost and utility. Designs can be "fuzzily" efficient if they are within a percentage of the range of the cost and utility data. This percentage is the FPN. A truly Pareto efficient design will require 0% fuzziness to be on the Pareto front, whereas a design that would require a percentage change in cost or utility to be on the Pareto front would have an FPN of that percentage. The FPN of a design d is the smallest percentage K for which the design is in the fuzzy Pareto set P_K :

$$FPN(d) = \min\{K | d \in P_K\} \quad (4-5)$$

It is useful to view the tradespace in several different forms, such as color-coding every point to represent a change in the level of one variable at a time. For each design variable, examine the tradespace with a different color for each possible level. Patterns may emerge in the tradespace from unstructured tinkering with the visualization options. This is especially handy in early iterations of the models to see whether the limits of the design variables are being tested at all, or if the design vector should be expanded to account for a wider range of design variable levels.

4.4.3 Key Assumptions and Limitations

Building tradespaces with a large set of design variables and high-fidelity models requires massive amounts of computing power. While computer models to analyze point designs on the tradespace may take fractions of a second to run on an ordinary machine, each design variable and possible value that design variable can take increases the size of the tradespace exponentially. Simple models are applied up front and then more complex and time-consuming models are applied to the Pareto-optimal designs found in the full tradespace enumeration. Essentially, an inherent limitation of MATE is the computer time and power required. While there are other ways to explore the tradespace, such as design of experiments, parameter studies, or optimization algorithms, these methods limit the size of the tradespace so the actual set of Pareto-optimal designs may be still unknown.

The weaknesses of MATE are particularly detrimental when exploring architectures with multiple satellites because of the inherent limitations in MAUT if MATE is using MAUT as an attribute aggregation method. Once again, the static context of environmental parameters and stakeholder preferences are based on limits how the value proposition can be accurately shown in the face of contextual uncertainty. Most of the previous studies that used MATE to compare architectures were unable to account for how these emergent capabilities and lifecycle properties add to the value proposition of distributed satellites.

4.5 Phase 4: Epoch-Era Analysis

System success depends on the system meeting *expectations* within a given *context*; both the expectations and the context can change, and if either one changes, the epoch changes. The context changes with context variables, while the expectations change depending on the stakeholder needs. As discussed in Chapter 2, Epoch-Era Analysis is a method that addresses uncertainties and sensitivities in a design long before the detailed design process starts, not just in the performance model but also in the value model. EEA consists of four major processes: Single Epoch Analyses, Multi-Epoch Analysis, Single Era Analyses, and Multi-Era Analysis.

4.5.1 Contextualizing Phase 4 for Distributed Satellite Systems

In previous studies, stakeholders have reported that they definitely want multi-epoch analysis but that era analysis wasn't useful for them. This may be true for inflexible, monolithic systems, but the value proposition of distributed systems deals heavily with their operations and the options they give stakeholders over time. Without considering the timelines and decisions that could occur in future scenarios, it is difficult to show how the inherent changeability of a DSS is valuable.

While the main goal of this work is to show what can be done with DSS to achieve new scientific goals that are impossible or overlooked under a monolithic systems paradigm rather than show how these lifecycle properties add value, it is still very relevant to show how these lifecycle properties add value among different DSS designs and changeability options. Because these options are present in the case studies, it is important to model them as so.

Furthermore, the remaining weaknesses of MAUT must be addressed. New techniques have been developed that can be added to the processes of EEA to better rank designs using the insights of the value models applied to science missions that were discussed in Section 4.2.4. There are a number of special

considerations that can be made for space science missions that change the way one would apply the general case of EEA. These changes will be discussed after each process of EEA is introduced.

4.5.2 EEA Process 1: Single-Epoch Analyses

Single-Epoch Analyses are very similar to tradespace exploration, only with many (or all) of the contexts enumerated, evaluated, and explored. If there are n distinct contexts and m sets of stakeholder expectations, then Single Epoch Analyses can be conducted on $n \times m$ distinct tradespaces. This set may be reduced if a context variable affects both the performance and the set of expectations. FPN is also an important metric in this process.

4.5.3 EEA Process 2: Multi-Epoch Analysis

Multi-Epoch Analysis considers patterns across all of the (single-epoch) tradespaces. The output from this process has been impactful in previous RSC studies because the key metrics that are outputs of this process help to define value robust solutions. These metrics are generally simple for stakeholders who are unfamiliar with TSE to understand, and the analysis requires less data than era analyses.

The Pareto trace of a design is defined as “the number of Pareto sets containing that design.” Given $n \times m$ Pareto fronts, one for each epoch, the best tend to have the highest Pareto trace. Since the number of epochs can be unbounded and infinite, simply knowing how many epochs a design is Pareto-efficient is not enough. The Normalized Pareto Trace (NPT) is the Pareto trace divided by the total number of epochs. A design that is on the Pareto front of every epoch will have $NPT = 1$.

In addition to NPT, the fuzzy Normalized Pareto Trace (fNPT) is also an important metric used in Multi-Epoch Analysis. Given some fuzziness level K , the fNPT of a design measures the number of fuzzy Pareto fronts that design is within. This can be expressed by choosing a few different values of K (such 1%, 5%, 10%, 20%) or by analyzing fNPT as a function as K varies from 0 to 1.

4.5.4 EEA Process 3: Single-Era Analyses

Single-Era Analyses consider the designs and how they can evolve over the course of the mission lifecycle. In this process, individual eras are constructed by building a timeline sequence of epochs and defining how long each of those epochs last. Then, every design is run through this epoch; designs with change options may result in multiple futures depending on the strategies that can be employed by designs within the era. Eras can be constructed via narrative or computational means. Additionally, disturbances and change options can be modeled to show how utility is gained or lost as a result of different strategies that could be employed during the mission. See Figure 2-3 and Figure 2-4 for examples of how eras can be constructed.

4.5.5 EEA Process 4: Multi-Era Analysis

Multi-Era Analysis is a generalized version of Single-Era Analyses where eras are automatically constructed based on rules for how to build eras and options for changeability. Monte Carlo analysis can be used to analyze a single design over thousands or more possible lifetimes to examine how a design could possibly perform. The elements required for Multi-Era Analysis are discussed in Section 2.3.3.1.

The weakness of Multi-Era Analysis is that utility is only the perception of a design’s ability to deliver value; it is not a measure of the value delivered over the mission lifecycle. This weakness will be addressed in the remaining subsections of Section 4.5.

4.5.6 Cumulative Value Distribution and Value at Opportunity

The primary justification for being able to use cumulative value as a metric to evaluate design alternatives is to examine the cumulative value distribution to compare both the opportunity and the risk between alternatives. This is especially important when comparing the natural inflexibility of monolithic systems with the value-added changeability of distributed systems.

One metric that is often used in the financial industry to rate investments is Value at Risk (VaR). This is used to measure the risk of a loss on an investment. The simplest expression of this is the probability that an investment will *lose* value rather than gaining it. However, in the design of complex systems, stakeholders have specific performance targets in mind and expect a certain return on their investment; VaR shows what they have to *lose*, not what they have to *gain*.

In contrast to Value at Risk, Value at Opportunity (VaO) explores the opposite end of the cumulative value distribution function. For purposes of this work, VaO is defined as the probability that a design will *exceed* its expected design value over the course of its mission lifecycle as it is expressed through EEA. A comparison of VaR and VaO is shown in Figure 4-15. On the left, VaR is the probability that the return on investment will be below the original investment (less than 0% gain). On the right, VaO is the probability that a design in a tradespace will exceed the value that is expected in the primary design context with no epoch shifts (value found through EEA divided by the expected value from TSE).

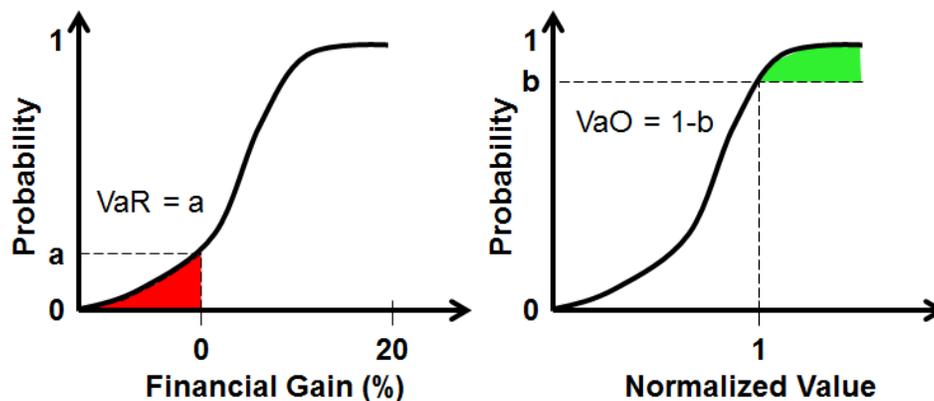


Figure 4-15: Comparison of (left) financial VaR and (right) systems VaO.

4.5.6.1 Value at Opportunity in Science Missions

Given the discrete nature of change options to capitalize on opportunity that may arise during a mission, the cumulative value may not necessarily take the form of a standard statistical distribution. Data from a case study in this dissertation comparing three different designs and their value distributions is shown in Figure 4-16. These three designs all have the same probability of discovering an opportunity that they could take advantage of to attain higher cumulative value than would be expected by that designed in the epoch of the primary value model. (In this case, the expected design cumulative value is represented by “1” on the x-axis, and the axes are not necessarily the same on the three different graphs.) The designs differ in the risk of failure in the design and the number of assets available.

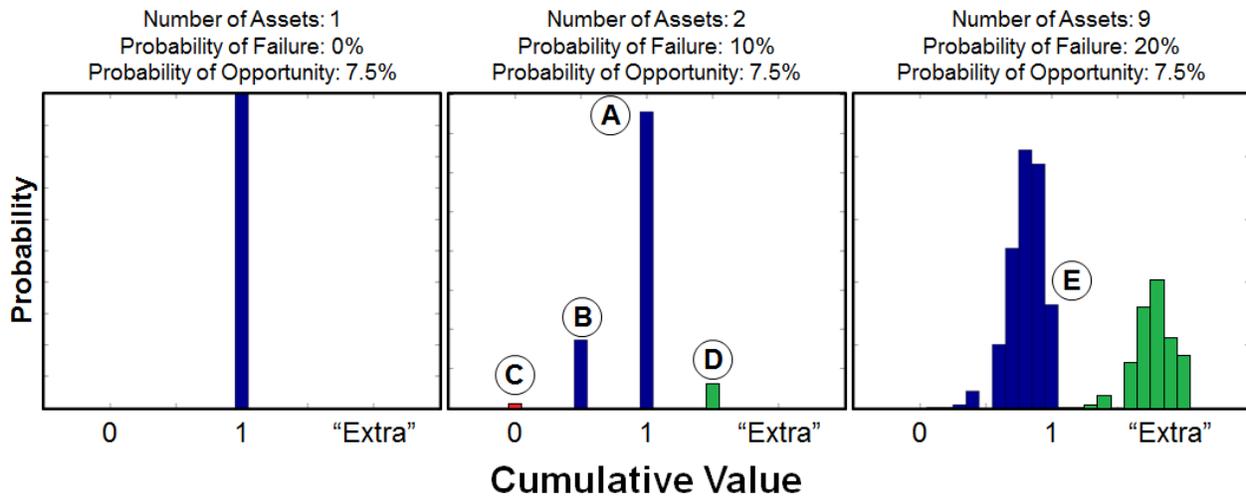


Figure 4-16: Comparison of the probability distributions of value for designs with different numbers of assets and probabilities of failure given the same probability of discovering opportunity.

The left side of Figure 4-16 shows a monolithic system with a probability of failure of 0%. This could represent a system that has been rigorously designed and tested at great expense because the cost of failure jeopardizes the entire mission, so the stakeholder spent the extra money to design out the risk. Even if this asset makes a discovery that would warrant a follow-up measurement, there is no possibility that it could capitalize on it because there are no assets remaining to do so if that payload is restricted to one area in particular, and its value probability distribution collapses to a point.

The center of Figure 4-16 shows a design with two assets, and each has a probability of failure of 10%. Now, there are several scenarios that could occur that are motivated by the possible outcomes as discussed in Table 4-2: (1) both assets survive but discover nothing noteworthy (location (A)), (2) One asset fails and the other asset discovers nothing noteworthy (location (B)), (3) both assets fail (~1%) (location (C)), (4) the first asset discovers something noteworthy and the second asset can choose to capitalize on that opportunity (location (D)), or (5) the first asset does not make a noteworthy discovery while the second asset does, but because there are no remaining assets, no additional value can be captured and the stakeholders may be left wishing they had another asset to deploy even though the mission delivered exactly what was expected from the design process (also location (A)).

The right side of Figure 4-16 shows a design with nine assets that each has a probability of failure of 20%. Because of the high failure rate, this design has the *lowest* probability of achieving exactly the value of what is expected from the design (location (E)), but there is a *much* higher probability of capitalizing on the opportunities that may arise over the course of the mission due to information gained after the deployment of each asset.

The reason why VaO is such an important metric, especially when comparing monolithic systems to distributed systems (which is not the goal of this work, but it heavily builds on those efforts), is that risk, opportunity, and cost can now be compared on equal terms. A monolithic mission would invest heavily in ensuring risk is reduced significantly, whereas a distributed system can demonstrate added value even with reduced investment in risk reduction. As shown in Figure 4-16, there is a *negligible* probability that the mission with the most assets with the highest failure rate will deliver an unacceptable cumulative value, and it is the one that allows for a much greater probability of seizing opportunity that arises.

If a performance model can incorporate design variables that show the relationship between cost and risk reduction, *distributed systems will be able to definitely show their added lifecycle value* compared to monolithic systems and also to distributed alternatives in the same tradespace with VaO. This will also help designers optimize the number of assets available versus the investment in risk reduction per asset in a design depending on how the unit cost and added cumulative value per asset change.

Value at Opportunity is especially useful for stakeholders that are risk-tolerant and can assume a mentality similar to stock brokers or gamblers. A risk-averse stakeholder would prefer a low-risk mission even if it means missing out on opportunity that may arise during the mission. Unfortunately, this is typically the mindset of such stakeholders; one reason why space technology takes so long to qualify is that few mission managers are willing to add technological risk to the mission by using untested technologies. However, if a mission manager or decision maker is willing to take risks, VaO in EEA gives them the ability to express that risk much more clearly than previous methods. Even risk-averse decision makers may be willing to compare alternatives that could deliver more opportunity.

4.5.6.2 Value at Opportunity in Other Distributed Systems

By using VaO as a lifecycle metric, researchers could revisit previous research on fractionated systems to better determine their lifecycle value. After the cancellation of F6, many researchers blamed the lack of a clear business case and prime contractors to lead the development of fractionated systems, but the answer could have simply been that fractionated satellites weren't valuable enough *at that time*. Rather than casting blame or wishing the program hadn't ended that way, researchers can use VaO and the cumulative value distribution to definitively show the cases where fractionated systems would or would not have outperformed monolithic systems, and the answer would have either helped justify continuing the program or prevented it from starting and all and saving hundreds of millions of dollars.

It is critical to specify "*at that time*" because of how rapidly the scales can tip once a certain threshold is reached in a particular market. Consider the example of the choice between solar panels and fossil fuels for supplying power to homes. The cost of electricity from fossil fuels has up to now been lower than the cost of solar panels plus power storage, but within the next three years, that paradigm could shift [265]. If this happens, it would open the floodgates to a rapid shift in power generation and storage compared to what has dominated the market since electricity first became a utility.

Likewise, in spacecraft development, there are key assumptions that may still make fractionated satellites less valuable than their monolithic alternatives, even with the added lifecycle benefits that can be shown through EEA and VaO. However, one should never assume that because one conclusion is valid today that it will still be valid tomorrow, the next year, or the next decade; new technologies change the relationships between both the cost *and* the value gained, and designers must be able to adjust for the fluidity of these factors in design processes and problem formulations and locate sensitivities that could change the conclusions of a design study entirely.

EEA is specifically formulated to analyze those sensitivities up front rather than at the end of the design process when there may be little (if any) funding remaining to conduct those analyses. Addressing whether or not fractionated systems *are* more valuable is less important from a research perspective than showing *how* and *under what circumstances* they are more valuable; by exploring the sensitivities in various parameters that affect the cumulative value distribution and VaO, researchers can show *when*

fractionated or distributed architectures are more valuable for a given task and whether the shift occurred in the past, will occur in the future, or will never occur.

4.5.7 Value Delivery over the Mission Lifecycle

Now it is time to revisit the inherent limitations of MAUT and how they are addressed in EEA. In era analysis, “value” is no longer simply “utility at cost” as it is in TSE and the first two processes of EEA, because now context variables and value models can change in time. Value is more akin to “the integration of utility over time,” since utility is “the perceived ability” to deliver value, but that definition is oversimplified when perturbations and shifts are brought into the analysis. Space science missions deliver value in one of two ways: continuous or discrete value delivery. This section proposes a new method for dealing with both of these while maintaining the value model structure of MAUT. The goal is to be able to use the cumulative value distribution to evaluate missions and let stakeholders understand the complexities associated with individual design choices.

4.5.7.1 Continuous Value Delivery

Consider the case where scientists are interested in long-term observations (~5 years) of a target. One attribute of a design would be “Coverage Time” or “Mission Lifetime”. Suppose that 2 years of observations represented the minimum satisfaction criteria (“0” utility) and that 5 years represented the maximum perceived value (“1” utility).

Now consider the MAU that several different designs may have and how it may change over the mission lifecycle. An example timeline of several designs and the utility they deliver over the course of the mission lifecycle is shown in Figure 4-17. Both the dark blue and light blue designs were launched with maximum (projected) design utility, $MAU = 1$, but in the fourth year of operation, the light blue design experiences a perturbation that causes it to completely fail. This design’s “experienced utility” or “static utility” during the fifth year drops to zero and no value is being delivered. The red design starts with $MAU = 0.75$ but due to a disturbance its experienced or static utility drops to $MAU = 0.25$ for the remaining two and a half years of the lifecycle (let’s say another attribute was “Resolution” and this perturbation decreased the resolution of the red design significantly). The green and black designs deliver constant $MAU = 0.5$ and $MAU = 0.25$ respectively for the entire mission lifecycle.

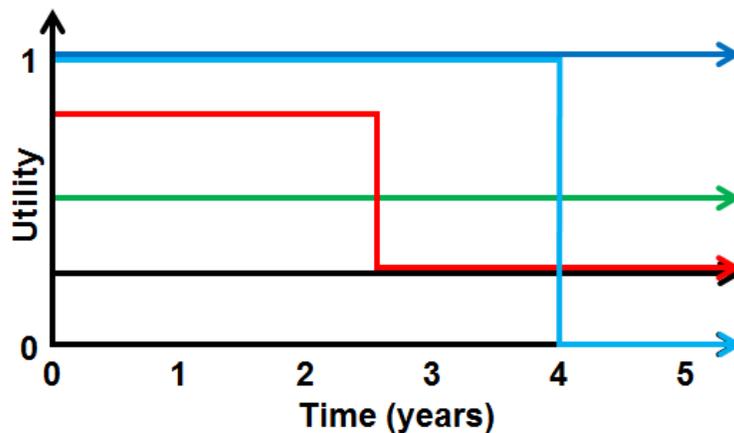


Figure 4-17: Example experienced or static utility during the mission lifecycle for five different designs.

If “value” in the EEA context were simply the integration of utility over time, the red and green designs would deliver the exact same value. However, when the attribute for coverage is taken into account and how it varies according to the stakeholder value model, a more accurate model for how value is delivered comes forward.

First, let us consider an attribute curve with constant slope with minimum acceptable and maximum perceived levels of two and five years and the same designs shown in Figure 4-17. Given that anything less than two years of observations is unacceptable, one could argue that any mission does not even begin to deliver value until its second year of observation. This utility curve and how its value delivery system would work are shown in Figure 4-18.

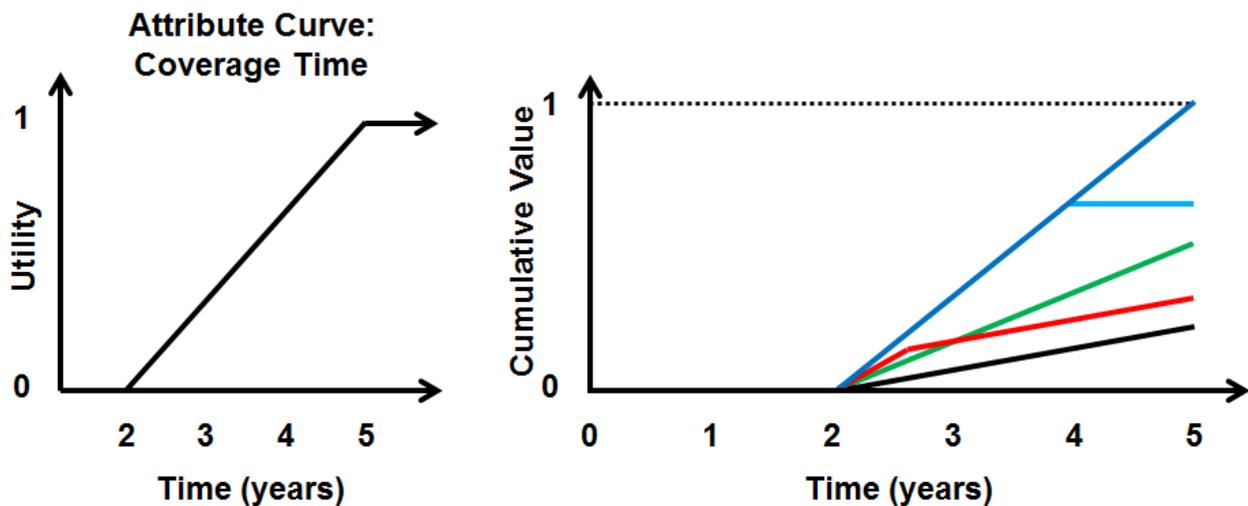


Figure 4-18: (Left) Constant attribute curve definition for coverage time. (Right) Cumulative value delivery model with a constant attribute curve definition.

The dark and light blue designs deliver maximum utility and thus have the steepest slope of all the lines. The dark blue design continues to deliver value, but the light blue design only delivers 66% of the value of the dark blue design, *as if it were designed* to operate for only four years. The green and black designs deliver constant value and thus their slope does not change. The red design begins delivering value starting in the third year just like every other design, with a slope in between the blue and green designs, but midway through the third year its utility drops to the same as the black design. For this reason, the black and red lines are parallel after the 2.5-year mark.

It is easy to see why the light blue design only managed to deliver 66% of the value of the dark blue design; all other factors being equal, the light blue design only operated for four years, and with the constant attribute curve that’s the utility that would be gained in a static context.

It is more difficult to say with certainty that the red design exhibits the value delivery behavior seen in Figure 4-18. Clearly the red design delivered more value than the black design because for 2.5 years, it delivered higher-resolution data, but for the final 2.5 years, it delivered the same resolution as the black design. We also see that the green design delivered more value to the stakeholders because more value was obtained later in the mission due to the utility curve starting at the two-year mark instead of at time $t = 0$. However, whether or not the red design truly delivered less value than the green design as Figure

4-18 shows (as opposed to the constant integration assumption based on the graph in Figure 4-17 where they deliver the same value) is dependent on how the stakeholders would perceive the data sets.

On the other hand, the underlying assumption of MAUT is that the attributes have *both* preferential and utility independence (see Section 2.3.2.2); gains from one attribute can offset losses from another. If the stakeholder is of the opinion that the green design actually delivered less value than the red design, then the problem may lay in the attribute definitions themselves rather than the value aggregation presented here. This is why iterations on the stakeholder value models and follow-up interviews are so important to this process; without feedback to compare how a stakeholder would perceive differences between designs. If the stakeholder was truly indifferent between the red and green designs, the constant utility integration would be correct, but one could also argue that the attribute definitions minimum acceptable level should have started at $t = 0$ instead of $t = 2$ years.

Now consider the same set of designs but with different attribute curves: one where utility gains arrive earlier in the lifecycle (convex curve, Figure 4-19), and one where utility gains arrive later (concave curve, Figure 4-20).

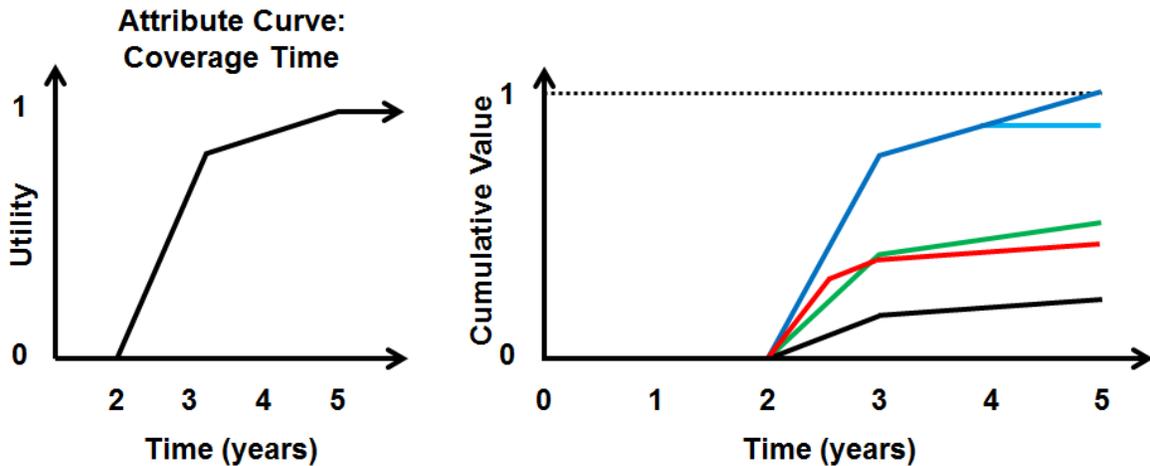


Figure 4-19: (Left) Convex attribute curve definition for coverage time. (Right) Cumulative value delivery model with a convex attribute curve definition.

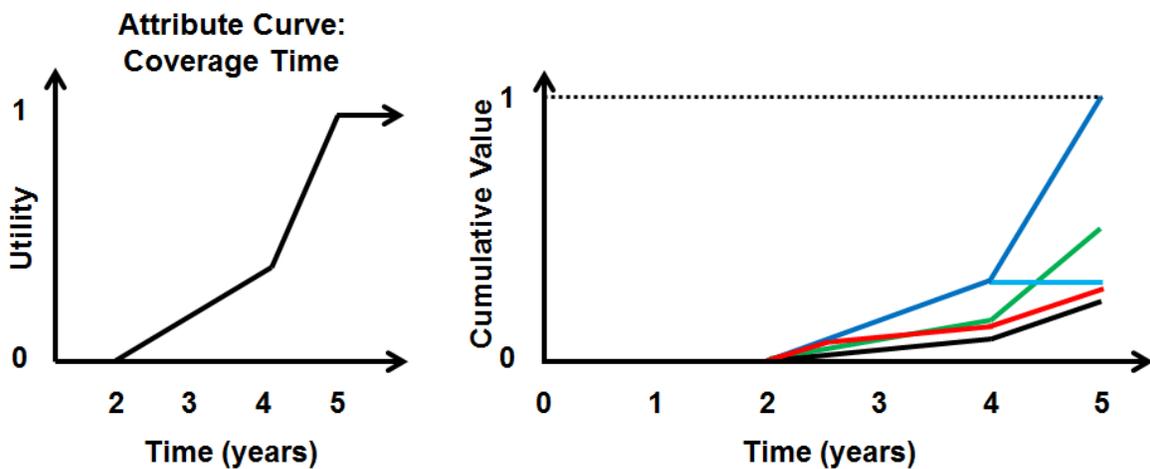


Figure 4-20: (Left) Concave attribute curve definition for coverage time. (Right) Cumulative value delivery model with a concave attribute curve definition.

In the convex case, value is delivered earlier than the later phases (specifically between years two and three), so not only are the cumulative value slopes higher before the three-year mark, but also the relative penalty for failure near the end is less. In Figure 4-19 we see that the light blue design takes much less of a penalty compared to Figure 4-18; assuming three years of operation gives $MAU = 0.75$, the light blue design delivers 87.5% of the value of the dark blue design. We also see that the red design delivers value closer to that of the green design, and still follows the same slope of the black design after it experiences a drop in utility at the 2.5 year mark.

In the concave case, value is delivered later than the earlier phases (specifically after the fourth year), so not only are the cumulative value slopes lower before the four-year mark, but also the relative penalty for failure near the end is more. In Figure 4-20 we see that the light blue design takes a much bigger penalty compared to Figure 4-18; assuming four years of operation gives $MAU = 0.33$, the light blue design only delivers 33% of the value of the dark blue design, which is even less than the 50% that the green design delivers. We also see the red design delivers even less value than the green than compared to Figure 4-18; in fact, in this model, it delivers the same value as the light blue design because of the harsh penalty the light blue design experiences. The red design once again follows the slope of the black design.

Mathematically, the value of a design *within an epoch* can be represented as

$$CV = MAU(t_{min}) + \int_{t_{min}}^T \frac{dMAU(t)}{dt} dt \quad (4-6)$$

where CV is the cumulative value, $MAU(t)$ is a function representing the multi-attribute utility of a particular design over the course of the mission lifecycle (see Figure 4-17), $U(t)$ is an attribute curve associated with mission lifetime, and T is the end of the expected mission lifecycle. As is the case in utility theory, differences in cumulative value can be used to compare and rank designs; in the event that there are no disturbances and $MAU(t)$ remains constant, then according to this definition, $V \equiv MAU$.

4.5.7.2 Discrete Value Delivery

Now consider the case where assets are deployed and perform their primary mission activities over a short period of time relative to the mission lifecycle, such as in the case of regolith sampling with penetrators or several-day visits to asteroids during a multi-year campaign. In this case, the mission's static utility at launch is relative to how much value the mission could *potentially* deliver, but true *value* is delivered in discrete chunks later throughout the mission lifecycle.

Suppose that “no assets” is unacceptable and “four assets” is the maximum perceived utility. Consider four design options, with 1 (red), 2 (orange), 3 (green), and 4 (blue) assets to deploy, respectively. If “Number of Sampling Locations” is a desired attribute (which conveniently maps directly to the design variable controlling the number of assets) and all sampling locations are perceived as equal, then utility curve would be linear and the designs would be evenly spaced in their utility over time. What this value situation would look like is shown in Figure 4-21.

Assuming there are no failures when deploying assets, all four designs deliver 25% of the maximum perceived mission value after the first deployment, three deliver another 25% of the total value after the second, two another 25% after the third, and finally only one of the designs attains the maximum value because only it carried enough assets to achieve maximum utility from the start.

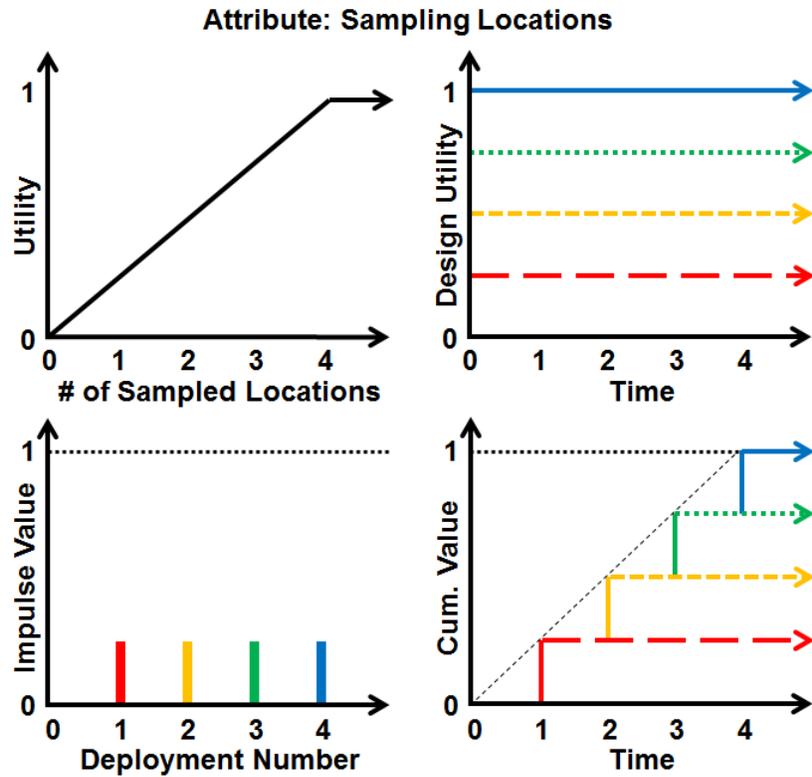


Figure 4-21: Linear returns on an impulsive value attribute. (Top left) Utility function. (Top right) Utility of four designs in time. (Bottom left) Impulsive value delivery. (Bottom right) Cumulative value delivery.

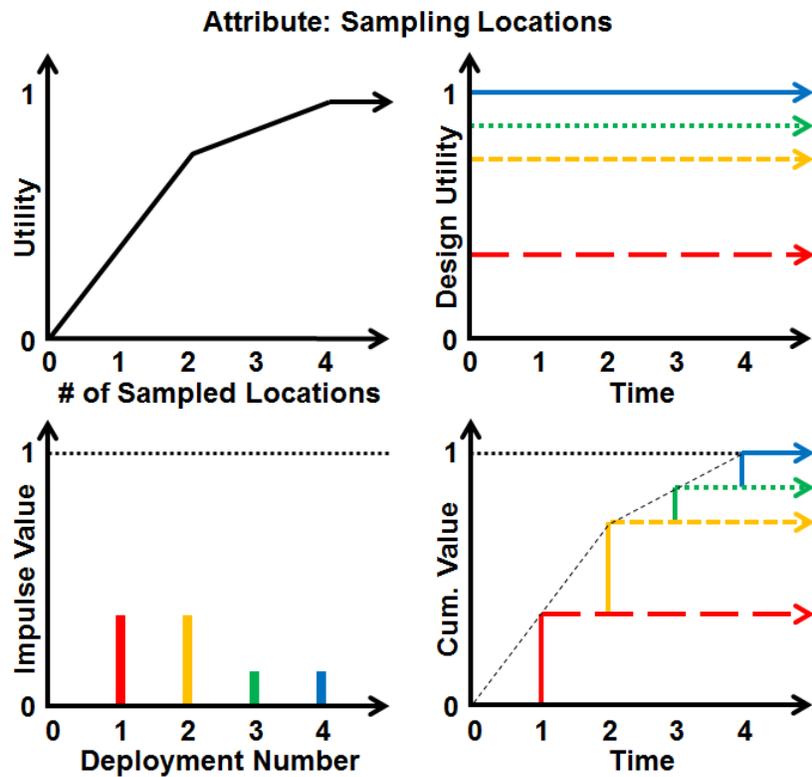


Figure 4-22: Diminishing returns on an impulsive value attribute. (Top left) Utility function. (Top right) Utility of four designs in time. (Bottom left) Impulsive value delivery. (Bottom right) Cum. value delivery.

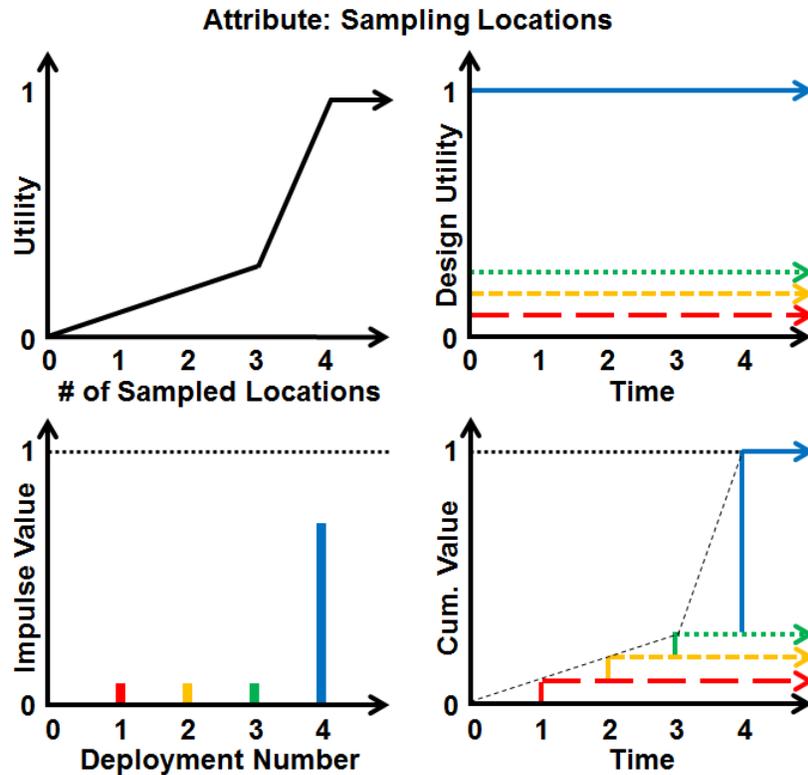


Figure 4-23: Increasing returns on an impulsive value attribute. (Top left) Utility function. (Top right) Utility of four designs in time. (Bottom left) Impulsive value delivery. (Bottom right) Cumulative value delivery.

A more likely situation is that there would be diminishing returns on the sampling locations. The logic here is that not all landing sites have the same priority, or that later ones would be slightly redundant in comparison to the first ones. The highest priority landing sites would be visited first in case there are asset failures, so a mission that was designed to sample four sites but only sampled three would still capture the most possible value. This value situation is shown in Figure 4-22.

However, in the event that there is a ground-breaking discovery at one location, there is a possibility for the type of epoch shift that was mentioned in Section 4.2.4.2. If suddenly a discovery were made, the stakeholders would have a vested interest to use the *next* asset to make a follow-up measurement. Such an epoch shift could result in a value model that looks like the situation shown in Figure 4-23. The value of this additional measurement to a stakeholder may justify sacrificing one planned landing site if there are not enough assets to visit all planned landing sites and take a second measurement in the interesting site. Furthermore, such an outcome could easily help justify to a stakeholder the added benefit at cost of adding additional assets to a mission. The increase in cost can result in far greater returns depending on the nature and probability of making such a discovery.

Epoch shifts like these are not accounted for in MATE or other static design contexts alone; this is where EEA shines as a way to show the value of the opportunity that can be captured with changeable systems including DSS. The value models for these alternative scenarios would need to be solicited in the value-modeling phase but finally become applicable here in era analysis.

4.5.7.3 Combined Value Delivery

Finally, consider the case where assets are being deployed to perform both short-term *and* long-term measurements, such as the case would be with ground-penetrators that also record weather data. What this may look like with a particular design is shown in Figure 4-24.

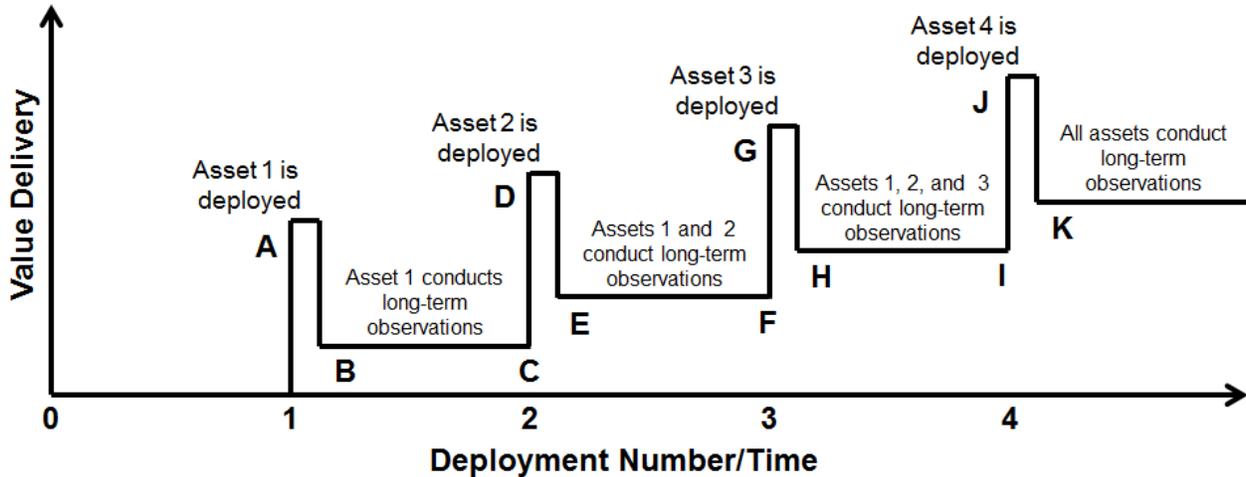


Figure 4-24: Model of how value is delivered with both short-term and long-term measurements.

However, this illustrative sketch is incredibly naïve because it makes many assumptions. First, it assumes that all short-term data is weighted the same, meaning that even one asset can deliver acceptable value such as in the case shown in Figure 4-21. In reality, the y-axis value at point A may be different than the difference in values between points C and D, F and G, and I and J. Depending on the nature of the value models, these impulsive value delivery modes may look more like Figure 4-22 or Figure 4-23.

This illustrative sketch also makes three assumptions about the long-term measurements: (1) that value starts being delivered from $t = 0$ (unlike in the cases described in Section 4.5.7.1), (2) that the attribute curve is constant (such as in Figure 4-18 but starting at zero instead of two years), and (3) that all assets deliver value precisely additively (such as shown in Figure 4-21, only for a time-dependent case).

The first assumption lets the design begin creating value as soon as it is deployed, rather than after some point later once the minimum acceptable value has been delivered, as is the case in Figure 4-18 through Figure 4-20. The first and second assumptions combined result in the perfect flatness of the line segments BC, EF, and HI and beyond point K.

The third assumption is especially bad in the case where multiple assets work together to perform *simultaneous sampling*. For instance, if the long-term measurements are seismometer data, one or two assets would deliver very little value before the other assets were deployed; in this case, the y-axis values of line segments BC and EF may still be zero or close to zero, even if the first assumption were true, because the fidelity of the seismometer network may not be acceptable before more than two assets are deployed.

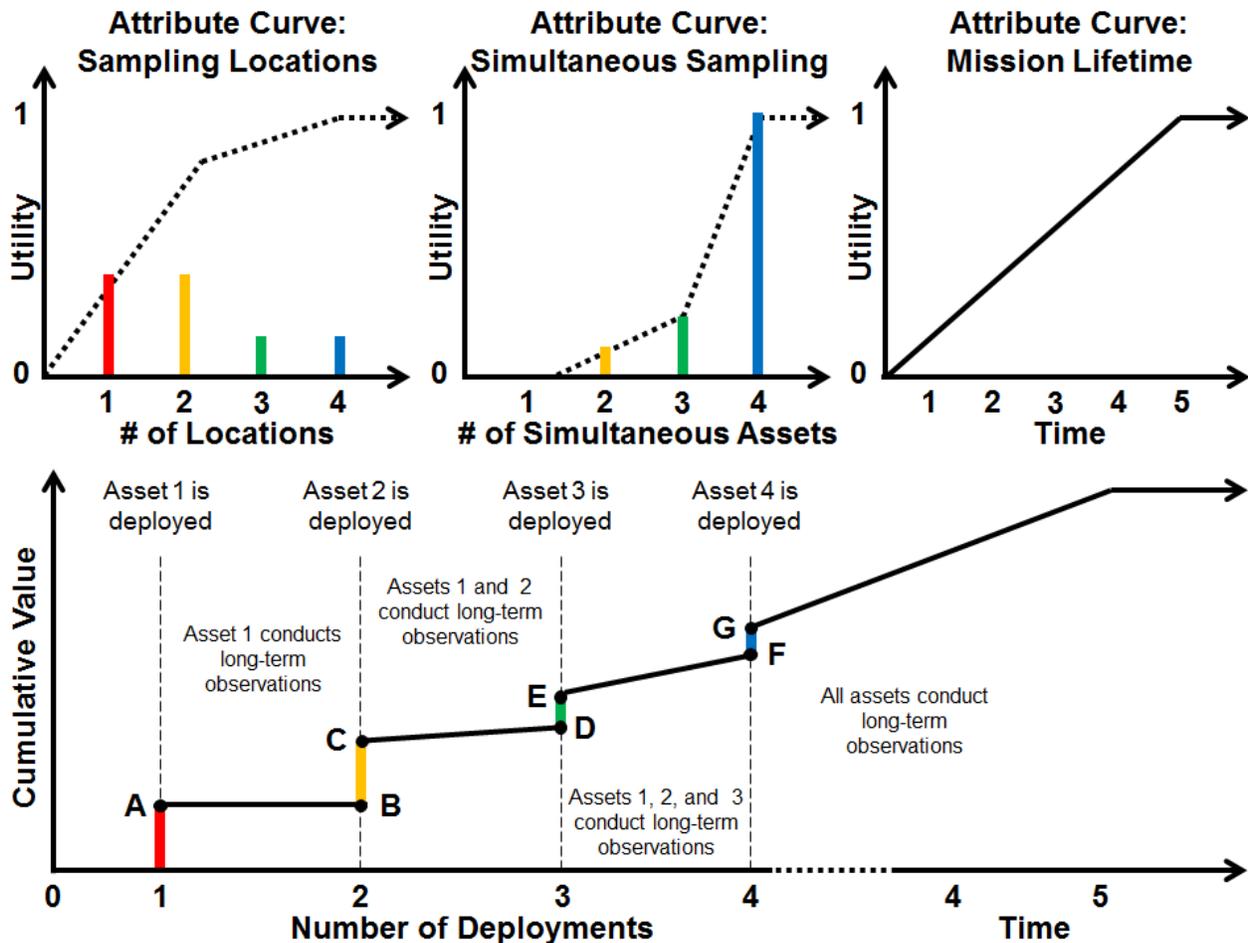


Figure 4-25: Less naive time value model that accounts for both discrete and long-term mission value and shows how value builds as more assets are deployed.

Figure 4-25 shows a combined value model that is less naïve than the one shown in Figure 4-24. Given the attribute curves shown at the top for number of sampling locations (discrete value), number of simultaneous samples (time-dependent value), and a simple linear mission lifetime attribute curve, the cumulative value function is shown on the bottom of the figure.

Even this value model is slightly simplified for the reader to see the relationship. The mixing of axes of the bottom graph does not take into account the amount of time between distributed deployments. If all four assets are deployed simultaneously, the end result at time equal to 5 would look the same as it does in the figure. However, if the deployments of the assets were spaced out in time significantly relative to the mission length, the ending would look different.

Now consider the case where the assets are deployed over a period of time rather than all at once, and the value delivery due to the time-dependent measurements is delayed because the minimum acceptable mission lifetime is greater than zero. The complete value delivery over time for a single design with four assets when both discrete and time-delayed continuous sampling will deliver value when assets are deployed at different times is shown in Figure 4-26.

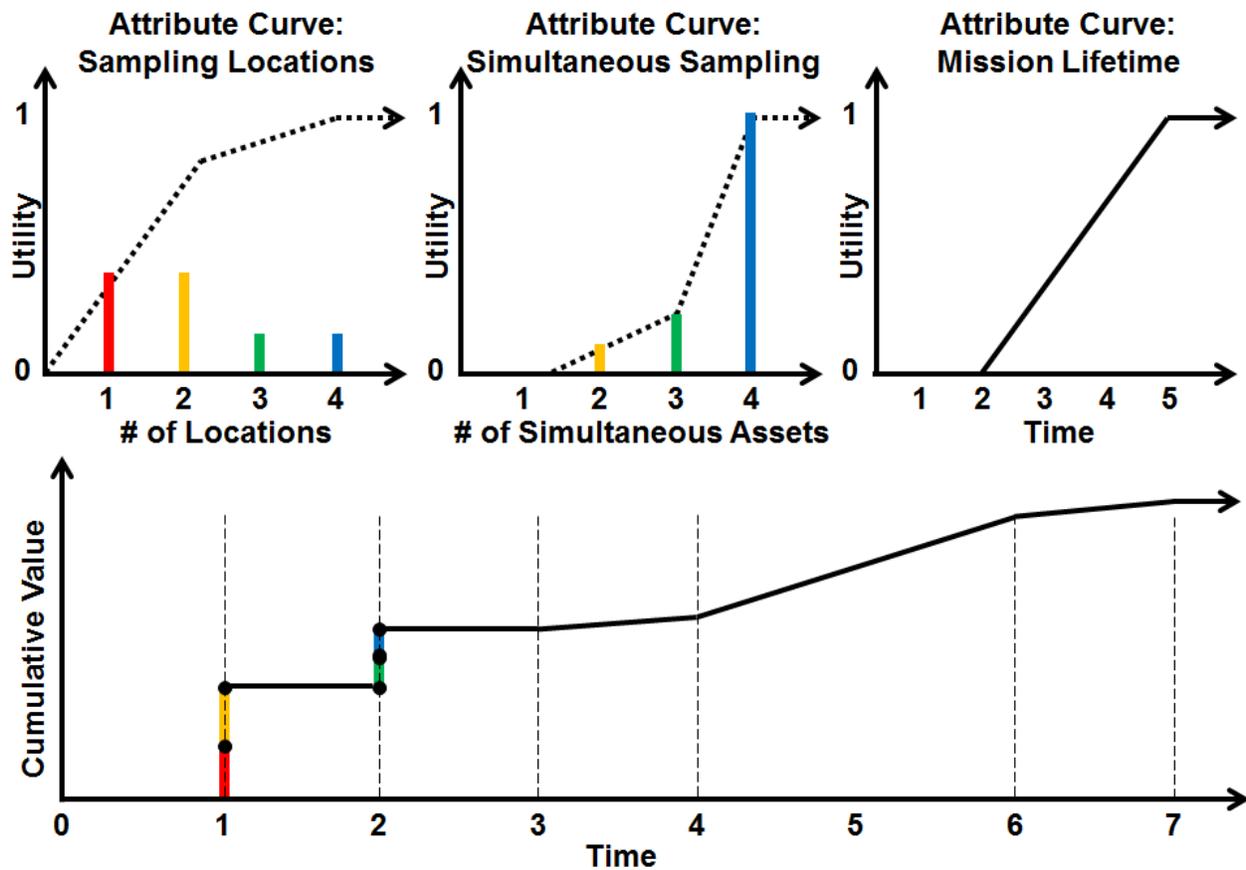


Figure 4-26: Complete model of time value delivery with delays in time-dependent value returns.

In this scenario, the first two assets are deployed a year after the mission launches and deliver discrete value upon their deployment. The cumulative value delivered stays constant, because value has not been delivered by these assets due to the delay caused by the minimum lifetime attribute. After two years, the next two assets are deployed, and they deliver their discrete value. Between 2 and 3 years, value also stays constant. After 3 years, value starts to be delivered, but the slope of the cumulative value at this time is low because only two assets have been contributing to the simultaneous sampling attribute for at least two years at this point. It is not until 4 years that the higher utility of all four assets sampling for a minimum of two years kicks in and drives the slope higher.

Between years 4 and 6, the cumulative value slope is increasing at its highest rate because the simultaneous sampling utility is maximized while the mission lifetime grows. If this curve were not strictly linear as shown in the upper right of Figure 4-26, there would be more intricacies in the slope changes. After 6 years, the slope decreases because the first two assets will have observed for the maximum lifetime and thus value is no longer added. The slope of the line between 6 and 7 years is related to the difference between the utilities of simultaneous sampling with four assets versus two. The total value would be the same after 7 years with this delayed deployment as it would have been had they been deployed all at once after a total mission lifetime of 5 years.

If, suddenly, the first two assets failed after 6 years, the slope of the line between 6 and 7 years would be zero. This is because under this value model, the maximum value of having only two assets sampling simultaneously would have been achieved, whereas sampling with four assets would only have been

achieved up to 4 years and not 5. However, a scientist may still find value in continuing to operate the remaining assets for the final year of the timeline. If this were the case, it is not reflected in this value model, so a new value model would need to be applied. This would result in an epoch shift, which does not apply to what has been discussed thus far.

4.5.7.4 Value Delivery Summary

Understanding how the value is delivered in time is especially important in EEA because part of Era Analysis is understanding what happens to the value delivery as assets start to fail over periods of time. While monolithic systems generally treat risk and value delivery separately because value delivery typically falls to zero after the spacecraft fails, distributed systems need to incorporate value delivery and risk together to predict what the mission will deliver, especially when mission designers choose to cut costs by increasing the acceptable risk because the failure compensation costs have also dropped as a result of having multiple satellites.

These examples show how value is delivered *within a single epoch* over time in the form of an era with multiple value delivery sources. How value model shifts are compared across multiple epochs in an era will be discussed in the next section.

4.5.8 Value Model Epoch Shifts in Era Analysis

The reason why value cannot be aggregated simply as the utility integrated over time (in addition to the reasons discussed in the previous section) is that the value model and stakeholder expectations may change over the system lifecycle. With the MAU scaling that has been chosen, $MAU \geq 1$, but an epoch shift could result in a stakeholder suddenly seeing a number of different changes in value.

A trivial example of why value cannot be defined simply as the integration of utility over time is shown in Figure 4-27. As the expectations change due to a shift from Epoch #1 to Epoch #2, the red design becomes infeasible or unacceptable to the stakeholder. The yellow design moves from $SAU = 1$ to $SAU = 0$, the green design drop to $SAU = 0.25$, and the blue design remains at the highest utility.

In this epoch, the red design is temporarily unacceptable while the yellow design is acceptable, but because the utility is zero, no value is being added over that period of time. Technically, it is more valuable to a stakeholder to have an acceptable design than an unacceptable one in that epoch, but there is no mathematical difference in that epoch. It may be acceptable to compare the yellow, green, and blue designs this way, but unless the red design is removed from consideration entirely due to the lapse in acceptability, this representation would skew the results improperly. An easy way to avoid this is to assume that if a design is unacceptable at any point in the era, it is unacceptable in the entire era.

Calculating value across epochs is difficult when a system may be switching epochs due to rapidly changing stakeholder expectations or mission needs, but this is rarely the case in space science missions. Generally, in space science missions, the number of sets of stakeholder expectations is small, and expectations do not return to the same set as they were in the first epoch if they have already changed once during the mission lifecycle.

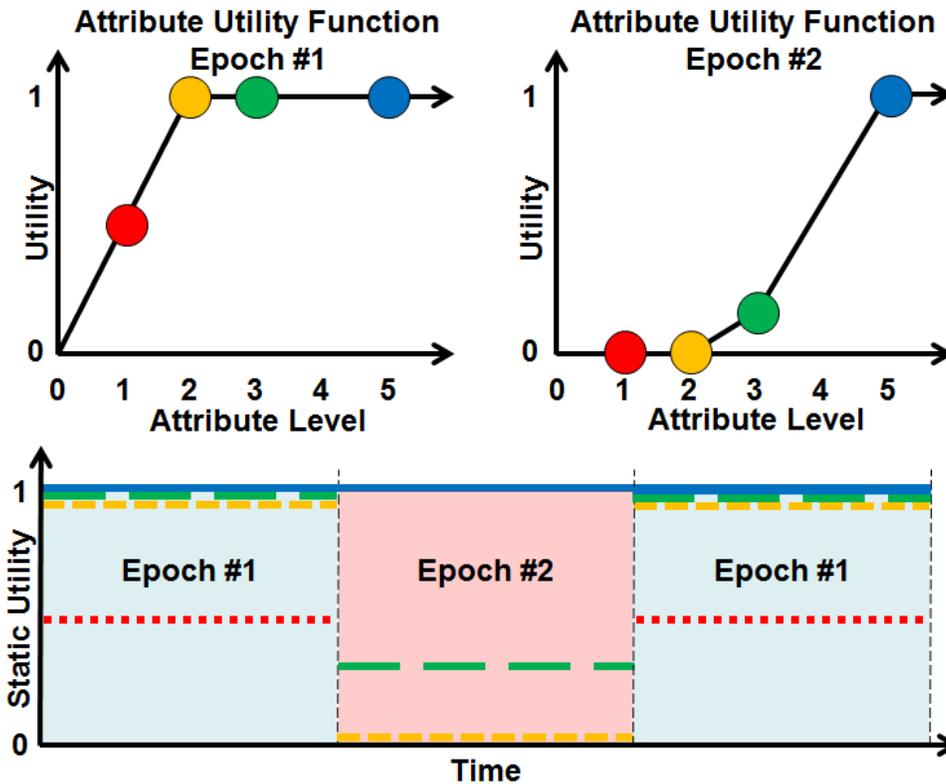


Figure 4-27: Example era with epoch shifts resulting in value model changes on an attribute.

Under these assumptions, which do not apply to the general case of EEA such as the US Navy case in Schaffner’s thesis [58], it is still appropriate to use the time integration method, but only when viewing the entire era *retroactively* [46]. Retroactive era analysis is useful in examining how the value of space science missions can drastically change due to discoveries made over the course of the mission.

4.5.8.1 Retroactive Era Analysis

An example was brought to the author about applying a value model retroactively to a person taking vacations. Suppose a person takes the best vacation of his or her life. This vacation is so good that they cannot possibly imagine a better one. If MAUT had been applied to model the value of the vacation, this particular vacation’s design would have been assigned $MAU = 1$. Now suppose the person took another vacation a year later that was *even better* than the previous vacation, but one of the major fun activities was canceled due to weather and a tour guide was mean. The person may assign this vacation a design utility $MAU = 0.93$ due to these design deficiencies because their expectations were not fully met.

What does that mean for the utility assigned to the first vacation? Obviously, the value model must have changed, and therefore the utility assigned to the first vacation would no longer be at the maximum level after the second vacation. This *does not* mean that the value of the first vacation was diminished by the second vacation; it means that the value of the same vacation after the second vacation has changed as a result of a shift in expectations, so the utility of the first vacation can no longer be $MAU = 1$.

Now reconsider the case of a space science mission with multiple deployable assets attempting to discover something unknown, and take into consideration the possible outcomes listed in Table 4-2. If one asset makes an interesting find, the value model may change significantly. Suppose a PSS, PI,

mission manager, or group of stakeholders initially wants to explore three locations or deploy three assets. An attribute utility curve may look something like what is shown in the top left of Figure 4-28. Designs that have more than three assets would be penalized in cost but not rewarded with extra utility or possible added value but would still exist in a well-populated tradespace.

Now suppose one of the assets makes a major discovery after deployment. There is no doubt that the expectations of the mission would change after this point; although the scientists and stakeholders would be pleased with the mission, there would be great interest in exploring that particular location even more. With new attention brought to the mission due to the announcement of the discovery, both endogenous and exogenous pressures can change the strategy of a mission after this epoch shift. A follow-up measurement, if possible very soon after the initial discovery rather than after several more years with a new mission, possibly under a new PI who could take scientific credit for a better follow-up confirmation, would be the new top priority for the PSS.

As a result of this new information, an epoch shift would occur and the attribute utility curve would look like something shown in the top right of Figure 4-28. The center and bottom of the figure show the static utility and the cumulative value delivered over the era by the various designs. The switch in experienced or static utility as shown makes sense and follows the trends of previous MATE and RSC studies. However, the value delivered by the first three assets is not suddenly diminished by this discovery; if anything, the scientists would be even happier because of the results.

Pre-Phase A design studies are *prospective* looks at what value will be delivered *retroactively* at the end of the mission. What matters in space science missions is how well the product performed at the end of its lifecycle. A design that pleases a stakeholder immediately before launch but fails to meet changed expectations later has failed its mission. The correct value model to apply in evaluating the success of a mission is the value model of the epoch that is active *at the end of the mission* when the assumption that the value model cannot change back to the value model used at the start of the mission holds true.

Retroactive value modeling does not apply to designs in rapidly shifting contexts such as the previously cited Navy missions, because value is being delivered *presently* in each epoch. In science missions, what matters are the results that have been delivered up to that point, whether the expectations change or not.

4.5.8.2 Added Value Era Analysis

In the case of mission extensions, value is clearly added beyond the initial expectations, though exactly how that value is added is dependent on what else changes in the set of expectations. MAU is an affine function that can be scaled linearly if the relative MAU weights' and attributes' minimum and maximum levels stay constant in an epoch shift that would occur in a mission extension. In the event that these expectations do not change, the method discussed here and retroactive value measurements are the same.

In the event of a mission extension, the expectations of the stakeholders have changed, but usually not significantly. The total value delivered over the lifecycle must be scaled down so that the new maximum value is scaled to meet the expectations at the end of the mission. This scaling is shown in Figure 4-29. The solid line represents the cumulative value as defined in that epoch, while the dotted line represents what the value should have been in the previous epoch. This linear scaling is possible because expectations on what has already been delivered have not changed.

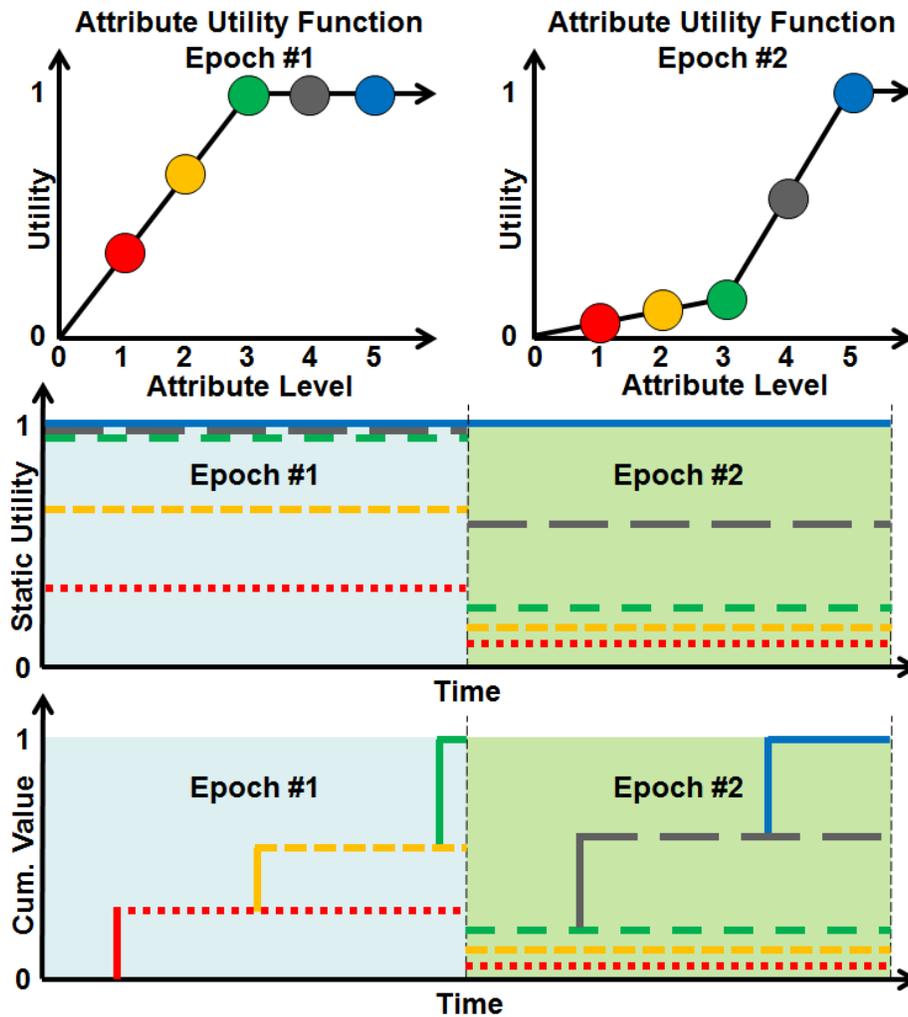


Figure 4-28: Epoch shift showing how expectations can change. Although the value delivered in the first epoch is not diminished by the epoch shift, the expectations and highest levels of satisfaction have been raised significantly as a result of the epoch shift.

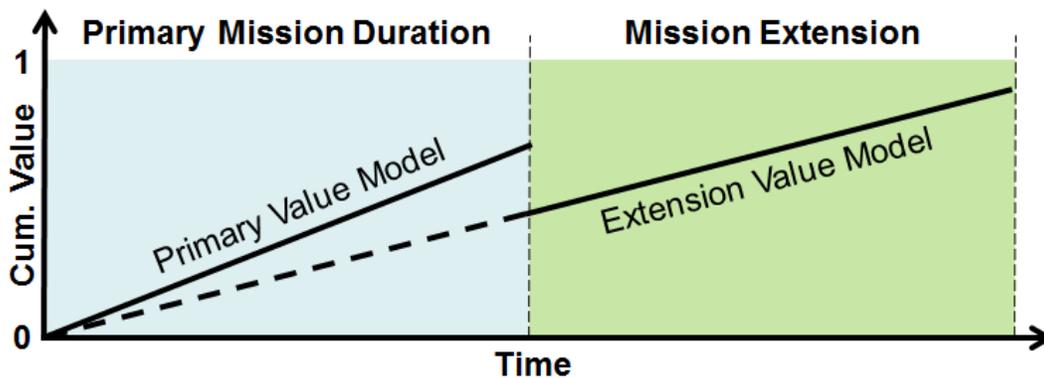


Figure 4-29: Affine scaling of the value model in the event of a mission extension with no changes in the components of the value model. The slope of the line decreases under the extension as the value is scaled down but does not change otherwise.

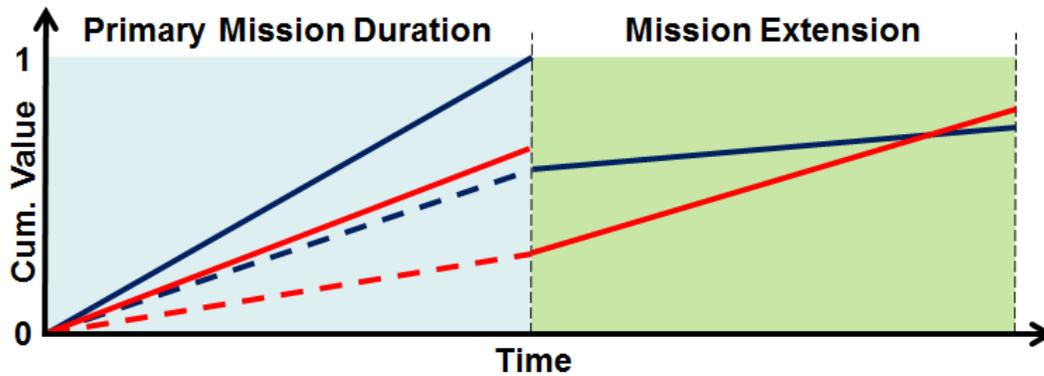


Figure 4-30: Affine scaling of the value model in a mission extension with changes in the components of the value model. The slope of the function can be different on one side of the epoch shift compared to the other if the components change.

However, if the mission extension does more than simply extend the lifetime and the mission goals change slightly, two different designs may deliver value differently than in the previous epoch. Consider the two designs shown in Figure 4-30. The blue design outperforms the red design under the expectations of that epoch. If the mission is extended, the cumulative value is scaled down at the end of the epoch for both designs, but the blue design does not perform as well in the epoch of the mission extension as the red design, so more value is captured by the design that was considered inferior before.

4.5.8.3 Comparison to Previous Work

The expression of cumulative value as an integration of utility over time is not dissimilar from previous work by Richards et al. [45] that examined time-weighted utility over system lifecycles as a value-centric metric for survivability. In fact, in cases where time is not related to any attributes and satisfaction is perceived as how the system operates at any given point in time, the time-weighted average utility and the cumulative value metric as shown in Equation 4-6 produce identical results.

Richards clearly distinguished between “design utility,” which is the calculated utility of the system from static MATE and what would be expected from the design in the absence of disturbances, and “experienced utility,” which is the perceived value of the system at a given point in its lifecycle. Richards then defined the time-weighted average utility \bar{U}_t as the integration of the experienced utility $U(t)$ over the design lifetime T_{dl} as

$$\bar{U}_t = \frac{1}{T_{dl}} \cdot \int U(t) dt \quad (4-7)$$

The statistical distribution of \bar{U}_t can be found using Monte Carlo analysis with different probabilities and amounts of disturbances a design can experience (and react to with changeability). Richards’s work was oriented around *survivability*, or maintaining value before, during, and after disturbances. Specifically, Richards explored how to maintain the utility above some set minimum before a change option is enabled and how different change options or acceptable minima affect the time-weighted average utility and total cost.

However, the time length of observations (especially when those observations cannot be interrupted) is sometimes critical to mission success or the perception of value that the system provides; in this case, Richards's metric is compatible with attributes like "operational response time" or "time to first observation," but it cannot integrate the possible relationship between the perception of value in time observations and other attributes; this relationship is more accurately captured by what has been presented in this chapter with the integration of the derivatives of the MAU function shown in Equation 4-6.

When observation lifetime is an attribute, it must be examined continuously within the utility calculation over the course of the era so that the effects on value caused by disturbances can be objectively compared even in changing stakeholder expectations. For example, assume the attribute for "observation lifetime" of a value model had minimum acceptable and maximum perceivable levels of three and five years, respectively. If a system was designed to last for 5 years, the system's "design utility" would be 1 for that attribute, and if all other attributes were 1, the MAU would be 1. However, if that system failed after 2.99 years due to a disturbance, it would not have met the minimum acceptable criteria despite operating at the highest desired capacity for that long. Such a situation would correctly show an unacceptable design when using the metric shown in Equation 4-6 but not with the one in Equation 4-7.

One consequence of Richards's formulation is that the MAU function must be universal or constant over the period of integration, meaning stakeholder expectations cannot change. The formulation of cumulative value derived in this chapter imposes a similar restriction that the value model must apply to the entire lifecycle. The value model can change, but the comparison is only objective when the system is evaluated using the value model at the *end* of the system lifecycle, because the stakeholders would know after the fact what their value model was after the mission. This restriction allows cumulative value to be incorporated into an EEA framework even with changing stakeholder expectations because of the retrospective point of view of the calculation.

4.5.9 Key Assumptions and Limitations

While the proposed modifications to the formulation of EEA as shown in the literature addresses some of its major weaknesses, there are still lingering limitations and new ones that are a result of the changes.

4.5.9.1 General Assumptions and Limitations of EEA

While TSE may seem like it requires a lot of computing power, in-depth EEA requires far more for each individual design. Displaying all the data that can be gleaned from EEA in an effective manner is an active area of research [266].

4.5.9.2 Assumptions and Limitations of Proposed Modifications to EEA

In the event of a mission failure, the retroactive value model also changes. For instance, if a mission lifetime attribute had a minimum acceptable level of 2 years, but the mission failed after 1.99 years, scientists would still see tangible value in that data. The data products or quality would be below their expectations, but it would not be "worthless" to them. Going back to the discussion in Section 4.2.1, scientists are unusual "customers" in the traditional sense in that they tend to have grand desires, but after a premature mission failure they have no negotiating power and would settle for what they could get.

Examining how a scientist "settles" for data at the end of a disappointing mission is difficult for many reasons. First, eliciting such a settlement model may place distrust in the designers for even considering

failure to be an option. Second, the end of the mission is far enough away into the future that the scientists may be unable or unwilling to imagine that reality. Finally, with a mission failure, the scientists really has no choice but to settle with what was gained, so they have no leverage to demand anything or work to find a “Best Alternative To a Negotiated Agreement” (BATNA).

Suppose that during a follow-up interview to validate the stakeholder’s value model against a set of calculated designs, a stakeholder agreed that a design with a projected utility $MAU = 1$ that operates for 1.99 years before prematurely failing *is more desirable and is still acceptable* compared to a design with a projected utility $MAU = 0.01$ that is designed to operate for exactly 2.01 years (which is, by the definition of the original value model, acceptable). This could mean that the lifetime attribute’s minimum acceptable level is simply not correct and needs to be readjusted, *or* it could mean that the attributes that govern lifetime are not *entirely* independent from other attributes, which breaks one of the fundamental assumptions of MAUT. Such changes would affect *all* possible designs and require the value model to be retooled. For these reasons and to control the ballooning scope of this work, these dependent value models will not be explored in this work, nor will cases of premature mission failure beyond the lower bound of the lifetime attribute definition.

One of the reasons why NPT, fNPT, eNPT, and efNPT are useful metrics to compare design alternatives in multi-epoch analysis is that they characterize a design across all epochs using a *single number*. Likewise, VaO represents the opportunity a design can capture with a single number. However, the added value may not behave according to a Gaussian distribution, so although the probability of achieving higher value is captured by VaO, the magnitude of that value is not. Probability distribution curves can compare two designs directly, even if VaO can be useful for comparing many designs simultaneously.

4.6 Method Synthesis and Summary

This section revisits the known weaknesses that were previously identified in the chosen VCDM and addresses how the proposed modifications strengthen the method, then summarizes the major processes of the proposed implementation of RSC and describes their inputs and outputs.

4.6.1 Revisiting Known Weaknesses

The strengths and weaknesses of MATE as they apply to the design and evaluation of space science missions are reexamined in Table 4-3 (second column). Adding EEA on top of MATE allows MAUT to overcome the weakness of not being applicable in a context change. Additionally, new research using TSE in multi-stakeholder negotiations is showing promise of how these otherwise irreconcilable mental value models can be normalized relative to the stakeholders. While this is not the same as a social welfare function, it can serve the same purpose by transforming the value model scaling so they are able to be mixed. However, the traditional formulation of EEA in RSC still does not address the weakness that utility is not necessarily the same as value delivered over time, nor does it explicitly address how synergy affects the marriage of science goals in a complex, multi-goal mission in the context of a greater scientific program or agency.

The formulation of the time-dependent value model described in Section 4.5 helps eliminate this weakness when there are no major changes among the relationships of the attributes in different epochs, which is more likely the case in science missions compared to previous MATE studies involving military systems that must operate in highly complex and dynamic contexts.

Table 4-3: Summary of the properties of typical space science missions and how well variations of MAUT can be used to assess their value

Properties of Typical Space Science Missions	MAUT	MAUT+EEA	Modified RSC
Value can be derived from multiple sources	✓	✓	✓
Value is NOT discounted cash flow	✓	✓	✓
Mutual Utility Independence	✓	✓	✓
Stakeholder decision is under uncertainty	✓	✓	✓
Value is an ordered, non-ratio comparison	✓	✓	✓
Context may change over time	X	✓	✓
Applicable with multiple stakeholders	X	O	O
Value is delivered over time	X	X	✓
Value can be delivered synergistically	O	O	✓

4.6.2 RSC Summary and Flowcharts

The major steps and proposed modifications to the RSC method that will be followed in the case studies are summarized in the following figures. A chart specifically for TSE has been omitted since it is a degenerate case of Single-Epoch Analyses, though it should not be considered nontrivial because of this.

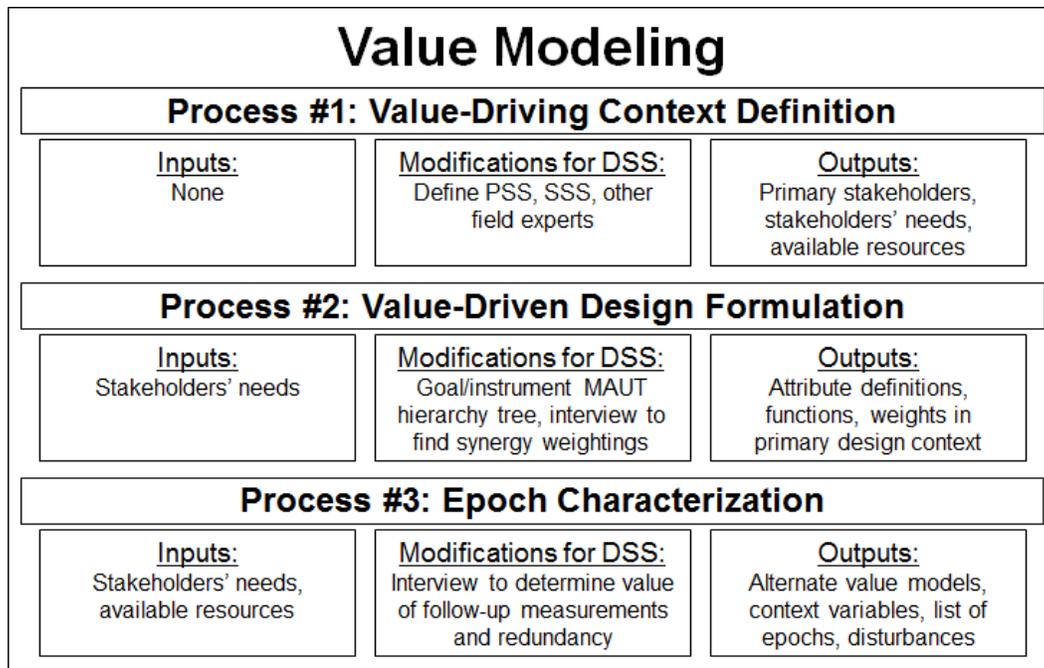


Figure 4-31: Summary of RSC Phase 1: Value Modeling

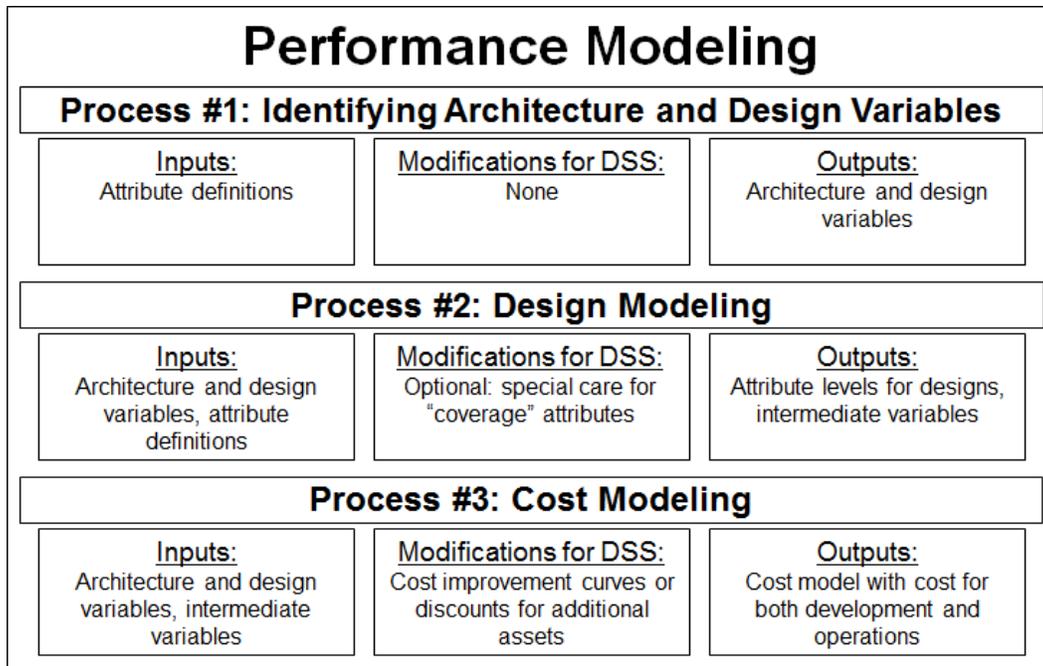


Figure 4-32: Summary of RSC Phase 2: Performance Modeling

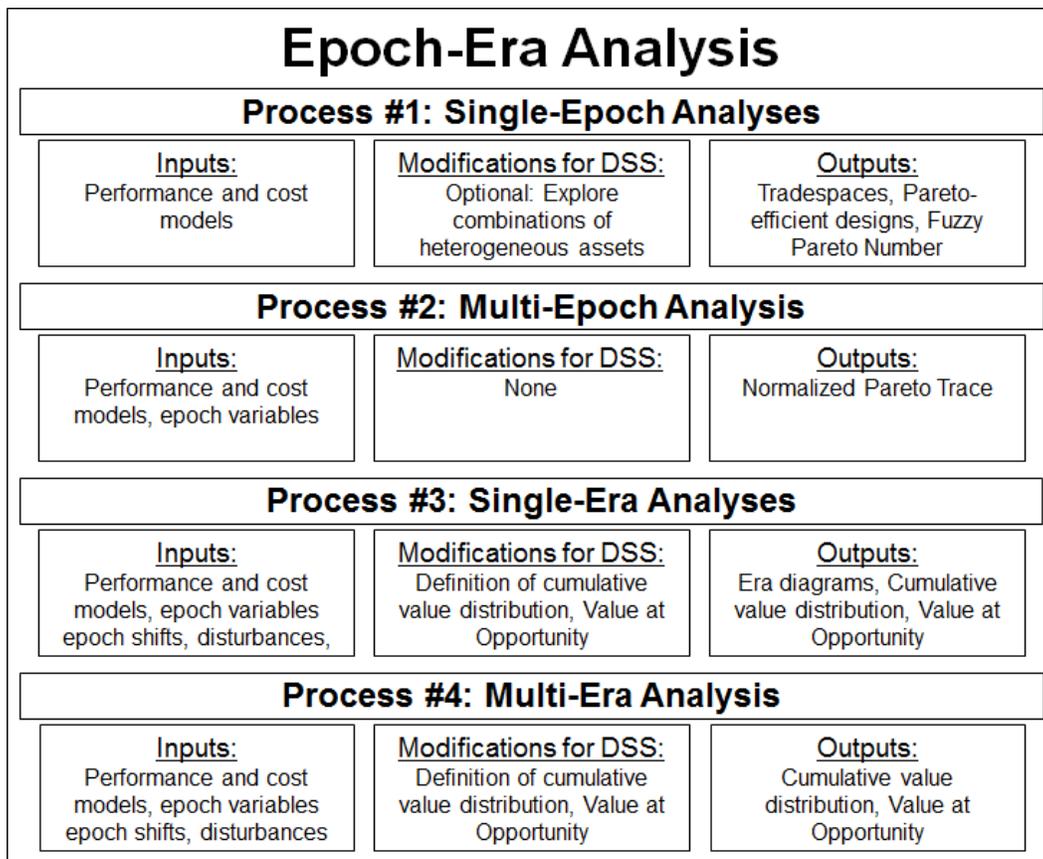


Figure 4-33: Summary of RSC Phase 4: Epoch-Era Analysis

4.7 Introduction to Case Studies

The primary thesis objectives will be met through the exploration and analysis of case studies. Empirical evidence from case studies can be used to argue that the value proposition for distributed satellite systems in space science missions is communicated more accurately compared to previous research with the methods described in this chapter, and that scientific opportunity that has been previously overlooked can be seized under the DSS paradigm.

The four case studies that will be examined here leverage one or more of the emergent capabilities of DSS described in Chapter 3. Additionally, these case studies have been intentionally chosen to cover a broad range of target types, target numbers, destinations, technologies, instrument complexities, and emergent capabilities, summarized in Table 4-4.

Most importantly, these case studies all address open areas of research in space science and scientific questions that have not yet been answered partially because they *cannot* be answered in a monolithic systems paradigm. Each case study on its own could form the basis of a genuine mission proposal that *would contribute significant value to the fields they address at a reasonable price*.

For each mission, a respected expert in the field relevant to the science mission of the case study has been chosen to serve as the PSS and is either (a) the PI for the existing mission, and is therefore the person whose needs must actually be satisfied with the mission, (b) is already studying the exact phenomenon being explored, or (c) works for the group that created the science goals the mission was chosen to satisfy.

Table 4-4: Summary of the important differences among the four case studies.

Mission/Case Study Name	ExoplanetSat	HOBOCOP	ÆGIR	GANGMIR
Operational Destination	Low Earth Orbit	Heliocentric Orbit	Near-Earth Orbit/ Asteroid Belt	Mars Orbit/ Mars Surface
Primary Science Goal	Find Earth-sized Exoplanets	Understand Solar Dynamics	Catalog Space Resources	Search for Biohazards
Target Type	Nearby, Sun-like Stars	Solar Current Sheet	Asteroids (All Classes)	Mars Regolith/ Mars Atmosphere
Target Quantity (10^x)	3-4	0	4-6	1-2
Instrument Complexity	Medium	Low/High	Medium/High	Very High
Leveraged Capability	Shared Sampling	Simultaneous Sampling	Census Sampling, Self-Sampling	Staged Sampling, Stacked Sampling
Value Delivery Over Lifetime	Continuous Value Only	Pseudo-Continuous Value	Mostly Discrete Value	Both Discrete and Continuous Value
Propulsion System Type	None	Electric	---	Chemical
Architectural Elements	Small Satellites	Small Satellites, Relay Satellites	Fractionated and Medium Satellites	Penetrators, Small Satellites
Maximum Number of Assets Considered	120	160	30	33

CHAPTER 5

ExoplanetSat



Original artwork © Judas Priest and Columbia Records, 1982. Parody modified by the author.

5. Case Study: ExoplanetSat

*“Always in focus, you can't feel my stare.
I zoom into you, but you don't know I'm there.
I take a pride in probing all your secret moves.
My tearless retina takes pictures that can prove...
Electric eye, in the sky, feel my stare, always there
There's nothing you can do about it. Develop and expose.
I feed upon your every thought, and so my power grows!”*

JUDAS PRIEST, “ELECTRIC EYE”

The first case study is the mission that operates closest to Earth. ExoplanetSat is a constellation of CubeSats that will search for previously unknown, Earth-sized exoplanets around Sun-sized stars. The hardware development for this mission is already under development at MIT under the management of Dr. Sara Seager with help from several graduate students, JPL, and Draper Laboratories; it is the only case study in this dissertation that has actual hardware tested and is not strictly a concept design study.

Plans for a technology demonstration mission using one CubeSat to observe known transiting exoplanets are underway; however, while the individual units of the fleet were being designed, there was no design plan for the full constellation, nor was there a plan to go forward to maximize the value and scientific returns from the constellation. The author was asked to explore the options for deploying and maintaining the full constellation; the end goal of this case study is to identify valuable architectural and design options and recommend an exact initial design for the final phase of this mission, as well as to identify strategies to preserve or enhance mission value over the course of the mission lifecycle. This mission heavily leverages **shared sampling** to achieve its primary scientific objective.

5.1 Case Study Introduction and Motivation

This case study builds off several years of work on the development of the technology of a single ExoplanetSat CubeSat unit. This section will describe the primary mission goals and stakeholders, summarize some of the design and development that has taken place already, and discuss the results of a NASA JPL Team X study to conduct mission operations.

5.1.1 Primary Mission Goals and Stakeholders

The primary science goal of ExoplanetSat is to find **Earth-sized** planets around **nearby, Sun-sized** stars using the transit technique (described in Section 5.2.1). This is challenging because there are three requirements to using the transit method: photometric precision, large coverage of the celestial sphere, and long-term observations. More about these requirements will be discussed in Section 5.2.

Once a transiting exoplanet is detected, additional follow-up observations, using ground-based or orbital telescopes, can be conducted to learn more about the planet. Specifically, scientists are interested in conducting spectroscopic observations to better understand the atmosphere of the planet if it has one.

Important properties of the planet, such as whether or not the planet may be habitable, can be deduced from such follow-up observations.

The approximate time of the transit must be known to conduct such operations because telescope time must be reserved in advance, and typically time on powerful telescopes is expensive, so long-term surveys to detect and spectrally characterize these planets at the same time is infeasible with the resources available to exoplanet hunters. ExoplanetSat is an attempt to lower the cost of detection so that follow-up observations are on known schedules with high priority and precise known timing.

The primary stakeholders of this mission are MIT, JPL, and Draper Lab, and these three groups are already collaborating on this mission concept. Dr. Sara Seager serves as the Primary Science Stakeholder (PSS) since she is the principal investigator (PI) for the mission already. Several of her students have been developing and testing mission hardware that makes this mission possible.

Due to the nature of the data products, there is potential for Secondary Science Stakeholders (SSS) in the field of stellar astrophysics to have interest in this mission (see Section 5.2.1.2); however, initial value modeling shows that the data products are already far more valuable than the data they currently have, and that the value can only be changed with design variables that are not being considered in this study (see Section 5.3.2). Additionally, this case study is intentionally being kept simple to demonstrate the proposed methods with one science goal and one PSS, so the value of this mission to the field of asteroseismology will not be considered in this case study. A future iteration could include their inputs.

5.1.2 Mission Concept of Operations

This section introduces how operations on the ExoplanetSat mission are conducted at both the single-satellite and multi-satellite levels.

5.1.2.1 Single Satellite Model

Three different models for an individual ExoplanetSat asset have been designed: 3U, 6U, and 12U CubeSats. All three use similar optics and control hardware that make it possible to point such a small spacecraft with the precision necessary for the mission to be feasible.

An exploded view of the first design, a 3U CubeSat, is shown in Figure 5-1. It uses many off-the-shelf components currently available from standard CubeSat retailers. The optics is a lens that is available to professional photographers that has been modified for spaceflight. Stability is achieved through a two-stage system; the course pointing control is achieved by a standard reaction control wheel unit, while fine pointing control is achieved by a piezocontroller. This combination achieves a 3σ pointing precision below 8 arcseconds, more than an order of magnitude better than the precision achieved in the BRITE constellation or any other satellite in its weight class [267]–[270].

The focal plane of the satellite, shown in Figure 5-2, consists of two detectors: one that conducts the primary science mission and one that tracks stars to stabilize the pointing [271]. The pixels on the science detector are larger than the full-width half maximum (FWHM) of the photons hitting the detector because in order to achieve the required photometric precision, all the photons from a single star must land on a single pixel. This controls for a number of variables that reduce noise. The star tracker detector is intentionally defocused and uses smaller pixels because star centroiding is made accurate to within $1/20^{\text{th}}$ of a pixel using this technique [272].

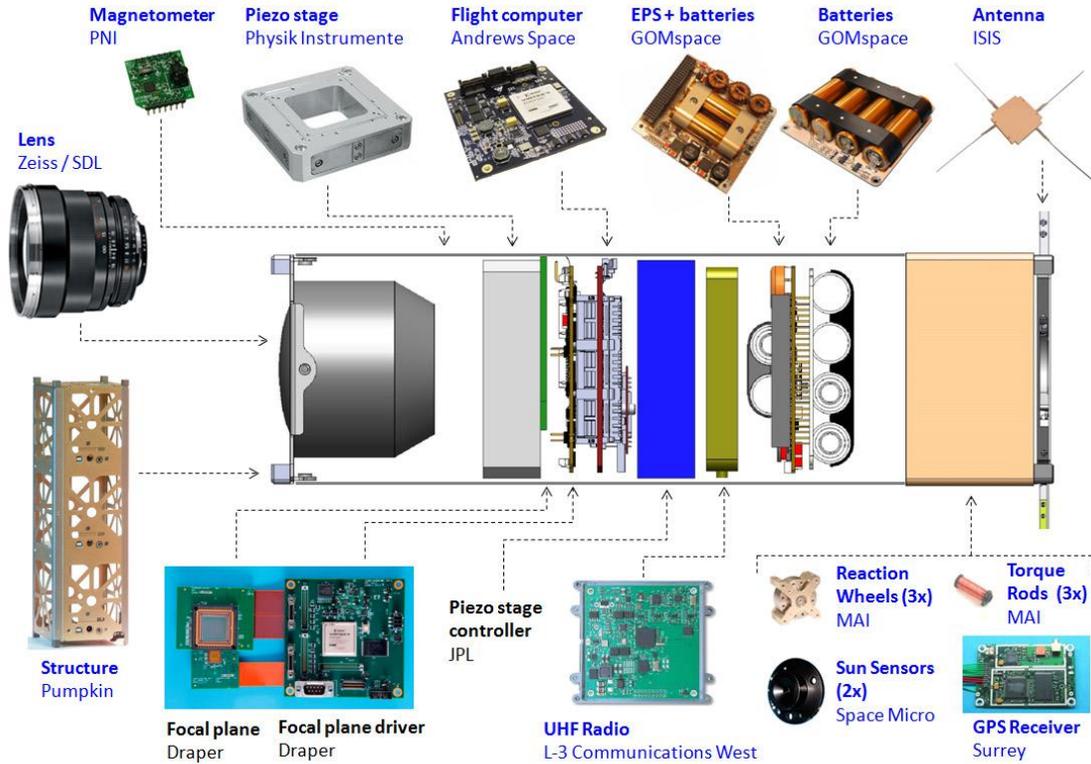
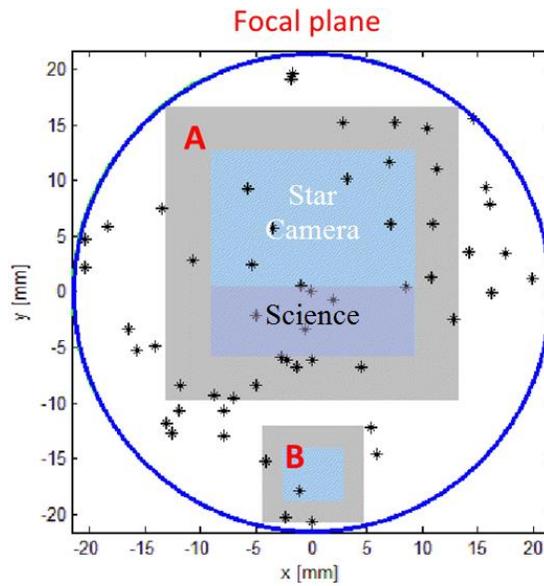


Figure 5-1: Exploded view of the interior components of a single ExoplanetSat satellite (reprinted with permission from [273]).



- A: Science / Star Camera CMOS (18 μm pixels)
- B: Star Tracker CMOS (5.8 μm pixels)

Figure 5-2: Focal plane showing the detectors on ExoplanetSat (reprinted with permission from [271]. Image credit: Matt Smith)

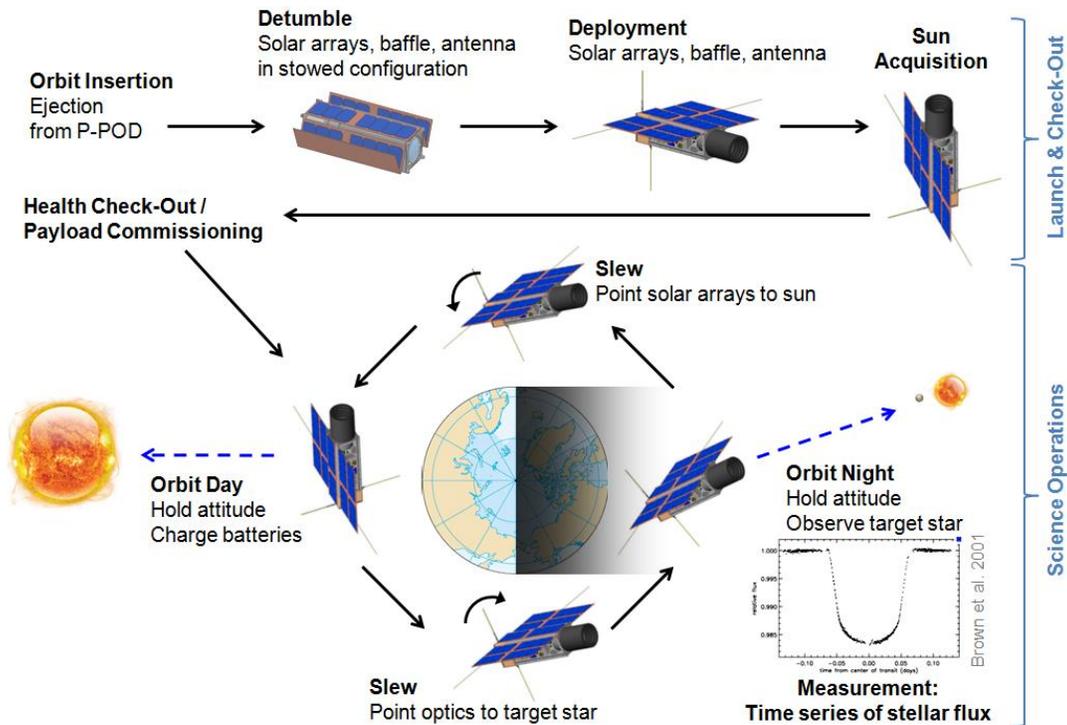


Figure 5-3: Concept of operations for single Exoplanetsat asset (reprinted with permission from [273]).

5.1.2.2 Single Satellite Operations

The concept of operations for a single asset of ExoplanetSat is shown in Figure 5-3. After launch and check-out operations have been completed, the satellite enters the science operations phase of its mission. As the satellite transitions from orbit night to orbit day, it points the solar arrays toward the Sun to recharge its batteries. The satellite will also be scheduled to downlink data during orbit day. As the satellite transitions into orbit night, it will slew to point at a target star and hold that position for the duration of the orbit night. This process repeats every orbit for a single target.

In order to observe a transit, the target star must be visible for approximately 20 minutes per orbit throughout the entire year. This is a challenging design goal for a number of reasons; not only do other photometric requirements limit ExoplanetSat’s pointing ability (Solar Exclusion Angle, Earth Grazing Angle, Lunar Exclusion Angle), orbital precession and seasonal variations make it impossible for a single satellite in any orbit to observe a single star for at least 20 minutes at a time for an entire year (see Section 5.3.2). This is why multiple satellites must work together to detect transiting exoplanets.

5.1.2.3 Full Mission Operations

ExoplanetSat is expected to be deployed in three phases. Phase 1 is meant to be a technology demonstration only: launch a satellite, observe known transiting exoplanets with short period orbits to characterize the performance of the satellite, and prove that the concept can work. Phase 2 is to launch more satellites to conduct follow-up measurements of known exoplanets that may or may not be transiting exoplanets; they have been identified by radial velocity measurements, but they have not been confirmed to eclipse the star relative to Earth.

Phase 3 is when the detection of *new* exoplanets will occur. The design of the full constellation of Phase 3 is the primary goal of this case study. Phase 3 is split into two segments, Survey Mode and Target Mode.

Survey Mode, which is expected to last anywhere from two to six months, is the mode wherein as many target stars as possible will be observed for months at a time to constrain the stellar inclinations of the target population. Typically these observations need to be made over a period of several stellar rotations; for reference, the Sun's rotation period is 24.5 days.

The stars that are determined to be within an acceptable inclination range will then be shifted to Target Mode, where they will be observed for approximately three years. The constellation will be designed in such a way that gaps in observations will be less than a satellite's orbital period.

5.1.3 Team X Study and Results

In August 2013, Dr. Sara Seager and several of her graduate students traveled to JPL to meet with Team X to conduct a concurrent engineering study of the ExoplanetSat mission. The goal was to explore design options and complete a configuration design with a detailed parts list and cost estimates.

The two-day design session proved valuable in many ways. Team X validated the feasibility studies the graduate students on the team had published. Most notably, the pointing accuracy and mass models were validated to show that such a mission was not only possible but also fit within the proposed budget. Team X also modeled all the subsystems of an individual ExoplanetSat satellite to characterize many of the engineering constraints, including thermal cycling, link budget, power budget, and orbital lifetime.

At the end of the two-day session, Team X presented two design options to the PSS, one with eight 3U satellites in a terminator, Sun-synchronous orbit, the other with six 12U satellites in a 45 degree orbit. For the 3U option, the terminator orbit was chosen because thermal fluctuations in the orbit day-night cycle could affect the CMOS detector and increase the thermal noise; the terminator orbit practically eliminates thermal fluctuations so detector noise can be easily characterized and subtracted from the total signal.

Team X also provided cost estimates for these two options. While these cost estimates for a single satellite may be realistic, their assumptions about how cost scales with the number of units are up for debate. Data suggest that the learning curve discounts they used may be too low, meaning the final cost estimate may be too high, especially as the number of satellites in the constellation rises.

However, one critical activity Team X failed to properly conduct was science value modeling to understand how exactly value is being derived in this mission. The orbits Team X chose make perfect sense from a traditional systems engineering perspective; in fact, these are perfectly acceptable orbit choices for both Phase 1 and Phase 2 of the constellation design, wherein target stars only need to be covered by a satellite during scheduled, known observation windows. However, Team X failed to understand how the pointing constraints of the instruments and the orbit choice fundamentally affected the annual coverage of the sky and therefore the value delivered to the PSS in Phase 3.

A three-year timeline of the coverage of a selection of stars by the constellation that Team X proposed is shown in Figure 5-4. The selected stars are located along the celestial prime meridian from the north pole (RA = 0, Dec = 90, top of figure) to the equator (RA = 0, Dec = 0, bottom) and are spaced in ten degree increments. Moving left to right, the colored portions represent times of the year that the star is observable for at least 20 minutes per orbit by any satellite in the constellation.

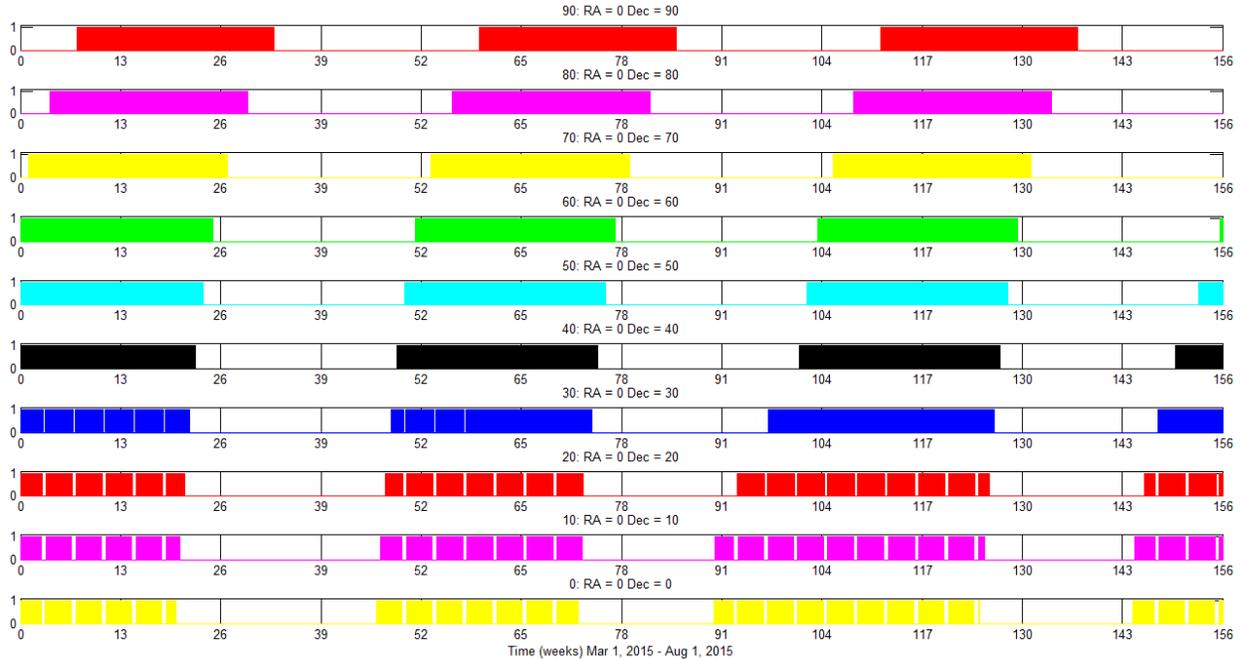


Figure 5-4: Three-year coverage of several stars by a constellation in a terminator, Sun-synchronous orbit (adapted from a model originally developed by Mary Knapp).

There are two important patterns to observe in Figure 5-4. First, the stars closer to the equator have gaps in coverage approximately once a month that last for a few days. This pattern is caused by the Moon and the Lunar Exclusion Angle in the definition of coverage. Stars with higher declinations will never experience gaps in coverage due to the Moon because the Moon stays near the ecliptic plane.

The second pattern is that there are large (~6 month) gaps in coverage for every star in this selection. This pattern is caused by the Sun and the Solar Exclusion Angle in the definition of coverage. Stars on the opposite side of the celestial sphere ($RA = 180$) would have gaps in coverage shifted by six months relative to these stars. This shift can be seen by how the coverage drifts to the right at higher declinations.

Even though Figure 5-4 only shows ten stars in particular as they are covered by one of the proposed constellations, *every* point on the celestial sphere experiences these large gaps in coverage from the proposed constellation. In fact, *there is not a single star in the sky that can be covered for a minimum of 20 minutes per orbit from any individual orbital plane given ExoplanetSat's known pointing constraints with Solar exclusion angle, Earth grazing angle, and Lunar exclusion angle (more on these constraints will be discussed in Section 5.3.2.).*

Because exoplanet transit timings are unknown, there should not be coverage gaps longer than the expected duration of a transit (~12 hours), and certainly not on the order of half a year. The coverage of both of Team X's proposed constellations proved to be insufficient for observing stars for long enough throughout the year to detect transits of any periodicity.

Additionally, the mission Team X proposed was only budgeted for one year of operations. This is not enough time to observe the minimum number of required transits (three) to confirm the existence of a

transiting, Earth-like exoplanet. Even if a single transit could be detected given the gaps in coverage, the probability of detecting the transit a second or third time would be very low.

5.1.4 Motivation

While the proposed Team X designs for a single satellite and for a constellation *do* satisfy the PSS for Phases 1 and 2 of the ExoplanetSat mission, both design options that Team X presented would not accomplish the science goals of Phase 3 and deliver no satisfaction at all to the PSS. The motivation for this case study is therefore to find feasible design options for constellations that *do* satisfy the PSS and can detect unknown transiting exoplanets.

The Responsive Systems Comparison (RSC) method was used to explore the design space and measure how effective those designs would be at satisfying the stakeholders' needs, not only in the context of what the stakeholders think they want right now and other factors the designers do not have control over, but also in other contexts in which the mission may operate instead of what is expected today and over the course of the mission lifecycle, given possible perturbations and shifts that may occur as a result of exogenous factors outside of the control of the designers or stakeholders.

The goal of RSC is not to provide “just one” design solution, the results from the method allow the stakeholders to choose among a much smaller sample set of designs and understand how those designs compare to each other, how they can change over time, and the risks and opportunities that are present in each one. However, the PSS has expressed a desire to know what the “right” solution is, one that maximizes potential and minimizes budget. For this reason, specific designs will be recommended in the conclusions for this case study but not the others in this dissertation.

5.2 Case-Specific Literature Review

There are a number of topics in the published literature that are relevant to this specific case study that were not included in Chapter 2, including the science field of exoplanets as well as the development of CubeSat technologies that have enabled this mission.

5.2.1 Exoplanet Detection

Exoplanets are detected via a variety of techniques, summarized by Seager in *Exoplanets* [274]. The two most successful techniques are the transit method and radial velocity (RV) measurements; as of February 10th, 2015, 1,889 planets in 1,188 planetary systems have been detected. 1,188 planets in 659 planetary systems have been confirmed by the transit method and 594 planets in 448 planetary systems have been confirmed by RV measurements. These two methods alone have been used to confirm 95% of all exoplanets (63% for transits, 31% for RV).

5.2.1.1 The Transit Method

An exoplanet “transits” the star it orbits when it comes directly between our viewpoint from Earth and the star, essentially eclipsing a small fraction of the star and causing a small dip in the total brightness as seen from Earth. Characteristics about the planet such as orbit semimajor axis, eccentricity, bulk mass, and bulk radius can be inferred from carefully studying the star's light curve (as shown in Figure 5-5). If an alien astronomer was detecting Earth with the transit method, Earth would transit once every 365 days, the transit would last 13 hours, and the Sun's flux would decrease by approximately 83 ppm.

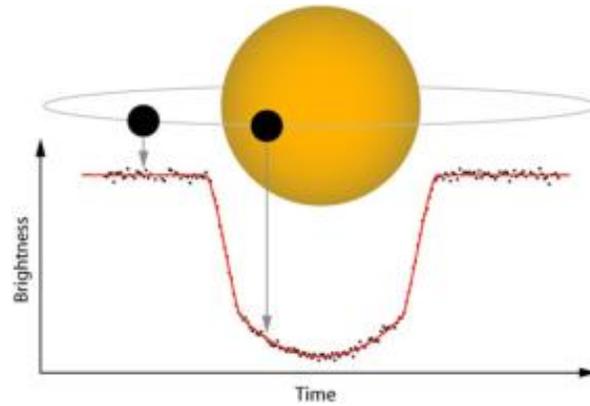


Figure 5-5: Example light curve from an exoplanet transit (reprinted with permission from [271]. Image credit: John Johnson).

This small change in the brightness of a star would be difficult to measure from ground telescopes due to atmospheric distortion even if the exact time of the transit was known. “Hot Jupiters” are detected much more easily via the transit method because the change in brightness is closer to 1% because the ratio of the projected area of the planet to the projected area of the star is much bigger. The change in brightness can be estimated as

$$\Delta B = \frac{A_P}{A_\odot} = \frac{\pi R_P^2}{\pi R_\odot^2} \quad (5-1)$$

This equation is only an estimate of the maximum possible change in brightness; there are several caveats to its use. The equation does not account for limb darkening or planets that transit off the central line-of-sight relative to the observer. It also does not account for stellar variability, planetary atmospheres, or potential moons around the exoplanet.

5.2.1.2 Exoplanet Detection Requirements

Although this technique sounds simple in theory, in practice it is incredibly difficult for many reasons that will be summarized here. The transit method only works if the planet around a given star (a) actually exists, and (b) actually transits, or partially eclipses the star relative to the view from Earth. One way to estimate whether a star could be a candidate for observation is to measure its inclination; if the stellar rotation axis is perpendicular to the line-of-sight, there is a much better chance of a transit because planets typically orbit close to the plane of stellar rotation (for instance, Earth’s ecliptic plane is ~7 degrees from the Solar rotation axis).

Asteroseismology provides a way to determine stellar inclination using photometric observations similar to those that would be used to find transits [275]; however, the inclination i is not measured directly, only $\sin(i)$, so precisely measuring stellar inclination within $90^\circ \pm 5^\circ$ is challenging and can be imprecise. This means it is easy to rule out stars for observation if their stellar rotation axes are outside of this range, but because uncertainty grows as the inclination gets closer to the desired inclination, many candidates must be selected for long-term observations.

Asteroseismology is both a blessing and a curse; stellar inclination is typically measured by variations in brightness as sunspots move across the star [276]. Those same variations in brightness that make it possible to downselect a list of target stars also make it more difficult to determine whether or not a

change in brightness is due to a transiting planet or to normal fluctuations in stellar brightness. Only stars that have low variability make good candidates for finding transits.

Mallen-Ornelas et al. developed a model for estimating the likelihood of finding a hot Jupiter using ground telescopes observing for 12 hours at a time over a period of a few months for the Extrasolar Planet Occultation Research (EXPLORE) project. [277]. Beatty and Seager expanded this work and developed a statistical model that estimates the probability of detecting a transit given an uncertain measurement in stellar inclination $\psi \pm \sigma$, an estimate of the ratio of the stellar radius to the planet semimajor axis R_*/a , and a distribution of possible differences between the stellar inclination and planetary orbit λ [278]. They determined that in order to have 95% confidence of finding an exoplanet with $R_*/a = 1/215$ given a standard error in inclination measurement $\sigma = 5^\circ$ over a possible range of planet inclinations $\lambda = \pm 7.5^\circ$ in the minimum possible time, 830 stars would have to be observed photometrically to determine inclination via asteroseismology, of which 118 would be possible candidates for long-term observations.

Finally, in order for a planet to be confirmed as detected, the transit must be observed a minimum of three times. This means that for an Earth-sized planet, it would require up to three years of continuous observation to confirm a planet's existence.

5.2.2 Similar Science Campaigns and Missions

ExoplanetSat will not be the first mission dedicated solely to exoplanet detection, nor is it the only one currently in development, but previous missions and other missions currently under development lack the capability to do what ExoplanetSat has the potential to do.

5.2.2.1 Previous and Current Missions

The Microvariability and Oscillations of Stars (MOST) telescope is Canada's first space telescope and was launched in 2003 [279], [280]. It is the first spacecraft dedicated to the study of asteroseismology, but it was also useful in the study of exoplanets because of the high-precision photometry it could perform. Data from MOST was used to measure the albedo of planets and detect "Super Earths," planets that are only a few times larger than Earth [281], [282].

The CoRoT (Convection, Rotation and planetary Transits) minisatellite was a mission led by the French Space Agency that was launched in late 2006 to find short-period transiting exoplanets and to conduct asteroseismology by measuring solar-like oscillations in stars [283]–[285]. In addition to detecting many hot Jupiters, this mission helped characterize acoustic modes of stellar atmospheres of stars with unknown inclinations [286]. Although this mission could perform high-precision photometry, the shorter observation times needed for asteroseismology and hot Jupiters were not conducive for finding Earth-sized planets, nor could this satellite cover large areas of the sky.

The Kepler mission is arguably the most well-known and successful mission to observe transiting exoplanets. The monolithic satellite carried a telescope into an Earth-trailing, heliocentric orbit and observed a 100 square degree patch of the sky, or about 0.25% of the whole celestial sphere [287]. However, there were over 223,000 target stars within this small patch of sky with stellar magnitudes of $m_v < 14$. About 136,000 are main sequence stars, and that list was further downselected by eliminating giants and white dwarves, leaving approximately 100,000 stars for study [288]. Asteroseismology measurements characterizing the inclinations of stars further limited this target pool after the mission began [289], [290].

While Kepler is capable of performing high-precision measurements for long periods of time, the area of the sky it views is small compared to the celestial sphere, and many of the candidates are too dim for follow-up spectroscopy to study exoplanet atmospheres.

The Gaia space telescope was launched in December 2013 and aims to build a three-dimensional model of the galaxy using photometry data from one billion stars collected over a five-year period. Many hot Jupiters are expected to be detected as a result of this data, but the gaps in coverage make the probability of finding an Earth-sized planet low [291].

The critical traits for detecting Earth-sized planets and which current missions are capable of delivering those traits are summarized in Table 5-1. ExoplanetSat is the only mission that can deliver all three of these traits.

Table 5-1: Comparison of current exoplanet detection missions and requirements to find Earth-sized planets.

Detection Requirement	MOST	CoRoT	Kepler	GAIA	ExoplanetSat
High Photometric Precision	✓	✓	✓	✓	✓
Long-Term Observations	X	X	✓	X	✓
Large Area of Sky Covered	X	X	X	✓	✓

5.2.2.2 Future and Proposed Missions

Future planned exoplanet detection missions include the Transiting Exoplanet Survey Satellite (TESS), the CHAracterizing ExOPlanets Satellite (CHEOPS), and PLANetary Transits and Oscillations of stars (PLATO). TESS is scheduled to launch in August 2017 and will also search for planets using the transit method [292]–[294]. Unlike Kepler, this mission will conduct a 2-year, full-sky survey of the brightest 100,000 stars; however, this survey will be conducted in 27-day segments, with the northern celestial hemisphere surveyed in the first year and the southern in the second. While this is enough time to find hot Jupiters, it is not enough to confirm the existence of Earth-sized planets with periods close to one year.

CHEOPS is also scheduled to launch in 2017 and will examine known transiting planets on known bright, nearby stars [295]. This mission aims to accurately measure the radii of exoplanets, though it works most effectively when the transit times are known, so while this mission will accurately characterize known exoplanets, its primary purpose is not to detect new ones. It will also conduct follow-up measurements of the exoplanet candidates detected by TESS.

PLATO is a proposed ESA mission that will search for exoplanets in a similar fashion as ExoplanetSat, but rather than using distributed satellites, a single satellite will have 34 small telescopes packed into one spacecraft bus [296], [297]. These telescopes will have a larger field of view than Kepler’s telescope and will observe up to one million stars in one region of the sky. PLATO’s goal is exactly the same as ExoplanetSat’s and it is therefore considered a significant monolithic competitor. ESA selected PLATO as a Medium-class mission, meaning its budget will be near one billion dollars [298], [299].

The key desired traits for finding Earth-sized planets and which future missions are capable of delivering those traits are summarized in Table 5-2. Only PLATO has the capability of achieving the same goals, but at a much higher price than what has been proposed for ExoplanetSat, and without the lifecycle properties that make a DSS mission beneficial for long-duration missions.

Table 5-2: Comparison of future exoplanet detection missions and requirements to find Earth-sized planets.

Detection Requirement	ExoplanetSat	TESS	CHEOPS	PLATO
High Photometric Precision	✓	✓	✓	✓
Long-Term Observations	✓	X	X	✓
Large Area of Sky Covered	✓	✓	X	✓

5.2.3 Performance Modeling in ExoplanetSat

Because the individual assets of this mission have already been designed and cost modeling was conducted by Team X at JPL, little additional literature review is relevant that is specifically related to mass and cost modeling on ExoplanetSat that has not been covered in Chapter 2. In comparison, Chapters 6-8 will include more literature review because the mass and cost models in those case studies were built from the ground up rather than being inherited. However, there is still literature related to the performance of ExoplanetSat, which will be reviewed here.

Over the past five years, Dr. Seager’s students have conducted a number of technology tests and simulations to prove that a satellite as small as 3U (10 cm x 10 cm x 30 cm) with a total mass of only 5kg has the capability to perform the photometric measurements required for detecting exoplanets via the transit method. Among the many challenges they tackled were the pointing control, detector sensitivity, communications, and configuration design.

One graduate student demonstrated that high-precision pointing and attitude estimation could be achieved in small, hardware-constrained spacecraft in his doctoral dissertation using ExoplanetSat as his primary case study [300]. Through his work, he demonstrated that sufficient pointing accuracy could be achieved using standard CubeSat reaction control wheels in conjunction with a high-precision piezoelectric control device. Using specialized control algorithms developed for this purpose, he demonstrated how an ExoplanetSat unit could achieve 90 x 30 arcseconds (3σ error) course pointing with reaction control wheels. With a second-stage piezoelectric controller, this could be improved to below 8 arcseconds (3σ).

Another graduate student developed a coverage tool in STK that links to MATLAB to find the availability of target stars for 20 minutes per orbit over long periods (~3 years) for a single ExoplanetSat asset in a given orbit [301], which was based off an earlier model built for the ExoplanetSat feasibility studies [268]. This model helped identify when different stars with known transiting planets would be available during Phase 1 of the ExoplanetSat mission, which is the technology demonstration to prove the concept. This coverage tool was heavily modified in the RSC Performance Model.

5.3 The RSC Model for the ExoplanetSat Mission

The Responsive Systems Comparison Model for ExoplanetSat considers only the design of the constellation; since the individual satellite has passed the critical design review, the author was not given permission to try to change those variables, so they will not be present in the tradespace. While this is not ideal from the perspective of developing a mission from the ground up using the RSC method, it does highlight how real-world problems with many additional constraints can still be solved or explored further with this method. This section details the model development process and explicitly states the assumptions that are being used either to simplify the problem or to cover up for missing data that cannot be provided yet (but can be in the future).

It is important to remember that these models went through several iterations over the past two years, especially as a result of feedback between the value and performance models. Although this case study's design space is the simplest of the four being explored, the stakeholder ambiguity, computational complexity, and pre-existing constraints make this case study nontrivial and valuable for demonstrating how the RSC method can be applied in real-world scenarios in addition to only Pre-Phase A studies.

5.3.1 ExoplanetSat RSC Phase 1: Value Modeling

Value modeling for ExoplanetSat was not a simple process. Since this was the first case study conducted in this work, there was a significant learning curve to using this VCDM and eliciting stakeholder preferences. Additionally, because no known literature exists on year-round observations of stars from constellations of satellites, there was no basis on which to communicate with the stakeholders without a working performance model. There were many iterations between the value modeling and performance modeling stages, more than any of the other case studies, both due to this being the first case study and due to the lack of published literature with how this mission measures its success.

In the first major phase of the RSC method (see Figure 4-2), stakeholders' needs and resources were identified, a value model was constructed that represents the stakeholder's perception of value of the system, and additional contexts that could change either the performance of the design or the stakeholder value model were identified. In the end, the multi-tiered MAU structure presented in Chapter 4 (Figure 4-5) degenerated into a SAU value model. The single attribute that completely captures the stakeholder's perceived value of the mission is the Expected Number of exoplanets detected, *EN*. This section describes how the value model was derived and what outputs flow into the performance model.

5.3.1.1 Value Modeling Process 1: Value-Driving Context Definition

In this step of the value modeling phase, the stakeholders, primary mission goals, available resources, and stakeholders' needs that drive the value of the mission were identified. Unlike with the other case studies, this mission is not just a concept study. Stakeholders had already been identified, a prototype satellite has been built, and several students have earned masters and Ph.D. degrees based off their work on this project. Dr. Sara Seager is the PSS, and Draper Labs and JPL also have interest in being the first to launch a constellation strictly for astrophysics.

The primary goal of the mission is to detect **Earth-sized** planets around **nearby, Sun-sized stars**, especially planets that are in the "habitable zone" that make them suitable candidates for harboring life. This is the only goal of the mission and is well-defined. "Nearby, Sun-like stars," for purposes of this mission, include spectral types F, G, K, and possibly M; luminosity classes IV and V (subgiant and main sequence stars); and apparent magnitude m_v between 4 and 8. Another requirement due to the nature of the transit method is that the stars must be non-variable. There are 3,769 stars that fall within these parameters and are considered targets for this mission.

The initial proposed budget for ExoplanetSat was \$30 million. This is not a hard upper limit, however, and it was requested that the case study explore options with hundreds of satellites to see just how much a mission might cost as the number of assets increased dramatically compared to the cost ceiling that the Team X design study assumed.

Despite the feasibility studies on the hardware design, there is still much uncertainty in the quality of the data based on how long a star is observed per orbit. The initial requirements stated that observations

needed to be at least 20 minutes at a time assuming a 90-minute orbit for the entire year, but this is not stated explicitly in published research nor has it been derived yet from an SNR calculation based on the hardware design.

Additionally, others on the team argued that *complete* coverage was necessary, meaning a camera must be on the star 24/7/365. This is because although the transit of an Earth-sized planet may take 13 hours, the ingress and egress times, i.e. the amount of time it takes for the full diameter of the planet to move across the limb of the star, are on the order of 7 minutes. The ingress and egress light curves are essential to *characterizing* an exoplanet, but not for detecting or confirming it. For this reason, “complete coverage” of a star, i.e. the star has a satellite focusing on it nonstop for the entire period of the mission, was deemed unnecessary. Instead, “once per orbit coverage” for a given amount of time per orbit was considered a driving design consideration that could change depending on the final SNR calculations.

5.3.1.2 Value Modeling Process 2: Value-Driven Design Formulation

In this step of the value modeling phase, the stakeholder value model was explicitly elicited that transposes the outputs of the performance model onto a scale to rank the designs by how much value they deliver to the stakeholders. This model started as a complex MAU but eventually degenerated into a simpler SAU that satisfactorily assigns a utility to each design in the design space.

Stakeholder Interviews

The initial interviews with Dr. Seager shed some light on the characteristics of the data that the mission hoped to achieve. It was well known that year-round coverage was required, and that 20-minutes per orbit was definitely acceptable, but that perhaps less time per orbit would also suffice. Full-sky coverage was desired, but it was known early on that this was impossible due to the pointing constraints of the satellite and the precise nature of the photometric measurements being used. However, the reason full-sky coverage was desired was because there was a better probability for finding targets. It was also known that the mission needed to last approximately three years in order to see three transits of an Earth-sized planet near the habitable zone of a Sun-like star.

Initially, three attributes were identified that seemed to capture the needs of the PSS: number of targets available, number of targets that could be observed at once, and mission lifetime. These were presented to Dr. Seager to sketch the single-attribute utility (SAU) curve definitions and to find the weights among the three attributes. However, during the following interview, Dr. Seager was clearly uncomfortable giving firm, numerical answers on her perceptions and how they changed with different attributes levels. This is because in reality, these attributes did not properly reflect her perception of how value was being derived from this mission.

The stakeholder interviews were further complicated when the performance model was upgraded to include time variations in sky coverage. Shorter coverage time meant larger sky coverage, so there was a tradeoff between coverage time per orbit, which was not an attribute to begin with, and sky coverage. This caused even more problems and the value model was made even murkier because of the improvements in the performance model. Something was clearly not right.

After several more improvements in the performance model, a revelation occurred: Dr. Seager couldn't properly express her satisfaction with changes in these attributes, not because she couldn't grasp the subtle complexities of risk analysis and answer questions with the ASSESS software, but because these

attributes do not represent her mental value model. The *only* thing Dr. Seager wants from this mission is to find exoplanets. Everything else was simply details. If a systems engineer or mission manager who was unfamiliar with the scientific literature had tried to solicit stakeholder opinions and construct a value model, he or she probably would have not made this connection.

Rather than trying to build a mathematical value model with attributes that combined to represent the stakeholder's *mental value model*, it was more appropriate to use a *constructed value model* that more accurately represented the goal of the mission. In this case, the constructed value model already exists. A paper by Beatty and Seager showed how to calculate the probability of detecting at least one Earth-sized exoplanet with known parameters, although their paper stopped short of mathematically defining the *expected number of exoplanets* that would be detected [278], which is the *only* attribute Dr. Seager cares about as the PSS. This attribute can be calculated from what had already been identified as value-driving attributes to the system: the number of possible targets, the number of observable targets, the schedule for observations, the length of the mission, and other factors that can be considered context variables.

Once this was realized in a later stakeholder interview, Dr. Seager was able to quickly plot her satisfaction with the ASSESS software. The utility function from the final stakeholder expectations solicitation interview are shown in Figure 5-6.

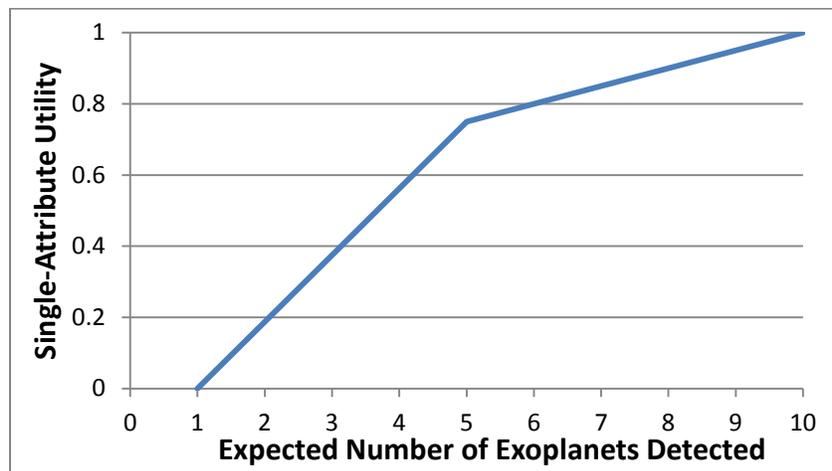


Figure 5-6: Single-attribute utility for the expected number of exoplanets detected by ExoplanetSat.

From this utility function we see that any mission that has an expected value for the number of exoplanets found that is less than one is unacceptable. The function also shows diminishing returns past the point where *EN* is greater than five in comparison to first five planets detected.

Although the MAU structure shown in Figure 4-5 was already reduced because there was only one science goal, it was reduced further because there is only one attribute, so MAU degenerated to SAU. The interview process still resulted in a proper value model, just one that was not expected. MAUT was being used to develop a metric for a satisfaction, when in reality the metric for calculating this satisfaction was a constructed model that was already mostly developed. The following section shows the mathematical formulation of the constructed model.

Constructed Value Model

From Beatty's and Seager's paper, we know that the probability of finding *at least* one transiting exoplanet in a long-term study of stars is

$$P_{>0} = 1 - Bi(n_i, 0, P_{det}) = 1 - \binom{n_i}{0} (P_{det})^0 (1 - P_{det})^{n_i-0} \quad (5-2)$$

where the binomial distribution $Bi(n, s, p) = \binom{n}{s} (p)^s (1 - p)^{n-s}$, $s = 0$ is the probability that zero exoplanets will be found based on the number of targets observed for long periods of time n_i and the probability of detecting an exoplanet P_{det} . The expected number of exoplanets that will be found by a given constellation design is therefore expressed as

$$EN = \sum_{i=0}^{n_i} Bi(n_i, i, P_{det}) \cdot i \quad (5-3)$$

The probability of detecting an exoplanet around a star P_{det} is

$$P_{det} = P_{exist} \cdot P_{trans} \cdot P_{sens} \quad (5-4)$$

where P_{exist} is the probability that a planet exists around the star, P_{trans} is the probability that that planet's orbit is aligned so that it transits the star relative to the line of sight from Earth, and P_{sens} is the probability that the instrument has a sufficient SNR to detect the transit if it does occur. From gravitational microlensing data, it is safe to assume that $P_{exist} = 1$ [302]. The probability that a planet will actually transit if the stellar inclination is unknown is

$$P_{trans} = \frac{R_*}{a} = \sin(\theta) \quad (5-5)$$

However, if the stellar inclination ψ is known within some Gaussian uncertainty $\pm\sigma$, the probability is

$$P_{trans} = \frac{\int_{90-\theta}^{90+\theta} \sin(\psi) \cdot f_{Gauss}(\psi_m|\psi) d\psi}{\int_{-180}^{180} \sin(\psi) \cdot f_{Gauss}(\psi_m|\psi) d\psi} \quad (5-6)$$

with the Gaussian distribution

$$f_{Gauss}(\psi_m|\psi) = \frac{1}{\sqrt{2\pi\sigma^2}} \exp\left[-\frac{(\psi - \psi_m)^2}{2\sigma^2}\right] \quad (5-7)$$

There is also an additional probability factor that relates the probability of detection to the speed at which the survey is conducted. In order to minimize survey time, candidates could be excluded prematurely. While this may be a goal when the target selection is abundant, when the area of the observable sky is limited and therefore target selection is limited, this can be ignored.

The probability that the instrument will observe a transiting planet if it exists P_{sens} is

$$P_{sens} = P_{photo} \cdot P_{life} \cdot P_{sche} \quad (5-8)$$

where P_{photo} is the probability that photometric precision of the instrument is high enough to capture the transit, P_{life} is the probability that the long-term observations will be long enough to capture the desired number of transits to confirm the planet, and P_{sche} is the probability that the schedule of observations is

such that the satellite will actually be able to observe the transit. Since P_{photo} is determined by design variables outside of the control of the tradespace, it is assumed that $P_{photo} = 1$.

Although there is still an ongoing debate about how much time is necessary per orbit to observe a star in order to detect the transit, the PSS has maintained throughout this process that 20 minutes per orbit with one satellite is enough observation time to detect a transit. However, when multiple satellites are sharing the responsibility for making observations in a single orbit, there is going to be added noise because not only do different cameras have different sensitivities, different pixels on the same camera can have variations in sensitivity. As a result, any design that requires more than one satellite to observe the light curve per orbit reduces the probability of detecting the transit.

A heuristic model for probability of detection due to the schedule observations per orbit P_{sche} as a function of observation time and number of satellites was constructed and approved by Dr. Seager. This assumption is being made in the absence of better data that will come later in the project lifecycle and can be treated as a context variable if there is still ambiguity between various probability models.

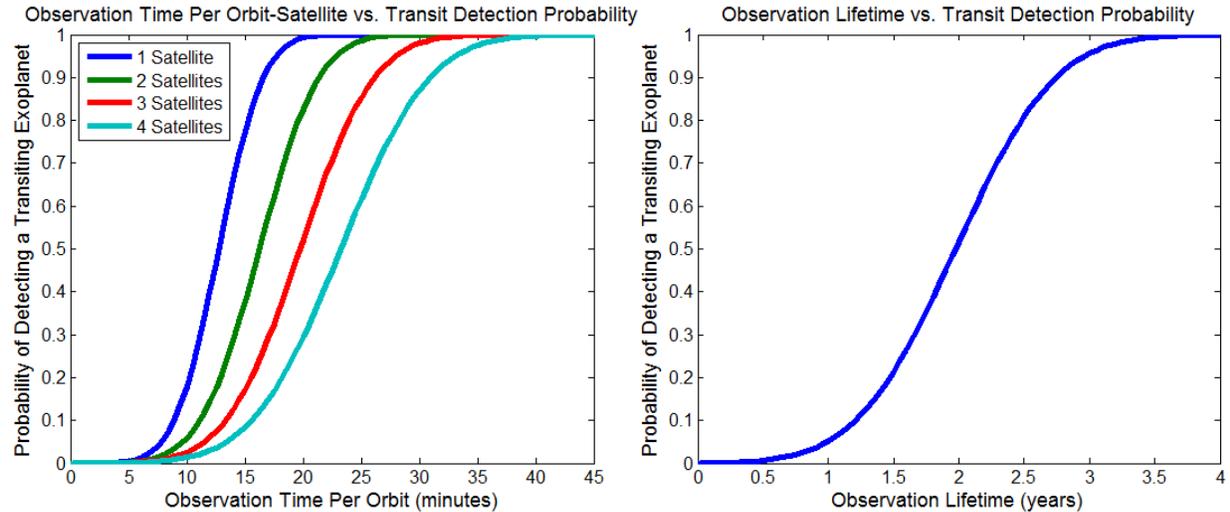


Figure 5-7: Heuristic functions for probability of transit detection for observation time per orbit and number of shared satellite sample times (left) and observational lifetime (right).

The lifetime versus probability curve to determine the value of P_{life} was also developed heuristically based on the assumption that 3 years of observations gives 95% confidence that the transit will have been observed at least three times. Both functions for P_{sche} and P_{life} are shown in Figure 5-7.

Finally, the last term in the expected number of exoplanets EN calculation is the number of targets that are observed for long periods of time, n_i . This is a fraction of the total targets that are available for long-term observation, n_{obs} , defined as

$$n_{obs} = n_i \sin(\varphi) \quad (5-9)$$

where $90 \pm \varphi$ is an acceptable range of acceptable stellar inclination measurements, and φ is based off ψ , σ , and λ . As Beatty and Seager showed, in order to have 95% confidence in finding an Earth-sized exoplanet in a minimum survey time, $n_{obs} = 830$ and $n_i = 118$. This is how the design of the ExoplanetSat constellation finally affects the value model; a well-designed constellation will have many

targets to choose from, while a less-acceptable one will not be able to observe as many targets, leading to a lower expected value of planets detected.

The difference between n_{obs} and n_i also represents the difference between Survey Mode and Target Mode as described in Section 5.1.2.3. A constellation will have the ability to observe some of the n_{obs} total stars in the observable area of the sky. During Survey Mode, stars are observed for several months at a time to constrain their inclinations. If the inclination of a star is too high for a transit to be likely, the star would be removed from the list of observations. If the inclination can be constrained within $90 \pm \varphi$, the star would be placed on the list of candidates to observe for long-term observations in Target Mode.

Better constraints on inclination and a lower φ mean more stars are removed from the list of possible candidates, so n_i/n_{obs} is lower, but it also means that the probability of detection P_{det} is higher, so resources aren't wasted on stars where transits are unlikely in Target Mode.

5.3.1.3 Value Modeling Process 3: Epoch Characterization

Epochs are the unique combination of design variables and stakeholder expectations. In this phase of value modeling, context variables and alternative stakeholder expectations were explored and possible answers for known unknowns were created. These alternative contexts and value models are used in the final phase of the RSC method, Epoch-Era Analysis, and are summarized in Table 5-3.

Performance Context Variables

As discussed in the previous section, there are multiple unknowns associated with the probability of detecting transiting exoplanets. Among the many variables that could be tested at different values that are outside of the designer's control are the precision of stellar inclination measurement, the distribution of planetary inclinations relative to a target star, the added noise from multiple satellites combining observations, the long-term time scales, and whether or not a planet actually exists.

The variable that has a large, known influence on the detection probability is the ratio of stellar radius to planet semimajor axis R_*/a . For a planet exactly like Earth orbiting a star exactly like the Sun, $R_*/a = 1/215$. However, the list of target stars includes not only Sun analogs but also stars that are smaller. Smaller stars have habitable zones that are closer to the star; therefore, the average value of R_*/a for the target set of stars could be higher. A change in R_*/a also alters the lifetime probability curve (Figure 5-7, right) because planets closer to their stars have shorter orbital periods, thus the length of time that a star must be observed in order to observe three transits goes down.

A contextual unknown that is related to the design is the ability of a single satellite to study multiple stars at once. Although many stars can fit within the science window of the detector plane (see Figure 5-2), the chances that more than two or three stars that are determined to be within an acceptable inclination range being close enough to each other to be studied simultaneously goes down as the total number of targets goes down. Since the number of target stars observed for long periods n_i greatly affects the expected number of planets found and thus the value proposition to the PSS, the number of targets per satellite is an important context variable to consider.

One context variable that could be considered is the probability that the satellite's photometric precision is high enough to detect a planet, P_{photo} . The photometry requirements for finding a planet around a bright star ($m_v = 4$) are much more relaxed than a dim one ($m_v = 8$). However, because the set of

possible candidates is still in flux, it is difficult to separate the probabilities of detecting a transit for individual stars visible to a given constellation. A future iteration of this tradespace could include multiple sets of target stars separated by magnitude, but because the proposed design requirement for the satellite’s photometric precision is to be capable of observing transits on stars as low as $m_v = 8$, it is assumed they will be detected in this implementation.

Cost Context Variables

There are many assumptions within the cost estimates provided by Team X. Their cost models assume flat rates for additional units on not only the manufacturing of additional assets but also on wrap factors such as management and systems engineering costs that would not scale at the same rate as manufacturing costs in the production of additional assets. Both learning curve slopes for production and wrap factors can take different values to see how the value proposition and cost per unit may vary as the number of units goes up. A custom model was also built as an alternative cost estimating tool that is based off the Team X model but changes the scaling rates for labor and uses a launch cost estimator based on estimates from launch providers rather than a flat-rate discount for additional launches (see 5.3.2.3 for details).

Additional context variables related to cost were also considered in early versions of the context vector, such as launch provider, bulk launch discount as a function of number of satellites launched, and cost margins, but these did little to alter the relative spacing among designs on the tradespace and these changes affected all designs relatively equally. Those variables were therefore eliminated since they only change the positions of each design along the x-axis of the tradespace.

Context Vector

The context variables and the values that they can take are summarized in Table 5-3. There are three context variables, each with three possible values; therefore, there are $3^3 = 27$ unique contexts to explore.

Table 5-3: Context variables and their values for ExoplanetSat

Context Variable	Discrete Values
Cost Model	[JPL Model, Custom 1, Custom 2]
Stellar Radius vs. Planet Semimajor Axis	[1/215, 1/150, 1/100]
Targets Per Satellite	[2, 3, 4]

If the PSS wishes to add an additional context variable that changes the probability of detection based on the observation time per orbit (see Figure 5-7 (left)), it is quite easy to do.

The *primary design context* that will be discussed in Section 5.3.3 and used as a baseline in Section 5.3.4 is when the context vector is [JPL Model, 1/215, 3].

Alternative Stakeholder Expectations

During the stakeholder interviews, the PSS was urged to consider alternative future scenarios that would alter her perception of the value model as shown in Figure 5-6. Additional variables that do affect her perception of the value of the system include the acceptance of other competing exoplanet detection missions, but although those may change the probability of the mission being accepted and implemented, they do not change how she views the system relative to the single attribute utility. As a result, there are no additional value models to consider.

5.3.2 ExoplanetSat RSC Phase 2: Performance Modeling

In the second phase of RSC, a model was built that explores many possible designs for ExoplanetSat and evaluates their performance. First, a list of design variables and possible values for those variables were selected to build the design space. Next, the model calculated the expected number of detected exoplanets EN , which is the single attribute by which a constellation is ranked according to the value model elicited in the previous phase. Finally, the model provides a cost estimate so that all the designs can be explored through tradespace exploration and EEA in the third and fourth phases of RSC.

This case study has a relatively simple design space due to the fact that the hardware design is not part of the tradespace; since the hardware design had already been mostly finalized with some work done on a prototype, there was no need to include those design variables in the tradespace (although some sensitivities were explored in the performance model to demonstrate how the design could be improved to increase the performance). However, while the design space is the simplest of the four case studies, the performance model is the most complex and requires the most computing power. This is because the performance is based primarily on coverage for long periods of time, which requires intense calculations.

5.3.2.1 Performance Modeling Process 1: Identifying Architecture and Design Variables

The first process of performance modeling is identifying what design variables drive the value of the mission as it was elicited from the value modeling phase of RSC. The focus of this mission is on designing the constellation, so variables that change the performance of the constellation are what are being explored in this phase. Unlike the other case studies in this dissertation, there are no architectural variables because only one asset is being used (a CubeSat), and there are no design variables that affect the design of an individual satellite because it has already been designed.

Quality Function Deployment (QFD) would have been used to identify the preliminary variables that would have the most effect on the design space; however, because no published research is known to exist on the subject of all-sky, long-term coverage, a QFD model would not have been very accurate. Instead, the known design variables of a Walker constellation (number of planes, number of satellites per plane, and difference in right ascension from one plane to the next), orbital altitude, and required observation time per orbit were chosen based on previous experience in the initial feasibility studies.

Iterations on the coverage model changed the design variables significantly. At first, the focus of the model was on maximizing the area of the sky that could be covered. It was soon discovered that difference in right ascension between planes had no affect whatsoever on the coverage. Altitude within the range of acceptable values also did little to affect the model outputs. The starting altitude for exploring options was 500 km. Options for lower altitudes resulted in designs with lifetimes that were far too short to be acceptable, and altitudes higher than this resulted in shrunken coverage areas that reduced the number of possible target stars to observe (altitudes above 700 km were also considered unacceptable because of radiation exposure from the Van Allen belts).

A later iteration of the performance model showed that even though Team X's suggestion of a terminator, Sun-synchronous orbit left large gaps in yearly coverage, other Sun-synchronous orbits could still observe some stars year-round, but for less than 20 minutes per orbit. No other type of orbit besides Sun-synchronous orbit can do this, but because the observation time is shorter, multiple satellites have to work together to observe the same star during the same orbit. This is in contrast to Walker constellations

wherein a satellite in one orbital plane can be completely responsible for observations of a single star for weeks at a time before transferring that responsibility to another satellite on another plane.

In addition to Walker constellation variables, the variable that related to Sun-synchronous orbits, the right ascension of the ascending node (RAAN), was added to create additional design options for single-plane constellations. More on how the design space evolved with improvements in the coverage model will be discussed in the next section. The final list of design variables, of which there are 4,000 possible combinations, is shown in Table 5-4.

Table 5-4: Design variables and their discrete values for ExoplanetSat

Design Variable	Discrete Levels
Orbital Inclination	[30, 38, 44, 52, Sun-Synch]
Number of Planes (excluding Sun-synch)	[2, 3, 4, 6, 8]
RAAN (Sun-synch only)	[180, 202, 214, 225, 236, 248]
Minimum Observation Time	[5, 10, 15, 20]
Satellites Per Plane	[1, 2, 3, 4, 6, 8, 10, 12, 16, (more for Sun-synch)]
Redundancy Per Target	[1, 2, 3, 4]

A brief interlude is warranted to explain the discrete choices for these variables. Initial models showed that constellations with either low or high inclinations (excluding Sun-synchronous orbits) covered very little of the sky, and that there must be some inclination where coverage would be maximized. The inclination options represent a swath of possible inclinations that were shown to provide some level of coverage and bound the inclination where coverage may be maximized.

The options for RAAN represent a range of non-degenerate options for Sun-synchronous orbits that goes up to but does not include the terminator orbit (RAAN = 270). Other options from different quadrants of the unit circle would give the same results as the levels shown. Minimum observation time was discretized at one-minute levels in the coverage calculations, but little value was seen in examining every discrete minute separately, so observation time was discretized in five-minute increments to reduce the size of the tradespace.

Redundancy per target represents the number of satellites observing a target from the same orbit. This way, shorter observations can be stacked to improve coverage time per orbit. However, what this makes up for by increasing the number of observable targets n_{obs} is offset by a decrease in the number of targets observed n_i .

5.3.2.2 Performance Modeling Process 2: Design Modeling

Process 2 of ExoplanetSat performance modeling builds the engineering models to evaluate how well a design satisfies the stakeholder value mode. Because ExoplanetSat does not have a particularly diverse design vector, only two performance modules are required: the coverage module and the lifetime module.

Coverage Module

By far the most extensively used, most accurate, and most complete module of the system performance model for ExoplanetSat or any of the four case studies in this work is the ExoplanetSat MATLAB/STK coverage model. This module was developed to calculate the number of possible stars a constellation can observe for long periods, n_{obs} , which is one of the key parameters for calculated EN . The model was originally developed by Matt Smith and improved upon by Mary Knapp before being extensively

modified, upgraded, expanded, and integrated into the performance model presented herein. This section details the evolution of this model and how it works.

As shown in the ExoplanetSat concept of operations (Figure 5-6), the satellite only operates during orbit night to reduce noise from the Sun. Its baffle is not effective at reducing direct sunlight noise if it is pointed within 90 degrees of the Sun, so there is a strict Solar Exclusion Angle constraint. There are also constraints for how close it may point to Earth and the Moon. These constraints were treated as parameters, and although they were varied individually to understand their sensitivities, they were not included in the tradespace because they are satellite design variables outside of the control of the tradespace. These fixed parameters are listed in Table 5-5.

Table 5-5: Important parameters for the STK coverage module of ExoplanetSat

Simulation Parameter	Value
Solar Exclusion Angle	90 degrees
Earth Grazing Angle	22 degrees
Half Angle Field of View (FOV)	14.5 degrees
Lunar Exclusion Angle	FOV + 15 degrees

The resolution of both the coverage map and the time of coverage evolved gradually over the course of exploration of the tradespace as new insights were gained from understanding the way coverage changes with changes in the design variables. The evolution of the set of test points used to understand how the area of the sky that can be covered changes with changes in the design vector is shown in Figure 5-8.

The first stage of building the model was simple, unstructured exploration of coverage of one line of stars. Different individual satellite and full constellation options were tested on stars from 0 to 90 degrees declination (from the equator to the North Pole) in 10 degree increments along the celestial prime meridian, or right ascension (RA) of 0 degrees, as shown Figure 5-8(1). The output timeline was similar to what is shown in Figure 5-4; the coverage shown represents times of year where the satellite is observable for at least 20 minutes during orbit night.

During the course of this exploration, it was noted that different orbital inclinations produce different coverage patterns over the course of a year. Satellites in low inclination orbits can observe medium-declination stars uninterrupted for approximately 6 months at a time before they are obstructed from observation by the Sun, whereas satellites in higher inclinations observing the same stars had completely different coverage patterns.

A coverage timeline for a single satellite at an inclination of 28 degrees is shown in Figure 5-9 (exactly the same stars as the timeline in Figure 5-4). Stars with lower declination (bottom of the figure) can be observed for at least 20 minutes per orbit for six months at a time, but there are small interruptions due to the position of the Moon approximately once per month. Stars at higher declinations do not experience these interruptions because they are outside of the Lunar Exclusion Angle for the entire year. However, as the declination of the target star increases (top of the figure), coverage is no longer continuous for several months at a time. This is due to natural orbital precession.

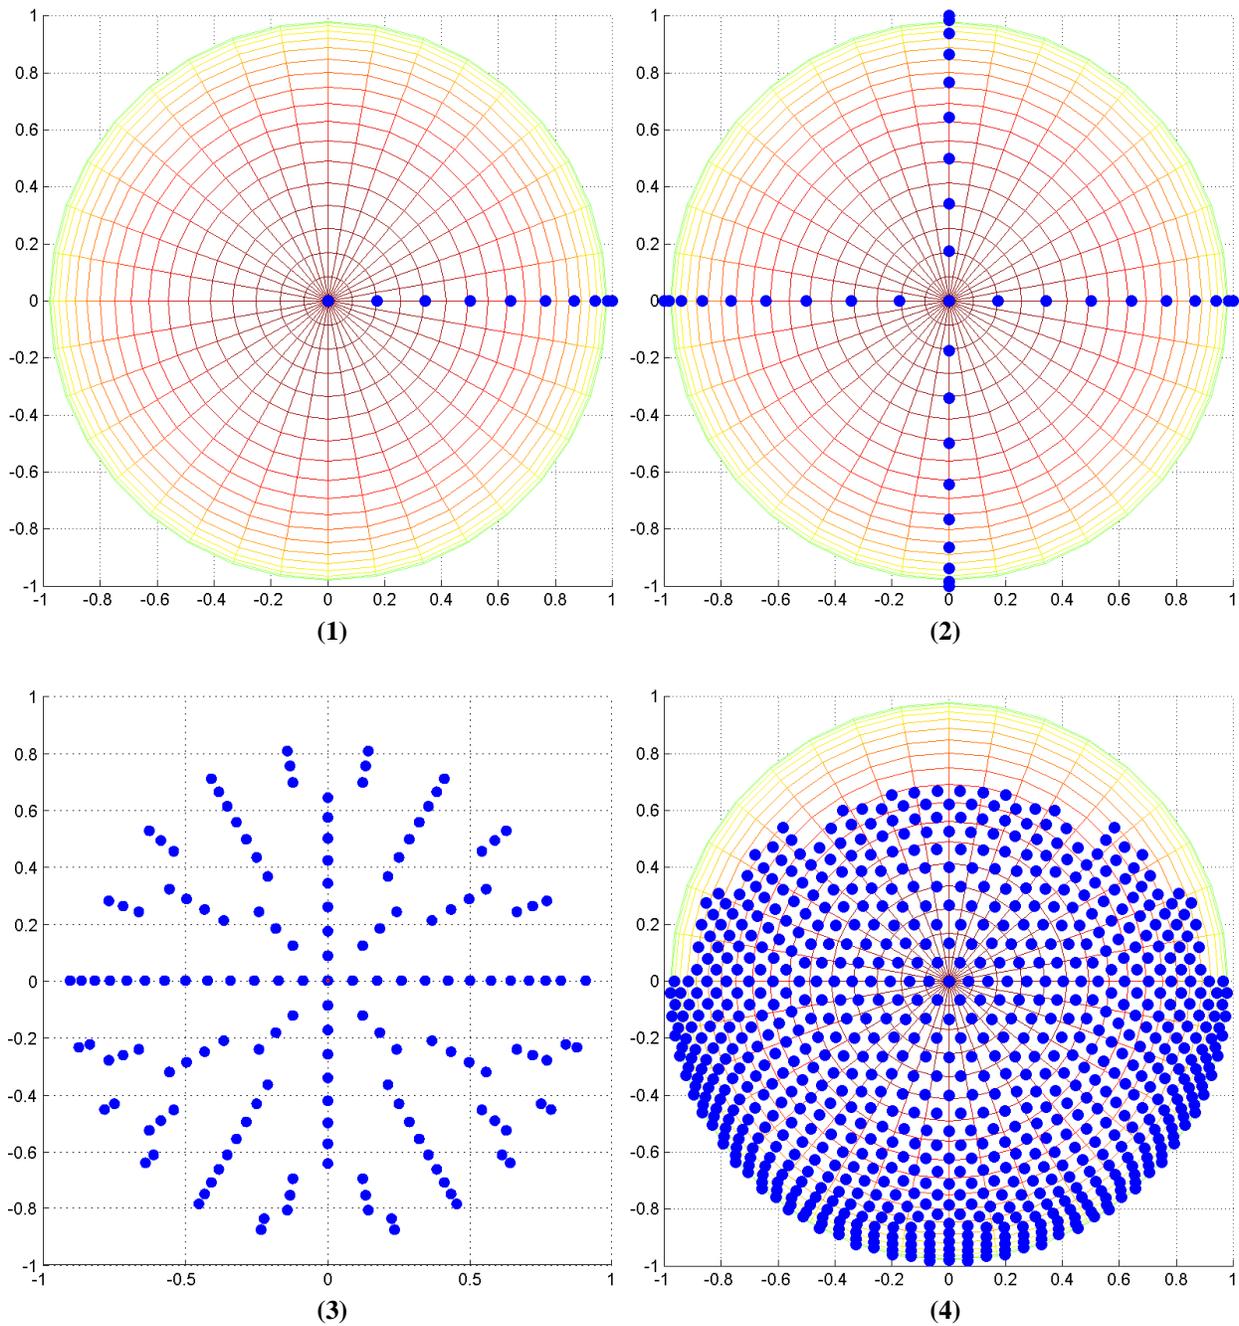


Figure 5-8: Evolution of the test stars on the celestial sphere as viewed from the celestial North Pole for (1) initial low resolution scan during Team X study, (2) low resolution 4-direction scan, (3) medium resolution scan, and (4) high resolution scan for complete coverage mapping.

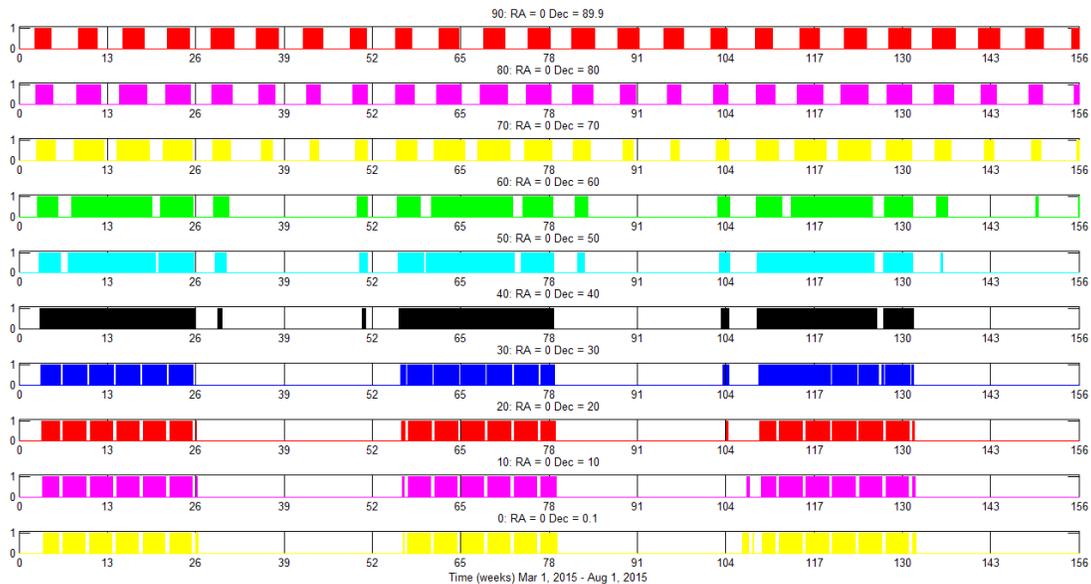


Figure 5-9: Coverage timeline for a single satellite at 28 degrees inclination viewing stars along the celestial prime meridian from the equator to the North Pole.

The star at the top of the timeline in Figure 5-9 is Polaris. A satellite in an orbit with an inclination equal to 28 degrees will be able to see Polaris most easily and for the longest time during its orbit night when the ascending node is at the Earth's sunset line (meaning the entire night half of its orbit is spent above the equator). However, the satellite does not stay in this position; at 28 degrees, a full period of nodal precession lasts 53 days, or approximately 8 weeks. The gaps in coverage correspond to the times when precession has moved the half of the orbit that is north of the equator onto the daylight side of Earth, and the satellite's view of Polaris during orbit night is obstructed by the Earth.

These coverage gaps show precisely why **shared sampling** is critical to this mission. A simple graphic of how satellites in multiple planes can share the responsibilities of observing targets throughout the year is shown in Figure 5-10. At one point, a satellite in one plane can observe the star, while satellites in other planes cannot because the Sun is too close to the target or Earth is obstructing their view for the entire orbit. However, as the orbits precess, the satellite whose plane was aligned to observe the star best is now incapable of observing the star. The satellites in other planes are now positioned to make these observations instead. During the course of the year, these satellite will continuously trade off responsibilities for observing that particular star (video demonstrations generated in STK that show this tradeoff were created to show both [short-period](#) and [long-period](#) trades [303], [304].)

For a concrete example, the coverage timeline for a *single* star (Dec = 68, RA = 0) by satellites in a Walker constellation with six evenly spaced orbital planes (the difference in RAAN is 60 degrees between each plane) is shown in Figure 5-11. No individual satellite (or set of satellites in the same orbital plane) can observe this particular star uninterrupted over the course of three years, as demonstrated by the broken coverage in the first six lines of the timeline. However, *at least one satellite* in the constellation *could* observe the star at any point in time, which is shown on the final line of the timeline. As the orbits of the satellites precess (together, at the same rate, because they are all at the same inclination), they can hand off responsibilities for observing that particular star and observe the next star that another satellite or set of satellites had been observing previously.

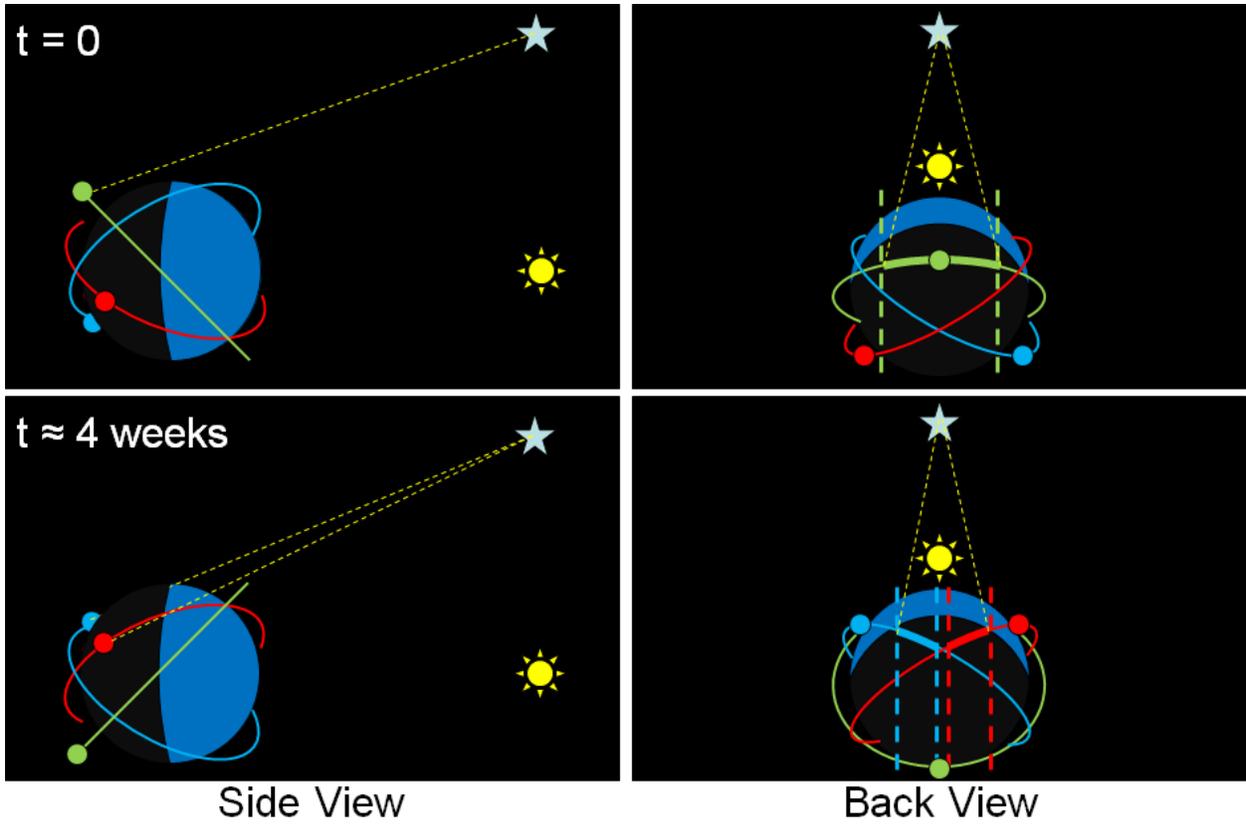


Figure 5-10: Illustration of how orbital precession prevents a single satellite from observing a target year-round, and how shared sampling can be leveraging to

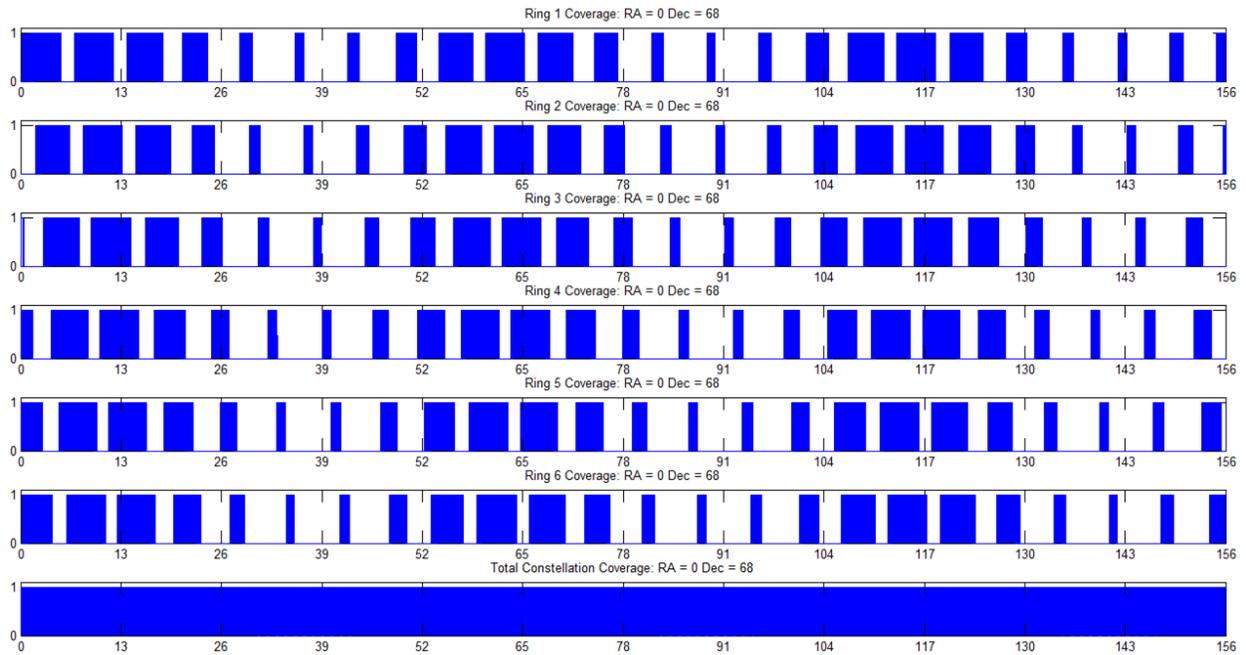


Figure 5-11: Example coverage timeline of a particular test star (RA = 0, Dec = 68) by a 6-plane Walker constellation of ExoplanetSat assets.

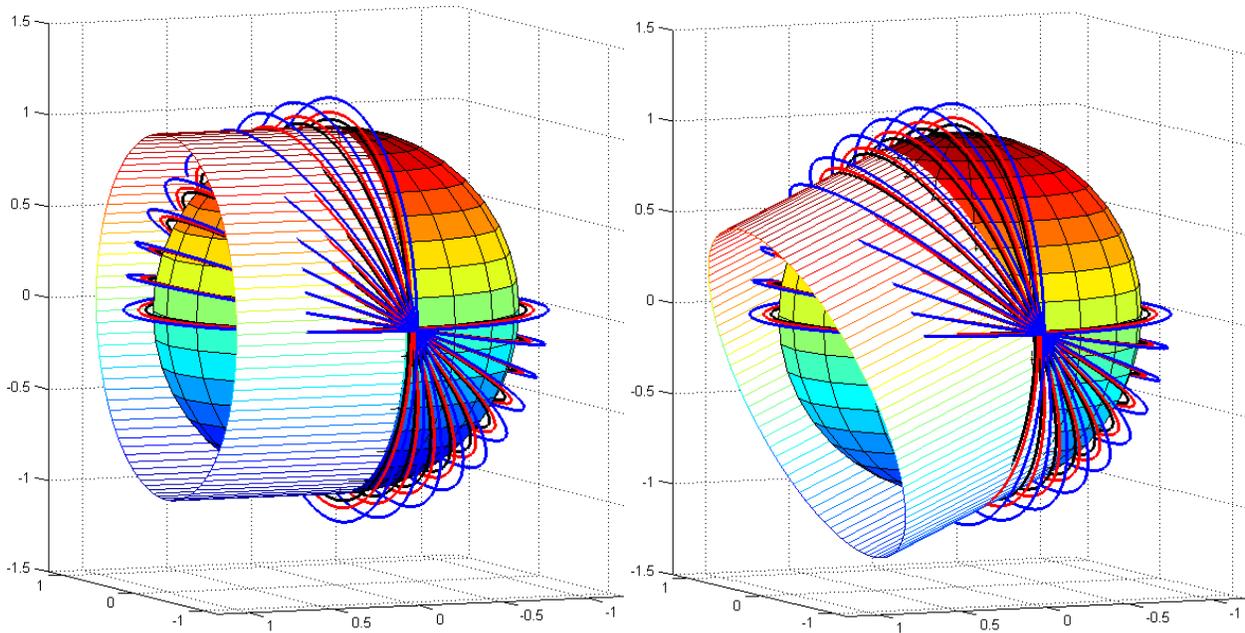


Figure 5-12: Examples of how the path length within Earth's umbra and the length of orbit night changes with orbital inclination and altitude during (left) the Summer Solstice and (right) either seasonal equinox. (Model code originally written by Niraj Inamdar)

In the second phase (Figure 5-8(2)), more stars were tested with varying RA. The coverage patterns were similar for low inclinations; the only differences were that the 6-month periods of coverage changed by 3 months for every 90 degrees added to RA. This makes sense because these stars are visible at different times of the year and changes in RA. However, orbits at the same inclination of the International Space Station (ISS) experienced coverage gaps during periods when the patterns from previous coverage timelines would predict that there should be coverage.

A simple MATLAB model was constructed to understand this phenomenon. A variety of orbits were plotted with variations in altitude and inclination with the RAAN at the Earth sunset line. An approximation of Earth's umbra modeled as a cylinder was also plotted to represent orbit night. Orbits above a certain inclination would not have an orbit night during this time, which is what happens in the sun synchronous terminator orbit; however, what was not considered was how the angle of Earth's umbra changes over the course of the year due to seasonal variations. The 23.5 degree tilt of Earth's rotational axis causes orbits that are completely within Earth's umbra at one part of the year to be completely sunlit during another. The difference between the position of Earth's umbra during a seasonal equinox and during the Summer Solstice is shown in Figure 5-12.

Since coverage of a target can only occur either during orbit night or when the Sun is greater than 90 degrees away from the target, constellations with orbits at high inclinations observing stars at RA near 90 degrees above the equator experienced gaps near the Summer Solstice because they were no longer in the shadow of the Earth. This explained the unusual gaps that were seen in the coverage maps for stars located along the line RA = 90. This phenomenon is also present for stars below the equator at RA = 270 during the Winter Solstice.

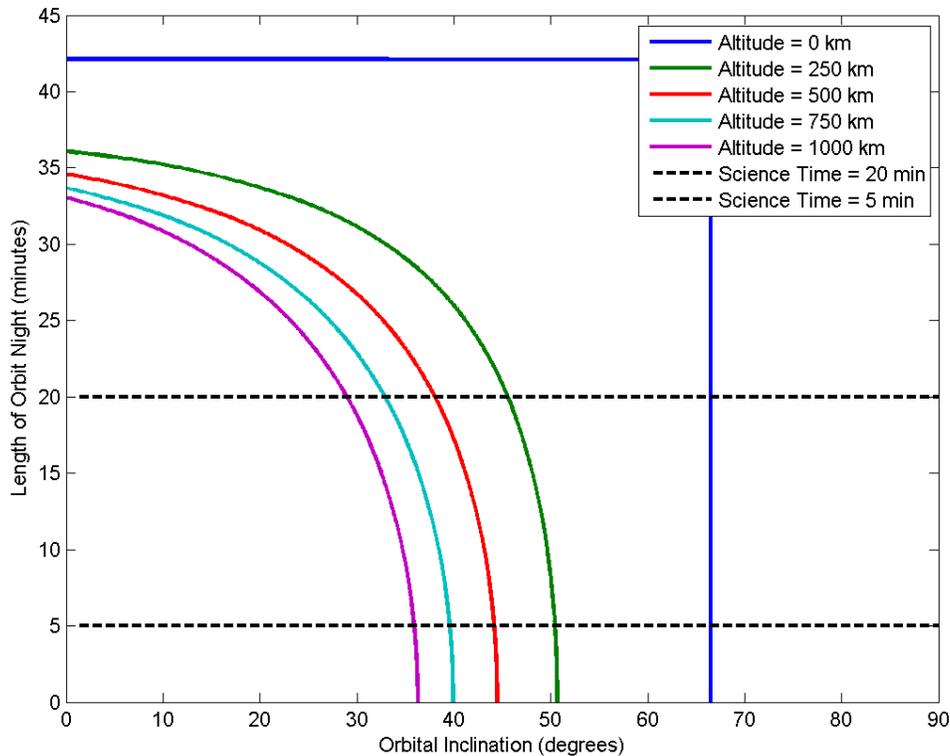


Figure 5-13: Graph showing how changes in orbital altitude and inclination affect the length of orbit night when the line of nodes is perpendicular to the Earth-Sun line during the Summer Solstice. (Model code originally written by Niraj Inamdar)

In addition to the visual representation, the length of time in orbit night was also calculated for a small tradespace of inclinations and altitudes. How the length of orbit night of an orbit with its line of nodes positioned perpendicular to the Earth-Sun line varies as a function of altitude and inclination at a solstice is shown in Figure 5-13. At sea level (where altitude is equal to 0 km), the edge of the Arctic Circle (latitude 63.5°N) is barely lit at midnight, so this serves as a convenient way to validate the model. As both altitude and inclination go up, the length of orbit night shrinks.

When the combination of altitude and inclination is high enough, the time in the umbra sinks beneath the minimum observation time of 20 minutes. For an orbit at an altitude of 500 km, the highest inclination that will experience an orbit night that is at least 20 minutes long is just above 38 degrees. This is defined as the “critical inclination” and is a function of altitude. Although it limits the available targets that can be observed without interruption over a period of years, it is not a death sentence to any constellation design with inclination above 38 degrees, but at this stage of the modeling process it was thought to be.

After identifying previously unknown variations in coverage due to changes in RA, a higher-resolution sampling of the celestial sphere was needed to understand how coverage changed with the design variables and star locations. By this point, stakeholder value modeling had determined that a useful metric for measuring how well a constellation performed compared to others was the total number of target stars that could be covered or the area of the sky that could be covered. For this reason, the test points shown in Figure 5-8(3) were tested.

Additionally, there was still major ambiguity on whether or not multiple satellites could combine observations within a single orbit. Even if a single satellite could not observe a target for 20 minutes per orbit, it may still be possible to observe it for a *combined* 20 minutes per orbit with multiple satellites sharing responsibility for a target over a period of an orbit rather than a period of precession from one plane to the next. Stars were already being tested for coverage in both 20- and 7-minute intervals; the model was then updated to record coverage by 1-minute resolution steps between 4 and 23 minutes.

In order to drastically reduce the required computation time to calculate coverages for the test stars, several key assumptions were made:

- A star that is completely covered over a one-year period is also covered over a three-year period.
- The southern celestial hemisphere's coverage map will be the same as the northern hemisphere's, only rotated about the axis of the Vernal Equinox (0-180 degrees RA)

The first assumption was made because there appeared to be very little variation in the year-to-year pattern on a single three-year timeline for any star. If there was a gap in coverage in the first year, it would also be present in the second and third years. This assumption may incorrectly lead to targets being labeled as covered when in fact they are not covered past the first year. It is possible, due to precession and the initial starting RAAN of the orbits, that stars covered by a constellation with fewer planes would line up well during one equinox but poorly during the next. However, after testing several cases and seeing no exceptions in the coverage timelines, all coverage calculations were shortened to one year instead of three.

The second assumption was made because stars tested in the southern hemisphere had similar coverage characteristics as their counterparts on the opposite side of the celestial sphere. While there are variations in Earth's orbit around the Sun due to eccentricity, these variations are small and would not be noticeable at the resolution on which these scans are being covered. A few constellations were checked against identical test star maps in the southern hemisphere and they were covered in the same way.

An example of the data generated from coverage calculations of a single star by a single constellation design is shown in Figure 5-14. Each bar shows the periods of the year at which that star is covered for at least x minutes per orbit by a single satellites, where x varies from 4 to 13 minutes on the left and 14 to 23 minutes on the right. For this particular star being observed by this particular constellation, we see that it is observable for the entire year for at least 12 minutes per orbit. There is a tiny gap in coverage if the requirement is 13 minutes, which can be seen in the last bar on the left side. As the required observation time per orbit continues to rise, that gap grows bigger and additional gaps become present.

After calculating the coverage of every test star in the celestial hemisphere by a constellation and recording the time data shown in Figure 5-14, coverage maps that show which areas of the sky can be covered for how long can be produced. For each image, MATLAB scans the timeline and searches for gaps. If there is complete coverage, it is marked as a "good" star and colored black. If the coverage is less than 100% but greater than 90%, it is marked as "almost" for qualitative purposes and colored blue. Any star covered for less than 90% of the year is marked "bad" and colored red.

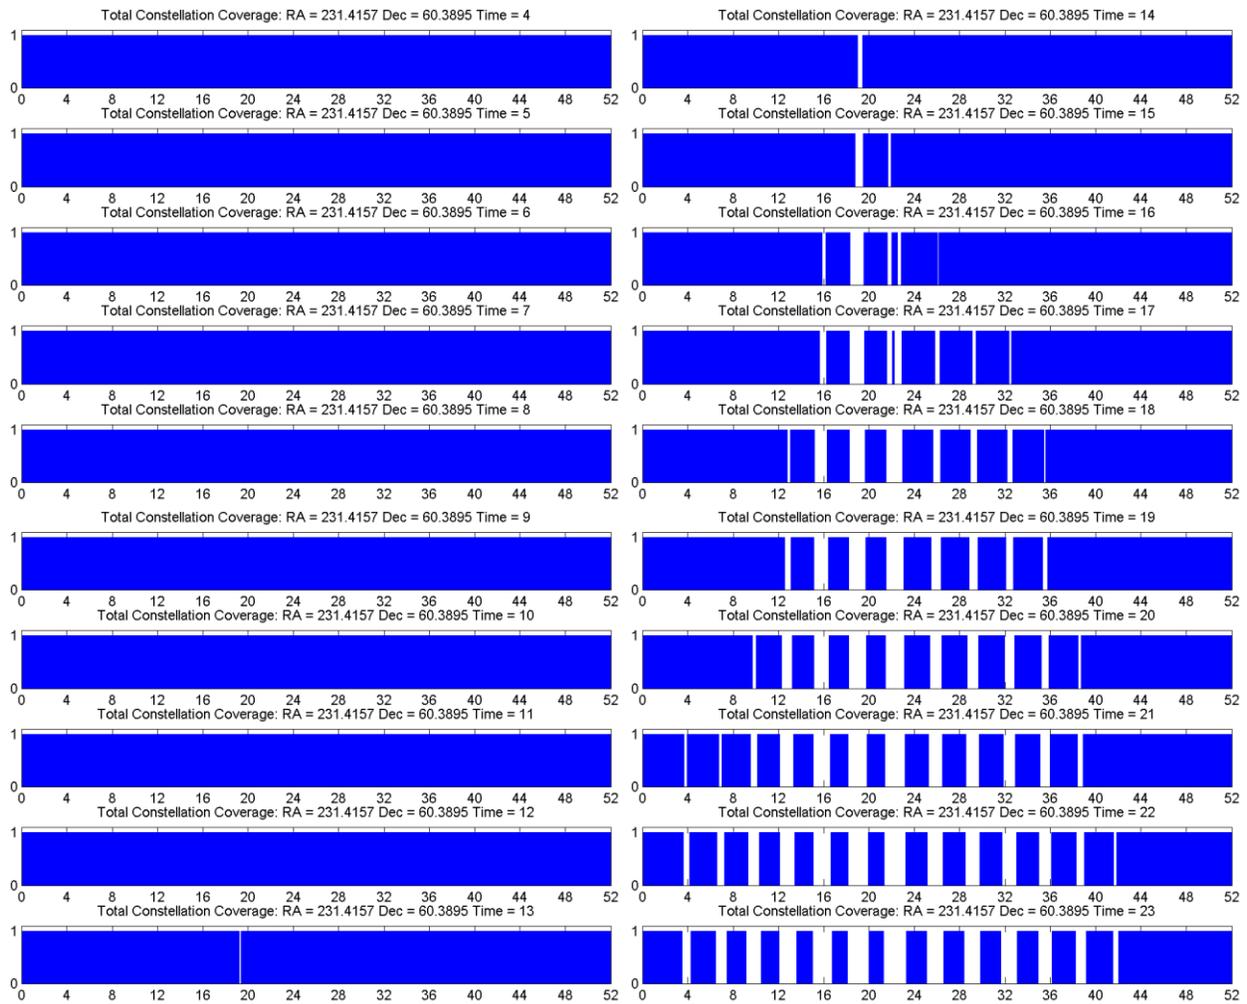


Figure 5-14: Coverage of a single star (Dec 60, RA 231) by one unique constellation design (3 planes, 52 degrees inclination, 500 km altitude) separated out by minimum length of coverage per orbit by a single satellite. (Left) Coverage from 4 to 13 minutes per orbit. (Right) Coverage from 14 to 23 minutes per orbit.

Once a coverage map is complete with the test stars, a list of actual target stars and their positions on the celestial sphere are loaded. A technique called support vector machines (SVM) is used to identify which target stars are coverable by the constellation based on the locations of the test stars marked “good” compared to the ones marked “almost” or “bad”. For constellations with a high number of planes, the area that could be covered extended beyond the limits of the test stars in Figure 5-8(3), and in some cases the resolution was not high enough to differentiate between designs and minimum coverage times per orbit. The highest resolution map of test stars in Figure 5-8(4) was then created from a gridded sphere of 2,562 evenly spaced points around the celestial sphere [305].

Coverage maps for stars that can be covered for a minimum of 10 minutes by a Walker constellation made of 3 planes inclined at 30 degrees at an altitude of 500 km are shown in Figure 5-15. The left side shows the raw data, where black points are test stars that can be covered year round, blue points are test stars that can be covered for at least 90% of the year, large red dots are stars that cannot be covered for at least 90% of the year, and small red dots are points that were not tested in this particular constellation because we know from experience that they would not be covered.

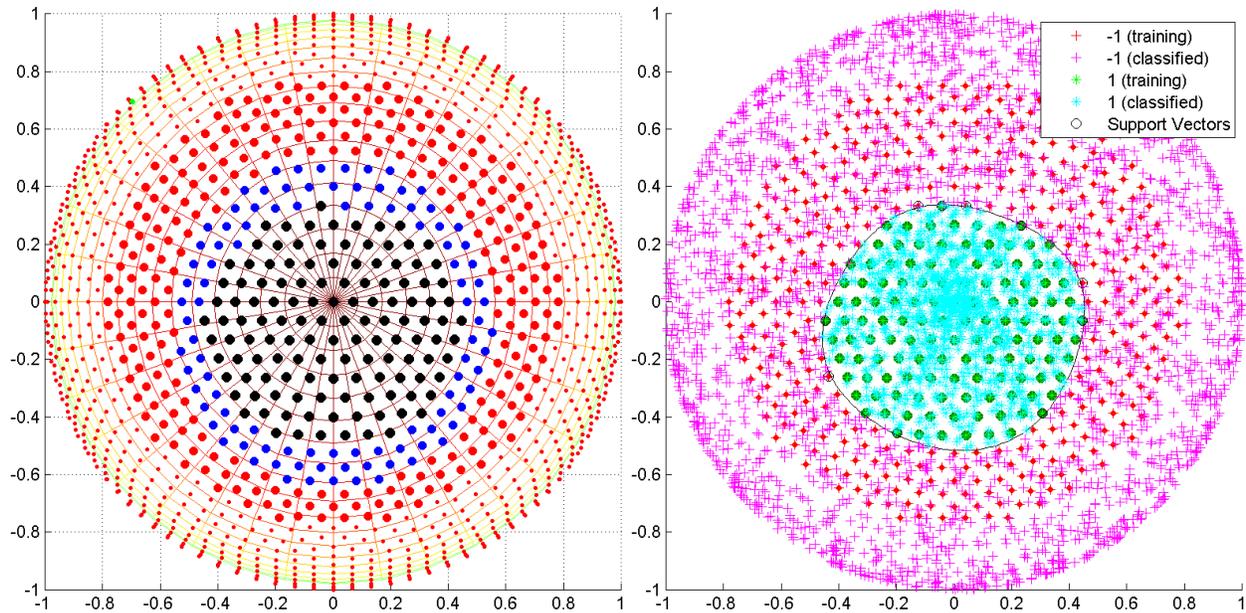


Figure 5-15: (Left) Coverage map of a particular constellation over the northern celestial hemisphere. Black dots are covered the entire year, blue dots are covered at least 90% of the year, red dots are covered for less than that. (Right) SVM modeling that defines the area covered and counts the number of actual target stars within the covered region. In these figures, the Vernal Equinox (RA = 0 degrees) is to the right, and during the Summer Solstice, the Sun is at the top (RA = 90 degrees)

The right side shows the 100% covered test stars in green, while the actual stars that could potentially be targets for this particular ExoplanetSat constellation design are shown in cyan, while the ones outside the coverage range are shown in purple. This includes stars in the southern hemisphere that were rotated about the axis of the Vernal Equinox so that stars in both hemispheres that would be fully covered by the constellation would also be counted.

There are many important observations to consider from the resulting coverage maps. The first and most obvious conclusion is that the strict pointing and exclusion angle requirements of ExoplanetSat make it impossible for the entire sky to be covered year round. This was already known because the Sun is always on the ecliptic plane, and therefore any stars behind it cannot be observed even with relaxed solar exclusion angle constraints.

However, these maps disproved an assumption made in the early modeling phases that the covered area would be always symmetric about either the North Pole or the imaginary axis about which the Sun would rotate if the Sun rotated around the Earth (an imaginary line passing from the points Dec 66.5 RA 270 to Dec -66.5 RA 90). Low-inclination constellations appear to exhibit symmetry about the rotational pole of the Earth, but as inclination is raised above the critical inclination, not only does the apparent center move down in declination along the line RA = 270, the shape of the area that is covered mutates as well.

It was also assumed that the area of the sky that could be covered would shrink as the required coverage time grew. This is true, but the nature of how the area shrunk was dependent on the inclination of the constellation. Once again, low inclination constellations' coverage area shrunk rather symmetrically with increasing time, but high-inclination constellations' coverage area started out as nearly circular areas and

bended into a shape reminiscent of the Rosenbrock function (or a banana, in layman’s terms) on the medium-resolution test map. Examples of different coverage areas for constellations that only differ by changes in inclination are shown in Figure 5-16.

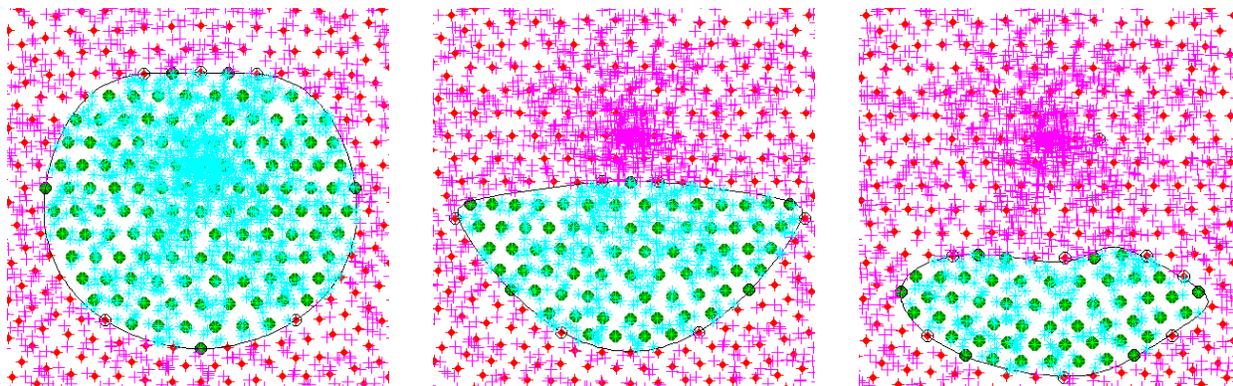


Figure 5-16: Coverage maps for constellations that only differ in inclination. All constellations are 6 planes at 500 km altitude observing targets for a minimum of 20 minutes per orbit. From left to right, orbital inclinations are 38, 44, and 52 degrees.

Even though Figure 5-12 shows that constellations above the critical inclination do not experience an orbit night near the time of the solstice if their line of nodes is perpendicular to the Earth-Sun line, they can still cover targets for at least 20 minutes per orbit when they are greater than 90 degrees away from the Sun. This is why the most northern star along the line RA = 270 that a constellation can cover for longer periods of time appears to converge near Dec = 66.5 degrees as its inclination is raised, which is exactly 90 degrees from the Sun during the Summer Solstice.

Although the constellation design that Team X suggested, a Sun-synchronous orbit along the terminator (RAAN 270 degrees), has no stars that are observable year-round for any length of time at all according to its coverage maps, other Sun-synchronous orbits at differing RAANs had relatively large coverage areas for low observation times. This is why these options were also included in the tradespace. The RAAN with the largest area of coverage is RAAN = 180, but because the target stars are not evenly distributed about the celestial sphere, other options were considered in the design vector (see Table 5-4).

Another observation that came from analyzing the area of coverage as a function of observation time was that in general, constellations below the critical inclination had smaller coverage areas than one above it at low observation times per orbit. However, higher inclinations were more sensitive to changes in observation time, and the coverage area decreased faster as observation time increased in comparison to lower inclinations. This is especially true for Sun-synchronous constellations, because although they can cover large areas for short observation periods, no Sun-synchronous orbit can observe any star for more than 17 minutes per orbit for an entire year (and the coverage area appears to shrink to a point near Dec 66.5, RA 270 for a Sun-synchronous orbit at RAAN = 180).

Thanks to SVM, the number of stars within the area of coverage n_{obs} could quickly be counted. Since the final stakeholder interview showed how critical the number of possible targets a constellation can observe is to the calculation of the expected number of exoplanets that could be found, this was the only metric that was necessary from this portion of the performance modeling and simulation.

However, because it is impossible to know yet how many targets will be candidates for long-term observations and what those targets' positions will be, it is impossible to know precisely how many targets will be able to fit onto the science viewing area of the detector on an individual ExoplanetSat (see Figure 5-2). For this reason, the number of targets viewable per satellite is being treated as a context variable (see Section 5.3.1.3). The possible values for this context variable have been estimated based off of the average density of candidate stars in the catalog over the entire celestial sphere multiplied by the probability of selecting a star for long-term observation, but some areas of the sky have higher densities than others, and there may be intelligent ways to assign observation duties to maximize this on a case-by-case basis that is too detailed for tradespace exploration and is therefore outside the scope of this project at this phase.

Lifetime Module

In addition to the area of the sky that could be covered, the lifetime of a constellation also affects the value proposition for this science mission. However, there was very little to model here. STK can estimate the lifetime of a satellite given its mass, ballistic frontal area, and coefficients of drag. Although the mass and frontal areas of the satellite are known, high estimates on the coefficients of drag were used for worst-case scenario calculations.

Atmospheric decay time is usually modeled mathematically as an exponential function with respect to altitude. In a pessimistic scenario, the lifetime of the satellite at 300 km is 13 days; for 400 km, it is 4 months; and for 500 km, it jumps to almost 4 years. Beyond that altitude, the orbital lifetime begins to stretch well beyond the expected operational lifetime of any CubeSat ever built in addition to the planned operational lifetime of the mission and into the order of human lifespans.

More optimistic calculations assume lower coefficients of drag and variations in the frontal area exposed to atmospheric deceleration forces since the spacecraft will not always be pointed with its solar panels directly into the wind. These optimistic estimates were confirmed with STK and show that any altitude choices below 500 km would risk cutting the mission lifetime short, and any choice above 500km risks being a space debris hazard that would need to be tracked for generations.

Additionally, because the coverable area of the sky is relatively insensitive to variations in altitude (but only shrinks with increasing altitude), there appears to be no point in going beyond this altitude. Originally, constellations at 700km altitude were tested along with 500 km, but once the resolution of the test star map went up, the time required to complete a coverage map was too high to justify the computational expense given these factors. For this reason, "Orbital Altitude" was removed from the tradespace as a design variable and simply set to a parameter at 500 km.

5.3.2.3 Performance Modeling Process 3: Cost Modeling

While other case studies will use some approximation of cost, such as total mass, adjusted mass, or complex mass, as the cost metric by which the utility is compared to in MATE, ExoplanetSat had the benefit of being analyzed in a concurrent engineering design studio by Team X. While they did not provide exact details on their cost estimates, nor do they publish anything regarding their proprietary cost models, they did provide two cost estimates for two different designs that were broken down into subcomponents. From this information, a reverse-engineered cost model was constructed that used real estimates for costs of labor, management, payload testing, engineering models, launch, and operations.

Team X Cost Model

One of the major faults of the Team X cost model, and other space systems cost estimating tools, is that they assume that cost line items that are typically called “wrap costs” scale the same as other labor costs. Not only are the labor costs multiplied as a flat rate after the first unit rather than with a cost improvement function, they apply equally to the wrap factors. Because they did not break the cost down strictly into recurring and non-recurring costs, however, it is difficult to apply a cost improvement curve to these numbers accurately. Additionally, the element that *does* attempt to follow a learning curve, hardware costs, only takes stepwise increments of 0.95 times the cost of the first unit for units 2 through 5, and 0.85 for additional units. In reality, this incremental change would be constant and decrease with additional units. Finally, the launch cost model is somewhat optimistic compared to the estimates provided by SpaceFlight Services (SPS). The constant scaled costs are summarized in Table 5-6.

In addition to scaling costs, there were fixed costs in the estimate. These are shown in Table 5-7. The remaining line items, hardware costs and mission operations/ground support follow different models. Hardware follows the 0.95/0.85 steps, and ground support is fixed based on the cost of the Near Earth Network (NEN) for communications per satellite per year. These are shown in Table 5-8. Team X estimated that a mission lasting one year with eight operational satellites and three engineering units (no payload, only hardware costs) with 25% reserves for margin would cost a total of \$28,681,910.

Table 5-6: Team X’s ExoplanetSat cost model line items with multiplying factors for additional units.

Exoplanet Scaled Cost Factors	Cost for First Unit (\$k)	Additional Units (\$k)	Cost Factor
Project Management	\$579.33	\$115.87	0.2
Project Systems Engineering	\$333.71	\$66.74	0.2
Mission Assurance	\$46.17	\$9.23	0.2
Science and Technology	\$200.00	\$40.00	0.2
Payload	\$2,588.04	\$854.05	0.33
Spacecraft Labor	\$1,383.08	\$276.62	0.2
Launch Vehicle Costs	\$331.73	\$199.04	0.6
ATLO	\$146.74	\$48.42	0.33

Table 5-7: Fixed costs in Team X's ExoplanetSat cost model.

ExoplanetSat Fixed Cost Factors	Cost (\$k)
Software Development	\$462.10
Testing Platform	\$126.02
Education and Outreach	\$120.00

Table 5-8: Costs for hardware and ground support operations, which do not scale the same way as other cost elements listed in the previous tables.

Other ExoplanetSat Costs	Cost for First Unit (\$k)	Cost for Units 2-5	Cost for Units >5
Hardware	\$230.70	\$219.16	\$196.09
MOS and GDS	\$384.36 (/satellite-year)	--	--

Although their estimates on scalability due to mass production and launch costs may be inaccurate, the uncertainty, variation, or use of a completely different estimation schedule can be treated as variations in the context variables. Additionally, the JPL cost estimates for CubeSat launch costs were slightly optimistic compared to estimates obtained from launch providers. However, there are no official estimates yet on costs associated with delivery of satellites to multiple planes. For now, this cost model assumes

that all launches are equal price, and that payloads can be delivered to the proper True Anomaly within their respective orbits for the cost included in the launch. The total cost per satellite depending on the operational lifetime is summarized in Table 5-9.

Table 5-9: Summary of total costs per satellite in the ExoplanetSat constellation.

Total Costs	Cost for First Unit (\$k)	Cost for Units 2-5	Cost for Units >5
Fixed Costs	\$708.12	--	--
Hardware Development	\$5,839.50	\$1,829.10	\$1,806.10
Operations (per year)	\$384.36	\$384.36	\$384.36
Total Per Unit Cost (1 year)	\$6,932.00	\$2,213.50	\$2,190.50
Total Per Unit Cost (4 year)	\$8,085.10	\$3,366.50	\$3,343.50

Custom Cost Models

Two custom cost models were also constructed based on the Team X cost model. All three models are exactly the same at the unit cost levels for all the components except for the launch costs, which are based on estimates from launch providers. The differences between the custom models and the Team X models are in the cost factor discounts for each additional unit compared to the values shown in Table 5-6.

The wrap costs associated with project management, systems engineering, mission assurance, and science and technology have been reduced to 0.05 times the first unit cost. Payload development follows a tiered schedule similar to the Team X hardware costs that is a factor of 0.33 for units 2-5 and 0.2 for additional units. Spacecraft labor, ATLO, MOS and GDS have the same cost factors, and all fixed costs remain the same. Additionally, because hardware prices can be estimated using actual cost improvement curves (mathematically defined Section 2.4.3.5), these will be used to calculate the hardware costs. These changes between the Team X and Custom Cost Models are summarized in Table 5-10 and Table 5-11.

Table 5-10: Custom Cost Models #1 and #2 line items with constant multiplying factors for more units.

Exoplanet Scaled Cost Factors	Cost for First Unit (\$k)	Cost Factor
Project Management	\$579.33	0.1
Project Systems Engineering	\$333.71	0.1
Mission Assurance	\$46.17	0.1
Science and Technology	\$200.00	0.1
Hardware	\$230.70	S = 0.85
Spacecraft Labor	\$1,383.08	0.2
ATLO	\$146.74	0.33

Table 5-11: Custom Cost Models #1 and #2 for payload and ground support operations.

Other Costs	Cost for First Unit (\$k)	Cost Factor for Units 2-5	Cost Factor for Units >5
Payload	\$2,588.04	0.33	0.2
MOS and GDS	\$384.36 (/satellite-year)	--	--

For modeling launch costs, launch service providers were contacted for more information. SFS provided launch cost estimates (not quotes) for bulk launch purchases to different inclinations. Unfortunately, these prices are proprietary and cannot be published in a table here, but these estimates will be used in the custom cost models to determine the launch price. Additionally, because not all the inclinations in the

design space were quoted, the costs for the missing inclinations were linearly extrapolated from the data that was provided.

Unfortunately, the SFS launch cost estimates do not take into account transport of satellites to multiple planes. While there are a number of options that may be available in the near future for launching ExoplanetSat assets into multiple planes, there are no solid cost estimates for how much these services will cost. To make up for this difference and realistically penalize constellations with multiple planes, Custom Cost Model #2 includes a 10% cost increase on launch prices for each additional plane. For instance, if the launch cost per satellite in Cost Model #1 is X , a constellation of four satellites in four planes would cost $4X$ in Cost Model #1, but it would cost $1X + 1.1X + 1.2X + 1.3X = 4.6X$ in Cost Model #2.

5.3.2.4 Performance Model Limitations

There are a number of ways that the performance model could give incorrect estimations on the performance of ExoplanetSat. The coverage model is limited by the resolution of the star maps used. Doubling the resolution of the star map would do little to provide additional information at the TSE level and would double the length of time it took to calculate the coverages, and some uncertainty in the exact number of stars that can be covered year-round is acceptable.

Additionally, the accuracy of the coverage calculation is limited to the resolution of the images showing the coverage timelines such as the ones in Figure 5-14. If a star loses coverage for less time that is represented by the width of one pixel on the coverage timeline, it will be accidentally counted as a star that is covered for 100% of the year. Since the width of these coverage timelines are 928 pixels, a star could theoretically be uncovered for a period of time represented by 1.5 pixels, or $1.5 * 52 \text{ weeks} / 928 \text{ pixels} * 7 \text{ days/week} = 0.58 \text{ days}$ but still be listed as 100% covered.

This limitation can be fixed by storing the actual data files instead of images of graphs, but that requires far more memory for each coverage calculation and the gap times are much more difficult to parse out. The limitation was ignored because of how inconsequential it would be to lose coverage for approximately the length of a transit exactly once per year, which is the worst case scenario. Even so, if the star does not have coverage for the number of required minutes per orbit, then it would still have coverage for one minute less per orbit, meaning the SNR drop would be insignificant.

Another limitation is the accuracy of the SVM counting procedures. A vector that contains all of the covered test stars and none of the uncovered test stars may miscount the number of actual target stars. The error increases and the area of the sky that is covered decreases, but that only means undesirable designs would have significant error due to this effect.

Finally, the launch cost models may be inaccurate for multiple planes, and the lower inclinations (the critical inclination and below) may be unreachable for any piggy-backed payload, so a more expensive, dedicated launch may be necessary. There are other orbit delivery options, such as the SHERPA vehicle, Super Strypi, and Firefly small launch vehicles, but cost metrics have not been released for these options because they are so new. Additionally, using orbital precession alone to change RAAN could increase the radiation dose significantly while the satellites are in the transfer orbit. This could affect the operational lifetime of the satellites. One final limitation to these cost models is that they do not take into account the operational complexity associated with multiple planes.

5.3.3 ExoplanetSat RSC Phase 3: Tradespace Exploration

When the value model and the performance models are brought together, tradespace exploration can begin. Before jumping directly to the utility versus cost curve, it is useful to examine the attribute versus cost curve first. The tradespace in Figure 5-17 shows all 4,000 designs plotted versus the expected number of exoplanets detected EN color-coded by inclination. A similar tradespace is shown in Figure 5-18 but with a different metric on the y-axis, the probability that at least one exoplanet will be detected.

The first observation to point out is that the Pareto front is linear near the left side of this diagram. This is due to two reasons. First, cost scales directly with the number of satellites, so the total number of satellites is the only factor that affects the cost axis of the tradespace. If more information on a cost model included information on how costs could change with the other design variables (e.g. more planes requires more complex asset delivery operations but the price per satellite delivered to a plane goes down with more satellites per plane), then this would not be the case.

Second, EN scales with n_i , which also scales directly with the total number of satellites up until the point where a constellation is observing all possible targets within the observable area of the sky. This “saturation point” happens when $n_i = n_{obs} \sin(\varphi)$ becomes less than the number of satellites times the number of targets a single satellite can observe simultaneously. This relates to the second observation, that designs with a range of inclinations are on the Pareto front near the left of the diagram, but only designs at the critical inclination $i = 38$ dominate after this point. That is because these constellations are the ones that have the potential to cover the largest area of the sky, and thus are able to observe more targets before becoming saturated.

The tradespace plotted against the SAU is shown in Figure 5-19 through Figure 5-21. Notice that no design that costs less than \$50 million is acceptable. This is because these designs don’t have enough satellites to cover enough stars for EN to be brought above 1. Also notice from Figure 5-17 and Figure 5-19 that none of the satellites in a Sun-synchronous orbit are on the Pareto front. The reason for this is better illustrated in Figure 5-20 and is due to three factors. First, the inherent noise that comes from adding orbit-by-orbit observations together decreases the probability of detection. Second, when satellites are redundantly viewing the same targets they are also decreasing the total number of targets they could be viewing, which lowers n_i . Third, the Sun-synchronous orbit is incapable of viewing any stars for more than 17 minutes at a time throughout the year, so all designs in this orbit *must* use redundancy *or* be penalized for observing for less time per orbit.

If the probability of finding an exoplanet were higher for lower observation times, i.e. the probability curves in Figure 5-7 (left) were shifted further to the left, designs in Sun-synchronous orbits would also be on the Pareto front. This could be explored with additional context variables, but the stakeholder was adamant about these already being optimistic, and so lower times with higher probabilities were not considered in this iteration of the performance model. Once more information from instrumentation and modeling is available, this issue can be revisited quite easily.

Figure 5-20 also shows that all designs on the Pareto front observe stars for 20 minutes at a time, no less, and do not share responsibilities for viewing stars on the same orbit with other satellites. This maximized n_i which in turn maximizes EN .

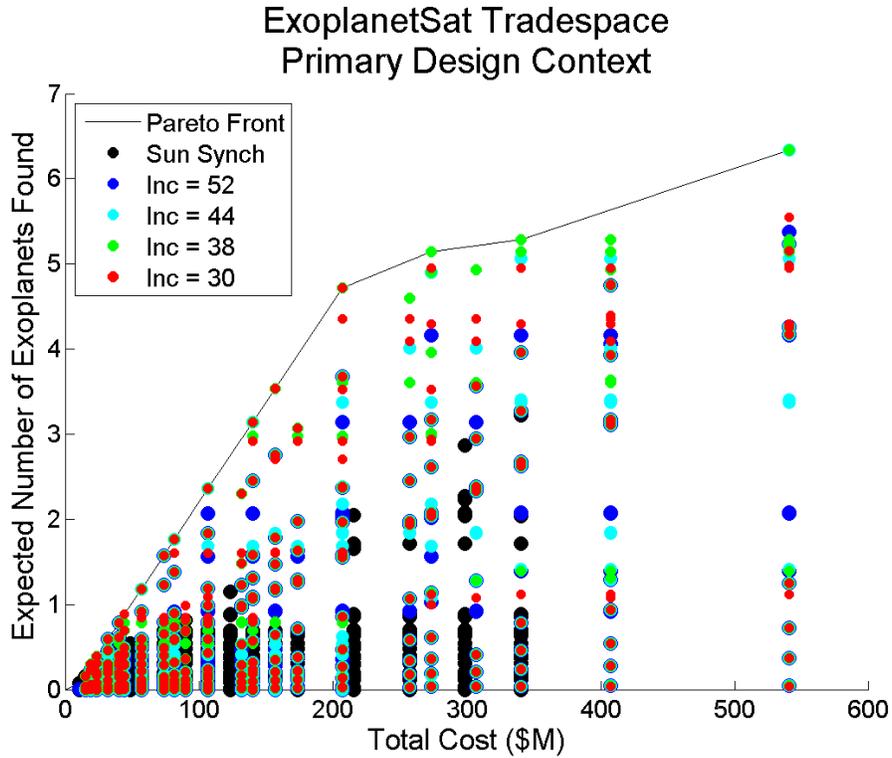


Figure 5-17: ExoplanetSat tradespace comparing total cost to Expected Number of exoplanets detected, which is the primary attribute for stakeholder satisfaction.

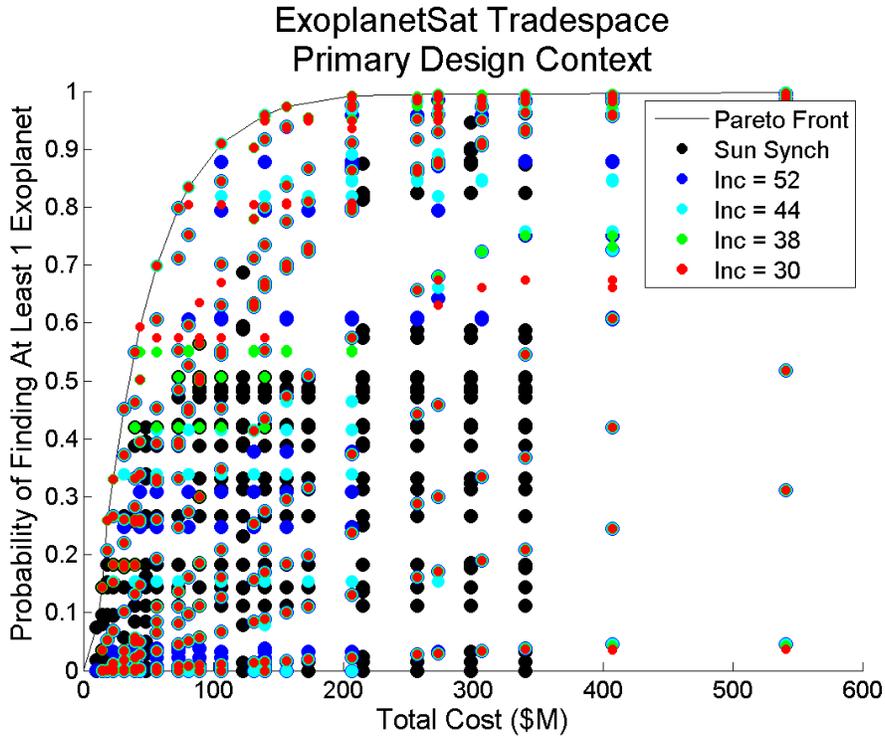


Figure 5-18: ExoplanetSat tradespace comparing total cost to probability of detecting at least 1 exoplanet. This is an alternative way for the stakeholder to view how value can be gained from the mission.

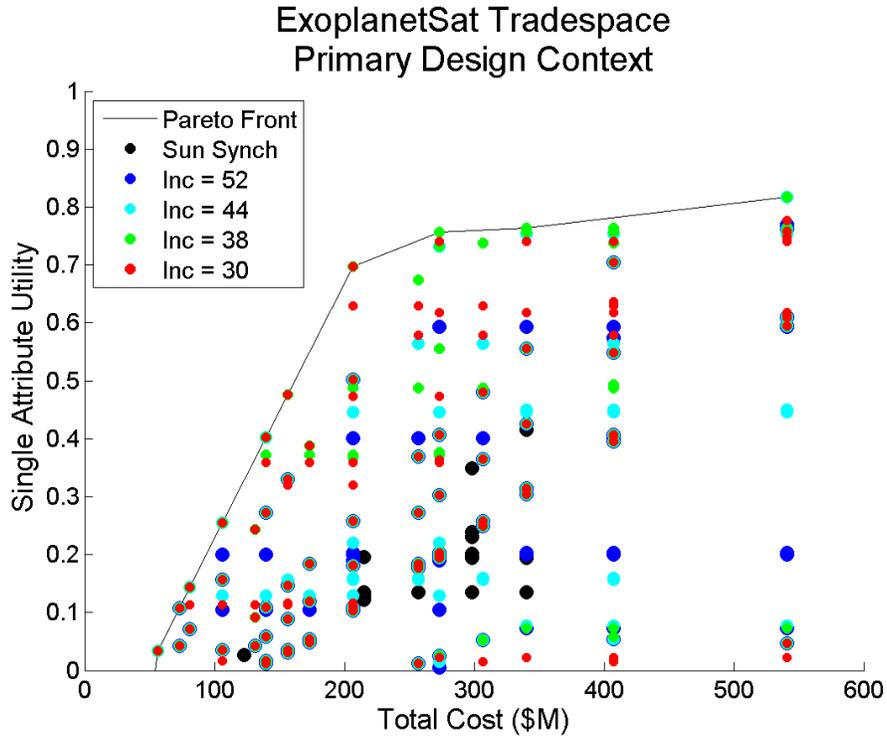


Figure 5-19: ExoplanetSat tradespace comparing total cost to single attribute utility.

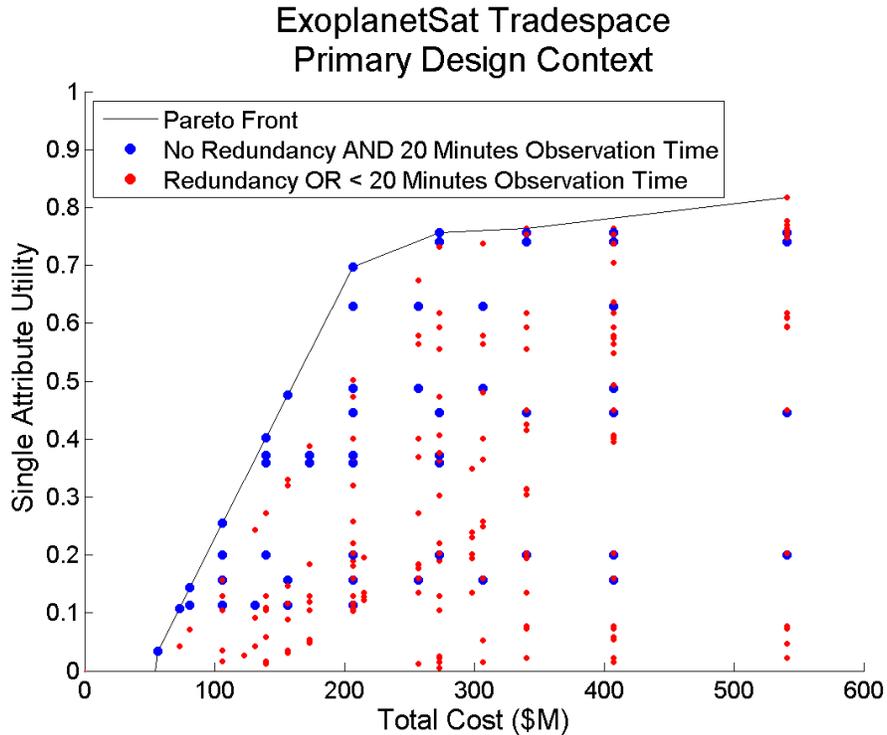


Figure 5-20: ExoplanetSat tradespace comparing total cost to single attribute utility. From this, we see that only designs that have no redundancy and observe for 20 minutes per orbit are on the Pareto front up until the point where the number of possible targets is less than the number of observable targets.

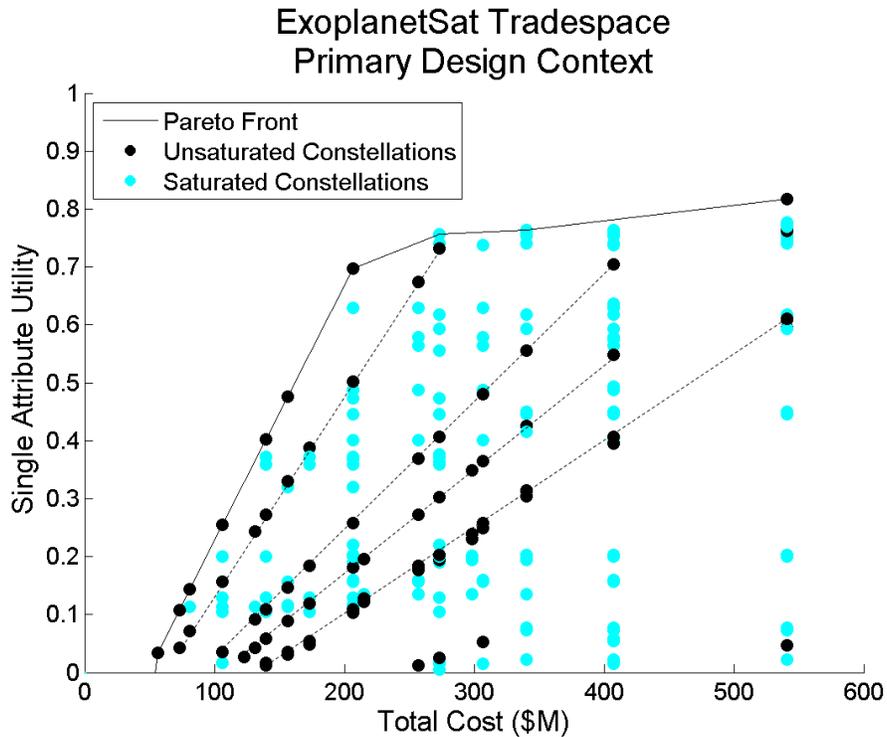


Figure 5-21: ExoplanetSat tradespace comparing total cost to single attribute utility. Constellations that are “unsaturated” meaning there are more targets that can be observed than there are satellites times targets per satellite, tend to not be on the Pareto front.

Figure 5-21 highlights the saturated versus unsaturated constellations. Here it is evident that it almost doesn’t matter which inclination is chosen if there are fewer satellites, because as long as there are more stars in the observable part of the sky than the constellation can view at once, the constellation is operating at its maximum capacity and therefore is delivering maximum utility for its cost. Inclination only matters with large constellations; a large observable area of the celestial sphere is necessary in order to capture more utility, in which case inclination choice becomes important. Additionally, if the system were to be upgraded after a certain period, the inclination cannot be changed, so a design choice with no consequence in the early stages could have consequences later in the program lifecycle.

This analysis hints at how strongly the number of targets per satellite affects the utility of the system. While this is not a design variable that can be adjusted in this tradespace because the individual satellite design parameters are not being varied in this study, it does cast light on how the system could be improved before it is implemented. If the science viewing area of the focal plane (see Figure 5-2) were able to be expanded further with little additional cost, it would clearly be valuable for the mission. The effects of this increase will be studied in EEA.

5.3.4 ExoplanetSat RSC Phase 4: Epoch-Era Analysis

While many characteristics of the design space have already been identified, including showing once again that the design that Team X proposed does not satisfy the PSS during Phase 3 of the mission, there is still more that can be learned by conducting Epoch-Era Analysis. However, EEA is rather trivial in its formal implementation on this case study for a number of reasons.

First, cost varies only as a function of the number of satellites built and the length of time that they operate. While varying cost models with different values for learning curves would alter the tradespace and add convex curvature to the Pareto front that we see in Figure 5-19, it would not cause any designs to switch places along the cost axis.

Second, there are no additional value models to consider that would cause an epoch shift. There is only one PSS, one scientific goal, and one attribute of that goal. The value model was able to be expressed so simply that very little could change it. Other case studies do not have this advantage (although it is a problem for the Ph.D. candidate trying to show how it works).

Finally, all epochs identified by unique context variables are mutually exclusive, so there is no possibility of an epoch shift, and any era generated in era analysis will consist of a single epoch. Context variables that affect cost would shift before the development process, and there cannot be two prices assigned to a design simultaneously. The price of operations and accessing the NEN may change over time, but this change is trivial and would still affect all designs equally. Context variables that affect anything besides cost are mutually exclusive; the average R_*/a of planets around the candidate stars can't take multiple values simultaneously, and it won't shift during the course of the mission lifecycle. Likewise, the average number of targets a satellite can view simultaneously will not suddenly change due to a shift.

5.3.4.1 EEA Process 1: Single-Epoch Analyses

Figure 5-22 shows how the tradespace changes if the value of R_*/a is not the same as it is for Earth. In these figures, moving from left to right, R_*/a rises from 1/215 (left) to 1/150 (center) to 1/100 (right). We see a significant change in the slope of the Pareto front as a result of higher possible values for EN . In the bottom three graphs that show the SAU, more designs, especially cheaper designs, that were once unacceptable because $EN < 1$ are now acceptable given the value model of the PSS. Additionally, with increased R_*/a , some designs are capable of completely satisfying the PSS, unlike in the primary design context shown in MATE in the previous section.

If we take a higher estimate on what R_*/a would be, the tradespace changes quite a bit, but the properties still remain similar; Pareto designs are still unsaturated and observe for 20 minutes with no redundancy because satisfaction is still dependent on how many stars can be viewed. In these cases, however, there is a higher probability that these stars will have a transiting exoplanet around them, so the ratio of the number of targets that are deemed observable for long periods to the number of stars within the area of the sky that can be covered n_i/n_{obs} is much lower.

However, there is little else to learn from changes in this context variable. The relationships between the designs and the utility remain the same as they were in the primary context. Figure 5-22 shows that even as R_*/a is increased, the designs on the Pareto front are still non-redundant, observe for 20 minutes, and are unsaturated beneath the \$200 million mark. Once again, orbital inclination is not important for small constellations besides Sun-synchronous orbits.

In the primary context, the least costly acceptable constellations cost \$56.2 million. If in reality $R_*/a = 150$ then the least costly acceptable options are \$39.5 million, and if $R_*/a = 1/100$, options that only cost \$22.7 million would still be acceptable.

Figure 5-23 shows the variations in the tradespace as the number of targets per satellite is changed. In these figures, the left side shows the pessimistic example where a satellite can only view two targets simultaneously; the center shows the primary context where that value is three; and the right side shows when it is four. The variations aren't as drastic as variation in R_*/a , but this scenario does illustrate to the designers that finding a way to increase the science area of the detector focal plane will add to the value of the mission in a nontrivial way.

The slope of the linear part of the Pareto curve changes slightly with varying target capacity, meaning if more targets can be viewed at once, cheaper constellations become acceptable. When target capacity is two, there is a kink in the Pareto front that is especially visible in the bottom left of Figure 5-23. This occurs because it is much harder to saturate a constellation; therefore an unsaturated constellation that would have dominated a saturated constellation is no longer producing as much utility and drops below the curve. This effect would not be present with better sampling of the total number of satellites. However, the unsaturated design that is on the Pareto front in this context is still non-redundant and observes targets for 20 minutes at a time.

The least expensive constellations that are acceptable when target capacity is equal to four cost \$39.5 million. If target capacity is only two, the least expensive acceptable constellations cost \$72.9 million. This change in price is huge considering how small of a difference it is in the hardware specifications. This fact alone should force the PSS to reconsider the design options in the focal plane. Such a revelation would not have been seen without implementation of EEA.

Finally, the differences between the JPL Cost Model and the two Custom Cost Models are shown in Figure 5-24. There is admittedly little variation among these three models. What the Custom Cost Models gain by assuming better discounts on repeated units, especially on the higher end of the cost spectrum, they lose in more realistic launch costs that factor in inclination and number of planes. The Pareto fronts in the center and right side of Figure 5-24 are no longer linear because of the inclination penalty associated with going to inclinations below 52 degrees, meaning the better constellations closer to the critical inclination are slightly more expensive rather than being the same price. This is evident by the nearly vertical leaps the Pareto front makes, showing that for very modest increases in cost a design achieves much better utility, so the price of going to those inclinations is clearly worth it.

Not all contexts are shown here, but it is clear that the most optimistic context vector for ExoplanetSat is [Custom Cost Model 1, $R_*/a = 1/100$, 4 Targets per Satellite] and the most pessimistic is [JPL Cost Model, $R_*/a = 1/215$, 2 Targets per Satellite].

5.3.4.2 EEA Process 2: Multi-Epoch Analysis

Multi-epoch analysis is especially useful in this case to show definitively what has been qualitatively described by single-epoch analyses, that the Pareto front is always populated by saturated constellations.

The fuzzy Pareto number was calculated for all designs in all epochs. In addition to the 27 epochs listed in Table 5-3, and additional variable on the operational lifetime was added to better show how lifetime can affect performance, making a total of 81 epochs that can be analyzed in multi-epoch analysis. Nine designs along the Pareto front in the primary design context were examined in detail, and the results are shown in Figure 5-25.

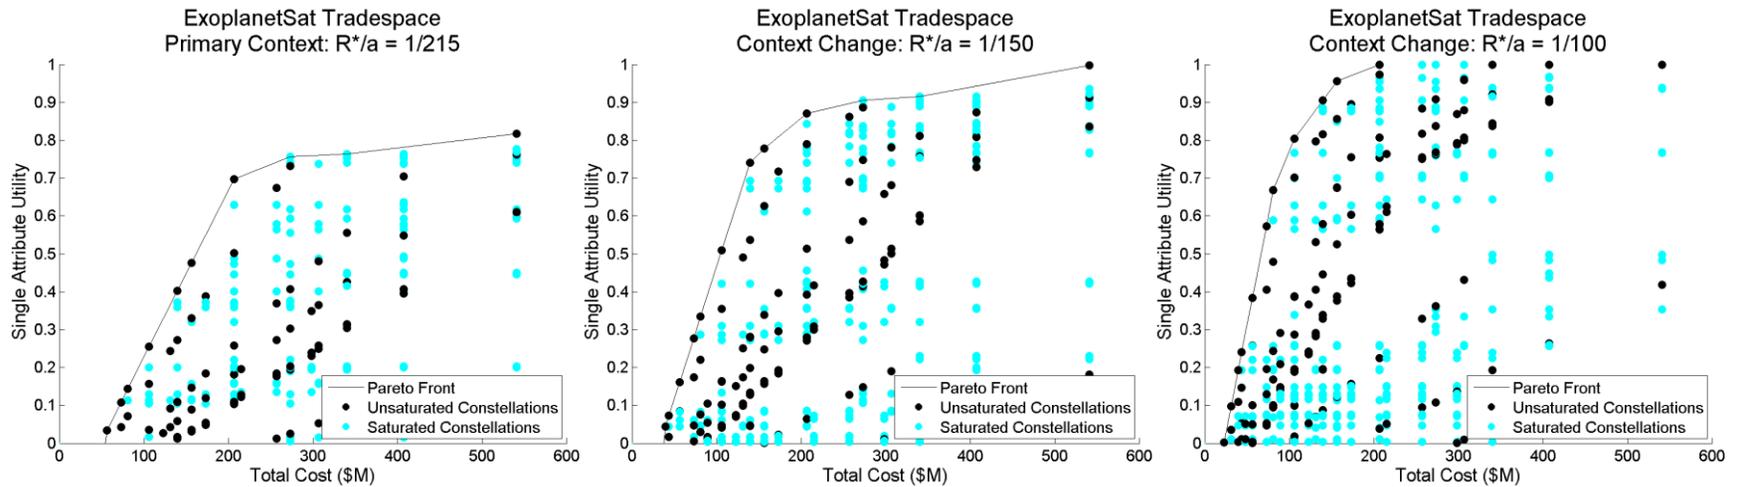


Figure 5-22: Tradespaces showing variations in the Single-Attribute Utility in contexts ranging from R^*/a equals 1/215 (left), 1/150 (center), and 1/100 (right). As the ratio gets higher, utility tends to rise for all designs, though the Pareto front does not change significantly on the lower end of the curve. Even in these contexts, only unsaturated constellations with maximum observation time and no redundancy are on the Pareto front.

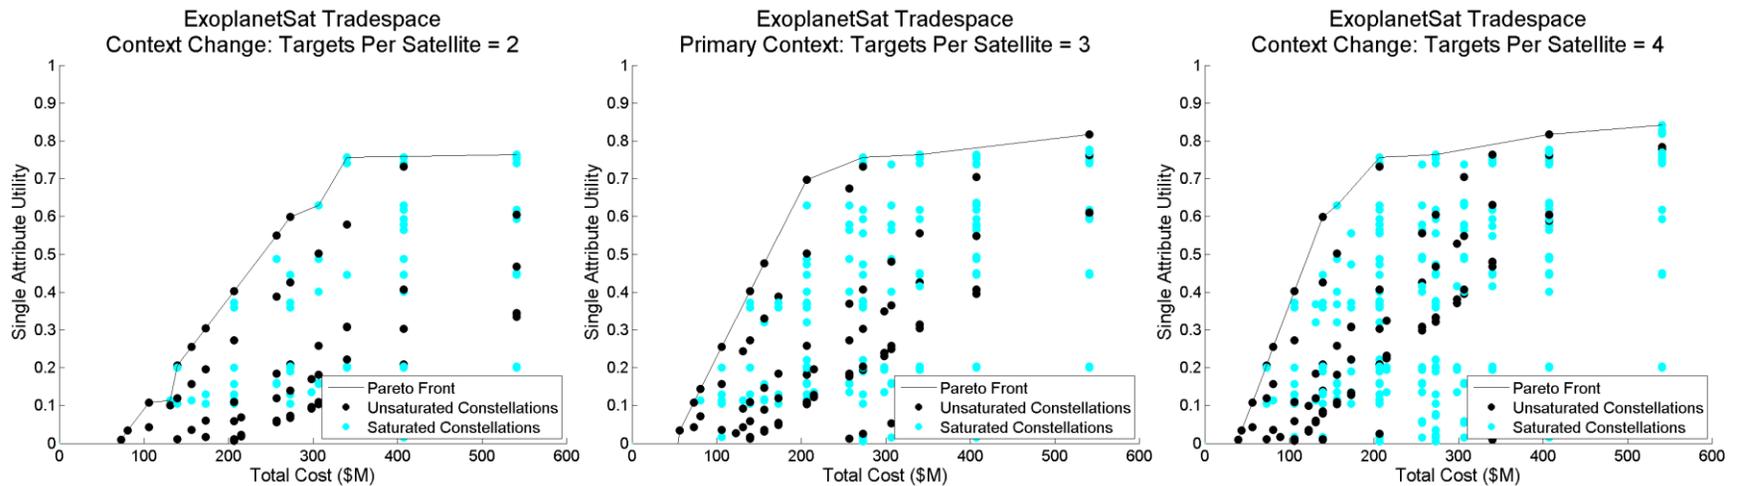


Figure 5-23: Tradespaces showing variations in Single-Attribute Utility in contexts ranging from number of targets per satellite equals 2 (left), 3 (center), and 4 (right).

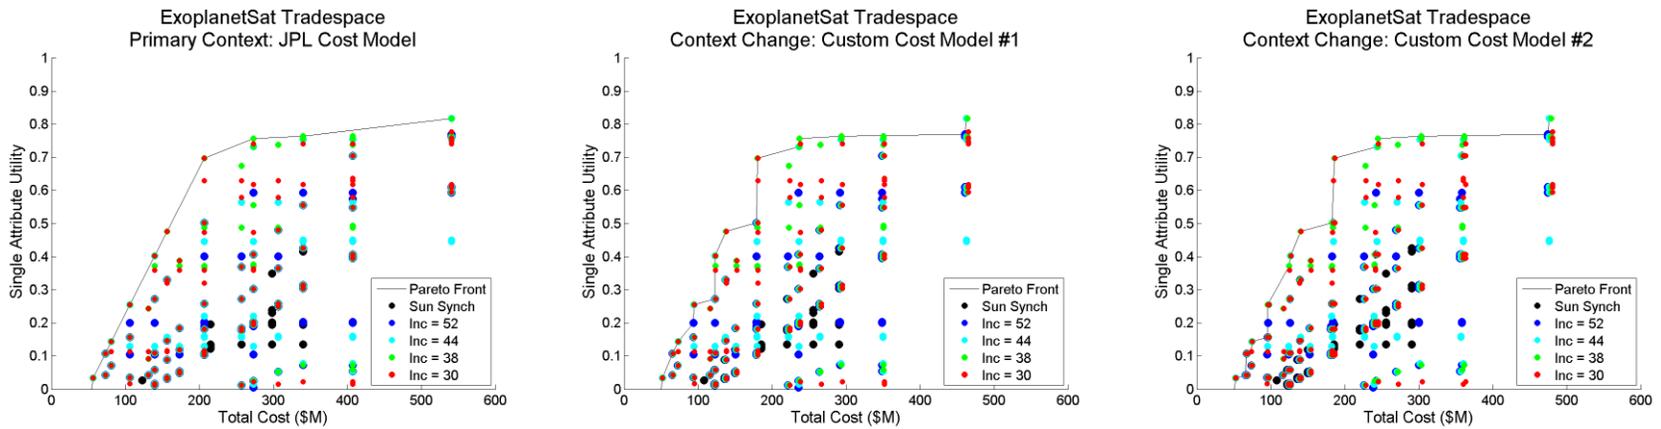


Figure 5-24: Tradespaces showing variations in the Expected Number of exoplanets detected (top) and Single-Attribute Utility (bottom) by using the JPL Cost Model (left), the Custom Cost Model #1 (center), and the Custom Cost Model #2 (right).

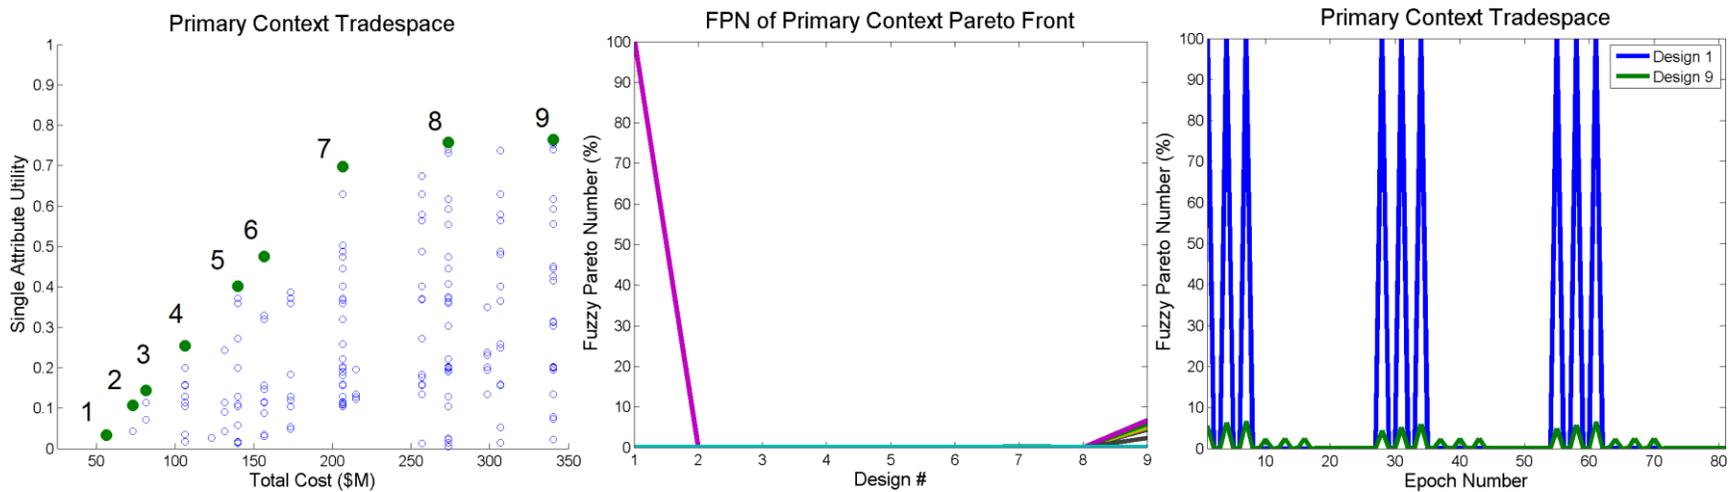


Figure 5-25: Multi-epoch analysis of (left) nine designs on the Pareto front examining (center) the FPN of each design in all 81 epochs, and (right) a closer look at Designs 1 and 9 and how they leave the Pareto front.

Seven of the nine designs marked along the Pareto front on the left of Figure 5-25 are Pareto dominant in all 81 epochs. Only designs #1 and #9 leave the Pareto front, as shown in the center of Figure 5-25. Design #1 can leave the Pareto front in certain epochs because it no longer achieves the minimum required satisfaction, which is why the only other fuzzy Pareto number it can have is 100. This happens when the number of possible targets that can be viewed at once is two and $R_*/a = 1/215$, which occurs in epochs 1, 4, 7, 28, 31, 34, 55, 58, and 61 as shown on the right side of Figure 5-25. Design #9 loses Pareto-optimality in the same epochs plus ones where $R_*/a = 1/150$, but the drop in Pareto efficiency is small. Either the fuzzy Pareto number is between 5% and 7%, or it is between 2.2% and 2.4%.

5.3.4.3 EEA Process 3: Single Era Analyses

As described in Chapter 2, one of the optional inputs to era analysis is a set of perturbations. Perturbations can either be *disturbances* or *shifts*. As described in Section 5.3.4.1, there are no ways for an epoch shift to occur because they are all mutually exclusive. However, there can still be disturbances, since disturbances affect the design variables and not the epoch.

The only real disturbance in this case study is the loss of a satellite. This is detrimental to the value because once there is a gap in coverage due to a loss of the satellite observing a set of stars, the data from those stars past that point is lost. If two transits have not already been detected, there is no way to know that no transit will occur around a star unless observations started over again.

Losing a single satellite affects the value delivered by the constellation in two ways. The first is that the total number of targets being observed for long duration n_i is decreased by the number of targets that the satellite is observing simultaneously. If there are many satellites, this is not detrimental to the mission, though it does lower the expected number of exoplanets detected.

The second way that losing a satellite affects the value delivery is if losing that satellite eliminates an entire *plane* from the constellation (assuming a Walker constellation). The consequences of this are much more difficult to predict, and are mitigated by having more planes in the constellation. The exact consequences are dependent on which plane is lost, and whether or not that plane's line of nodes is perpendicular to the Earth-Sun line during a solstice for the remaining operational lifetime of the constellation. Losing a plane relatively early in the mission would more likely limit the area of observable sky compared to later in the mission when there are fewer solstices left and thus less of a chance that the plane was critical to stars located near the edge of the observable sky.

The consequences of losing an entire orbital plane during the long-term observation phase are dependent on how many planes are in the constellation. A general rule is that once a plane is lost, the observable area of the sky drops to the same area as the exact same constellation with half as many planes. If an 8-plane constellation lost one plane, it would effectively be a 4-plane constellation because the one gap in RAAN is the same as all the gaps in the 4-plane constellation (see Figure 5-26). A 6-plane constellation would have the same observable area as a 3-plane constellation. For Walker constellations with two or three planes, losing a plane could effectively end the mission, or, if there was an option to upgrade the system again, reset all long-duration observations. If this happened during the phase where stars are observed to constrain stellar inclinations, it would only delay the start of long-duration observations.

The best way to mitigate failure of an entire plane is to add more satellites per plane and not cluster too many observations near one section at the edge of the observable sky. If only one satellite remains in a

critical plane, it can only observe one patch of sky, so if it is the only plane *capable* of continuing observations here, it may have to drop potential targets that are still observable only because it does not have the capacity to observe them all.

It is possible to upgrade the system over time by launching additional units. There are two ways to do this. The first way is to add units to the constellation in the same orbital planes that already exist. This could be done to raise the number of targets being observed simultaneously up to the point where the constellation becomes saturated. Adding any satellites to a saturated constellation will increase redundancy and preserve existing observations if other units fail, but there is no opportunity to be gained because all of the available targets within the observable area of the sky are already being observed.

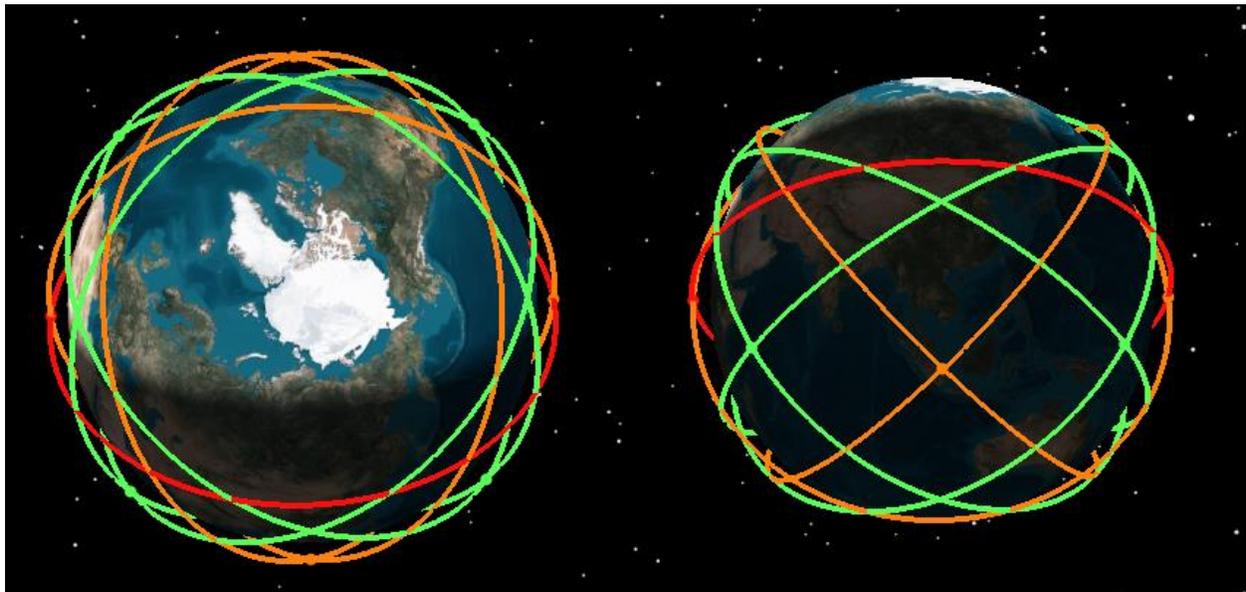


Figure 5-26: (Left) Overhead and (right) angled back view of an 8-plane constellation illustrating how losing a plane perpendicular to Earth-Sun line (red) at the Summer Solstice lowers the coverable area of the sky to that of a 4-plane constellation (green).

The second way to upgrade a constellation is to add more orbital planes. This would increase the observable area of the sky, thus increasing the number of possible targets to observe. However, although the *utility* of the constellation would increase immediately, the *value* that is delivered would not immediately increase because value is obtained only through long-duration measurements of candidate stars. Additionally, the option to add more planes is dependent on how many planes there were at the start; adding two more planes to a 6-plane constellation to make it an 8-plane constellation is ineffective because the planes would not be evenly spaced. The observable area of the sky *may* increase, but in an unpredictable way that shouldn't be relied upon during this Pre-Phase A study. Adding three planes to a 3-plane constellation such that all planes are evenly spaced will create a robust 6-plane constellation that is guaranteed to increase the observable area of the sky and thus the number of observable targets.

The time probability curve in Figure 5-7 (right) shows how value is delivered over an era as a function of utility. This is a case wherein the value model is a function of a single attribute utility, which is an even

simpler and degenerate case of one where epoch shifts do not change the relative weights among multiple attributes. As a result, the total value of data from the mission as a function of time takes the shape of this curve. However, this is *only* in the case of all data being delivered from the beginning of the mission; in the case of future upgrades to the constellation, that data doesn't begin to be delivered until later in the mission.

An example era that an ExoplanetSat design may undergo with both satellite failures and upgrades that extend the mission lifetime is shown in Figure 5-27. The top of the figure shows the instantaneous SAU as determined by the design vector. However, SAU (or MAU) is a metric of the *perceived ability to deliver value*, not the value that is delivered over the mission lifecycle. The bottom of the figure shows how different levels of static utility increase *EN*, which is the primary metric of success for PSS.

In the top of the figure, the black line represents the SAU at launch. After two years, some of the satellites fail, reducing the instantaneous utility. Although the probability curve is lower in those first two years, the mission doesn't suddenly lose those observations entirely. On the bottom, the dashed black curve shows how *EN* rises with the launch level SAU, while the dashed red curve shows how *EN* rises with the constellation as if it was deployed with only the units that survived after two years. Satellites launched at the 3-year mark instantly add utility to the constellation, but will not begin to affect *EN* until at least the 4-year mark, as shown by the dashed blue line.

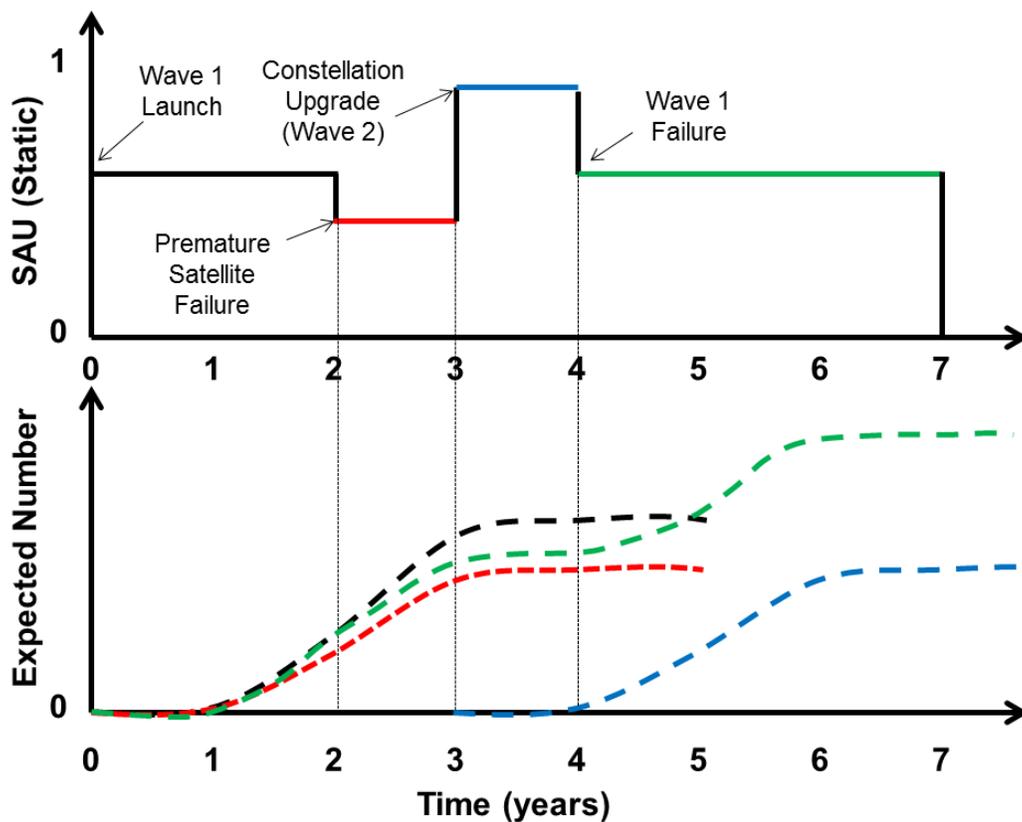


Figure 5-27: Example era with both failures and upgrades to demonstrate how static utility is not always the best metric for a system's value delivery.

The green dotted line shows how the EN grows as a function of time. Essentially, it's the addition of the black curve from year zero to two, the red curve from year two to four, and the blue curve from year three and beyond. Even though the mission ended with less utility than it started with, the EN was higher than it would have been had it not been upgraded. Additionally, the green line never goes *down* as a result of failure; this can only happen with an epoch shift, where the stakeholder would have a different set of preferences on how the mission delivered value. Since there are no epoch shifts in this example era, the total value that has been delivered by the mission will only increase.

The EN at the *end* of the era should be the one that is used at the beginning of an era to judge how well an era meets the stakeholder's needs. However, in reality, at the end of the mission, the success will be determined not by a probabilistic expectation but by the actual number of detections.

5.3.4.4 EEA Process 4: Multi-Era Analysis

Calculating the precise effects of single-plane failure for every constellation for as many era that could be constructed using Monte Carlo analysis is far beyond the computing power available for this project, but estimating the degradation of value of an unsaturated constellation as assets fail without an entire plane failing is possible. Multi-era analysis was conducted by assuming a probability of failure per year of satellite operation for each satellite and conducting 10,000 randomized trials on eras to see what the effect on the final SAU would be. Upgrades were not considered in multi-era analysis because they would have to be measured on a case-by-case basis.

An example of how EN and SAU fall given different satellite failure probabilities is shown in Figure 5-28. A 15-satellite, unsaturated constellation has $EN = 1.48$ and $SAU = 0.09$ when launched. However, if satellites fail, these numbers will drop. These figures represent 10,000 eras that were generated based on different failure probabilities. The curves present the launch case where no satellites fail (blue line), the average case from 10,000 simulations (black line), and +/-1 standard deviation (green and red lines). With a low probability, EN falls very little, and SAU does not substantially change. However, in the case of 10% probability of failure per year, EN drops to 1.08 on average, bringing SAU down to 0.015. This is dangerously close to being unacceptable to the stakeholders; a 14-satellite constellation with the same failure rate would on average be unacceptable because $EN_{av} = 0.98$, which is below the SAU curve threshold. While a stakeholder may begrudgingly accept a mission that has already launched and fallen just below the initial threshold, the stakeholder should know the risks up front of how value can be lost as a result.

While Figure 5-28 shows how a 15-satellite constellation delivers value over time, Figure 5-29 shows n -satellite constellations and their expected EN and SAU at the end of the mission. From here we can see the average and standard deviations in the drop in EN on the top graphs given a failure rate and the corresponding SAU on the bottom graphs. It is interesting to note that from these graphs we can see that designs that MATE alone would have labeled as acceptable (blue lines) suddenly appear unacceptable. Given that CubeSats typically last no longer than 2 years due to the radiation environment, this analysis shows why it is important to understand with more certainty the failure rate of small spacecraft and how to mitigate this failure when long-term observations are what deliver value to stakeholders.

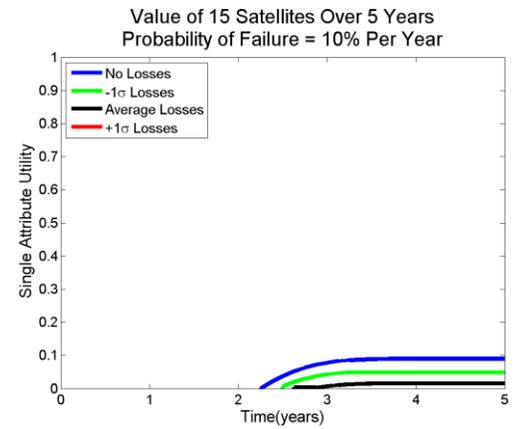
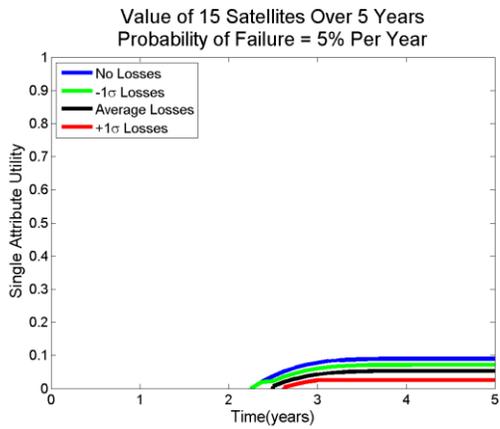
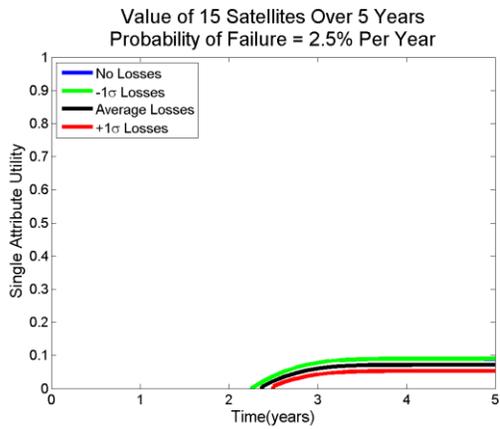
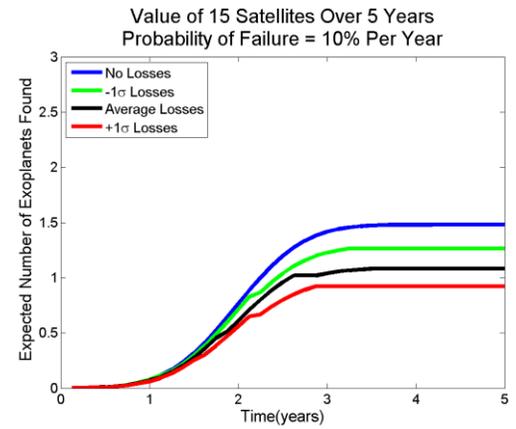
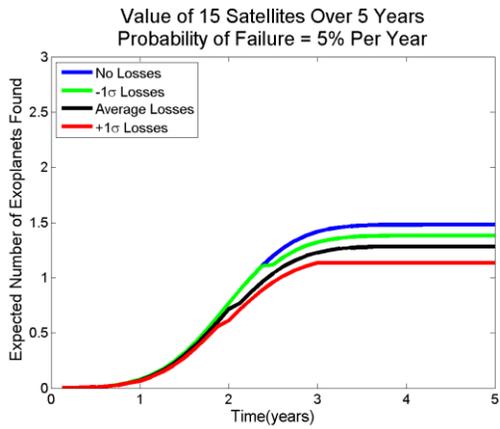
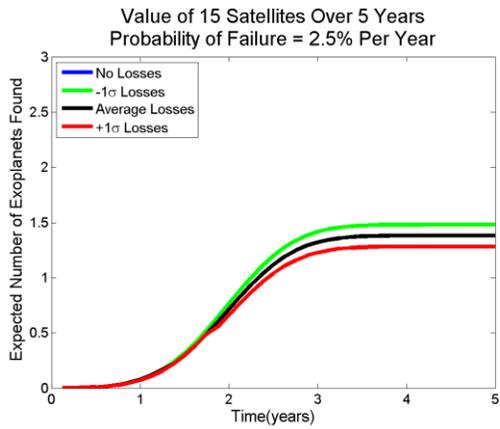


Figure 5-28: Drop in EN (top) and SAU (bottom) versus operational lifetime for an unsaturated constellation of 15 satellites given satellite failure rates of 2.5% (left), 5% (center), and 10% (right) per year of operation.

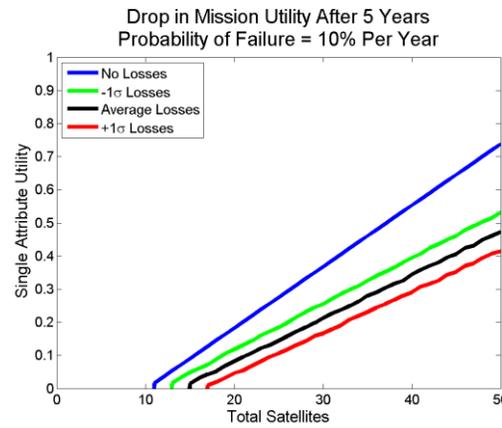
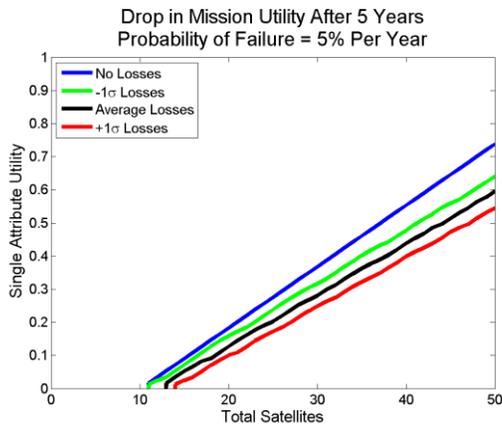
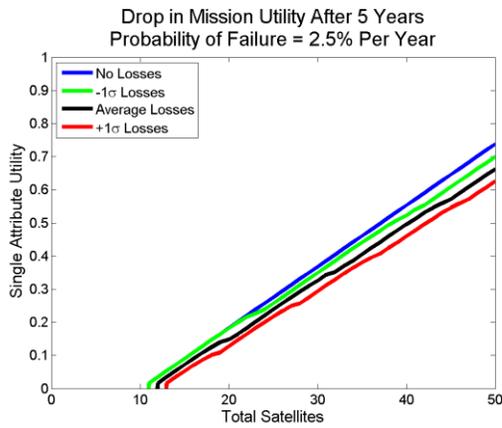
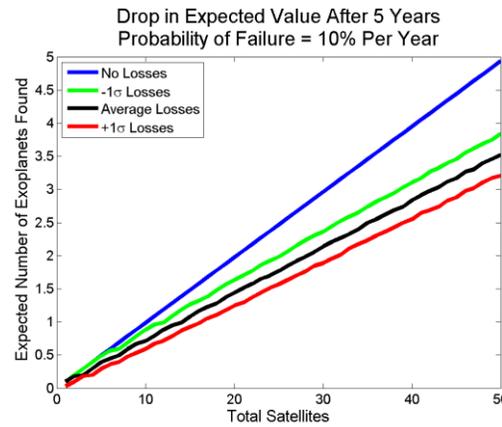
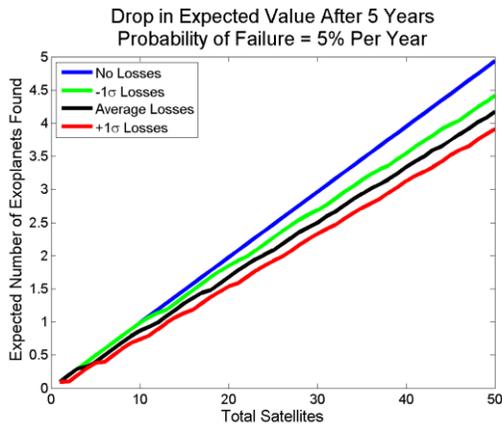
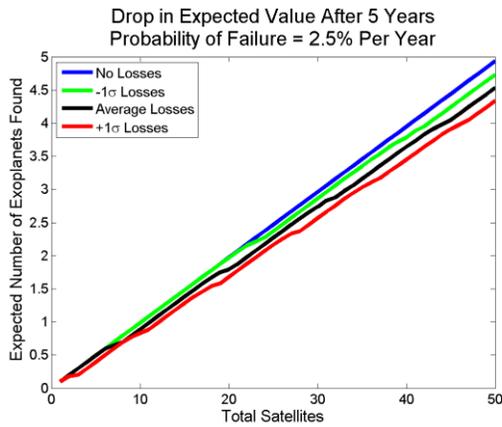


Figure 5-29: Drop in EN (top) and SAU (bottom) versus number of satellites in an unsaturated constellation given satellite failure rates of 2.5% (left), 5% (center), and 10% (right) per year of operation.

5.4 Discussion and Conclusions

ExoplanetSat serves as a powerful example of the RSC method in action despite the fact that it was applied on a limited tradespace as a demonstration and simpler case study compared to the other cases in this dissertation. Additionally, this mission has already undergone significant design work; therefore this work is more impactful because the results show concrete recommendations and insights that were previously unknown and may not have been uncovered without this work.

5.4.1 Discussion of RSC Method Results

The data discussed in the previous sections as a way to illustrate the method in action is only a small fraction of what can be learned from this case study. A discussion of additional findings, observations, and recommendations follows here.

5.4.1.1 Stakeholder Value Modeling

The value model that was elicited for ExoplanetSat accurately represents the PSS's perception of value from the mission and was validated through examination of the results. While the value model does not match the hierarchy structure proposed in Section 4.2.3.2, it would not have been formulated if the value modeling approach with MAUT had not been followed from the beginning. This case study especially helped illuminate the complexity of the value modeling process and reveal the cerebral schism between scientists and systems engineers.

Creating the value model required two crucial elements: extended interviews and conversation with the PSS, and a deeper knowledge of the science goal that went beyond the basic concepts and into the published literature. Without both of these elements, an incorrect measure of how the system obtains scientific gains for its cost would have been used to compare design alternatives.

There are many other approaches that may have resulted in creating a value model similar to the one that was created here. Even the PSS's own paper came close to explicitly deriving the SAU function. However, the methods used in concept studies may not try to explore other mathematically constructed value models, or even conduct a series of interviews to elicit how the stakeholders perceive value to the same degree that was taken in this approach. Surveys, science traceability matrices, and other methods are geared towards setting *requirements*, not eliciting *perceptions*. Additionally, with those methods, framing bias may be present that would go unexamined; with the interviews conducted through RSC, multiple questions were asked that showed inherent flaws in the initial assumptions that helped lead to the construction of the correct value model.

This case study can also show how concurrent design engineering teams like Team X can improve the products they deliver to their customers, even outside the scope of distributed satellites in space science missions. To their credit, Team X did not ignore the PSS's needs; Team X just did not fully understand how value was gained, nor did the PSS explicitly communicate it. The value modeling process forces a dialogue between the science and engineers roles that can foster deeper understanding to develop figures of merit for the designs that better meet the stakeholders' needs.

5.4.1.1 Optimal Constellation Design

The results from this study show that there are in fact optimal constellation designs to choose for this mission. The constellations at the critical inclination (38 degrees) or just below are Pareto-dominant,

while constellations with multiple satellites per plane are less susceptible to large losses in n_{obs} if individual satellites fail. Furthermore, it is never optimal for multiple satellites to observe the same target in the same orbit, not just because the added noise from multiple sensors reduces the probability of detection (which was known in advance), but also because lowering n_{obs} is such a strong component in the PSS's determination of value.

If a tradespace that considers many design options (including known “bad” options) had not been built, nor had a value model not been meticulously elicited from the PSS and the scientific literature, the conditions for optimality in this mission would not have been uncovered. The amount of effort required to build the performance model to analyze many designs was much higher than had it been designed to analyze only a few, but the knowledge gained as a result of those efforts more than offsets the cost that would have been wasted had this mission been launched without this knowledge.

Many of the parameters and assumptions used in this case study may change with better data. Among those changes are especially likely include the probability functions for detection (Figure 5-7). This would certainly have an impact on the conditions for optimality and may change the conclusions about which designs are optimal designs. Whether or not changes in these conclusions occurs is irrelevant; what is important is that the framework can easily incorporate these changes to determine whether or not those conditions change as new data comes forward over time.

Additionally, the performance modeling work in this study shows a concrete formulation for evaluating other LEO constellations for astrophysical science missions. Although ExoplanetSat's observational “requirements” are much stringent than most other constellations, the problem formulation is applicable to other science goals that have strong exclusion angle constraints, require some minimum observation times per orbit that is constant for a long period of time, and leverage shared sampling to meet the minimum acceptable attribute levels.

5.4.1.2 Value Degradation and Constellation Replenishment

When choosing a design that uses multiple assets, it is important to understand how the performance of the system degrades as assets fail over time and how the science returns can be affected by those degradations. The results from EEA showed how this can be built into a particular design's value proposition up front.

Single-era analysis shows that there is a significant delay between assets are launched and when the scientific benefits are accrued. Without first building the value and performance models on top of which to conduct EEA, this may not have been seen during the concept development stage. Unlike military applications, where responsive systems are deployed to operate constantly at high utility, value from this mission comes from the collection of extended observations. Losing coverage on one target for a brief period means the observation must be restarted in order to gain additional value from that target.

There are many caveats to how value is lost if assets fail that have not been considered here. For instance, a coverage gap of one week on one star due to one satellite failing may still be acceptable given that the probability of missing an exoplanet's transit would be ~2%; however, the timescale for launching an additional satellite and placing it into the appropriate orbit to continue observing that star is much longer than that. When a satellite fails, the star it was observing will most likely be lost, which is what has been shown in this analysis.

Multi-era analysis shows that the minimum number of assets to achieve minimum satisfaction is 16, given how the number of satellites correlates to SAU (Figure 5-29). Additional satellites create redundancy, but the results of MATE show that redundancy only reduces *EN* and does not help the mission under the current assumptions about detection probability.

Future iterations of this case study with more advanced cost and radiation modeling modules can use RSC to help designers and stakeholders understand the relationships among failure rate, mass production, and unit cost. If reducing the probability of failure causes an increase in cost, RSC offers a method to find the balance point between some acceptable failure rate and the cost per unit.

5.4.2 Recommended Designs Solutions

From the results of the MATE, it is clear that the original budget of \$30 million is not quite sufficient to ensure stakeholder satisfaction; however, this conclusion is very dependent on the accuracy of the JPL Team X cost estimates. Payload development is the single most expensive line item in the budget at \$2,588,040 for the first unit, more than 40% of the entire budget for the first unit alone. If this line item could be decreased significantly, more options would fit in this budget. Additionally, hardware and labor costs appear high given how small the satellite is.

Given the primary context (which is already a relatively pessimistic outlook compared to other contexts), a 5% failure rate per year, and no option for future upgrades, a robust choice that maximizes gains while mitigating risks is a 4-plane Walker constellation with 4 satellites per plane at an inclination of 38 degrees with each satellite observing a target for 20 minutes at a time with no per-orbit redundancy. Unfortunately, this constellation would cost more than double the original budget constraint, but unlike the constellation recommended by Team X, it would actually be able to find transiting exoplanets and meet the PSS's minimum expectations.

If a higher price point is a possibility, a constellation of 24 satellites in six planes, with four satellites per plane, would be recommended for a budget of \$100 million. This allows for more stars to be considered for targets due to the increased area of the sky that can be covered while limiting the odds of an entire plane being removed and reducing the available targets. Furthermore, it may be unrealistic to deliver satellites to more planes, and the increase in the number of targets as a result of more planes begins to diminish.

Luckily, a decision on the final design does not need to be made now. As development on the payload continues, more information will be known so that the PSS can make a more informed decision about the final constellation design in the future. Additionally, if the context outlook improves with new information, less expensive designs will still be able to deliver scientific value and satisfy the PSS.

5.4.3 Suggested Improvements for ExoplanetSat

Part of the TSE process is gaining a deeper understanding of the various trades that are involved in the design process that can lead to greater value. Although all satellite design variables were frozen in this tradespace, the effects that minor changes in these variables could have on the value proposition of this mission became increasingly evident.

5.4.3.1 Detector Field of View

The FOV of the ExoplanetSat optics is remarkably wide for a system of this size. However, the area of the CMOS detector that is dedicated to the primary science mission is woefully small in comparison. We saw minor changes in the SAU depending on how many stars that could be packed into the science portion of the window. Those estimates were based on the average density of stars in the areas of the sky that could be covered. During mission operations, there may be more clever ways to pack the long-duration targets into the window as it is now.

The context variable controlling the number of targets in the science window was maximized at four, but if instead the number was closer to six or even ten, EN would be higher for all options; less expensive missions with fewer satellites could still deliver satisfaction. It would be worthwhile to explore how much the cost of the system would increase if this science area of the CMOS detector was capable of using the entire CCD array of Detector A in Figure 5-2.

5.4.3.2 Detector Sensitivity

There are 41,137 stars with visual apparent magnitude $m_v < 8.000$, of which there are 3,769 that are considered “Sun-like” and non-variable. There are 83,114 stars with $m_v < 9.000$, and if the ratio of Sun-like stars is the same, there should be approximately 7,600 potential targets on the celestial sphere. While the real driver of EN is n_i , doubling n_{obs} doubles the density of stars, thus approximately doubling n_i without even conducting a trade on the detector’s science area size.

Reaching another magnitude further down in brightness may be much more difficult than simply increasing the science window FOV, especially given that the satellite is already fighting thermal noise induced by day-night effects, but knowing that the value proposition could nearly double make at least understanding the cost tradeoff for upgrading the payload hardware a worthwhile pursuit.

5.4.3.3 Solar Baffle

Easily the most critical improvement that can be made to the satellite design is the telescope baffle. The Solar Exclusion Angle as reported in Table 5-5 is 90 degrees. This constraint is the primary reason why not a single star in the sky is observable throughout the year for more than 20 minutes at a time, not even circumpolar stars, which by definition are observable year-round (though the area of the sky that is defined as circumpolar changes with observer latitude).

If a better baffle or solar shield could be designed and implemented on ExoplanetSat, it could completely eliminate the need for multiple planes; even a terminator, Sun-synchronous orbit would be able to view polar targets year-round with minor improvements in Solar Exclusion Angle. This would have a noticeable effect on n_{obs} but not n_i . Additionally, there would be no particular need to space out the satellites’ True Anomalies in an orbit, nor would there be a need for the satellites to share responsibilities for certain targets because a single satellite could observe a target for its entire lifetime without having to switch targets. If this were the case, then the Team X design *would* be a feasible solution, though it would still be less than satisfactory simply because more than eight satellites would be required for $EN \geq 1$.

5.4.3.4 Operations Cost

One of the largest cost factors in ExoplanetSat is operations. At \$500 per pass with two passes per day for each satellite plus a \$25,000 startup fee to use the NEN, four years of operations for a single satellite costs more than \$1.5 million. If there are other ways to reduce the downlink frequency or cost per downlink,

this will significantly decrease the lifetime costs of the mission. This is yet another disadvantage of distributed satellites compared to monolithic systems; operations costs are non-negligible, especially for long-duration missions. A DSS that can reduce the reliance on expensive operations components will benefit during the operations phase.

Another sensitivity analysis can be performed that factors in the cost of building a dedicated ground station or network of ground stations. If it is less expensive to build and operate ground stations than it is to utilize the NEN, it could be part of the tradespace, but this is outside the scope of this iteration of the case study. Additionally, given the large number of satellites available, there may be a way to negotiate lower prices for bulk purchases of downlink time to many satellites. This is more feasible with private companies than it is with NASA or other government agencies, which typically operate on fixed pricing schemes and have no incentive to reduce prices or negotiate with customers.

5.4.4 Comparisons to Monolithic Systems

While the aim of this dissertation is to show how new science can be conducted and how scientific stakeholder goals can be met under different definitions of value rather than strictly comparing DSS to monolithic equivalents, in this case, there is a monolithic equivalent that is being planned by ESA.

The PLATO mission comes with a steep price tag of one billion dollars. Even the most expensive ExoplanetSat constellation options under the most pessimistic cost models do not have a price that high. However, PLATO has the potential to outperform ExoplanetSat because it has a solar shield that allows it to view stars at least once per orbit for some minimum amount of time throughout the year and its cameras have higher sensitivity, allowing it to view dimmer and therefore more targets at once.

These two constraints may not be possible for ExoplanetSat to overcome, and were not considered part of the TSE process for this reason, but TSE shows clear evidence that satisfaction can still be achieved for a much lower cost with a DSS compared to its monolithic competitor. A potential design that costs an order of magnitude less than PLATO is still expected to find more than two Earth-sized planets and have over a 90% chance of finding at least one because it leverages **shared sampling**.

5.4.5 Future Work

During the later stages of this case study, the scope of the ExoplanetSat mission changed significantly due to the selection of the competing TESS and PLATO missions. As a result, the utility of a mission to search specifically for exoplanets diminished. Had the information about the stakeholder's perceptions on the value of the mission under these contexts been understood earlier in the case study, a more detailed epoch-era analysis could have been conducted that includes the mission's new science goals.

ExoplanetSat was selected as the first case study primarily for its simplicity as a first-order setup to practice implementing RSC: there is only one science goal; the design variables relate to the constellation only, meaning none of the design variables of the satellites themselves were changed; and there is only one set of stakeholder expectations that mattered.

Now, ExoplanetSat has changed its mission significantly enough that it has been renamed the Arcsecond Space Telescope Enabling Research in Astrophysics (ASTERIA). Rather than aiming to achieve one science goal, this new mission's science goals include transit searches, asteroseismology, stellar rotation period observations, radial velocity follow-ups, reverberation mapping, ultraviolet transits, and near-

infrared spectroscopy [306]. The goals of this “science menu” require instrument tradeoffs and careful value elicitation to balance the weights among these goals.

A constellation of heterogeneous small satellites appears to be the best solution for accomplishing all of these goals, which requires instruments that operate on multiple spectral bands (infrared, visible, ultraviolet). A question that must be addressed is what is the most cost-effective way to have all of these instruments in operation? A naïve designer may assume that packing them all into one spacecraft and making the fleet homogeneous is the best solution. This would increase the size of each satellite, therefore the launch costs, but at the cost of some coverage metric since fewer satellites are affordable. Additionally, not all payloads are required for all science goals (only reverberation mapping requires all three spectral bands, and these observations are not simultaneous, meaning spatial distribution of these payloads is actually desirable), and the coverage requirements vary for all of them.

With the addition of these extra science goals and alternate payloads to consider as design variables, the tradespace of ASTERIA becomes much richer compared to ExoplanetSat. Now more than ever for this mission, a method like RSC is important to understand how value is perceived from which stakeholders, what design variables drive both cost and utility, and what the risk and opportunity are for a constellation of distributed satellites over the course of its mission lifecycle.

5.5 Chapter Summary

With this case study, the value of leveraging the emergent capability of **shared sampling** was examined in a DSS mission with CubeSats to detect transiting exoplanets. A world-class concurrent engineering studio conducted a design study on the same mission, but the results that were returned were clearly incapable of satisfying the stakeholder’s needs when examined through RSC.

A value model showing what actually would satisfy the stakeholder and find fundamentally significant scientific discoveries was constructed based on utility theory and published literature, though this value model was a *constructed model* rather than purely a *mental model*. A performance model was created that determines the area of the sky a constellation can cover for some minimum amount of time per orbit. This performance model is the first of its kind and can be used for future astrophysics constellations where long-term observations are critical for mission success. The model can also be easily modified to work for specific parts of the year or find acceptable coverage for some fraction of the year instead of a full year.

Tradespace exploration identified design options that would deliver value to the stakeholder and helped characterize the sensitivities of the tradespace to changes in those variables. Valuable knowledge from this study was also gained through EEA, where sensitivities in contextual variables and survival rate were explored. From these new perspectives, information about how value was delivered over the course of the mission lifecycle was uncovered, as was how the expected failure of satellites over time could affect the value proposition. These effects would not have been discovered without the application of EEA.

With this information, the PSS can make more informed decisions about the course of the mission development and the final constellation design in order for the science goal to be achieved. Given a limited budget to launch N satellites, the PSS can choose a specific constellation design that maximizes the value of the mission while still maintaining robustness to mitigate losses as some of these small satellites inevitably fail over long periods of time.

In order to guarantee that the system will achieve the minimum satisfaction criteria, $EN = 1$, the constellation must launch at least ten satellites; however, if the characteristics of failure are known and can be predicted, more satellites must be launched. If the probability of failure per year is as high as ten percent, an additional six satellites would be required to prevent the mission from falling below the minimum threshold to preserve value. However, this number may change if other context variables also change. Given the information gathered in this case study, the results of changing those context variables are also known and can be accounted for.

There happens to be a proposed monolithic mission (PLATO) with similar goals that can be directly compared to ExoplanetSat. While it does achieve complete stakeholder satisfaction, its budget is nearly twice as high as the most expensive options that were explored herein. Possible modifications to ExoplanetSat were identified that were beyond the scope of this study that could make ExoplanetSat even more competitive compared to that monolithic mission.

CHAPTER 6

HOBOCOP

Heliocentric-Orbiting Baseline-Optimized
Cosmos Observation Paradigm



Robocop character © Orion Pictures, 1987. Parody artwork by reddit user /u/Pan_Goat.

6. Case Study: HOBOCOP

*“I’m stranded in space, I’m lost without trace, I haven’t a chance of getting away
Too close to the sun, I surely will burn, like Icarus before me or so legend goes
I think of my life, reliving the past, there’s nothing but wait ‘til my time comes
I’ve had a good life, I’d do it again, maybe I’ll come back some time, my friends
For I have lived my life to the full, I have no regrets
But I wish I could talk to my family to tell them one last goodbye”*

IRON MAIDEN, “SATELLITE 15... THE FINAL FRONTIER”

The second case study moves beyond low Earth orbit and develops the concept of the Heliocentric-Orbiting Baseline-Optimized Cosmos Observation Paradigm (HOBOCOP), named for self-sufficient traveling workers. This mission concept is an entirely novel idea developed by the author with no other governing authority or principal investigator. The primary science goals are to understand the Sun’s magnetosphere, the solar wind, and the heliospheric current sheet with more detail than ever before, but there are additional commercial and scientific gains that can be achieved through the course of this mission’s lifecycle by a number of stakeholders. These include an inner solar system internet network and small plug-and-play payloads and satellites that can connect to it. This mission heavily leverages **simultaneous sampling** to achieve its primary mission objectives, but the value proposition also benefits from leveraging **stacked sampling** of multiple scientific goals.

6.1 Case Study Introduction and Motivation

While CubeSats generally achieve adequate data downlink rates to ground stations from LEO, communications systems design become increasingly critical in returning data to Earth, and the physical limitations become insurmountable as the distance between a satellite and a ground station is more easily measured in astronomical units than it is in kilometers.

Consequently, as more satellites are launched into deep space that are difficult to communicate with, the most effective way to communicate with them is the Deep Space Network. Unfortunately, this system is outdated, underfunded, and already oversubscribed; suddenly adding a large number of satellites that require a long time to downlink their data would not only be difficult and expensive, it would detract from every other mission that requires the DSN to communicate with Earth.

The core of this mission is based on the HOBO[®] sensor concept [307]. On Earth, HOBO[®] sensors are used in scientific investigations in harsh environments where it is difficult for humans to be present or retrieve data on a regular basis. HOBO[®] sensors can operate autonomously for months or years at a time collecting data before they are retrieved. This data is usually time-insensitive, meaning there is no rush to retrieve the data.

Space science missions are increasing in autonomy as advances in computing make smaller satellites more independent than ever before. In missions where there is no need for continuous uplink, command, or control of a spacecraft, the spacecraft may be capable of going for some time without the need to

communicate with Earth. This gives several advantages to the design of a small spacecraft, most notably in the design of the communications hardware. Satellites that operate beyond Earth orbit tend to have antennas larger than most small satellites simply because of the physical limitations surrounding radio communications. If small satellites did not need to communicate with Earth at all, or could communicate periodically when they made close approaches to Earth, the size of the communications subsystem can be much smaller.

However, if a satellite will never make a close approach to Earth, it would either have to have a large antenna or be forced to communicate at a slow rate. If the ExoplanetSat mission was suddenly moved from LEO to a distance of 0.01 AU and communicated with the same ground systems, its data rate would drop from 2 Mbps to less than 1 bps. While there certainly are better communications hardware options than omnidirectional antennas, once a small satellite leaves Earth's sphere of influence it needs to either have a large antenna or communicate with a relay.

Having a large antenna may not always be an option. In the case of a satellite entering a challenging environment, it may be infeasible or at least very difficult to design a small spacecraft that can survive at all, let alone while carrying a high-gain antenna. As antenna size and thus total mass increases, the number of assets that a mission leveraging the emergent capabilities of DSS can afford goes down, lowering the potential gains from the mission. In the case of a satellite going close to the Sun, the antenna would cause problems with the thermal requirements and design that could swallow the mass budget.

In these cases, it may be better for the satellite to communicate through a relay instead of communicating directly with Earth. With relay communications, higher frequencies can be used that do not penetrate through Earth's atmosphere. Intersatellite links that communicate over the V-band (~40-70 GHz) take advantage of the fact that beam spreading is much lower, so antenna gain is higher and so is the data rate. Granted, it is virtually impossible for a satellite to compete with the 70 m diameter receiving antennas used by the DSN if they were at the same distance, but because the relays can be placed in strategic orbits to communicate from much shorter distances and communicate with higher frequencies, it is possible to lower the cost of DSS missions by leveraging the potential benefit of *path diversity*.

The HOBOCOP tradespace shows how design decisions as simple as an antenna size can drastically affect the value proposition of small satellites when they venture beyond LEO. This mission not only explores engineering trades involving small satellite operating beyond LEO, it also shows how groups of satellites can work together to achieve unique science goals by leveraging *simultaneous sampling*.

6.1.1 Primary Mission Goals and Stakeholders

The primary science goals of HOBOCOP involve studies of the Sun and the interplanetary environment. However, if there is the possibility of leveraging an interplanetary relay network to send time-insensitive data back to Earth, additional stakeholders such as the owners and operators of that network would need to have some say in the design process. For this reason, this mission will be viewed from many different angles and possibilities.

6.1.1.1 Primary Science Goals

The science goals of HOBOCOP involve studying phenomena that have already been studied extensively with monolithic satellites but only one dedicated DSS mission [308]. The magnetosphere of the Sun is a complex, 3-dimensional field that is rapidly changing compared to the length of time monolithic missions

can study it, and there is a desire to study it not only from more than one angle, but also from higher inclination orbits to study what happens at the rotational and magnetic poles throughout the 22-year solar cycle (the solar sunspot cycle has an 11-year period, but this period represents a 180 degree flip in the magnetic axis compared to the rotational axis. The full period of this cycle is therefore twice this length).

Three phenomena associated with the Sun’s magnetosphere have been identified that would benefit by being studied with DSS: the surface magnetic field, the solar wind, and the heliospheric current sheet. These goals and the instruments that can be used to achieve them are summarized in Table 6-1.

Table 6-1: Primary Science Goals in HOBOCOP

Primary Science Goal	Primary Instrument
Characterize Sun’s Surface Magnetic Field	Vector Field Magnetograph
Characterize Solar Wind	Fast Particle or Plasma Analyzer
Characterize Heliospheric Current Sheet	Magnetometer

In all of these cases, the science goals are to better understand phenomena so they can be modeled with higher fidelity. The basic phenomenon is well-understood, but the full picture is not. For science goals like these, one could argue that obtaining *more* data points is valuable, which means quantity is more important than quality, which means more, smaller instruments recording data at lower resolution and accuracy (to a point) are better than one, highly accurate instrument.

6.1.1.2 Primary Stakeholders

For this mission, a postdoctoral researcher at the Harvard-Smithsonian Center for Astrophysics served as the Primary Scientific Stakeholder (PSS). Additional opinions on the value of different goals were also solicited from a scientist at NASA JPL, a professor from the MIT Department of Physics, and a graduate student from the Earth, Atmospheric, and Planetary Sciences (EAPS) department at MIT. While there are many solar physicists in the community who could have been interviewed to elicit a value model, these were chosen due to personal connections within EAPS.

6.1.1.3 Secondary Mission Goals and Stakeholders

In addition to the science stakeholders in this mission, there is room for additional stakeholders involving a system of relay satellites. Inputs on how to measure the values of such a network were elicited from conversations with Dr. Alessandro Golkar from the Skolkovo Institute of Science and Technology.

6.1.2 Mission Concept of Operations

While the concept of operations for this mission is largely dependent on the design and context variables, this section will frame the general case and some of the options for how this mission might be structured.

6.1.2.1 Single Satellite Model

The idea behind this mission is that small, heliocentric-orbiting, baseline-optimized (HOBO) satellites can be delivered to separate orbits that each requires changes in velocity (ΔV) much higher than most other missions, even ones to the outer planets. An individual HOBO satellite consists of two modules: the primary spacecraft bus that contains a plug-and-play payload and other subsystems required for survival in the interplanetary environment; and an expendable electric propulsion module that will deliver the satellite to its intended orbit but may not be able to survive the thermal environment of the final orbit and can be released from the probe. The primary spacecraft bus is illustrated in Figure 6-1 and

the combination of the two units is illustrated in Figure 6-2. The “baseline-optimized” part of the acronym HOBO refers to a multidisciplinary optimization of the propellant, propellant tank, and solar shield masses (the solar panel area and power production unit size were independent variables).

The original concept of the HOBO satellite was to essentially be a 3U CubeSat with an attached heat shield (left, Figure 6-1). Originally, the only instrument that was being considered was a magnetometer, and similar 3U CubeSats have been launched with such a payload. Given the nature of the orbits that were being considered, a heat shield was absolutely necessary to prevent the spacecraft from overheating. However, if other instruments are also being considered, such as a telescope and vector magnetograph, the front of the spacecraft would need an unobstructed view of the Sun (right, Figure 6-1). Such a design would require more extensive thermal analysis, but given how close other missions in the design phase plan to get while having unobstructed views, it is still possible.

The expendable propulsion module consists of a set of electrospray thrusters, a fuel tank, and a deployable solar panel. Normally, small spacecraft power budgets would never be able to support a Solar Electric Propulsion (SEP) system, but for this mission there are no other options; the orbits that need to be achieved are expensive in terms of ΔV even for large satellites, and previous inner solar system missions have relied on complicated flyby maneuvers to achieve their orbits. This is not an option due to the timing desired for simultaneous close passes of multiple spacecraft.

6.1.2.2 Single Satellite Concept of Operations

While there are a number of ways to leverage planetary flybys (“slingshots”) to reduce the required ΔV to enter an orbit, and as a result lower the total mass of the system, such strategies would require timing that would prevent satellites from being at opposition to each other or making simultaneous perihelion passes as desired by the PSS, which will be discussed in the next section.

If the satellite used chemical propulsion, or if the thrust maneuver were short enough, it could be safe to assume that orbital aphelion would be close to 1AU and that certain orbits would periodically put the satellite close to Earth. While this turns out to be the case in some designs and contexts, orbits that maneuver the satellites closer to the Sun with electric propulsion require long burn times that can significantly lower the aphelion radius, and thus the link distance between Earth and the satellite will never be short.

The HOBO satellite will start out in a near-Earth, heliocentric orbit as shown in Figure 6-3. It will then perform an electric thrust maneuver (1) to slow its velocity to enter an elliptical orbit (2). At this point, the spacecraft may jettison its propulsion module. The spacecraft will begin its scientific investigation and take data along its orbit, and when it reaches its closest points to the Sun (3), it can retract its solar panels so they stay safe from the increased solar heat flux. When it reaches a point higher in its orbit, it may communicate with a relay satellite (4) if one exists; if none exist, it would have to communicate with Earth, which may never come closer than 0.25AU, depending on the orbit.

The PSSs were also interested in achieving polar heliocentric orbits. Inclination maneuvers are typically the most expensive in terms of ΔV , and achieving a polar heliocentric orbit has never been achieved without a flyby of Jupiter. With SEP and specifically timed orbits, it is possible to achieve this with Earth flybys, but the required ΔV (~42 km/s total ΔV from ecliptic to polar orbit, not including the ~11 km/s required for Earth escape) is still higher than any mission that has ever been launched before.

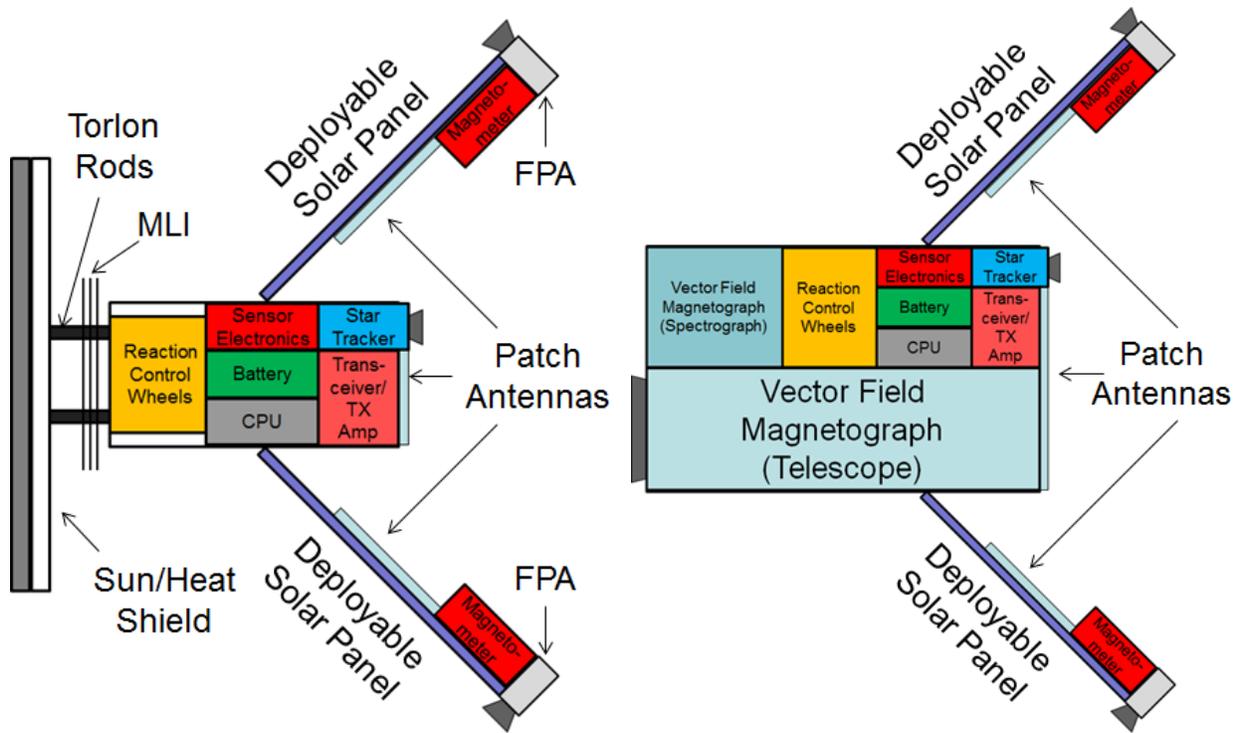


Figure 6-1: Preliminary concept sketch of a HOB0 satellite primary spacecraft bus. (Left) Close-approaching HOB0 with sun shield. (Right) Polar-orbiting HOB0 with VFM.

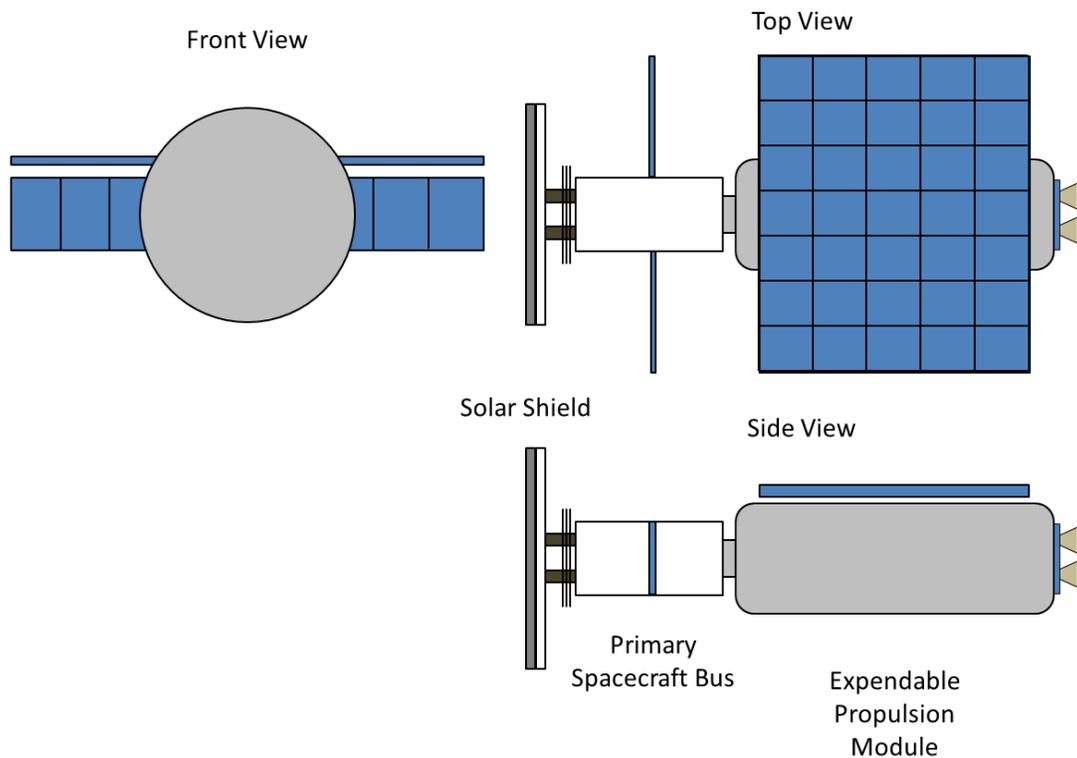


Figure 6-2: Preliminary concept sketch of a HOB0 satellite primary spacecraft bus and expendable propulsion module.

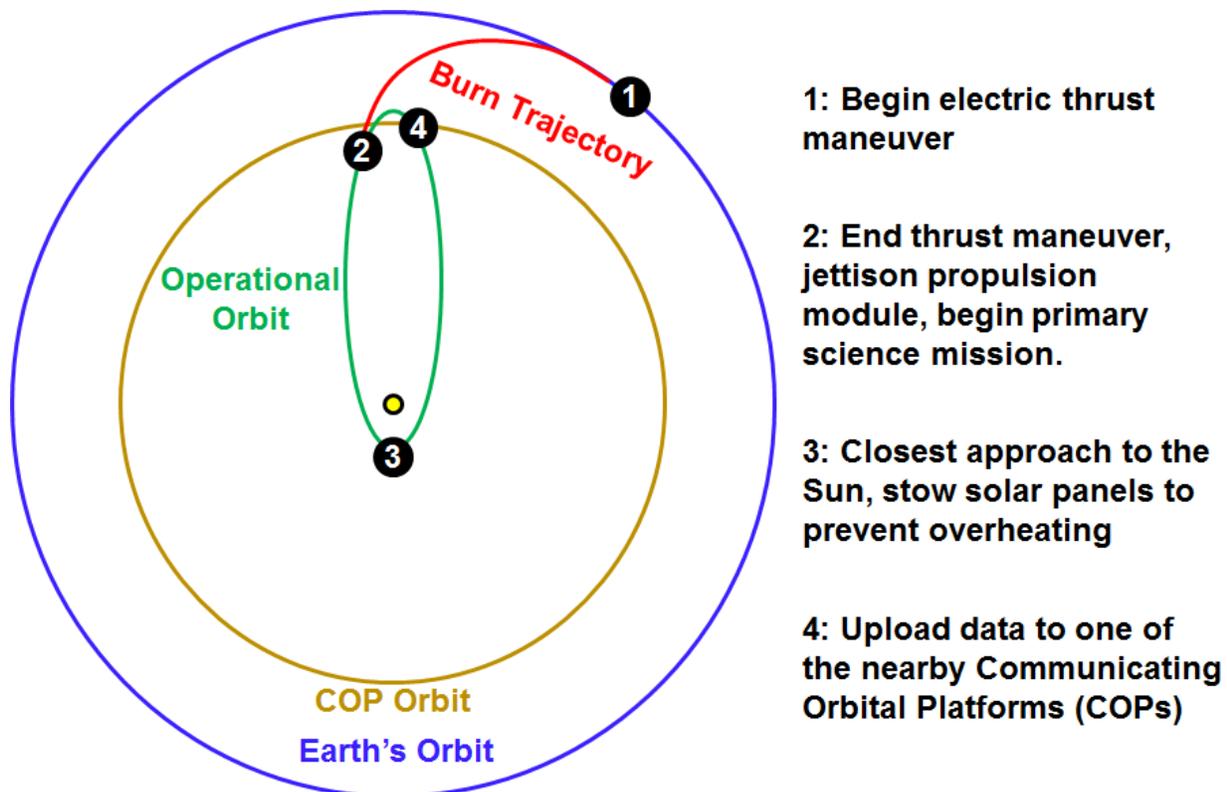


Figure 6-3: An individual HOBOSatellite concept of operations (for close-approaching HOBOSatellites).

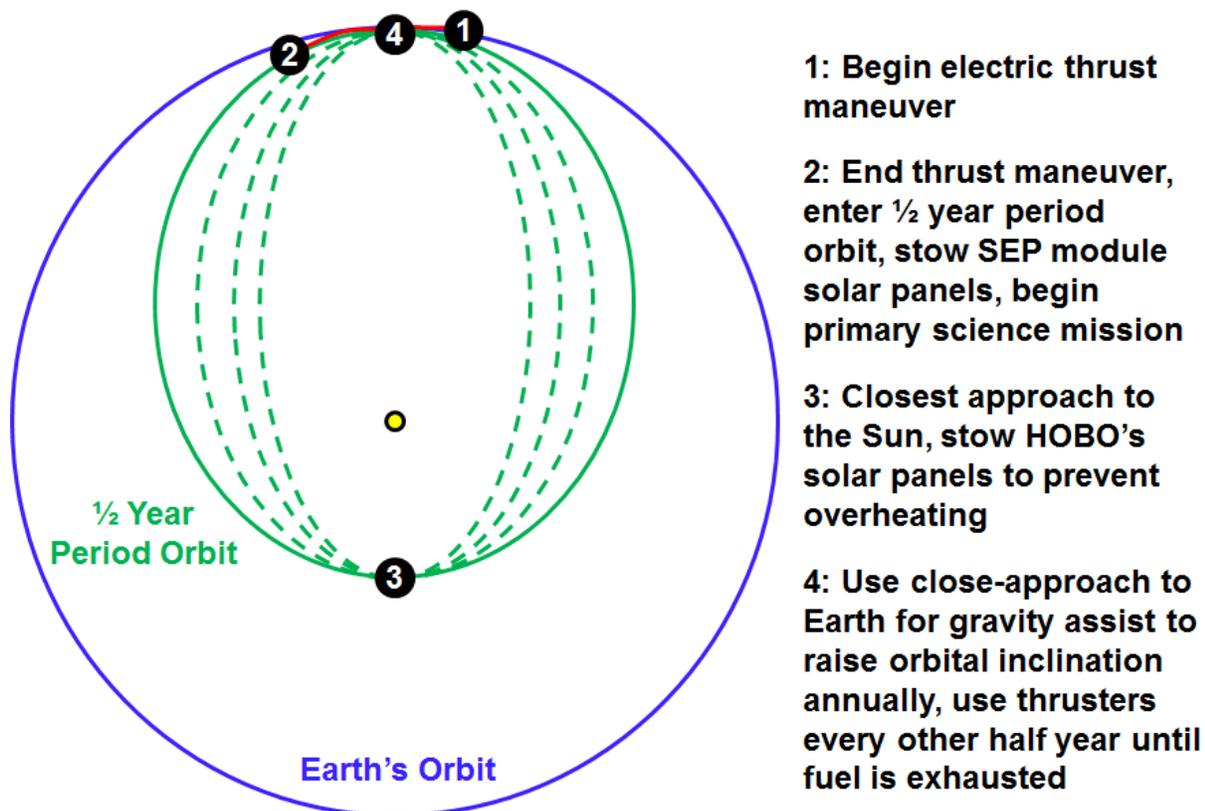


Figure 6-4: An individual HOBOSatellite concept of operations (for polar orbiting HOBOSatellites)

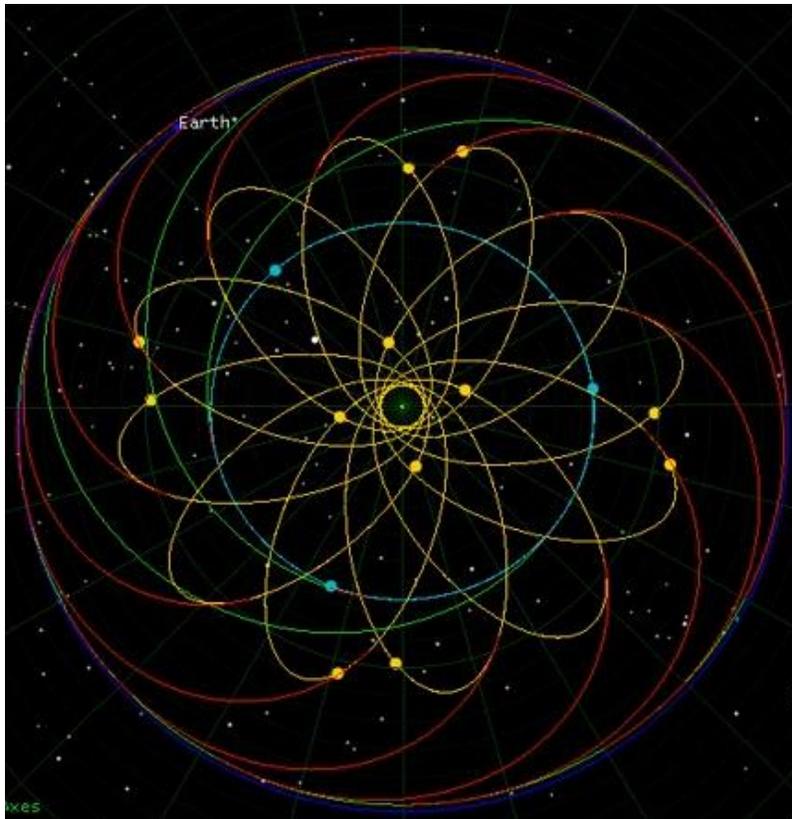


Figure 6-5: Concept of Operations of a constellation of close-approaching HOBOS in $\frac{1}{4}$ -year orbits with COPs. Royal blue is Earth's orbit. Red curves are HOBOS' burn trajectories. Yellow orbits are HOBOS' operational orbits. Green curves are the COP burn trajectories. Cyan represents possible COP orbits.

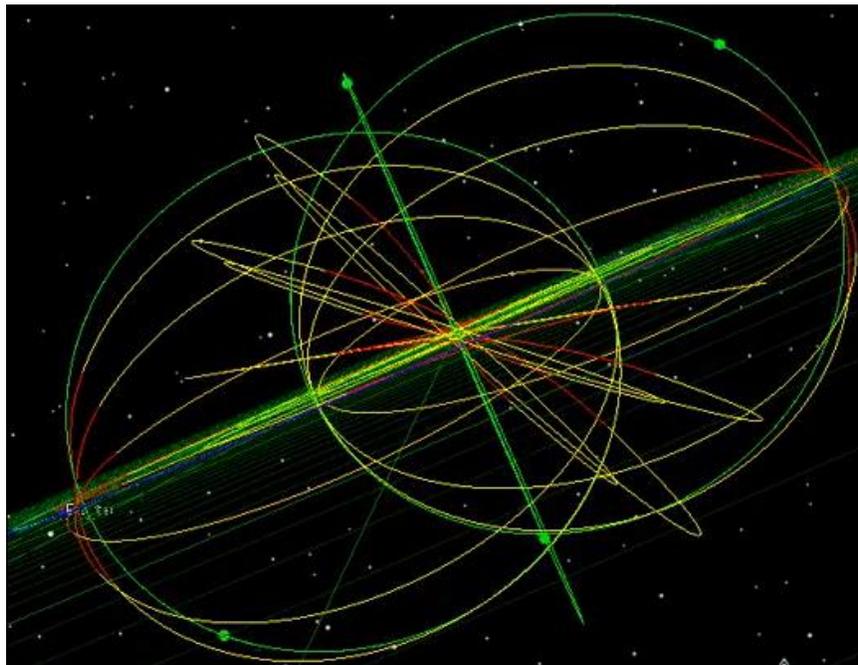


Figure 6-6: Concept of Operations of polar-orbiting HOBOS in $\frac{1}{2}$ -year period polar orbits with no COPs. Red lines are the inclination-changing burns. Yellow lines are coasting orbits. Green lines are operational orbits.

The method this mission will use to achieve polar heliocentric orbits is illustrated in Figure 6-4. First, the satellite will thrust to slow down and enter a ½-year period orbit. Rather than jettison its propulsion module, it will retract the solar panels while the science mission starts. Once the satellite reaches aphelion for the first time, it will perform an inclination change maneuver. This will only change the inclination by a small fraction, so additional maneuvers are needed. During the second orbital aphelion, exactly one year after the initial orbital insertion, Earth will be in range for a gravity assist. Every odd ½ year, the satellite will use thrust only, while every even ½ year the satellite will use a combination of thrust and a gravity assist, so that within a few years, the desired heliocentric inclination would be achieved.

6.1.2.3 Full Mission Operations

The initial concept of HOBOCOP was to place dozens of 3U CubeSats in the same heliocentric orbit as Earth spaced out evenly with different sets of Communicating Orbital Platforms (COPs) in eccentric orbits to periodically fly close enough to establish a communications link. Two separate graduate course final projects were dedicated to this concept. The first was a [network model of the communications system](#) combined with a Keplerian orbital propagator [309]. The second was an MDO study to minimize total system mass of all spacecraft [148].

However, not only was this implementation impossible due to the fact that achieving those orbits for dozens of satellites would take longer than the expected lifetime of a human being, but the positioning of the satellites was not satisfactory to the stakeholders to achieve the science goals that were initially set out. These models and optimization techniques nonetheless set the basis for modeling within the tradespace exploration phase once the PSS opinions were elicited.

One possible implementation of HOBOCOP is shown in Figure 6-5 (a [video demonstration](#) is also available for viewing [310]). In this example, sets of HOBOS make simultaneous passes of the Sun in ¼-year orbits. All the satellites start off together just outside of Earth’s sphere of influence, and HOBOS that are released three months apart swing near the Sun simultaneously as a result. In this figure, there are three sets of four HOBOS (yellow dots) that make perihelion passes together and then communicate with COPs (blue dots). This is analogous to the concept of operations in Figure 6-3.

An alternative implementation that uses polar orbits with no COPs is shown in Figure 6-6 (a [video demonstration](#) is also available for viewing [311]). This implementation only works if the HOBOS are in ½-year orbits and have propulsion systems capable of surviving for the entire mission. Rather than utilizing a gravity assist from Jupiter to achieve a polar orbit in one pass, these HOBOS leverage a series of Earth flybys once a year combined with SEP thrusts at aphelion to gradually raise their orbital inclination. The advantage of this concept is that pairs of satellites can orbit the Sun on opposite sides to observe the poles simultaneously, something that has never been done before. This is analogous to the concept of operations in Figure 6-4.

Another possibility for this approach is to keep the satellites in 1-year orbits and raise their inclination from an Earth-like heliocentric orbit. This way, Earth flybys can be accomplished twice per year, raising the inclination faster at the cost of better resolution and closer passes to the Sun. However, pairs of satellites would not have directly opposing views of the Sun for their entire orbit, only once every half year, because they cannot be spaced out on opposite sides of a 1-year orbit. This also removes any radial Doppler shifting in the spectral measurements throughout the satellite orbit, meaning the detector’s design

can be simpler and require less spectral bandwidth. One way this is better is that the opportunity to downlink data is more frequent because of the closer passes to Earth on a more regular basis.

The tradespace of HOBOCOP will explore several concept options including the ones shown above.

6.2 Case-Specific Literature Review

There are a number of topics in the published literature that are relevant to this specific case study that were not included in Chapter 2, including the science fields that study the Sun as well as the development of technologies that would make this mission more feasible now than at any point in the past two decades.

6.2.1 The Sun and Solar Phenomena

There are many phenomena associated with the Sun. The ones that HOBOCOP will consider as potential targets of interest are the magnetic field of the Sun's photosphere, the solar wind, the interplanetary magnetic field (IMF), and the heliospheric current sheet.

6.2.1.1 The Solar Magnetic Field

The Sun's large scale magnetic field is generated by a dynamo process powered primarily by differential rotation in the sun's convective zone. Because the Sun has no solid surface, the rotation period is different at the equator than it is at the poles – 25.1 days at the equator, 34.4 days at the poles. Smaller scale magnetic structures (e.g. sun spots) are generated by convection and torsion of plasma that carries magnetic field lines.

While the Sun is composed of matter that is a gas under standard atmospheric temperature and pressure, the matter is actually in the form of plasma, a different state of matter entirely. Plasma is composed of ionized atomic nuclei and the electrons that have been stripped away from those nuclei. Because the oppositely charged particles are separated, the motion of plasma creates and disturbs electric and magnetic fields. As a result, magnetic field lines tend to follow the plasma.

The Babcock Model describes the 22-year solar magnetic cycle [312]. At solar minimum, the magnetic field is a dipole aligned with the rotational pole of the Sun. Because of the differential rotation rate, the magnetic field lines are wrapped up by 20% every ~27 days. After many rotations, the magnetic field lines become twisted and bundled, which increases the intensity. The increased intensity creates a buoyancy effect that lifts the material in the magnetic bundle up to the solar surface, while centrifugal force moves it near the equator. This creates a bipolar field that appears as two sunspots, which explains why sunspots come in pairs on opposite sides of the equator.

The leading spot of this bipolar field tends to migrate to the equator while the trailing spot moves to the pole of the opposite hemisphere. Because this side is the magnetic opposite of that hemisphere, the result is a net reduction of the magnetic dipole and a tilt in the magnetic axis. This process of sunspot formation and migration continues until the dipole axis flips, which happens every 11 years. The full cycle of magnetic dipole rotation is thus 22 years.

The magnetic field of the spot still at the equator sometimes weakens. This can cause an influx of plasma that can form a magnetic bubble. If this bubble bursts, a Coronal Mass Ejection (CME) can occur. CMEs are a major component of high-speed solar wind and are a major event in space weather that can damage satellites and cause aurora on Earth.

The magnetic field of the solar surface can be observed remotely due to spectral line splitting caused by the Zeeman effect [313], [314]. In a strong magnetic field, the photons emitted from electron transitions are shifted due to the magnetic dipole moment associated with orbital angular momentum. Electrons with different spins are transitioning to and from different energy states compared to when no external magnetic field is present. The magnetic field vector (magnitude and direction) can be inferred from spectral measurements of plasma moving in a magnetic field.

6.2.1.2 The Solar Wind and Interplanetary Magnetic Field

The solar wind is constant stream of energetic particles that radiate from the Sun. This plasma consists mostly of electrons and protons [315], [316]. The temperature, density, and velocity (speed and direction) can vary significantly due to a number of factors. At 1 AU, the speed of the particles in the solar wind varies from ~320-710 km/s (468 km/s average), and the density ranges from 3-20 particles per cubic centimeter (8.7 /cm³ average).

As the distance to the Sun increases, the direction changes from radial to azimuthal flow. The Voyager spacecraft have studied the solar wind throughout their journeys beyond the Sun's magnetosphere [317]. Recent data has shown that near the heliopause, the flow of the solar wind is entirely azimuthal as particles make their way toward the tail of the heliosphere [318].

The supersonic speed of the solar wind is thought to be achieved by magnetohydrodynamic contraction and expansion at some critical radius from the Sun in a process similar to how supersonic flow is achieved in rocket engines [319]. Closer to the Sun (< 10 solar radii), the particles are subsonic and sub-Alfvénic, or below the Alfvén velocity, which is determined by the strength of the magnetic field and the plasma density. Between 10 and 20 solar radii, the solar wind cools and transitions from subsonic and sub-Alfvénic to supersonic. No spacecraft has yet sampled the solar wind inside this radius. The Alfvén velocity v_0 is defined as

$$v_0 = \frac{B}{\sqrt{\mu_0 \rho}} \quad (6-1)$$

where B is the magnetic field strength, μ_0 is the permeability of vacuum, and ρ is the total density of charged plasma particles. When $v_0 \ll c$, the velocity of the plasma is approximately the Alfvén velocity. However, as magnetic field strength increases or plasma density decreases, v_0 approaches the speed of light and the plasma behaves more like an ordinary electromagnetic wave.

The Wang-Sheeley-Argge (WSA) ENLIL Cone Model is based on data that comes from daily updates from several satellites around the solar system, including STEREO, the Solar Dynamics Observatory, the Solar and Heliospheric Observatory, and from remote observations of Mercury, Venus, and Mars as their atmospheres and surfaces interact with the solar wind [320], [321]. The WSA-ENLIL as it appeared during a CME event that hit Earth on July 14th, 2012 is shown in Figure 6-7.

Most observations of the solar wind are made from low latitudes and heliocentric orbital inclinations. This is because achieving higher inclinations is expensive due to the energy required to change inclination in a heliocentric orbit. There are currently no satellites in operation that are taking measurements of the solar wind to provide data for predictive models at locations other than the ecliptic plane at 1 AU.

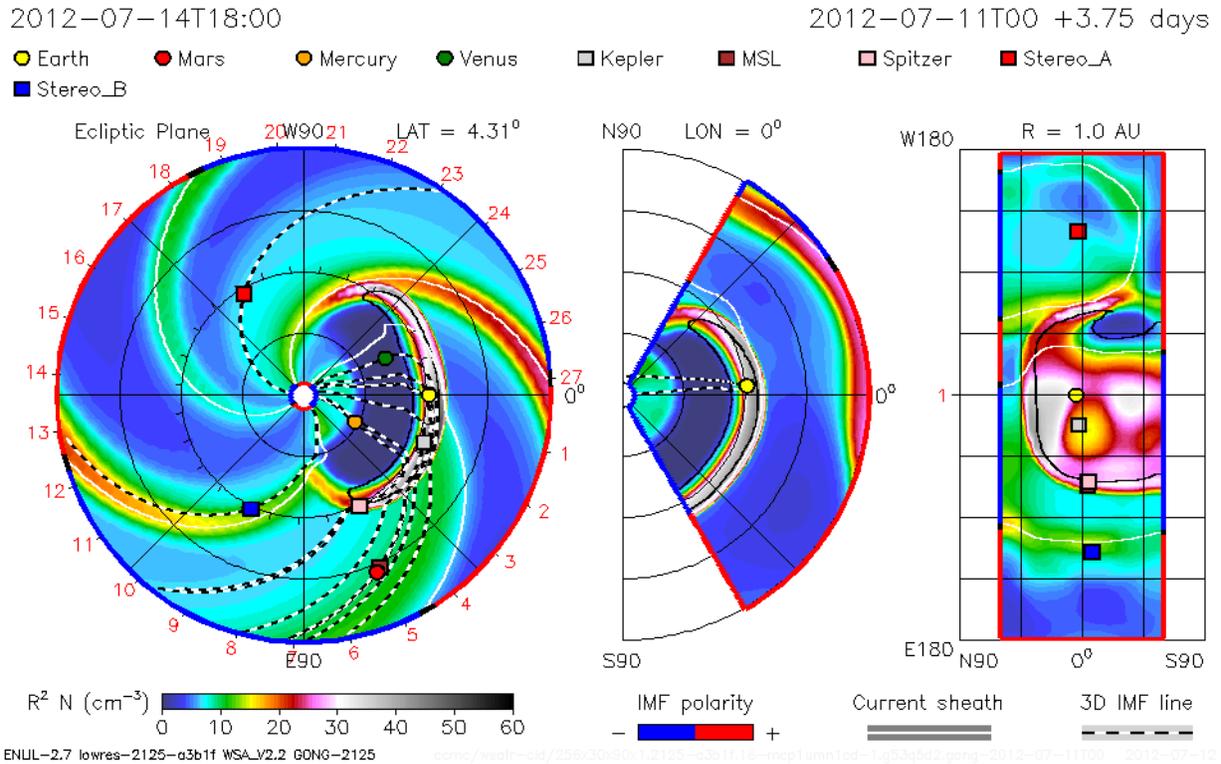


Figure 6-7: WSA-ENLIL Cone Model of the Solar Wind. On July 14th, 2012, a CME hit Earth, which is shown here (reprinted under Fair Use. Image credit: NASA Goddard Space Weather Laboratory).

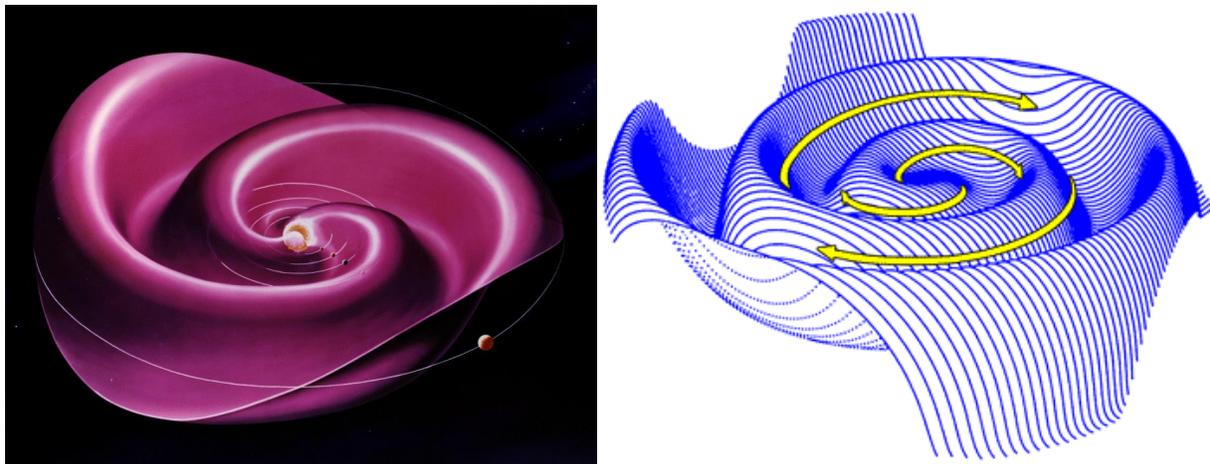


Figure 6-8: Artist renditions of the heliospheric current sheet (reprinted under Fair Use from [322])

ENLIL and other space weather models can be greatly improved with more data points and observations from around the solar system, especially at inclinations that are not the same as Earth's. In particular, in-situ observations of multiple, impacting CMEs and co-rotating interaction of fast and slow streams bound by shockwaves are desired [316].

6.2.1.3 The Heliospheric Current Sheet

A noteworthy component of the solar wind is the heliocentric current sheet, which is located on the plane where the polarity of the Sun’s magnetic field switches from north to south [323]. This current flowing through the plasma radiates from the magnetic equator of the Sun and changes significantly over the course of the solar cycle. During solar minimum (magnetic and rotational axes closely aligned), the sheet is relatively flat, but as the Sun transitions into solar maximum and the magnetic axis moves further from the rotational axis, the sheet twists into a spiral-like shape called a Parker spiral as the Sun rotates as shown in Figure 6-8.

As the radial speed of the solar wind decreases with increasing solar distance, the distance between waves of the spiral decreases. When a CME or other disruptive event occurs, parts of the sheet can collapse on each other. This can be seen in Figure 6-7; the white line representing the current sheet near Earth is disconnected as the CME hits it. The total current in the sheet is roughly 3×10^9 A, or 10^{-10} A/m². The magnetic field of the sheet near Earth is approximately 5×10^{-6} T, which is five orders of magnitude higher than what would be expected if the magnetic field would be if the Sun were a simple dipole.

The structure of the current sheet can be better known with more passes through the sheet. Only a DSS can make in-situ measurements of the magnetic field with enough distributed samples to have a fine enough resolution to see the finer structure of the sheet. This data can be used to understand the internal dynamics of the Sun in addition to better informing space weather models like ENLIL.

6.2.2 Similar Science Campaigns and Missions

This section summarizes some of the campaigns to study the Sun and Earth’s magnetic fields.

6.2.2.1 Campaigns to Study the Sun

There have been many satellites launched that study the Sun and other phenomena associated with the solar magnetosphere, solar wind, and other high-energy particles that contribute to interplanetary space weather. Some of the major and upcoming missions are summarized in Table 6-2 (also see [324]).

Table 6-2: A sample of the campaigns launched or planned to study solar phenomena.

Mission	Launch Date	Agency	Orbit/Location
Ulysses	10/6/1990	NASA/ESA	80.2° Inclined, Heliocentric
SOHO	12/2/1995	ESA/NASA	L1 Lagrange Point
Advanced Composition Explorer	8/25/1997	NASA	L1 Lagrange Point
RHESSI	2/5/2002	NASA/GSFC	LEO, 38°
Hinode	9/22/2006	JAXA	LEO, Sun-synch
STEREO	10/26/2006	NASA	Sun-centric, 1AU
IBEX	10/19/2008	NASA	Elliptic Earth Orbit
PICARD	6/15/2010	CNES	LEO, Sun-synch
Solar Dynamics Observatory	11/2/2011	NASA	GEO
DSCOVR	2/11/2015	NOAA	L1 Lagrange Point
Solar Orbiter	2017 (scheduled)	ESA	Eccentric, 60 Solar Radii
Solar Probe Plus	2018 (scheduled)	NASA/APL	Eccentric, 8.5 Solar Radii

The Ulysses spacecraft carried two radio/plasma antennas and several detectors for ions, electrons, and other particles, and to date is the only vehicle to ever travel in a high-inclination solar orbit [325]–[328]. It achieved an orbital inclination close to 80 degrees after a flyby of Jupiter, which gave it enough ΔV to

perform this maneuver. Because it spent the majority of its life far from the Sun, it was powered by a radioisotope thermoelectric generator (RTG).

The Solar and Heliospheric Observatory (SOHO) carried twelve instruments, each capable of individual and coordinated measurements with other instruments. [329]–[331]. The three primary objectives were to investigate the surface and outer layer of the Sun, observe the solar wind, and conduct helioseismology to probe the interior of the Sun.

The Advanced Composition Explorer was launched to better understand the composition of the solar wind and other matter in the interplanetary medium [332], [333]. It carries nine instrument packages, including several isotope and charged particle spectrometers and a magnetometer. It is still in operation and has recently been joined by DSCOVR at the L1 Lagrange point.

The Reuven Ramaty High Energy Solar Spectroscopic Imager (RHESSI) was launched to study the physics of solar flares [334], [335]. This spacecraft uses a unique technique of selectively blocking X-rays to construct X-ray images of the Sun.

The Hinode mission was launched to explore how the magnetic field of the Sun interacts with the corona, and the results have given improved understanding of solar prominences and other space weather events [336]. The spacecraft carries an optical telescope, an X-ray telescope, and an extreme ultraviolet imaging spectrograph.

The Solar Terrestrial RElations Observatory (STEREO) is a fractionated mission that stereoscopically images the Sun and other solar phenomena from multiple perspectives [308], [337]–[340]. Two spacecraft are in 1AU orbits, one trailing and one leading Earth, and they move further apart as time goes on. Each spacecraft carries four instrument packages that include cameras, particle detectors, and radio detectors.

The Interstellar Boundary Explorer (IBEX) is a small spacecraft that has overturned decades of assumptions about the boundary between the solar system and interstellar space. It discovered a ribbon of energetic neutral atom emission that was different in composition from the Sun, suggesting it comes from interstellar sources [341]. IBEX also failed to detect a bow shock [342], which overturns the textbook definition of the heliosphere, although this claim is still being debated in peer-reviewed literature [343]. More recent results show that the heliosphere has a four-lobed tail [344].

The Solar Dynamics Observatory (SDO) was launched in an effort to understand how the Sun affects life and society on Earth by studying many wavelengths simultaneously and at short time scales [345]–[347]. The spacecraft carries three primary instruments, the Heliospheric and Magnetic Imager, the Atmospheric Imaging Assembly, and the Extreme Ultraviolet Variability Experiment.

The PICARD satellite is dedicated to the measurement of total and spectral irradiance of the Sun to determine how the Sun functions internally and how it influences the climate of Earth [348], [349]. The satellite carries radiometers, photometers, and an imaging telescope.

The Deep Space Climate Observatory (DSCOVR) is a space weather monitoring satellite that will observe Earth's climate from L1 while also monitoring solar wind [350]. It carries three suites of instruments: a plasma magnetometer to study the solar wind, and a radiometer and an imaging camera to study Earth [351].

Solar Orbiter is a mission that will investigate the inner heliosphere and solar wind from the highest solar inclination achieved since Ulysses [352]. The final orbit will be within Mercury's perihelion distance. The mission will carry several instruments, including a solar wind analyzer, an energetic particle detector, a magnetometer, and an X-ray spectrometer. The most relevant instrument to HOBOCOP is Solar Orbiter's Polarimetric and Helioseismic Imager (PHI), which provides high-resolution measurements of the photospheric vector magnetic field. This information is of major interest to the PSS [353].

Solar Probe Plus is planned mission to study the solar corona with in-situ investigations [354]–[356]. This mission will use a series of planetary flybys of Earth and Venus to enter a highly eccentric orbit and swing as close as 8.5 solar radii from the surface of the Sun. This mission will examine structures of the magnetic field, determine what mechanisms accelerate energetic particles, and explore the plasma environment of the corona among other goals.

There are several observations that can be made about this list of Sun-studying missions. With the exception of STEREO, none of these spacecraft have leveraged DSS to achieve their missions or study solar phenomena simultaneously from spatially separated locations. With the exception of Ulysses, none of them have traveled into polar orbits. With the exception of Solar Orbiter and Solar Probe Plus, which have not been launched, none have been closer to the Sun than the L1 Lagrange point.

6.2.2.2 Campaigns to Study Earth's Magnetic Field with DSS

The HOBOCOP mission intends to study the Sun's magnetosphere with DSS in a similar way that other missions have studied Earth's magnetosphere. The Cluster II mission, which was launched in 2000 and recently given a mission extension to 2016, uses four satellites flying in a tetrahedral formation to study Earth's magnetic field [357]. This mission observed magnetohydrodynamic (MDH) turbulence in the plasma sheet as it passed through the magnetotail region [358]. These satellites fly in a ~135 degree (or 45 degree retrograde) inclination eccentric orbit.

The Time History of Events and Macroscale Interactions during Substorms (THEMIS) mission, launched in 2007, is a constellation of five satellites that studies energy releases in Earth's magnetosphere such as auroras and other substorms [359]. These spacecraft carry flux gate magnetometers, electrostatic analyzers, electric field instruments, and solid state telescopes. Two of the five satellites were moved to lunar orbit and are now designated a separate mission called Acceleration, Reconnection, Turbulence and Electrodynamics of the Moon's Interaction with the Sun (ARTEMIS).

The SWARM constellation, launched by ESA in November 2013, consists of three satellites in polar orbits, two flying in formation at 460 km and the third at 530 km [360]–[362]. These satellites carry vector field and absolute scalar magnetometers among other instruments and build off the experience from previous ESA campaigns of the past two decades.

The Magnetospheric Multiscale Mission, launched in March 2015, will build on the Cluster II mission and launch four satellites, also flying in a tetrahedral formation, into an eccentric orbit inclined at 28 degrees [363]. These four satellites each carry a suite of instruments including plasma analyzers, magnetometers, energetic particle detectors, and electric field instruments.

6.2.3 Performance Modeling in HOBOCOP

As space science missions move beyond LEO, modeling their performance characteristics may require additional considerations beyond what was discussed in Chapter 2. Elements that are specific to modeling performance in HOBOCOP are discussed in this section.

6.2.3.1 Electric Propulsion Systems

Unlike with chemical propulsion, where thrust times are low compared to the orbital period so that impulsive maneuvers can be assumed, electric propulsion requires long thrust times, and as a result, the astrodynamics calculations are not straightforward and are difficult to approximate. A satellite using chemical propulsion to raise its orbit does so by thrusting once in the initial orbit to create an elliptical transfer orbit that connects at apoapsis to the final orbit. However, orbit raising maneuvers using electric propulsion may take place over many orbital periods. As a result, an orbit transfer looks more like a long spiral than a single half ellipse [212].

For orbit lowering, such as in the case with HOBOCOP, as satellites fly closer to the Sun, the orbital trajectory could take the form of a spiral if the thrust time was on the order of years, but thrust time is dependent on the power available to the thrusters, the total mass of the spacecraft, and the specific impulse (I_{SP}) of the thruster. For finite burns, the assumption that the aphelion length of the final orbit will be equal to 1AU is not valid. However, an additional thrust at perihelion could raise the final orbit again if it is necessary for aphelion to be at 1AU for an Earth flyby.

6.2.3.2 Planetary Flybys

Many spacecraft have used planetary flybys as a means of gaining ΔV without using any fuel. Examples include Ulysses, MESSENGER, Voyager 1 and 2, New Horizons, and Cassini. During a flyby, the spacecraft imparts or steals momentum from a planet as it travels near it, and the closer it passes the more the momentum changes. Battin explains the mathematics of using patched conic sections to simplify the three-body problem into a series of two-body problems to calculate how much ΔV a spacecraft can potentially gain from a planetary flyby [209].

A close flyby of Earth could net a spacecraft a change in velocity up to 3.6 km/s, depending on how close the spacecraft gets to the atmosphere. This would change the inclination of a ½-year orbit with aphelion radius equal to 1AU by approximately 17 degrees. A series of annual flybys would put a spacecraft into a polar, heliocentric orbit. Jupiter, on the other hand, can impart a ΔV of nearly 11 km/s. At that distance from the Sun, that is enough ΔV for the spacecraft to achieve a polar orbit in one flyby.

6.2.3.3 Radiation Environment

Since this mission will take place outside the relative protection of Earth magnetic field, radiation plays an important component in determining the lifetime of individual satellites. There are a number of tools to estimate radiation dose received during an interplanetary spaceflight.

Radiation data from MSL's interplanetary flight from Earth to Mars shows that the radiation dose from the Sun and galactic cosmic radiation (GCR) is on the order of 500 $\mu\text{Gy/day}$ [364]. The data also shows that solar particle events (SPEs) contributed approximately 15 days' worth of radiation doses each. SPEs happened three times during the 9-month journey, or approximately four times per year. This data provides an estimate on the radiation dose a spacecraft may receive in interplanetary flight.

There are several advanced radiation modeling software suites, including CREME96 [365]–[367] and GEANT4 [368]. However, these suits require precise knowledge of the hardware being used and are not particularly appropriate for tradespace exploration. In contrast, the Space Environment, Effects, and Education System (SPENVIS), developed by the Belgian Institute for Space Aeronomy and ESA, allows for simple calculations of the radiation environment for a given circular heliocentric orbit [369]. This can be used to get worst-case estimates for radiation doses of spacecraft in eccentric orbits.

STK has a built-in Space Environment and Effects Tool (SEET), but it is not included with the license that the MIT SSL has.

6.2.3.4 Thermal Environment

There are a number of resources for modeling the thermal characteristics of a satellite. SMAD and SME provide information to calculate the approximate temperature of a spacecraft given different material properties including albedo, emissivity, and thermal conductivity [370], [371]. Radiative and conductive heat transfers are the primary sources of spacecraft thermal fluctuations. The calculated temperature of the spacecraft can also be checked against the known maximum operating and surviving temperatures of the spacecraft components being flown aboard satellites.

Another important factor in thermal modeling is that the quality of reflecting coatings degrades over time due to high-energy particles that impact the coating and change its chemical properties. This chemical transformation causes a physical transformation that alters the absorptivity and emissivity of the coating. While this degradation cannot be predicted precisely, end-of-life (EOL) estimates can be made so that the spacecraft still survives, or lower-quality materials can be used as substitutes for worst-case scenarios.

6.2.3.5 Instrument Modeling

HOBOCOP will consider a number of instruments, some of which have flown on CubeSats, and others which are larger with no known small satellite equivalent.

Magnetometers, Electrostatic Probes, and Energetic Particle Analyzers

Section 6.2.2 outlined many of the traditional missions to measure the Sun and other magnetic fields, but more recently there have been a number of programs that use CubeSats to study magnetic fields and energetic particles. These missions have used magnetometers, particle traps, Faraday cups, and other devices to make measurements similar to what HOBOCOP will be making.

Some of the options available for magnetometers used in small satellites are listed in Table 6-3. UCLA developed a digital fluxgate magnetometer for use in a lunar magnetometer network that is lightweight and has high accuracy [372], [373]. SSBV Aerospace & Technology Group developed a magnetometer that is commercially available for CubeSats [374]. Utah State University developed a magnetometer for their Dynamic Ionosphere Cubesat Experiment (DICE) mission [375]. These options are more massive compared to others that are typically embedded in other CubeSat components but have much higher dynamic range and are therefore more appropriate for science investigations rather than attitude control.

Table 6-3: Available CubeSat magnetometer options.

Manufacturer	Mass	Power	Resolution
UCLA	150 g	100 mW	1 pT, 50 Hz
SSBV	200 g	400 mW	13 nT, 10 Hz
Utah State (TAM)	< 100 g	110 mW	1.5 nT, 70 Hz

There are a number of instruments that can be used to study the solar wind. The simplest is a Faraday Cup, which catches charged particles and determines the number of ions or electrons hitting the cup to measure the plasma density. This and similar electrostatic probes are on Solar Probe Plus and can be easily shrunk down to fit on CubeSats [376]. However, unlike with magnetic field measurements, there is more than one measurement to make in the solar wind, so a variety of detectors must be considered for the payload of HOBOCOP.

Several options being developed for CubeSats to study energetic particles are listed in Table 6-4. The University of Colorado at Boulder’s Relativistic Electron and Proton Telescope integrated little experiment (REPTile) can measure both electrons and ions of varying intensities simultaneously, although it is the most massive of these options [377], [378]. The UC Berkeley SupraThermal Electrons, Ions and Neutrals (STEIN) satellite uses a particle detector that detects electrons between 4-100 keV and can separate between ions and neutral particles up to 20 keV [379], [380]. Utah State’s DICE satellite has electrostatic probes including a Langmuir probe in addition to its magnetometer payload [375], [381], [382]. UCLA’s ELFIN satellite can measure electrons and ions at the expected energies of the solar winds as well [383]–[385].

Table 6-4: Available CubeSat energetic particle analyzer options.

Manufacturer	Mass	Power	Notes
CU Boulder (REPTile)	1.25 kg	< 1 W	4 bins, >2.9 MeV electrons, 40 MeV max ions
UC Berkeley (STEIN)	<0.5 kg	550 mW	4-100 keV, separates ions and neutrons < 20 keV
Utah State (FFP/ILP)	<0.2 kg	100 mW	Small electrostatic probes
UCLA (ELFIN)	<1.5 kg	< 1.5 W	Electrons 50 keV to 4 MeV, Ions 50 keV to 500 keV

Vector Field Magnetographs

Unlike magnetometers, particle analyzers, and electrostatic probes, Vector Field Magnetographs (VFM) are not common among small satellites. A VFM measures spectral line splitting due to the Zeeman Effect with a telescope and spectrograph combination. Both the longitudinal (line-of-sight) and transverse (in-plane) directions are measured separately. The line-of-sight wavelength is red- or blue-shifted, because photons emitted from atoms with different polarizations gain or lose energy depending on their orientation to the magnetic field. The amount of Doppler shifting increases with increasing magnetic field. The relative shifts and line spreading of a spectral band give the transverse direction of motion. These two components are combined to calculate the magnetic field vector.

However, there is symmetry in the transverse measurement, so there is 180-degree ambiguity in the transverse direction. This means some other information must be used to constrain the direction. Usually this is done by simply observing other parts of the magnetic field to deduce the direction of many areas at once. Typically, spectral lines of emissions from iron atoms are observed [386], [387]. Because the instrument uses both a telescope and a spectrograph, it usually does not fit within the mass budget of CubeSats, and no known models exist for small satellites.

The Polarimetric and Helioseismic Imager (SO/PHI), the VFM on Solar Orbiter, has a mass of 35 kg, requires 30 W of power, and generates 20 kbps of data [352], [353], [388], [389]. It was designed with heritage from the Sunrise telescope’s Imaging Magnetograph eXperiment (IMaX) [390]. While its mass is relatively small for a monolithic mission, its data rate is the critical component due to limited windows to

communicate with the spacecraft during its mission. For this reason, it cannot operate throughout the mission and must be selective with its data return.

PHI consists of two telescopes: a Full Disk Telescope with 2 degree FOV, 3.5 arcsecond resolution, and 17 mm aperture diameter; and a High Resolution Telescope with 16 arcminute FOV, 0.5 arcsecond resolution, and 140 mm aperture diameter. The spectrograph package can measure the magnetic field with up to 15 Ga resolution. The orbit of Solar Orbiter is not circular, however, so the instrument is built with additional components to handle the thermal environment. These elements may or may not be needed depending on what orbit a VFM is sent to in HOBOCOP. Although smaller versions of this instrument are not available to build a parametric size model, the sizing equations in SMAD can be used to estimate the mass of a scaled version of PHI for use in HOBOCOP (see Section 2.4.2.1).

6.3 The RSC Model for the HOBOCOP Mission

The Responsive Systems Comparison model for HOBOCOP, like Exoplanet's, has gone through several iterations over the past two years as stakeholders, stakeholder needs, and the core concept of the mission have changed. Various efforts to characterize and optimize the communications side of the mission came long before the scientific stakeholders were consulted, which resulted in most of the initial work being thrown out or otherwise heavily modified before becoming a part of the version presented here. This section will trace the development process shown in Figure 4-2, starting with value modeling and performance modeling before moving into tradespace exploration and EEA.

6.3.1 HOBOCOP RSC Phase 1: Value Modeling

One of the advantages of value modeling in HOBOCOP compared to ExoplanetSat is that the goals of this mission are related to modeling of known phenomena. This is definitely a case where MAUT is an appropriate VCDM to apply to the value model of this mission because there is little chance that a constructed value function exists (like it did in ExoplanetSat). The PSS has opinions on the quality of the data products based on independent attributes that can be elicited from a typical stakeholder interview.

On the other hand, when considering the needs of secondary stakeholders — such as a private corporation operating a network of satellites — MAUT is less appropriate than economics-based VCDMs like NPV. Unfortunately, explicitly calculating the monetary components of many possible futures and customers is beyond the scope of this work. System attributes can still be identified and constructed within a MAUT framework in this stage, however, and in the absence of a revenue model to assess benefits, MAUT can still be used to assess the utility of such a constellation.

6.3.1.1 Value Modeling Process 1: Value-Driving Context Definition

Although this is a concept study generated entirely by the author with no budgeting authority whatsoever, there are a number of possible stakeholders and several different goals that can be achieved given varying levels of resources available to those stakeholders. Stakeholders include scientists studying the Sun and heliosphere. Several scientists were approached to elicit opinions of the system. The author intended to gather information about how a system could achieve these goals, and which attributes drive the value of designs so design alternatives can be compared against each other.

Other possible stakeholders include those that wish to build an interplanetary satellite relay network for federated satellite systems. Their goals would be built around maximizing revenue, but without a

monetary revenue function or other known customers, a MAU value function can be used to evaluate a system they would build. Developing such a stakeholder value model is beyond the scope of this case study, but the advantages of a satellite network will be explored in further detail.

6.3.1.2 Value Modeling Process 2: Value-Driven Design Formulation

HOBOCOP will address three scientific goals: characterize the magnetic field on the surface of the Sun, characterize the magnetic field in interplanetary space, and characterize the solar wind. Here, the value models were elicited from experts in the field who could serve as stakeholders on this mission.

Science Goal 1: Characterize Surface Magnetic Field

Goal #1 involves the most complex of the three instruments being considered for the HOBOCOP mission. For this reason, the attributes that drive value in this instrument are more important to understand. Over the course of several interviews, the PSS provided insight into how value was delivered in the form of multiple VFMs. The three attributes of the instrument data that are important to her are the spatial, temporal, and magnetic resolution. The minimum acceptable level, maximum perceived level, and relative weight of the attributes as they were elicited from an interview are summarized in Table 6-5. The attribute utility functions that were also elicited are shown in Figure 6-9.

Table 6-5: Single Attribute Utility weights for the Vector Field Magnetometer instrument.

VFM Attributes (Units)	Minimum	Maximum	Weight (k_{1ij})
Spatial Resolution (km)	7000	100	0.50
Temporal Resolution (min)	30	0.133	0.10
Magnetic Resolution (Ga)	50	1	0.60

From these utility curves, we can see that spatial resolution below 7,000 km is desired so that the finer structure of the particle events can be observed (note the x-axis of all three graphs is logarithmic), but structure beneath 100 km is not important. Sampling a few seconds apart is also highly desired in order to capture events that happen quickly. Sampling more than half an hour apart is unacceptable due to how quickly the magnetic field changes. Finally, the magnetic resolution must be at least 50 Ga, otherwise the magnetic structure and direction of the magnetic vector will be too muddled to be of any value.

In addition to the attributes of the instrument, the other attributes related to the science goal was elicited. The PSS noted that viewing inclination and simultaneous viewing points were highly desired. Figure 6-10 shows the utility curves that were elicited for the science goal attributes one level above the instrument attributes on the multi-tiered MAUT structure (see Figure 4-5). From these curves it can be deduced that although a low inclination is acceptable, value is not strongly derived until the assets can see a good view of the poles of the Sun. The threshold shape of the curve is also due to the fact that a satellite does not need to be *directly* above the poles to view them, but a better picture of an entire polar hemisphere results from higher inclinations.

Finally, the relative weights of the science goal attributes were elicited. These are shown in Table 6-6. There are several observations that can be made here. First, notice that the sum of the weight is only 0.3, far below 1. This shows that the PSS is relatively risk-averse with her selections or that she is uncompromising in her desire to have a combination of attributes. This means that more value can be delivered by satisfying all attributes compared to satisfying one attribute in particular.

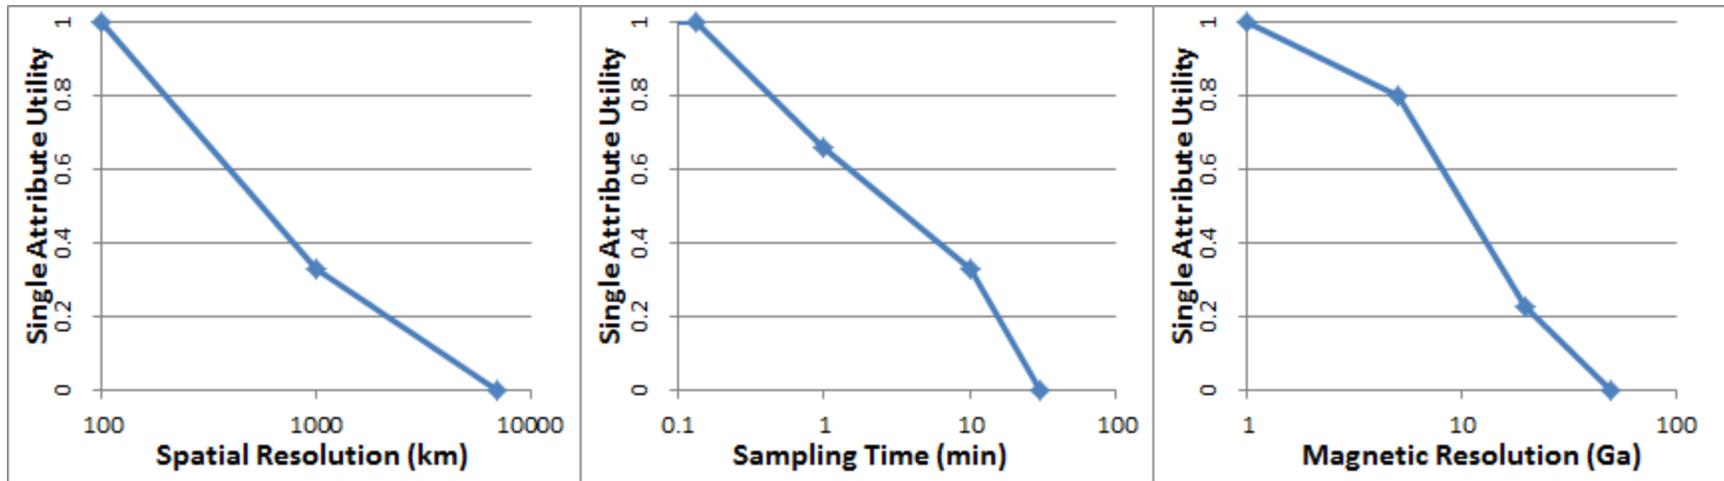


Figure 6-9: Single Attribute Utility functions for the Vector Field Magnetometer instrument.

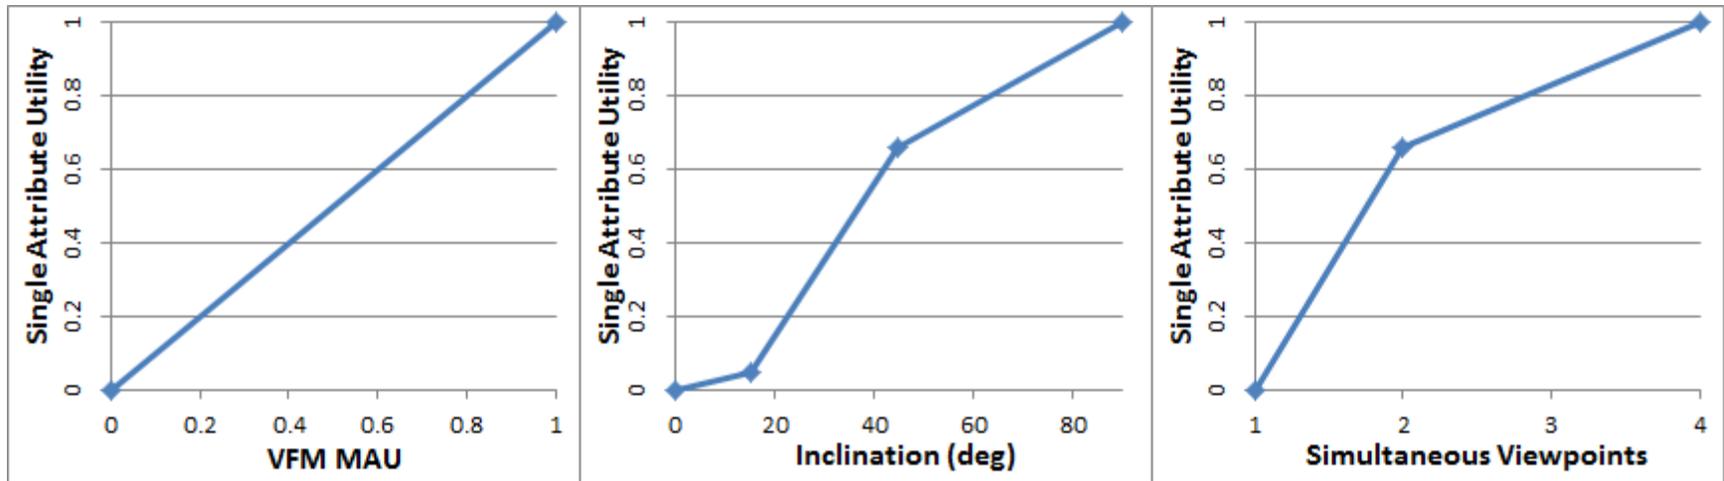


Figure 6-10: Single Attribute Utility functions of HOB COP Science Goal #1

Table 6-6: Single Attribute Utility weights for HOBOCOP Science Goal #1.

Science Goal Attributes (Units)	Minimum	Maximum	Weights (k_i)
VFM MAU	0	1	0.20
Inclination (degrees)	0	90	0.07
Viewing Angles	1	4	0.03

The relatively low weighting on the final science goal attribute was unexpected. This shows that the number of viewing angles is the least important science goal attribute of all and that the PSS’s preference is weighted heavily towards a robust instrument in a similar fashion to a monolithic system. However, MAUT still shows that even with the best VFM orbiting at the highest inclination, the maximum utility that a monolithic system can deliver is $MAU = 0.5816$.

Science Goal 2: Characterize Solar Wind

There are many qualitative factors that describe a system that better characterizes the solar wind in the interplanetary medium. The MIT physics professor was interviewed to solicit his opinions about instruments and locations of spacecraft to meet this goal.

Having many spacecraft is clearly better so that more locations in the solar wind can be sampled simultaneously. However, “number of spacecraft” is more akin to a design variable; what a scientist really wants is *simultaneous sampling resolution*. This can be defined as simultaneous samples per unit area or per unit volume. The area or volume sampled is also very important; since solar wind measurements come only from points along the ecliptic plane at 1 AU, varying sample distance from the Sun and from the ecliptic plane is desired. Although the third dimension could be wrapped into the simultaneous sampling resolution as samples per volume, it is better in this case to consider the area and height separately; a complex, well-populated constellation operating entirely in the ecliptic plane would have no sampling resolution per unit *volume* at all, but a single satellite with 1 degree inclination would have finite sampling volume and score higher on this very important attribute. For this reason, *vertical sampling range* more accurately describes this characteristic.

Another important modeling caveat is to determine up front what sampling area is being used to measure resolution. If the area is simply the annulus where the spacecraft fly, a set of satellites in circular orbit would have infinite sampling resolution because they are only tracing out a line. A constellation that sends a satellite further out into the solar system instead of closer to the Sun would lower the sampling resolution of an entire constellation. This problem is similar to that of many coverage metrics, where the statistical metrics can mask or hide value in some systems and make inferior systems seem superior.

Two more critical qualitative descriptors are how close to the Sun samples are taken and the maximum distance between samples. Ultimately, it would be of maximum value to sample the solar wind close to the Sun to observe it at sub-Alfvénic velocities, but any measurements closer to the Sun than we have now are still valuable to the scientific community. Additionally, the range of samples taken, or the difference between maximum and minimum sample points, is important to see the evolution of the solar wind as it transitions from primarily radial motion to primarily azimuthal motion.

Another qualitative aspect that should be captured is simultaneous azimuthal measurements. Two satellites along the same magnetic field line would observe the same parcel of the solar wind, and two satellites along the same radial line would observe how much azimuthal acceleration is affecting the solar

wind. This is a difficult aspect to characterize quantitatively, but can be captured with threshold returns in simultaneous sampling resolution because at higher resolution the satellites will begin to natural overlap in their orbits.

Furthermore, there is great desire for multiple satellites to observe from opposite sides of the Sun simultaneously. Simultaneous close approaches are of the greatest value here because close approaches sample a large orbital angle in a short period of time, and having a complete set of measurements from all around the Sun is of far greater value than only having one arc.

Finally, the instrument itself is important. The instrument should be able to sample a wide range of particle energies and types. The interview with the physics professor showed that there were many aspects to the instrument that were important, but the most critical was the maximum resolvable energy. The complete list of attributes, their minimum acceptable and maximum perceived levels, and the attribute weights are shown in Table 6-7. The utility functions for each attribute are shown in Figure 6-11.

Table 6-7: Single Attribute Utility weights for HOBOCOP Science Goal #2.

Science Goal Attributes	Minimum	Maximum	Weights (k_{2i})
Minimum Perihelion Distance (AU)	0.046	0.6	0.15
Radial Sampling Range (AU)	0.5	1	0.05
Radial Sampling Resolution (samples/AU ²)	2	50	0.35
Vertical Sampling Range (AU)	0.2	1	0.05
Simultaneous Perihelion Flybys	1	5	0.2
Particle Energy Range	<400 keV	<5 MeV	0.1

Science Goal 3: Characterize Heliospheric Current Sheet

The qualitative attributes that help characterize the heliospheric current sheet are very similar to the previous goal: more samples and higher sampling resolution are better, bigger range is better, more height samples are better, simultaneous flybys are better, and closer measurements are better.

One of the key differences between the two goals is that although higher orbital inclination is better to obtain a larger vertical sampling range, the heliospheric current sheet stays closer to the ecliptic plane, while the solar wind radiates in all directions. A polar orbit is undesirable for observing the current sheet because it will only ever pass through the sheet twice per orbit; a lower inclination orbit will have varying numbers of passes through the sheet depending on what part of the solar cycle is being observed.

This means there are both *increasing* and *decreasing* returns on orbital inclination, something that MAUT does not handle well without complicated combinations of attributes or double attributes because all utility functions must be strictly monotonic. This could be avoided by simply constraining the design space, but doing so would hurt the achievable utility of the other goals. A double attribute will be used in this case to solve this problem.

The complete list of attributes for characterizing the heliospheric current sheet, the minimum acceptable and maximum perceived levels, and the attribute weights are shown in Table 6-8. The utility functions for each attribute are shown in Figure 6-12. The attribute functions for Minimum and Maximum Vertical Sampling Range have been combined into one graph on the bottom left to save space, but they are separate functions that are both monotonic.

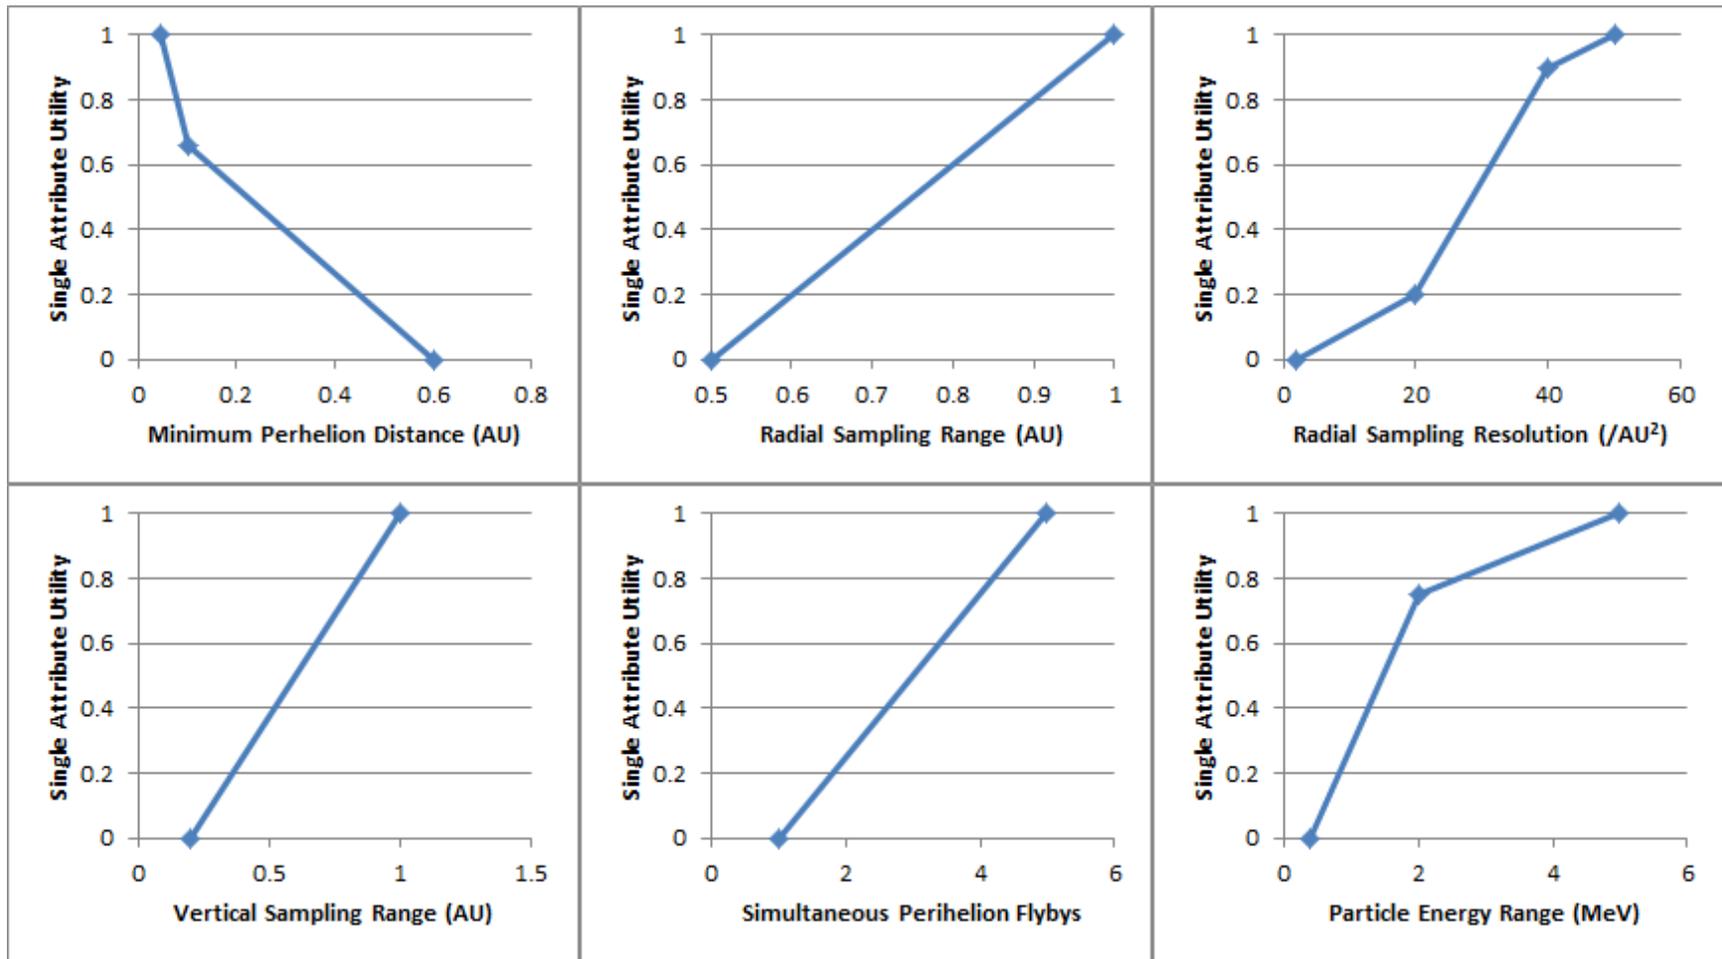


Figure 6-11: Single Attribute Utility functions of HOBOP Science Goal #2

Table 6-8: Single Attribute Utility weights for HOBOCOP Science Goal #3.

Science Goal Attributes	Minimum	Maximum	Weights (k_{3i})
Minimum Perihelion Distance (AU)	0.1	0.3	0.1
Radial Sampling Range (AU)	0.5	2	0.05
Radial Sampling Resolution (samples/AU²)	3	50	0.2
Minimum Vertical Sampling Range (AU)	0.2	0.5	0.05
Maximum Vertical Sampling Range (AU)	0.6	1	0.03
Simultaneous Perihelion Flybys	2	8	0.1
Magnetometer Resolution (nT)	30	0.5	0.05

Total Mission MAU

Two key aspects left out of the previous discussion were the total mission lifetime and the age of data. HOBOCOP assumes that the data is time-insensitive, so it can be retrieved whenever it is convenient or feasible. It is not required to be downlinked daily, and very few commands need to be sent to the spacecraft due to their high autonomy and the simplicity of their mission operations once they are in orbit.

However, mission lifetime is important because these observations hope to capture variations of the phenomena being observed over the course of the solar cycle. Both the PSS and the MIT professor independently agreed that this mission should be launched near or immediately after solar minimum and observe the Sun as it transitions to solar maximum. While the full cycle is from minimum to minimum is eleven years, only the upward swing would be required. This means the mission would need to operate for a minimum of two and a half years, but up to a maximum of seven. This primary mission attribute curve is shown in Figure 6-13.

With these three goals, there is an incredible amount of synergy, especially between Goals #2 and #3. This is clear because the sum of the mission MAU weights shown in Table 6-9 is less than one, meaning the value of the mission increases significantly by addressing all three goals instead of optimizing toward completely satisfying one of them.

Table 6-9: Science Goal MAU weights for HOBOCOP

Science Goals	MAU Weights (k_i)
Science Goal #1 MAU: Characterize Photosphere Magnetic Field	0.25
Science Goal #2 MAU: Characterize the Solar Wind	0.2
Science Goal #3 MAU: Characterize the Heliospheric Current Sheet	0.2

One of the fundamental assumptions of MAUT is that all attributes are independent of one another. A scientist may see a significant relationship between Goals #2 and #3, meaning they are not independent. If this is actually the case, an alternative MAUT hierarchy structure should be considered that addresses this. This structure is shown in Figure 6-14, and the attribute weights on these levels is shown in Table 6-10

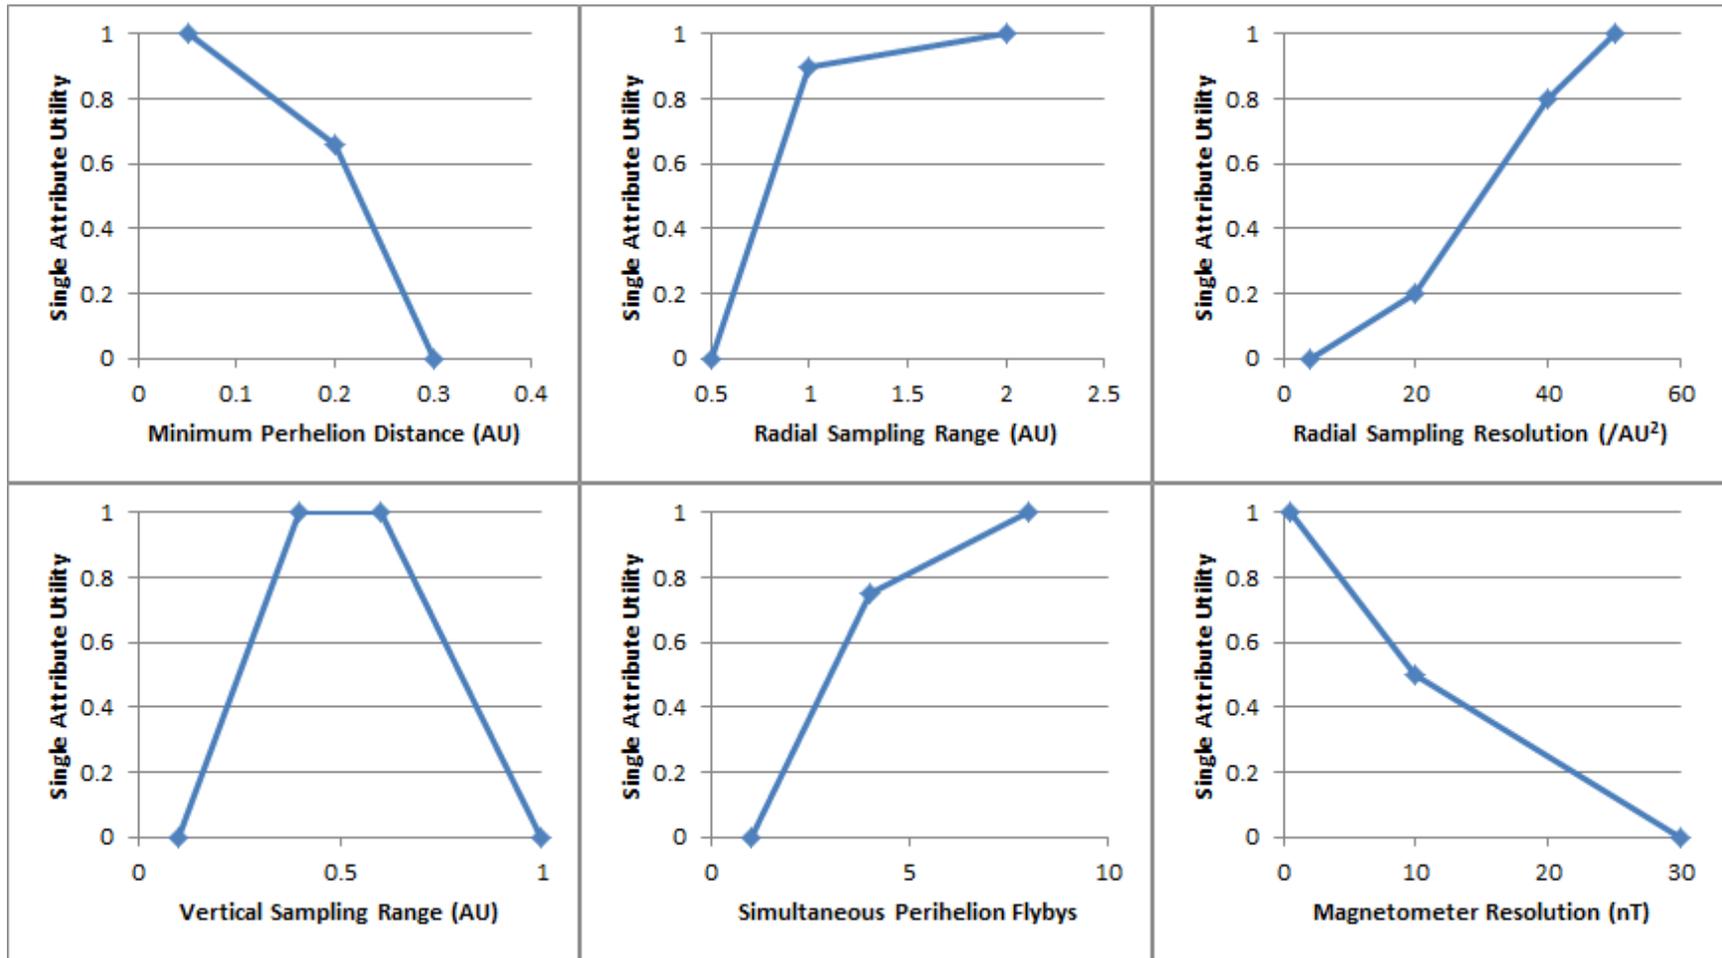


Figure 6-12: Single Attribute Utility functions of HOBOP Science Goal #3

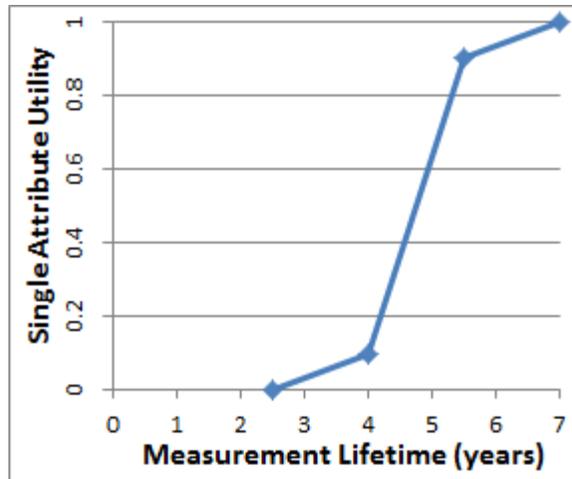


Figure 6-13: Single Attribute Utility function of HOBOCOP Mission Lifetime.

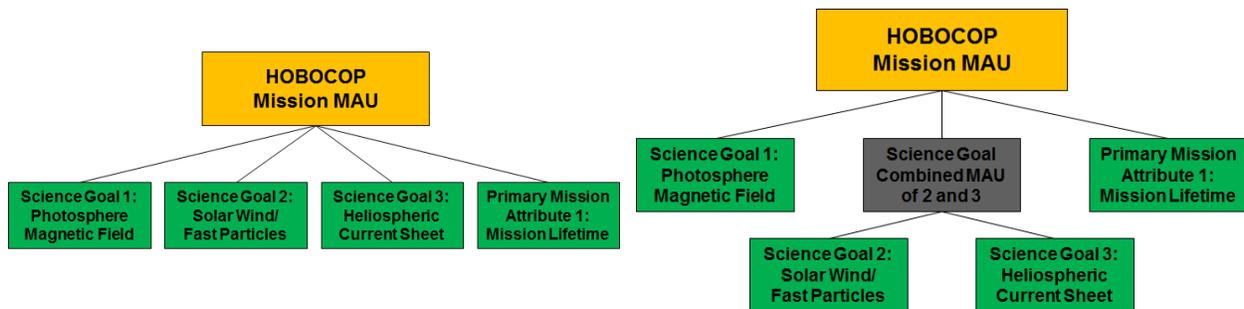


Figure 6-14: (Left) Primary MAUT structure similar to what is shown in Figure 4-5. (Right) Alternative MAUT structure that better represents an enhanced relative synergy between Science Goals #2 and #3.

Table 6-10: Alternative science goal MAU weight for HOBOCOP with split synergy.

Science Goals	Sub MAU Weights	MAU Weights (k_i)
Goal #1: Characterize Photosphere Magnetic Field	--	0.25
Goal #2: Characterize the Solar Wind	0.25	0.4
Goal #3: Characterize the Heliospheric Current Sheet	0.25	

6.3.1.3 Value Modeling Process 3: Epoch Characterization

In this phase of value modeling, context variables and alternative stakeholder expectations were explored, and alternative levels for parameters with high uncertainty were created. These alternative contexts and value models are normally used in the final phase of the RSC method, but in this case study they represent important alternatives that can be considered even in the tradespace exploration phase.

Performance Context Variables

While none of the technologies presented in this mission are science fiction, the immaturity of some of them make it difficult to be certain that the performance model accurately predicts the necessary mass and cost of any point design. In particular, the mass of the SEP subsystem can vary wildly depending on the numbers assumed for the area-to-mass ratio of the solar panels, the yield stress-to-weight ratio of the propellant tank, and the mass-to-power ratio of the power production unit. Due the extreme sensitivity of

the performance of a design to these parameters, it is important to understand how they vary. (Why they are sensitive will be discussed in Section 6.3.2.)

Table 6-11 lists several of the technology assumptions that need to be made to model this mission and two sets of values the parameters showing these assumptions can make. The assumptions that result in a reduced mass of the system are in the “Cutting Edge/Optimistic Technology Epoch.” These numbers are based either on theoretical limits (such as the mass of the propellant tank) or recent data solicited from experts in the fields (such as the power production unit mass as a function of power). The assumptions that result in higher mass are in the “SMAD/Pessimistic Technology Epoch.” The values of these parameters come from established literature and parametric models available in SMAD that do not necessarily represent some of the theoretical best performance that can be achieved by these parts.

Table 6-11: Comparison of parameters used in both the Optimistic context and the SMAD-Based context for spacecraft mass estimations.

Component Mass or Ratio	Cutting Edge/Optimistic Epoch	SMAD/ Pessimistic Epoch
Propellant Tank	Stress-limited, safety factor = 2	SMAD Parametric Model
Solar Panel W/kg	500 W/kg (Spectrolab/Emcore)	200 W/kg (SMAD)
Power Production Unit	300 W/kg (Lozano)	200 W/kg
Heat Shield Coating	$\alpha = 0.08, \epsilon = 0.79$ (Quartz/Silver)	$\alpha = 0.20, \epsilon = 0.75$ (Degraded)
Structural Mass Fraction	7%	12%
Solar Panel Efficiency	0.30	0.28.3% (Spectrolab)
Electric Thruster Mass	0.250 kg	0.500 kg
Electric Thruster Efficiency	90%	70%
Margin	20%	20%
System Noise Temp (Radio)	Low (80/30/1000 K)	High (300/150/2000 K)

Other important unknowns include the availability of a network of federated satellites within the inner solar system and the communications technology used by the satellites. For the first, there are three possible options: do not use relay communications at all and only communicate with Earth; send additional satellites to serve as relays on the mission’s budget; or communicate with a federated network at regular intervals in addition to Earth.

For the communications technology, a system could use the X-band, the Ka-band, or not use radio at all and instead use optical communication. The choice between the X-band and Ka-band represents the next step small step in communications technology development; the choice between the Ka-band and optical communications represents the next giant leap.

Cost Context Variables

At this time, no additional cost variables are being considered for modeling costs in HOBOCOP.

Alternative Stakeholder Expectations

A graduate student was also interviewed to give an alternative viewpoint to the PSS’s expectations for the first goal. While the PSS had a very low risk tolerance for everything except for the attributes associated with the instrument (as demonstrated by the relative weights in Table 6-6), the graduate student was much more willing to take risks on the instrument performance to increase the sampling inclination and simultaneous viewpoints of the Sun. As a result, the weights on those attributes are much higher. One could consider this value model a case of eliciting the value model of the instrument from an

instrumentation scientist and then eliciting the value model of the science goal itself from a scientist higher on the MAUT hierarchy. The two sets of science goal attribute weights are shown in Table 6-12, and a graphical representation of this effect is shown in Figure 6-15.

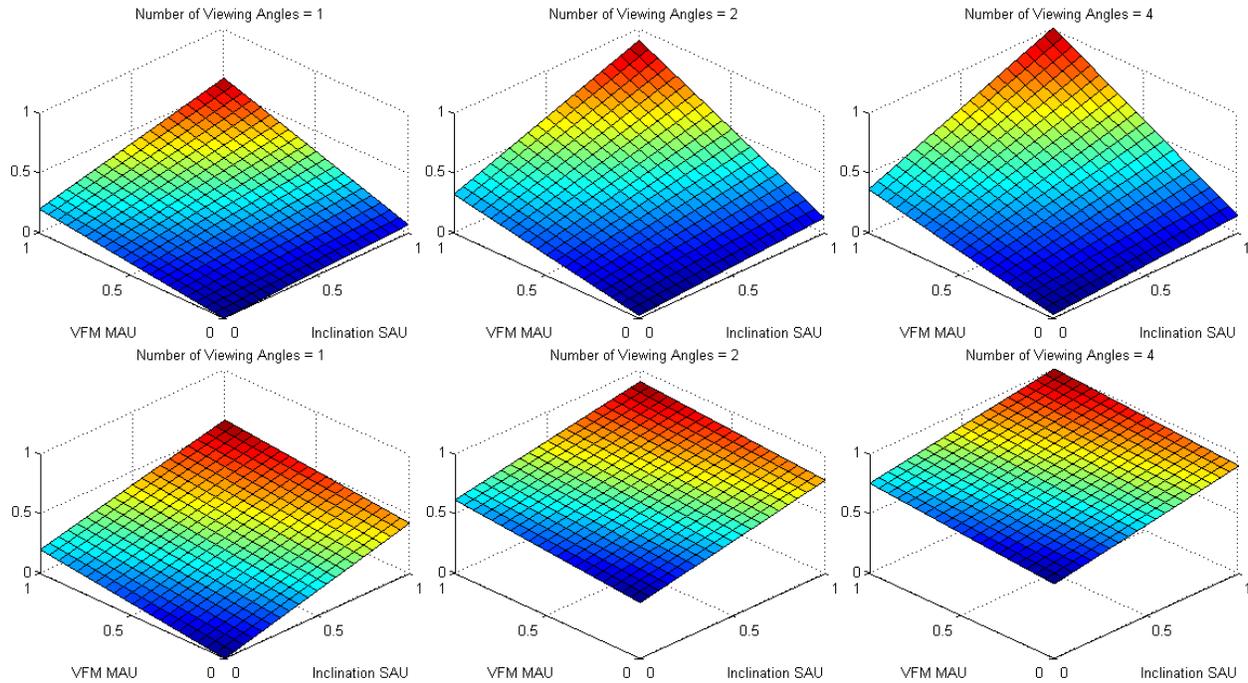


Figure 6-15: Comparison of two different weighting of attributes and their effect on MAU of Science Goal #1.
(Top) PSS weightings and how utility increases as number of viewing angles increasing from left to right.
(Bottom) Alternative weightings showing higher weight on inclination angle and number of viewing angles.

Table 6-12: Comparison of differing stakeholder opinion on attribute weights of Science Goal #1.

Science Goal Attributes	PSS Weights	Alternative Weights
VFM MAU	0.20	0.20
Inclination	0.07	0.42
Viewing Angles	0.03	0.62

One of the largest assumptions being made up front is that the data being returned by the HOBO satellites is *time-insensitive*, meaning there is little demand to retrieve the data in a timely manner. This assumption could change if the space weather community suddenly had significant interest in this project. Scientists running space weather models are used to receiving daily data downlinked from satellites observing the Sun, and there is especially high demand for data during a CME or other significant solar event. If this were the case, the mission would be practically infeasible with standard radio technology (as will be shown in the performance model) because of the limitations created by requiring a massive communications system. However, if a network of relay satellites is present, or if optical communication is an option, then it would definitely be feasible for the satellites to communicate more than once a year each and deliver data back to Earth on a regular basis.

One could construct a new value model entirely for this situation, but no experts were available for elicitation of a value model on the time sensitivity of the data. It is simpler to require that data be feasibly

returned at any point except during periods when the satellite is close to the Sun in the absence of a formal stakeholder elicitation. In this scenario, designs that are unable to communicate regularly are simply unacceptable.

Finally, an additional value ranking for the combination of the goals was also constructed to evaluate how sensitive the result is to synergy as shown in Figure 6-14 and Table 6-10. There are two choices for each of the three value variables being considered, giving a total of eight possible values models that can be used to determine the science value delivered by HOBOCOP.

Epoch Characterization Summary

The contexts variables and stakeholder value models that form the epoch space of HOBOCOP are summarized in Table 6-13. Only the stakeholder value model and relay network availability can change during the mission lifecycle, so rules will need to be in place during Era Analysis to ensure these stay fixed. There are 18 possible contexts and 8 stakeholder value models for a total of 144 possible epochs.

Table 6-13: Context variables and their values for HOBOCOP

Epoch Variables	Possible Values
Stakeholder Value Model	[Goal 1(A/B), MAU(A/B), Time Sensitivity(A/B)]
Relay Network Availability	[None, Self-Built, Commercial]
Technology Mass Modeling	[Optimistic, SMAD-based]
Communication Technology	[X-Band, Ka-Band, Optical]

While there are many assumptions that need to be made that could lead a designer to incorrectly lead a decision maker to a less-than-desirable outcome for a mission, EEA helps mitigate the uncertainty by allowing the designer to see the sensitivities to certain factors up front as context variables. This way, as new information accumulates during the design process, those uncertainties can be eliminated and decisions can be made on the realistic datasets. The assumptions that go into the performance model will be made as explicit as possible, and a complete list of all modeling parameters can be found in Appendix B. All relevant code is available on the supplementary DVD.

6.3.2 HOBOCOP RSC Phase 2: Performance Modeling

In the second phase of RSC, the metrics of satisfaction flow down to the design variables that drive the attributes. The performance model of HOBOCOP consists of several modules that have been combined under a rough system optimizer to minimize the mass of the individual designs. This is important because of the technical constraints involved with close orbits of the Sun and the large ΔV s involved in this mission. A spacecraft that is heavy or underpowered will not be able to achieve the desired orbit before it must jettison its propulsion module because of solar panel thermal constraints. The mass budget may increase more than three grams for every gram added to the satellite due to fuel and power requirements.

6.3.2.1 Performance Modeling Process 1: Identifying Architecture and Design Variables

The original HOBOCOP concept used QFD to build a set of critical variables based on the needs of a communications network, rather than the needs of a science mission. After the stakeholder interviews and the scientific needs were identified, the variables that changed the value of the SAUs described in Section 6.3.1 were added to the tradespace. Most of these already existed in previous versions, but were applicable to eccentric relay satellites while the HOBOSAT satellites were in circular orbits; after the

stakeholder interviews, it was apparent that the HOBOS needed to be in eccentric orbits while the relays could be in whatever orbit worked best.

The literature review in Section 6.2 gave examples of instruments that could be used to satisfy the science goals. Given how small the FPAs and magnetometers are compared to the VFMs, only five payload combinations were chosen (rather than all eight possible combinations) since the relative mass difference between a payload with only a VFM versus a payload with all three options is small. This also leaves out the possibility of a HOBOS having no payload at all, which has the exact same scientific value of launching a pile of cash into space. These five payload options are shown in Table 6-14.

Table 6-14: Payload Packages available for HOBOS satellites.

Payload Option	Payload Package
1	Magnetometer Only
2	Fast Particle Analyzer Only
3	Magnetometer AND Fast Particle Analyzer
4	Vector Field Magnetometer Only
5	All Three Instruments (Magnetometer, FPA, and VFM)

The complete list of architectural and design variables is shown in Table 6-15. The primary architectural variables are the first five in the table. For purposes of this work, an “arc” is defined as a cluster of satellites that reach orbital perihelion simultaneously. More arcs means higher frequency of perihelion passes. The number of petals per arc corresponds to the number of azimuthal angles from which a given arc observes. It also corresponds to the inverse of the orbital period of the petals; if there are four petals, the HOBOS are in ¼-year orbits.

Table 6-15: Architectural and design variables and their discrete values for HOBOS

Design Variable	Discrete Values
Number of Arcs (sets of Petals)	[1:20] (for payloads 1, 2, 3), [1:10] (payloads 4, 5)
Number of Petals per Arc	[2, 3, 4] (for payloads 1, 2, 3), [2] (payloads 4, 5)
Number of HOBOS per Petal	[1, 3, 5] (depends on Z-Thrust Angle)
Payload Option	[1, 2, 3, 4, 5] (see Table 6-14)
Z-Thrust Angle (Payload 1, 2, and 3 only)	[0, 10, 20, 30, 40] (2 petals), [0, 5, 10, 15] (> 2 petals)
FPA Type (Payload 2, 4, and 5 only)	[FFP/ILP, STEIN, REPTile]
VFM Diameter (Payload 4 and 5 only)	[40 mm, 70 mm, 100 mm]
Target Inclination (Payload 4 and 5 only)	[0, 30, 60, 90]
Propulsion Solar Panel Area	Varies sizes and steps between 0.2 m ² and 8 m ²

A brief interlude is warranted to explain the discrete choices for these variables. The total number of HOBOS is a function of the number of arcs, the number of petals per arc, and the number of HOBOS per petal. The number of arcs is limited only by the mass and volume constraints of the launch vehicle. The number of petals per arc is the inverse of the orbital period; in order to achieve simultaneous perihelion sweeps, these periods must be integer fractions. The maximum level of this design variable is four because no designs were found that could achieve 1/5-year period orbits.

The number of HOBOS per petal is dependent on the inclination of the HOBOS within the petal. There is no point in sending two HOBOS to the exact same orbit, but if they are at corresponding positive and negative inclinations in the same orbit as viewed from above, this provides more observations from

differing stellar latitudes. The bigger the difference in inclination, the more effective multiple viewpoints from the same azimuth become. This is illustrated in Figure 6-16. Finally, Payload Option determines which payload package will be onboard the individual satellite.

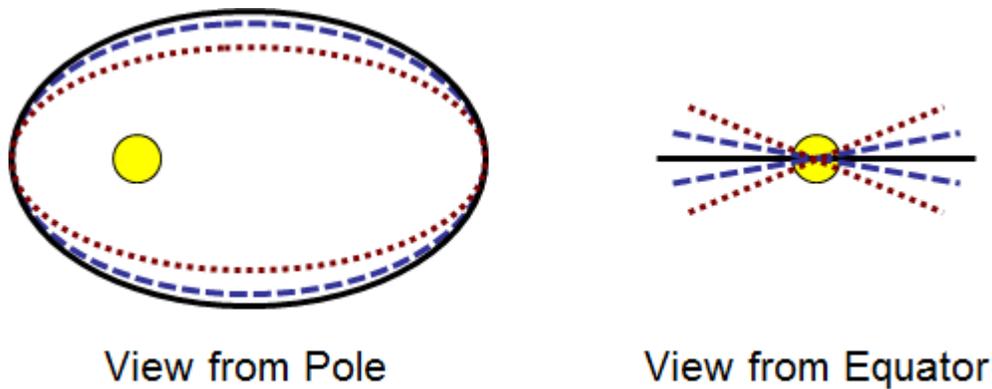


Figure 6-16: Illustration of multiple satellites within the same orbital petal with different inclinations.

For the satellites that do not have VFMs and only perform one thrust maneuver, rather than specify an inclination, the thrust vector is angled out of the orbital plane by some amount. This angle is given by the variable “Z-Thrust Angle.” The final inclination that can be achieved is dependent on a number of other factors (such as total mass, thruster efficiency, and solar panel angle relative to the flight path angle), so the inclination of this satellite is an output of the performance model rather than an input. This is acceptable because “inclination” is not an attribute associated with Goals #2 or #3, but rather “vertical sampling range” in units of AU is. In contrast, satellites with VFMs can perform multiple burns, and “inclination” is an attribute associated with Goal #1, so specific inclinations can and should be targeted. This is represented by the design variable “Target Inclination.”

The levels for Z-Thrust Angle range between zero and an upper limit where designs beyond this lever were infeasible. For two-petal designs, this limit is much higher; designs that require longer burn times to achieve shorter-period orbits cannot afford to allocate as much out-of-plane thrust to change the inclination, hence the limit is smaller for these designs.

If an FPA is present (as determined by Payload Option), it can take one of the three designs from Table 6-4, since those correspond closely with minimum, middle, and maximum levels of satisfaction from the stakeholder attribute curve associated with that instrument. If a VFM is present, several values for aperture diameter were chosen to represent a range of possible performance outputs while still keeping the instrument small enough for use on these satellites.

6.3.2.2 Performance Modeling Process 2: Design Modeling

HOBOP’s performance model consists of a variety of modules that ensure that a design is both feasible and (somewhat) optimal. In comparison with ExoplanetSat’s performance model, HOBOP’s performance model requires many more modules to compute utility and cost, though each module is not as detailed and exhaustive as ExoplanetSat’s performance module. Each module will be introduced in the following subsections, and each will be discussed in more detail in the System Optimizer section so the reader can understand how these modules interact and apply constraints to one another.

Known Small Satellite Components

The performance and mass characteristics of most of the components on the HOB0 spacecraft can be modeled based on known components from CubeSat suppliers. The parts that were used in HOB0COP are listed in Table 6-16.

Table 6-16: Known components of the HOB0 satellite, suppliers, and masses.

Component	Supplier	Mass
Star Tracker	Sinclair Planetary	0.158 kg (with protective case)
Reaction Control Wheels	Maryland Aerospace Inc.	0.360 kg + 0.120 kg/6 kg wet mass
CPU (Radiation Hardened)	GOMSpace	0.100 kg x 2 (Rad penalty)
Solar Panels (x2)	Clyde Space	0.135 kg x 2
Radio Transceiver (Rad Hardened)	GOMSpace	0.130 kg x 2 (Rad penalty)

The mass of critical computer components is shown to be doubled as a penalty for radiation hardening. While extensive radiation doses can be calculated using advanced models, it is also appropriate to assume system redundancy or that some amount of mass per square centimeter can be used to protect the individual components [391]. Table 6-16 also lists which components are penalized with radiation hardening and the penalties associated with them. There are additional mass penalties for radiation protection on payloads; these will be discussed later in this section.

Astrodynamics Module

Initial Keplerian orbits were modeled in MATLAB with a custom propagator. Other simulations and video demonstrations were created in STK to test the feasibility of the propulsion system, but because STK's Astrogator module is slow to modify elements of many satellites, a second custom MATLAB propagator was written that included finite propulsion elements. This model was validated against STK's results to ensure accuracy. This module is heavily dependent on the design of the expendable propulsion module and requires elements from the thermal, power, and thruster models, and it will be discussed in further detail later in this section.

Communications Module

Communications were modeled with a link budget used in previous class projects at MIT that has been validated against many other communication systems. Expert opinion was solicited from Dr. Alessandra Babuscia regarding mass estimates of current technologies being developed for small satellites in the X-band. A Ka-band system that used the same components for power conversion and transmission as the X-band system was assumed to have a similar mass.

Many weeks were spent attempting to characterize an optical communications subsystem to create better performance and mass estimates. Due to limited and proprietary information and a lack of expert validation, a simplified version was created that assumes a small transmitter mass and uses the same link budget with a transmission wavelength of 1.5 μm instead of a frequency on the order of GHz. System noise temperature was given a high estimate because of the lack of published information on how to calculate this parameter in optical communications (see Table 6-11). Additionally, because the DSN is not an optical network, a smaller receiver was assumed.

Thermal Module

A simple, one-dimensional finite element model was constructed to verify that the spacecraft heat shield would be sufficient to keep the temperature of the other components below an acceptable level. A more detailed thermal analysis of the HOBO satellites is beyond the scope of this work. A sketch of this model and the components of heat transfer that were modeled with it are shown in Figure 6-17.

The model consists of three nodes: the heat shield front, the heat shield back, and the spacecraft structure. The front of the heat shield is coated with a highly reflective material like quartz-over-silver or Silverized Teflon. This coating could degrade, however, and become less reflective and emissive, reflected in the pessimistic technology epoch (Table 6-11). The heat shield material is a carbon-carbon material like CFOAM that has a low thermal conductivity. The back of the heat shield is coated with white epoxy to maximize emissivity. Torlon rods connect the heat shield to the structure because they provide the lowest thermal conductivity for support elements. The spacecraft body is also covered in a white epoxy material to maximize effective radiated heat. The MLI was not modeled but would only improve the spacecraft's maximum internal temperature.

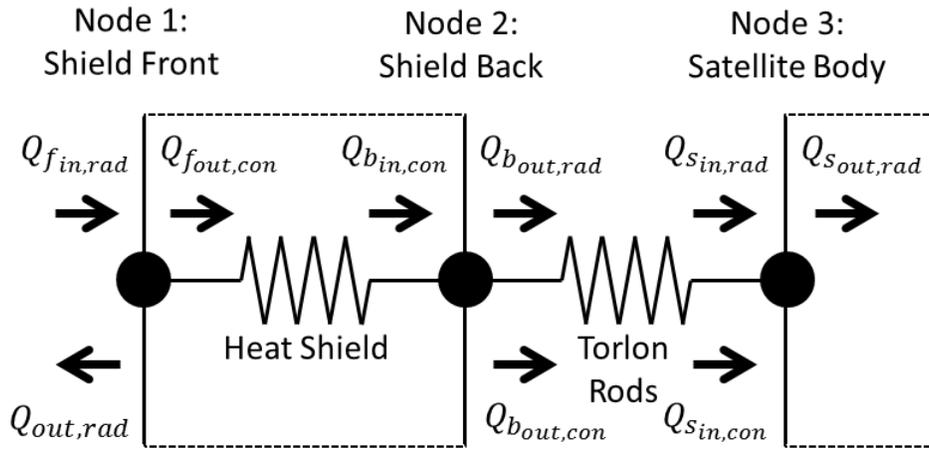


Figure 6-17: Sketch of the one-dimensional thermal model to ensure survivable temperature of a HOBO satellite with a heat shield.

While technically it is more accurate to run a time-dependent model that accounts for the varying distance of the spacecraft from the Sun during its orbit, the model showed that there was little difference between the maximum spacecraft temperature in the time dependent model with orbit propagation versus the steady-state equations. For this reason, the more conservative steady-state assumption was used in the tradespace model because it is much faster to compute.

The solar panels on the propulsion module are also a concern for heating. As the satellite decreases its orbital period, it is still moving closer to the Sun. While temperature is a cause for concern regarding reduced power, the increase in solar flux negates this because of the increase in available power. However, the maximum operating temperature of the solar panels cannot be exceeded during this time. A check was set in place to ensure that the temperature of the solar panels stayed below the maximum:

$$T_{SP} = \sqrt[4]{\frac{W_0}{r^2} \left(\frac{\alpha_{SP}}{\epsilon_{SP}} \right) \frac{1}{2} \cos \varphi} < T_{SP_{max}} \quad (6-2)$$

where $W_0 = 1365 \text{ W/m}^2$ is the solar flux at 1AU, r is the distance from the Sun at the time, α_{SP} and ϵ_{SP} are the absorptivity and emissivity of the solar panels, and $\varphi = \theta - \gamma$ is the difference between the True Anomaly and the flight path angle. The factor of $\frac{1}{2}$ is a result of the solar panel being two-sided so it can radiate from double the surface area that it is absorbing solar radiation.

This is actually one of the major limitations of achieving lower-period orbits with an expendable propulsion system. For long thrust times used in SEP maneuvers, the limit for how long a thruster can operate is usually limited by the fuel. In this case study, the failure of the solar panels from heat is a major limitation, especially when there is *too much* fuel, which causes the craft to accelerate too slowly to achieve its final orbit if it is trying to reach a low period.

An additional limitation built into the thermal model is the melting temperature of the coating on the heat shield. Although carbon-carbon can survive high temperatures, quartz-over-silver and Silverized Teflon will melt. A check is in place to ensure that the front face of the shield stays below the melting point in the heat shield mass calculation.

Power Module

While the HOBOSAT satellites will spend no time in the shade, the heat flux near the Sun is high enough to cause failure of the solar panels if left exposed for the entire flight. For this reason, a battery would need to be installed to help minimize the amount of time the solar panels would need to be exposed. A battery life of 2 days was estimated as necessary due to the low duty cycle that could be afforded during the closest approach to the Sun.

The small deployable solar panels on the HOBOSAT satellite are not mathematically minimized in the same way other subsystems components are because they are off-the-shelf components, but there is a check built-in to the integrated model to ensure that they provide ample power even at a distance of 1 AU. Larger payloads with more power draw would require more solar panel area to operate and would be added on the front and back sides of the model shown in Figure 6-1.

The power production unit (PPU) of the SEP system is sized to match the capacity of the larger solar panel on the expendable propulsion module. Initial simulations used both the solar panel area and PPU capacity as independent variables. This was because as the satellite moves closer to the Sun, the power generated by the solar panels would increase due to the increased solar flux, and a larger PPU could handle the extra power to provide extra thrust compared to when it is at 1 AU while keeping overall mass minimized. However, astrodynamics principles show that thrust is most effective at the orbit perihelion. Simulations showed that the propulsion system had a better chance of surviving and not overheating before reaching a desired lower-period orbit when the PPU and solar panels were matched at 1 AU.

Propulsion Module

The HOBOSATs will use Xenon gas as the primary propellant. One of the many advantages of using a propellant with a high molar mass is that the pressure and volume are lower compared to propellants with lower molar mass. The mass of the propellant tank scales as a function of pressure multiplied by volume, and the ideal gas law can be used to calculate this factor.

SMAD provides parametric models for estimating the mass of propellant tanks made of stainless steel or composite-reinforced metal. While these models may be useful for larger systems, they are not

necessarily the best for estimating the mass of smaller tanks. The SMAD model for a composite reinforced system was used for estimating the propellant tank mass in the pessimistic technology epoch.

In the optimistic technology epoch, the mass of the propellant tank was estimated using the stress equations for an inflated sphere. The tank is assumed to be composed of Kevlar, which has a strength-to-weight ratio four times higher than steel. This provided mass estimates that were a fraction of the SMAD estimates, although with custom tanks the cost would realistically be much higher than heavier, off-the-shelf equivalents.

The mass of the electrospray thrusters was assumed to be constant depending on the context (see Table 6-11). While research over the past five years has drastically reduced the mass of electrospray thrusters that use ionic liquids as propellants, more conservative estimates were made because these systems have not been tested in space environments yet, and the cost of procuring ionic liquids in bulk for this mission is not known compared to the availability of Xenon.

Instrument Module

Table 6-3 and Table 6-4 show the possible instruments that can be used for magnetometers and FPAs. Rather than building a parametric model to sample different options within these areas, the design space will simply use these exact instruments. It is evident that the performance of the SSBV magnetometer and the ELFIN energetic particle detector are not Pareto-optimal and will not be considered. For the remaining options, a mass multiplier penalty of 1.5 will be applied for radiation protection.

For the VFM, a comparative model was built using the SMAD scaling equations and the known properties of SO/PHI (35 kg, 30 W, 140 mm aperture, 200 km resolution at closest perihelion). The scaling equations are

$$R = \frac{A_i}{A_0}, \quad W_i = KR^3W_0, \quad P_i = KR^3P_0 \quad (6-3)$$

where $K = 2$ because the aperture diameter is less than 50 cm. For an aperture diameter of 80 mm, the ratio $R = 0.5714$, and the total mass is $W_i = 6.53$ kg, and total power $P_i = 5.60$ W.

System Optimizer Module

Finally, the individual modules were brought together to calculate the total mass required. The optimizer starts with a given design and initial estimates on the final orbital radii of aphelion and perihelion, the required propellant, and the HOBO mass. The heat shield mass is first calculated, then the spacecraft bus mass, then the expendable propulsion module mass. Next, the orbit is propagated to ensure that it achieves the required orbital period while satisfying the thermal constraints.

If the satellite runs out of fuel before reaching the desired orbit, more fuel is added on the next iteration. If the spacecraft is still thrusting as the propulsion module's solar panels reach their maximum operating temperature, fuel is subtracted. If the optimizer flips back and forth between these two options, the design is infeasible because the spacecraft is underpowered and cannot reach the desired orbit. If the desired orbit is reached but there is fuel remaining, the remaining fuel is subtracted and the loop continues until the remaining fuel is below some small minimum.

Designs may reach their desired orbital period but still be rejected. The amount of data a payload generates could exceed the communication system’s capability to downlink the data in a feasible amount of time. If connection to the DSN requires more than 100 hours per satellite, the design is thrown out. If a design requires Earth flybys to achieve a desired orbital inclination but its aphelion distance after its initial thrust is less than 0.95 AU, it is assumed that it will not be powerful enough to bring its aphelion back to 1 AU to achieve those flybys.

Once the desired orbit is achieved and the mass is minimized, the estimated data rates are calculated. Total data volume is based on the payload.

Heliospheric Current Sheet Module

A basic model of the heliospheric current sheet was created in MATLAB to estimate the location of the current sheet in time and when which satellites in which orbits could be expected to cross the plane at which times. A visualization of this model as the current sheet changes from the period of Solar Minimum and moving into Solar Maximum is shown in Figure 6-18. A [video demonstration](#) is also available that shows the growth of the Parker Spiral [392]. Although it does not represent any quantitative metric of satisfaction, this is a useful tool for showing the stakeholders where the satellites will be at certain times relative to the shape of the current sheet.

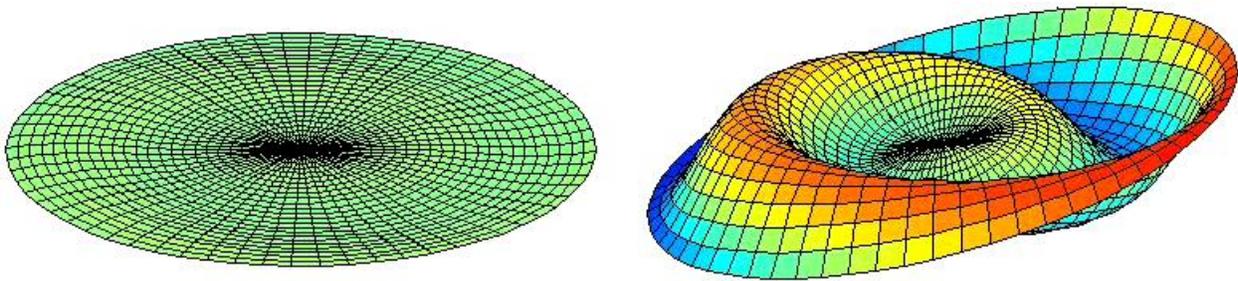


Figure 6-18: MATLAB model of the heliospheric current sheet at solar minimum (left) and near solar maximum (right).

6.3.2.3 Performance Modeling Process 3: Cost Modeling

Cost is calculated based on the complex cost presented in Section 4.3.4. Because a mothership carrier was not explicitly designed for this mission, a fraction of the total mass m_C was assumed for whatever spacecraft would carry all the units beyond Earth orbit and into heliocentric orbit rather than an actual mass M_C . The equation for the cost of the spacecraft was modified slightly as a result and is shown here.

$$C_{SC} = c_C m_C \left(\sum_{i=1}^m \left(c_S M_{Si} + c_{P_i} M_{Pi} + \sum_{j=2}^{n_i} (c_{PD_j} c_S M_{Si} + c_{SP_j} c_{P_i} M_{Pi}) \right) + c_F M_F \right) C_{CC} \quad (6-4)$$

An additional modification is a gross cost per “complex kilogram” of mass C_{CC} , which assumes that after all the discounts have been applied to the payload and fuel, the adjusted mass is a flat rate per kilogram. This multiplying factor converts the design from kilograms to dollars. The values of all the cost estimating parameters in the formula above are listed in Table 6-17.

Table 6-17: Cost Parameters for HOBOCOP Cost Estimate

Cost Parameter	Name	Value
Mothership Carrier Cost/Mass	c_C	1
Mothership Carrier Mass Fraction	m_C	0.15
Spacecraft Cost/Mass	c_S	1
Payload Cost/Mass	c_P	10
Spacecraft Cost Discount	c_{SD}	0.1
Payload Cost Discount	c_{PD}	0.2
Fuel Cost/Mass	c_F	0.05
Cost Per Complex Kilogram	C_{CC}	\$500,000

The cost of accessing the DSN was also added. In FY2015, the cost of accessing the DSN is approximately \$6,500 per hour, and the satellites only make one connection with the DSN per year as they reach perihelion near the time Earth swings by. Given the data rates and the required link time calculated in the communications module, an operations cost estimate could be made and added to the system.

6.3.2.4 Performance Model Limitations

Myriad assumptions were involved in the construction of this performance model. While most of them were made based on the information available at the time, new information can always be used to update these assumptions in future iterations.

The scaling equations for spatial resolution in optical instruments do not necessarily apply in the same way to spectral resolution in spectrometers. Because there were no other sources of information, magnetic resolution of the VFM was scaled the same way the spatial resolution was. A more detailed instrument performance model would warrant the inclusion of at least one additional design variable to characterize the magnetic resolution because it is an important attribute to Goal #1.

The performance of the system and the achievability of certain orbits are heavily sensitive to the assumptions on mass and efficiency of the components related to the propulsion system. For example, 1/4- and 1/3-year orbits are difficult to achieve with larger systems, and small changes in critical assumptions like solar panel maximum operating temperature could drastically change the number of designs that can feasibly reach these orbits.

The astrodynamics module assumes Keplerian orbits for the HOBOS. It assumes that solar radiation pressure is negligible, which may be fine for this stage as there is room for uncertainty because the precision of the orbit is not important. The module also ignores the effect of general relativity, which in contrast, may have much more effect on the orbit. For example, Mercury’s orbit is affected enough by general relativity that the perturbations were observed long before general relativity was proposed, and the HOBOS will achieve orbits within Mercury’s perihelion distance.

Noise calculation in communications systems and radiation exposure are two areas that are of concerns in this model. While it is easy to assume robust values for both components to ensure a feasible system, that is not the best option when dealing with multiple satellites because the penalties are multiplied by the number of assets in the system. Being off by a small factor could cause a cascade of designs to be less attractive or feasible depending on the number of assets involved.

Finally, a major limitation is determining how much data is worthwhile if not all of the data over the course of a year can be downlinked. Even if the range is far and the communications system is substandard, valuable data can still be downlinked to Earth. Whether or not a system has a communications system that is feasible was simply set so that all designs that require more than 100 hours of time on the DSN at their closest approach would not be satisfactory. If a satellite is *not* recording data at full capacity throughout its entire orbit, it would not generate as much data, but it would also have gaps in observation time. Modeling how these gaps would affect the value proposition is beyond the scope of this case study, but the value of other missions like Solar Orbiter and Galileo are restricted by how little data they can downlink compared to how much they are capable of recording. Setting a specific cutoff for DSN time on the higher end serves as a compromise to prevent mostly feasible designs from being eliminated solely because of their communications system even if 100 hours of contact time per satellite per year on the DSN is unrealistic because the observations that would be downlinked can be prioritized before contact is made.

6.3.3 HOBOCOP RSC Phase 3: Tradespace Exploration

In the third phase of RSC, the results of value modeling and performance modeling combine and the tradespace of options becomes visible. There is now enough information to select designs that meet the needs of the stakeholders.

However, before showing this tradespace, it is important to identify what context should be considered the “primary context” for this analysis. After examining the tradespaces of the different possible epochs from a high level, it is apparent that the communications technology is a key component in determining whether or not a system is feasible. The data generation rate of the VFM is high enough that the small satellite communications system used on each satellite is too small to make continuous operation of the instrument feasible when operating in the S-band. The same system operating in the Ka-band does have feasible results. Additionally, the optimistic technology epoch and the initial stakeholder value model (which assumed time-insensitive data and no relay satellites) have served as a baseline for much of this work. For these reasons, this epoch variable combination will serve as the tradespace that is explored in this phase.

First, the design vector as shown in Table 6-15, with homogeneous assets only, was explored. The MAU for each science goal and for the entire mission is shown in Figure 6-19. Based on the value model, it was obvious that designs that used a VFM would be the only ones that satisfy Goal #1, designs that used a FPA would be the only ones to satisfy Goal #2, and designs that use a magnetometer would be the only ones to satisfy Goal #3. That means that the only designs that had any hope of partially satisfying all three goals were the ones that carried all three payloads, which corresponds to Payload Option 5 in the design variables (black dots in the Figure 6-19).

However, Payload Option 5 is restricted to ½-year orbits because it does not have a heat shield, meaning those HOBOS will never fully satisfy the goals of measuring the solar wind at sub-Alfvénic velocities nor achieve close-range simultaneous flybys. This is why there are no black dots in the top right of Figure 6-19; none of these designs are satisfactory. The black dots in the bottom left are present because the minimally acceptable radial distance is larger for that science goal. Additionally, the substantial cost increase of the VFM compared to the other instruments means that it is more expensive to fly additional assets, so the in-situ spatial sampling resolution would be reduced in comparison to other designs.

The Pareto front is dominated on the lower end by the designs that satisfy Goal #1 because fewer assets are required to obtain minimum satisfaction levels for Goal #1 compared to Goal #2 and Goal #3, even if the assets themselves are much heavier. Designs that carry both a magnetometer and an FPA (Payload Option 3) are not on the Pareto front of Goal #1 (top right) or Goal #2 (bottom left), but because they satisfy both, they dominate over Payload Options 1 and 2 in the combined utility ranking (bottom left).

There is clearly tension between the design options, but there is also known synergy between the science goals that is lost if only homogeneous assets are used. This is why the total mission MAU does not go much higher than 0.4 in Figure 6-19.

Heterogeneous sets of assets working in the same campaign offer a solution to capturing more value than any single asset type. However, choosing combinations of design sets to evaluate is extremely cumbersome. A total of 1,320 *individual* satellite designs were evaluated before being combined with the architectural variables, which increased the size of the tradespace to 395,750 *possible* (not necessarily feasible) architectures in eight contexts (additional contexts can be evaluated without the performance model since they are copies of the same set of data that are applied to different value models).

Rather than simply picking every single combination of two or three options, some intelligence was applied in selecting the set of heterogeneous designs. First, the Pareto fronts of each payload option in each goal and the combined MAU of all goals were identified (20 Pareto fronts total). The set of designs on all of these fronts were then chosen in combination in each epoch, resulting in an additional ~40,000 designs depending on the size of the Pareto fronts.

The difference in the value proposition between homogeneous and heterogeneous assets is clearly noticeable. By comparing Figure 6-19 with Figure 6-20, it is evident that there are major advantages to using more than one type of asset.

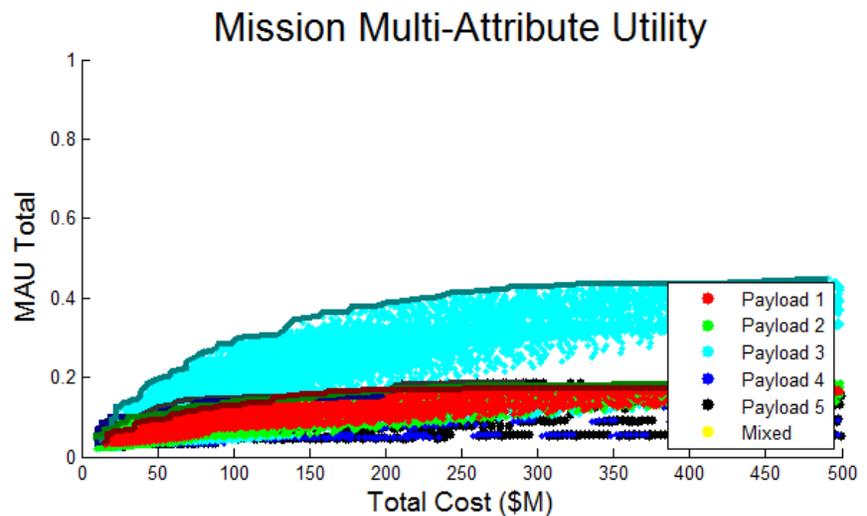
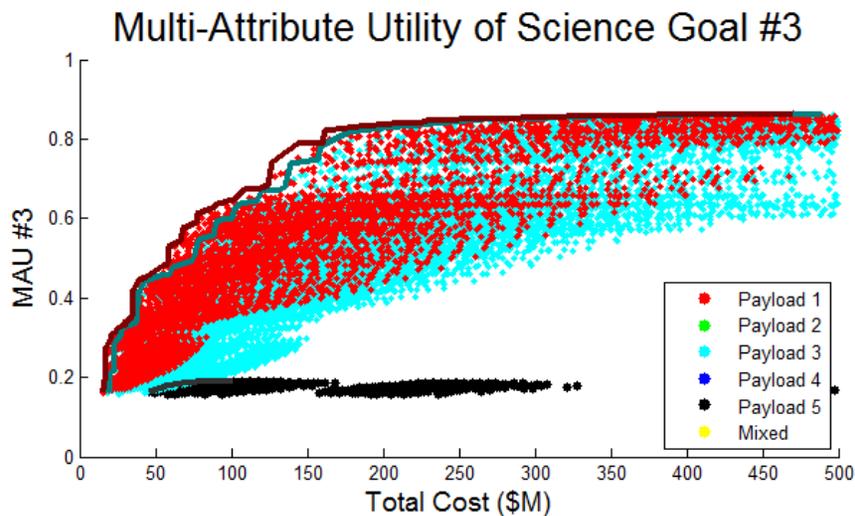
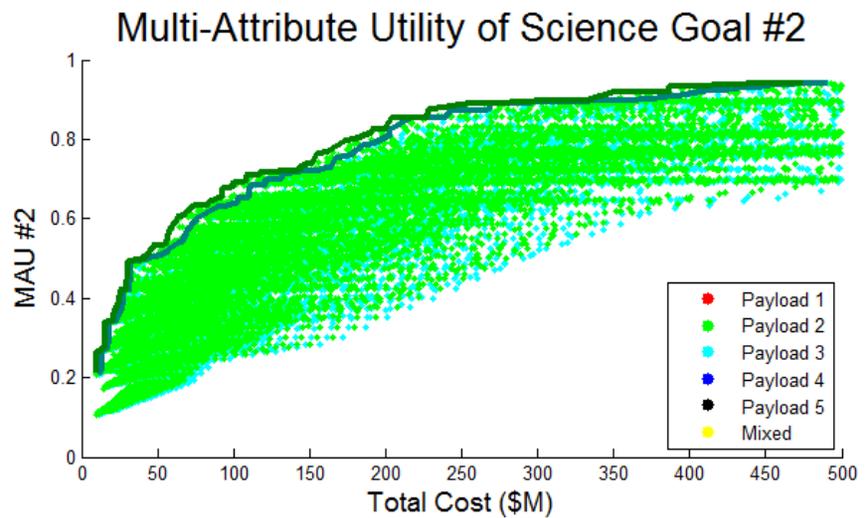
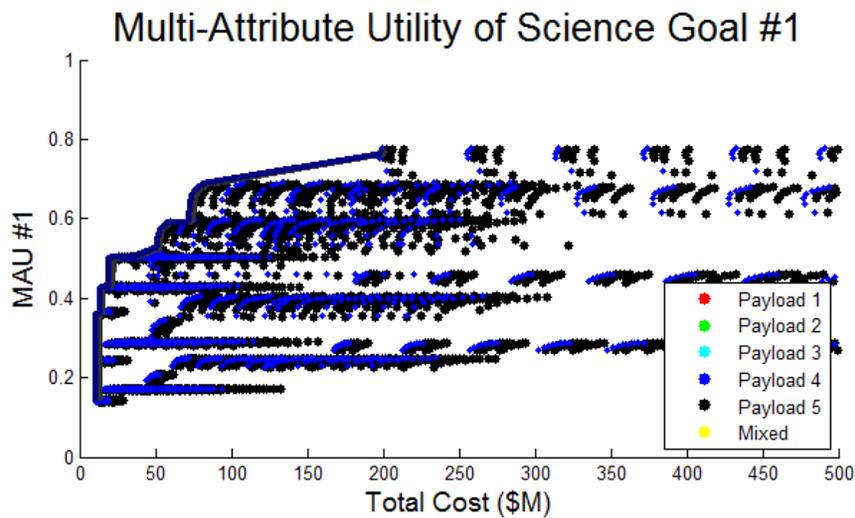


Figure 6-19: Tradespace of HOBOCOP with only homogeneous assets, color-coded by payload option. (Top left) MAU of Goal #1. (Top right) MAU of Goal #2. (Bottom left) MAU of Goal #3. (Bottom Right) Total Mission MAU.

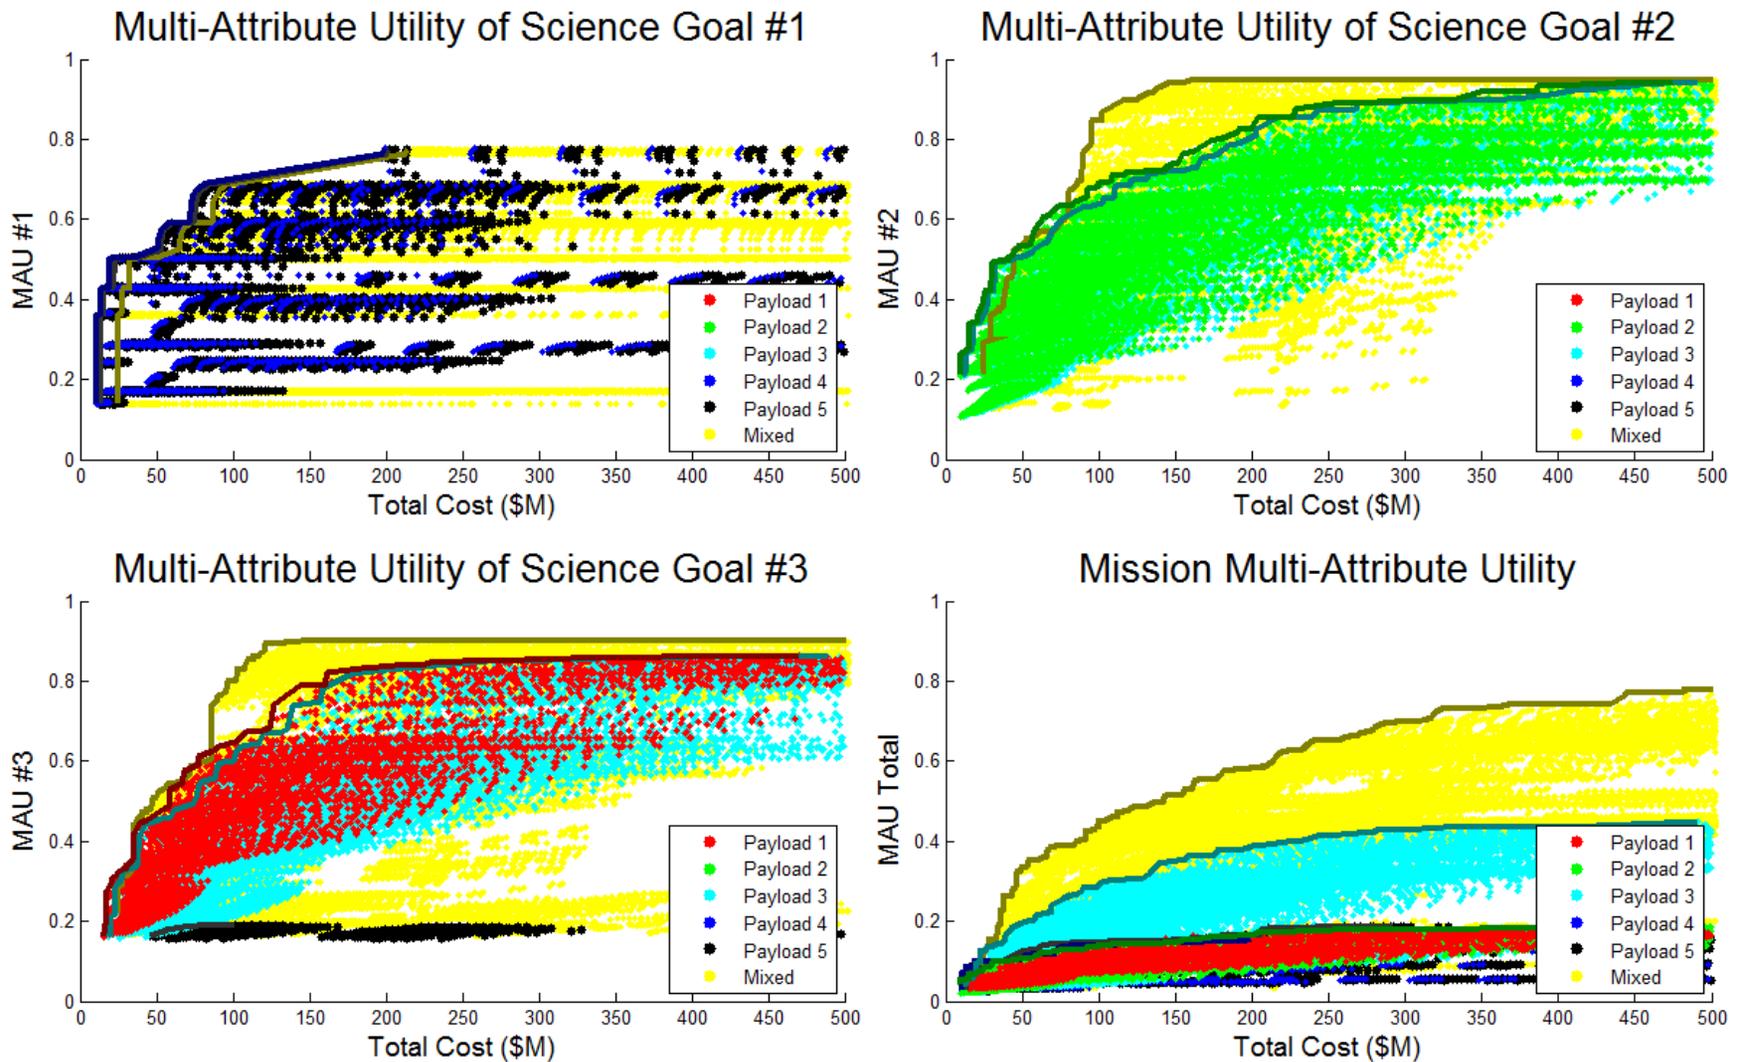


Figure 6-20: Tradespace of HOBOCOP with both homogeneous and heterogeneous assets, color-coded by payload option. (Top left) MAU of Goal #1. (Top right) MAU of Goal #2. (Bottom left) MAU of Goal #3. (Bottom Right) Total Mission MAU.

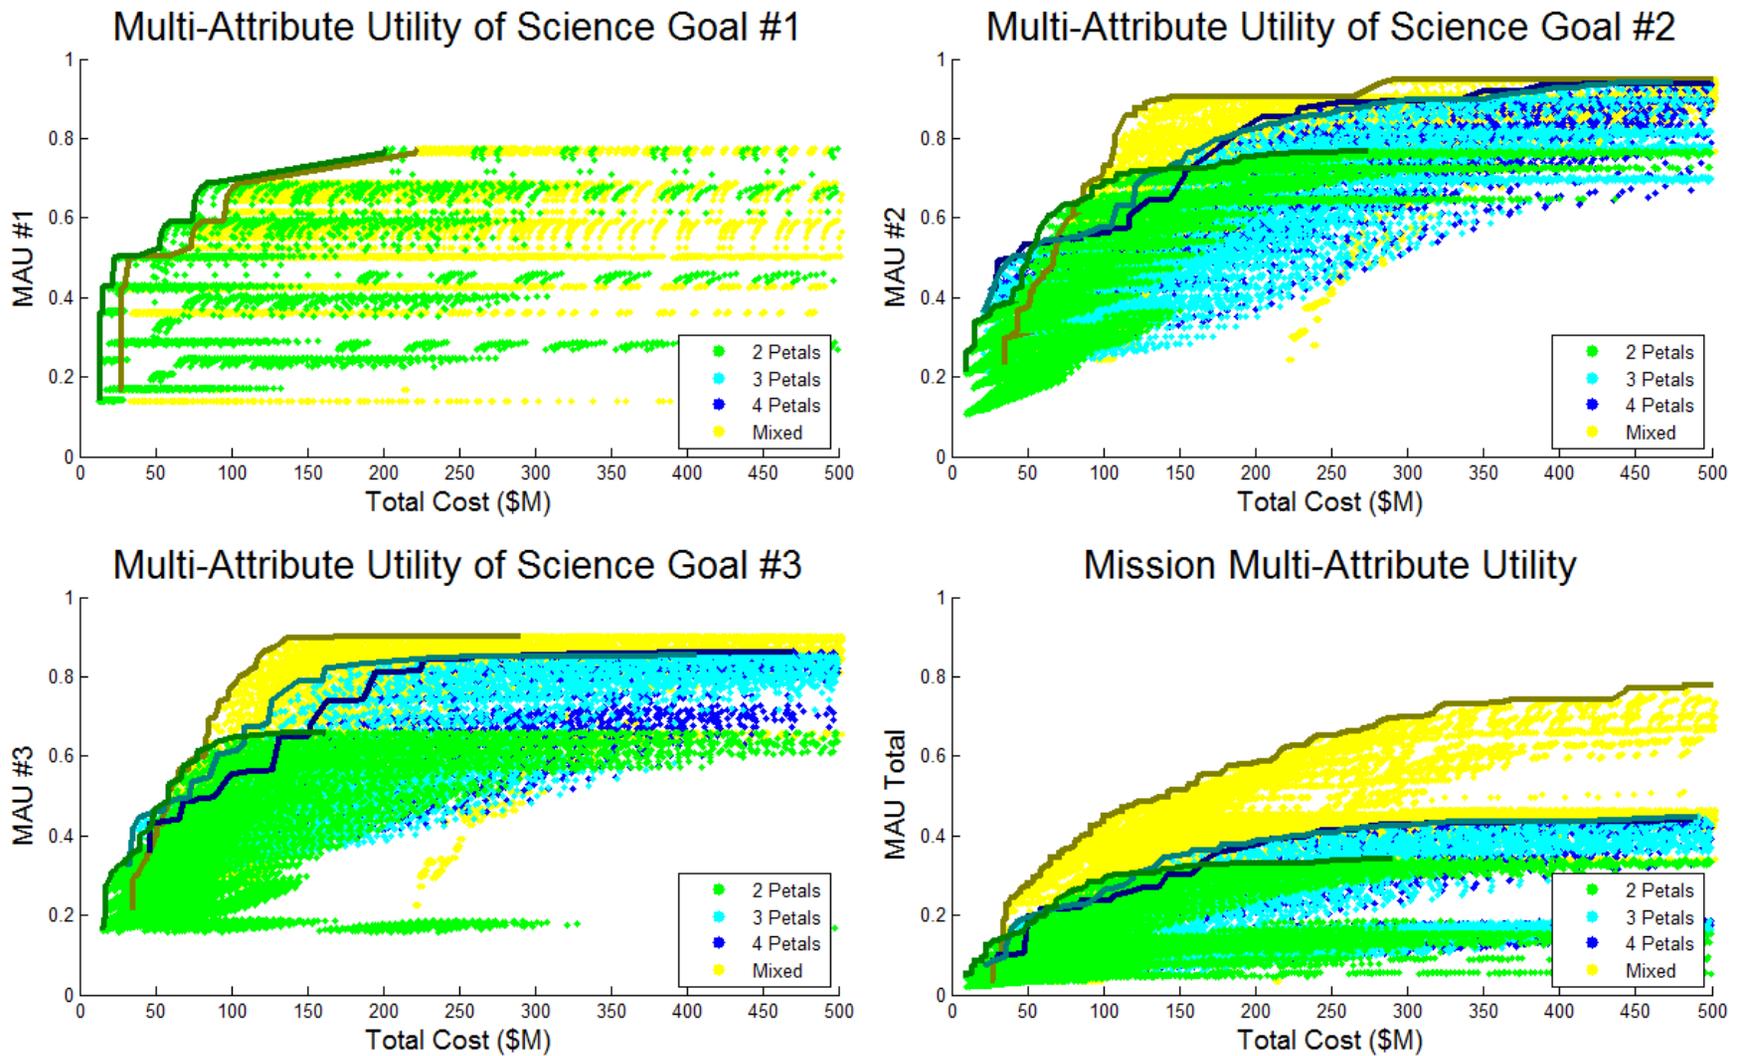


Figure 6-21: Tradespace of HOBOP with both homogeneous and heterogeneous assets, color-coded by number of petals. (Top left) MAU of Goal #1. (Top right) MAU of Goal #2. (Bottom left) MAU of Goal #3. (Bottom Right) Total Mission MAU.

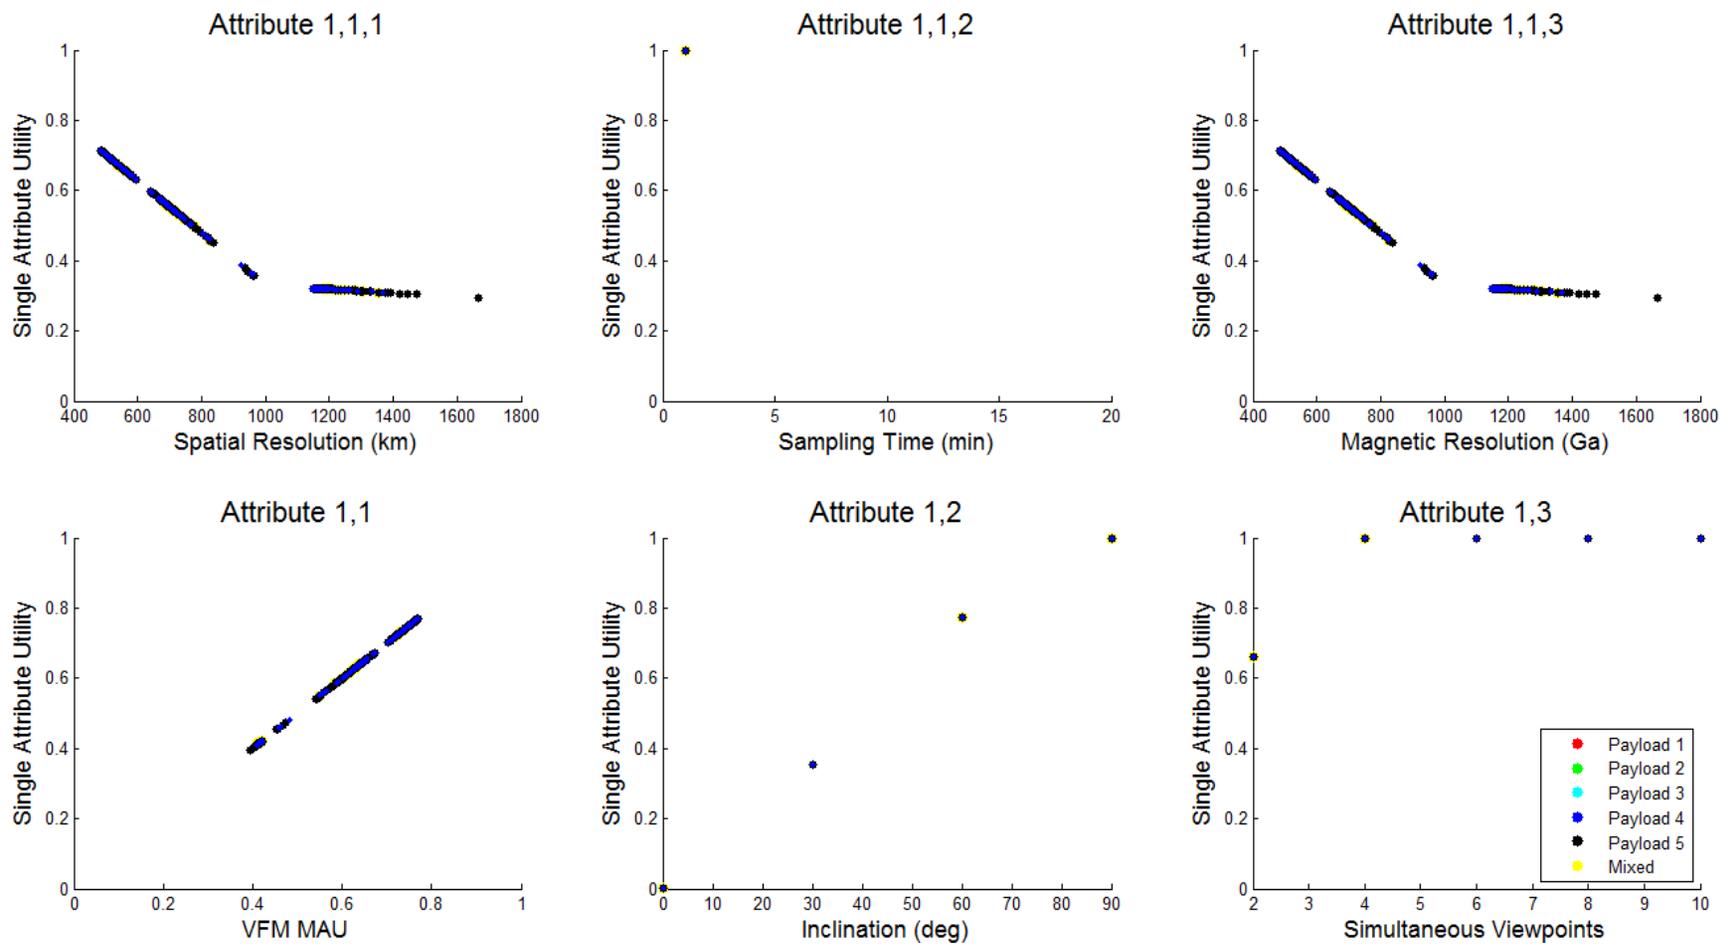


Figure 6-22: Attributes and utility levels for Science Goal #1.

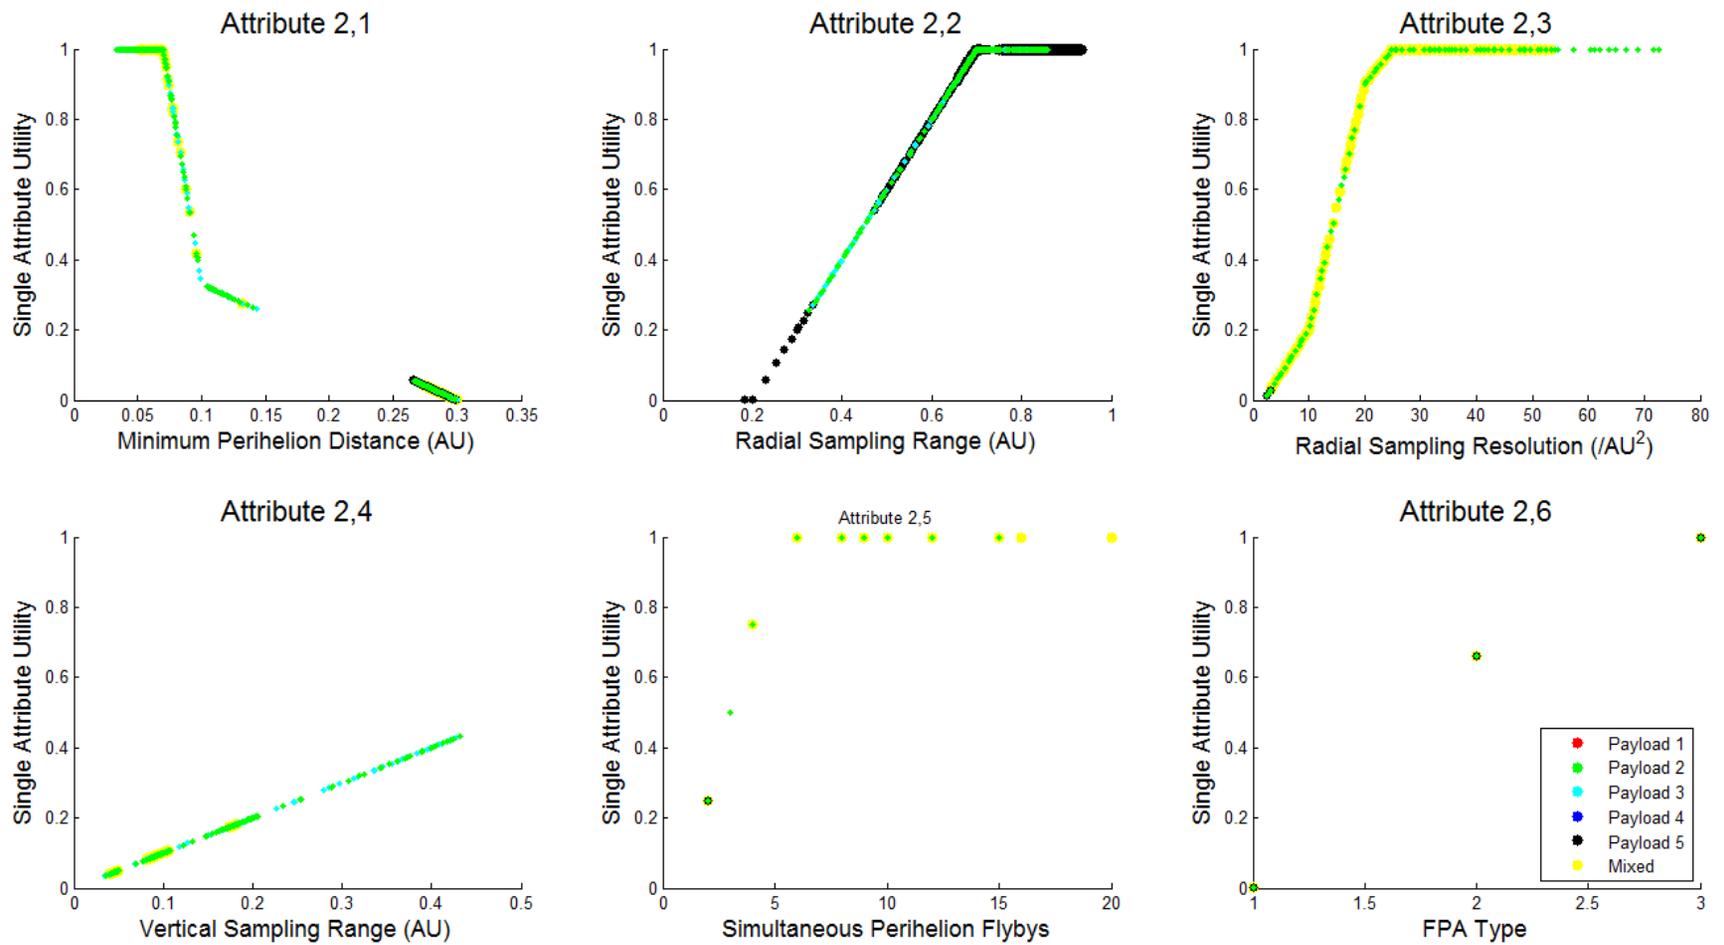


Figure 6-23: Attributes and utility levels for Science Goal #2.

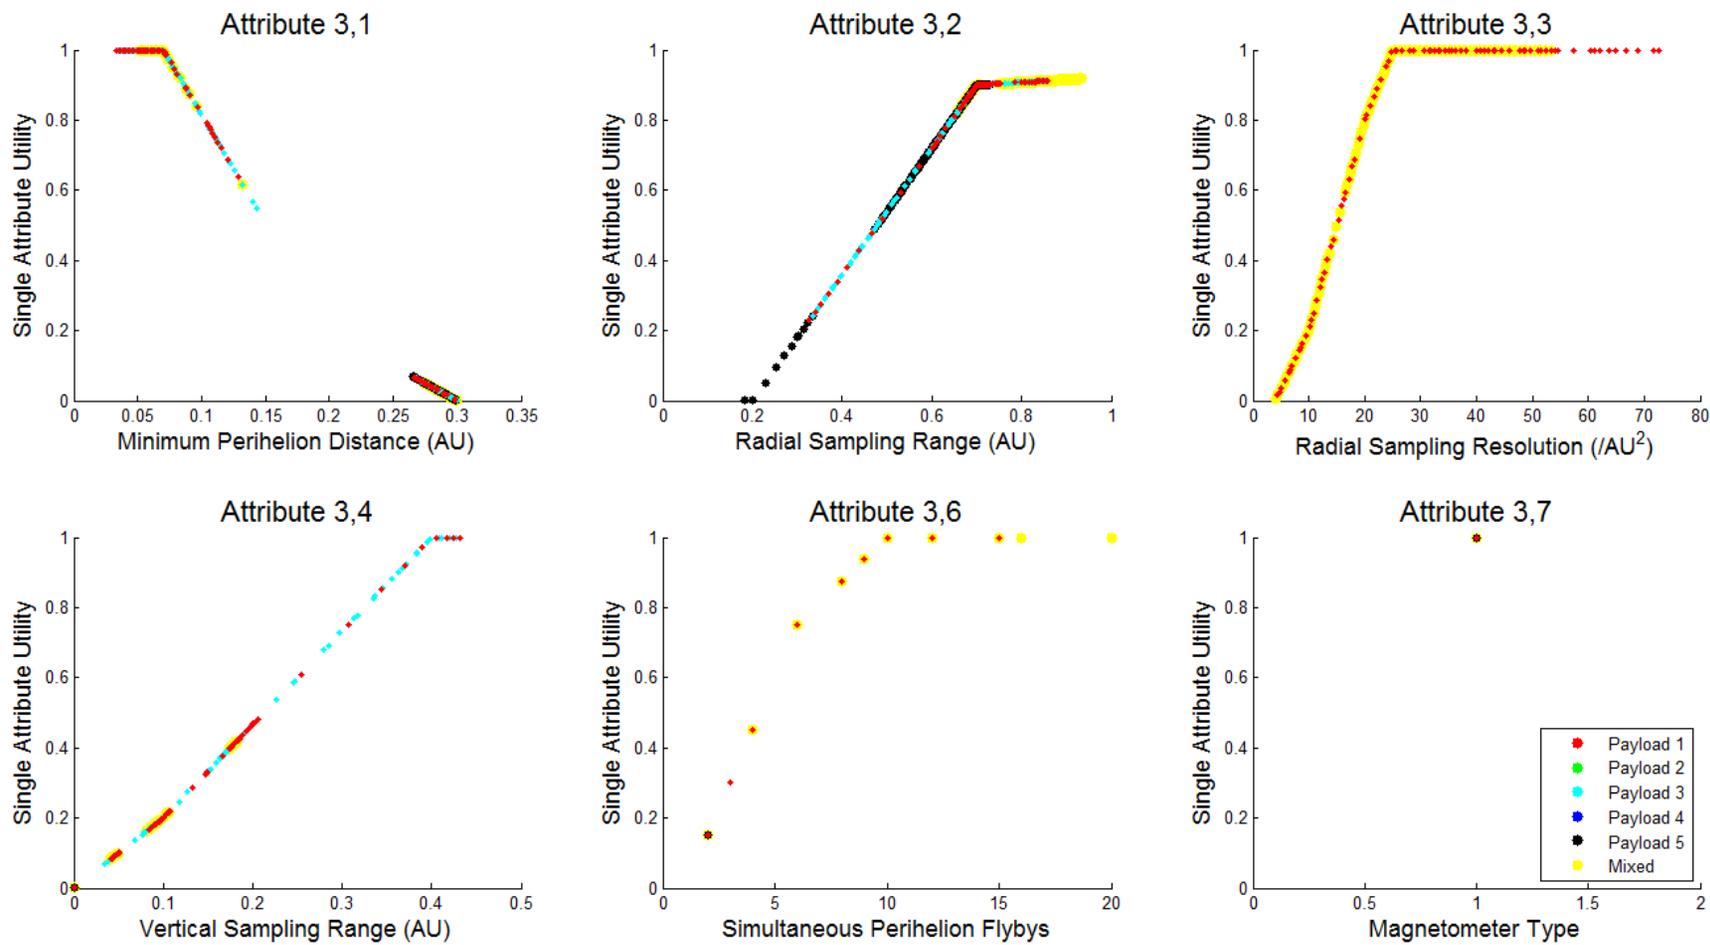


Figure 6-24: Attributes and utility levels for Science Goal #3.

There are peaks of heterogeneous designs in both Goal #2 and Goal #3 because of the advantages of using different sets of orbits; assets in 1/3- and 1/4-year period orbits require much more mass to achieve higher inclinations compared to assets in 1/2-year period orbits, but the assets in polar orbits can sample the solar wind and magnetic field at greater latitudinal range. Additionally, some assets combine so that their combined perihelion and aphelion range is greater than either one individually. As a result, both the “Vertical” and “Radial Sampling Range” attributes of Goal #2 and Goal #3 are better satisfied according to the utility functions shown in Figure 6-11 and Figure 6-12.

The tradespace with both homogeneous and heterogeneous assets is shown in Figure 6-21 again, but color-coded by the highest number of petals in the design. The tradespace of Goal #1 looks no different compared to the previous figure as expected; no designs with three or four petals carry VFMs, and any heterogeneous assets only add dead weight to achieving this goal.

The tradespace is more complex for Goals #2 and #3. In the top right of Figure 6-21, the designs with two, three, and four petals weave amongst each other along the Pareto front. Designs with two petals are in 1/2-year orbits and require the least propellant mass and dominate the lower end. Further along, designs with three and four petals take over the Pareto front because of the importance of the minimum perihelion distance attribute combined with the threshold returns on sampling resolution. Even further along the Pareto front, designs with two petals return to dominance as the threshold returns on sampling resolution start improving rapidly with each additional asset. Finally, after costs increase above \$150 million, three- and four-petal designs return to being dominant and remain there.

When heterogeneous designs are included, there is a large increase in utility over the Pareto front of the homogeneous designs starting near \$100 million. At this point, heterogeneous designs continue to dominate the Pareto front for the rest of the tradespace because the heterogeneous assets of a design make up for what each homogeneous asset is lacking. A similar pattern is seen in Goal #3 for the same reasons, but heterogeneous assets are entirely unmatched because only heterogeneous assets can achieve the radial sampling range to maximize that attribute.

The attribute levels of Goal #1 for all the designs are shown in Figure 6-22. From these curves, some comments can be made about the design space. The spatial resolution never achieves the maximum attribute level because the aperture diameter of the VFM does not go high enough nor does any option for an orbit go close enough to the Sun to achieve a spatial resolution better than 400 km (maximum utility is at 100 km). No design variables affect VFM sampling time, which is why there is no variation in that attribute’s utility. Simultaneous viewpoints are meant to be discretized, but there could be more options for targeted inclination to examine the value of that attribute in more detail.

The attribute levels of Goal #2 for all the designs are shown in Figure 6-23. There is a large gap in the minimum perihelion levels because of the natural difference between 1/2-year and 1/3-year period orbits. There were only three FPAs to choose from, hence the three points for the FPA type attribute. Otherwise, the design space samples these attributes well.

The attribute levels of Goal #3 for all the designs are shown in Figure 6-24. There is a similar gap in minimum perihelion distance levels like in Goal #2. Because there is only one magnetometer in the design space, there is only one point for the magnetometer type attribute.

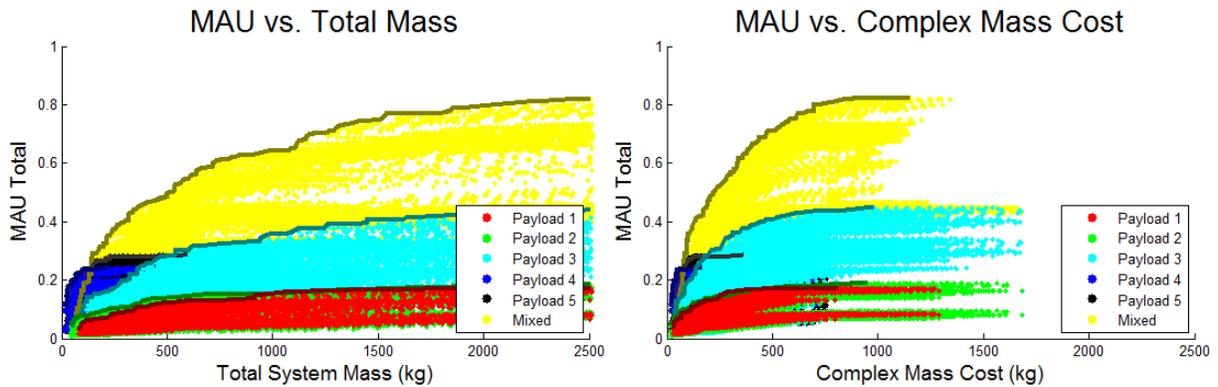


Figure 6-25: Comparison of tradespaces showing (left) total mass versus (right) complex mass.

The effects of separating payload costs and fuel costs from other subsystems costs when estimating the total spacecraft cost can be seen in Figure 6-25. The two x-axes are the same scale, but because of learning curves, a cost estimate based purely on mass would be incorrect. The slope of the Pareto front on the lower end of the cost spectrum is higher with complex mass than it is compared to bulk mass.

There are many more intricacies of the design space to be explored in Epoch-Era Analysis, but from tradespace visualization alone, several key components of value can be seen. Although it was already obvious from the sampling resolution attributes alone in Goals #2 and #3 that this case study leveraging **simultaneous sampling** to achieve objectives that no monolithic satellite could, the benefits of **stacked sampling** by leveraging heterogeneous distributed assets and by synergies between science goals are more apparent because of the value model as it was elicited from experts.

6.3.4 HOBOCOP RSC Phase 4: Epoch-Era Analysis

The final phase of RSC examines how the mission can change due to factors outside of the control of the designers and how the mission can change over time as a result of disturbances and other change options. Although many interesting properties of the design space have been examined, there are a number of critical sensitivities that must also be examined in order for the decision makers to be able to choose a design that will fit their needs and remain valuable over the course of the system lifecycle.

6.3.4.1 EEA Process 1: Single Epoch Analyses

The context variables and alternative stakeholder value models were also applied to the design vector to examine the tradespace under changing epochs. Eight epochs' tradespaces are shown in Figure 6-26 and Figure 6-27.

The first major finding from examining these tradespaces is how value is more easily obtained from designs that satisfy Goal #1 under the alternate stakeholder value model for aggregating the goal-level attributes of that goal. This can be best seen by comparing the top two graphs in Figure 6-26 with the bottom two. The blue and black Pareto fronts representing payload design options #4 and #5 achieve higher levels of satisfaction at lower costs; these designs are Pareto-efficient further along the full Pareto front all the way until they reach their maximum utility. This is most evident in the bottom right of the figure, where the Pareto front plateaus briefly before heterogeneous designs (yellow dots) take over with rapidly increasing MAU. As a result of this value model change, heterogeneous designs that incorporate payload option #5 achieve higher utility overall and higher utility at lower cost in comparison to the original value model. This change in value model will be revisited later in this section.

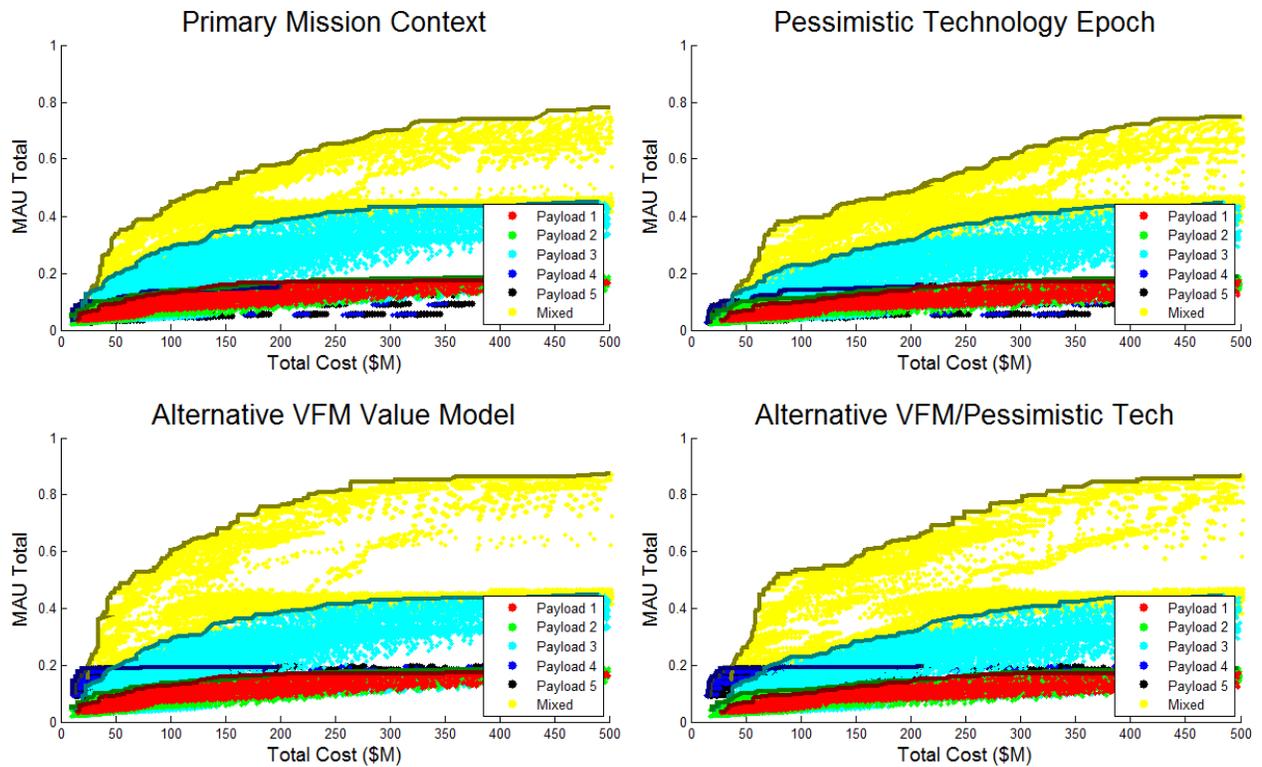


Figure 6-26: Epoch comparisons of Mission MAU. (Top left) primary mission context. (Top right) pessimistic technology epoch. (Bottom left) alternate stakeholder value model changing weights on Goal #1. (Bottom right) alternate stakeholder value model and pessimistic technology epoch.

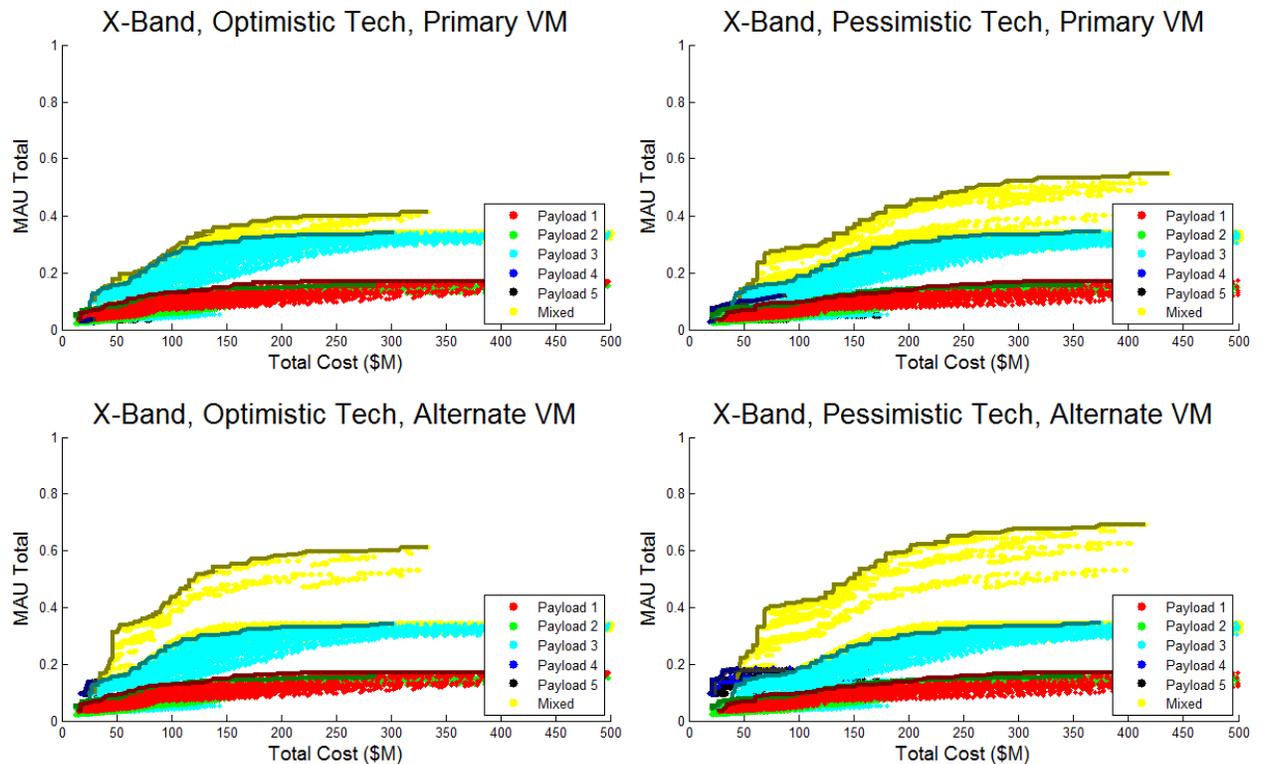


Figure 6-27: Same as previous figure, but with X-Band technology instead of Ka-band in each epoch.

Another obvious change between the epochs shown in Figure 6-26 is the change in the cost of designs near the “knee” of the Pareto front due to changes in technology outlook. The left graphs in the figure show the optimistic technology context, while the right graphs show the pessimistic technology context. The “knee” of the Pareto front occurs near a price of \$50 million under the optimistic technology context, whereas it occurs near \$75 million under the pessimistic technology context.

The difference between available communications technology is shown in Figure 6-27. Here, the same four epochs are shown similarly to the previous figure, but with a system that uses the X-band instead of the Ka-band. In these epochs, overall satisfaction is much lower and the cost to achieve the same satisfaction is significantly higher. It is not obvious from inspection, but there are far fewer feasible designs that carry payloads #4 and #5 because they generate the highest data rates and cannot feasibly communicate it all back. This is why there is such a small difference between the heterogeneous and homogeneous Pareto fronts in the top left; heterogeneous designs don’t have as many feasible options to utilize payloads #4 and #5.

There is a much bigger difference between the heterogeneous and homogeneous Pareto fronts in the bottom left because of the increased satisfaction from achieving Goal #1 under the alternate value model. This means any heterogeneous designs that do utilize the few feasible options with payloads #4 and #5 gain added value in this context compared to the primary context. Comparing the left and right sides of Figure 6-27 also results in a similar conclusion as it did in Figure 6-26; the pessimistic technology epoch required higher costs to achieve the same levels of satisfaction.

The effects of this change in value model on Goal #1 specifically are examined more closely in Figure 6-28. Comparing the same technology context to the two different value models (left side), the Pareto front rises far more quickly under the alternative value model. An MAU close to the maximum possible level can be achieved for ~\$25 million, whereas a design on the Pareto front under the original value model with the same cost would only achieve an MAU slightly higher than 0.5.

What is most apparent from this epoch comparison at the goal-level is the difference in the number of feasible designs between the communications technology availability contexts. The left side of Figure 6-28 shows two well-populated tradespaces using the Ka-band, but the right side showing designs that use the X-band looks empty in comparison. Under the epochs with X-band communication, it is difficult to downlink the full data volume of the VFM, and no other figure shows it better than this one. The only feasible designs are the lightest ones (smallest aperture diameter); given how much more important the instrument attributes are in the primary stakeholder value model, it makes sense why the MAUs in the top right are so much lower than the ones in the bottom right.

With the exception of Figure 6-28, it is not readily apparent from these tradespaces that the number of feasible designs changes significantly from epoch to epoch depending on the technology available. The fractions of feasible designs over the total number of designs are summarized in Table 6-18.

The first finding this data shows is the dramatic difference in feasibility between S-band and Ka-band communications systems of the same size. Many of the S-band designs were eliminated due to the amount of time required to link to the DSN to downlink a year’s worth of data in one pass. Of the instrument options, VFMs generate the most data, which is why feasibility drops from 84.6% to 3.1% with this change in communications wavelength.

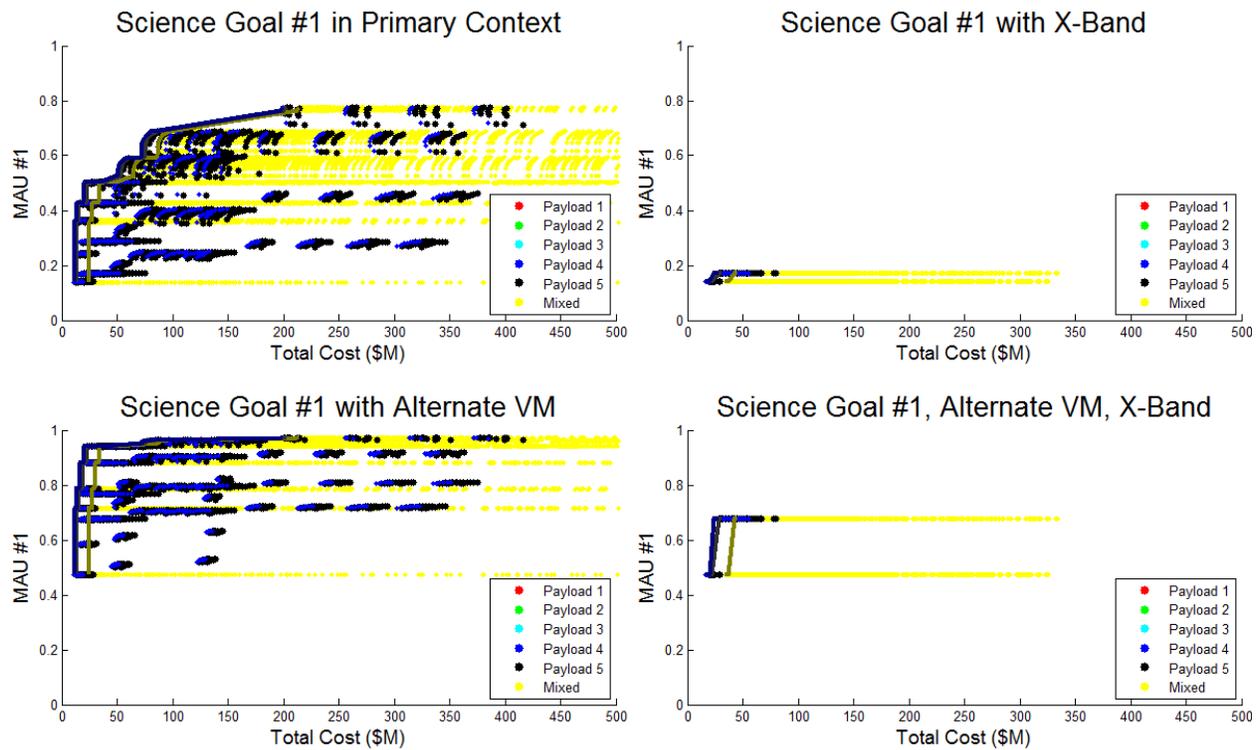


Figure 6-28: Epoch comparison of Goal #1. (Top left) original stakeholder value model, Ka-band (Top right) Original stakeholder value model, X-band. (Bottom left) Alternate stakeholder value model, Ka-band. (Bottom right) Alternate stakeholder value model, X-band.

Table 6-18: Comparison of the fraction of feasible designs in different technology epochs.

Epoch #	Comm. Type	Tech. Outlook	All Designs	VFM Designs	2-Petal Designs	3-Petal Designs	4-Petal Designs
1	Ka-Band	Optimistic	79.2%	84.6%	97.3%	79.0%	44.3%
2	Ka-Band	Pessimistic	71.4%	88.1%	90.5%	66.1%	37.1%
3	X-Band	Optimistic	33.2%	3.1%	62.4%	11.9%	2.2%
4	X-Band	Pessimistic	32.1%	15.6%	60.0%	9.6%	2.3%

The results in Table 6-18 also show that in the pessimistic technology epoch, fewer designs are feasible, although it does not have nearly as large of an impact on feasibility as communications does. However, assets carrying a VFM were surprisingly *more* feasible in the pessimistic technology era, by a small percentage. After seeing this, a mistake was discovered in the formulation of the design vector. All designs should be independent of epoch variables (though in this case, the communications technology epoch rightfully represents a design change depending on the availability of certain technologies). However, in the system optimizer module, not only was the initial guess on the required fuel doubled, so was the design variable controlling the solar panel area. This means *all* designs in the pessimistic technology epoch have double the available solar collecting area and double the PPU maximum capacity.

Technically, designs with the same designation in the optimistic and pessimistic epochs are different as a result of this oversight. However, this mistake highlights even more why accurate performance modeling

is critical for systems with high ΔV requirements; even doubling the power availability does not completely make up for the losses incurred by components with higher mass or less efficient designs. Doubling the solar panel area was an intelligent decision to produce more feasible designs in that epoch and avoid wasting time calculating infeasible designs; the fact that the differences are as small as they are is a testament to how much intuition was gained by modeling many designs instead of a few.

One could also argue that this oversight also means that the pessimistic technology epoch is unfairly penalized in cost as a result; however, the range of possible solar panel sizes analyzes a swath of possibilities that typically includes infeasible designs on the lower end. In order to guarantee this, it would have to be examined in detail, but a cursory glance at the raw results shows that the smallest solar panel areas in both epochs are infeasible most of the time.

6.3.4.2 EEA Process 2: Multi-Epoch Analysis

Multi-epoch analysis is not as applicable in this case study as it is in other responsive systems because many of the epochs are cost- or performance-related and are mutually exclusive within the same era. Mission managers would know in advance which of the assumed performance parameters would change as time goes on and would be able to adapt the designs to fit the needs of the mission.

As a result of these very different performance modeling assumptions, many of the designs that are Pareto-dominant in one epoch are infeasible in another epoch. The design with the highest utility in the primary design context may have a low normalized Pareto trace (NPT) because they are infeasible in other epochs. Designs with high NPT tend to be the ones carrying only a magnetometer payload because they have the lowest data rates and are the lightest designs. These designs are great for maximizing the satisfaction of one goal, but not all goals across the entire mission.

Furthermore, it is difficult to conduct a true multi-epoch analysis with the HOBOCOP tradespace because of the vast number of heterogeneous designs there are. The combinations of heterogeneous assets shown in the above tradespaces were generated from the Pareto-optimal designs for each science goal and the MAU *within each epoch*, so those combinations would be different in each epoch. The number of possible designs exceeds the computational limits that are available.

Multi-epoch analysis may still give insight on the comparison among epochs with different value models, even if it's in a limited scope. As was shown in single-epoch analysis on the left side of Figure 6-26, designs with VFMs can gain utility by having more than one asset under the alternative value model.

6.3.4.3 EEA Process 3: Single-Era Analyses

Value from HOBOCOP is dependent on long-term observations because the stakeholders wish to study the Sun and the interplanetary environment as the solar cycle changes from solar minimum to solar maximum (see Figure 6-13 for the lifetime attribute), and era analyses will help better understand the risks associated with the scientific value delivered by a particular design. Despite the fact that value is “sensed” continuously with the instruments, value is “delivered” discretely because the satellites are not in constant communication with Earth. If the satellites are communicating only with Earth through the DSN, data is delivered only once per year when the satellite makes a close pass with Earth.

Due to the intense radiation environment as the satellites sweep through perihelion, assets are more likely to fail during or after the most scientifically valuable part of their orbit. There is a significant chance that

data will be lost as a result. Furthermore, satellites that make more perihelion passes per year will have a higher risk of losing more data; satellites in 4-petal configurations ($\frac{1}{4}$ -year orbit periods) could therefore have four perihelion passes of data on board if they fail on their way to aphelion as Earth approaches.

Systems that use optical communications or are otherwise able to communicate for most of their orbital periods mitigate this risk. As a result, it is appropriate to model the value delivered by these systems as continuous rather than discrete in Single-Era Analyses.

Just like in the ExoplanetSat case study, the epochs that have been identified are somewhat mutually exclusive, and no epoch shifts would occur during a mission's lifetime. It may still be possible to experience an epoch shift due to changes in stakeholder expectations, such as the alternative value models or a change in the time sensitivity of the data, but these changes are more likely to also be precluded by technology changes that would affect the hardware design and would not occur after launch. For this reason, only single-epoch eras will be considered in era analysis.

There are also no change options built into this system nor are there ways to capture opportunity as it comes during the course of the mission. The only disturbance that will be assumed in the construction of eras is the failure of an individual satellite.

If a single satellite of the constellation fails, the design vector of the design being tested in era analysis changes in the following ways. First, the total number of HOBOS decreases by one. The number of HOBOS per petal also drops by a small factor depending on the number of arcs and the number of petals per arc. If the number of failed HOBOS exceeds the number of HOBOS per arc, the total number of arcs decreases by one. The total number of HOBOS can never be less than the total number of arcs, nor can the number of petals per arc decrease below one. Once these have been accounted for, the MAU can be calculated using the modified set of design variables. Finally, because there is no need to contact failed satellites, there is no additional added cost to the system for operations costs when they are recalculated.

Due to the simplistic nature of single era analyses after the previous rules for constructed eras have been identified, multi-era analysis will be conducted in favor of addressing single era analyses in more detail.

6.3.4.4 EEA Process 4: Multi-Era Analysis

In multi-era analysis, the remaining contexts will be addressed in more detail and how changes in those contexts affect the overall value that a system can deliver.

Individual Design Results

Results from multi-era analysis on two specific example designs are shown in Figure 6-29 and Figure 6-30. In each of these figures, the top left plot shows a distribution of cumulative value based on Monte Carlo simulations of 200 eras. The static utility of these eras is shown in the top center. In both plots, the thick black line represents the design if no assets failed during the course of the mission. The cumulative value and static utility are calculated annually in these plots because that is how often the HOBOS communicate with Earth. The top right shows the probability distribution of value, while the bottom left translates this into a cumulative value distribution. The bottom center shows the probabilistic tradespace of the cost and cumulative value of the design, and the bottom right shows how many satellites are remaining as a function of time.

Design: 112110 Value Model: 1 Probability of Failure: 0.05 Contacts Per Year: 1

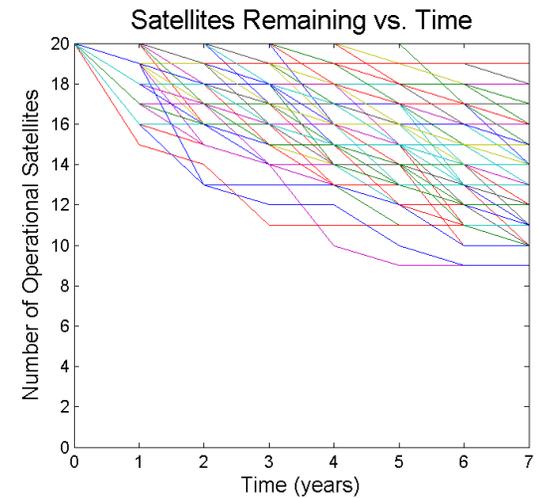
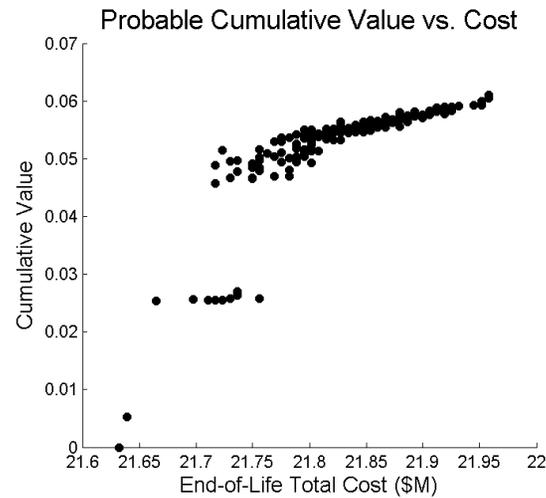
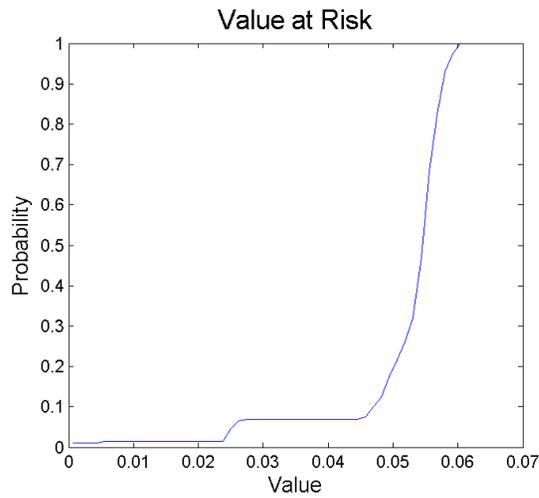
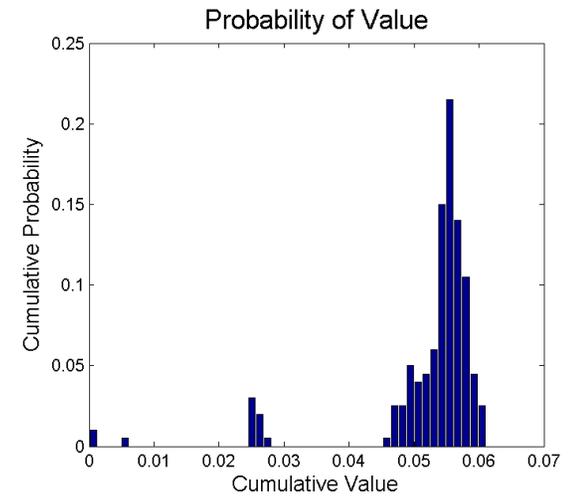
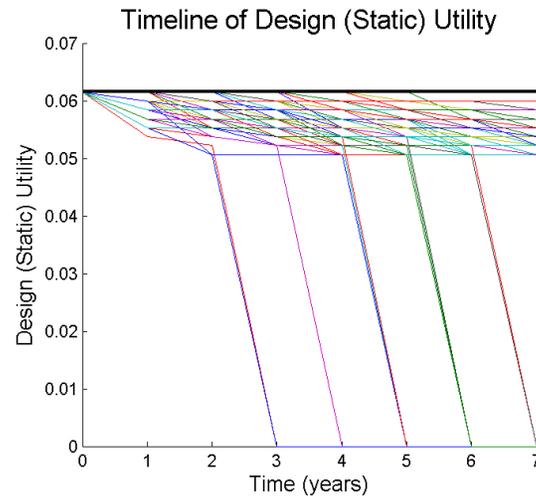
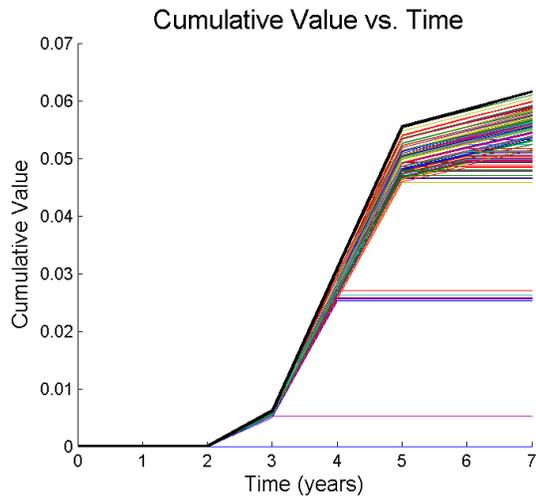


Figure 6-29: Multi-era analysis results for Design 112110.

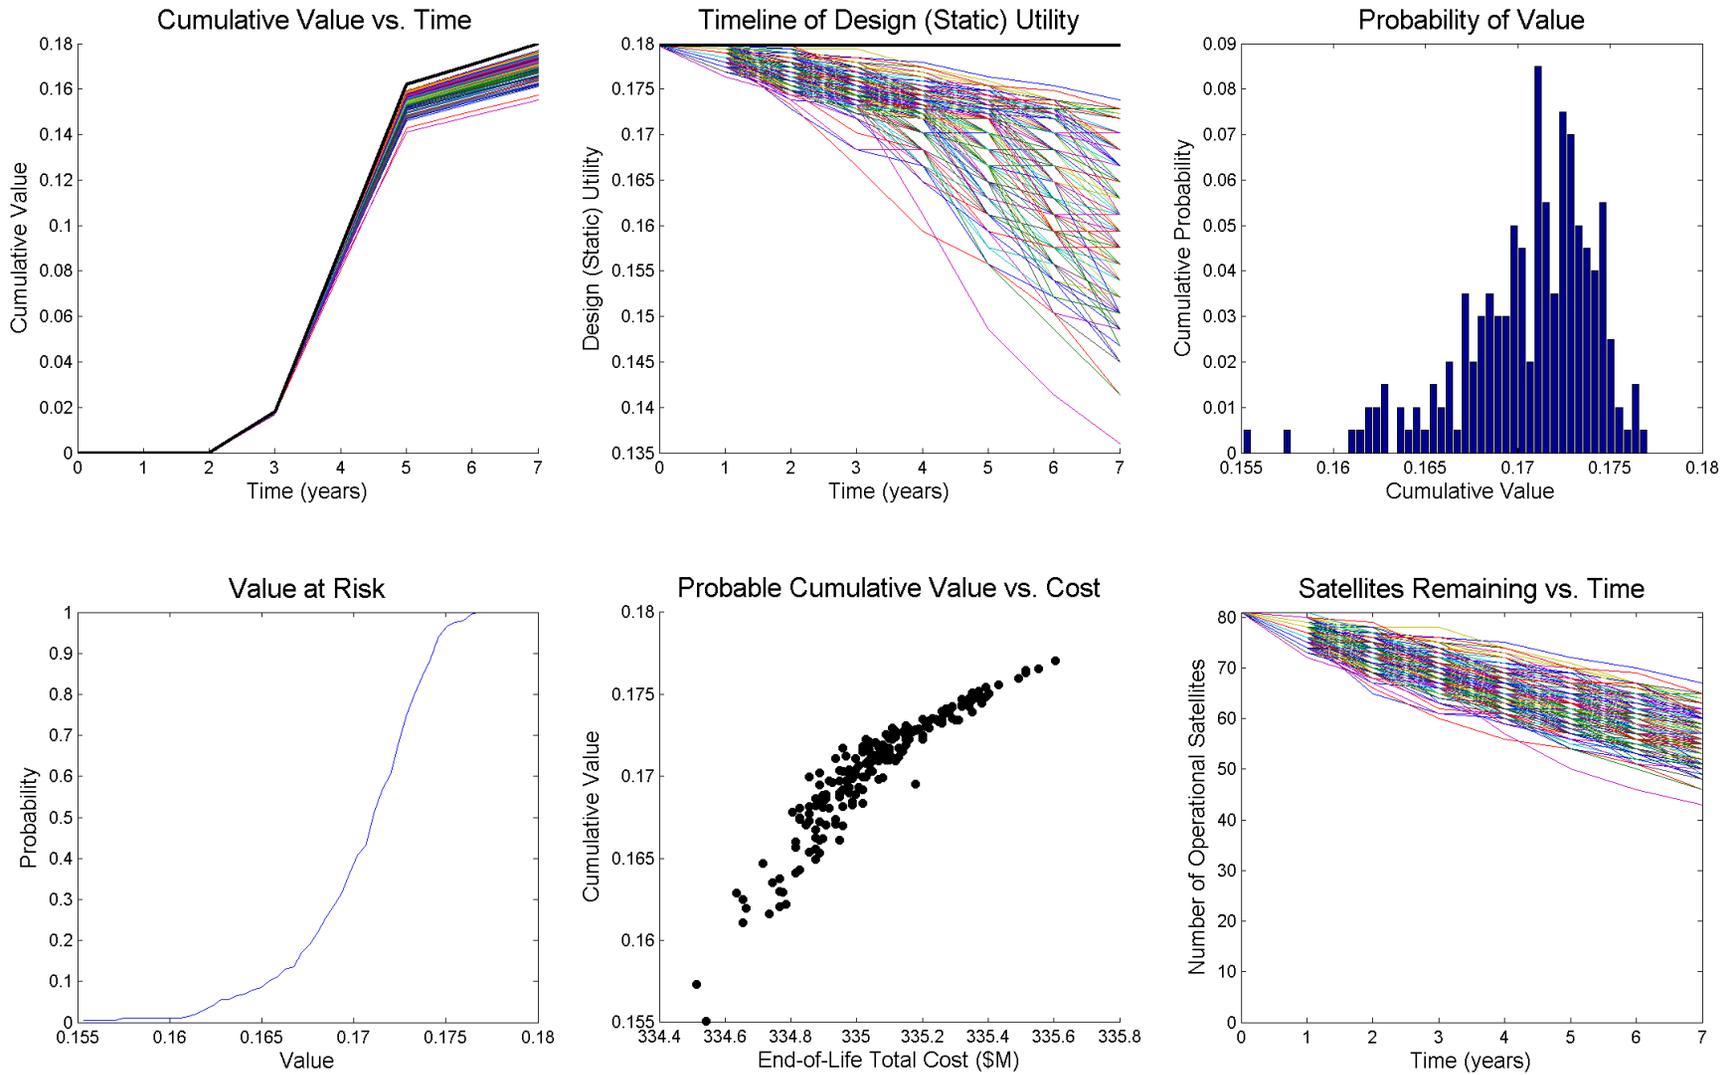


Figure 6-30: Multi-era analysis results for Design 144387.

An easy way to check that definitions of coverage are accurate is to examine the static utility over time diagram (top center). Assuming satellites are failing, the utility of the system must be going down. However, if there are any points where the utility increases, there could be an error in either the definition of coverage, which would indicate a problem in the value or performance model, or an error in the rules for evaluating designs in era analysis. While developing the initial set of rules for how satellites could fail in HOBOCOP in era analysis, there were several instances of some designs increasing utility. Luckily, these errors were related to the era construction rather than a larger, systemic problem.

One important observation from these figures is the probability distribution of value (top right). For some designs, especially ones with a lower number of satellites, the probabilistic distribution of value *does not* have a Gaussian distribution (Figure 6-29). As satellites fail, their static utility may fall in different ways depending on which satellites fail and when. Designs with more assets are more robust to failures of individual assets, and the cumulative value distribution tends to be more Gaussian in nature (Figure 6-30).

These figures show the results for two designs under the primary stakeholder value model where satellites have a probability of failure of 5% per year. There are an additional 6,376 full-page figures that show how these results vary for every design given four different failure rates. Examination of these figures shows that the spread of operations costs can be several million dollars for large constellations depending on how quickly satellites fail, but the spread in cumulative value can be within +/- 0.01 of the mean.

Communications Networks

A brief sidetrack is necessary here to explain how multi-era analysis will address the options for communicating with relay satellites. Up to this point, there has been little discussion of how the COPs (Communicating Orbital Platforms) affect the value of the mission. However, multi-era analysis can illustrate the added benefit of being able to communicate with a network of satellites.

Much of the original work on HOBOCOP involved optimizing a network of relay satellites by minimizing mass, minimizing length between contacts, and maximizing the total data throughput of the network various lifecycle lengths. However, this work took place before scientist stakeholders were identified, and the problem setup was inherently flawed. The HOBOS are in eccentric orbits rather than the COPs; finding the new COP orbits was a problem that was delayed for future analysis.

With a tradespace that better reflected the value model, the performance modeling in HOBOCOP found a simpler, and arguably more elegant, solution for possible placement of any COP satellites. When the number of HOBOS was increased beyond what is feasible in a single launch, and all HOBOS orbits were propagated in time, the instantaneous organization of the constellation took a form that was completely unexpected. The individual satellites behaved according to the laws of gravitation, but the curve connecting all the satellites in the same set of petals rotates with a period of one year. A static image of this pattern is shown in Figure 6-31 (a [video demonstration](#) is also available for viewing [393]).

From this configuration, it is evident that a system leveraging assets with two petals (aquamarine dots) would benefit from having a relay satellite placed on the opposite side of Earth's solar orbit. However, a satellite in this location would be nearly impossible to communicate with from Earth; as of this writing, both STEREO-A and B are near this location and cannot communicate with Earth using their primary antenna because of the high noise associated with trying to communicate on a line of sight so close to the Sun. As a result, they are currently using low-gain antennas and communicating at limited data rates.

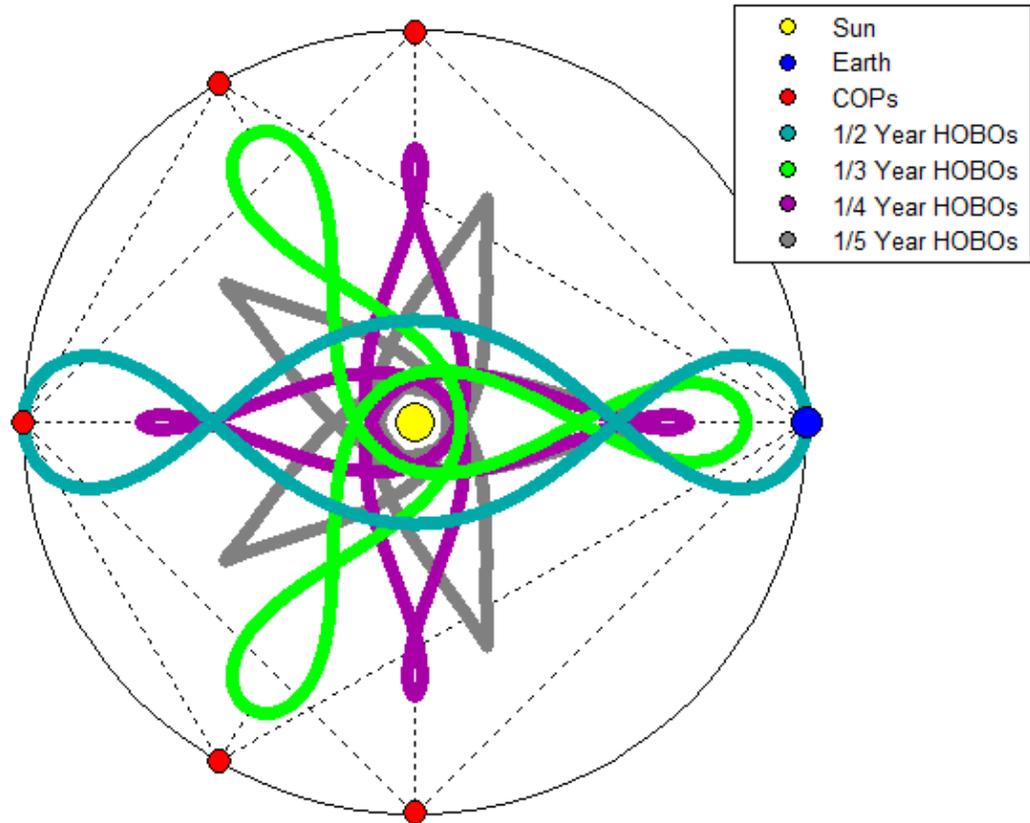


Figure 6-31: Instantaneous distribution of HOBO satellites in an extreme case with many satellites with multiple petal configurations (2, 3, 4, and 5 are shown) and potential locations for COPs.

A system leveraging assets with three petals (green dots) would benefit from two relay COP satellites placed at locations separated by 120 degrees in mean anomaly from Earth, while a system leveraging four petals (purple dots) would benefit from three relay satellites placed at 90, 180, and 270 degrees. In either of these configurations, a satellite placed 180 degrees in mean anomaly from Earth would be able to relay data to any of the other satellites as shown by the light gray dotted lines in Figure 6-31.

(Note: although the figure shows a pattern for five-petal systems (1/5-year period orbits) with gray dots, no designs could reach this orbit with a single maneuver. A tradespace that would consider these would require multiple maneuvers and possibly planetary flybys. Additionally, the distance between this orbit's aphelion and 1 AU is large, so communicating with these satellites would always be more difficult than it would be for systems with longer orbital periods).

Figure 6-31 also provided the PSSs with a different view of the instantaneous distribution of the satellites; the STK model showing a limited number of satellites (Figure 6-5) focused on the orbits of the individual satellites but did not show the emergent behavior of the orbits of many satellites. This is relevant because the definition of "Sampling Resolution" in the value model was dependent only on the number of samples generated by the system; it did not take into account the instantaneous size of the field being sampled.

A stakeholder may choose to redefine the attribute for "Sampling Resolution" to account for this or decide that "Instantaneous Sampling Area" is an additional attribute that needs to go into the value model.

A set of HOBOS in two petals covers a rectangle 2 AU long but only ~0.4 AU wide. Three-petal configurations cover a triangularly shaped area ~1.5 AU on each side, and four-petal configurations cover a square-shaped area ~1 AU on each side.

This picture also demonstrates the importance of both feedback loops among the different phases of RSC and continuous stakeholder involvement. When there is no constructed model for an attribute, a mental model could be changed entirely with new information that comes about during tradespace exploration and performance modeling. In this case, an alternate model was not developed, but a future iteration with different stakeholders may prefer to take “Instantaneous Coverage Area” into account along with (or as opposed to) “Sampling Resolution.”

Multi-era analysis can show the added value or risk mitigation of having the ability to communicate with COPs during the times of year when an individual asset is incapable of communication with Earth given some probability of failure of an asset. Experiments were conducted on homogeneous, Pareto-efficient designs for each payload option in the primary design context (i.e. the designs along the solid red, green, cyan, blue, and black lines in the bottom right of Figure 6-19). Although the previous section showed that cumulative value does not follow a Gaussian distribution, the average from 200 iterations of Monte Carlo analysis for four different failure rates on each design were compared. This process also assumed that contact with a relay satellite costs the same as contact with the DSN; while this is most likely a bad assumption given the cost of putting a relay satellite into heliocentric orbit, this analysis focuses entirely on the difference in *value*, not cost.

Homogeneous designs in two-petal configurations were given the option to communicate either one per year (Earth only) or twice per year (once with Earth, once with a COP at 180 degrees MA). Designs in three- and four-petal configurations were also given the option to communicate three or four times per year with COPs at appropriate locations. The results from this analysis are summarized in Table 6-19.

Table 6-19: Comparison of the difference in cumulative value between satellites that communicate only with Earth and satellites that can communicate multiple times per year through relay satellites.

Probability of Failure Per Satellite Per Year	Change in Value: 2 Contacts/Year	Change in Value: 3 Contacts/Year	Change in Value: 4 Contacts/Year
2.5%	0.942%	0.870%	1.68%
5.0%	2.10%	2.41%	3.21%
7.5%	4.28%	5.86%	8.76%
10.0%	9.09%	8.94%	13.0%

These results show that all designs experience a net increase in average cumulative value no matter what the failure rate is, but the increase in this average rises as both the failure rate and the number of opportunities to communicate rise (the exception being 2.5% failure rate and two versus three contacts per year, but this could be attributed to not enough Monte Carlo iterations). It is important to note that these percentage changes represent *ratio comparisons* of *cardinal* numbers; these does not imply that a certain percent more value is gained by the stakeholder, and when choosing between individual designs the *difference* should be compared. This is being shown to demonstrate the effect of leveraging COPs.

A concrete example of how more contacts can mitigate the risk involved in an individual design is shown in Figure 6-32. This particular design places 18 satellites in three-petal configurations and assumes a high failure rate (10%/year). With only Earth communication (left), there is a 30% chance that the mission will

not meet the minimum lifetime requirements and therefore will deliver an unacceptable amount of value to the stakeholder. With at least one other contact per year (center), that risk drops below 20%, and by communicating with two COPs, that risk is below 15%.

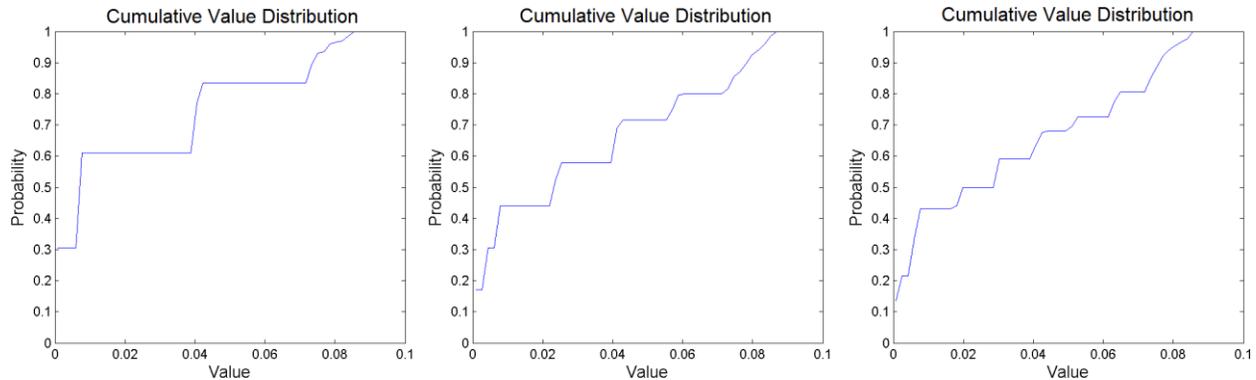


Figure 6-32: Cumulative value distribution comparison of a single design that downlinks data (left) once per year, (center) twice per year, (right) three times per year.

Whether the added cost of building these relay satellites as part of the mission itself, or whether purchasing downlink time from relay satellites is worth the price remains to be seen. It is outside the scope of this case study to model those systems in detail, but future collaboration with other groups working on federated satellites could continue using HOBOCOP as a case study to examine this possibility.

Time Sensitivity of Data

Finally, one of the important value drivers that has not been discussed up to this point is the time sensitivity of the data. The spirit of the commercially available HOBOS® sensors this satellite network borrows from assumes that data is time-insensitive and that there is no rush to collect the data, because the purpose of the data is to analyze solar activity over long periods of time. Tradespace exploration has shown that designs exist that are satisfactory and do not leverage COPs or optical communications.

However, in the event that a different set of stakeholder wishes to use data from the HOBOS to monitor space weather in real time, this paradigm changes. Only the satellites that are nearby Earth during the time of a solar weather event can provide instantaneous data effectively, and this data is arguably the least valuable given how well other satellites near Earth are already monitoring space weather.

Unfortunately, despite best efforts, a concrete optical communications model was not able to be completed in time to compare the value of this technology with other communications technologies. Additionally, without knowledge of what a space weather stakeholder would want, or how much data is necessary at times of solar events, or ways to predict the ability of getting access via the DSN or other optical communications network during daytime, this question becomes even more ambiguous. Satellites using the Ka-band system can still downlink a significant amount of data from far away compared to the X-band, but how much is enough?

Satellites carrying magnetometers only can easily downlink one day’s worth of data in a matter of minutes from a distance of 1 AU, but the link time stretches to several hours or days for the FPA. A much

more thorough look is required to properly analyze the value of communications systems for space weather stakeholders than can be conducted with the resources and expertise available to the author.

6.4 Discussion and Conclusions

HOBOCOP expands the scope of what the RSC method can accomplish compared to the ExoplanetSat case study and considers more science goals, fully utilizes the proposed MAUT value model hierarchy, and varies not only the operational design variables but also the spacecraft design variables. The results lead to many interesting conclusions, not only for the design of a mission to study the Sun but also for examining the effectiveness of the RSC method in a more complex case study.

6.4.1 Discussion of RSC Method Results

The data discussed in the previous sections as a way to illustrate the method in action is only a small fraction of what can be learned from this case study. A discussion of additional findings, observations, and recommendations follows here.

6.4.1.1 The Value of Heterogeneity

The initial tradespace results shows that heterogeneous assets carrying different payloads and operating in different locations drastically increases the utility that can be achieved by a particular design compared to the utility that can be achieved by launching strictly homogeneous assets. Had the scope of this study been limited to a few key design variables and payloads, this would not have been seen. Instead, because of the breadth of the performance model combined with the realization of several key points in the value model, combinations of assets can achieve the highest utility.

There are two key components of the value model that explain this. First, the attribute levels at the bottom of the MAUT hierarchy cannot be satisfied by a single asset alone. One spacecraft can achieve a low perihelion radius to sample sub-Alfvénic solar wind at the expense of a wider total sampling range. Likewise, reducing perihelion distance also requires higher ΔV that could have otherwise been spent increasing the orbital inclination, thereby increasing the vertical sampling range. By considering more than one type of orbit, i.e. not using Walker constellations exclusively but using more than one shape of orbit, these attribute levels do not have to be compromised.

Second, the stakeholders perceive synergies among the science goals, meaning the designs that utilize all three payloads have the potential to generate the most value. However, these payloads do not all need to operate in the same locations; the VFM payload achieves the maximum utility at a polar orbit, whereas the other two payloads achieve their maximum utility as close to the surface of the Sun as they can get. Not only is sending a VFM to an orbit with a close perihelion technically challenging, a highly-eccentric polar orbit would drastically reduce the amount of time the satellite spends above the poles observing, which is a value-driving attribute that was not explicitly captured in the value model. Rather than building one spacecraft to try to achieve everything, multiple spacecraft can be used to split sensitive payloads from non-sensitive ones to build more complete data sets for less overall cost.

Admittedly, it may have been possible to realize that heterogeneous assets may be able to achieve greater science returns in a single science campaign without the structure of RSC. However, RSC offers a method to quantify the benefits of different allocations of heterogeneous assets. Without a value model elicited from stakeholders or a performance model that considers many design alternatives, it would be

impossible to objectively rank design alternatives. Allocating resources without such a method more than likely would result in a decision maker selecting a design that does not achieve as much as other designs could for the same price.

An unfortunate limitation in the exploration of heterogeneous assets is the number of combinations that are possible given a tradespace of homogeneous assets of this size. For n homogeneous designs, there are $(n + 1) * n/2 - n$ possible heterogeneous designs that combine exactly two designs. With 55,600 homogeneous designs, there would be over 1.5 billion additional designs. The analysis on heterogeneous designs only considered a few hundred designs that were Pareto optimal on the total mission MAU or an individual goal MAU among alternative designs with the same payload option. Because the Pareto fronts change depending on the epoch, the set of designs change, therefore multi-epoch analysis cannot be conducted in the same way on these assets.

Additionally, combinations of three designs would generate even more value but require even more computational capacity to explore properly; the author was limited by the available space of the 32-bit version of MATLAB that was used to run these models, which is why not even two designs could be combined in all possible combinations. However, even with these limitations, TSE can give designers insights on what *drives value*, so even without being able to completely enumerate the design space, intelligent observations can be made by exploring the results.

6.4.1.2 Decision-Making with Cumulative Value Distribution

The RSC method showed that some designs have cumulative value distribution functions that are non-Gaussian. In this case study, these designs are typically the ones that achieve low overall utility and have low redundancy within the number of petals, but the results also show which designs are most robust to failure than others. Since distributed missions *should* count on assets failing prematurely and be prepared to mitigate losses due to failures, this output is important when examining the individual designs.

The top left of Figure 6-29 shows a large gap between possible cumulative values. This particular design starts its system lifecycle with 20 satellites, but there is still a nonzero chance that the mission will deliver an unacceptable amount of value. Designs that carry even fewer assets are at a much higher risk as shown by the cumulative value distributions in Figure 6-32. Designs that are at much higher risk such as the ones that make closer approaches to the Sun have a much higher probability of failing before an acceptable level of value is delivered.

Although other methods can be used to assess the failure rate of spacecraft, measuring the degradation of scientific value can only be done with a value model to compare it to. Eliciting this model and understanding that the perception of value delivered over time is definitely nonlinear according to the attribute utility curves is essential for predicting how the delivered value degrades over time. Other methods may assume a Gaussian distribution for the uncertainty; reporting such a distribution only for the mission to suffer from a discrete failure that drastically reduces the value is irresponsible for engineers and could result in further consequences. Accurately communicating the failure modes to the stakeholders mitigates the disappointment when things go wrong.

6.4.1.3 The Value of Communications Path Flexibility

For the first time, the added value of a responsive distributed space system compared to a similar unresponsive system can be objectively quantified. The RSC method was able to precisely show the

added value of communications path flexibility in the HOBOCOP mission if the option to communicate with a satellite relay network was available when assets are known to fail. The comparison of the cumulative value distribution between assets that downlink data once per year compared to assets that downlink data multiple times per year in Figure 6-32 *conclusively* shows how communications path flexibility mitigates the loss of potential value in a system.

If any of the attributes specified a time frame for downlinking data at certain intervals, the utility would be more apparent and would have been captured at an earlier stage; however, because no attribute depends on how often data from individual satellites must be downlinked, MATE alone would not have articulated the added value the way RSC with EEA did.

However, this analysis is still rather limited in scope and immature. Although this analysis only used a crude form of cost estimation, the goal was to show the added value to the system as a result of the path flexibility. Without accurate inputs for the added cost of building a relay network or tapping into a federated network, it is impossible to show the cost difference (left to right on a tradespace diagram) in addition to the value difference (top to bottom) for the responsive or federated systems.

Future research on FSS will go a long way towards helping make these decisions, but no matter what the answer is, the RSC method shows the most promise for articulating the added value for such a network and the price points that must be achieved to make the added value better compared to a more capable system operating without communications path flexibility. Such research is beyond the scope of this case study, but future iterations of the model can incorporate estimates based on more credible research than can be gathered currently.

6.4.2 Recommended Design Solutions

Rather than choose a few specific designs to recommend, it may be more helpful to propose a strategy for implementing designs to reach a certain price point. This reduces the programmatic risk if the budget situation changes significantly during the course of the mission; if the budget is suddenly halved, a valuable mission can still result from wherever in the design process the mission is.

First, there should be at least two satellites with Payload #5 in polar orbits. These quickly add value to Goal #1 while reaching the higher ends of some of the attributes with the other goals. The solar panels and fuel mass should be minimized to reach the polar orbit while still being capable of using Earth flybys to reduce the total propellant mass.

After that, satellites with Payload #3 should be deployed into 3-petal orbits. The exact distribution between number of arcs and the number of assets in a petal should be balanced depending on the trajectory of the Pareto front. As more satellites are added, the risk of a large discrete drop in total value goes down. If cost overruns occur, the total number of satellites can be scaled back down along the same route. While different tracks along the Pareto front require different target inclinations, the size of the propellant tanks can be easily adjusted *during the design phase*; once the mission passes CDR, hopefully no additional changes will be made, but the results from the RSC study can be used to make programmatic decisions about how to adjust the mission if changes occur after CDR.

At a certain point while adding more satellites until the maximum budget is reached, it will be wise to add another satellite with Payload #5. Because this satellite will be more expensive than others, a careful

analysis on exactly when this would occur would need to be conducted. The PSS is fairly satisfied even with two satellites; adding more satellites with this payload does arguably more to reduce the risk from single-point failure than it does to add to the value proposition, but this may be enough to sway certain decision makers.

A valuable design combination of three distinct designs may consist of a few satellites with payload option #5 in polar heliocentric orbits, several (admittedly) heavier satellites with payload option #3 in three- or four-petal configurations, and many, much lighter satellites with payload option #3 in two-petal configurations at higher inclinations. These satellites would all share the same expendable propulsion module because the ΔV requirements would be similar between high inclinations versus low perihelion (and even if not, less propellant can be loaded in those tanks). Sharing the same hardware reduces design and production costs, even if the satellite's propulsion system is considered "suboptimal" as a result.

6.4.3 Comparison to Monolithic Missions

The value proposition for using DSS in the HOBOCOP mission is strong because the proposed designs heavily leverage **simultaneous sampling**. This is self-evident from the attribute definitions in two of the goals; the scientific value comes from being able to sample in multiple locations at the same time across vast distances. It is unfair to even compare a monolithic system to these goals because monolithic systems are inherently incapable of satisfying them.

On the other hand, a monolithic system containing multiple payloads can still leverage **stacked sampling** to raise its value. This is nothing new; modern satellites almost always carry more than one instrument because it reduces overall costs compared to launching multiple independent satellites with redundant subsystems. However, even though one goal can achieve minimum satisfaction with a single satellite, the value of all three goals is strengthened significantly by leveraging multiple assets that even for that goal a highly capable monolithic system still could not deliver as much value as two less-capable assets for a similar price.

One reason why multiple satellites are needed to achieve the minimum satisfaction criteria is because previous missions have already discovered much of what can be discovered with these instruments. There is little added value in repeating those measurements, but there is significant value in repeating those measurements across large distances simultaneously to more accurately characterize and model the phenomena associated with the Sun.

6.4.4 Future Work

The work in this chapter focused mostly on one half of the HOBOCOP mission, the science side. The original work focused on designing and optimizing an interplanetary communications network. Future work on HOBOCOP would bring the two together more completely with better inputs from both sets of stakeholders.

Ongoing research in federated satellite systems will also converge with science missions in the near future. The author was originally working more closely with students and faculty at the Skolkovo Institute of Science and Technology who are currently researching a number of topics related to federated satellites, but there was significantly less collaboration than there needed to be to include a detailed analysis on the value of federated satellite systems in the HOBOCOP mission. Future research on this

mission may yield more interesting results for the field of interplanetary federated satellites, especially when that research is combined with research in miniaturized communications subsystems.

Other stakeholder objectives that must be examined for an interplanetary communications network are the availability of future customers. This may require market research on heliophysics missions and other commercial ventures beyond Earth orbit such as asteroid mining and prospecting. The owners of such a network would need to be able to capitalize on other opportunity, such as future missions to Mars where the network would be vital when Mars and Earth are at opposition and the Sun inhibits high-speed connections.

6.5 Chapter Summary

With this case study, the value of leveraging **simultaneous sampling** was self-evident by the minimum satisfaction criteria of the spatial sampling resolution attributes, but also the value of leveraging **stacked sampling** by operating heterogeneous assets in different environments was made clear in a mission to study important phenomena related to the Sun and the interplanetary medium.

A value model was constructed that explored metrics for evaluating the utility of measurements from multiple locations in a large field where the current state-of-the-art measurements using existing solar observatories is not sufficient for exploring phenomena like the solar wind and the heliospheric current sheet in detail. The proposed modifications to the method helped express various stakeholders' opinions on how multiple data sets can create synergies that increase their perception of value of the data products relative to the value of the sum of the products.

A performance model was constructed with the built-in ability to substitute different technologies to explore what the options for this mission are using today's standard technology versus the near future's cutting-edge technology and how that affects the value proposition. The RSC method is particularly well-suited for such comparisons, and sensitivities can be explored up front that affect the value delivered and cost of the system in the Pre-Phase A phase rather than after the system has been designed. Designs that may be feasible and Pareto-efficient during one technology paradigm may be infeasible in another due to the extreme performance required to achieve the desired orbits in this mission, and analysis on these designs was made possible because of the RSC method.

Tradespace exploration showed the value of heterogeneous assets, especially when goals require assets to enter operational areas where they cannot take some instruments in particular. This mission in particular is well-suited for heterogeneous assets because of the simplicity and protectability of magnetometers and FPAs compared to VFMs, but also because the range that can be swept out by multiple types of orbits compared to a single type.

This case study also demonstrated a method to calculate the added value of a communications network to downlink data from high-risk satellites. Although a full cost model was not applied to additional relay satellites, a federated satellite company could use this method to negotiate with stakeholders who would wish to utilize such a network to determine under what conditions it would be worthwhile to invest in the development and implementation of that network.

Concrete examples of how the expected value of a design can change significantly due to the discrete nature of asset failures and non-uniform distributions of probable cumulative value. Risk is obviously

mitigated with more assets and redundancy, but the consequences are cannot necessarily be treated as continuous depending on the number of assets and the nature of the attributes and how they are satisfied. This case also shows how designs that are feasible and Pareto-efficient but also high-risk with fewer assets could still have a high probability of delivering an unacceptable cumulative value. This demonstrates the importance of incorporating design variables that accurately characterize the lifetime and cost of satellites so that alternatives can be compared equally, not only during tradespace exploration but also especially during multi-era analysis.

CHAPTER 7

ÆGIR

Asteroid Exploration and Geohydrology
for In-situ Resources



Original artwork by Lars Grant-West, © Wizards of the Coast, 2002. Parody modified by the author.

7. Case Study: ÆGIR

*“Swirling masses of rock careening through the black
An orbit unseen by the distant eyes of man
Floating through space, a **distant planetoid**
Collision course, waiting for the moment to destroy*

VEKTOR, “ASTEROID”

*We call out to the beasts of the sea to come forth and join us, this night is yours
Because, one day we will all be with you in the black and deep
One day we will all **go into the water***

DETHKLOK, “GO INTO THE WATER”

The third case study blurs the line between traditional science goals and commercial exploration in the near future. Over the past three years, two separate asteroid mining corporations have been formed and have begun working towards the goal of utilizing the vast resources of space. While these resources include precious metals and rare Earth elements that can be sold for high prices on Earth, asteroids are also rich in water that can be used to support both public and private space exploration. However, before this industry can be self-sustaining, these companies must know what resources are available and how they are distributed in the solar system.

The Asteroid Exploration and Geohydrology for In-situ Resources (ÆGIR) mission, named for the Norse god of the sea and king of sea creatures, aims to characterize the water content of asteroids by conducting in-situ measurements of many asteroids. This mission serves as the *initial exploration* phase to a larger, later mission that will conduct mining operations with knowledge gained during the exploration phase. The goals of the asteroid science and asteroid mining communities overlap in many ways in this mission, and both groups are interested in similar data for different reasons. There are many potential stakeholders and value models, some with notable time dependencies different from other science missions. This mission heavily leverages **census sampling** to achieve its primary mission objectives, but the value proposition also benefits from leveraging **self-sampling** of fractionated spacecraft.

Due to the high computing power necessary to properly study this mission’s orbital tradespace, this mission will forego the final three phases of the RSC method, and instead focus on value modeling and how different stakeholders would perceive value. A performance model can be built with more subject matter expertise on asteroid trajectory modeling and rendezvous, but a credible performance model to make valid conclusions and design recommendations on the results of such a model would require more talent and resources than are available to the author.

7.1 Case Study Introduction and Motivation

On one hand, asteroids have the potential to destroy humanity, just as an asteroid impact may have caused a mass extinction at the end of the Mesozoic Era. The threat of asteroids should not be ignored given how

much damage can be done even by medium-sized impacts; an impact of an asteroid only 1 km in diameter can cause a global catastrophe [394]. Defense against such threats has been studied [395], though no global plan is currently in place to carry out any specific defense strategy or mitigation [396].

On the other hand, asteroids are compelling targets for resources in space that can be leveraged to fuel a new space age. The value of the raw materials inside a single asteroid could be in the trillions of dollars [397]. Meteor samples show high concentrations of valuable resources such as platinum group metals (PGMs) and other rare Earth elements. Others have theorized that space colonies can be built inside asteroids so humanity becomes more than a one-planet civilization [398].

Although this chapter does not intend to explore scenarios that are centuries away from realization, these possibilities do motivate today's entrepreneurial explorers with long-term visions of extended human presence in space. Asteroid mining is a major step towards creating a viable, self-sustaining, space-based economy that benefits not only the people who may be living beyond Earth, but also the people on Earth. In the past three years, two different asteroid mining companies have been founded: Planetary Resources Inc. and Deep Space Industries [399].

In addition to PGMs, water is another valuable resource that is also abundant in some asteroids. Carbonaceous chondrites have lower mass fractions of precious metals (but high water mass fractions) compared to asteroids that would be better targets for PGM mining. Water can be used for human exploration and sustainment as well as propellant for longer journeys. Although it is abundant on Earth's surface, transporting vast quantities of water into space is too expensive and energy-intensive, and mining asteroids for water could be a more energy- and cost-efficient way to supply resources to enable other missions, programs, and exploration. The value proposition for mining water for space-based use appears to be better than the value proposition of mining metals for delivery to Earth in the near-term future, which is why water is the focus of this study.

There are limits on which properties of asteroids can be detected via remote sensing with telescopes from Earth. Reflectance spectroscopy only provides details on the surface of the asteroid, not the composition of the core; although asteroids of different spectral classes are likely to be composed of different materials, there is little evidence that asteroids of the same spectral class are composed of the same material with similar mass fractions [400]. Most asteroids show signs of significant weathering, so any part of the surface that was once volatile-rich could be depleted due to thermal cycling, but the interior could be water-rich. In order to understand the desired characteristics of asteroids, they must be studied up close. This case study examines the value model from the perspective of an asteroid mining company studying many asteroids in order to statistically characterize their properties for future investigations.

(Note: this case study does not deal specifically with the mineral engineering or business side of mining asteroids, which many consider centuries from coming to fruition, but there are complementary goals from which both pure science researchers and asteroid prospectors can benefit.)

7.1.1 Primary Mission Goals and Stakeholders

ÆGIR is a two-phase program to characterize the water content on asteroids. The first phase is an initial exploration phase. Multiple small spacecraft will be sent to survey many asteroids that could be targets for future mining operations. Data from this phase will be used to inform and motivate the design of the second phase. The second phase is a pilot program that will send fewer spacecraft with more advanced

equipment for more in-depth investigations of promising asteroids identified from the first phase. The focus of this chapter is on Phase 1, the Exploration Phase. An approximate timeline of the entire ÆGIR program is shown in Figure 7-1, and key dates for each phase are summarized in Table 7-1 and Table 7-2.

The timeline of this mission is approximate, but it does serve as a good set of bounds on which to base the mission and determine the usefulness of data over time. Ordinarily, scientists may not express strong opinions on the time sensitivity of the data they are pursuing with space missions. In contrast, an asteroid mining company would have investors as stakeholders who desire returns on those investments, though these investors are fully aware of the risk and long-term nature of pursuing such lofty business goals. Understanding the return on that investment is outside the scope of this case study and this chapter, but the desire to complete Phase 1 and move onto Phase 2 does have an impact on the stakeholder perceptions and how stakeholders derive value from this mission. It is important to characterize the time scales that this mission is operating on and how they affect the value derived from Phase 1 not only through the eyes of the investors but also through the eyes of the designers of Phase 2, who are using data to determine what instruments to bring to which targets.

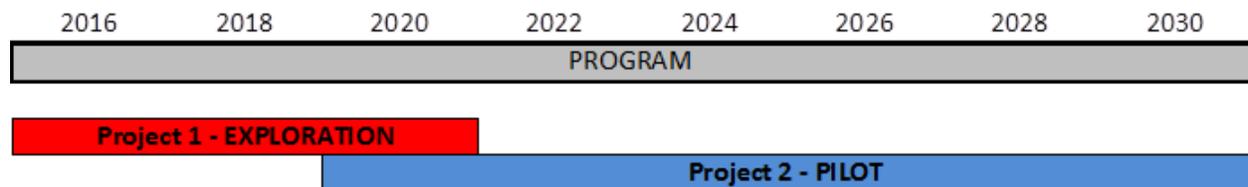


Figure 7-1: Proposed timeline of the exploration and pilot projects of the ÆGIR mission (timeline created by Jared Atkinson).

Table 7-1: Timeline of events during Phase 1 of ÆGIR (schedule created by Jared Atkinson).

Phase 1 - EXPLORATION		
Phase	End Defined By	Expected Duration
Formulation		
Concept Studies	Start of technology funding, preliminary requirements released, program initiation	6 months
Concept & Tech. Dev.	Start of program funding	6 months
Preliminary Design & Technology Completion	Formal requirements release	3 months
Implementation		
Final Design & Fabrication	Assembly begins	9 months
Assembly; I&T; Launch	Lift off	6 months
Operations & Sustainment	Mission complete	2 years
Closeout	Spacecraft dies or decision to put to sleep	6 months
Total		5 years

Table 7-2: Timeline of events during Phase 2 of ÆGIR (schedule created by Jared Atkinson).

Phase 2 - PILOT		
Phase	End Defined By	Expected Duration
Formulation		
Concept Studies	Start of technology funding, preliminary requirements released, program initiation	1 years
Concept & Technology Dev.	Start of program funding	2 years
Preliminary Design & Technology Completion	Formal requirements release	2 years
Implementation		
Final Design & Fabrication	Assembly begins	1 years
Assembly; I&T; Launch	Lift off	1 years
Operations & Sustainment	Mission complete	4 years
Closeout	Spacecraft dies or decision to put to sleep	1 years
Total		12 years

7.1.1.1 Primary Science Goals

There are many science goals that could be accomplished by visiting asteroids that would satisfy a large number of scientific and commercial stakeholders. Asteroids are among the most primitive bodies in the solar system, and their study provides insight into the early formation, origin, and evolution of the solar system. Studying their exteriors also provides data on space weathering, small-body perturbations and astrodynamics, and the evolution of the solar system.

This case study examines a specific set of goals related to understanding the water content of asteroids. These primary science goals of Phase 1 of ÆGIR are shown in Table 7-3.

Table 7-3: Primary Science Goals in ÆGIR

Primary Science Goals	Description
Water Bulk Mass Percentage	How much of the bulk mass of the asteroid is water?
Water Distribution	How is water distributed around the asteroid?
Water Composition and Structure	What physical state is the water in? What is it mixed with?

These three goals are necessary for asteroid miners because it motivates the value proposition of their business model. For the first goal, it is necessary to know how much water is available in order to estimate how much can be extracted from a target. For the second goal, it is necessary to know the distribution and possible variability of the water so that the spacecraft used in Phase 2 has information on where to conduct deeper investigations. For the third goal, it is necessary to know what form the water is in so the energy requirements of the mining equipment can be better estimated.

These goals are also interesting to scientists. Understanding the water content of asteroids may help scientists understand where water on Earth came from. The distribution and form that water may take within an asteroid also gives insight into its thermal evolution and impact history.

Additionally, studying many asteroids up close can help constrain the estimates on material composition and other properties based on reflectance spectra. In the event that one spectral class of asteroids shows low variability in the properties that the spacecraft in ÆGIR would measure, there is a higher chance that reflectance spectroscopy data can be used to make assumptions about other asteroids that have not been visited up close. On the other hand, if the water content of asteroids of the same spectral class has a high variance, a company should not use spectral data alone to make decisions about future mining targets.

7.1.1.2 Primary Stakeholders

ÆGIR is a mission concept that is not formally endorsed by any company, institute or agency, but assumptions have been made on what companies, institutes, and agencies could have a stake in the mission and its outcome. A list of potential stakeholders for the ÆGIR mission is shown in Table 7-4.

Table 7-4: Potential stakeholders and their relative importance in the ÆGIR mission (modified from an original table created by Jared Atkinson).

Level	Example Stakeholder	Objective
Primary	Planetary Resources	Determine water content, composition, distribution, and to demonstrate enterprise viability to investors.
Secondary	Honeybee Robotics	Demonstrate technological capabilities related to water content analyses and designing extraterrestrial mining equipment.
Tertiary	NASA	Cooperate and collaborate with space companies, supporting public perception that asteroid mining is a viable future prospect
Quaternary	MIT	Explore asteroid compositions and advance research into these bodies, demonstrate feasibility of small satellite operations

7.1.1.3 Secondary Mission Goals and Stakeholders

Valuable data can be returned to scientists, not just asteroid miners, using instruments already included in the payload. The same camera, radar, or LIDAR system that would be used for navigation can also be used to map the surface of an asteroid. Rendezvousing with an asteroid for any period of time can help constrain the magnitude of the Yarkovsky–O'Keefe–Radzievskii–Paddack (YORP) effect on asteroid orbits [401]–[404], [396]. Data from instruments built to detect water can also detect other material that is of interest to scientists.

This mission does not focus on PGMs and precious metals; therefore the instrument package may be different from one an asteroid mining company specifically interested in PGMs would use. A different set of instruments could be used if there were an epoch shift that changed the mission goals, but that is outside the scope of this study. However, the methods discussed herein would be identical for formulating the value proposition of this alternative scenario.

7.1.2 Mission Concept of Operations (Single Satellite)

The primary goal of Phase 1 of the ÆGIR mission is to conduct a survey of asteroids by sending a number of spacecraft into heliocentric orbit. Options for deploying the satellites include launching them all at once into an Earth-trailing orbit, launching them all at once into LEO, and launching them one-by-one. When a target asteroid makes an energetically feasible approach, one spacecraft will maneuver to rendezvous with it. Once at the target, it will conduct its scientific investigation; once the investigation is concluded, it will maneuver to the next energetically feasible target. A single asset's mission will end

when it uses all of its fuel and cannot continue traveling to new asteroids, cannot maintain its orbit around an asteroid it is already investigating, or when operations for the mission cease.

The satellite can take several forms. Concept architectures examined in this case study include single-satellite orbiters (as part of a fleet), penetrators, hoppers, and fractionated systems. Because these satellites are working to characterize a large population of targets quickly, the data from multiple visits can be used to draw conclusions about other aspects of the population; this is how **census sampling** is leveraged in ÆGIR. The fractionated satellites can also leverage **self-sampling** to conduct radar measurements that allow for deeper depths of investigation into the interiors of asteroids.

(Note: because a full-scale performance model is not being developed, this chapter will discuss which concepts are likely to be more effective at satisfying stakeholder expectations, which will help limit the scope of such a performance model for Phase 1.)

7.2 Case-Specific Literature Review

There are a number of topics in the published literature that are relevant to this specific case study that were not included in Chapter 2, including the asteroid families and compositions, remote sensing of water, previous missions to study asteroids, and statistical methods for evaluating populations from a small number of samples.

7.2.1 Asteroid Classes and Compositions

As of May 8th, 2015, more than 680,000 minor planets have been detected [405]. These include comets, near-Earth objects (NEOs), and Kuiper belt objects. The vast majority of them are Main Belt objects between the orbits of Mars and Jupiter (1.5 to 5 AU), though more than 12,000 are classified as NEOs, whose perihelion distance is a maximum of 1.3 AU.

Most of what is known for the classification of small bodies comes from meteorite samples and spectral surveys. Asteroids can generally be classified into three primary categories: C-type (carbonaceous), M-type (metallic), and S-type (silicates) [406]. Other subcategories and classifications can be made based on spectroscopic signatures and other observable patterns.

C-type carbonaceous asteroids are the most common type of asteroid, making up 75% of known asteroids. Their name is derived from their high carbon content, which may suggest that asteroids could have delivered organic compounds to Earth. Some subtypes of C-type asteroids have similar chemical compositions to the Sun and the primitive solar nebula except for the absence of volatiles like hydrogen and helium, and hydrated minerals have also been found in meteorite samples [407].

Measurements of C-type asteroids suggest that they are the most porous type, with porosities ranging between 0% and 20% or higher. If water ice in any form is present in asteroids, it will most likely be found beneath the surface, shielded from solar radiation. Pure water ice would be found within the pores between minerals, whereas hydrated minerals could exist even in low-porosity structures. For this reason, ÆGIR will target the C-type asteroids in particular for in-situ measurements.

M-type metallic asteroids are primarily composed of nickel and iron. This suggests that these asteroids were part of a parent body that accreted during protoplanetary formation and was large enough to undergo core formation or some level of differentiation [408]–[411]. This phenomenon is thought to be the result

of an unstable isotope of aluminum (^{26}Al) radioactively decaying into a stable isotope of magnesium (^{26}Mg), which causes heating and melting of metals, followed by gravitationally driven separation by buoyancy. Formation via this mechanism explains why M-type meteorite samples are rich in PGMs compared to Earth's lithosphere and other asteroid types. Although these asteroids may be prime targets for asteroid miners specifically interested in PGMs, they are not the focus for the *ÆGIR* mission.

S-type stony asteroids are the second most common type of asteroid, making up 17% of known asteroids. They are composed primarily of magnesium and iron silicates, which suggest they come from parent bodies that have undergone partial or full differentiation. Meteorite samples of this type are split into two groups, pallasites and mesosiderites. Pallasites are composed of olivine in a nickel-iron matrix, suggesting a more complete melting and differentiation. Mesosiderites are breccias with irregular texture.

Other taxonomies for classifying asteroids will be discussed in Section 7.2.3.1.

7.2.2 Measuring Water on Asteroids

Measurements of water on celestial bodies are typically performed using one of two tools (or both in conjunction): neutron spectroscopy and near-infrared spectroscopy. These techniques measure varying properties of the body at different penetration depths and respond to different physical phenomena related to the presence and form the water is in. New research on the dispersion curves at ultra-low and medium frequencies (100Hz to 1 MHz) has shown that radar may also be useful for studying the composition of larger (kilometer-sized) asteroids.

7.2.2.1 Neutron Spectroscopy

Neutron spectroscopy techniques are used to determine the bulk mass percentage of hydrogen on planetary surfaces. Neutron spectrometers measure the gamma radiation being emitted from, and the relative neutron output of, the surface of a planetary body. As galactic cosmic rays enter the surface, inelastic scattering and radiative capture produce gamma rays that are radiated outward [412]. Because a hydrogen nucleus has roughly the same mass as a neutron, a single collision between the two can cause the neutron to lose the majority of its energy. In contrast, a collision with a carbon atom can reduce the energy by a maximum of 28%. The energies and fluxes of the received gamma rays act as a spectral signature that can be used to identify the residual nucleus of the atoms.

Neutron spectrometers have been flown on previous planetary missions, including Lunar Prospector, MESSENGER, and Mars Odyssey [413]–[415]. Currently, Dawn and the Mars Science Laboratory are using neutron spectroscopy to map the hydrogen distribution near the surface of Vesta, Ceres, and Mars [416]–[419].

The depth of investigation of a neutron spectrometer is typically on the order of centimeters to tens of centimeters, though the exact depth depends on the density and location of hydrogen that is present (i.e. lower hydrogen content allows for deeper investigation because there are fewer neutron collisions). This is a far greater depth of investigation compared to X-ray spectroscopy, which is limited to depths on the order of microns.

While neutron spectroscopy does provide an indication of the amount of water that is present near the surface of a planetary body, it does not provide information on whether the water is bound as ice or within hydrated minerals. To discern this, other measurement techniques are required.

7.2.2.2 Near-Infrared Spectroscopy

Near-infrared (NIR) spectroscopy characterizes elements and compounds based on the frequency dependence of molecular absorptivity, which is a signature for identifying molecules and not just atoms. NIR typically operates between 0.5-3.5 μm wavelengths. Water and minerals that contain hydroxyls have absorption features throughout this wavelength range, but the most prominent bands occur at 0.7 μm , 3.0 μm , and 3.1 μm [420]. The absorption features at 0.7 μm and 3.0 μm appear to indicate hydrated minerals, and the 3.0 μm feature is particularly associated with phyllosilicates. The absorption feature at 3.1 μm is strongly indicative of water ice that is not bound to other minerals [421].

However, NIR reflection measurements only sample the upper microns of the surface of the planetary body and do not give a good estimate on the composition of the material immediately below the surface. It is possible due to space weathering and other effects that the material beneath the surface is completely different from the surface, especially regarding volatiles like water that may have sublimated and outgassed due to thermal shifts.

NIR spectrometers have flown on a variety of missions, including Dawn, Hayabusa, MESSENGER, Cassini, NEAR Shoemaker, among others [422].

7.2.2.3 Radar Observations

Radar studies offer a way to increase the depth of investigation on planetary bodies [423]. As radio waves transmit through the asteroid, the material properties of the interior can cause refraction and dispersion, while boundary layers can create reflections. Spacecraft on either side of an asteroid can send and receive carefully timed radio pulses; the received signal will have changed phase, velocity, and magnitude depending on the properties of the material attenuating the signal [424].

The Comet Nucleus Sounding Experiment by Radio wave Transmission (CONSERT) leverages this technique with the Rosetta spacecraft and its small lander Philae [425]. Rosetta transmits radio pulses at 90 MHz that are received by Philae. As Rosetta orbits the comet, the signals received by Philae can be used to build a map of the interior of the comet. As of this writing, results from this experiment have not been published because Philae is stuck without power, although it should return to duty soon [426].

Other monolithic spacecraft have used radio signals from Earth to conduct similar studies. The MESSENGER spacecraft observed radio signals generated by the Arecibo radio telescope that were bounced off of Mercury [427]. The regions that were radar bright also corresponded to locations where neutron spectroscopy data showed evidence of water, providing even more confirmation that ice could exist in the craters of Mercury [414]. This is a clear example of **stacked sampling** between synergistic data sets that creates more value than the sum of the sets.

One of the major differences between the radar responses of hydrated minerals and water ice is the dispersion curves due to differences in the dielectric constants of the materials at different wavelengths [428], [429]. Radar sounding at frequencies ranging between 100 Hz and 10 MHz can distinguish between hydrated clays and water ice due to differences in dielectric dispersion. CONSERT only uses one frequency and therefore cannot distinguish between ice and hydrated clays, but a system that can generate and receive a larger range of frequencies has a better chance of constraining the composition. Generating such low frequency signals may be more feasible with arrays of distributed small satellites compared to

monolithic satellites with larger antennas. One downside is that longer wavelengths at lower frequencies reduce the sampling resolution that is possible.

7.2.3 Similar Science Campaigns and Missions

There have been a number of observation campaigns and science missions dedicated to the study of asteroids. Other missions have made close flybys of asteroids on their way to their primary target of study but still conducted notable scientific investigations.

7.2.3.1 Asteroid Classification Taxonomies

Giuseppe Piazzi discovered the first and largest asteroid, Ceres, in 1801 (though it is now designated as a dwarf planet). Since then, many campaigns to find and classify asteroids have been conducted, resulting in a number of classification taxonomies.

One of the most widely used asteroid taxonomies is the Tholen classification, proposed in 1984 [430]. This classification is based on broad band spectral observations of 978 asteroids from 0.31 μm to 1.06 μm observed in the Eight-Color Asteroid Survey (ECAS) [431]. This survey also measured asteroid albedos. In the Tholen taxonomy, there are 14 types, with most falling into one of three broad categories: C-group (carbonaceous), S-group (stony), and X-group (which includes metal).

Another taxonomy is based on results from the Small Main-Belt Asteroid Spectroscopic Survey (SMASS) classification proposed in 2002 [432], [433]. This survey observed 1,447 asteroids with much higher spectral resolution than ECAS but with a smaller total bandwidth (0.44 μm to 0.92 μm) and no albedo considerations. Once again, there are three broad categories that most asteroids fall into, but there are 24 individual types.

The Bus-DeMeo taxonomy was proposed in 2009 based on more recent spectral observations [434], [435]. This survey observed 371 asteroids on a spectral range of 0.45 μm to 2.45 μm and has 25 classes. The categories from this most recent classification are used as the basis for the value modeling process later in this chapter and are shown in Figure 7-2.

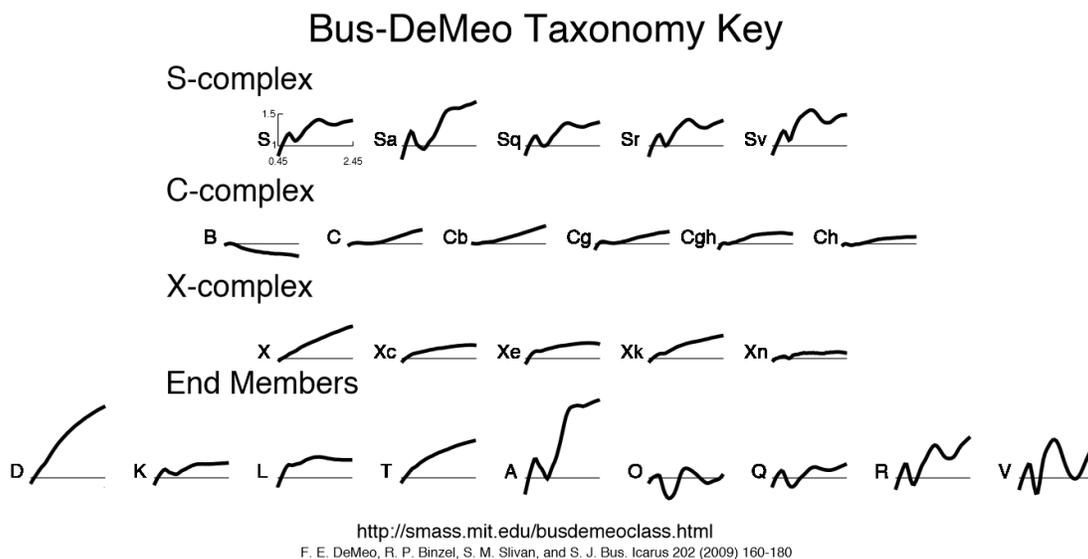


Figure 7-2: Bus-DeMeo taxonomy of asteroids (reprinted with permission from [435]).

Within the Bus-DeMeo taxonomy, the C-complex asteroids are the most likely to contain water, though there is strong evidence that D-type asteroids (one of the End Member types) may contain organic molecules. The X-complex consists of mostly metals, and the S-complex is mostly silicate asteroids.

7.2.3.2 Asteroid Ground Observations

When asteroids pass by Earth at sufficiently close ranges, radar telescopes can be used to conduct more in-depth investigations in comparison to infrared and visible observations. The Arecibo Observatory has studied dozens of asteroids over the past three decades [403], [436], [437]. These studies include mapping asteroid surface structure and placing constraints on the strength of the YORP effect.

It is difficult to detect NEOs whose orbits are inside of Earth's orbit because their backsides are not lit by the Sun, so their signal is weak, and because the Sun generates too much noise. Many of the asteroids that are only discovered when they make close approaches to Earth come from Earth's day side. In order to give Earth more warning, the B612 Foundation has proposed building an infrared telescope to be sent to a near-Venus orbit [438]. This orbit allows the telescope to conduct observations of the Sun-lit side of asteroids and observe with a synodic orbit relative to the asteroids that is different from Earth, which would eliminate the possibility of decade-long waits to conduct follow-up observations of potentially hazardous objects (PHOs) if their semimajor axes are close to 1 AU.

The information that can be gathered through remote sensing may not be sufficient for some science goals. High-quality observations can deduce the spectral class, spin rate, and sometimes the shape and size. Without other data obtained through special circumstances (e.g. close approaches to Earth where radio telescopes can use active sensing combined with NIR spectroscopy data to deduce surface properties), remote observations do not give an accurate estimate of mass, structure, or interior composition (though mass can be reliably estimated by observing interactions with planets or other asteroids). This is a strong case for why asteroids need to be studied up close with spacecraft.

7.2.3.3 Asteroid Visit Missions

Many spacecraft have been in proximity of asteroids, either while conducting a close flyby or rendezvous. The Galileo spacecraft was the first to fly by an asteroid [439], [440]. On its way to Jupiter, it passed two asteroids, Gaspra and Ida. Ida has its own small moon, named Dactyl, the first ever discovered case of such a phenomenon. This mission also provided the first space-based observations of a comet colliding with another planet [441], [442].

The Cassini spacecraft flew within 1.6×10^6 km of the asteroid 2685 Masursky [443]. At this range, it could spatially resolve the asteroid, which provided its size. Cassini also measured the electric charge of interplanetary dust grains in the asteroid belt on its way to Saturn [444].

The NEAR Shoemaker spacecraft was the first mission dedicated to studying asteroids. Its primary objective was to study the bulk properties, composition, mineralogy, morphology, internal mass distribution, and magnetic field of Eros [445], [446]. The spacecraft made a soft landing on Eros, and it also flew by the asteroid Mathilde.

The Stardust spacecraft sampled dust from the coma of the comet Wild 2 and other cosmic dust and returned these samples to Earth [447]. It was the first sample return from beyond cislunar orbit and successfully returned on January 15, 2006.

The Hayabusa spacecraft was the first to return a sample from an asteroid, Itokawa, though the sample consisted mostly of small dust grains blasted from the surface by tiny projectiles (that may have failed to deploy at all) [448]–[450]. The spacecraft conducted investigations on Itokawa’s shape, terrain, reflectance, mineral composition, and gravity. The spacecraft had a small rover, MINERVA, which failed to reach the surface.

The Deep Impact spacecraft flew by and impacted comet Temple 1 [451]–[453]. The spacecraft also flew by Hartley 2. Follow-up observations of Temple 1 were conducted through a mission extension of Stardust, which was renamed to New Exploration of comet Tempel 1 (NExT) [454], [455]. This mission was made possible because Stardust had propellant reserves that were leveraged to capitalize on the opportunity to observe the comet from close range again after the Deep Impact mission.

The Rosetta spacecraft is currently conducting investigations of the comet 67-P [456]. En route to 67-P, it flew by the asteroids Lutetia and Šteins. It also carried a lander called Philae to conduct in-situ testing of drills and harpoons in addition to the CONSERT instrument [457]. However, contact was lost because the lander became stuck in a shaded area, and it could not recharge its batteries via its solar panels.

The Dawn spacecraft is the first spacecraft to have orbited (not just flown-by) two extraterrestrial bodies and has investigated two of the largest asteroids in the solar system, Vesta and Ceres [458], [459]. The spacecraft is notable for its use of ion thrusters – it has a ΔV budget of more than 10 km/s after separation from the launch vehicle, far more than any previous spacecraft.

Two notable missions have been proposed to continue asteroid exploration. Hayabusa 2 will improve on Hayabusa’s mission by visiting asteroid 1999 JU3 and returning a sample to Earth [460]. The OSIRIS-REx spacecraft will also return a sample to Earth from the asteroid Bennu [461]. Both of these missions are scheduled to launch in 2018.

Only 23 of the known 680,000 asteroids and comets have been visited by spacecraft in some form (15 asteroids, 10 comets). More asteroids have been studied because they flew close enough to the Earth to conduct active radio observations than have been studied by intentionally sending spacecraft to asteroids. Even among those 15 visits by spacecraft, some were at a distance that was so far away that the quality of observations was no better than what could be achieved from the ground.

7.2.3.4 Estimated Number of Asteroids in the Solar System

While approximately 680,000 small bodies have been discovered in the Solar system, the actual number is far greater. Estimates for the total number of asteroids are based on the rate of discovery of new asteroids and the observational limits of the surveys that have been conducted previously. A 2010 report on planetary defense estimates there may be as many as 30 *million* asteroids with diameters greater than ten meters [395]. The discovered (as of 2009) and assumed populations are compared in Figure 7-3 along with survey limits and select impact energies for comparison.

Knowing that the population of potential targets will change during the course of the mission lifecycle is crucial in understanding the value proposition. As more asteroids are discovered, responsive systems will be able to capitalize on the opportunity to change their trajectories to visit more valuable targets for exploration. Designs with little flexibility (e.g. only enough fuel to visit the originally planned asteroids, fewer spacecraft overall, etc.) will not be able to capitalize on these opportunities and adjust their mission

accordingly. Designs that can respond dynamically will need to quantitatively express the added value to justify the added cost of this flexibility in order to be chosen by decision makers.

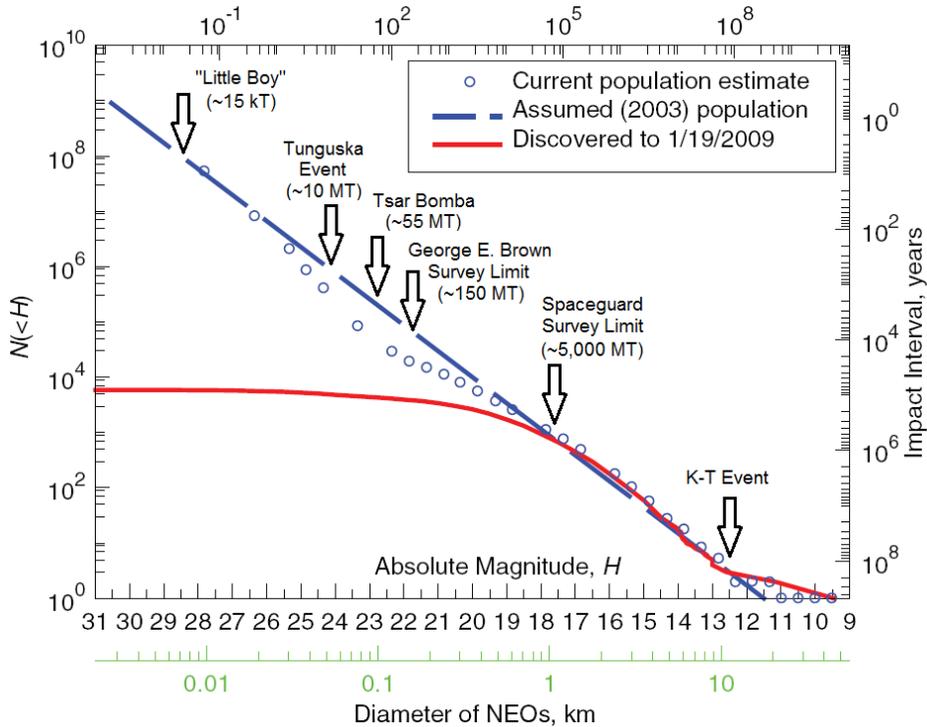


Figure 7-3: Comparison of number of the number of discovered asteroids versus the number of asteroids that are assumed to exist (originally from [395], modified in [396] and reprinted with permission)

7.2.4 Performance Modeling in ÆGIR

Although constructing a performance model for ÆGIR is beyond the scope of this case study and requires more resources than are available, there are several key aspects of performance modeling that can be addressed as they relate to the overall problem.

7.2.4.1 The Traveling Asteroid Prospector Problem

One of the most famous and intensely studied problems in optimization is the Traveling Salesman Problem [462]. The goal is for a salesman to travel to every city on a map exactly once and return to the starting point in the minimum travel distance or time. The problem is NP-complete, meaning there is no known solution that can be found quickly, and the time for an algorithm to find the solution increases very quickly as the size of the problem increases.

While it is completely infeasible for a single spacecraft to visit all 680,000 minor planets, many have worked on problems involving low-thrust trajectories among sets of asteroids, and the optimization techniques required are similar [463]. This problem has been studied since the 1970's [464]. Trajectories leveraging planetary flybys have been examined in detail [465], [466]. Others have developed an automated search procedure to generate optimal low-thrust trajectories to tour the Trojan asteroids [467].

The Global Trajectory Optimization Competition (GTOC) has twice examined paths among asteroids. The 4th competition, held in 2009, had participating teams find trajectories to perform as many flybys of NEAs as possible in a ten year period before rendezvousing with a final asteroid [468]. The competition

assumed a spacecraft with a dry mass of 500 kg, propellant mass of 1,000 kg, I_{SP} of 3,000 s, and a maximum thrust of 0.135 N. Moscow State University's team, led by Ilia S. Grigoriev, won the competition by finding a route that visited 44 asteroids in the ten-year period. However, this competition was for flybys only, not rendezvous. The goals for this case study can only be accomplished with rendezvous and some amount of time orbiting the asteroid.

The 7th GTOC in 2014 examined exploration of the Main Belt using distributed satellites [469]. Given a mothership and three probes, the goal was to find trajectories to rendezvous with as many asteroids as possible and explore them for 30 days at a time. The probes were required to dock back with the mothership within 12 years of the initial launch, and no probe could be gone from the mothership for more than 6 years at a time. A team from JPL, led by Anastassios Petropoulos, won the competition by visiting 36 asteroids with these three spacecraft. One caveat of this solution is that the target data set consisted of only 16,256 possible targets; there are more than 600,000 objects with semimajor axes between 1.8 and 4.8 AU. Additionally, this competition did not take into account varying solar power availability as a spacecraft moves further from the Sun.

7.2.4.2 Other Asteroid Prospectors

The instrumentation aboard the spacecraft described in Section 7.2.3.3 can serve as the basis for some parametric models of the instruments required for ÆGIR. There are a few other notable examples that can also be included.

The Clementine mission is considered somewhat of a prospecting mission to study the Moon [470]. However, a primary computer caused it to use up its propellant and enter a spin, so it did not continue its mission to study the near-Earth asteroid Geographos.

A group at Princeton has developed a model for a small NEO prospector [471]. This 30 kg spacecraft carries a multi-spectral imager as its primary payload. It starts in LEO and uses electric propulsion to escape Earth's orbit using a spiral trajectory over a period of one year. It communicates over the S-band frequency range.

7.3 The Value Model for the ÆGIR Mission

The value model for the ÆGIR mission is more complex than the previous case studies and does not strictly follow the exact MAUT hierarchical structure discussed in Section 4.2.3.2. Modifications to this structure will be discussed where appropriate. Additionally, there are a number of alternative scenarios that could change stakeholder expectations drastically that must be accounted for in epoch characterization. For this reason, a responsive system that can react to shifts of such a mission is far more likely to be valuable compared to rigid design with no ability to change course during the course of the system lifecycle.

The value model was elicited from interviews and discussions with two stakeholders in particular: a person acting the part of the head of an asteroid mining company, and a leading scientist in the field of asteroid studies who acted as the Primary Science Stakeholder (PSS).

7.3.1 Value Modeling Process 1: Value-Driving Context Definition

Table 7-3 summarizes the top-level goals of ÆGIR, and Table 7-4 lists potential real-world stakeholders associated with this mission. The top-level goals are discussed in more detail in Section 7.3.2. The

resources available to asteroid miners are not insignificant; billionaires are backing these companies hoping to reap trillions of dollars in long-term returns on investment.

Because these investors expect a return on their investment during their lifetime, there is a strong desire to finish the mission sooner rather than later; if a mission like Voyager was a business venture, the time delay in monetary returns would more than likely provide a negative net-present value (NPV) and would not have been an intelligent investment.

Additionally, the time component for scientific data returned from Phase 1 to inform the design of Phase 2 creates a stronger constraint for the non-scientific stakeholders of this mission than would otherwise be applied through the natural formulation of MAUT. This will be made clear in the next section.

7.3.2 Value Modeling Process 2: Value-Driven Design Formulation

Interviews were conducted with the stakeholder representing the mining corporation to define and characterize the top-level goals of the mission, to identify the possible instruments that could be used and compared to one another to satisfy those goals, and to propose a set of attributes that could be used to accurately characterize and rank the utility of a design across a number of possible concept architectures. The results of work are summarized in Table 7-5.

The standard MAUT hierarchy that was proposed in Section 4.2.3.2 was used to organize the top-level goals, goal-level attributes, and instrument attributes. A hierarchy showing all subcategories of the first science goal is shown in Figure 7-4. The hierarchies beneath the other two goals look similar and have been omitted to save space.

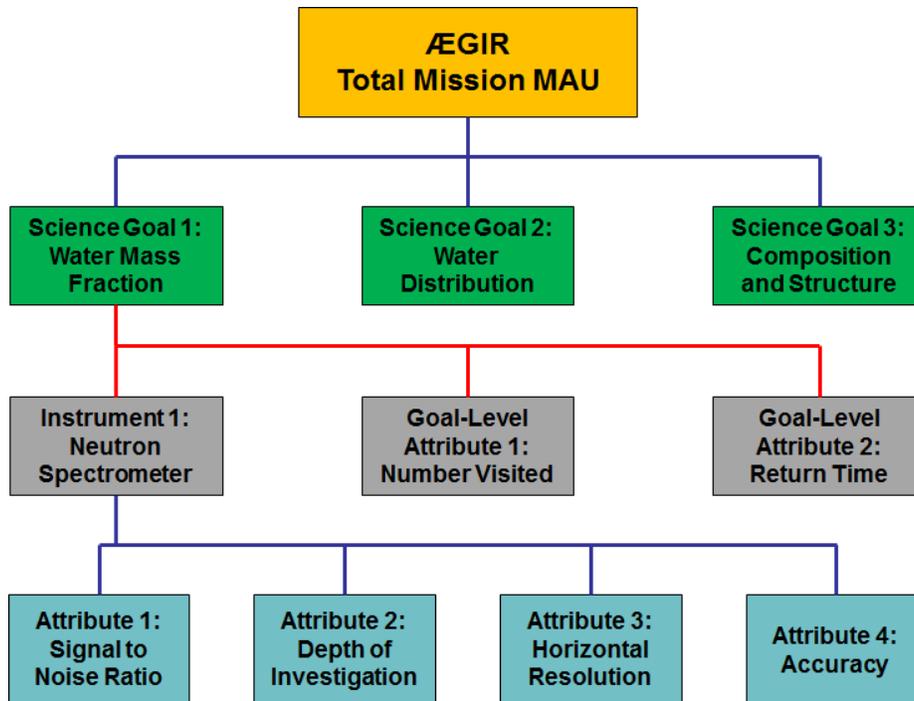


Figure 7-4: MAU hierarchy for ÆGIR, showing all levels for Science Goal #1. The parts of the tree connected by blue lines are easily integrated into a standard MAU hierarchy. The parts of the tree connected by red lines are not, and an alternative value-ranking method must be developed.

Table 7-5: Potential concept architectures, instruments, and instrument attributes for the Exploration Phase of ÆGIR (created in conjunction with Jared Atkinson).

Science Goal	Potential Vehicles	Potential Instruments	Source	Attributes
Bulk Mass Percentage	Orbiter	Neutron Spectrometer	Passive	SNR
				DoI
				Horizontal Resolution
	Penetrator	Neutron Spectrometer	Active/ Passive	Accuracy
				SNR
				DoI
Structure	Orbiter	Near IR Spectrometer	Passive	Horizontal Resolution
				Field of View
				Resolving Power
		Gamma Ray Spectrometer	Passive	SNR
				SNR
				DoI
	Penetrator	Near IR Spectrometer	Active	Horizontal Resolution
				Accuracy
				Resolving Power
		Gamma Ray Spectrometer	Active	SNR
				SNR
				DoI
		Reflectance Spectrometer (Sample Grab)	Active	Vertical Resolution
				Accuracy
				SNR
Distribution	Orbiter	Radar	Active	Spatial Resolution
				Precision
	Hopper	Radar	Active	Bandwidth
				Resolution
				Center Frequency
	Rover	Radar	Active	Bandwidth
				Resolution
				Center Frequency
	Penetrator	Radar	Active	Bandwidth
				Resolution
				Center Frequency

However, subsequent interviews where the discrete nature of the data return and how returns added later in the mission were seen as far less valuable but still acceptable makes ranking a single design difficult. For instance, two exact copies of the same design could conduct different trajectories: one that was front-loaded with data return on low-value targets; another that was back-loaded with data return on high-value targets. The standard formulation of MAUT may not accurately model the stakeholder’s perceptions of where the balance is between these extremes.

Additionally, the stakeholder was unsure about which targets were valuable or how many needed to be visited within each of the interesting types in the Bus-DeMeo taxonomy. The PSS was interviewed later to learn more about this (see Section 7.3.2.4). However, the mental value of one stakeholder that ranks the time-sensitivity of the data return is on the same level as the mental value of another stakeholder that ranks the importance of visiting multiple targets. While it is possible to have multiple stakeholders provide inputs for different levels of the hierarchy (see Section 6.3.1), it is ill-advised to have them provide inputs for the same level.

A standard MAUT model could be applied and (possibly) validated with a true performance model, but because of the statistical nature of trying to learn more about a large population by sampling a few, it is more likely that a value model can be constructed with previously defined statistical models to create some other metric that better maps the design space to the stakeholders' value models. However, the MAUT hierarchy is still applicable at the top and bottom levels, represented by the blue lines in Figure 7-4. The center level, connected with red lines, will be defined using a different method. The next two subsections present the results from the interviews for the top and bottom levels.

7.3.2.1 Top-Level Goal Attribute Weights

The stakeholder representing the mining corporation was interviewed using the ASSESS software to find the relative weights among the three top-level goals. The results are shown in Table 7-6.

Table 7-6: Goal-level attribute weights elicited from the initial stakeholder interview using ASSESS.

Science Goal	Initial Attribute Weight k_i
Water Mass Percentage	0.67
Water Composition	0.44
Water Distribution	0.25

The attribute weights on these top-level goals show that there are some substitution effects among these goals because $\sum k_i > 1$. However, a later interview showed that when the stakeholder saw the mapping between different SAU values for these goals, there was some hesitation to claim it represented the way he felt about how having multiple goals well-satisfied.

Instead, an alternate question was posed to the stakeholder: “Imagine you are a broker selling science goals. You could sell individual goals at the highest level of satisfaction (utility equal to 1) or combinations of goals to buying scientists. Let’s say that, based on the previous interview, Goal #1 cost \$67 (where “\$” is some imaginary unit of science currency), Goal #2 cost \$44, and Goal cost \$25. How much would you sell a package that contained the highest satisfaction of Goal #1 and Goal #2 for? What about #1 and #3? What about #2 and #3? What about all three goals?”

Ideally, if the MAUT interview went perfectly, the answer to the last question would be “\$100,” because satisfying all three goals to their individual highest levels of satisfaction should result in the definition of the highest level of satisfaction in MAUT, which is $MAU = 1$.

However, the stakeholder’s opinions for the final question did not map to $MAU = 1$. Instead, the combinations of goals mapped to the values shown in Table 7-7. The differences that are expected from simple addition are also shown, as is the expected MAU of some combination of goals if the sum of the

value of each pair of goals was divided by the sum of the value of all goals. The initial k_i weights were then linearly scaled by the sum of all three goals to find the final k_i weights, shown in Table 7-8

Table 7-7: Results from hypothetical “science broker” interview selling combinations of goals.

Goal Combinations	Summed Value	Perceived Value	Difference	Normalized Sum MAU
Goal #1 + Goal #2	\$110	\$150	\$40	0.83
Goal #1 + Goal #3	\$92	\$110	\$18	0.61
Goal #2 + Goal #3	\$69	\$80	\$11	0.44
All Science Goals	\$136	\$180	\$49	1.00

Table 7-8: Goal-level attribute weights scaled from the results of the “science broker” interview.

Science Goal	Final Attribute Weight k_i
Water Mass Percentage	0.37
Water Composition	0.24
Water Distribution	0.14

From this additional line of questioning and its results, there appears to be synergy among the science goals, which is now reflected because $\sum k_i < 1$. This means that the returns on accomplishing an additional goal go up for each unit of investment into that goal more than they would if no other goals were being satisfied at that point.

These new weights for MAU match up relatively well with the expected MAU from the previous exercise. This correlation is shown in Figure 7-5. The left graph shows the total mission MAU as a function of Goal #2 and Goal #3 MAU when Goal #1 MAU = 0, and the right graph shows the total mission MAU as a function of Goal #2 and Goal #3 MAU when Goal #1 MAU = 1. The labeled points correspond to what the normalized sum of goals is expected to be as shown in the rightmost column of Table 7-7. The sum of the maximum utility of #2 and #3 is within 0.01 of the expected utility, and the sums of the maximum utilities of #1 and #2 and #1 and #3 are within 0.09 and 0.03, respectively.

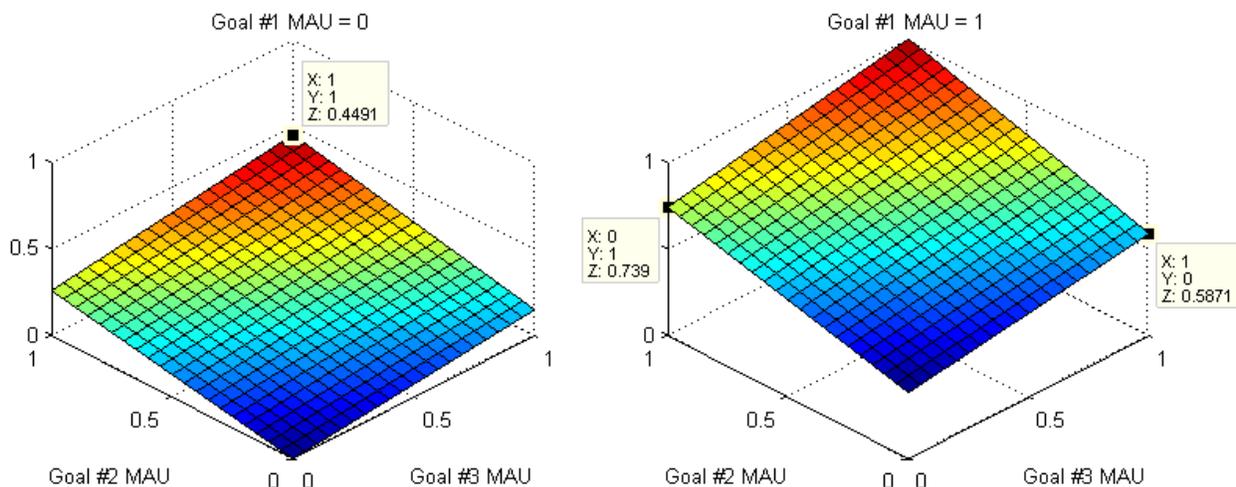


Figure 7-5: 3D utility functions showing the Total MAU when (left) Goal #1 MAU = 0, and (right) Goal #1 MAU = 1

Regardless of whichever interpretation is more correct, these weights show that there is little relative value in radar studies to characterize the distribution of water on asteroids. A series of fractionated systems that leverages **shared sampling** to conduct single-frequency radar occultation mapping and no other instruments would have a maximum $MAU = 0.14$, whereas individual satellites carrying the most satisfactory neutron spectrometers could reach $MAU = 0.37$. In this scenario, justifying the use of an expensive, complex system of free-flying spacecraft to accomplish the least important goal is a losing value proposition; leveraging shared sampling does not appear to increase the net benefits enough. However, the potential value of shared sampling will be revisited in Section 7.3.3.1.

7.3.2.2 Instrument-Level Attributes

It is advised to begin interviews with the subject imagining a black box from which scientific returns come through, so that the subject can describe what aspects of the returns make it valuable; however, in this case, the stakeholder started with a broad range of concepts to understand what capabilities of these concepts could be used to describe differences in the returns from different architectures so that the stakeholder could understand the attributes that drove the perception of value. The potential concepts and instruments are shown in Table 7-5.

The interview then shifted focus to qualitatively describe what was important about the measurements. It quickly became clear that for Phase 1 of the ÆGIR mission, some concepts could provide novel, ground-breaking data, but it was not the right data for the time. Revelations like these could help constrain the design space by limiting the number of concepts that are generated and modeled because it places constraints on what is and is not acceptable without mathematically altering the value model.

One particular constraint is that the goal is to study the *whole* asteroid, not just a subsection of it. For this reason, penetrators *alone* would not be capable of providing the minimum satisfaction for studying the bulk properties of the entire asteroid. This was not identified as a goal-level attribute because it was assumed that the entire asteroid would be studied in Phase 1. There still may be use for penetrators that carry neutron spectrometers because the depth of investigation, which was identified as being a value-driving instrument-level attribute in Table 7-5, can be much greater compared to an orbiter.

However, once the minimum acceptable and maximum perceived utility levels were identified, it was clear that a neutron spectrometer in an orbiter would be sufficient to achieve the maximum desired depth of investigation. From this perspective, a penetrator does not seem to be a smart choice; it can only study one target and is not reusable, it does not study the entire target, and it does not provide any greater information compared to a cheaper alternative like an orbiter or hopper.

If one were using these assumptions, it would be wise not to waste time developing a model for a penetrator concept. However, these assumptions could change; epoch-era analysis was made to handle changes like these, and as such, more should be invested in the concept exploration phase of the design process to ensure that there are not situations where suddenly penetrators can boost the value of the mission more so than other concepts. Such changes will be discussed in Section 7.3.3.1.

Aside from mentioning the maximum perceived utility for depth of investigation, the attribute weights and utility curves are not pertinent to any future discussion herein because there is no performance model to compare to. For this reason, they have not been included in this work.

7.3.2.3 Lifetime Attributes

Now that the top and bottom levels of the MAU hierarchy shown in Figure 7-4 have been defined, the attention turns to the center level where a different approach is necessary to model the effect of the number of asteroids visited in a given period of time.

Because this is a mission with implications for future business ventures, there is a great need to finish the exploration mission quickly so investors can have a better chance of seeing return on their investment that beats the NPV of other potential investments (though these investors are well aware of the risks and the long-term nature of this investment). Unlike the other case studies in this work, the utility of the attribute for mission lifetime *decreases* with increasing time. The attribute utility function that was elicited from a stakeholder interview is shown in Figure 7-6.

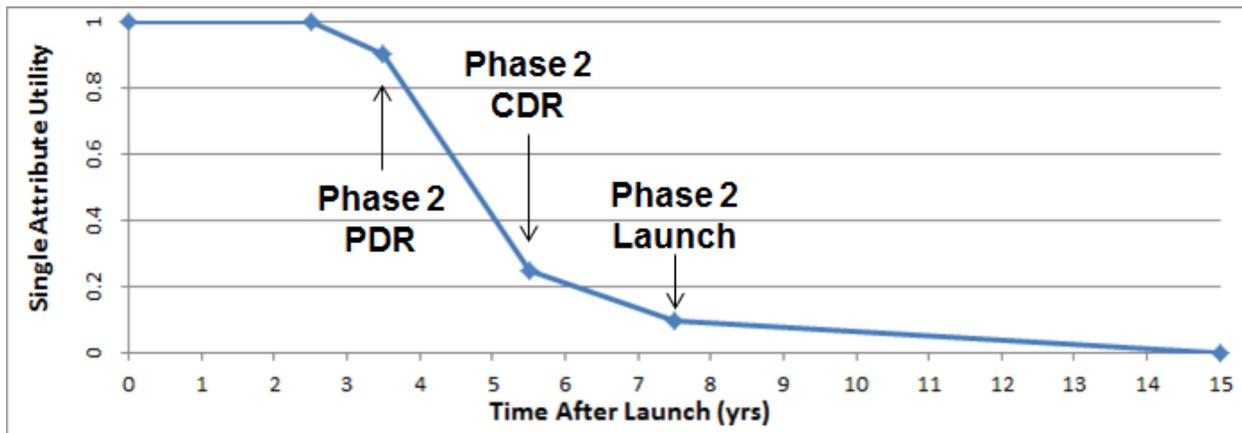


Figure 7-6: Lifetime attribute utility function for AEGIR Phase 1, labeled with Phase 2 milestones.

The points on this curve correspond to very specific moments in the system lifecycle after the launch of the exploration project (see Figure 7-1). Maximum utility is achieved when data returns within 2.5 years of the launch of the exploration project. The decrease in utility represented by the point $t = 3.5$ years is the date of the Preliminary Design Review (PDR) of the pilot program. Information received before this point can help shape the design in effective ways up until that point. However, after PDR, the ability to use the data to change the plan or design of the pilot program decreases significantly because design changes after this point become increasingly expensive.

The Critical Design Review (CDR) of the pilot project (Phase 2) is scheduled for 5.5 years after the launch of the exploration project (Phase 1). After this point, the information from the exploration project is even less valuable in aiding the direction of the pilot program, and it is nearly valueless after 7.5 years when the pilot program launches. However, the utility does not go to zero because there is still some, albeit low, value in gathering information for other scientists from the perspective of the company, and there is still a possibility of gathering information for later programs that may affect the business case of an asteroid mining company.

7.3.2.4 Statistical Attributes Among Spectral Types

Finally, the point where the value model for this mission differs drastically from a traditional exploration program has been reached—understanding the balance in the perception of value between high-fidelity observations with large, expensive instruments visiting few targets versus low-fidelity observations

visiting many targets is the key to choosing the best architecture for a census sample. For information on how value in this area is perceived, the PSS was interviewed.

Perceived Utility Among Spectral Types

The PSS prefers the Bus-DeMeo taxonomy shown in Figure 7-2 as a means of classifying asteroids, and this was used to identify the spectral types worth investigating for this mission. The PSS suggested that for a water-finding mission, all spectral types within the C-complex were approximately equal in scientific value, and that S-complex and X-complex asteroids would not be worthwhile targets of investigation due to their supposedly low water content (although a company mining for metal instead of water would be interested in only X-complex targets). Additionally, one spectral type in the set of end members, D-type, was especially interesting.

In the interview, the PSS ranked the spectral types of note for the mission. All members of the C-complex were equally important, and the D-type was slightly more important than the others. The ranking of each type and adjusted weights so the sum of the weights is equal to one are shown in Table 7-9

Table 7-9: Desired spectral types of asteroids to study and their relative importance as they were elicited from stakeholder interviews.

Spectral Type	Relative Weight	Adjusted Weight (w_j)
B	1.0	0.1389
C	1.0	0.1389
Cb	1.0	0.1389
Cg	1.0	0.1389
Cgh	1.0	0.1389
Ch	1.0	0.1389
D	1.2	0.1666
(Sum)	7.2	1.0000

For classifications where the scientist has a more complex opinion on the relative value of different spectral types, a weighting system derived from Analytic Hierarchy Process may also be appropriate to use to find the adjusted weights if they are not this easy to elicit [182]. MAUT could also be used to assign weights, but this is a much more challenging process for the stakeholder when the number of categories is high. It is easier for the stakeholder to mentally process the relative merits between two categories over and over again than it is to process the relative merits among them all.

Perceived Utility Within Spectral Types

Next, the PSS was asked about perceptions on visiting many targets of the same spectral type versus visiting as many spectral types as possible. The PSS strongly felt that “variety is the spice of life” and that it would be more valuable to visit many spectral types than visit too many of a smaller number of types.

When asked to try to put this into a more mathematical perspective to measure the relative satisfaction that would be perceived by visiting one versus infinity of a certain type, the PSS felt that satisfaction would be halfway achieved with a visit to one asteroid. The PSS then reached for a statistics textbook to determine what z-score on the normal distribution that represented to map out how much more satisfying subsequent visited within the same spectral type would be. Since one asteroid visit would provide half of the total satisfaction, that is represented by a z-score on the normal distribution of 0.675. Each subsequent

asteroid within the same spectral class that is visited would increase the satisfaction by the same z-score. The results are summarized in Table 7-10.

Table 7-10: Relative satisfaction for number of subsequent visits per spectral class.

Number of Visits within Spectral Type	Z-Score (+/- Standard Deviations)	Cumulative Satisfaction	Satisfaction Weight (w_{kj})
1	0.675	0.50	0.50
2	1.350	0.84	0.34
3	2.025	0.96	0.12
4	2.700	0.99	0.03
+5	3.375	~1.00	0.01

Follow-up conversations with the stakeholder representing the mining company echoed the PSS's assertion that approximately three asteroids per class would be mostly satisfactory for the purposes of the exploration mission given what is known or assumed.

7.3.2.5 Value Model Synthesis

With the two goal-level attributes defined, an alternative formulation to the standard MAUT hierarchy for the central tier of the structure in Figure 7-4 can be proposed that handles the discrete and time-sensitive nature of the data returns. The utility of a goal can be defined as

$$G_{MAU} = \sum (I_{MAU}) \cdot (T_i) \cdot (w_j \cdot w_{kj}) \quad (7-1)$$

where G_{MAU} is the goal-level MAU that feeds into the upper level of the standard MAU hierarchy, I_{MAU} is the instrument multi-attribute utility fed from the lower level, T_i is the time-weighted value of the visit to an asteroid as shown in Figure 7-6, w_j is the adjusted weight of the spectral class of asteroid visited, and w_{kj} is the satisfaction weight of the k th visit to an asteroid in the j th spectral class.

With this structure, the highest possible goal-level MAU that can be achieved is 1, which is also the highest utility allowed in MAUT. Therefore, the goal-level MAUs can be used within a standard top-level MAU hierarchy to calculate the overall, final MAU of a specific design depending on that design's mission plan and how many asteroids in which classes can be visited.

The value elicitation interview with the scientist showed that there were no variations in the opinions of the spectral type adjusted weights w_j and satisfaction weights w_{kj} between the three science goals. The structure shown in Figure 7-4 is for the general case where these may be different; a possible alternative structure for when they are all the same is shown in Figure 7-7.

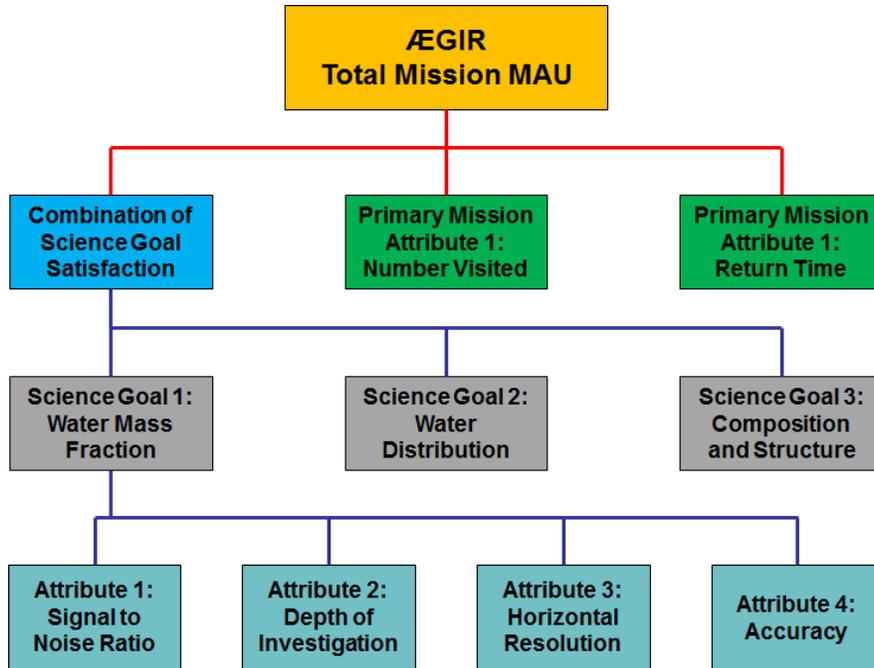


Figure 7-7: Alternative special MAU hierarchy for degenerate case where no differences are perceived among the science goals for the spectral type adjusted and satisfaction weights. This structure inverts the top two levels compared to Figure 7-4.

In this case, the goals are defined explicitly by the instrument-level attributes and can be elicited through the standard MAU formulation (as shown by the blue lines of the tree), but the highest level MAU is governed instead by the equation

$$MAU = \sum (G_{MAU}) \cdot (T_i) \cdot (w_j \cdot w_{k_j}) \quad (7-2)$$

The red lines in the tree show where the standard MAU structure is not applied, using the same color notation as Figure 7-4.

An example of how this works is shown in Figure 7-8. In this timeline three different designs are compared: one design has an instrument-level $MAU = 1$ and visits asteroids once per year; one design has an instrument-level $MAU = 0.5$ and visits asteroids twice per year; and one design has an instrument-level $MAU = 0.25$ and visits asteroids four times per year. Because the first visit to an asteroid of a particular spectral class is the most valuable, the strategy for visiting asteroids is to not repeat a spectral class until all the relevant ones have been visited (Note: this is only true in this example. Because the value model has not yet put a form to the value concept, the discussion of ΔV budgets and optimal strategies is beyond this discussion).

The timeline shows that initially, all three designs achieve similar utilities early on despite the differences in instrument-level MAU. Once the mission reaches 3.5 years, all designs begin to see the penalty from the lifetime utility function (Figure 7-6). However, the design that has visited the most asteroids up to the point also begins to see diminished returns because several member of each type have already been

visited and the satisfaction weights are low. This is why the design that visits asteroids at higher frequency (denoted by green triangles) suddenly drops away from the other two.

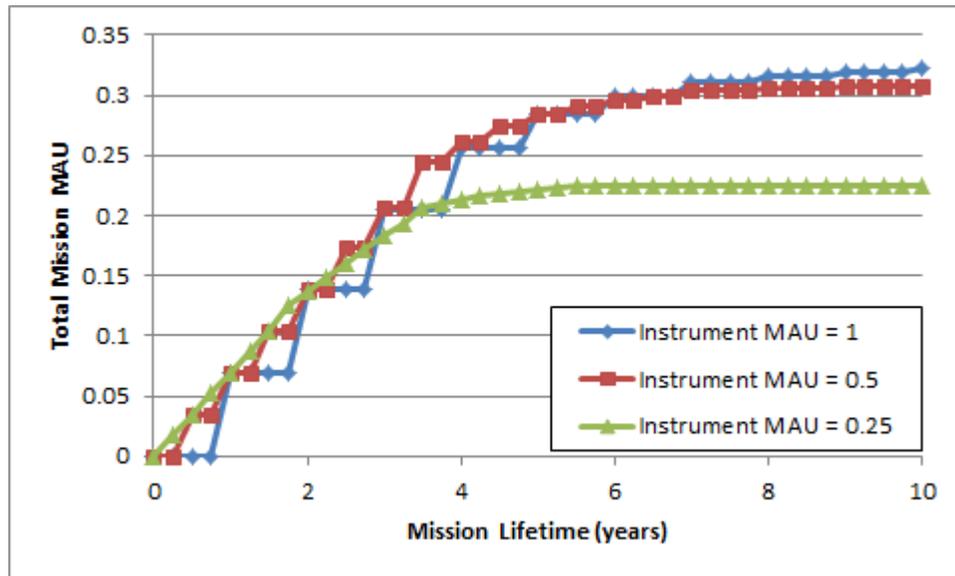


Figure 7-8: Timeline comparison of missions that trade frequency of asteroid visits with instrument capability.

In summary, this three-tier value model uses two standard MAU formulations for the top and bottom levels, but because of the inherent limits of quantifying returns in MAUT, MAUT is not an appropriate method for aggregating the center level of the hierarchy. An alternative approach has been proposed that quantifies the science returns for each goal in a way that accounts for some of the cardinal nature of the science return gained through a census.

7.3.3 Value Modeling Process 3: Epoch Characterization

There are a number of different ways the value models for the ÆGIR mission can change. Characterizing the potential shift in stakeholder expectations is critical to building a robust architecture that can deliver satisfaction under high uncertainty. This section outlines other possible ways to view, calculate, and rank value that can be gained from this mission.

7.3.3.1 Alternative Perceptions of Utility in Time

A scientist interested solely in asteroid exploration for scientific purposes with no business reasons to apply exogenous pressure to the lifetime attribute curve will still desire data to arrive sooner rather than later. A discount function similar to what is used in NPV to account for inflation plus other investments can be applied to the utility curve to represent one way the experienced utility of a design would drop over time even if it continues producing data for many years beyond its original target.

Two possible mental representations of how utility could be discounted in the eyes of a scientific stakeholder using NPV are shown in Figure 7-9.

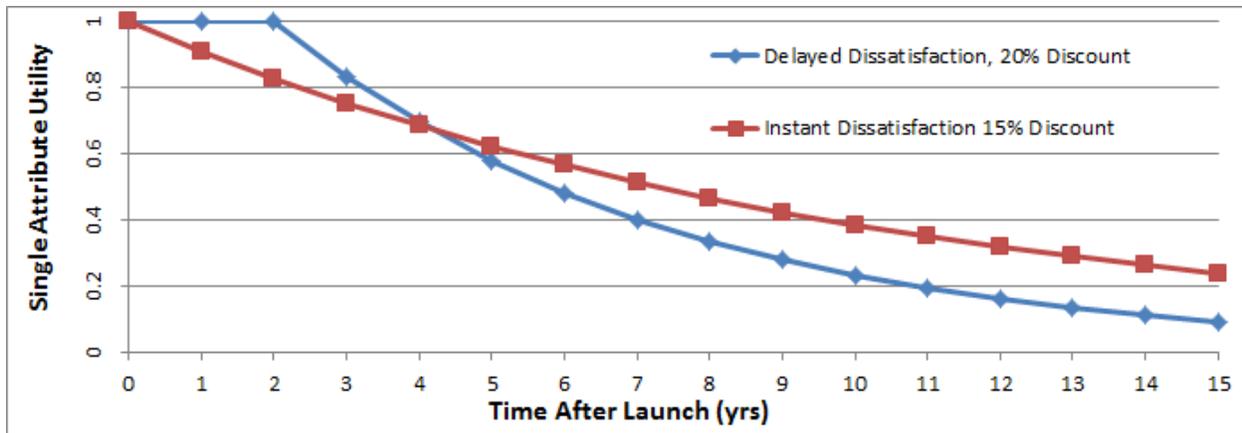


Figure 7-9: Alternative lifetime utility function with discounted utility over time.

Even though there is no specific monetary reason to obtain data faster, there are many other reasons why a scientist would perceive value in having data sooner rather than later. There may be other competitors with missions that could obtain the same data and publish sooner, thereby scooping the discovery. People in general are impatient and want results now instead of later. There may not be many years left in the scientists' career, leaving less time to capitalize on the boost in prestige that comes with these discoveries. Although Figure 7-9 shows two possible utility curves, a better curve can be elicited through interviews.

7.3.3.2 Alternative Stakeholder Perceptions of Statistical Attributes

While the asteroid mining stakeholders are primarily interested in C-complex and D-type asteroids, the scientific community in general is interested in *all* spectral types for study. A science stakeholder's adjusted weights for each type may all be $w_j = 1/25$ if all types are considered equal, since there are 25 classes. There may be other relative variations that make one type more important than another; for instance, because D-type asteroids have been studied less than other types and may contain organic molecules, they may have a slightly higher adjusted weight.

Even without considering the different spectral types that could be relevant, there are a number of ways to interpret data that was gained through census sampling to see how satisfactory it is to answering high-level science goals.

Alternate Mental Interpretations

There are other ways to interpret the value of multiple measurements of the same spectral type. For the values shown in Table 7-10, the PSS thought for a moment and decided that whatever z-score corresponded to 50% of the normal distribution would be sufficient for describing his satisfaction. Another scientist using this same method may instead choose 25% as the baseline for the first measurement.

Another example timeline can be made where the satisfaction weights follow the progression [0.1, 0.15, 0.25, 0.25, 0.15, 0.05, 0.03, 0.01, 0.01] instead of the progression shown in Table 7-10 ([0.50, 0.34, 0.12, 0.03, 0.01]). Such a progression shows higher utility when there are other asteroids within the same spectral type to compare to. Now, not only has the value model changed, this schedule for value delivery heavily influences the trajectory of the mission, since there would be different strategies for maximizing value.

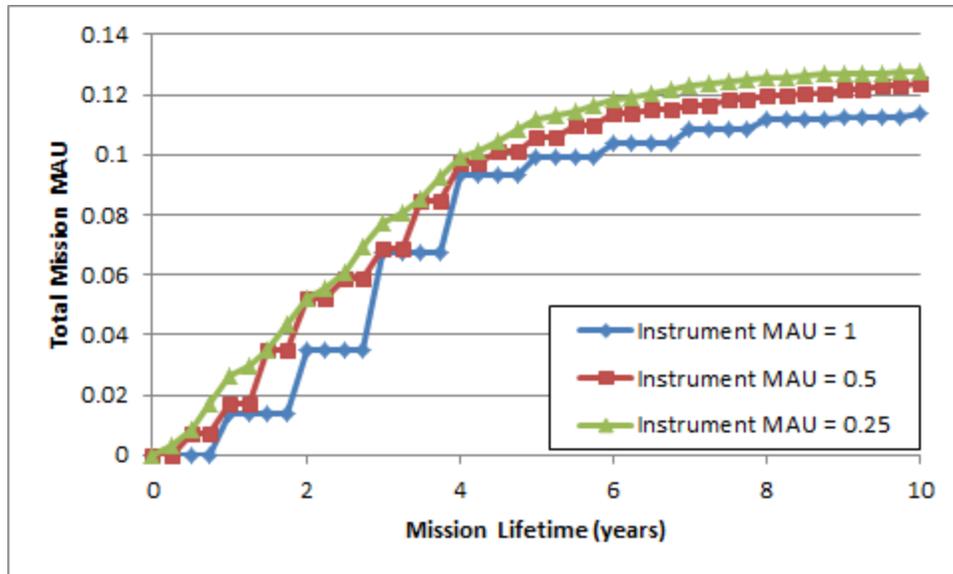


Figure 7-10: Alternative timeline comparing three designs with a different set of satisfaction weights.

An updated timeline showing the same designs from Figure 7-8 with a different progression of satisfaction weights is shown in Figure 7-10. Now, the design with the lowest instrument utility achieves the highest mission MAU as a result of it being able to visit more asteroids and capture more utility before the penalty from the lifetime utility function kicks in.

This shows how crucial understanding how the scientist perceives the value of the census as a whole over the value of an individual instrument, especially when constrained by a limited budget where difficult decisions need to be made about how to proceed once Pre-Phase A is over.

Constructed Models

Another scientist may not even perceive statistical satisfaction in the same manner; rather, a scientist may interpret the number of visits as a means of expressing a confidence interval. If this were the case, a statistically constructed value model that incorporates the number of measurements *and* the quality of the instrument, much like the value model in Chapter 5, may be a more appropriate value model.

The PSS thought that visiting even one asteroid within a spectral class would provide 50% of the total satisfaction that could be achieved for a given spectral class. When using small number statistics to make conclusions about a large number of bodies, a single measurement gives *zero* satisfaction for determining the mean of a sample, since the number of degrees of freedom when using a t-score to find the confidence interval of the mean is zero if the broader science goal is to understand the statistics more thoroughly.

If the first top-level science goal was instead posed as “constrain the mean of the bulk water content of asteroids of different spectral types” rather than “measure bulk water content,” there are objective ways to compare and rank design alternatives that include the instrument attributes as sources of measurement error. Here lies a more rigorous approach that may better reflect a scientific stakeholder’s true thoughts on how the mission benefits the greater scientific good; with a statistically defined objective, the balance between high instrument fidelity compared to high number of measurements becomes apparent.

A comparison of various instrument uncertainties and sample standard deviations is shown in Figure 7-11. Assuming a normal distribution for the bulk water content as a percentage of total mass and the standard deviation of a given measured sample s , one method for expressing the confidence interval is

$$\bar{X} - CV < \mu < \bar{X} + CV, \quad CV = \left[t_{\alpha/2} \left(\frac{s}{\sqrt{n}} \right) + E \right] \quad (7-3)$$

where \bar{X} is the sample mean, s is the sample standard deviation, $t_{\alpha/2}$ is the t-score using a normal distribution for a sample size less than 30, n is the number of samples, μ is the true mean of the population, and E is the measurement bias or error.

From here, it is evident that instruments with higher error can still constrain the mean of the population better than instruments with lower error when the standard deviation of the sample size is large. In the figure, the design that has the highest error is still better than the design with the lowest error if the one with the lowest error is only capable of visiting three targets and the one with the highest error is capable of visiting four or more except in the case where the standard deviation is 2.5%.

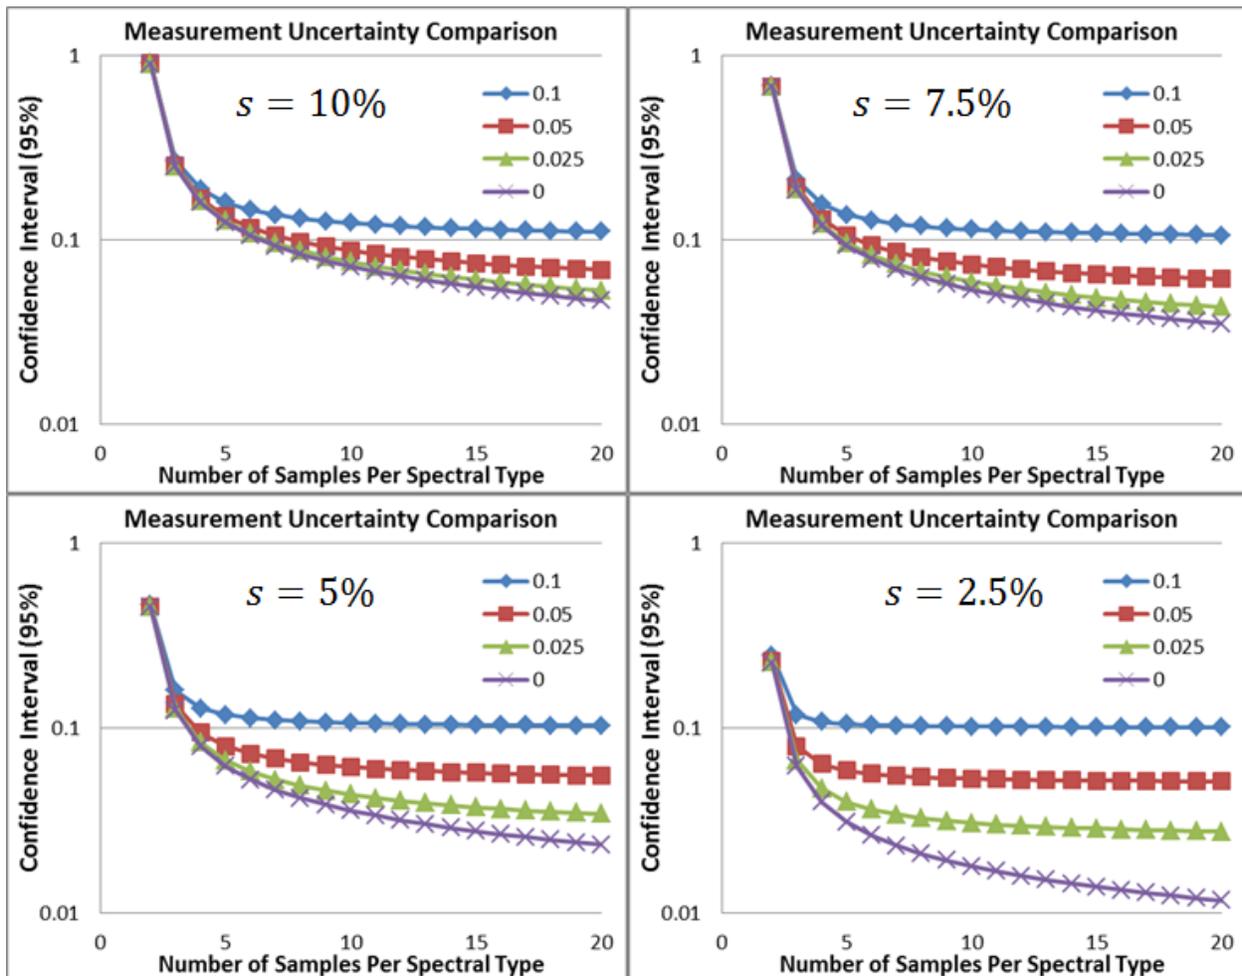


Figure 7-11: Comparison of varying levels of measurement uncertainty and varying sample means with the number of samples per spectral type and the relative size of the 95% confidence interval.

7.3.3.1 Alternative Goal-Satisfying Instruments

Previous space-based radar measurements using fractionated assets such as Rosetta and Philae have used single-frequency scans to map the interiors of asteroids, but recent research on dielectrics and measurements in a range of frequencies has shown promise that radar studies can also study the composition of asteroids.

It is obvious from the attribute weights at the highest levels of the value model shown in Table 7-6 that mapping the distribution of water is the lowest priority of the three goals for the ÆGIR exploration project. Although this is an important goal overall that will be addressed more by the pilot project, the low weighting means that the value returned from a fractionated architecture would be heavily dominated by single-satellite prospectors that carried instruments to satisfy the other two goals.

However, current research has shown that a fractionated radar instrument that scans across frequencies has the ability to satisfy all three science goals. With this technology available, it is possible that, depending on many other design considerations, a fractionated architecture that leverages **self-sampling** to satisfy a variety of goals can dominate any single-satellite prospector.

The availability of an alternate technique may not only change the *form* of the design; it may also change the stakeholder's perceptions about what is and is not acceptable for instrument attributes. For instance, the maximum perceived utility for "depth of investigation" for the first science goal is only a few meters. However, if there is a shift in expectations due to the demonstration or prevalence of better technologies, the utility of any measurements that can be accomplished with neutron spectrometers, passive or active, may be dwarfed by the capabilities of a frequency-scanning radar system.

Even though penetrators are expendable assets, not only can they give greater depth of investigation for neutron spectrometers, they can also be used to conduct radar studies by acting as a signal receiver. This eliminates many of the formation-flying and tracking requirements that small, fractionated satellites would need while increasing the precision of the radar measurements. Furthermore, they have the ability to conduct long-term investigations that may be considered alternate science goals that have not been accounted for yet, such as seismology investigations, tracking beacon missions for planetary defense and YORP Effect measurements, and even minor drilling investigations (which are not relevant to the ÆGIR Phase 1 mission, but may be relevant for other stakeholders if the technology is right).

If the initial assumptions from the value-driving context definition process had driven the concept selection process and epoch characterization had not been conducted, penetrators and fractionated satellites that leverage **self-sampling** would have been prematurely eliminated from the tradespace without considering what could be gained through more research and development.

7.3.3.2 Other Value Model Epoch Shifts

There are other factors that can contribute to a change in the value model that would need to be captured in the epoch characterization process. Some changes include unusual results and adding attributes that have not been accounted for but could be perceived as important later.

Utility Dependence on Results

One of the major scientific unknowns that will be solved in this census of asteroids is whether or not spectral data from remote observations can be used to estimate the water characteristics of asteroids

without traveling close enough to conduct neutron or NIR spectroscopy or radar observations. Due to the nature of the error involved when estimate population means, variances, and confidence intervals from a small sample set, statistical metrics would be affected by the accuracy of the instruments.

Most critically, however, is if the observations from the first few asteroid visits *do not* correlate to the spectral classifications of the Bus-DeMeo taxonomy. For instance, if there is extreme variation between three asteroids of the same spectral type, it may mean an entirely new classification is needed.

If results from this census reveal that the interiors of asteroids are quite different than what we expect them to be from remote spectral measurements, value models for different stakeholders will take wildly different turns. On one hand, scientists tend to be the most excited when (a) a wild prediction of an unknown phenomenon is proven true, such as was the case of the Higgs boson, or (b) when something that has been assumed to be the case for a long period of time is suddenly disproven. The second case is applicable here, and the unexpected result would prove to be scientifically interesting and open up new research questions.

However, this is bad news for asteroid miners. Without being able to know what targets are ripe for mining before they start traveling to them, it may mean there is no way to know the value of a target without visiting it up close. Granted, the goal of ÆGIR is to find an acceptable target for mining since the Phase 2 spacecraft would most likely visit a target that had already been visited by a spacecraft in Phase 1, but it would make future prospecting much more difficult.

Considering how an epoch shift like this would affect the value model of an asteroid mining company is beyond the scope of this project, but in a real-world mission, such an outcome would need to be taken into consideration early on to justify raising the private capital to invest in an exploration mission like ÆGIR.

Additional Attributes

In the interviews with the stakeholders, there appeared to be no relationship between the size of the asteroid visited and the value that could be gained from visiting it. This is not to imply that size is completely irrelevant; an asteroid mining company would gain very little from prospecting an asteroid only ten meters in diameter, even though it would serve usefulness for scientific endeavors. The total volume of resources would likely be too small to turn a profit before the mining equipment would move on to another asteroid, which would more likely require either a large ΔV expenditure or a long wait time.

7.4 Discussion and Conclusions

ÆGIR delves deeper into the value drivers of an atypical science mission and illuminates how complex a mission truly can be when there are thousands or millions of targets to choose from and when no single person knows with certainty how many targets should be visited. The value models can help inform how a performance model for such a mission should be built to accommodate all of the possible value-driving variables. This is especially necessary given the computational power required to optimize interplanetary spacecraft trajectories when the list of possible targets is changing and designs have the ability to respond to new targets that are discovered during the course of the mission.

7.4.1 Discussion of RSC Value Modeling Results

The value model discussed in the previous sections as a way to illustrate the method in action is only a small fraction of what can be learned from this case study. A discussion of additional findings, observations, and recommendations follows here.

7.4.1.1 Challenges of Value Modeling with Multiple Targets

There are a number of challenges when applying MAUT to this case study. The success of science goals should be measured by some mathematical metric, but MAUT aggregates the attributes of a mission based on a stakeholder's feelings and opinions, which are subjective. Because so few asteroids have been studied up close, there is great scientific value in having advanced instrument packages to explore targets. On the other hand, the goals of a census are to study more objects, not fewer, in order to gain more certainty in the population as a whole.

A scientist who would rather have high quality data as opposed to high quantity data would weigh the instrument attributes of the MAU tree higher than the other goal-level attributes. In this case, a statistical function would not properly represent the scientist's perception of value, even if it does ideally represent a metric for satisfying the goals as they have been articulated.

There is clearly some stakeholder ambiguity when it comes to determining the value per sample, especially when multiple asteroids classes are attempting to be visited. Understanding the value difference between fewer high-resolution measurements and more low-resolution measurements is the key to choosing a mission profile and the appropriate number of copies of the asset. Proper epoch characterization can help sort these alternatives out.

An asteroid mining company in their exploration phase may think "*find some potential targets to mine later*" is an acceptable goal. Scientists who have little leverage may also be perfectly fine with this, as their negotiating power is too low to bother thinking about what is *really* needed. A scientist looking at this mission who wants to accomplish broader, more fundamental science goals may rephrase it as "*improve our confidence for knowing the water content of asteroids of a given spectral type.*" The difference between these two mindsets drastically affects the relative satisfaction of each subsequent visit within the same spectral type; the more focused thinker sees only diminishing returns, but the broader thinker sees threshold returns. Somewhere in between, the argument between "go for variety of spectral types" versus "go for quantity within a type" blurs, but the argument for "go for more, period" is only strengthened.

The systems engineer may not have the responsibility to tell the scientist that their perception of value is wrong or to assume that the scientist does not know about their own field; however, it is the author's opinion that the systems engineer has the responsibility to present the scientist with alternative viewpoints for measuring value, just like the systems engineer present alternate designs to accomplish the mission via tradespace exploration. This way the scientist at least has the ability to rethink their position, especially when the results can be drastically different when viewed through the lens of distributed satellites leveraging emergent capabilities that involve statistical satisfaction along with traditional instrument satisfaction criteria.

If a company had begun design work without conducting appropriate stakeholder interviews to understand the value proposition of their own mission, they would have set requirements and selected a

concept long before the full tradespace would have been explored, which drastically reduces the chances of designing a mission that lies on the true Pareto front of stakeholder expectations. The performance model that would have been used to assess trades among different design variables, whether or not all options were compared on the same tradespace, would not explore the full range of possibilities.

With this known value model, not only can additional concepts be explored, pitfalls can be avoided in the design process. There are a number of possible operational strategies that need to be assessed and calculated depending on how the stakeholders perceive the value of each asteroid that explored within a spectral class. The trajectory optimization problem is an enormous computational challenge, which made building a performance model infeasible for a graduate student project, but it will be a necessary component that needs additional inputs that would not have been built without epoch characterization.

7.4.1.2 Payload and Bus Design Synergies

There are several synergies between the payload and the spacecraft bus that are notable. One side effect of using electric propulsion is that the spacecraft has an abundance of electrical power that can be put into other instruments when the spacecraft is coasting or orbiting. While the general focus has been to lower the power budgets of instruments for small satellite missions, electric propulsion allows for high-powered instruments to operate when the thrusters are not.

Higher available power is good for instruments that use active power, such as active X-ray fluorescence spectrometers, even though this is not an instrument that is being considered for this mission. However, if the focus of this mission shifted from searching for water to characterizing PMGs, which will eventually be the goal considering the high value of PMGs, X-ray fluorescence is a candidate instrument for achieving valuable goals.

7.4.1.3 Adapting to New Technologies

A well-constructed value model should be independent of the limits of the instrumentation that are available at the time. However, scientists' expectations may change with technology in the same way that a car owner's idea of luxury improves as more features are added to a newer model car. With the possible emergence of radar as a technique for accomplishing all three science goals, scientists' perception of the minimum acceptable and maximum perceivable levels for penetration depth can change drastically.

The synergy between propulsion and instrument power discussed in Section 7.4.1.2 may be especially relevant for fractionated satellites in formation flight around an asteroid. Small distributed assets may not have the power capacity to use electric propulsion, but would still need to maneuver around an asteroid. A performance model may include designs that use power beaming from a mothership to a fractionated piece of the satellite to provide the power for propulsion. If this concept were suddenly included in the tradespace, it should not affect the value model directly.

7.4.1.4 Developing Operational Strategies Based on Value Models

This case study especially shows that the operational strategy of a flexible mission with many possible targets is heavily influenced by the value model that is being used to rank designs. With the value model elicited from one of the stakeholders, variety is always better, so revisiting an asteroid of the same class that has already been explored is always less valuable than exploring a new type (unless all types have the

same number explored already). However, under a different value model, more visits within a single class are more valuable because of the improved confidence in the statistical mean of the observations.

The operational strategies for achieving the highest utility or value, capturing the most opportunity, are dependent on the structure of the value model. This is especially true when trying to balance ΔV limitations and the increasing target set as new asteroids are discovered over time. For a situation like this mission, a metric like normalized Pareto trace delivers critical information to the decision makers because it shows which designs are robust and responsive to changing mission needs depending on the definitions of success.

7.4.1.5 Alternative Elicitation Techniques

Two value elicitation methods were used to obtain the same k_i weights for the primary science goals. These interview techniques resulted in wildly different weights for the same interview subject. Using ASSESS, the subject reported some substitution effects, but through the method that was created on the fly as a way to ensure the ASSESS numbers were correct, the subject explicitly stated that he perceived synergy and was more satisfied as more science goals were accomplished. More research in applications to science goals is needed to determine whether the ASSESS software unfairly creates substitution effects in value models compared to other elicitation techniques.

7.4.1.6 Implications for Performance Modeling

A high-fidelity performance model for ÆGIR would require an advanced astrodynamics simulation software package to sort through many possible trajectories of sets of spacecraft. Such operational options would need to be revisited in detail once the mission had launched as more asteroids were discovered that could be on more efficient value-driving trajectories.

One possible way to simplify this is to create estimates of the required ΔV and time between asteroids for bulk guesses. This would be a lower fidelity option, but it would allow for a logical, methodic comparison of spacecraft propulsion, power, and propellant tanks. Such comparisons would be especially useful in epoch characterization; for instance, in one epoch, the estimated ΔV to move between asteroids of one class could be some fraction of the estimated ΔV to move to any random asteroid, but in another epoch both of those numbers have changed. These numbers are also location dependent; the average ΔV to travel between NEOs is higher than it is for the main-belt asteroids, but getting to the main belt is more difficult and costly.

Some comparisons and alternative scenarios may be better suited for later in the design phases when comparisons between fewer designs are more computationally feasible and higher-fidelity models are required, but at least some sensitivity analysis must be performed to determine whether or not low-fidelity estimates are good enough.

7.4.2 Comparison to Monolithic Missions

There are few missions (either real or concept studies) that can compare to the ÆGIR mission. The Dawn spacecraft was launched in 2007 and has properly visited only two objects. Dawn is the first spacecraft to orbit two extraterrestrial bodies, and it does so by leveraging solar electric propulsion (SEP); Dawn has the largest ΔV budget in the history of space exploration. No real mission has orbited or rendezvoused with more than two asteroids.

The results from the 2009 GTOC competition show that it is possible of one spacecraft to *fly by* enough asteroids to meet some minimum standards in ÆGIR, but a flyby is not enough for the scientific observations of the ÆGIR mission. Additionally, the competition did not discriminate between different classes of asteroids or weigh visits to one class over others for penalize for visiting too many of the same class. Furthermore, this competition was an idealized study to demonstrate optimization methods more than true mission feasibility. The 2014 competition is much closer in needs and scope to the ÆGIR mission, but it also (rightfully) assumes that a DSS is better suited for this mission.

For a mission to adequately leverage census sampling to characterize a population of targets via in-situ measurements, multiple assets are needed. Additionally, given the intense ΔV required to reach asteroid targets, the most cost-effective way to complete a survey like ÆGIR would be to leverage SEP.

7.4.3 Future Work

Due to time and resource constraints, the ÆGIR mission was not given as much time and effort compared to the other case studies. Before the literature survey of Chapter 7 was conducted, one of the planned contributions of this dissertation was going to be a method to estimate the minimum required ΔV and time to do tours of asteroids, both near-Earth and in the main belt. Some of the work from Petropoulos and others has already sought to tackle this problem [466]. More research and tools are needed to not only understand what the tradeoffs between thrust, available power, I_{sp} , total mass, travel time, and the number of targets a single asset can visit in one tour, but also to synthesize these tools into something tradespace modelers can use to quickly evaluate architectures to reduce computational complexity.

The emergence of two asteroid mining companies in the past five years has also brought up many interesting hypothetical situations that the results of a full RSC treatment of ÆGIR could help characterize more accurately. Both companies are currently planning to send fleets of small satellites to explore asteroids, but the details about how they evaluated their plans are publicly unknown. If either company sees the results of a full tradespace study of ÆGIR rather than a simple value model study, they may be inclined to reevaluate their future plans.

7.5 Chapter Summary

This chapter explored the value proposition of a mission to study many targets of a large population by leveraging **census sampling**. Stakeholder satisfaction is achieved through timely completion of a survey of asteroids using multiple spacecraft. Even the most capable monolithic satellite system would not be able to visit as many asteroids in the desired time period to compete with less-capable distributed systems.

Although the technology shows great promise for addressing more goals than it has in the past, at this time there does not appear to be enough return on investment to leverage **self-sampling** in order to study the distribution of water compared to the importance of the other goals that can be accomplished with simpler instruments. A future iteration of this case study once the stakeholders fully understand the potential of dispersion signatures and how to use them in radar studies may change this in the near future.

CHAPTER 8

GANGMIR

Groups of Atmospheric Nanosatellites and
Ground-penetrators for Mars In-situ Reconnaissance



Original artist unknown. Parody modified by the author and John Alexander.

8. Case Study: GANGMIR

“Shooting for the stars, cruise the speed of light.

***Glowing god of Mars**, body burning bright.*

Well I’m riding, riding on the wind! Yes I’m riding, riding on the wind!

Thunderbolt from hell, shattering aloud.

*Screaming demons yell, **bursting through the clouds**.*

Well I’m riding, riding on the wind! Yes I’m riding, riding on the wind!”

JUDAS PRIEST, “RIDING ON THE WIND”

The fourth and final case study examines one of the most traditional exploration targets in space science: Mars. NASA has been studying Mars with satellites since Mariner 4 launched in 1964, and current NASA policy has plans to send humans to Mars in the mid-2030s. Before NASA can send humans there, however, a large number of science questions must be answered first. The Mars Exploration Program Analysis Group (MEPAG) has outlined and prioritized a number of science goals related to not only the traditional science of Mars but also to the technical studies that will help drive design work for spacecraft to send humans there.

This case study develops the Groups of Atmospheric Nanosatellites and Ground-penetrators for Mars In-situ Research (GANGMIR) mission, named after the spear of the Norse god Odin in part because of the nature of ground penetrators falling from the sky. This mission shows ways to achieve some of the highest-priority science objectives laid out by MEPAG using DSS. This mission heavily leverages **staged sampling** (or **sacrifice sampling**) to achieve its primary mission objectives, but the value proposition also benefits from leveraging **stacked sampling** of heterogeneous assets in multiple environments and possibly several other emergent capabilities of DSS.

8.1 Case Study Introduction and Motivation

This case study was built upon a student research project at Politecnico di Torino that examined several science goals for studying Mars. Three student teams focused on orbital, atmospheric, and ground goals and identified potential architectures that could be used to accomplish those goals. The ground goals were of the highest importance because they have the highest difficulty and relatively high priority for eventual human missions to Mars.

Fabio Nichele, a Ph.D. candidate at Politecnico di Torino, helped oversee the student project. He also helped with this case study by identifying science goals laid out by the MEPAG and prioritizing them according to scientific value and urgency. Other masters’ students involved in the class project also contributed to this case study by providing MATLAB code and checking spacecraft mass and performance models.

Independent of the Italian study, other professors guiding this project are interested in subsurface scientific investigations on Mars. In contrast to the previous and ongoing Mars missions, these scientists are interested in studying subsurface depths that no Mars rover or lander has currently been able to reach

with the equipment they have. Among these possible subsurface investigations include studying chemical chirality, searching for peroxides or other biohazards, and searching for liquid water. These investigations will be at depths more than 1 m below the surface.

Meanwhile, there has been growing interest in the engineering community to use ground penetrators to explore solar system bodies. These penetrators borrow from military technologies like missiles and bunker busters, which can penetrate through many layers of concrete to deliver an intact payload to a specific location. In the case of the space science community, the payload is *not* an explosive warhead but rather science instruments. Previous mission concept design efforts have investigated sending penetrators to the Moon, Mars, and Europa among other targets.

An advantage that penetrators have over rovers is that they do not carry the extra mass required to make a soft landing on a gravitationally strong body. This means that there can be up to several hundred m/s of ΔV taken out of a propellant budget, or that a parachute system is not needed, which saves mass and cost. If subsurface investigations are desired, penetrators do not require the heavy drilling equipment a rover or lander would need to carry – in fact, the deeper a penetrator is designed to go, the less it needs to slow down (though it would need to be stronger and therefore heavier to withstand the higher impact speed).

There are also disadvantages for using penetrators. One disadvantage is that they are confined to their landing site, just like a lander; once the penetrator has been deployed, it can only conduct its investigation from that spot. It cannot move to another location and take another sample; at best, it could take another sample from a different spot in the same hole. There is also higher risk if the target location's surface properties are unknown. If a penetrator impacts a hard surface at the speed expected for a soft one, it could be destroyed. If a penetrator impacts a soft surface at the speed expected for a hard one, it could be buried too far under the surface to communicate.

Additionally, the instruments inside a penetrator must be designed to withstand high g-forces from impact. Larger instruments are more susceptible to be damaged in an impact, so penetrators benefit from the miniaturization of instruments in the same way CubeSats do. This usually means sacrifices in other instrument properties such as resolution, sample size, and power.

8.1.1 Primary Mission Goals and Stakeholders

Unlike the other case studies, GANGMIR is somewhat of a “value by committee,” where inputs and suggestions from many sources have been combined into an area where there are many scientific goals and many ways to interpret their satisfaction.

8.1.1.1 Primary Science Goals

The science goals of GANGMIR are based on the Humans to the Martian System Summary of Strategic Knowledge Gaps report published by the Precursors Strategy Analysis Group (P-SAG), which is jointly sponsored by MEPAG and the Small Bodies Assessment Group (SBAG) [472]. The published matrix identifies 79 goals and assigns a priority and an estimated time of needed completion to each. Fabio combined these priority and timing scores to create a combined priority score and ordered the goals by this metric.

Originally, the case study intended to address the top ten goals, but there were too many unknowns associated with some of them and no meaningful analysis could be performed, so they were dropped from

consideration. The top ten goals identified by Fabio are listed in Table 8-1. The ones for which there was not enough subject expertise to make valid design considerations have been crossed out. Seven goals were deemed sufficient for a single mission given the limited manpower available. The two items in bold will be given special treatment in this case study because of how they can individually demonstrate the potential value of DSS.

Table 8-1: Top Ten science goals of the MEPAG and P-SAG

Rank	Science Goal	Sampling Location	P-SAG Code
1	Observe Global Temperature Field	Orbit	A1-1
2	Observe Global Aerosol Composition	Orbit	A1-2
3	Measure Global Surface Pressure	Ground	B1-2a
4	Observe Local/Regional Weather	Ground	B1-2b
5	Search for Subsurface Biohazards	Ground	B2-1
6	Global Observations of Wind Velocity	Orbit	A1-3
7	Observe Orbital Particulate in High Orbit	Orbit	A3-1
8	Observe Dust and Aerosol Activity	Orbit	B1-1
9	Measure Electrical Properties of Atmosphere	Ground	B6-1
10	Measure Presence of Ground Ice	Ground	B5-1

The “P-SAG Location” in Table 8-1 is the location from where measurements can be taken as identified by P-SAG, not the location that measurements *could* be taken. For instance, the global temperature field can be measured with radiometers from orbit, but knowing the ground temperature as measured by an asset on the ground could provide valuable information for this goal. There is great potential to show **stacked sampling** with heterogeneous assets with some of these goals. These goals will be discussed in more detail in Section 8.2.1.

8.1.1.2 Primary Stakeholders

Like HOBOCOP, inputs from a number of scientists have been solicited or inferred from the study conducted at Politecnico di Torino. Among these scientists include: an engineering professor who served as the faculty advisor to the Italian student groups; an MIT professor who gave inputs for the chirality studies; an atmospheric scientist at JPL who works on Mars atmosphere research; and members of the Mars Climate Sounder (MCS) team at JPL. The P-SAG matrix of science goals also shows objectives for certain attributes of the goals, such as resolution and timeline of observations.

8.1.1.3 Secondary Mission Goals and Stakeholders

This mission is strictly a traditional science mission, so there are no secondary mission goals or non-scientific stakeholders that have been identified at this time. There could be stakeholders involved with other penetrator missions that wish to see the technology developed for their purposes, but they have little say here. Those in the human spaceflight community may have specific thoughts on the importance of certain science goals, but those inputs have been captured by the prioritization work Fabio has done.

8.1.2 Mission Concept of Operations

There are many assumptions that are being made at the start of this mission that help constrain the problem. The tradespace assumes that a single mothership launched from Earth will deliver the mission assets to Mars. This mothership will use a combination of aerobraking and propulsive maneuvers to enter a high-inclination, eccentric orbit around Mars.

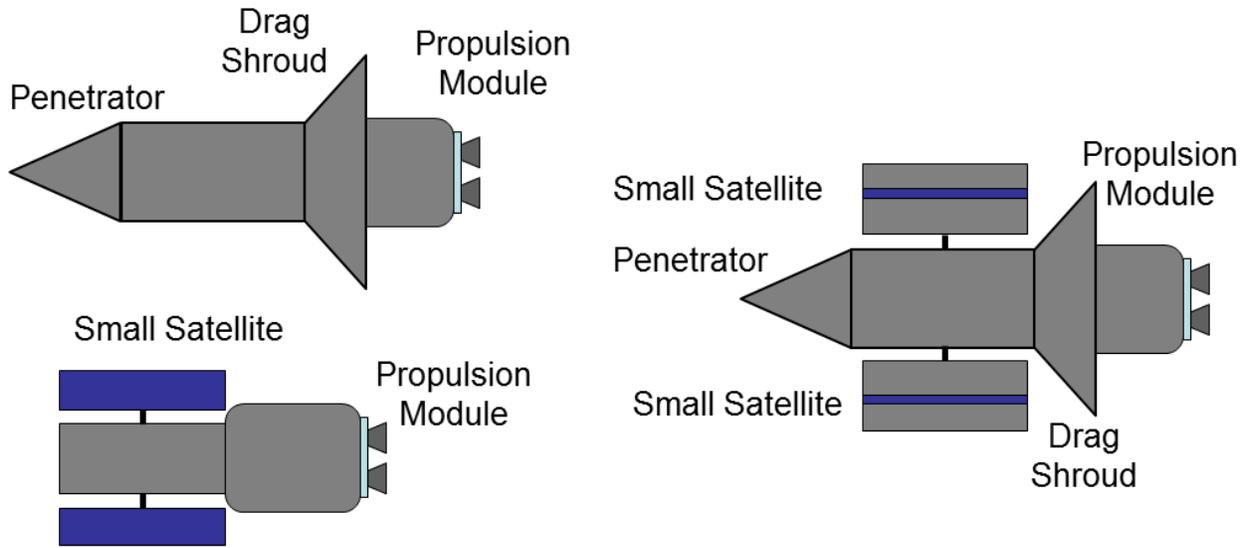


Figure 8-1: GANGMIR mission assets concept design. (Top left) Individual penetrator. (Bottom left) Individual small satellite. (Right) Combined set with two small satellites and one penetrator.

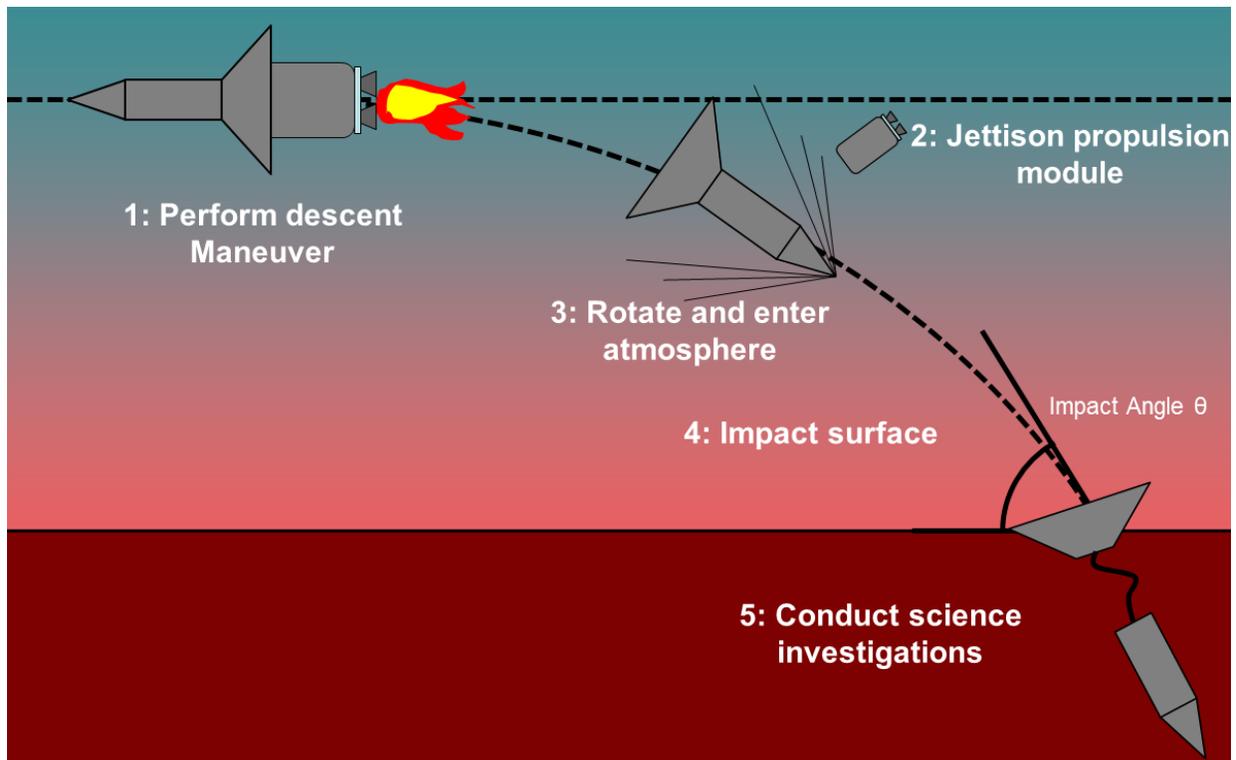


Figure 8-2: Concept of operations for a penetrator.

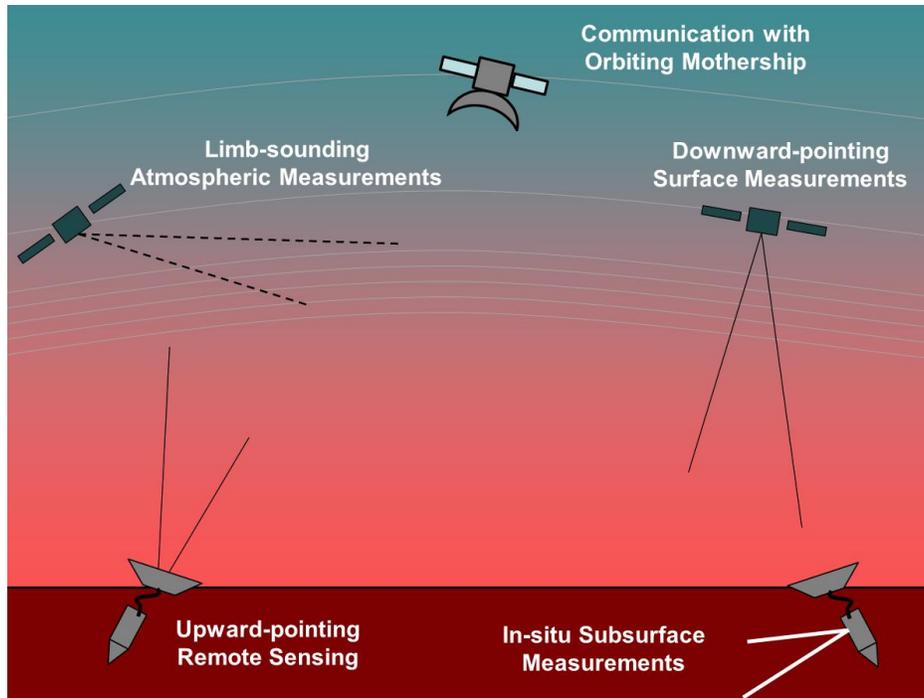


Figure 8-3: Combined operations in GANGMIR with small satellites and ground penetrators.

A high inclination was chosen to help satisfy both the ground and air science goals; all of the desired ground sampling locations can be easily reached (eventually) from a high-inclination orbit, and higher inclination gives more global coverage for observation satellites compared to low-inclination orbits. Since two primary, heterogeneous assets are being examined in this case study, there are three concepts of operations: one for each individual asset, and for the combination of the two assets.

8.1.2.1 Penetrator Operations

A complete ground penetrator asset consists of three components: the penetrator, the drag shroud, and the propulsion module. The mothership will release a penetrator at its orbital periapsis. The penetrator will perform a thrust maneuver to slow down, causing it to enter the atmosphere. The penetrator will release the propulsion module and turn around, since the thrust maneuver will be pointed in the opposite direction of the direction of motion.

The drag shroud will help slow the penetrator from hypersonic to subsonic speeds, which also causes its angle of attack to drop as it descends. On impact, the penetrator will physically separate from the shroud and bury itself into the surface. The shroud will also survive impact and remain on the surface. The shroud contains communications equipment and, if necessary, solar cells to generate power. The penetrator will be connected to the shroud via a cable so that it can send data through the shroud to the mothership.

Once the penetrator is underground, it will take a sample of the regolith via a drill and bring it inside the body of the penetrator, where other instruments will measure it. These instruments will be powered by a primary, non-rechargeable battery. This is the end of short-term mission operations for the penetrator.

If the penetrator is responsible for long-term investigations, it will continue to operate and communicate with the mothership. It will continue to be powered by solar panels inside the shroud that charge a

secondary battery. This secondary battery will power the other instruments conducting long-term weather investigations.

There is no guideline for how quickly the mothership releases the penetrator assets. The timetable only depends on when the target locations on the surface are along the flight path of the penetrator.

8.1.2.2 Small Satellite Operations

A complete small satellite asset consists of a satellite bus and a propulsion module. The mothership will release a small satellite at its orbital periapsis. The small satellite will then perform a thrust maneuver to slow down and enter a circular orbit. The small satellite will then conduct its primary mission. Data is relayed from the small satellite to the mothership when a link is available.

Because the small satellite is in a circular orbit while the mothership remains in an eccentric orbit with a larger semimajor axis, and because Mars is oblate (not a perfect sphere), orbital precession will separate the orbits of these two satellites. Over time, there will be larger angular separation between the two orbits. The mothership will then release another small satellite that will enter a circular orbit. Because this orbit will be the same as the other small satellite, they will precess at the same rate while the mothership continues to precess at a slower rate. The time between separations depends on the number of small satellite assets and the specific orbit of the mothership. A [video demonstration](#) is available to illustrate the difference in precession speeds between these orbits [473].

8.1.2.3 Combined Operations

If a mission uses both small satellites and penetrators, there is opportunity for shared subsystems assets. A penetrator and a small satellite or a pair of small satellites can be deployed from the mothership together. The set of assets will perform a maneuver to slow down and enter a circular orbit. The penetrator will release a small satellite into this orbit. If there are more small satellites being placed into this orbit, two more small maneuvers can be performed to gradually space out these satellites in True Anomaly within the same orbit. Once all the small satellites have been placed into their orbits, the penetrator will perform a final maneuver and enter the atmosphere, and both types of assets will conduct their mission operations as previously described.

In this type of operation, the penetrators can only be deployed with the small satellites, rather than at will, whenever the mothership's trajectory takes it over a landing site. This matters little to the short-term investigations, but will cause the value delivery of the long-term investigations to have a more gradual ramp-up. The exact timetable for precession-aided deployment is dependent on the mothership's orbit and the desired orbit of the small satellites, but it is on the order of one or two months. This is a shorter time period than a similar constellation around Earth because Earth is less oblate and precession is slower.

There are many advantages to using both sets of assets. Penetrators can perform in-situ subsurface investigations, which are valuable on their own, and they can conduct upward-pointing measurements that can work in concert with the small satellites gathering other atmospheric data. The small satellites conduct both limb-sounding and nadir-pointing measurements (see Figure 8-3). The combinations of these measurements allows for far more possibilities to be considered compared to monolithic missions.

8.2 Case-Specific Literature Review

There are a number of topics in the published literature that are relevant to this specific case study, most notably previous Mars missions and the development of penetrator technologies. This section will review the remaining literature that is relevant to this case study that has not been covered in Chapter 2 or in previous case study chapters.

8.2.1 Mars Science Goals

The Mars Exploration Program Analysis Group has identified a number of goals related to human exploration and the habitability of Mars in the past [474]. This section offers a brief overview of the goals relevant to GANGMIR. There are two distinct types of science goals involved in their studies. The first involves characterization of the atmosphere. This requires long-term, high-temporal resolution measurements of phenomena that have been measured extensively by previous missions to better characterize them for future missions.

The second type of goal involves the search for chemicals, both organic and nonorganic, that may be present on the ground and beneath the surface, but have never been detected before. These investigations can be conducted quickly over short periods of time; once a sample has been analyzed, it does not need to be continually analyzed over months or years.

8.2.1.1 Mars Atmospheric Science

There are a number of high-priority science goals relevant to human exploration of Mars involving atmospheric science. The two highest-priority goals identified by MEPAG are global measurements of temperature and aerosols. Knowledge of these phenomena is directly related to improving Mars entry, descent, and landing (EDL). Atmospheric temperature influences the drag profile and trajectory of incoming spacecraft, while understanding and being able to model aerosols is critical for guidance and navigation during EDL.

The observation time required to fully characterize atmospheric models to the fidelity that is needed for EDL of large vehicles is on the order of five years. The vertical resolution that is typically required for limb-sounding instruments that can measure temperature and aerosol profiles are on the order of two or three pixels per atmospheric scale height.

There are a number of lower-priority goals related to Mars' atmosphere, including measuring the global surface pressure, surface wind speeds, orbital particulates, dust and aerosol climatology, electrical properties of the atmosphere, and regional and local weather. For each of these goals, the P-SAG matrix of science goals outlines target requirements or desires that would improve on information that is already known from previous missions.

8.2.1.2 Mars Ground Science

One of the most important goals related to human exploration is determining whether or not microbial life exists on Mars, and if so, whether it harmful to humans. If there is none, other substances could be hazardous to human health, both during operations on the surface and possibly when returned to Earth if a sample of toxic, spreadable material is returned.

Ideally, a sample return mission would be able to achieve this science goal. Testing protocols for sample returns has been drafted to limit the possibility of cross-contamination [475], but there are other options

such as in-situ measurements that may provide most of the necessary information without the great expense of a sample return.

Data from the Mars Odyssey gamma ray spectrometer shows a high lower limit of the water mass fraction of the upper layers of the surface regolith in several regions, most notably the poles [476]. From this data, a lower limit of the mass fraction of water can be derived [477], [478]. This map is shown in Figure 8-4. From here it is evident that the mass fraction of the regolith at the poles is approximately one third or higher, suggesting that these areas could exhibit mechanical properties similar to permafrost.

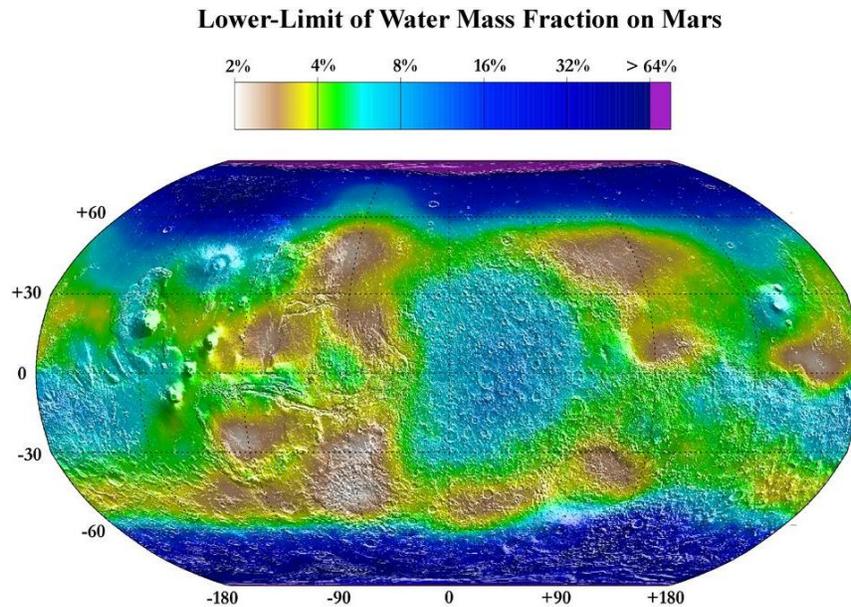


Figure 8-4: Mass fraction of water on Mars as determined by gamma ray spectrometer on Mars Odyssey (Image credit: NASA/JPL/Los Alamos National Laboratory [479])

More recent analysis from Odyssey’s gamma ray spectrometer shows that water may be bound to sulfates in the soil [480]. Understanding the sulfate content of different potential landing sites for human exploration is critical because sulfates can be deadly to humans. Additionally, extracting water from sulfate-rich regolith would drive the design of any equipment related to in-situ resource utilization (ISRU).

None of the assets that have been deployed to Mars so far have dug very deep into the surface. The Phoenix lander used a digging mechanism to examine regolith a few centimeters below surface. Phoenix’s wet chemistry experiment found evidence of perchlorates in the soil, which could also be hazardous to humans [481], [482].

8.2.2 Similar Science Campaigns and Missions

Mars is by far the most popular destination for planetary science missions. The first attempt to visit Mars was in 1960. Since then, more than half of all missions that were destined for the Red Planet have failed or only been partially successful. This section outlines aspects of some of the more notable aspects of missions that are relevant to the development of the GANGMIR mission.

8.2.2.1 Previous Successful Mars Campaigns

Several missions to Mars have been successful. These missions are summarized in Table 8-2.

Table 8-2: Notable successful Mars missions.

Mission	Launch Date	Agency	Type	Status
Mariner 4	11/28/1964	NASA	Flyby	Defunct
Mariner 6	2/25/1969	NASA	Flyby	Defunct
Mariner 7	3/27/1969	NASA	Flyby	Defunct
Mariner 9	5/30/1971	NASA	Orbiter	Derelict Orbit
Viking 1	8/20/1975	NASA	Orbiter/Lander	Defunct
Viking 2	9/9/1975	NASA	Orbiter/Lander	Defunct
Mars Global Surveyor	11/7/1996	NASA	Orbiter	Contact Lost
Mars Pathfinder	12/4/1996	NASA	Lander/Rover	Defunct
2001 Mars Odyssey	4/7/2001	NASA	Orbiter	Operational
Mars Express	6/2/2003	ESA	Orbiter	Operational
Spirit	6/10/2003	NASA	Rover	Defunct
Opportunity	7/8/2003	NASA	Rover	Operational
Mars Reconnaissance Orbiter	8/12/2005	NASA	Orbiter	Operational
Phoenix	8/4/2007	NASA	Lander	Defunct
Mars Science Laboratory	11/26/2011	NASA	Rover	Operational
Mars Orbiter Mission	11/5/2013	ISRO	Orbiter	Operational
MAVEN	11/18/2013	NASA	Orbiter	Operational

Mariner 4, 6, and 7 were flyby missions, while Mariner 9 was the first spacecraft to successfully orbit Mars. These early spacecraft carried a range of instruments including magnetometers, Geiger counters, cosmic dust detectors, IR and UV spectrometers, radiometers, and visual imaging systems.

Viking 1 and 2 were the first successful landers on Mars. They each carried three biological experiments: a Pyrolytic Release experiment (PR), a Labeled Release experiment (LR), and a Gas Exchange Experiment (GEX). Additionally, each carried a Gas Chromatograph/Mass Spectrometer (GCMS). Of these experiments, only the LR experiment gave positive results for signs of life on Mars. Images of the desolate surface led many to conclude that life did not exist on Mars.

The Mars Global Surveyor was the first spacecraft to visit Mars in over 20 years. It carried five scientific instruments including a Thermal Emission Spectrometer (TES). It also served as a signal relay for future spacecraft. The Mars Pathfinder carrying the Sojourner rover was launched one month later. Sojourner carried six primary experiments including an X-ray spectrometer. However, this rover had no capability to drill or collect soil samples.

After several spacecraft failures, the 2001 Mars Odyssey orbiter was the next successful Mars missions. Odyssey carried three primary instruments, including the Thermal Emission Imaging System (THEMIS). This instrument has a mass of 11.2 kg and requires 14 W of power [483]; it is a marked improvement over the TES instrument on the Mars Global Surveyor.

Mars Express is an ESA orbiter carries seven instruments: three for surface and subsurface measurements, three for atmospheric and ionospheric measurements, and one radio science experiment. The high resolution stereo camera can map the surface to a resolution of 10 m.

The Mars Exploration Rovers (MERs) Spirit and Opportunity carry a number of instruments, including a miniaturized TES for identifying interesting geologic features and a Rock Abrasion Tool (RAT) to expose fresh material for analysis. The RAT is capable of drilling a 45 mm diameter hole 5 mm deep. The Opportunity rover has traveled a total distance of 41.4 km since landing; the Spirit rover traveled 7.73 km before becoming stuck in sand.

The Mars Reconnaissance Orbiter (MRO) builds on the technology from Surveyor and Odyssey. The orbiter carries three cameras, two spectrometers, and radar. The most relevant instrument to GANGMIR aboard MRO is the Mars Climate Sounder (MCS), a sounding radiometer that measures thermal profiles in several near- and far-infrared bands (see Section 8.2.3.2).

The Phoenix lander was the first spacecraft sent to a polar region on Mars. The lander carried a robotic arm and shovel that could dig as deep as 0.5 m into the regolith to sample dirt and ice. Phoenix's Microscopy, Electrochemistry, and Conductivity Analyzer (MECA) included a wet chemistry laboratory. A meteorological station was also aboard, but the lander did not survive through the winter season, so long-term, ground-based meteorological studies could not be conducted.

The Mars Science Laboratory (MSL) builds on the technology of the MERs and carries several instrument packages, including several chemistry experiments and weather monitors. The rover, named Curiosity, has a robotic arm with a percussive drill capable of drilling up 5 cm deep into rocks.

The Mars Orbiter Mission (MOM), also called Mangalyaan, carries instruments for atmospheric, particle environment, and surface imaging studies. The Thermal Imaging Spectrometer (TIS) is mapping the composition and mineralogy of the surface and can operate day and night. This spacecraft does not conduct sounding observations to measure atmospheric temperature.

The Mars Atmosphere and Volatile Evolution (MAVEN) spacecraft is the most recent orbiter sent to Mars. Instruments aboard this orbiter include several solar wind analyzers and Langmuir probes, an imaging ultraviolet spectrometer, and a neutral gas and ion mass spectrometer package. This mission's primary scientific goals are related to studying the top layers of the atmosphere, which are not closely related to the scientific goals put forth by the MEPAG and P-SAG.

Despite so many successes, many science questions are left unanswered. Neither Phoenix nor Curiosity can dig beneath the surface into the permafrost layer, and MCS is incapable of mapping the thermal profile of the entire planet often enough for building high-accuracy atmospheric models. After these monolithic systems, the science questions that remain are best answered using DSSs.

8.2.2.2 Landing Site Selection Processes

Each time an asset is placed on Mars, it is sent to a location that has been debated for years before the mission is launched. This selection process involves poring over the details of each potential landing site and balancing the science goals that would be accomplished at a particular landing site, the capabilities of the asset that is being designed to land on Mars, and the risks associated with the landing site itself.

The Pathfinder rover's landing site selection process took over two and a half years and was narrowed down to three finalists [484]. Constraints on the landing site included: smooth, flat landing area; latitude between 10 and 20 degrees; altitude below 0 km; and moderate rock abundance. Area Vallis was chosen because it had greater scientific potential than the other two finalists.

The Mars Exploration Rovers landing site selection process took over two years and was conducted over a series of four open workshops. The final four candidates were downselected from a list of over 150 potential landing sites [485]. Constraints on the sites included: latitude between 10 degrees N and 15 degrees S for maximum solar power; elevation below -1.3 km for sufficient atmosphere for EDL; low winds to reduce landing error ellipse; low slopes; and a radar-reflective, load-bearing surface.

The Phoenix lander's landing site selection was a "multi-year effort," even though the mission was already constrained to explore high latitudes of Mars' surface [486]. Desired characteristics of the site included: location between 65 and 72 degrees latitude; elevation below -3.5 km; soil cover of at least several centimeters over ice or icy soil; and a benign wind environment. Green Valley was chosen from a list of 13 potential sites [487]

The Mars Science Laboratory's landing site selection process took five years and involved broad participation in the science community [488]. The initial list of over 50 sites was narrowed down to four finalists. Engineering constraints included: mid-latitude for thermal management of the instruments; low elevation so there would be enough atmosphere to slow the spacecraft upon entry; slopes of less than 30 degrees; moderate rock abundance; and a radar-reflective, load-bearing surface that the rover would not sink too deep into to move. All four the final sites had layered sedimentary rocks with evidence of phyllosilicates that addressed the mission objectives. Gale Crater was chosen over the other three finalists because it has greater scientific diversity and potential for habitability.

One common theme among the rejected landing locations is the increased risk associated with those sites, not the lack of possible scientific returns. Many of the scientifically interesting sites on Mars are in canyons, gullies, craters, and generally rocky areas that engineers want to avoid because designing a guidance system that can navigate such obstacles is difficult and expensive. The Viking lander did not know it landed in an area littered with boulders until after it touched down; the landing site selection team thought the area was incredibly smooth because the resolution of the radar data was too low to detect those boulders.

As Mars payloads have gotten bigger and the price of failure has gotten higher, scientifically interesting landing sites have been constrained due to the technical challenges and uncertainties. Because distributed systems have the benefit of minimized failure compensation costs and can leverage the emergent capability of sacrifice sampling, these landing sites can finally justify the risk required to explore them. Penetrators may not be able to provide the same level of detail as a rover or a lander from above ground, but now the value proposition of those sites can be evaluated fairly, not rejected from the beginning.

8.2.2.3 Distributed Deployment on Mars

In addition to the problem of selecting locations to land monolithic assets, Mars missions are also no stranger to the operational complexities that arise when leveraging **sacrifice sampling** or **staged sampling** to achieve science goals. These forms of sacrifice sampling have not been in the form of distributed systems but rather in the form of single-use payloads.

The MSL science goals to investigate habitability include assessing the biological potential by determining the inventory of organic compounds and identifying geologic features that may record biologically relevant processes [489]. An instrument suite on MSL called Sample Analysis on Mars (SAM) will be used to analyze samples collected during the mission [490]. There are nine sample slots for

wet chemistry analysis onboard SAM. These slots contain material to study wet chemicals and are expendable, meaning once they are used, they cannot be used to take another measurement. Early results show that the apparatus has leaks, which further confuses the results from the instrument [491], but the first results have been released [492].

There is ongoing debate over when to use these sample slots and when it is more valuable to take another sample at the same location where another sample has already been taken. This is the same problem that other distributed assets will face if they leverage sacrifice sampling – when is it appropriate to “pull the trigger” and use an asset that was originally intended for another target?

8.2.3 Performance Modeling in GANGMIR

While the literature reviews in previous chapters have all covered many of the aspects needed to model GANGMIR, two important areas have not yet been discussed. In this section, literature that is specific to the performance modeling of penetrators and sounding radiometers will be discussed.

8.2.3.1 Penetrators in Space Science Missions

Although bunker buster technology is heavily restricted by ITAR, there have been several other efforts to use penetrators in planetary missions that can be used to create a parametric model for penetrator mass and performance estimates.

Ground penetrators are applicable to a number of space science goals on a variety of targets [493], [494]. The key advantage of penetrators is that they do not need to waste as much fuel slowing down to achieve a soft landing on planetary bodies with no atmosphere, and they can use smaller parachutes or drag shrouds to control their descent on bodies with atmospheres. They also have the potential to be outfitted with autonomous control to achieve higher precision on landing, making them valuable delivery vessels for payloads that can withstand the impact deceleration.

The Mars Polar Lander carried two micropenetrators, the Deep Space 2 (DS2) probes, to characterize subsurface regolith properties. Although they successfully separated from the Mars Polar Lander, which crashed due to a serious software error during its descent, the micropenetrators did not establish contact with the Mars Global Surveyor after impact and were deemed lost [495]–[498]. It is possible that the penetrators landed successfully but that the impact melted water beneath the surface and flooded the capsule, short-circuiting the electronics. It is also possible that the flex cable snapped because the penetrator went deeper than expected into the regolith. A cross-sectional view of a DS2 penetrator is shown in Figure 8-5.

There is new interest in exploring Jupiter’s moon Europa with penetrators. ESA is in the early planning stages of exploring beneath the top layers of ice on the surface [499], [500]. Although this penetrator would come nowhere near sampling the possibly liquid ocean beneath the ice surface, it would collect a sample in one of the rifts where liquid water could have risen and refrozen. It would be possible to sample material from a later geological period in these locations.

To estimate the depth of penetration of a projectile, Young conducted a number of experiments that varied projectile mass, cross-sectional area, nosecone type, impact velocity, and soil type [501]. The parametric equations of penetration depth were derived from experimental data that will be used in this case study.

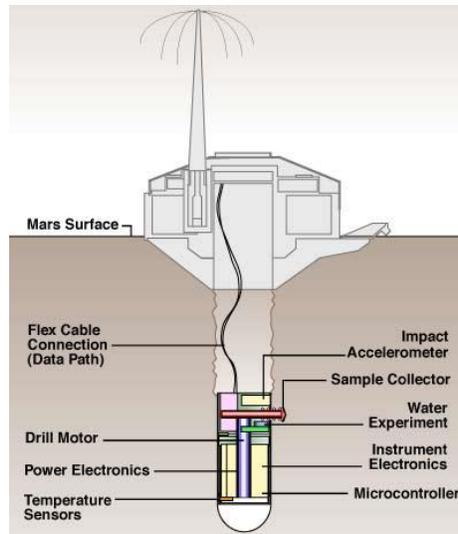


Figure 8-5: Cross-section sketch of the Deep Space 2 micropenetrator (reprinted under fair use from [502]).

The survivability of penetrators is limited by the deceleration they experience on impact, which is heavily influenced by the penetrability of the surface that the penetrator hits and by the impact velocity. These two factors also heavily influence penetration depth, so there is a tradeoff between the structural integrity (and therefore total mass) and other subsystems associated with the mission. A strong, heavy penetrator may survive an impact into a rock that a lighter penetrator would not survive; however, the heavy penetrator may bury itself too deeply to effectively communicate with a mothership if the ground penetrability is high and the radio attenuation of the regolith is high (which it would be with high water content). The lighter penetrator would not have as much regolith to communicate through and may be better on this particular ground type.

One particular area of uncertainty involved with estimating impact deceleration and penetration depth is the penetrability of ice or frozen saturated soil at different temperatures. While Young showed that separate equations dictate the penetration depth of dry and unfrozen saturated soil compared to ice and frozen soil, these equations do not take into account whether ice or saturated frozen soil is brittle or ductile. Colder, brittle ice has a higher chance of shattered upon impact, lowering the deceleration force and allowing the penetrator to go deeper into the surface, while warmer, ductile ice may behave differently due to the high strain rate from impact and cause higher forces that could damage the penetrator. This is a concern particularly because the surface temperature of Mars can dip below the transition temperature.

The UK Penetrator Consortium has researched using penetrators to explore the Moon [503]. The team's project name has changed several times (MoonLITE, LunarEX, and LunarNet), but more recent work has included specific instrument recommendations and engineering configuration models [504], [505]. The primary instruments that have been proposed for LunarNet include a heat flow probe, a magnetometer, a seismometer, a mass spectrometer, a tilt meter, a volatile detector, an X-ray spectrometer, and a descent camera. Other potential instruments include a microscopic imager, a dielectric/permittivity sensor, and a thermos-gravimeter. A small sample acquisition drill would collect a sample for the instruments to study. The mass of the LunarNET penetrator itself is 13 kg. When the penetrator delivery system is included

(which includes fuel to decelerate from lunar orbit), the total wet mass of the payload is 51 kg. A configuration design of the penetrator is shown in Figure 8-6.

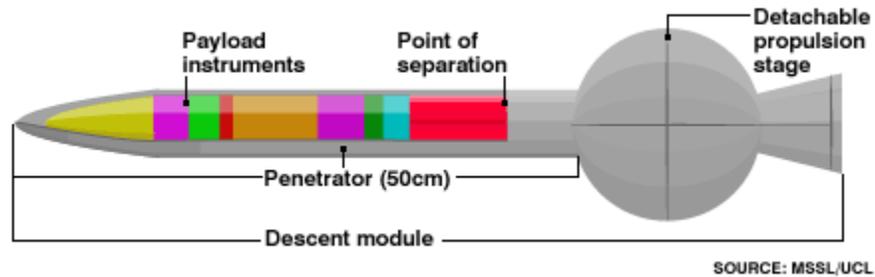


Figure 8-6: Cross-section of a planetary penetrator (reprinted under fair use from [506]).

Another team from Finland is designing a network of small penetrators for Mars weather observations called MetNet [507]. These penetrators carry a number of instruments that can link with an orbiting satellite to provide weather surfaces across the surface of mars, including thermometers, hygrometers, magnetometers, solar irradiance sensors, and panoramic cameras.

From these references, a list of potential instruments that could be included in a tradespace of penetrator instrument options was created. The proposed instruments for the LunarNet/LunarEX penetrators and their associated characteristics are listed in Table 8-3, and the proposed instruments for the METNET penetrators and their associated characteristics are listed in Table 8-4.

Table 8-3: Instruments designed for the LunarNet penetrator mission (data from)

Instrument	Mass (kg)	Power (W)	Energy (W-hr)	Data (Mbit)
Accelerometer	0.07	0.8-1.2	0.17	1
Descent Camera	0.16	0.160	0.015	~10
Heatflow Probe	0.30	0.025-0.3	--	0.5
Magnetometer	0.07	0.15-0.4	--	1
Mass Spectrometer	0.75	3-6	30	0.2
Seismometer	0.30	0.053-0.112	500	5
Tiltmeter	0.01	0.1	--	0.001
Water/Volatile Detector	0.75	3	15	TBD
X-Ray Spectrometer	0.26	4	24	0.1
Drill	0.10	2	1	--
TOTAL	2.77	--	--	--

Table 8-4: Instruments design for the METNET mission (data from [507])

Instrument	Mass (kg)	Power (W)
Barometer	0.100	0.015
Hygrometer	0.015	0.015
Panoramic Camera	0.100	--
Accelerometer	0.050	--
Solar Irradiance Sensor	0.065	--
Dust Sensor	0.200	--
Thermometer	0.015	0.015
Magnetometer	0.070	0.15
TOTAL	0.780	--

Other groups have developed advanced instruments to detect biological signatures that previous Mars campaigns would have failed to detect. An analytical chemistry experiment has been developed to study molecular chirality [508]. This experiment is the first demonstration of a completely automated capillary electrophoresis analysis on a single, fully integrated microfluidic device, and due to its small size, it would be perfect for use in penetrators. Due to a lack of expertise and a stakeholder to interview that has the necessary background in biology, this instrument was not considered for this study, but a future iteration with feedback from stakeholders who are invested in this instrument could examine its value as applied to Mars missions.

8.2.3.2 Small Satellite Instrument Modeling

The small satellites in GANGMIR will consider options for carrying up to two scientific payloads, a sounding radiometer and an optical camera.

Sounding Radiometer

The most important instrument for the orbital observations in GANGMIR is a sounding radiometer. Many satellites and air vehicles have carried sounding radiometers to measure atmospheric temperature profiles [509]–[511], but none have been small enough for a direct comparison or to make a parametric model for small satellites. However, the Mars Climate Sounder (MCS) launched on MRO serves as a single-point model to work from, especially given how small it is in relation to other instruments on MRO [512]. A performance and mass model of this instrument can be built off the known performance and mass of MCS using the equations presented in SMAD for observation payload sizing.

The MCS radiometer consists of two telescopes with 4 cm apertures connected to the same optics structure. The instrument samples infrared radiation in nine wavelength bands. These wavelength bands and their specific scientific purposes are shown in Table 8-5. The entire instrument has a mass of 9 kg and requires 11 W of power. This is a major improvement compared to its predecessor on the Mars Climate Orbiter, which had a mass of 44 kg and required 40 W of power.

As a consequence of using a wide range of bands, the detector’s pixels are sized to provide maximum resolution at the higher wavelength, but are not sized differently to achieve higher resolution at lower wavelengths. A more complex instrument could be developed that varies the size of the detector pixels to match the best possible resolution at difference wavelengths, but heterogeneous pixel sizes on the same array are not typical.

Table 8-5: MCS spectral channel bands and functions (data from [512]) arranged by band center.

Channel	Band Center (µm)	Measurement Function
A6	1.65	Polar radiative balance
A4	11.8	Dust and condensate (D&C) extinction from 0 to 80 km
A3	15.4	Temperature and pressure from 40 to 80 km
A2	15.9	Temperature and pressure from 40 to 80 km
A1	16.5	Temperature from 20 to 40 km
A5	22.2	Temperature from 0 to 20 km, D&C extinction from 0 to 80 km
B1	31.7	Temperature from 0 to 20 km, D&C extinction from 0 to 80 km
B2	41.7	Water vapor from 0 to 40 km, D&C extinction from 0 to 80 km
B3	42.1	Water vapor from 0 to 40 km, D&C extinction from 0 to 80 km

Optical Camera

While there are few small satellite models for radiometers, there are many options available for small satellites, including some that are commercially available for CubeSats. Three previously flown cameras were examined to build a parametric model to estimate the mass and power of a camera given the aperture diameter and focal length: GOMSpace, ExoplanetSat, and PlanetLabs.

GOMSpace's NanoCam CIU is a 35 mm camera designed for 1U CubeSats [513]. It provides an excellent capability for an instrument of its size. The camera on ExoplanetSat is a slight step up in quality but comes with mass and power penalties [270]. PlanetLabs has been hesitant to publish data about their Dove satellites, their instrumentation, and their mass capabilities, but some information was gathered through reverse calculations, estimates based on images available, information that has been published at small satellite conferences [228], and personal, anonymous communication with employees.

The necessary characteristics for building a parametric model to explore the small satellite optical camera tradespace are presented in Table 8-6.

Table 8-6: Summary of small satellite camera data for use in GANGMIR.

Supplier/Spacecraft	Aperture	Focal Ratio	Mass	Power
GOMSpace	35 mm	1.9	0.17 kg	0.6 W
ExoplanetSat	85 mm	1.4	0.86 kg	2.3 W
PlanetLabs (Estimates)	90 mm	12.6	2.3 kg	10 W

8.3 The RSC Model for the GANGMIR Mission

The Responsive Systems Comparison model for GANGMIR was constructed partially in conjunction with several other graduate students at Politecnico di Torino, including Fabio Nichele, Lorenzo Feruglio, and Enrico Cumino under the guidance of Dr. Sabrina Corpino. Dr. Corpino served as the PSS and offered guidance when expert opinion was otherwise unavailable for solicitation. The JPL atmospheric scientists was also consulted on the Mars atmospheric goals.

8.3.1 GANGMIR RSC Phase 1: Value Modeling

The first phase of the Responsive Systems Comparison builds a value model from information from the P-SAG goals and other expert opinions. Because the value of GANGMIR comes from many goals and many expert opinions, is it difficult to synthesize an accurate value model without having these experts on hand to consult with. The expert opinion of a variety of scientists has been gathered as much as possible from personal conversations that were not part of formal stakeholder interviews, as well as from the literature presented in the previous sections.

Not all of the numerical data used in the value model of GANGMIR comes from opinion elicited by these stakeholders because these stakeholders were mostly unavailable to sit down for a formal interview. When information is missing, the author explicitly states how those values were chosen based on his expert opinion of how the scientists had described the importance of certain attributes of the design.

8.3.1.1 Value Modeling Process 1: Value-Driving Context Definition

In this step of the value modeling phase, the stakeholders, primary mission goals, available resources, and stakeholders' needs that drive the value of the mission were identified. Because this is a concept study, some liberties had to be taken and assumptions made about the relative importance of the goals involved

in this study. The needs of this mission are driven by the desire to put humans on Mars, which is what drove the goal prioritization work that Fabio Nichele conducted based on the MEPAG and P-SAG documents.

Stakeholders

Given that this mission aims to fulfill the same role as traditional Mars missions have over the past twenty years, the stakeholders are assumed to be similar to NASA or another agency invested in planetary exploration. Since there are a number of goals, a variety of scientific stakeholders could be interviewed. The resources of this mission are also considered to be near the same level as typical planetary science missions to Mars.

GANGMIR Science Goals

The top ten goals identified in Fabio Nichele’s work that were initially considered for this work are shown in Table 8-1. However, because there was not enough subject expertise among the group working at Politecnico di Torino to address every goal, some were removed from consideration, and another was added. The final list of science goals that will be addressed by GANGMIR is shown in Table 8-7. These goals and their MAU weights will be discussed in more specific detail in the next process of value modeling. The first two goals are in boldface because they were the original two goals this mission initially sought to achieve before Nichele’s work on prioritization.

Table 8-7: Complete list of science goals being considered in GANGMIR.

#	Science Goal	Sampling Location	P-SAG Code	MAU Weight (ki)
1	Observe Global Temperature Field	Orbit/Ground	A1-1	0.3
2	Search for Subsurface Biohazards	Ground	B2-1	0.3
3	Observe Global Aerosol Composition	Orbit	A1-2	0.1
4	Measure Global Surface Pressure	Orbit/Ground	B1-2a	0.1
5	Observe Local/Regional Weather	Orbit/Ground	B1-2b	0.1
6	Observe Dust and Aerosol Activity	Orbit/Ground	B1-1	0.1
7	Measure Presence of Ground Ice	Ground	B5-1	0.1
8	<i>Obtain Atmospheric EDL Profiles</i>	<i>EDL</i>	<i>B1-4</i>	0.1

Rejected Science Goals

Not all of the top ten goals identified by Fabio Nichele were considered for GANGMIR, and the final goal was added only because the nature of the mission with penetrators actually satisfies the goal in a meaningful way. The rejected goals and reasons for their rejection are discussed here.

The goal “Global Observations of Wind Velocity” (P-SAG Code A1-3) was rejected because no known deployable anemometers that can survive impact inside a penetrator could be used as a basis to build a model for that instrument. How winds can be measured from orbit and what makes a quality measurement is also unknown, and expert solicitation on the subject was unavailable as this mission was planned.

The goal “Observe Orbital Particulates in High Orbit” (A3-1) was rejected because the orbital particulates of interest are in the orbits of Phobos and Deimos, which orbit with semimajor axes of 2.76 and 6.92 Mars radii, respectively. This area is far beyond the planned orbit for the mothership. Additionally, at the time this project started, the instruments associated with satisfying this goal were unknown. A future iteration of this cast study could easily include the option for the mothership to raise its orbit once the distributed assets have been deployed. If this happens, additional checks to ensure that communication

between the assets and the mothership is still feasible at the longer link distance, and more fuel is required to achieve the desired final orbit, but otherwise the mission would change little.

The goal “Measure Electrical Properties of the Atmosphere” was rejected because it was unclear what instruments could be used to measure these properties. Additionally, there was ambiguity in whether this should be accomplished from orbit or from the ground. If measured from the ground, the complexity and probability of survival upon impact is also unknown, as is whether or not these measurements would be contaminated by dust. The survivability of instruments in the penetrator’s shroud is much lower than that of the ones inside the penetrator because the shroud decelerates much more quickly on impact whereas the penetrator slows down over its entire penetration depth.

8.3.1.2 Value Modeling Process 2: Value-Driven Design Formulation

GANGMIR intends to address eight science goals relevant to future human exploration of Mars. The value models for each goal and the overall weighting of the mission were developed during this process.

During the development of this case study, there were many iterations between the value modeling and the performance modeling phases. The set of design variables changed many times and influenced how the performance attributes of several science goals are measured. Most of the instrument performance attributes were less dependent on the quality of the instrument but rather whether or not the instrument existed at all in the payload package. By combining multiple attribute descriptions into a single attribute, a significant amount of time was saved. All attributes will be summarized at the end of 8.3.1.2, and how these attributes are influenced by the design vector will be discussed in more detail in Section 6.3.2.

A major advantage of this case study is that every instrument that is being considered for this mission has high-TRL and is already miniaturized in some way. Ground penetrators themselves are a high-risk, low-TRL architectural variable with high complexity (as described in Table 4-4), but many of the instruments meant for spaceflight aboard penetrators have been tested already. More about these instruments will be discussed in Section 8.3.2.

Goal #1: Observe the Global Temperature Field

The most important goal as identified by the P-SAG is to measure the global atmospheric temperature with full diurnal coverage. The P-SAG document identified potential attributes and placed recommendations on the required levels of these attributes. Vertical resolution and measurement lifetime were listed as value-driving attributes, and they even note that this goal is “suitable for a mission of distributed satellites” – even they indirectly admit that this goal is best satisfied by a DSS leveraging **shared sampling** as opposed to a monolithic system that could never achieve the desired global coverage.

The P-SAG recommends that vertical resolution for remote sounding to measure atmospheric temperature be less than 5 km. An engineer not following the RSC method may assume this is a design requirement and not understand how this attribute affects satisfaction. Atmospheric scientists desire the resolution for limb-sounding instruments like a radiometer to be less than the scale height of the atmosphere. Ideally, they would like approximately three pixels per scale height. On Mars, the scale height is on average 8 km, but this level is dependent on temperature.

An atmospheric scientist at JPL, who was not working on the P-SAG document, provided more insight to this goal. Resolution could be as low as 10 km before being completely unacceptable, but there would be

very little added value in going beyond 2 km resolution. This set the upper and lower bounds for the attribute level in GANGMIR.

The definition of coverage for this goal was difficult to define with the expertise of the group in Italy (a proxy metric was used in early presentations with the caveat that this needed to be revisited). With the JPL scientist's inputs, it was determined that revisit time was important. Measuring the *global* temperature field as a function of time is of great interest to scientists. P-SAG recommended full diurnal coverage, but the JPL scientist noted that full-planet coverage every hour would be best, and better temporal sampling resolution would not be much more useful.

The design of MCS provides information on which wavelengths are necessary for observing temperature at different altitudes. The data in Table 8-5 shows which bands are necessary for measuring temperature and serves as the basis for how the value model for this goal was completed.

The P-SAG matrix states that the mission lifetime should be greater than five years. However, the JPL scientist suggested that five years is closer to the ideal lifetime for measurements, and that two years of data could still be valuable and acceptable. However, if there were a major dust storm or other global weather event, another year of observations would be highly desired. This possible scenario will be addressed as a context variable.

In addition to remote measurements, in-situ measurements can also help measure global atmospheric temperatures, especially as a way to confirm or calibrate measurements from satellites. In this case, the science goal cannot be accomplished with only in-situ measurements. Penetrators alone would not satisfy scientists because they only measure ground temperature, not atmospheric temperature profiles. Multiple ground sites would be required in spatially distributed locations to be useful for this goal.

The attribute utility functions for GANGMIR Goal #1 are shown Figure 8-7. These functions were agreed upon by the group in Italy (with the exception of "Maximum Revisit Time," which was originally substituted by "Number of Satellites") and examined by Dr. Corpino. The top three functions relate to the small satellites, the bottom two on the left relate to the penetrators, and the bottom right relates to the lifetime of these global atmospheric temperature measurements.

It is worth pointing out that the attributes "Number of Radiometer Bands" and "Penetrator Weather Package" shown in the top right and bottom center of Figure 8-7 do not seem like typical attributes. More on what these mean will be discussed in the final subsection of this section (Section 8.3.1.2).

The attribute weights are shown in Table 8-8, and the MAU hierarchy of the goal is shown in Figure 8-8 for clarification. These weights also show that **stacked sampling** may also increase the value proposition of the mission, though the low weighting of the ground assets could mean that the measurements are actually redundant and not quite valuable depending on the stakeholder. In this case, a best guess was made for these weights based on inputs from the scientists that were consulted, information available in the student reports, and Dr. Corpino's assessment of the importance of ground measurements.

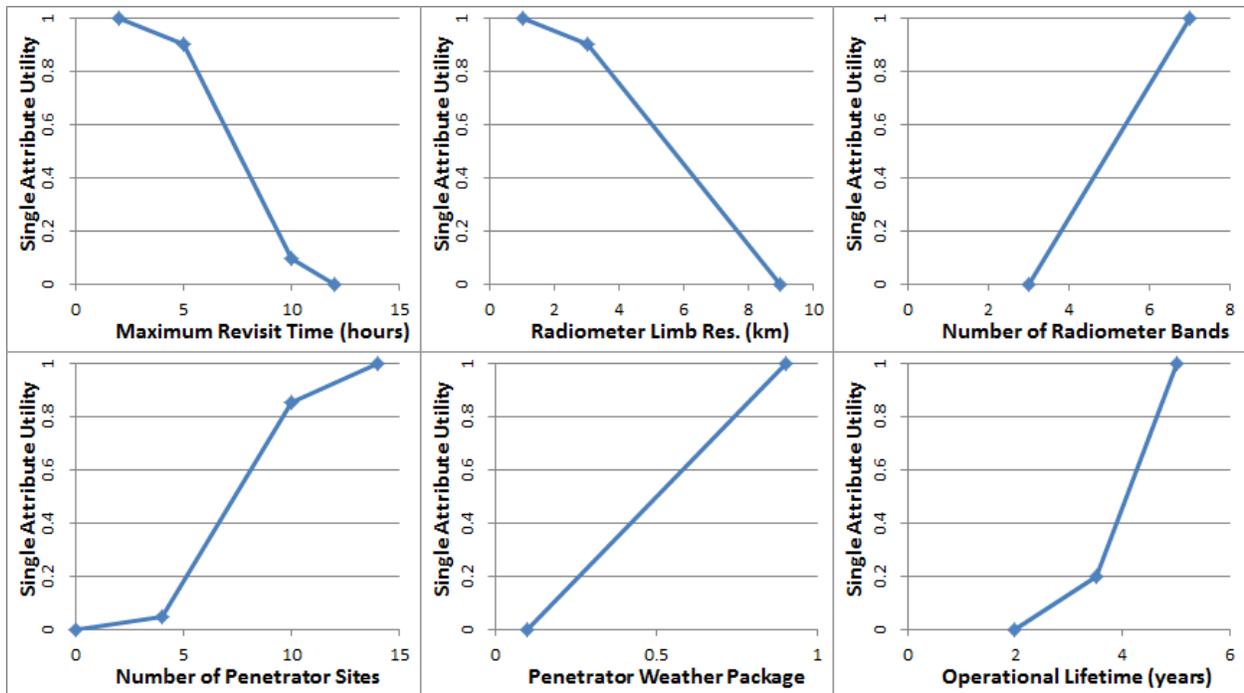


Figure 8-7: Single Attribute Utility functions for Science Goal #1 of the GANGMIR mission.

Table 8-8: Single Attribute Utility weights for GANGMIR Science Goal #1.

Goal-Level Attribute	Instrument-Level Attribute	Weight
Radiometer	---	0.5
---	Maximum Revisit Time	0.20
---	Radiometer Limb Resolution	0.15
---	Number of Radiometer Bands	0.10
Weather Station	---	0.05
---	Number of Penetrator Sites	0.35
---	Weather Package	0.05
Mission Lifetime	---	0.3

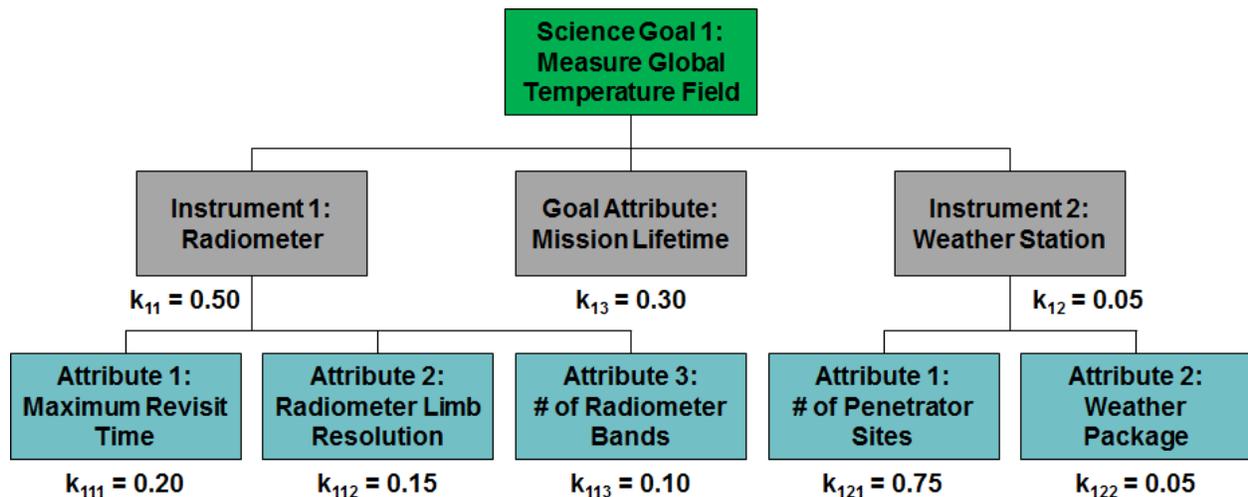


Figure 8-8: Hierarchy of the attributes of GANGMIR Science Goal #1: Observe Global Temperature Field

Goal #2: Search for Subsurface Biohazards

Previous missions have not detected any signs of life on Mars, but the possibility that dangerous chemicals or other forms of life could be present beneath the surface has not been ruled out entirely. By using ground penetrators to sample deeper than any lander or rover has previously, in multiple locations to better ensure a conclusive finding, scientists will know the risks of sending humans to Mars and whether or not the water that has been detected can be used for human consumption.

Landing Locations

The student group at Politecnico di Torino identified ten potential landing sites for subsurface sampling. These sites represent areas that scientists have argued to be scientifically valuable in previous campaigns, but due to limits of monolithic systems, most of these locations have never been visited before. The group also provided some justification for selecting these sites, which are listed in Table 8-9 and shown on the map in Figure 8-9.

Table 8-9: Landing sites identified by the students by the “Ground” group at Politecnico di Torino

Landing Site	Location	Scientifically Relevant Notes
Argyre Planitia	49.7 S, 316 E	Water may have flowed in channels here.
Green Valley	68.2 N, 234.3 E	Phoenix lander detected perchlorates. Mostly ice.
Nili Fossae	22.6 N, 76.8E	MRO detected carbonates. Possible methane source.
Promethei Terra	64.6 S, 243.8 W	Nearly pure ice. Rock formations may be moving glaciers.
Hellas Planitia	42.5 S, 70.0 E	Higher surface pressure, possible liquid water.
Amazonis Planitia	22.7 N, 196.9 E	Young terrain, hardened lava, sedimentary rock.
Deuteronilus Mensae	43.9 N, 337.4 W	Dusty surface, possible ice, interesting weather.
Chryse Planitia	28.4 N, 319.7 E	Viking 1 found rocks, evidence of water erosion.
Area Vallis	10.3 N, 334.4 E	Pathfinder site, possible channels from erosion.
Unnamed Site	70.0 N, 170.0 W	Interesting magnetic anomaly in this area.

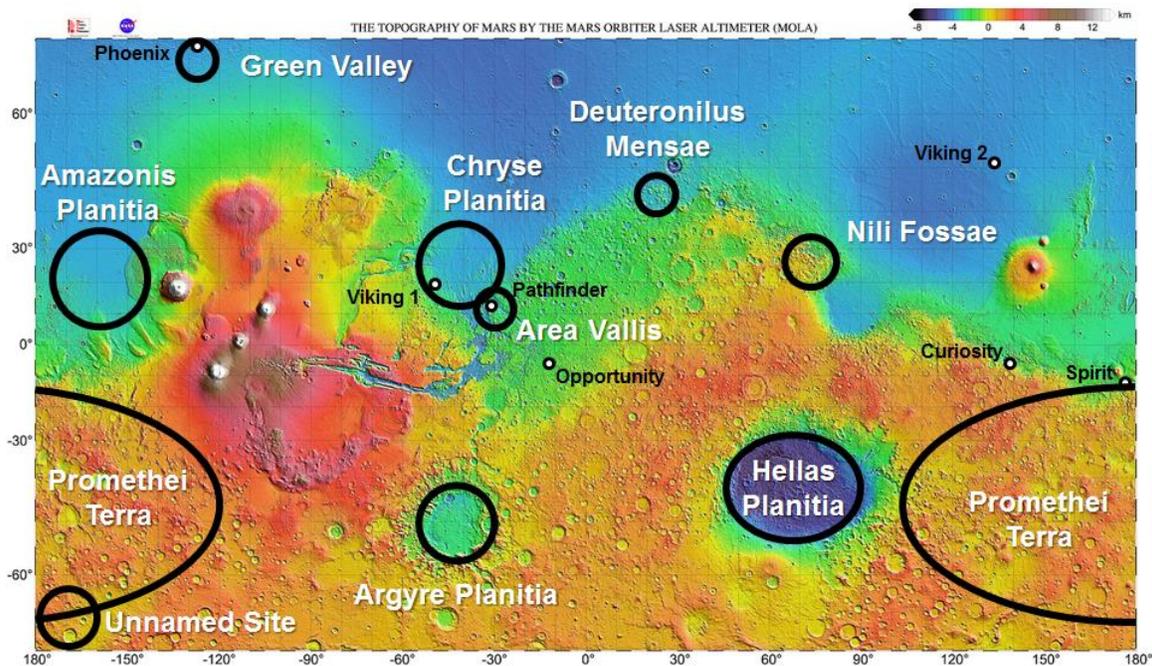


Figure 8-9: Map of desired landing sites for a Mars penetrator mission (data taken from [514]) and previous successful landing campaigns.

Seven of the identified landing sites have not yet been visited by landers or rovers. Argyre Planitia is an ancient, well-preserved impact basin that shows past evidence of liquid water from multiple time periods. It appears to be the largest drainage path in the solar system.

Nili Fossae appears to be the source of plumes of methane, raising the question of whether this source of biological or geological in nature. MRO detected carbonates in this region, and minerals in the clays could preserve organic materials. However, this region contains basaltic rocks that may make ground penetration risky.

Promethei Terra is a location of interest because of materials surrounding cliffs called lobate debris aprons (LDAs). Shallow radar data from MRO suggests that LDAs are glaciers covered by a thin layer of rocks. These glaciers may be a source of water for future human explorers.

Hellas Planitia is the largest impact basin on Mars. The altitude difference between the rim and the lowest depth is over 9,000 m. The pressure at this depth is above the triple point of water, meaning liquid water could exist naturally here under the right temperature conditions. However, evidence of a deeper rock layer suggests that penetration may be limited to less than one meter.

Amazonis Planitia appears to be the youngest geological area of Mars (only approximately 100 million years old). This area is smooth but covered in hardened lava that rests over layers of sedimentary rock.

Deuteronilus Mensae contains flat-topped knobby terrain that could have been carved by glaciers. These glaciers still exist there today, and some may have been formed as recently as 10,000 years ago.

Finally, the “Unnamed Site” is of interest because of the unusual crustal magnetism that is present at this site in addition to the ground ice present in other areas of Promethei Terra.

Three of the chosen sites have been visited by landers, but are still scientifically interesting enough to warrant revisits, though two of them are near the bottom of the priority list as shown in Table 8-9.

Green Valley is the site where the Phoenix lander is located. Data suggests that this site has the highest concentration of water in solid state outside of the poles, though the presence of perchlorates could make it a hostile environment for biological studies.

Chryse Planitia is the site of Viking 1 and shows evidence of past water erosion. The basin area is a boulder-strewn plain that no monolithic mission would ever be sent to again because of the landing risks that were unknown until the first Viking images were returned to Earth. There is evidence of basaltic rocks in this area, and penetration depth would be low with high risk.

Area Vallis is where Pathfinder landed. Its surface appears to have been excavated by currents, possibly in the form of liquid water, and the area is near the choking point where Valles Marineris would drain into Chryse Planitia.

Value-Driving Attributes

Interviews with Dr. Corpino suggested that although ten landing sites had been identified by the students in the Surface working group, full satisfaction of this science goal could be achieved with five or six successful penetrator impacts.

This science goal also requires a sample drill and spectrometers to study the sample. Full satisfaction is achieved with a combination of mass, X-ray, and volatile-detecting spectrometers, though some satisfaction is still achieved with only the mass spectrometer.

Initial work suggested that penetration depth was important to this goal, but the initial penetration depth models showed that achieving depths of more than one meter into solid basaltic rock would be impossible within the parametric space of the penetrators described in this chapter’s literature review. Later literature review revealed that most of the targets would be ice or ice mixed with soil, which is still a risky ground type to impact, but far more feasible in comparison. Additionally, the required depth was not as high as the initial assumptions, and there are diminishing returns beyond a depth of two meters. However, there is still ambiguity here that could be alleviated with stakeholder inputs from more knowledgeable sources.

The attribute utility curves for GANGMIR Science Goal #2 are shown in Figure 8-10. Like the attribute for “Weather Package” in Goal #1, “Spectrometer Package” will be discussed in more detail later.

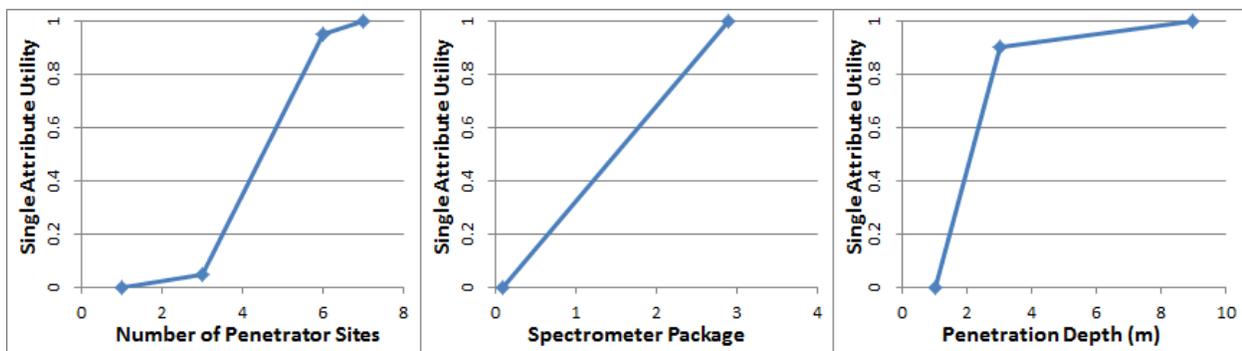


Figure 8-10: Single Attribute Utility functions of Science Goal #2 of the GANGMIR mission.

Although this science goal’s primary benefit will be shown through EEA by leveraging **staged sampling**, one could also argue that in some way, sampling many locations on a planetary surface is a crude form of **census sampling** considering the breadth of locations that can be sampling in a single mission.

The remaining goals will not be discussed in as great detail as the first two due to their lower relative importance, but the value-driving attributes and how they can leverage the emergent capabilities of DSS to increase stakeholder satisfaction will be discussed here. The full set of attribute utility curves and weights for each goal can be found in Appendix C.

Goal #3: Observe Global Aerosol Composition

Global aerosol composition can be observed with small satellites but not the ground penetrators. Observing at longer wavelengths is required (as shown in Table 8-5), meaning more observation bands are necessary. Global coverage as measured by maximum revisit time is far less critical for this goal compared to how critical it is for measuring the global atmospheric temperature, so fewer satellites are needed to achieve full satisfaction, but **shared sampling** as it related to coverage still drives value in this goal because no single satellite can achieve the required coverage necessary for full satisfaction.

Goal #4: Measure Global Surface Pressure

Measuring the global surface pressure by a network of ground stations conducting in-situ measurements shows an example of how **simultaneous sampling** can be leveraged to accomplish a goal that no single

asset can. Small satellites can leverage **shared sampling** to achieve global coverage but with lower precision than what can be achieved with ground sensors. The combination of these measurements shows how **stacked sampling** can add value to the mission.

Goal #5: Observe Regional and Local Weather

Regional and local weather observations do not require high-resolution cameras, so this goal benefits far more from increased coverage than it does from increased spatial resolution. Nadir-pointing cameras on a satellite and zenith-pointing sensors on the ground can combine observations to gain value through **stacked sampling**, though satellites alone leverage **shared sampling** while ground stations alone leverage **simultaneous sampling** to build bigger pictures of global weather.

Goal #6: Observe Global and Regional Dust Activity

Dust can be observed globally using radiometers and locally by the ground penetrators using zenith-pointing cameras and solar irradiance sensors. There is some balance between regional and global coverage observations that can be made, so different stakeholders would have different opinions on how these sets of assets contribute to satisfaction of the goal, whether it is synergistically or redundantly. Radiometers require longer wavelengths to measure the presence of dust and thus more spectral bands.

Goal #7: Measure the Presence of Ground Ice

In addition to organics, measuring subsurface ground ice is important for confirming gamma-ray spectroscopy data. It is also vital to know the water content of certain areas for future human missions to Mars that will rely on ISRU to provide sustenance to the astronauts. This goal can only be achieved with multiple ground penetrators that carry a volatile-detecting spectrometer.

Goal #8: Obtain Atmospheric EDL Profiles

This is the only goal that requires no instruments because an accelerometer associated with a flight control system will provide all the data that is necessary for this goal. For this goal, satisfaction is simply a case of “more is better” as the data from this mission can be used to help build atmospheric models just like the other science goals related to the atmosphere. The only attribute that drives value is the number of penetrators in the mission – in fact, if the penetrators are in communication with the mothership during the descent, landing does not even need to be successful to gather this data.

Goal Multi-Attribute Weights

The MAU weights for each goal on the right side of Table 8-7 were not elicited from a PSS; therefore, they may not represent the true relative importance of each goal relative to the set. The top two goals were deemed as the highest priority for this case study and were given a weight three times higher than the other weights. The remaining goals are all considered equally important. The sum of the goal-level weights $\sum k_i > 1$, meaning that a fictitious mission manager who selected these weights would experience substitution effects as more goals are satisfied and would not be willing to risk high satisfaction on many goals to achieve high satisfaction on *all* goals. However, the sum of k_i is not *much* greater than one, so the substitution effects are relatively small.

Additionally, this MAU hierarchy assumes that all goals are on the same level, meaning there is no particular synergy between two goals that can deliver more value as a set rather than individually (such as what is shown in Figure 4-12). This assumption may not be correct given the relationships among the weather goals and how some stakeholders may perceive the data to be related and stackable.

Value-Driving Attributes

Despite the large number of science goals, the complete list of performance characteristics that define the value of every goal can be condensed into a total of ten attributes, five for each type of asset. Those ten attributes are listed in Table 8-10. Not every goal requires all attributes, but these attributes encompass every science goal.

Table 8-10: Complete set of value-driving attributes of all science goals in GANGMIR.

Small Satellite Value-Driving Attributes	Penetrator Value-Driving Attributes
Maximum Revisit Time	Number of Ground Penetration Locations
Limb Sounding Resolution	Penetrator Weather Package
Number of Infrared Radiometer Bands	Penetrator Spectrometer Package
Nadir-Pointing Camera Resolution	Penetration Depth
Small Satellite Operational Lifetime	Penetrator Operational Lifetime

For the small satellite attributes, the number of radiometer bands corresponds to the wavelength bands listed in the order that they appear in Table 8-5. This implies that the more radiometer bands that are in a design's radiometer, the higher the maximum observable wavelength is, but also that smaller wavelength bands are still included. There are no options for radiometers with only the higher bands but not the lower bands. Complete satisfaction for Goal #1 occurs with seven bands (as shown in Figure 8-7) because those are the bands associated with measuring temperature. A radiometer with all nine bands would deliver no extra value towards Goal #1, but it would deliver value towards Goal #6 because minimum satisfaction requires five bands while full satisfaction requires all nine.

The penetrator weather and spectrometer packages are discussed in more detail but behave the same way as the radiometer bands; lower package attribute levels correspond to less capable packages, and while some goals may see continuing increases in benefits as these attributes rise, other goals will not meet the minimum satisfaction criteria until a higher level is met. At their cores, these attributes show whether or not these instruments are present in the design.

One observation to note from the value models for the goals related to weather than include penetrators is that a large number of penetrators are required to achieve the minimum satisfaction levels. This is because data from a single ground location is far less beneficial to describing global weather compared to an orbiting satellite, but the combination of ground stations operating in unison leverages **simultaneous sampling** to deliver more benefit overall. This also means that if the cost per penetrator is higher than the cost per satellite, the required investment in penetrator weather packages is greater than the investment for satellites and returns less value.

Furthermore, the k_i weights of the coverage attributes and the minimum acceptable and maximum perceptible levels vary with the goals. Only the first goal has a minimum revisit time of less than 12 hours; for the other goals, revisit time is not as critical, so the minimum level is 24 hours and no extra value is perceived for less than two hours.

8.3.1.3 Value Modeling Process 3: Epoch Characterization

The final process of the value modeling phase examines alternate levels for assumptions that have been made in the stakeholder expectations or variables that are outside the control of the designers.

Performance Context Variables

An early unknown in the performance model was the atmospheric conditions through which a penetrator would fly. Originally, atmospheric density and temperature were varied to check the sensitivity of the flight path to perturbations in one or both of those quantities. However, these variables were later scrapped because of how little sensitivity there is in the accuracy of the flight path and the landing site compared to the necessity of limiting impact velocity. At this time, no additional performance variables are being considered for modeling performance in GANGMIR.

Cost Context Variables

At this time, no additional cost variables are being considered for modeling costs in GANGMIR.

Alternative Stakeholder Expectations

The most obvious stakeholder expectation change occurs when information from a ground penetrator causes the stakeholders or mission operators to change their minds about the value of the science operation as it is being conducted. Because this mission uses ground penetrators that will be deployed separately, this mission can leverage **staged sampling** to seize opportunity that may arise over the mission lifecycle.

The method for how to measure changes in expectations is dependent on the type of discovery that is made. The first question a stakeholder should ask with new information gained over the course of the mission is “Is it more valuable to use an expendable asset to conduct a follow-up observation than it is to conduct an observation on a new target?” If the answer is “yes,” there are a number of positions that the stakeholder could find themselves in.

If there are many deployable assets and an interesting scientific find is made with the first or an early observation, and there are still more assets available that there are landing sites, it would be strategically advisable not to change the course of the mission yet because other landing sites on the original operation plan could also yield results that would warrant a follow-up.

This case study’s value model is particularly intriguing for several reasons. There are threshold returns on the number of sampling sites according to Figure 8-10, though satisfaction can still be achieved with one sample site according to this value model (other stakeholders for other missions may not be satisfied without more sites). In contrast, maximum satisfaction can be achieved with six sites, but there are ten sites that have been identified that are scientifically interesting. This creates interesting future scenarios in the event that there are no penetrator failures, no desire to revisit any sites, and spare assets that can be launched after the primary mission is completed.

However, the obvious operational decision is to go to the highest-priority site first, meaning the mission would derive the most value from that location, and each subsequent location would deliver less value than the previous. It is important to remember that “value” as defined in this thesis is not a currency or ratio comparison, but a ranking that can be used to analyze and compare designs. Just because the attribute utility curve exhibits threshold returns (more specifically, increasing returns in the beginning) does not mean these are “less valuable” than the ones that come after it; this simply means the stakeholder is barely satisfied with a lower number of sample sites and would prefer to collect more.

In general, the realm of future possibilities that could influence how a mission captures opportunity as a result of staged sampling is a spectrum, but for purposes of this case study, a single deployable, expendable asset can result in three possible epoch scenarios (excluding a crash landing or a failure, which is characterized as a disturbance, not a shift):

- The asset may discover nothing noteworthy (no shift)
- The asset may make a minor discovery that could be noteworthy and may warrant a follow-up observation on the order of value as one of the least valuable assets.
- The asset may make a major discovery that fundamentally alters the perceptions of value of the mission, and a follow-up to confirm this data is as valuable as or even more valuable than the previous mission assets combined

In the event of a minor discovery, an additional follow-up measurement to the same location early in the mission may not be worthwhile if other sites that are considered valuable have not been explored. However, if all the planned locations have been sampled and spare assets remain, there is opportunity to either conduct the follow-up measurements or decide whether to identify additional, unexplored landing sites. Such a consideration would be made at the end of the mission lifecycle, but adding landing sites for a mission extension will be considered out of scope for this case study.

In the event of a major discovery, the value of the *next* asset could be worth as much as the combination of all previous assets. In such a shift like this, if there are no additional assets, the situation shown in the bottom left of Table 4-2 occurs, and the stakeholders may feel like they have missed out on a serious opportunity to conduct follow-up measurements to confirm the discovery.

The epoch shifts described that could be applicable to **staged sampling** in GANGMIR (and would work the same way with **sacrifice sampling**) are summarized in Table 8-11 along with some possible values for beginning the exploration of how the value can change over eras.

Table 8-11: Summary of possible epochs in GANGMIR that could occur through staged sampling.

Value Model	Epoch Shift Occurs	Estimated Probability	Affine Shift in Value of Assets
Primary Model	Starting point	---	0%
Minor Discovery	Interesting result	10% or 20%	~20%
Major Discovery	Ground-breaking result	2.5% or 5%	~50%

The “affine shift in value of assets” can be thought of as the amount that the MAU function will shift. For instance, a groundbreaking discovery would “double” the maximum expected value, so an asset that cannot capitalize on the discovery but was initially assigned $MAU = 1$ would now be $MAU = 0.5$. Likewise, an interesting result that a follow-up measurement would add value to would lower the same design from $MAU = 1$ to $MAU = 0.8$.

It is important to note that it is irresponsible to assume some probability of a major discovery like finding concrete evidence of organic life, but the exercise of EEA will show that within the spectrum of possibilities, there is added value in carrying extra expendable assets, not just to mitigate risk but especially to capitalize on opportunity that presents itself during the mission lifecycle.

However, these epoch shifts cannot simply be estimate through single- or multi-epoch analysis; these shifts occur *because* of the mission design *through* a discovery made. If there are no assets to make a

discovery, no epoch shift occurs, and if there are no spare assets, no opportunity can be captured. Simply shifting the value model on every single design does not work in this case; for this reason, the standard implementation of multi-epoch analysis will not work.

8.3.2 GANGMIR RSC Phase 2: Performance Modeling

The performance model of GANGMIR is the only one out of the four case studies that was built in collaboration with others. While much of the background data had been gathered by the systems engineering class before the author arrived in Italy to work with other graduate students, little had been done to examine how to build the tradespace and evaluate and compare designs.

8.3.2.1 Performance Modeling Process 1: Identifying Architecture and Design Variables

The first process of performance modeling had already begun when the author joined with colleagues in Italy to begin working on this case study. Some Quality Function Deployment (QFD) had already been conducted to identify potential design variables, but in this case, the design space was *too* rich to constrain the tradespace to a small set of variables alone. The performance modeling phase went through many iterations as design variables that were initially considered to be important were taken out, while others that were uncovered later were added with each iteration of the value model. This section describes how the design space was built up to consider over ten billion unique designs and then reduced to the order of ten million.

Building the Design Space

The architectural variables were known from the beginning of this case study; one of the many goals of the mission was to explore how penetrators could be used to study Mars, and combined operations and sets of measurements required both ground and orbiting assets (as shown in Figure 8-3). It was natural to assume that there are two architectural variables: penetrators and small satellites. A mission could use one without the other, or both, leaving three possible options to consider.

A Vensim diagram was constructed to help better understand the relationships among the many variables that impact the tradespace. The most important outcome from this activity was discovering that the performance model could be constructed entirely without feedback loops, meaning that calculating the performance, mass, cost, and utility of a design would be computationally inexpensive compared to the other case studies in this dissertation and would require no optimization or algorithmic strategy to prevent errors in the code (this does not mean that each module would be without its computational challenges, however). The Vensim model diagram is shown in Figure 8-11.

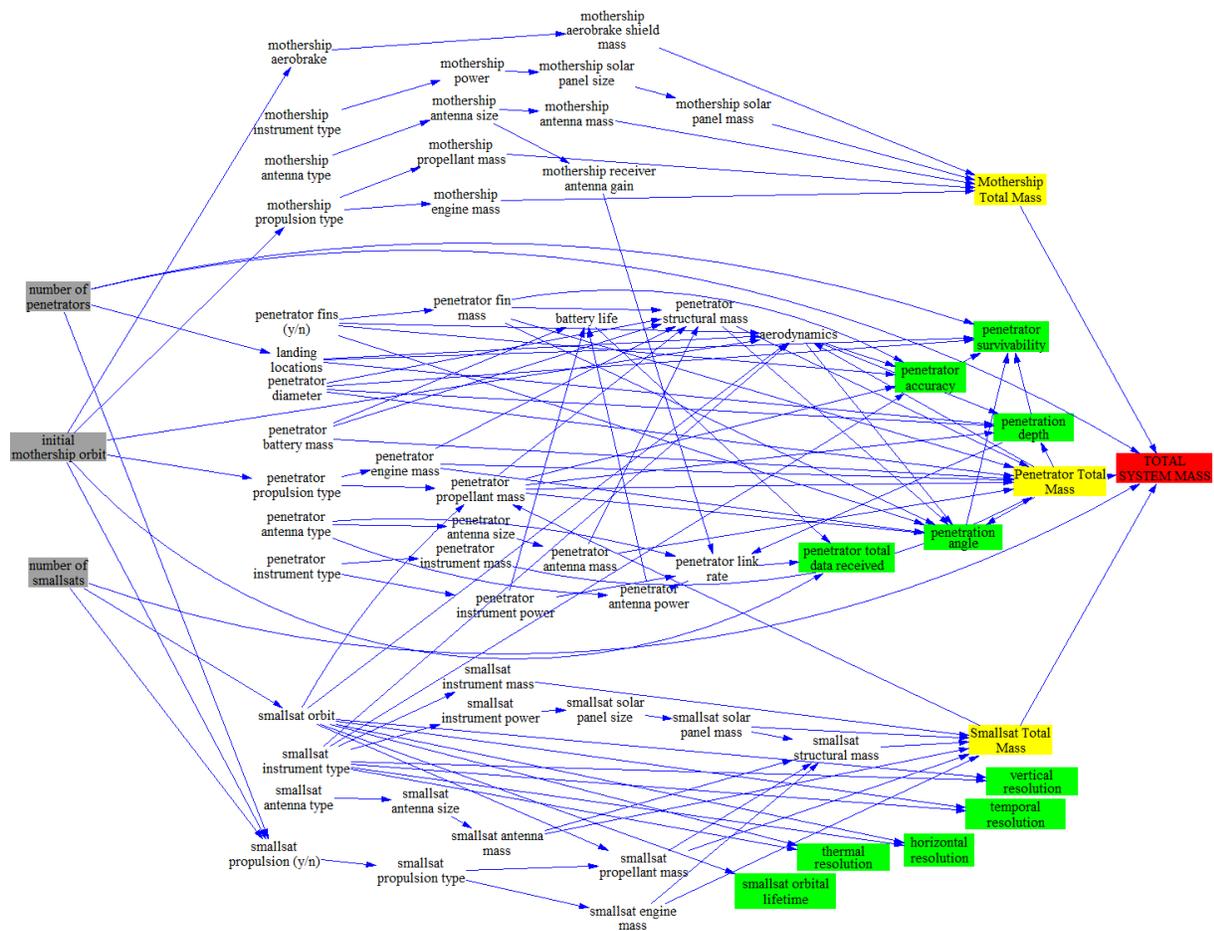


Figure 8-11: Vensim diagram showing the interactions of potential GANGMIR variables.

The Vensim diagram shows three major clusters of variables; from top to bottom, those groups related to the mothership carrier, the penetrators, and the small satellites. The connections between these groups show the relationships among these sets when all three architectural elements are present.

Further modeling and iterating on the performance model showed that what was originally thought to be a very important design variable was either inconsequential to the tradespace or better suited as a context variable since it related to an operational unknown rather than a design-driven characteristic that could be controlled on the ground.

Limiting the Design Space

If every design variable had been left in the tradespace and been allowed to take all of the original levels that had been assigned to it, the design space would have had 13.4 billion designs. This is too many designs to work with, but luckily many were able to be removed for a number of reasons. The final list of design variables and the values they can take are shown in Table 8-12.

Table 8-12: Design variables and their discrete values for GANGMIR

Design Variable	Discrete Values
Number of Penetrators	[0, 1:16]
Number of Small Satellites	[0, 1:16]
Penetrator Shroud Type	[3]
Penetrator Speed Decrease	[500]
Impact Ground Site Type	[Frozen Soil]
Small Satellites Per Penetrator (specified ratio)	[0 1 2]
Mothership Periapsis Altitude	[200 km, 300 km, 400 km, 500 km]
Mothership Orbital Eccentricity	[0.1, 0.2, 0.3, 0.4, 0.5]
Small Satellite Instrument Package	[Camera, Radiometer, Both]
Small Satellite Camera Aperture	[3 cm, 6 cm, 9 cm]
Small Satellite Camera Focal Length	[1.4, 2, 10]
Small Satellite Radiometer Aperture	[2 cm, 3.5 cm, 4.5 cm]
Small Satellite Radiometer Maximum Wavelength (Number of Radiometer Bands)	[17 μm , 23 μm , 32 μm , 43 μm] ([5, 6, 7, 9])
Penetrator Spectrometer Package	[0, 1, 2, 3] (see Table 8-13)
Penetrator Weather Package	[0, 1, 2, 3] (see Table 8-14)
Penetrator Lifetime	[0 years, 3 years, 5 years, 7 years]

A brief interlude is warranted to explain the discrete choices for these variables. It was obvious that the number of penetrators and small satellites could be large, so higher levels for the first two variables were chosen than are necessary to achieve full satisfaction of all the goals.

“Penetrator Shroud Type” refers to the shape and size of the drag shroud on the penetrator. In the end, the EDL module showed that only one of the chosen shapes provided enough drag to impact at an acceptable velocity and impact angle (see Section 8.3.2.2 “Penetrator Entry and Descent Module” for more). This module also showed that lower levels of deorbiting ΔV (“Penetrator Speed Decrease”) allowed for high impact angles and lower velocities while also requiring less propellant.

Young’s parametric equations for impact depth modeling show how ground penetrability plays an important role in modeling not only how deep a penetrator will drive itself into the regolith, but also in calculating the impact deceleration to ensure the penetrator will survive. No published literature was found on the subject of the penetrability of the regolith, but initial performance model results showed that a penetrator that would be buried too deep underground for communications to be possible in one ground type (such as soft clay or dirt) would be completely destroyed on impact with another ground type (solid rock or concrete). Since every penetrator would be designed with the same impact survivability, and because there is evidence that most of the landing sites have high water content and are cold enough for water to be in a solid phase, it was assumed that ground penetrability would be equal to that of a mixture of dirt and ice (see Section 8.3.2.2 “Penetrator Impact Module” for more). This is why “Frozen Soil” is the only option for “Impact Ground Site Type.”

“Mothership Periapsis Altitude” is a design variable because it is also the (circular) orbital altitude of the small satellites and the starting altitude of the penetrators. Lower altitudes for the small satellite orbits increases spatial resolution of the instruments, but it also decreases maximum orbit lifetime and thus mission operational lifetime. For the penetrators, higher altitudes means more propellant is necessary to reach the injection altitude of 100 km before the final maneuver.

“Mothership Orbital Eccentricity” determines the relative precession rate between the mothership and the small satellites, thus affecting the amount of time it takes to deploy a Walker constellation without using expensive inclination change maneuvers (which are not considered part of the design or operation space). Higher mothership eccentricity requires more propellant to circularize the small satellite orbits. The remaining design variables and the discrete levels they can be set to will be discussed in the appropriate subsections of Section 8.3.2.2.

Given this more limited design space, without degeneracies, there are still 93 million designs; however, because some design variables are not present when one asset is present, the design space shrinks. The design space can shrink significantly if one assumes that the ratio of small satellites to penetrators is fixed as determined by the “Penetrators Per Small Satellite” design variable (i.e. there can be 3 penetrators and 3 small satellites for a 1:1 ratio, or 3 and 6 for a 1:2 ratio, but not other combinations besides 3 and 0).

8.3.2.2 Performance Modeling Process 2: Design Modeling

The performance model for GANGMIR consists of a variety of modules that have been built through collaboration with students at Politecnico di Torino. These modules include models and simulations for: small satellite mass, performance, and lifetime; penetrator entry, supersonic descent, and penetration; observational payload performances; orbital dynamics; propulsion; coverage; power generation and storage; intersatellite communications links; and aerobraking. A summary of how the modules work along with a code map can be found at the end of this section.

Asset Modules

The two assets, penetrators and small satellites, were modeled separately but shared similar estimates for off-the-shelf components such as star trackers, reaction control wheels, transmitters, and thrusters. If a design calls for a penetrator to deliver a small satellite or two, then the models are brought together and the shared components are calculated.

Penetrator Instruments

The mass and power estimates for the penetrator instruments were taken from the LunarNET and METNET instruments in Table 8-3 and Table 8-4. One important note to make based on other literature that has been reviewed is that the LunarNet drill is appropriate for dry soil; it is *not* appropriate for soil that has significant water content because of how high the unconfined compressive strength of the terrain is because of the ice. For this reason, the drill’s expected mass and power were tripled from what is shown in Table 8-3 based on work by Zacny et al. [515]–[518]. The possible spectrometer instrument packages the penetrators can take are listed in Table 8-13, and the possible weather instrument packages are listed in Table 8-14.

Table 8-13: Instrument options for the penetrator subsurface package.

Subsurface Package	Instruments
Package 0	Accelerometer
Package 1	Package 0 + Drill, Mass Spectrometer
Package 2	Package 1 + X-Ray Spectrometer
Package 3	Package 2 + Water/Volatile Detector

Table 8-14: Instrument options for the penetrator surface package.

Surface Package	Instruments
Package 0	No Instruments
Package 1	Thermometer, Hygrometer, Barometer
Package 2	Package 1 + Panoramic Camera, Solar Irradiance Sensor
Package 3	Package 2 + Dust Camera

In addition to the instruments, the penetrators also contain a primary battery that is sized according to the energy requirements of the spectrometer package. This battery is not rechargeable. If a long-term weather package is present, a rechargeable lithium-ion battery with an energy density of 125 W-hr/kg powers the payload and is recharged with a small solar panel. This solar panel is assumed to be low-efficiency because it must be ruggedized for impact, and a low estimate is taken on the necessary size due to the possibility that it may be angle away from the Sun for most of its lifetime. Both batteries are sized to operate the payloads and transmit data back to the mothership.

Data on the instrument mass to penetrator impactor mass ratios of other penetrator programs was collected by Fabio Nichele (impactor mass denotes the mass of the components intended to survive impact, not the penetrator delivery system, which includes fuel, thrusters, and guidance packages). There appeared to be an exponential relationship between the mass of the instruments and the mass of the impactor; the lightest penetrators, DS2, had the highest instrument to impactor mass ratio, while the largest penetrator, the Mars 96 penetrator, had the lowest. However, a best-fit exponential parametric model using all of the available data is not necessarily the most accurate estimator. Although the DS2 micropenetrators were *designed* with the lowest mass of the available options, they also failed; using that design as a low-end baseline is not necessarily appropriate. The METNET penetrators are also designed to slow down significantly without penetrating very deeply; as a result, there is less mass in the structure but more mass in the descent module compared to other penetrators that impact at higher speeds.

Given the data from previous and planned missions and estimates of the impact forces from Young's equations, the group working in Torino agreed that the best fit for a tradespace model that produced reasonable estimates across the range of design variables was

$$M_{P,I} = 2.4(M_{P,pay})^{1.5} \quad (8-1)$$

where $M_{P,I}$ is the impactor mass and $M_{P,pay}$ is the total mass of the instruments.

The mass of the shroud was simply assumed to be 3 kg because of its size and the fact that it needed to house the communications system and shock cord to stay connected to the penetrator. Future modeling can improve on this estimate.

The remaining mass of the complete penetrator unit comes from an additional transmitter, additional computers for guidance, a star tracker for navigation, reaction control wheels, a small rocket motor, and the required propellant given the mothership insertion altitude, eccentricity, and information on whether or not the penetrator is also carrying small satellites.

Small Satellite Instruments

The equations presented in SMAD for sizing optical payloads form a solid basis for predicting the mass and power requirements of similar instruments; however, in the case of MCS, there are some caveats that should be taken into account before blindly applying these equations. First, MCS has *two* telescopes, whereas the radiometer for GANGMIR intends to only use one. However, the mass of MCS also included thermal protection, spectrometer hardware, and other structure that both telescopes shared, so multiplying by a factor of ½ seemed inappropriate. Additionally, the power requirements would not necessarily scale with telescope size, so it is unwise to assume SMAD’s power scaling law is appropriate here.

Instead of using the exact forms of the SMAD equations, the following method was used to estimate mass and power based on MCS. First, the two-telescope MCS mass was multiplied by 0.7 to adjust for a single telescope radiometer. Next, it was assumed that technology advances in the past 15 years would have reduced the required power by a factor 0.6. Finally, it was assumed that the mass would scale with both the aperture diameter *and* the number of spectral bands, and because power is unrelated to aperture diameter, power scaled with the number of spectral bands only. The aperture follows the same cube scaling law as SMAD, whereas the spectral bands were assumed to follow square root laws. The following equations were then used to estimate mass and power requirements for a sounding radiometer given an input aperture diameter and number of spectral bands used:

$$R_A = \frac{D_A}{D_{AMCS}}, \quad R_B = \frac{N_B}{N_{BMCS}}, \quad M = R_A^3 \sqrt{R_B} M_{MCS}, \quad P = \sqrt{R_B} P_{MCS} \quad (8-2)$$

Here, R_A is the ratio of the aperture diameter of the GANGMIR design over the MCS aperture, R_B is the ratio of the number of bands, M is the mass, and P is the power.

The mass and power requirements for the camera payload were estimated based on a linear interpolation between the aperture diameter and focal length of the three cameras that were used for comparisons and are shown in Table 8-6.

Penetrator Entry and Descent Module

The module that computes penetrator entry and descent to estimate distance error and penetrator impact velocity required the most effort to develop and test, but in the end, it is the least consequential module of the entire GANGMIR performance model.

Enrico Cumino, a graduate student at Politecnico di Torino, built a three-dimensional EDL simulation for Earth’s atmosphere that includes Coriolis rotation. Lorenzo Feruglio further modified it with a higher-fidelity atmospheric model. With Lorenzo’s help, this model was converted to a Mars EDL model with a one-dimensional model of Mars’ atmosphere serving as a baseline before a three-dimensional atmospheric model could be incorporated. The penetrator aerodynamics were modeled with experimental data of several candidate shapes tested in supersonic wind tunnels from several sources. [213]–[217]. This model accounted for the drag curve as these shapes decelerated through the transonic regime.

Sensitivity analysis showed that of all the penetrator drag shroud options in the design vector, only the largest drag shroud produced enough drag to slow to less than 300 m/s and achieve a high enough impact angle to not risk bouncing off the surface.

Furthermore, of the various options for the starting altitude and total ΔV of the penetrator's descent burn, the lowest altitude and smallest ΔV were the best options. Counterintuitively, a more drastic decrease in orbital velocity (higher ΔV descent burn) caused the penetrator to impact at a higher velocity compared to smaller ΔV burns. It was later discovered that this is because a larger impulse decreases the angle of attack, causing the length of the trajectory through the atmosphere to shorten, which decreases the total drag impulse before impact. In contrast, penetrators with a smaller impulse will fly a trajectory that takes them through a longer arc of atmosphere, slowing them down more before impact. These trajectories were decelerated to their terminal velocity. Because those options resulted in an impact at terminal velocity, there was very little variation in the impact speed between different penetrators. Furthermore, given the size of the shroud, the terminal velocity changes little as a function of penetrator mass.

It is relevant to note that development of the performance model began after the science goals had been defined but before more information was made available about the goals' attributes and their acceptable bounds, which is a mistake that is easy to make starting a case study without direct guidance from a PSS or governing authority.

After another set of interviews with students and an examination of the landing site selections, it was clear that the areas of interest for landing sites were larger than the designers originally had in mind. Although the lower-priority landing sites on the list may require precision navigation and aerodynamics control authority, which was descope from consideration early in this case study, the higher-priority landing sites were large. A sensitivity analysis was conducted by varying atmospheric temperature, pressure, wind speed, penetrator descent impulse, and landing site altitude. None of these variations caused the landing error ellipse to be large enough for concern.

Rather than run through a time-dependent differential equation to calculate a penetrator's impact speed, a process that takes approximately one second for each of the millions of designs that require this calculation, all penetrators are assumed to impact at 250 m/s. The actual range was calculated to be between 230-270 m/s for all designs, and the heavier designs impacted at higher speeds because of higher ballistic coefficients.

Penetrator Impact Module

Penetration depth is a function of several variables, including impact speed, impact angle, penetrator diameter, nosecone length, nosecone type, and ground penetrability. Ground penetrability is the only variable in Young's equations that is unknown and subject to uncertainty. Different parametric equations are applied to different ground penetrability values, but the equations vary linearly with this parameter. For loose clay and regolith, penetrability can be as high as 60, but for concrete and rock, it is 0.5. This not only affects the penetration depth but also the g-forces experienced by the penetrator on impact.

Because gamma ray spectroscopy shows evidence that much of the surface of Mars contains a high percentage of water, a good assumption to make is that the surface is similar to permafrost and other ground types with high frozen water content. In these ground types, penetrability is 2.75 +/- 0.5 (for pure ice, penetrability is 4.5 +/- 0.25). This represents a relatively conservative estimate on the penetration depth and gives a conservative estimate on the maximum g-forces experienced by the penetrator. For impact velocities higher than 61 m/s and a permafrost-like ground type, the penetration depth is

$$D_P = 0.0000046PN \left(\frac{m}{A}\right)^{0.6} (V_0 - 30.5) \log(50 + 0.29m^2) \quad N = 0.18 \frac{L_N}{d} + 0.56 \quad (8-3)$$

where P is the ground penetrability, m is the penetrator mass, A is the penetrator frontal area, V_0 is the impact velocity, L_N is the length of the nosecone, and d is the diameter of the penetrator. (Equation

There is not enough experimental data available to estimate the penetration depth as a function of impact angle, or at what impact angle the penetrator would bounce or slide instead of strike and bury itself into the ground. A natural feasibility check is to ensure that the impact angle is less than 30 degrees from zenith, and the vertical penetration depth is the cosine of the angle times the calculated depth.

Small Satellite Coverage Module

The satellite coverages were modeled exclusively in STK. Limb-sounding coverage was modeled as a 5-degree latitude/longitude grid raised 50 km above the surface ellipsoid. Coverage for both fixed-pointing and spinning radiometers were calculated for the different constellation options. If the rotation rate of the spinning radiometer is fast enough (~ 0.5 rev/min, or ~ 3 deg/s), the coverage can be approximated as the same coverage as a downward-pointing camera with a 180-degree FOV. Spin rates below this rate resulted in unpredictable coverage patterns, but because of the sampling time and FOV of the telescope, the assumption of a rotation rate of greater than 3 degrees per second holds for the radiometer designs being considered in this mission.

Nadir-pointing coverages were modeled with a 5-degree latitude/longitude grid on the surface of Mars. Field of view is an important factor with this coverage, and because it is much faster to calculate coverage for a fixed camera than a spinning radiometer, coverages for a variety of FOVs were calculated.

Coverages were calculated over a period of five days. Typical coverage even for single satellites resulted in complete coverage in approximately this time period, but the average revisit time had greater uncertainty as a result of fewer passes. This uncertainty nearly vanishes with multiple satellites, because the number of visits is statistically significant enough to not require more coverage calculation time.

Communications Module

There are no variables in the tradespace that change the performance of the communications system, but it was still necessary to ensure that the assumptions on communications subsystem sizing would be adequate for delivering data.

Penetrator Communications

Initially, there was great concern that a penetrator that impacted on loose, wet regolith would not be able to communicate with the orbiting mothership due to signal attenuation from the water. Helical antennas were initially chosen because they could spiral up the length of the penetrator, take up very little room inside the penetrator, and operate at lower frequencies that were less susceptible to attenuation due to water. However, after analysis with the drag shroud, and because of the high uncertainty associated with calculating signal attenuation, it was decided to simply put the antenna in the drag shroud.

Data volume from the penetrator's short-term instruments is well-characterized, and given the assumptions on the sizing, all of the data can be transmitted back to the mothership within a few passes, depending on the satellite's orbit relative to the landing location, and all data is returned within a few days, well within the lifetime of the primary battery.

The long-term instrument data volume varies depending on the weather package, specifically whether or not a dust camera is present. Even with a higher volume, adjustments in the solar panel size, duty cycle, and antenna power can accommodate the needs of the payload to within the bounds of the models.

Small Satellite Communications

The small satellite assumes a typical CubeSat communications subsystem. Initially, there were no concerns about data volume because radiometer data has a much lower volume compared to camera data. Data volume increases exponentially with high-resolution cameras, but since wide FOV cameras are more important to the science goals than high-resolution images, and the duty cycle of the camera can be scaled back as needed, no additional data volume modeling was conducted to ensure extreme accuracy.

Mothership Communications

The mothership is assumed to have a 2-meter diameter parabolic antenna in order to communicate with the DSN. This size is typical of other satellites in Mars orbit. The massive size of this antenna is why receiving data from the assets with smaller antennas is of relatively little consequence compared to other missions; if each asset was required to communicate with Earth directly, the mission would most likely be infeasible due to the amount of communications hardware that would be necessary.

Systems Integrator Module

Finally, all the system modules were brought together under one integrator that calculated the final performance metrics and total mass of the system. A code map is shown in Figure 8-12. As discussed in Section 8.1.2, the mission “starts” near Earth in a Mars transfer orbit, and each design in the design vector goes through this integrator module. The design variables were selected in such a way that there are no feedback loops associated with the design of the mission (see Figure 8-11). The red arrows represent information that feeds forward to calculate mass and cost, while the green arrows represent information that feeds forward to calculate attribute levels in Figure 8-12.

First, the integrator module calculates the required ΔV to reach Mars and be captured with aerobraking. Because aerocapture (achieving orbit with a single atmospheric pass) has not been demonstrated in a real mission to Mars yet, other missions were used to gauge how much of the required ΔV could be transferred via a combination of aerobraking (achieving orbit with many atmospheric passes) and thrust maneuvers. Approximately half of the required ΔV can be taken from aerobraking given a small heat shield, but since the payload mass is unknown, only ΔV is calculated in this step.

Next, the performance metrics and mass of the small satellites are computed. Rules are in place depending on whether or not the specific design calls for solo small satellites or piggybacking with penetrators to determine the mass of the propulsion system. The penetrator performance metrics and mass are then calculated. Once all of the assets have been sized, the total wet mass of all assets is calculated. Once the final payload mass of the mothership carrier is known, the mothership’s dry mass is calculated. This dry mass includes the communications system, a percentage of the payload mass for structure and mechanisms to deliver the assets, and a heat shield that is sized relative to previous Mars missions.

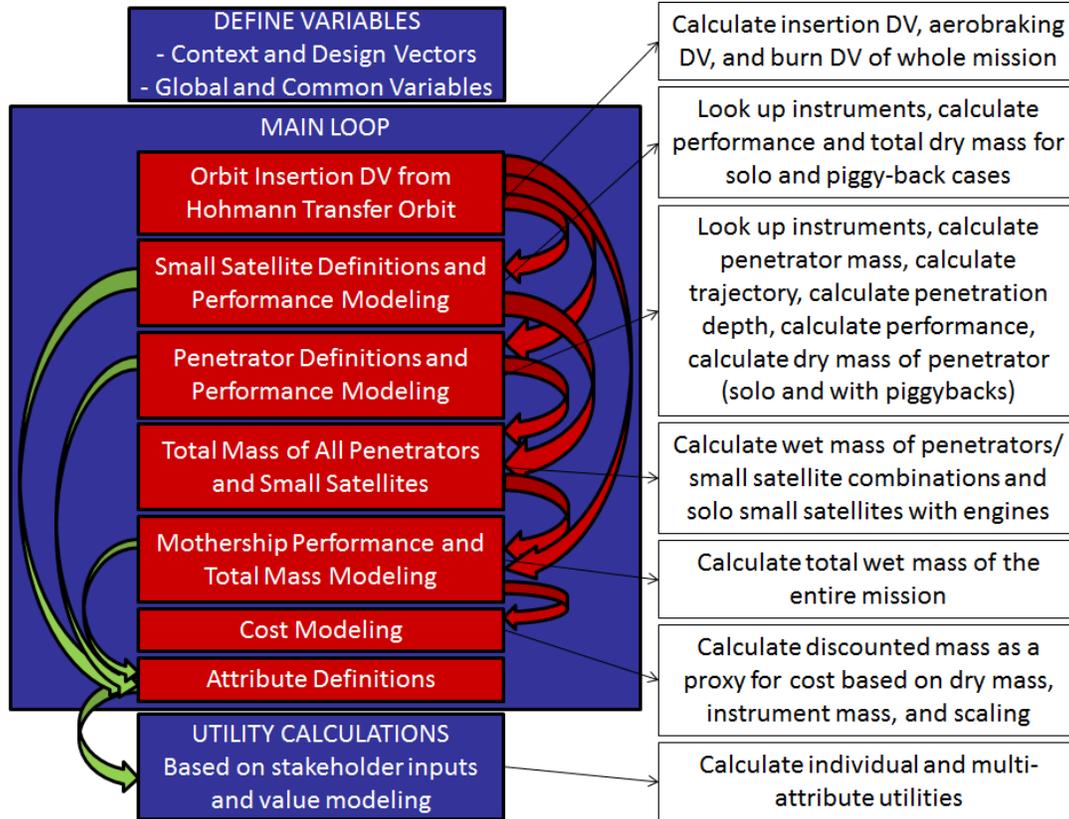


Figure 8-12: Code map of the system integrator in GANGMIR.

Finally, the information from the first block feeds to the final mass modeling block to calculate the wet mass of the total system. From here, the cost can be calculated.

8.3.2.3 Performance Modeling Process 3: Cost Modeling

Cost modeling in GANGMIR follows the same process as shown in Section 4.3.4. The formula for the adjusted spacecraft mass is:

$$M_{SC_A} = c_C M_C + \sum_{i=1}^m \left(c_S M_{Si} + \sum_{k=1}^l (c_{P_{ik}} M_{P_{ik}}) + \sum_{j=2}^{n_i} \left(c_{SD_j} c_S M_{Si} + \sum_{k=1}^l (c_{SP_{jk}} c_{P_{ik}} M_{P_{ik}}) \right) \right) + c_F M_F \quad (8-4)$$

and the values of the coefficients that are being assumed for this mission are shown in Table 8-15.

Table 8-15: Cost Parameters for GANGMIR Cost Estimate

Cost Parameter	Symbol	Value
Mothership Carrier Cost/Mass	c_C	1
Spacecraft Cost/Mass	c_S	1
Payload Cost/Mass	c_P	5
Spacecraft Cost Discount	c_{SD}	0.15
Payload Cost Discount	c_{PD}	0.3
Fuel Cost/Mass	c_F	0.1

For a mission to Mars of this magnitude, it does not seem unreasonable to estimate that costs would be on the order of \$1,000,000 per adjusted kilogram, especially given recent trends in the actual costs of missions to Mars. Thus, the total cost of the mission minus operations costs would be one million times the adjusted mass calculated using the formula above.

8.3.2.4 Performance Module Limitations and Scope

Although the design vector of GANGMIR is the largest of any of the case studies, there are still many variables that have been set as parameters to decrease the size of the design vector that may turn out to drive value, either due to changes in cost that arise during performance modeling or due to changes in stakeholder perceptions that were not captured during value modeling.

One system attribute that could have major effects on the perception of value is the size of the subsurface sample a penetrator can collect and analyze. While the information that was gathered by students at Politecnico di Torino and interviews with other scientists and engineers did not suggest that sample size was particularly important, other scientists may disagree or have perceptions that are based on the quality of the sample. No design variables drive the sample acquisition size, and no design variables change the number of samples a single penetrator can retrieve (for example, a single penetrator could have multiple drills, or the power system could be sized to allow for one drill to make multiple excavation attempts to gather a larger sample than the volume of a single borehole).

Because penetrator technologies are new, a significant mass margin was given to previous estimates for mass of certain systems to ruggedize them for impact. It is not known if there are any solar panels that can withstand the impact forces associated with ground penetrators, so the estimated mass for a spacecraft solar panel of the same size was simply doubled to account for this.

A drag shroud was assumed to be the method by which a penetrator would decelerate through the atmosphere to impact below the maximum survivable impact speed, but further modeling that is beyond the scope of this case study may show that parachutes provide the same relative drag with much lower mass. There may also be a tradeoff between the inherent chaotic uncertainty associated with ballistic flight using a parachute compared to a solid shroud that is less susceptible to perturbations due to turbulent eddies at supersonic reentry speeds.

A penetrator could also be given wings and flown like a drone and land like a dive bomber to achieve high-precision landing, which may be desired for landing in LDAs or even in the sides of cliffs along Valles Marineris (though this was not identified as a scientifically interesting landing site by the Italian group). However, modeling the dynamics and estimating the mass and cost of such a system is also beyond the scope, expertise, and available resources of this case study.

Since this is being viewed as a traditional Mars mission, there is already sufficient technology risk involved with penetrator EDL, and the ΔV budgets are relatively small compared to the previous two case studies, this mission will only consider chemical propulsion systems to maneuver the spacecraft.

Modeling subsurface communications proved to be difficult considering there is no way to know what the attenuation will be due to subsurface water content. For this reason, the drag shroud also needs to survive impact to communicate with the orbiting mothership. This also serves as a convenient platform for hosting the weather observation payloads, but the tradespace does not consider designs without a drag

shroud. Initial plans to use a helical antenna inside the penetrator were scrapped in favor of a patch antenna on the shroud once these difficulties were better understood. The patch antenna is a low-gain but robust option that represents an acceptable worst-case scenario for performance while still demonstrating realistic data rates from the ground to the satellite.

Also outside the scope of this mission are other exotic technologies, such as using electromagnetic rails to fire the penetrators from the mothership instead of using rocket propulsion.

The penetrators will be hot after impact, meaning that whatever organic matter that is in contact with them could be denatured or boiled away. Rather than calculating the distance at which the nearby regolith will be expected to boil, the sampling drill will simply extend a few centimeters into the regolith, where it is assumed that the boiling point was not reached. More precise modeling needs to be conducted to ensure that the chemicals being studied will not be destroyed.

8.3.3 GANGMIR RSC Phase 3: Tradespace Exploration

The design space of GANGMIR is too large to be easily displayed all at once; instead, multiple graphs will be used to examine subsets of this tradespace, and other graphs will show the Pareto fronts of the entire set of designs and several important subsets.

8.3.3.1 Homogeneous Asset Comparison

The instrument-level and goal-level MAUs for all eight goals are shown in Figure 8-13 and Figure 8-14. The top rows show the contribution of the small satellites, the middle row shows the contributions of the ground penetrators, and the bottom row shows the goal-level MAU from all designs that use *only one* of those two assets to illustrate the relative value of each (i.e. these points represent designs that use many small satellites but no penetrators, or designs that use penetrators but no small satellites).

Some of the subplots in these figures are empty; this occurs when that asset does not contribute to the goal-level MAU. For instance, in the top left center plot of Figure 8-13, the small satellites do not have utility greater than zero because Goal #2 is a goal that is satisfied entirely by ground penetrators. Likewise, the left middle plot of Figure 8-14 is empty because ground penetrators do not help observations of the global aerosol composition.

Because the bottom rows of Figure 8-13 and Figure 8-14 show designs that only carry one type of asset, the relative importance of either asset can be easily visualized. Observing the global temperature profile is done primarily with the small satellites, which is why the bottom left plot of Figure 8-13 shows two very distinct sets of points. The designs with only small satellites are in the top left (and make up the Pareto front entirely) while designs with only penetrators never achieve MAU greater than 0.2 because that is the weight k_{ii} that was assigned to the penetrators in achieving that goal, so penetrators can never achieve an instrument-level MAU higher than that.

The graphs representing Goal #3 show the opposite effect; measuring surface pressure from the ground is far more valuable than measuring it from orbit. The top center plot of Figure 8-13 shows that full satisfaction is easily achieved with the small satellites as evident by how far into the upper left corner the designs are on the plot. This is because the minimum revisit time for making these measurements is so much lower, so constellations with fewer satellites can still exceed the maximum perceptible level of

utility. However the total goal-level MAU two plots beneath that shows that the contributions to the goal are eventually dominated by the designs that use penetrators.

For observing weather in Goal #4, the small satellites and penetrators carried equal k_{ii} weights, and the sum of these weights was higher than zero. However, the small satellites satisfy their part of the goal for much lower cost because the small satellites weigh less than the penetrators and because building the same level of acceptability in ground coverage requires more ground assets operating simultaneously.

Small Satellites

The instrument-level MAUs for goals that are satisfied by the small satellites' radiometers are shown in Figure 8-15. The top row shows variations in the number of radiometer bands while the bottom row shows variations in the aperture diameter of the radiometer. From these graphs it is clear that the designs that are Pareto-efficient with high satisfaction for some goals are not Pareto-efficient for others. For instance, measuring pressure requires a lower number of radiometer bands, meaning a shorter observation wavelength, meaning the detector's pixels can be smaller, leading to higher resolution. However, observing dust and aerosols requires a higher number of bands, meaning lower spatial resolution. This tradeoff was expected, but now it is explicitly shown by the variations in utility across these different goals.

Variations in the aperture diameter follow an expected pattern; higher resolution is better, but larger apertures require heavier instruments. The differences in the spacing among these four graphs can be attributed to the relative weights k_{ijk} of the resolution attribute within each instrument MAU.

The instrument-level MAUs for goals that are satisfied by the small satellites' cameras are shown in Figure 8-16. In this figure, it is obvious by the fact that the higher aperture and higher focal ratio designs are dominated by smaller aperture and lower focal ratio designs that these goals do not require high-fidelity measurements. These goals are clearly better satisfied by leveraging more, low-capability assets compared to fewer, high-capability assets. Under this value model, and in these graphs, one of the primary hypotheses of this work is unquestionably proven. These graphs show that higher satisfaction can still be achieved with more capable assets, but the cost to move along the Pareto front toward those design options is much higher per unit increase in utility.

Tradespaces comparing different levels for the small satellite instruments' design variables for designs that only use small satellites (no penetrators) are shown in Figure 8-17. From these graphs, several conclusions can be drawn. First, it is evident by the Pareto fronts of these graphs that according to this value model, where higher weighting is placed on coverage that leverages shared sampling compared to the attributes that measure the performance of a single satellite, smaller apertures, lower focal ratios, and fewer wavelength bands provide more benefit at cost as opposed to the alternative.

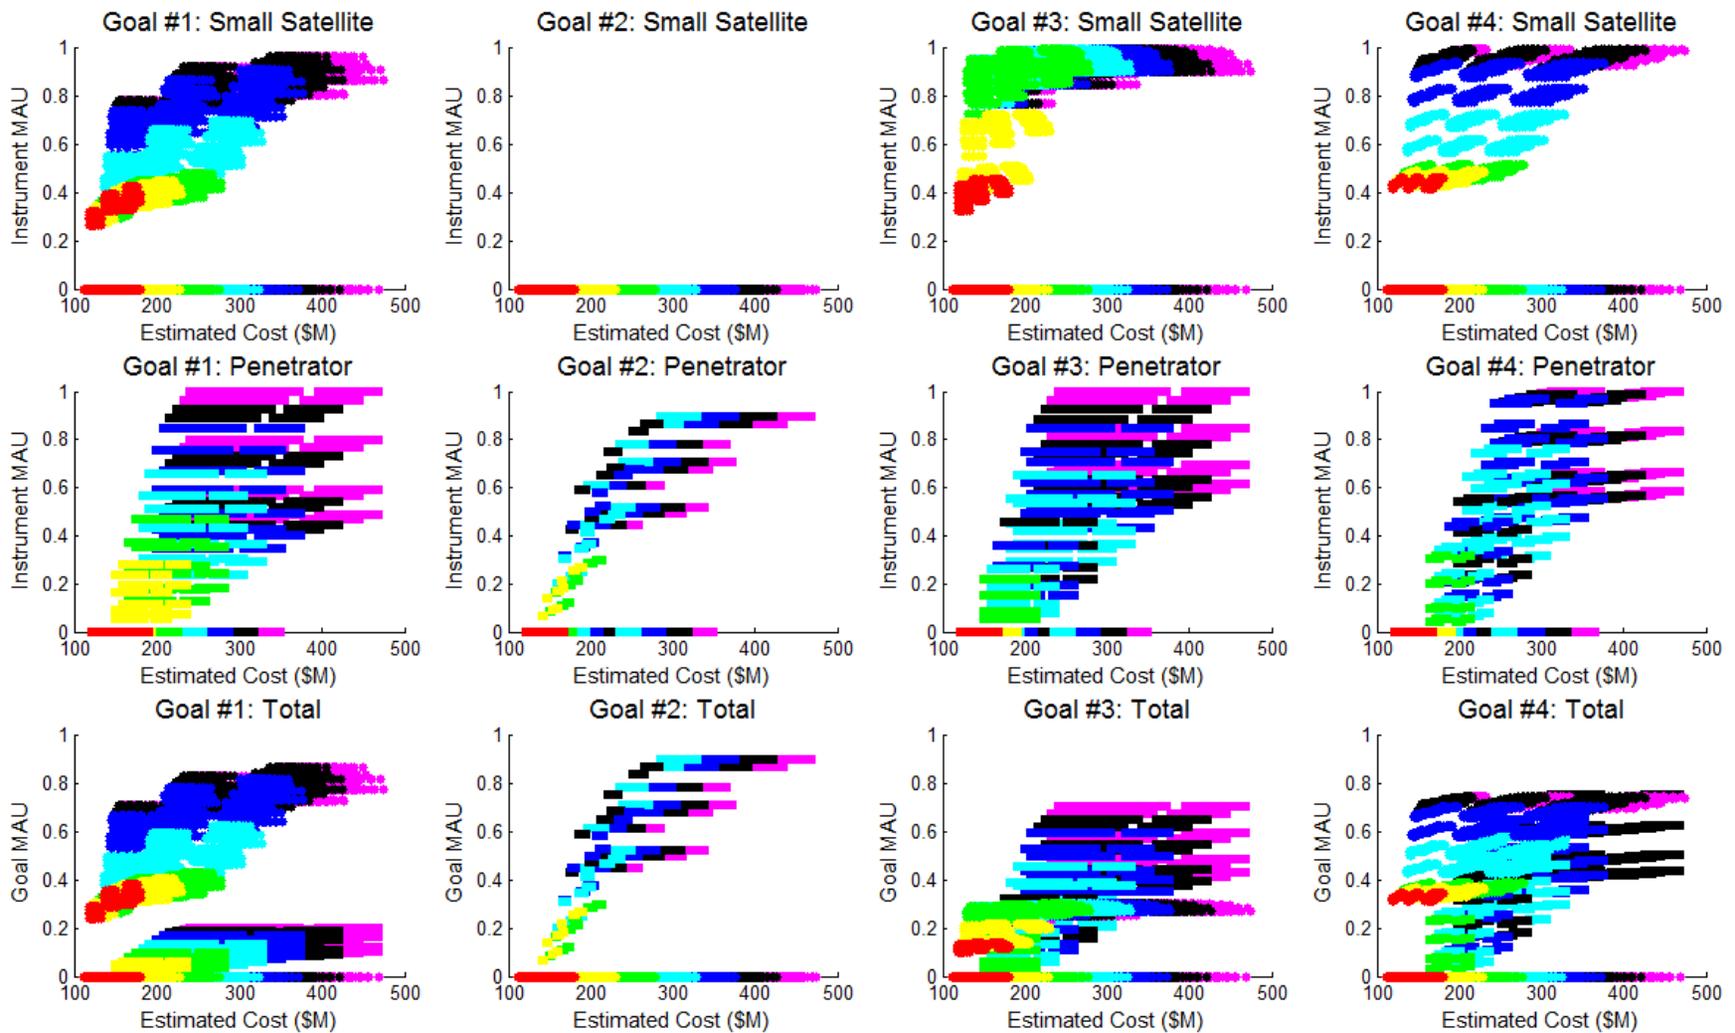


Figure 8-13: Tradespace for Science Goals 1-4 from left to right, showing the gained from (top) small satellites, (middle) ground penetrators, and (bottom) combined utility. Color coding: red = 1 or 2 assets; yellow = 3 or 4; green = 5 or 6; cyan = 7 or 8; blue = 9 or 10; black = 11 or 12; magenta = 13 or 14

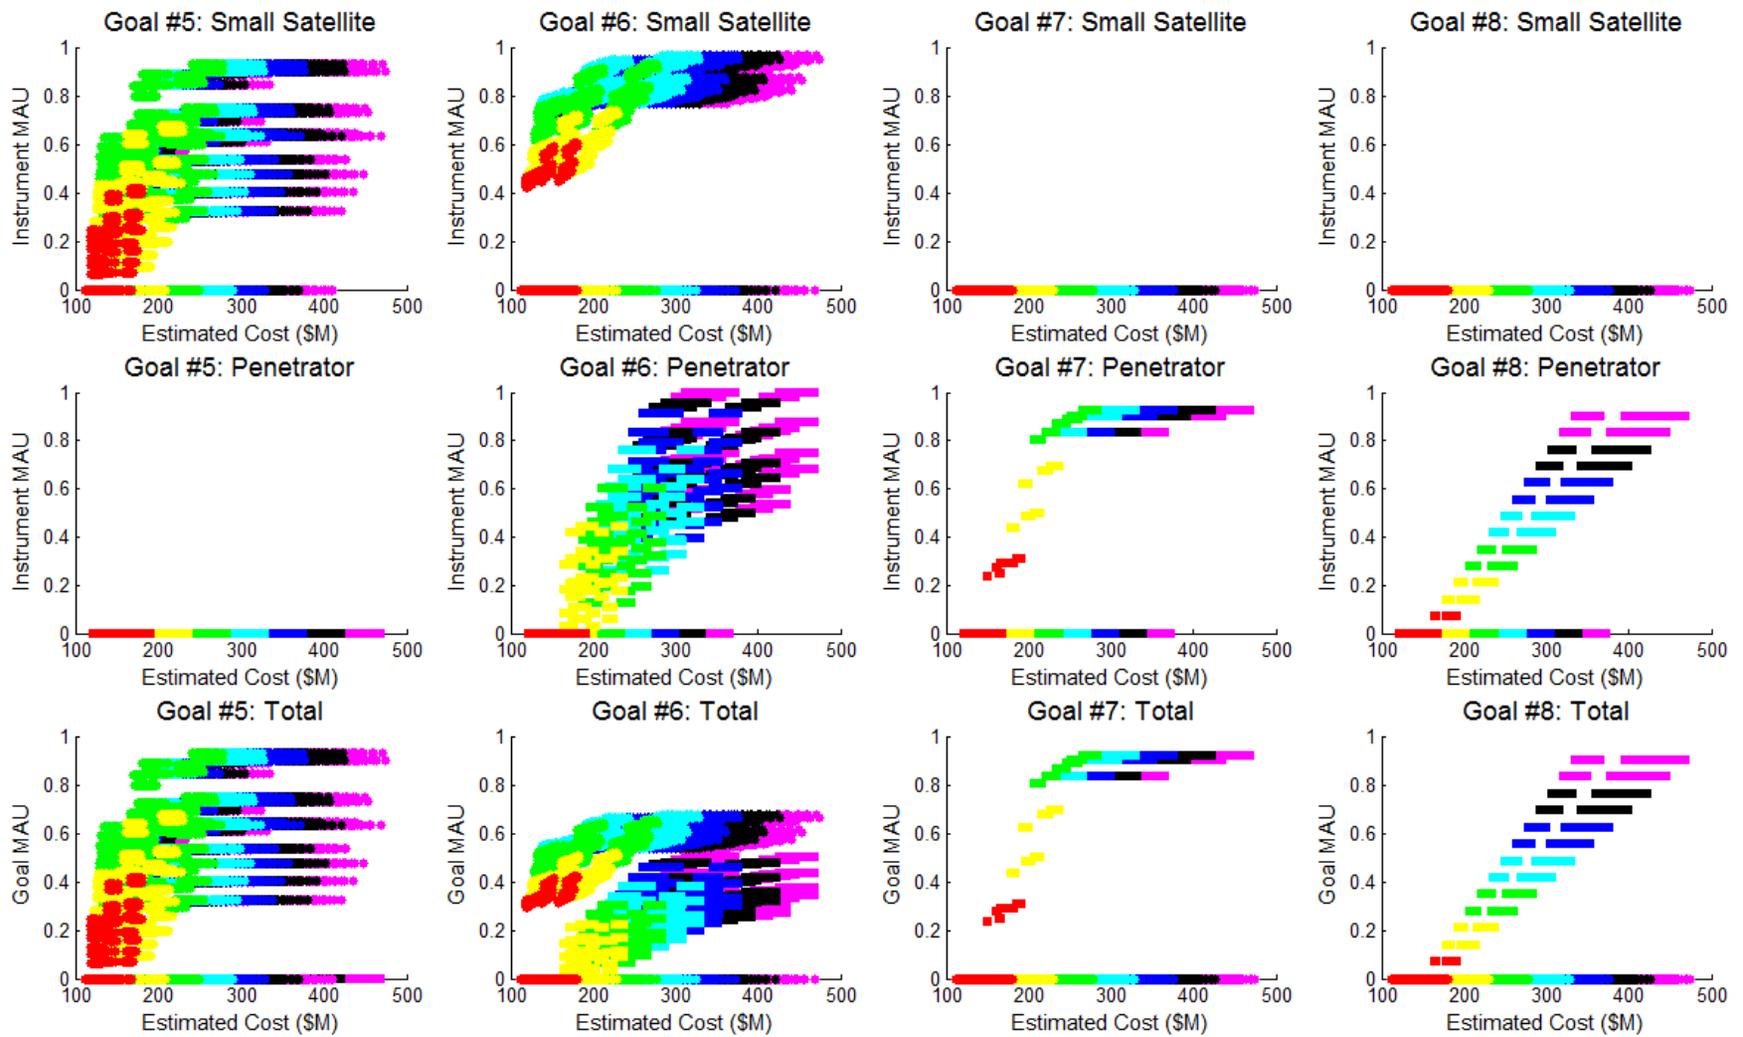


Figure 8-14: Tradespace for Science Goals 5-8 from left to right, showing the gained from the ground instruments (top), orbiting instruments (middle), and combined utility (bottom). Color coding: red = 1 or 2 assets; yellow = 3 or 4; green = 5 or 6; cyan = 7 or 8; blue = 9 or 10; black = 11 or 12; magenta = 13 or 14

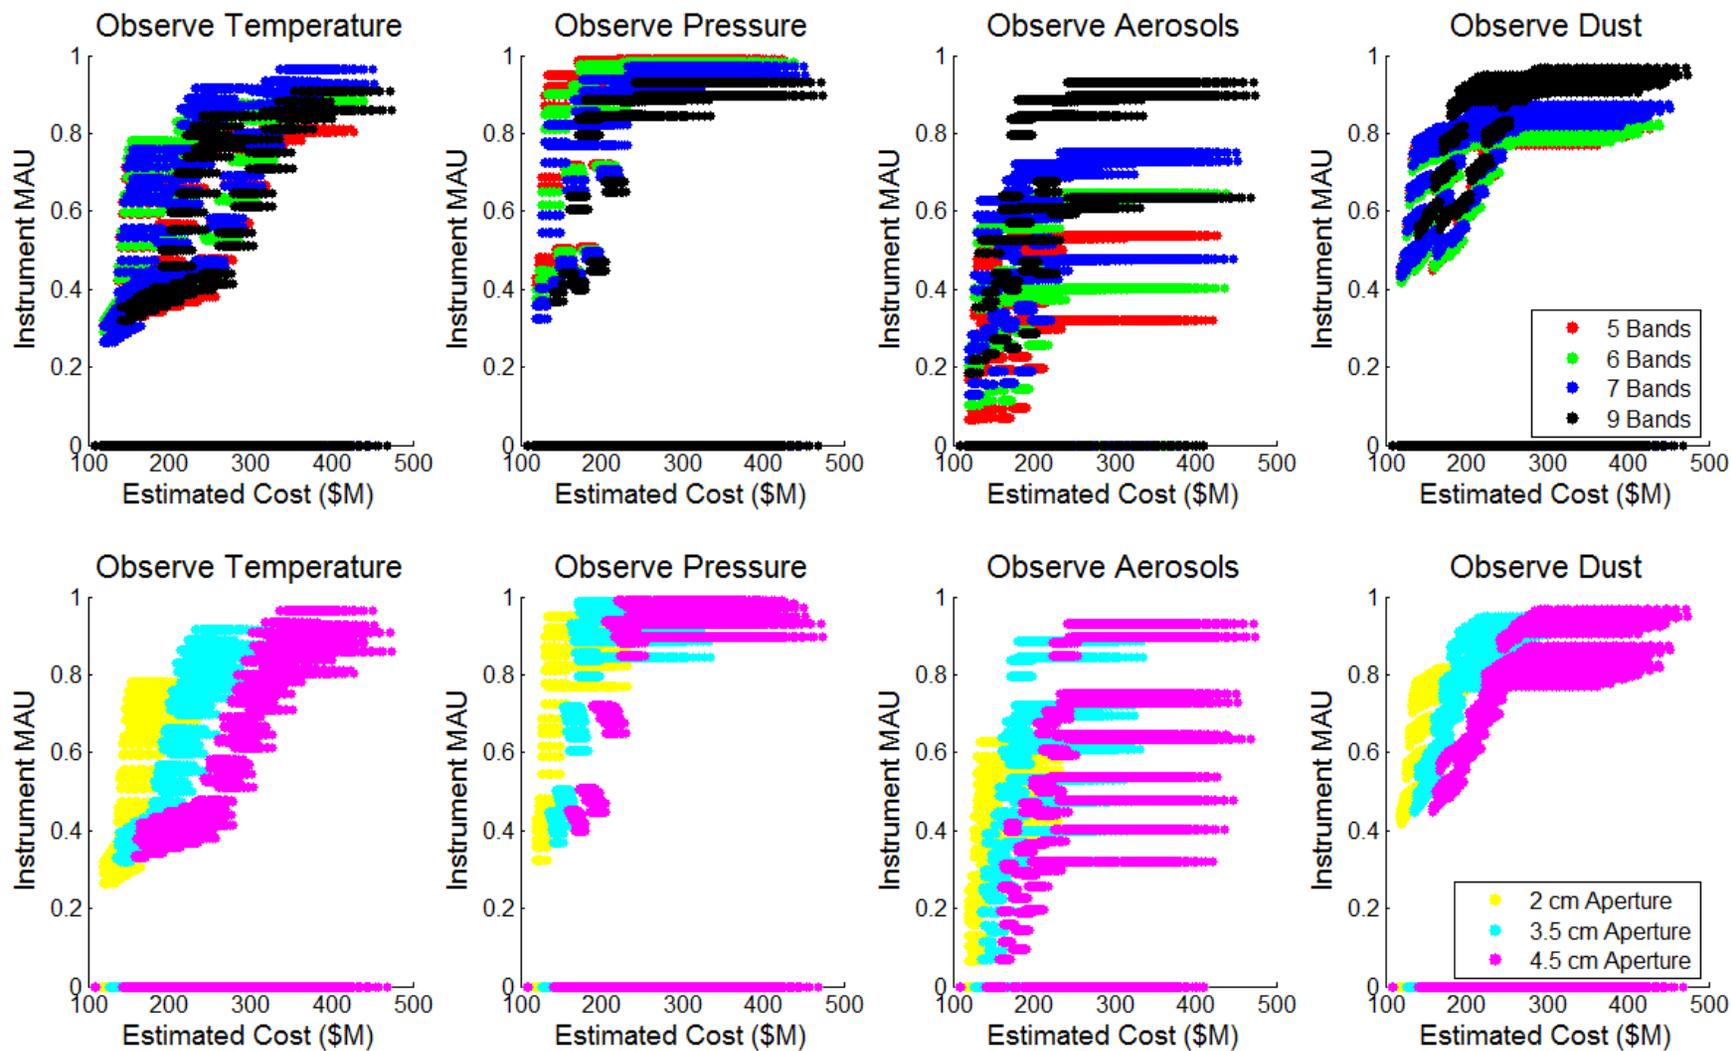


Figure 8-15: Tradespaces of instrument-level MAUs for goals related to the small satellite radiometer showing variations in (top) the number of wavelength bands and (bottom) aperture diameter.

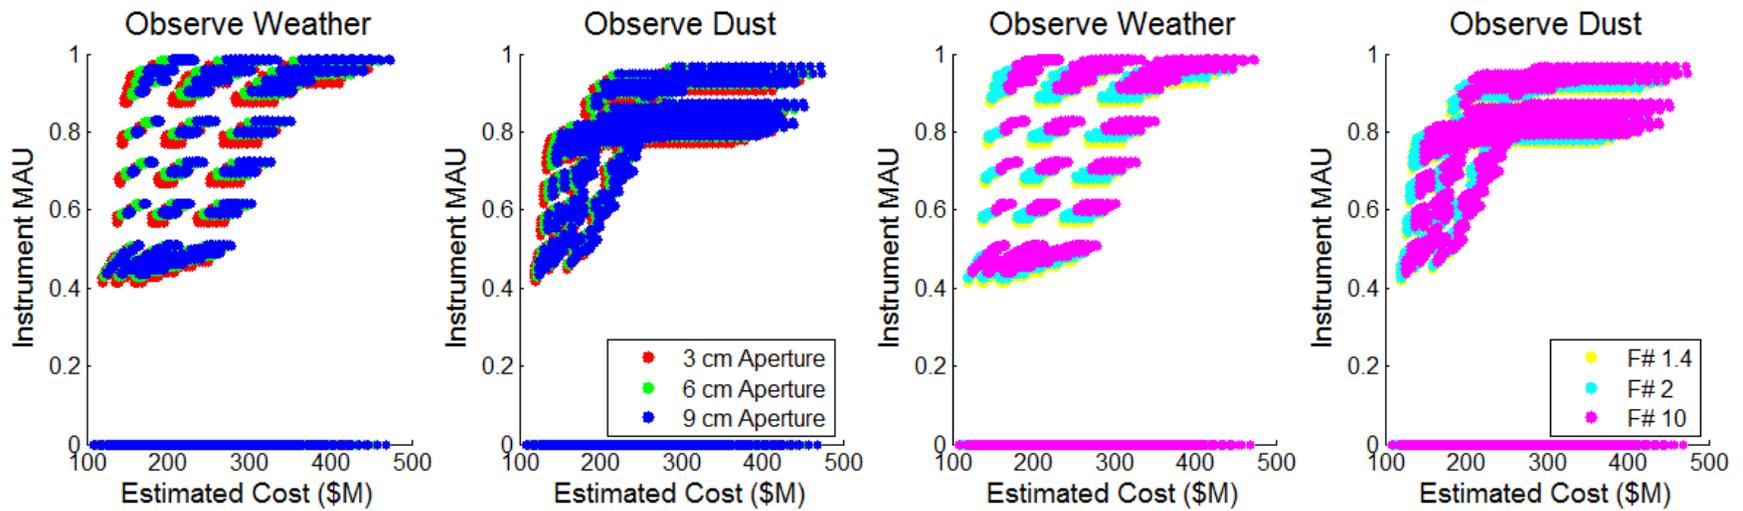


Figure 8-16: Tradespaces of instrument-level MAUs for goals related to the small satellite camera showing variations in (left) aperture diameter and (right) focal ratio.

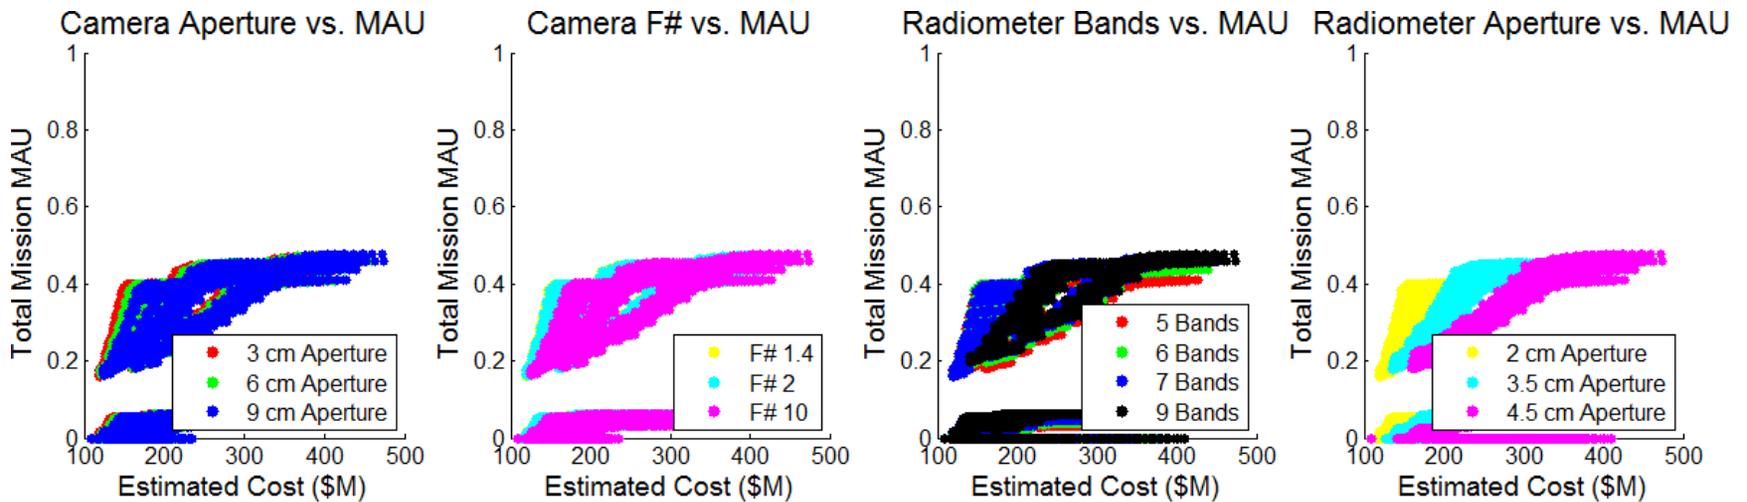


Figure 8-17: Tradespaces of total mission MAU showing how utility varies with (left) camera aperture diameter, (left center) camera F#, (right center) number of radiometer bands, and (right) radiometer aperture diameter.

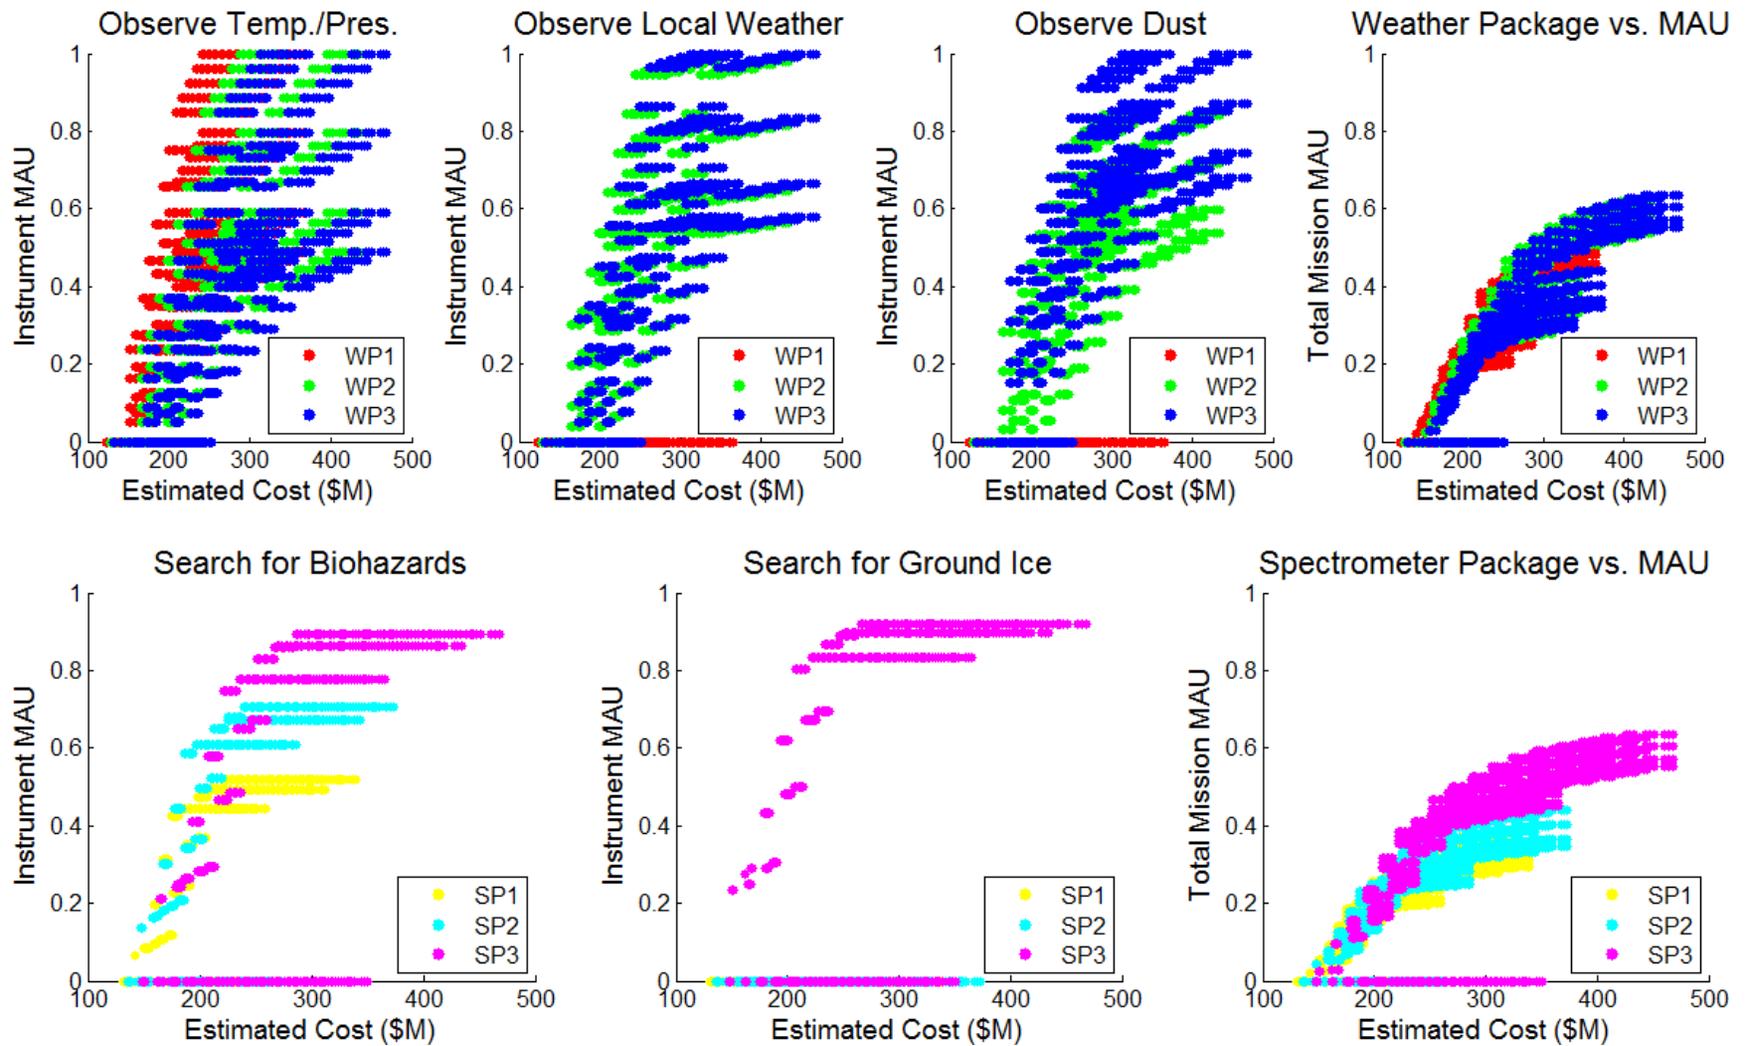


Figure 8-18: Tradespaces of instrument-level MAUs for goals related to the penetrators' (top) weather package options and (bottom) spectrometer package options. The right-most plots show total mission MAU

The left plot of Figure 8-17 shows that cameras with the smallest aperture diameter (red dots) are the most valuable because of how unimportant spatial resolution with a nadir-pointing camera is compared to coverage in the satisfaction of these science goals. This is also shown in the left center plot, where focal ratio is unimportant, though there appears to be very little difference in value between the two lower options. This was not modeled explicitly, but a narrow FOV would require more time to cover the planet, increasing the maximum revisit time and decreasing the utility, so the designs with focal ratio equal to ten should actually be penalized more.

The right two plots of Figure 8-17 shows that designs with five or six radiometer bands (green and blue dots, right center plot) are Pareto-efficient for designs beneath the knee of the curve, which is dominated by radiometers with 2 cm diameter apertures (yellow dots, far right plot). However, radiometers with nine bands and 2 cm apertures have considerably *less* utility according to the right center plot. This is because with such a small aperture using the largest wavelength, the sounding resolution is too poor to provide a few pixels per scale height, which is a desirable trait. The far right plot also shows that a large increase in cost is necessary for designs with larger apertures to become Pareto-efficient (~\$200M), but once this occurs, radiometers with nine bands become Pareto-efficient, thus satisfying the aerosol and dust observation goals better than the less expensive Pareto-efficient designs.

These tradespaces and variable options still do not show why there is such a large separation between two distinct groups of designs that can be seen in every plot in Figure 8-17. A tradespace with variations in orbital altitude plotted shows that the designs with the mothership's orbital altitude of 200 km are all in the lower section. This is because the orbital lifetime of smaller satellites at this altitude is lower compared to higher altitudes. The designs at an altitude of 350 km are Pareto efficient because they balance longer orbital lifetime with higher resolution that can be achieved by being closer to the target than designs at 500 km altitude.

Penetrators

Missions that only use ground penetrators do not appear to be as valuable (provide as much utility per unit cost compared to the small satellites given the structure and weights in this value model, even though three of the eight goals can only be accomplished with ground penetrators. This is because the ground penetrators tend to deliver a smaller relative satisfaction towards the weather goals than the small satellites (excluding the surface pressure goal) and require more mass to achieve that goal.

Tradespaces that show variations in the instrument packages of the penetrators and their relative effects on the science goals related to penetrators are shown in Figure 8-18. The top row shows how utility changes with different weather package options, while the bottom row shows changes due to spectrometer package options. The right-most plot on each row shows the total mission MAU for designs that only include penetrators.

From the top plots of Figure 8-18 it is obvious that the smallest weather packages can satisfy stakeholder needs for measuring temperature and pressure. This is because that package is so lightweight and cheap, and there is no added benefit from the larger packages for that goal. Observing local weather and dust requires a better package, but there is little added benefit for the added cost. The total mission MAU shows that weather packages #1 and #2 dominate the entire Pareto front, meaning that adding a dust camera to the penetrator weather package *never* increases the utility compared to other design options. A

sensitivity analysis in EEA can be used to show *under what conditions* it would be by adjusting the instrument-level and goal-level attribute weights k_{ii} and k_i .

The bottom row of plots in Figure 8-18 shows that there is some balance between spectrometer packages #1 and #2 for the biohazard goal, but the added benefit of having a volatile detector is not Pareto-efficient on this goal until much expensive designs are realized (> \$250M). However, a volatile detector is the only option that satisfies the goal to detect subsurface ice. The Pareto front of the total mission MAU shows that all designs with spectrometer package #2 are dominated; either the mission is less expensive with option #1, or it achieves more with option #3.

Asset Comparison

A tradespace showing the total mission MAU for homogeneous designs with *either* penetrators *or* small satellites is shown in Figure 8-19. In this figure, the small satellites are the designs in the upper left and they clearly dominate the penetrator designs represented by the long sloping curve on the bottom. Penetrator missions do not begin meeting the minimum satisfaction criteria until almost the \$200M point

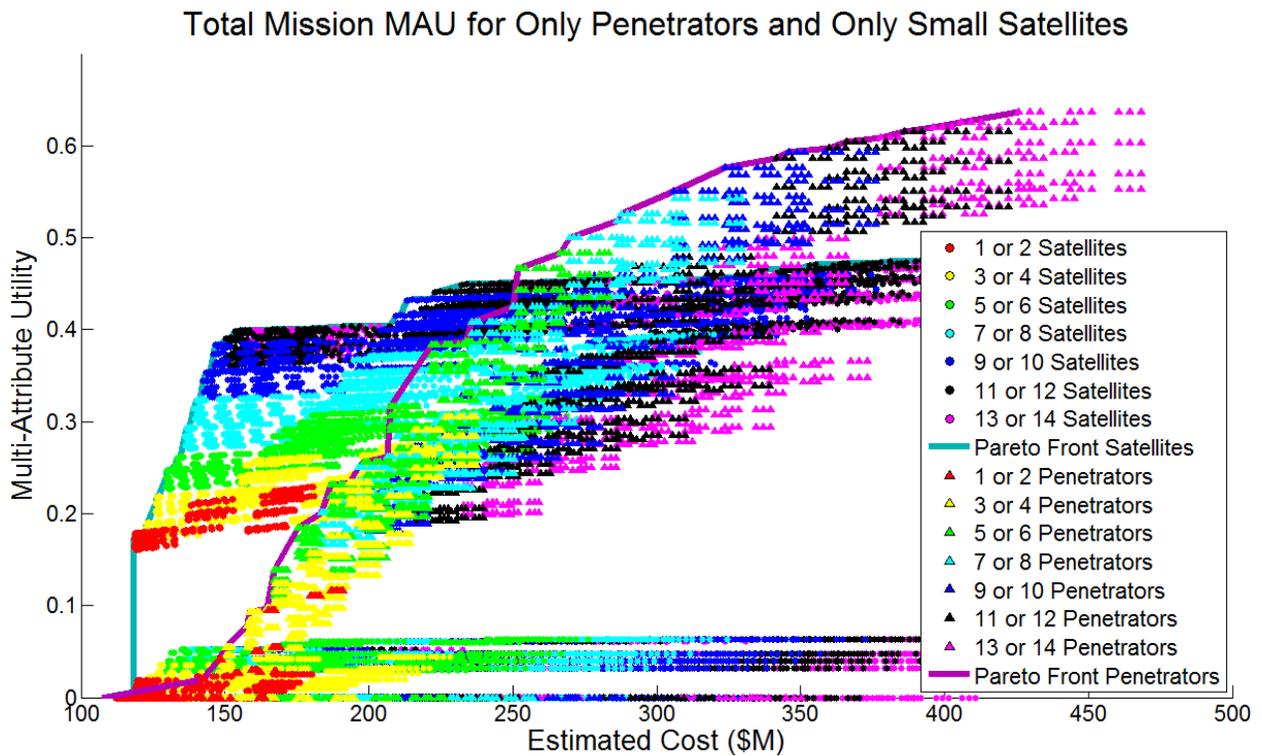


Figure 8-19: Tradespace of total mission MAU for designs with either smaller satellites or penetrators but not both.

Given a different set of top-level goal weights k_i , the relationship between these two sets would look different. If Goals #2, #7, and #8 were all weighted higher, the Pareto front of penetrator designs would start to rise. The priority of the goals that Fabio Nichele created from the MEPAG’s list of possible goals shows how much they value the weather data, but if a mission manager with a background in biology saw this tradespace, he or she may argue that the penetrators give far more value due to the nature of those in-situ measurements.

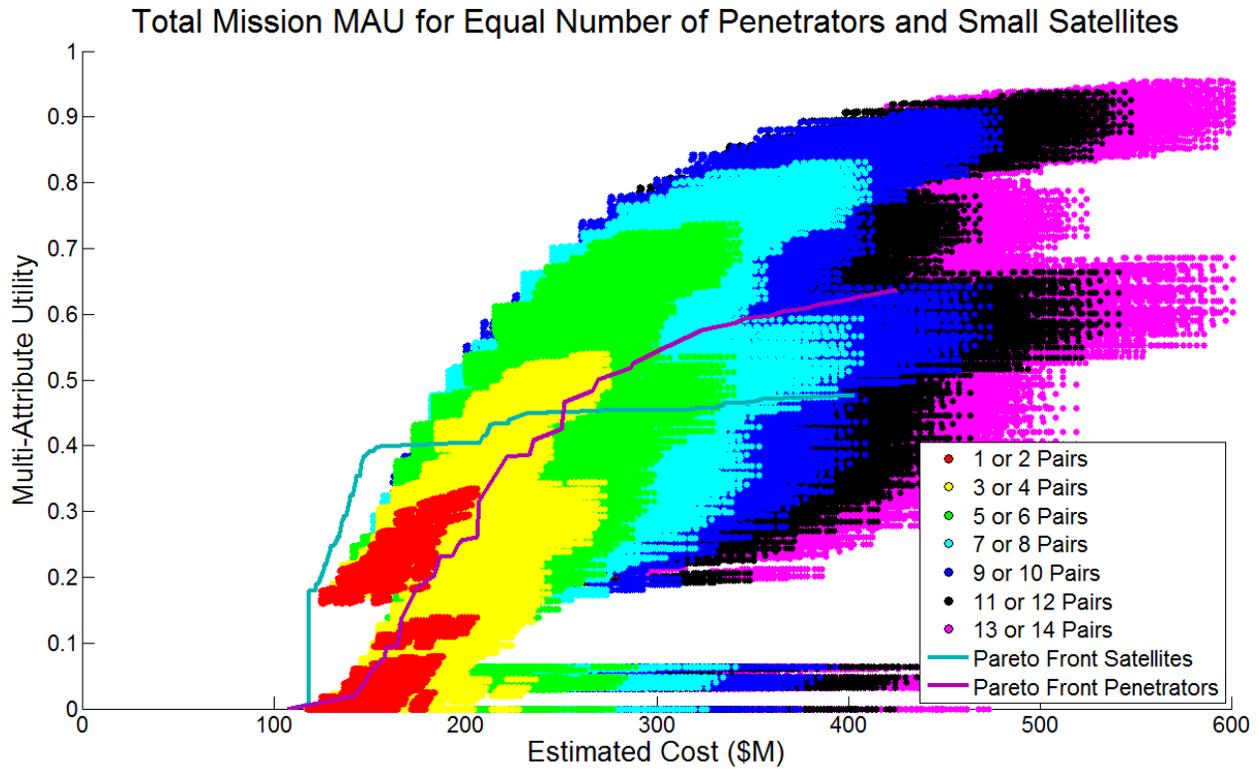


Figure 8-20: Tradespace of total mission MAU for designs with 1:1 ratio of small satellites to penetrators.

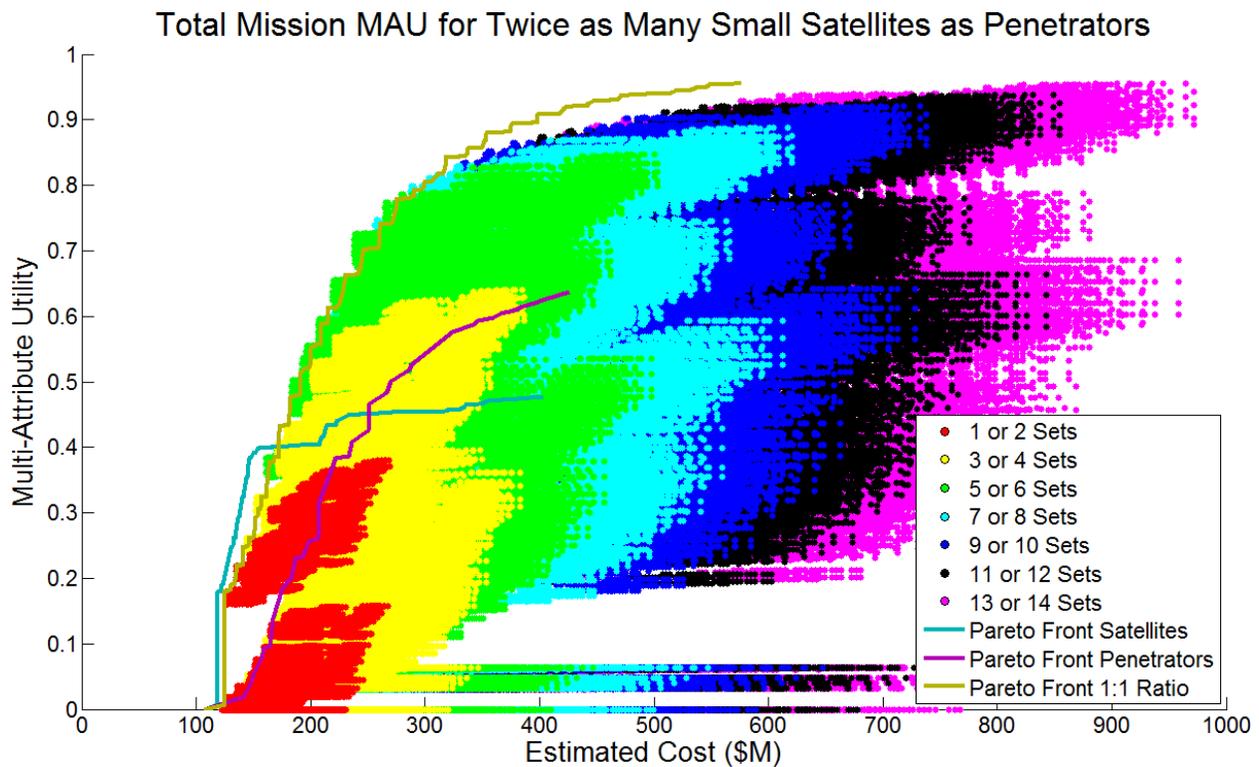


Figure 8-21: Tradespace of total mission MAU for designs with 2:1 ratio of small satellites as penetrators.

8.3.3.1 Heterogeneous Asset Comparison

A tradespace showing designs that share equal numbers of penetrators and small satellites is shown in Figure 8-20. In this tradespace, designs that achieve satisfaction levels higher than $MAU = 0.6$ begin to show. Now, assets are working synergistically not only to leverage **stacked sampling** but also to save mass on redundant subsystems because the penetrators can deliver the small satellites to their orbits. A tradespace showing designs that stack two small satellites onto one penetrator is shown in Figure 8-21.

8.3.4 GANGMIR RSC Phase 4: Epoch-Era Analysis

Even if a scientist thinks that the ground mission is undervalued in the value model shown in Section 8.3.1.2, the most interesting results come from understanding how endogenous epoch shifts can change the value proposition when a design has assets it can spare. This is shown in multi-era analysis.

8.3.4.1 EEA Process 1: Single Epoch Analyses

Because of the lack of tangible epochs in the normal problem setup, single-epoch analyses will not be conducted.

8.3.4.2 EEA Process 2: Multi-Epoch Analysis

Because of the lack of tangible epochs in the normal problem setup, multi-epoch analysis will not be conducted.

8.3.4.3 EEA Process 3: Single Era Analyses

Because of the lack of predictable epochs, epochs in the normal problem setup, single-era analyses will not be conducted.

8.3.4.4 EEA Process 4: Multi-Era Analysis

Multi-era analysis is where the true power of **staged sampling** can be shown. The potential epoch shifts in the epoch characterization process of the value-modeling phase (Section 8.3.1.3) show how the Value at Opportunity (VaO) changes as a result of findings that may occur during the mission lifecycle that could change how a mission manager would conduct the mission.

To review, there were two potential epoch shifts identified for GANGMIR: one that creates the desire to conduct a follow-up measurement resulting in added value that would shift the maximum perceived level of utility for a science goal by 10% (i.e. what was once considered the previous maximum is now $MAU = 0.9$), and one that creates the desire to conduct a follow-up measurement resulting in added value that would shift the maximum perceived level of utility for a science goal by 50% (i.e. what was once considered the previous maximum is now $MAU = 0.5$). These will be called “retake” and “eureka” moments. The probabilities for both of these epoch shifts occurring can be varied, as can the risk of penetrator failure, to study how VaO changes.

Despite efforts in Chapter 4 to explain how the previous value model would shift downward under a new epoch where the maximum perceived utility level was raised, the message is clearer when the higher level is allowed to instead move upward.

If the per-asset value relative to the total value of all assets is higher than the value that could be gained from a retake moment, it will *never* be valuable to send sacrificial assets to conduct follow-up

measurements. However, if there are a sufficient number of assets, there would be a tipping point where the value of a follow-up would outweigh the value lost by ignoring a planned target.

A Monte Carlo analysis was set up to consider the distributed deployment given some number of penetrators to start with n_i . A penetrator is deployed ($n = n - 1$), and a random number between 0 and 1 is chosen. If that number is above the probability of success, the penetrator fails and no value is added. If the penetrator is successful, value is added and another random number is chosen. If this number is below the probability of a retake moment, if there are penetrators remaining, *and* if the added value of a retake moment exceeds the value-per-asset ($V > 1/n_i$), another penetrator is launched. Another random number is chosen to determine whether *that* penetrator succeeds; if it does, the value is added. If the random number that controls the probability of a retake moment fails, another random number is chosen. If that random number is below the eureka moment probability and there are penetrators remaining, a penetrator is automatically launched, because the added value would exceed all the possible value any other target could achieve. A final random number is chosen to determine whether or not that penetrator succeeds; if it does, 1 is added to the total value. Once there are no penetrators remaining, the mission is over. This process is repeated one million times per unique set of number of starting penetrators, failure probability, retake probability, and eureka probability.

Several examples of how VaO changes as a result of more assets and even small probabilities of major success are shown in Figure 8-22 and Figure 8-23. These eight graphs show two different levels of probabilities for failure rate ([80%, 90%]), retake moments ([10%, 20%]), and eureka moments ([2.5%, 5%]). Fewer design choices were plotted to more easily display the variations between designs with different number of penetrators. It is important to note that the x-axis of these plots is *not* the utility or value, but rather the comparison between the expected utility and the actual value achieved at the end of the mission. Designs that leverage more assets will therefore have higher expected utility to begin with depending on the relative capabilities of each asset; the most important aspect of these plots is to observe how the VaO curve shifts to the right with more assets, meaning more value is capitalized on through increased opportunity.

The most obvious conclusion that can be drawn from these graphs is that designs that only carry one penetrator will never be able to capitalize on opportunity that may arise over the course of the mission. This is denoted by the blue lines on these graphs. The probability of achieving zero success is also equal to the failure rate for these designs, as expected.

One consequence of carrying multiple assets is risk mitigation. From these curves it is evident that the probability of a mission achieving no value whatsoever vanishes. Even with only four penetrators (green lines), the probability that all four penetrators fail is negligible on all eight plots. (Note: “achieving no value whatsoever” is different than “not achieving the minimum acceptable value.” These plots do not discern whether the design is valuable to the stakeholder, only the spread of possible value relative to the expected value of the design at the beginning of the mission.)

As the number of penetrators increases, the cumulative probability of the expected value occurring drops, meaning there is a higher probability that more value will be achieved. In the best scenario (Figure 8-23, right), a design with 13 penetrators has a 50% chance of exceeding its original expected value. Other less optimistic scenarios do not show as high a probability of further success; however, they do show significant risk mitigation compared to an analysis that did not consider emerging opportunities at all.

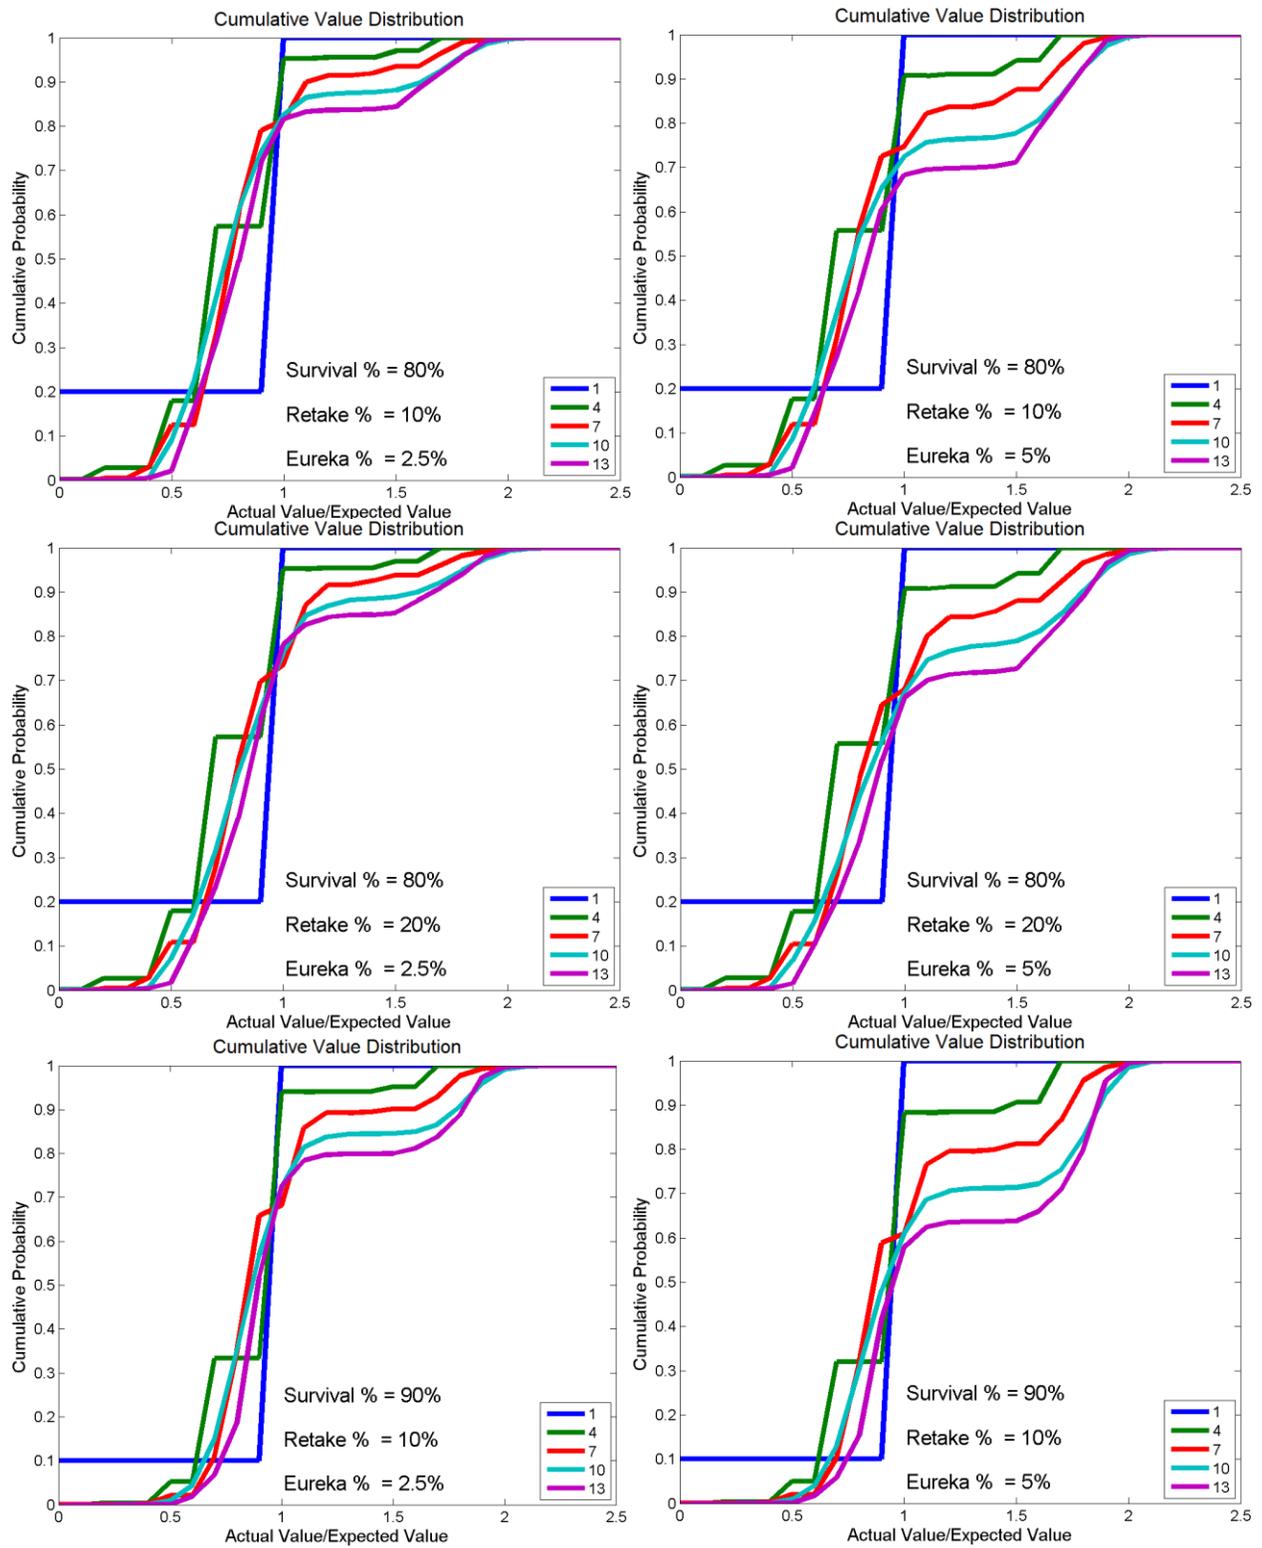


Figure 8-22: Cumulative value distribution curves for assets that leverage staged sampling to achieve opportunities that arise from endogenous epoch shifts.

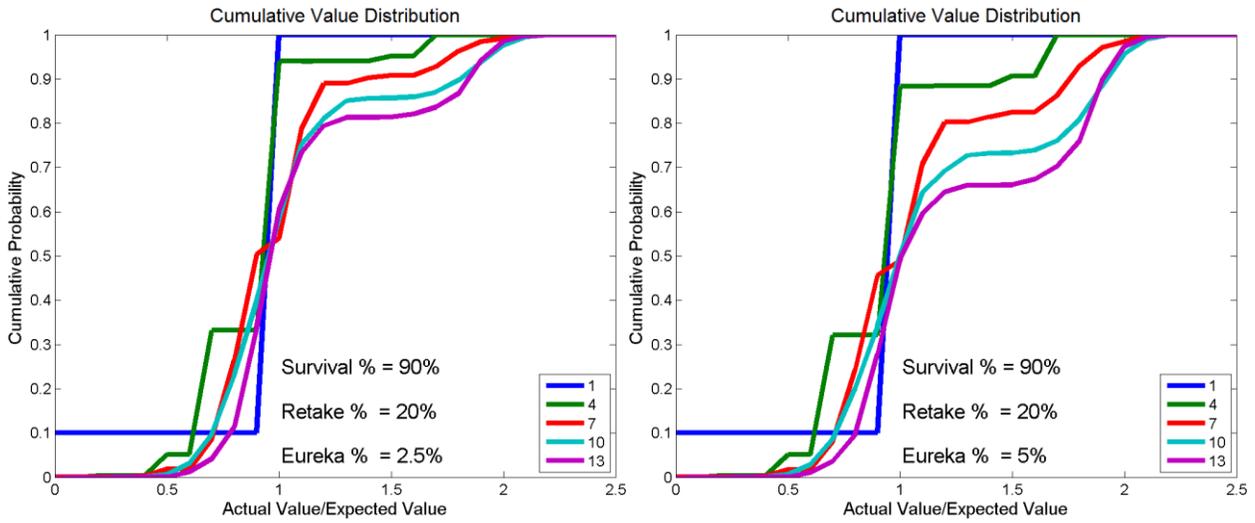


Figure 8-23: Additional cumulative value distribution comparisons.

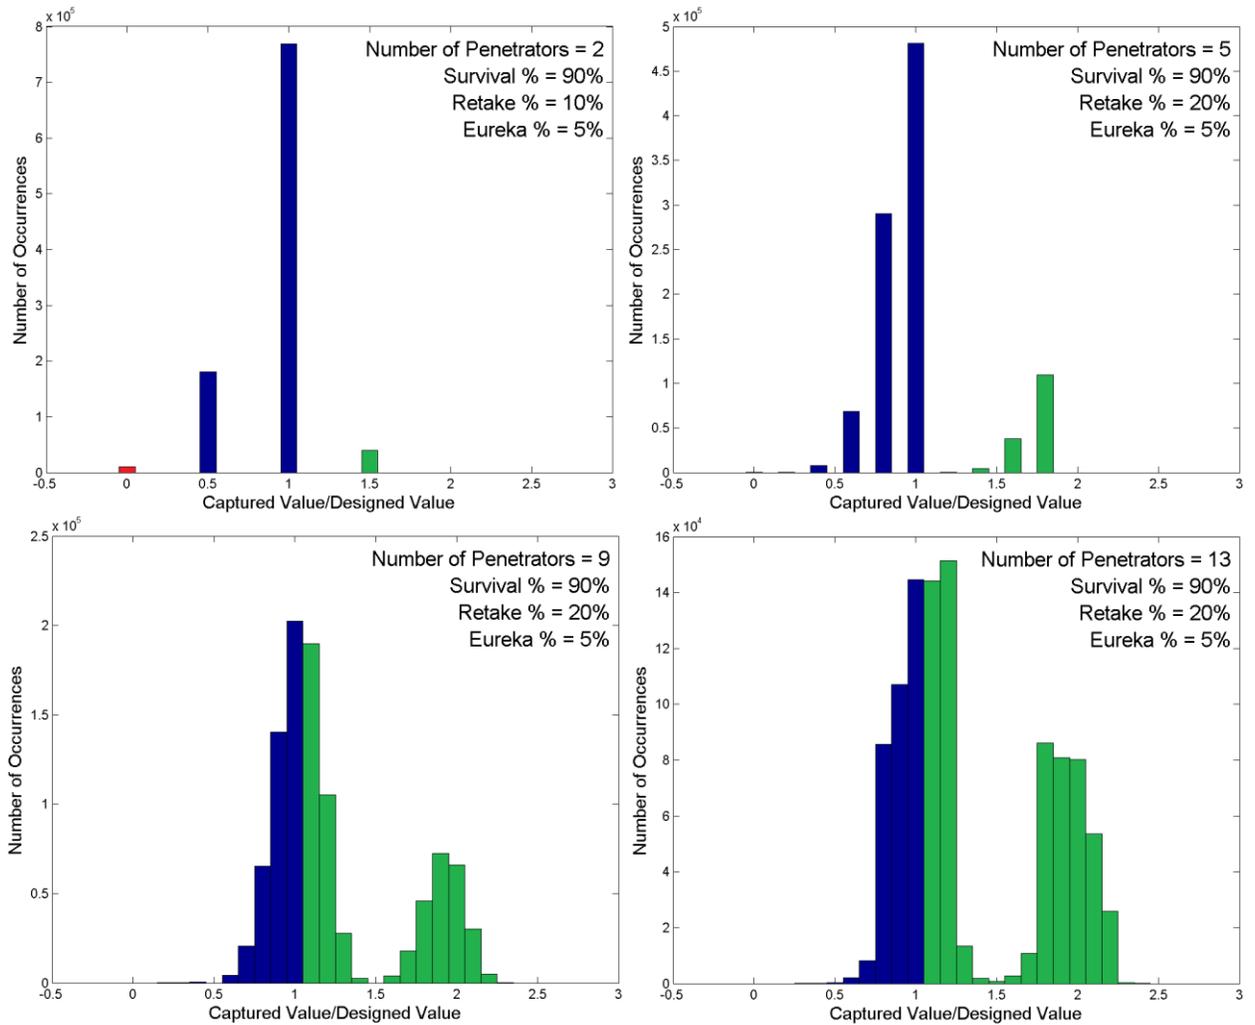


Figure 8-24: Example cumulative probability distribution functions for differing number of penetrators. Value at risk (VaR) is shown in red, value at opportunity (VaO) is shown in green.

Once again, it is important to point out that it is irresponsible to assign a definite probability that a major discovery could occur under a mission scenario. However, this technique shows that even with small probabilities (5% or less), significant impacts on the possible value of a mission could occur.

The cumulative value probability distribution of a few designs, all under the same assumptions on the probability of failure, probability of wanting a follow-up measurement, and probability of a game-changing result is shown in Figure 8-24. Each distribution is based on Monte Carlo analysis of one million eras. As the number of penetrators increases, the value distribution becomes less discrete, the average value of all possibilities increases, and at higher enough numbers, the distribution becomes two-peaked.

From these probability functions, two important metrics can be derived. The first is the ratio of the average value and the expected value for that particular design. This is the sum of all values divided by the number of Monte Carlo simulations. The second is the Value at Opportunity (VaO), representing the probability that the design will exceed expectations. These metrics are shown in Figure 8-25.

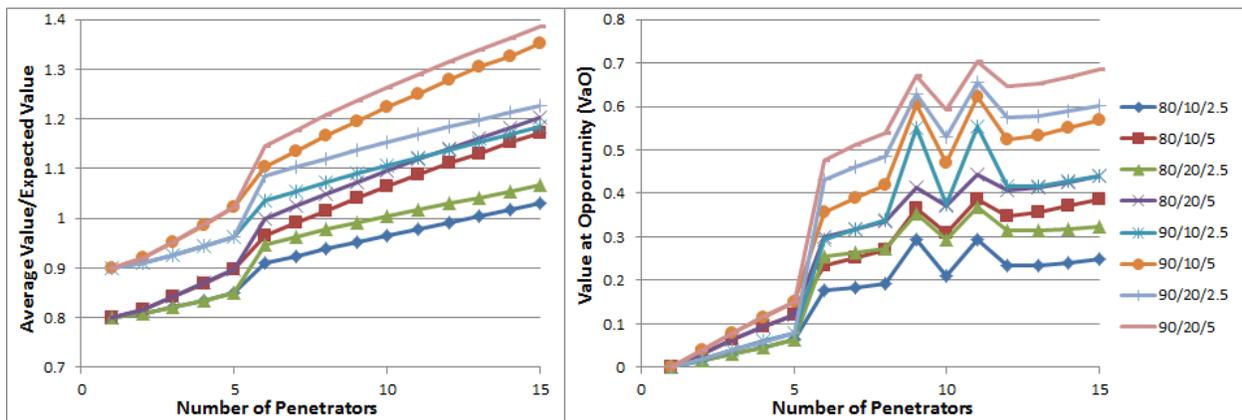


Figure 8-25: Comparison of (left) the ratio of the average and expected value and (right) value at opportunity (VaO) versus the number of penetrators.

There is a large jump in both of these metrics between five and six penetrators, because that is the threshold where the desire for a follow-up sample outweighs the relative value of conducting the mission as planned, which is a relative affine shift of 20% (as stated in Table 8-11). The relative scientific returns from each penetrator beyond the fifth one would be greater than the relative expected return for areas where follow-ups are desired, i.e. when the affine shift is greater than $1/n$, where n is the total number of penetrators, it makes more sense to attempt to capitalize on this value compared to when the mission has fewer assets.

The increase in VaO after the 5th penetrator is even more pronounced compared to the average versus expected value because even with some probability of failure, the 6th asset could be used to capture any opportunistic value that could arise during the mission.

These metrics unambiguously articulate and quantify the added value that can be achieved in a mission as a result of its ability to capture opportunity over its lifecycle. With these, a new version of the tradespace can be constructed that more fairly compares the value from the designs. A tradespace comparing the average value and the cost is shown in Figure 8-26.

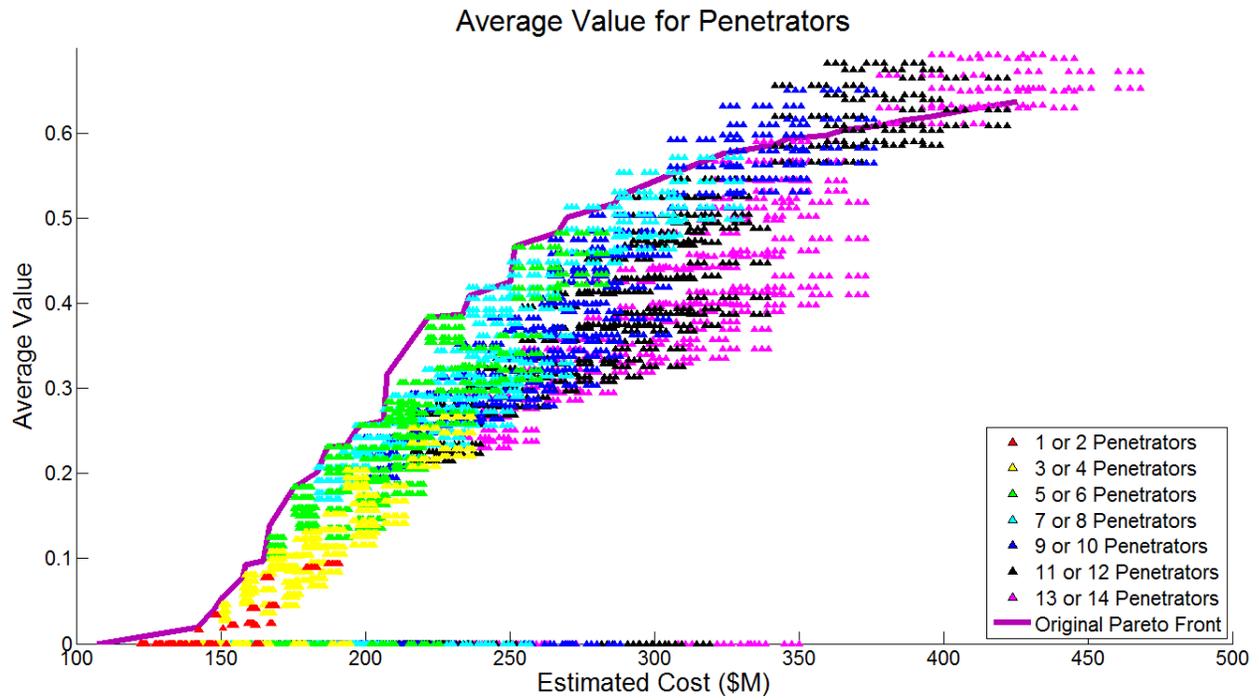


Figure 8-26: Updated tradespace incorporating shifts due to average value calculated from EEA.

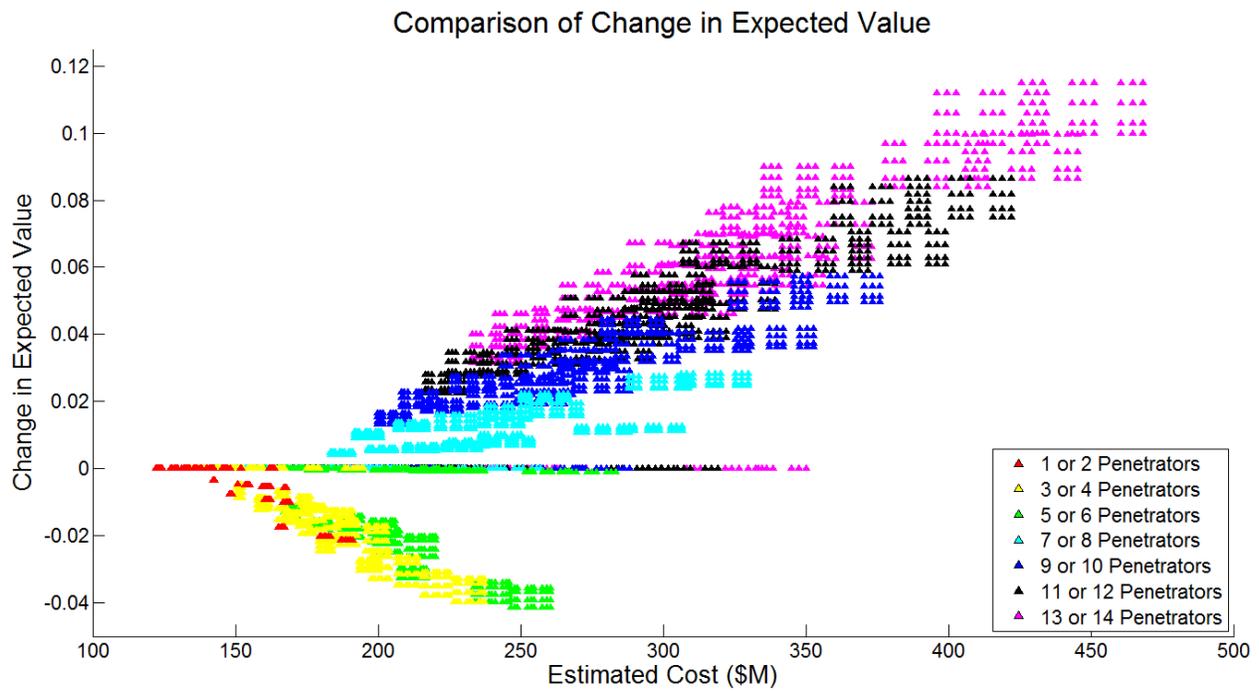


Figure 8-27: Change in the expected value for penetrators.

The tradespace in Figure 8-26 and be compared to the one in Figure 8-19. The original Pareto front has been drawn on the tradespace to make it clear how the value has changed from the results of multi-era analysis.

The changes are not readily apparent when viewing this many designs, so to illustrate this, Figure 8-27 shows the difference explicitly. From this figure, it is clear that the designs with more penetrators see a large increase in the average value due to the opportunity that arises over the course of the mission, whereas designs with fewer penetrators that cannot capitalize on the available opportunity but still face the same risk of failure have a lower average value.

8.4 Discussion and Conclusions

The results from GANGMIR show how many of the science goals identified by credible scientific groups in a wide variety of fields can be achieved using DSSs leveraging heterogeneous and expendable assets. This concept shows promise for a future mission proposal that would most likely never even be considered by traditional mission managers who do not possess the tools to understand how value is being delivered, how risk can be acceptable, and opportunity can be pursued; the RSC method and the modifications proposed in this work give mission managers the ability to weigh these options and how higher levels of satisfaction can be achieved in uncertain futures.

8.4.1 Discussion of RSC Method Results

The data discussed in the previous sections as a way to illustrate the method in action is only a small fraction of what can be learned from this case study. A discussion of additional findings, observations, and recommendations follows here.

8.4.1.1 Compromises Among Conflicting Small Satellite Design Variables

One of the interesting results from this study is how different goals may have competing priorities, and adjusting a design variable to better achieve one goal inherently decreases the overall value for achieving another goal. The RSC method provides a mathematical way to choose a design that provides the best overall compromise for the entire mission, even among other goals that are not in direct competition.

In GANGMIR, there are a number of tradeoffs and competing priorities with the sounding radiometer. It was known early in the modeling phase that there would be tension between the desire to use more wavelength bands (which requires a larger pixel size on the detector) and the desire to achieve higher resolution (which can be achieved with shorter wavelengths). This manifested itself in the struggle between number of wavelength bands, where more bands corresponds to a longer wavelength and larger pixel size, and aperture diameter, which is a major cost driver due to the scaling laws associated with optics. Increasing the aperture diameter mostly benefited one goal at high cost without benefiting other goals enough to justify the higher costs.

This case study also showed that because a long mission lifetime is desired, low orbital altitudes that will cause satellites to deorbit faster are far less desirable than higher orbits, which can sustain longer lifetimes, even if higher orbits result in lower resolution from the instruments on the satellites. This is supported by the two-swath distribution visible in Figure 8-19, Figure 8-20, and Figure 8-21 (altitude was not explicitly plotted in Figure 8-15 or Figure 8-16). All of the designs in the lower swath are at lower altitudes that decay too quickly to achieve great value.

In general, attempting to satisfy all the science goals will be expensive, but this GANGMIR study showed how trying to compromise in some areas can lead to overall decreases in utility instead of overall increases. The extent to which compromises would lower utility depends on the weighting among all of the science goals; goals whose weights $\sum k_i > 1$ are more likely to lose utility overall because of substitution effects, but when $\sum k_i < 1$, the synergy of accomplishing two goals even worse than either would have been accomplished before the compromise would result in an overall increase in utility.

A designer thinking in a monolithic systems paradigm may not understand why the satellites that carried cameras with the highest resolution were never on the Pareto front, but it is clear from the guidance in the MEPAG goals document that full global coverage is far more important for these science goals. Another reason why higher resolution cameras would most certainly be less valuable in comparison is that the FOV is much more limited, meaning it would take many more satellites to achieve the desired temporal resolution. GANGMIR is clearly a case where the temporal resolution a DSS can achieve outweighs the resolution a monolithic design utilizing a single high-performance camera can achieve.

8.4.1.2 Substitution Effects in Instrument Attributes

The RSC method also uncovered design variable tensions between instruments on different assets. It was known from the beginning that there would be some tradeoffs between maximum wavelength and aperture diameter due to the optics equations; what was *not* obvious until the TSE phase was that penetrator designs that carry the largest weather package are *never* Pareto-efficient.

This conclusion is because the particular goal that is well-satisfied by this package is also satisfied by the small satellites, and substitution effects diminish the role that penetrators would have if they are being deployed to achieve that science goal. Additionally, because the scope of some of these goals are global, having just a few point measurements is not nearly as valuable compared to a fleet of satellites that can achieve global coverage with low revisit time and make measurements that are almost as good as in-situ measurements.

Without first developing a value model, the substitution effects between these two instruments on entirely different payloads would not have been well-articulated or quantified. Because the stakeholder might perceive redundancy between these two payloads, the added cost of adding that payload to the penetrator moves the design further from the Pareto front. Had the stakeholder perceived synergy between these two payloads, the conclusion may be reversed, but this is not the case with this value model. Alternatively, one could argue that had more options for combinations of payloads on the penetrators been available, there might be some Pareto-efficient designs that do carry this payload, but given that the tradespace is already reaching the computational limits available for this study and the fact that the payload addressed a less-valuable goal than the others, it did not make sense to adjust the design vector at this stage.

Although nothing can be concluded without further value elicitation from the actual stakeholders who have an interest in understanding Mars' weather, these results show that the METNET concept for building a network of ground penetrators may be far less valuable than building a constellation of satellites to conduct the same weather observations. However, if the value model developed in this case study is incorrect and fewer in-situ measurements are much more valuable than many remote sounding measurements from orbit, METNET may still be a valuable option to achieve those goals.

8.4.1.3 Quantifying the Added Value of Stacked Sampling with Heterogeneous Assets

Just like in the HOBOCOP case study, the results from the RSC method show that the highest stakeholder satisfaction is achieved with a combination of heterogeneous assets rather than with only one type of asset. Figure 8-20 and Figure 8-21 show the Pareto fronts for homogeneous assets along with points representing heterogeneous designs. Although designs that are exclusively made of satellites dominate the lowest part of the overall Pareto front, the synergy between goals, stacked sampling, and shared subsystems of heterogeneous combinations of satellites and penetrators quickly overcomes the maximum possible utility that can be achieved with designs that use satellites only.

A mission that only considered one type asset and not another was not able to achieve full satisfaction because of the breadth of goals that are attempting to be achieved in this mission. Had only weather goals or only ground sampling goals been chosen, this may not be the case, but this mission attempts to satisfy many groups of stakeholders and does arguably more than any previous Mars mission.

Such a conclusion would have been obvious even without the RSC method; after all, there are many space science fields that focus on different aspects of Mars, including atmospheric science, geology, biology, and planetary formation, in addition to engineering fields related to future human exploration. However, the RSC method provides a framework to express the added benefit when assets operating in different environments can provide data that present a global picture of a phenomenon that cannot be captured by a single observation or set of observations from a single point.

GANGMIR's leveraging of stacked sampling is not the same as how it is achieved in HOBOCOP; while HOBOCOP's stacked sampling comes from the top level synergies among goals and heterogeneous assets working together to raise the bottom level attribute levels, GANGMIR's stacked sampling adds value within the goal-level attributes of the proposed MAUT hierarchy as instruments work synergistically to deliver more value together than they would separately.

8.4.1.4 Quantifying the Added Value of Staged Sampling with Expendable Assets

The obvious conclusion from multi-era analysis on this mission using deployable assets is that when there is a possibility that a discovery that could occur that would cause the stakeholders to desire a change in strategy, having more assets is better. Being able to pursue opportunity during the course of the mission rather than waiting a decade for another mission to address the potential findings of a mission is unprecedented in planetary science missions. This makes eliciting and characterizing what the added value could be difficult, because the PSSs may not have even thought about how they would feel about the ability to conduct follow-up measurements.

This is why good value elicitation requires the interviewer to have some understanding of the scientific field of study. With understanding of the current state of affairs and the latest findings, the interviewer can probe the PSS to elicit their opinions on scenarios they could consider changing their mind on and how they would perceive the added value of those change options. Multi-era analysis can then be used to show skeptical, risk-averse stakeholders that not only can risk be drastically reduced, but opportunity is much more achievable.

The differences between designs in the tradespace shown in Figure 8-26 are not as stark as they might be in other missions. This is because the utility of designs is more highly dependent on the number of penetrators than other factors. If there were more designs where a small number of penetrators was equal

or nearly equal in cost and expected design utility to designs with a large number of penetrators, the shift when examining the average value calculated through EEA would be much more apparent. However, the cost and utility increase rather proportionally along the Pareto front with increasing number of penetrators, so such knowledge may not be evident from these figures alone.

8.4.2 Recommended Design Solutions

The results from MATE alone provide adequate information for a designer to make informed decisions about which designs provide the most value across all science goals. It is easy to see at which price points a mission manager would choose to change instrument design variables (e.g. increasing aperture diameter) instead of adding more assets. The MATE tradespace shows that the Pareto front of this mission consists primarily of designs that use a 2:1 satellite to penetrator ratio, even though the Pareto front of designs that strictly have a 1:1 ratio of assets is just barely less Pareto optimal. This is primarily due to the fact that there are fewer redundant subsystems involved when three assets are deployed together instead of two.

However, the results from EEA show that more opportunity can be captured with more penetrator assets. For this reason, it is better to pick a design with a 1:1 ratio of assets rather than a 2:1 ratio. This choice also mitigates the risk in mission utility and value due to penetrator failures; with a 2:1 ratio, the relative risk to the mission value for each penetrator is higher. Even if no opportunity presented itself upon which to be capitalized during the mission lifecycle, a greater number of site samples and data points will provide more confidence in the non-results of the mission. Although non-results do not seem as valuable in the eyes of the public, they are still scientifically valuable.

Analysis of the effects that design variables have on total mission MAU show that the substitution effects in Goal #6 (observe dust) between the satellite's radiometer and penetrator's weather package have a strong influence on how a designer would choose among several variable levels. For instance, if the maximum number of radiometer bands is chosen to accomplish the dust observation dust, it is never Pareto optimal to include dust cameras on the ground penetrators because of the added mass on each asset. If dust cameras are included on the ground penetrators, it is never Pareto optimal to use the maximum number of bands because of the decreased resolution for observing the temperature profile of the atmosphere (Goal #1, a high priority), and high utility can be achieved with a smaller aperture diameter (and thus a lower mass and cost).

This tradespace does not consider heterogeneous penetrator assets. A future study may include three different penetrator options: weather packages alone, spectrometer packages alone, and both. To meet the maximum perceivable levels of the weather goals, more ground assets are needed compared to the maximum perceivable levels of the subsurface goals. This trade depends heavily on development costs, since the assets would have separate development needs and mass production characteristics that could actually add costs, but it is an important consideration.

8.4.3 Comparison to Monolithic Missions

The most successful Mars missions (in the eyes of the public) have been rovers. Unlike penetrators, rovers are not "sacrificed" in the sense that they can only take one sample and be done. If a rover finds an interesting result, the rover can take another sample from the same or a nearby location; under the alternate epoch value model discussed in this chapter, the rover would be able to capitalize on the

opportunity that would present itself. However, because the number of rovers would be limited to one in a monolithic mission, and because rovers have not been capable of traveling across vast distances to explore entirely different regions from the one they started their mission in, the chances of finding opportunity upon which to capitalize is low compared to many sacrificial ground penetrators.

On Mars, no rover has dug into the regolith as deep as a penetrator would be able to. Although a rover mission with a large enough drill could be launched (such as the drill proposed for the Mars 2020 mission [519]), this equipment would be heavy and expensive. Leveraging the kinetic energy from EDL to sample beneath the ground not only reduces equipment mass, it also reduces EDL subsystems mass.

A rover with a 99% chance of survival operating under the same epoch model would have a VaO of 23.7% and an average lifecycle value 1.079 times its expected value without using EEA. These are much smaller in comparison to the results from EEA for penetrators (see Figure 8-24 for comparisons). The value probability distribution for this analysis is shown in Figure 8-28.

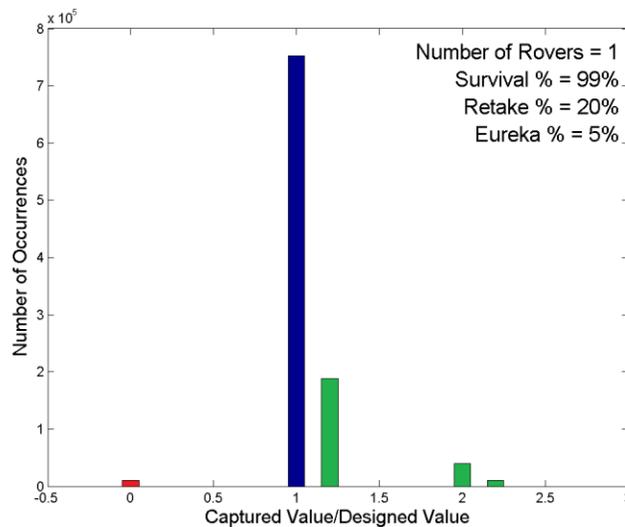


Figure 8-28: Value probability distribution for a rover. VaR is shown in red (1%), VaO is in green (23.7%)

Given the added cost of ensuring such a low probability of failure, as well as the added mass from EDL subsystems, the results of this study show that several penetrators are still more valuable than a single rover even if a single rover were capable of meeting the minimum satisfaction criteria for the science goals described in GANGMIR (and it is not).

8.4.4 Future Work

Collaborative efforts on GANGMIR ended in November 2014, so this mission may have no future beyond the design study presented herein and other presentations at conferences by the Italian groups.

However, this case opened up questions about how scientists value the ability to conduct follow-up observations after interesting results. Research into these perceptions is especially useful for methods like RSC to show how value is really gained in missions, even if stakeholders don't think there is value in such follow-ups initially.

Additionally, this case could continue to be updated with value models elicited directly from members of the teams that formally published the science goals rather than derived from assumptions based on those

publications. This mission could easily serve as a template for future Mars missions where a number of high-priority goals are being achieved regardless of the architectural options that are being considered.

8.5 Chapter Summary

A mission to the most visited planetary science target to achieve many of the highest priority goals was developed, using distributed and heterogeneous assets in orbit and on the ground.

The value model was built through inputs from multiple stakeholders, and assumptions were made based on scientific literature when no experts were available for direct elicitation of preferences. This value model took into account how measurements from multiple sources could stack so that more value would be perceived, and how substitution effects would diminish the value when similar data from multiple sources was received. This model also hypothesized what a principal investigator may perceive as the value gained from interesting results that were observed during the course of the mission and how their expectations would change if they had the ability to conduct follow-up observations.

A performance model that calculated the performance of many instruments was built to estimate the cost and utility of over ten million possible designs. This model simulated EDL, ground penetration, orbital insertion from a Hohmann transfer orbit, and orbital precession to deploy a Walker constellation.

Tradespace exploration showed how different design variable options impacted the total utility of the designs. Important tradeoffs were identified, especially between heterogeneous assets. The tradespace showed that the initial assumptions on the value of data from ground sampling are less valuable compared to observations from orbit. A skeptic could argue the value model is incorrect, but given the prioritization of goals from established science groups, it is also possible that those goals are not cost efficient to achieve.

Finally, epoch-era analysis showed how VaO changes drastically with an increased number of assets. This change and the possibility to achieve more than expected as information that comes naturally over the course of the mission could change stakeholder expectations. The added value of a change option was estimated based on conversations with scientists, but even with rough estimates, the added value is enough for even the most risk-averse scientists to weigh the possibility of a DSS mission.

9. Discussion and Conclusions

*“As the flames of warfare rage higher, we feel our destiny's embrace
We are ageing men of an old empire, now we can see Valhalla's gates
We! March again! First in line! To reach Valhalla's mighty royal gates!
We! March again! Give our lives! To reach Valhalla's mighty royal gates!”*

AMON AMARTH, “WARRIORS OF THE NORTH”

Four case studies have been presented as empirical evidence to show how the emergent capabilities of distributed satellite systems can be leveraged to achieve valuable scientific goals that could not have otherwise been achieved. This chapter will tie those case studies to the research questions proposed in Chapter 1 and synthesize the major takeaways of this work.

Other final areas of discussion include: a number of important observations regarding the development of space science missions that leverage the emergent capabilities of distributed satellite systems (DSS); some of the common themes among all four case studies and some of the problems a designer may face; specific observations relevant to both the systems engineering and space science communities; the nature of science goals and how different kinds of missions will have different value model structures in their mission proposals; open research questions that were either discovered during the course of this work or were deferred to limit the scope of this dissertation; and future work that could improve the value propositions as they have been presented by continually improving the methods and techniques for value modeling and tradespace exploration used herein.

9.1 Research Questions Revisited

Two research questions were posed in Chapter 1: one concerning the emergent capabilities of distributed satellite systems, and one concerning the metrics and analytical techniques that can be used to evaluate complex space science missions that leverage those emergent capabilities.

9.1.1 Research Question #1: Emergent Capabilities of DSS

The first research question is:

What emergent capabilities can be demonstrated to achieve science objectives with distributed satellite systems that cannot be achieved with monolithic satellite systems?

Chapter 3 answered this question by revisiting the potential benefits of DSS as described in published literature, and by examining how combinations of those benefits result in unique ways satellites can be used to achieve science goals. As a validation to this approach, numerous scientific investigations that accomplish science goals by using more than a single asset were examined. The examined investigations were not limited to only space systems. No technique or investigation that was examined fell outside of the defined set of emergent capabilities; although it is impossible to declare that the list of capabilities defined in this work is “complete,” it does encompass a wide breadth of distributed systems in science research applications.

The emergent capabilities were identified and defined by a critical examination and synthesis of the benefits of DSS and other scientific investigations that use multiple assets. Three distinct categories of emergent capabilities and seven individual emergent capabilities of distributed systems were identified:

Fundamentally unique emergent capabilities refer to operations that a monolithic system, *no matter how capable, optimized, or large it is*, cannot perform.

Shared Sampling: When multiple assets trade responsibilities for making the same measurement at different times, particularly when it is impossible for any single asset to make the measurement over the time period required for satisfaction.

Simultaneous Sampling: When multiple assets conduct a measurement of a common target at the same time from different locations such that the combination of the resultant data sets provides more detailed information that could not have been gathered by a single asset moving between locations and taking measurements at different times.

Self-Sampling: When multiple assets measure signals generated by each other, or the precise position and velocity of each other, to infer information about a target or phenomenon indirectly rather than measuring the phenomenon directly.

Analytically unique emergent capabilities refer to operations that an incredibly capable monolithic system could theoretically perform, but such systems are unrealistic or expensive.

Census Sampling: When multiple assets conduct multiple measurements of a subset (or the full set) of a collection of similar targets in order to have greater certainty in the variance of the desired properties or characteristics of the whole target population.

Stacked Sampling: When heterogeneous assets are deployed into different environments or locations to make measurements of the same phenomenon using different instruments such that the combination of data sets is more valuable than any individual data set.

Operationally unique emergent capabilities refer to operations that would be irrational or irresponsible for a monolithic system to perform.

Staged Sampling: When additional assets are deployed after knowledge gained by the first asset's or wave of assets' data has been received and analyzed such that the location or orbit of deployment can be re-evaluated and altered to seize opportunity that was previously unknown or provide similar measurements for more confidence in results.

Sacrifice Sampling: When assets are knowingly deployed to an unsafe operating environment, such that most if not all of the deployed assets will be destroyed, but the scientific returns can justify the development and launch of the entire mission.

All of these emergent capabilities create unique ways to sample data from observations of natural phenomena to achieve science goals that *cannot* or *would more likely not* be achieved using a monolithic system. By leveraging one or more of these capabilities, a space science mission may be able to capture scientific value that has been overlooked in past mission selection processes and answer fundamental questions that can lead to further breakthroughs in understanding and discovery.

9.1.2 Research Question #2: Analytical Techniques and Metrics

An investigator wishing to design a mission leveraging any of the emergent capabilities defined with the first research question must know how to analyze and rank alternative design solutions to better understand how the system delivers value and to maximize the potential gains from the mission. The second research question is:

Which analytical techniques and metrics should be applied or modified in the development of distributed satellite systems for science missions to more accurately characterize the overall benefit over the entire mission lifecycle?

A natural starting point to answering this question was an examination of other methods and frameworks that have attempted to analyze and evaluate distributed systems. Previous studies spent much effort making direct comparisons to monolithic missions, especially in situations where monolithic missions were perhaps more appropriate, using methodologies clearly designed with monolithic systems in mind. More recent work developed specifically to evaluate changeability over the system lifecycle still suffered from some potential weaknesses and did not explicitly address synergy among multiple science objectives. As a result, the true value of system lifecycle properties in DSS designs was not being articulated.

The discussion in Chapter 4 compared alternative methods and techniques and their weaknesses. This discussion showed how those weaknesses could be addressed or otherwise overcome, under assumptions that would be appropriate to make in space science missions but not in the general case for the application of those methods.

The modifications that were made to the Responsive Systems Comparison (RSC) method not only help set up the value model better for the specific case of DSS in space science missions, they also help better articulate the benefits that can be gained through system lifecycle properties across a range of applications. These modifications reduce the broad applicability of the method compared to the general case of RSC, but most system lifecycles are rigid enough to apply the assumptions that make these value metrics applicable and truthful.

The specific modifications to the RSC method developed and tested in this work are:

- *Phase 1: Value Modeling*
 - Identify additional stakeholders whose goals may have synergy with the PSS's goals.
 - Construct MAU hierarchy of goals, instruments, and instrument attributes.
 - Identify synergies between and among individual science goals.
 - Identify value-driving endogenous epoch shifts.
- *Phase 2: Performance Modeling*
 - Construct “coverage” models to match stakeholder value model accurately.
 - Apply cost improvement curves to the development of multiple assets.
- *Phase 3: Tradespace Exploration*
 - Explore combinations of heterogeneous concepts and assets.
- *Phase 4: Epoch-Era Analysis*
 - Use cumulative value to characterize value delivery over time.
 - Use Value at Opportunity (VaO) to show the potential for greater than expected value.

9.1.3 Combining Research Questions #1 and #2

While each of the research questions can be viewed separately and distinctly, answering both of them together created research synergies that help fill in major research gaps. The new methods and findings allow both distributed and responsive systems to be compared objectively among each other in addition to being compared objectively to monolithic or inflexible systems. In contrast, previous research undervalued responsive and distributed systems in direct comparisons to monolithic and inflexible systems. The conclusions from examining missions strictly suited for DSS also provided ways to reexamine previous research that focused on comparisons between monolithic and DSS designs. This section identifies research gaps that have been filled as a consequence of answering both research questions together.

9.1.3.1 Quantifying Certain Benefits for Space Science Missions Leveraging DSS

The real bridge between the two research questions is that the benefits that can be gained by leveraging the emergent capabilities of distributed satellites can be quantified and ranked against alternatives. Combining the emergent capabilities of DSS (#1) with the value modeling techniques of the modified RSC method (#2) allow mission designers to address space science goals that cannot be achieved with monolithic systems, no matter how capable such monolithic systems are. By developing a value model that can objectively compare how well vastly different missions satisfy science goals (and is not necessarily limited strictly to DSS), many concepts can be compared simultaneously before the design space is reduced in later development stages by requirements definitions processes.

All four case studies showed how science goals could be achieved by using DSSs that either could not be achieved well enough to meet the minimum satisfaction criteria by a monolithic system or could not match the value of DSS designs with a similar cost. The value modeling methods as they relate to the perception of value from the scientific stakeholders allows for objective comparisons among very different DSSs when the design space is rich. This is especially important when there are a number of measurement techniques that can be used to acquire the same data or achieve the same science goal.

More importantly, the science goals these missions addressed have either never been addressed in previous space missions or have been addressed in a way that fundamentally enhances our understanding of the science goal and has broader impacts on the scientific community. Not only can how well these goals are satisfied according to one person be quantified, the broader consequences to the scientific community can also be considered in the problem formulation to quantify the benefit across multiple scientific communities rather than one small section that would benefit the most.

9.1.3.2 Quantifying Uncertain Benefits for Responsive Systems

Another important connection between the two research questions is that in order to objectively compare systems over their lifecycles, the uncertain benefits must also be quantified. Uncertain benefits are anything that adds value to a mission (e.g. more instrumentation, extra propellant, changeable options) in ways that are not necessarily guaranteed or known at the beginning of the design process; as a result, they are typically not considered as part of the value proposition in the concept proposal phase.

Previous tradespace exploration research compared designs within the same context but did not have a method to objectively and explicitly model the added value that can be obtained through uncertain benefits over the mission lifecycle. By answering the two research questions together, a framework was

developed to quantify and objectively compare the uncertain benefits of designs in order to select the design that is truly the most suitable.

The uncertain benefits that are relevant when comparing DSS designs come from possible opportunities. The most important aspect of uncertain benefits in DSS designs is the *operational opportunity*, which is a change in the mission operations that benefits the mission, usually after launch but not exclusively, because the system was designed to be changeable so that it could pursue additional benefits. This can occur with or without any extra added cost.

The potential benefits of DSS (as discussed by Jilla [20] and listed in Table 3-1) do not always add utility or value (y-axis of the tradespace) to the mission *in the primary design context*; therefore, they have been seen as unattractive design choices in previous design studies even though they *do* add value over the lifecycle by mitigating failure or responding to operational opportunity. This is why monolithic and unchangeable systems have typically been chosen over distributed or responsive systems, because the added cost of changeability does not directly result in a change in a design's value proposition if alternative contexts are not considered in the value model.

The methods proposed to answer the second research question provide a means of numerically articulating this previously unarticulated uncertain value that arises as a result of a design's inherent changeability. By quantifying the certain and uncertain benefits of each design, the true value proposition of designs over their entire lifecycle can be expressed, rather than the value proposition of designs and the expectations and context they operate in only on launch day.

Furthermore, the modified RSC method of quantifying uncertain benefits is not only applicable to space science missions that leverage the emergent capabilities of DSS, but also to other quasi-responsive systems whose value can be quantified over the system lifecycle under similar assumptions (that have been discussed in Chapter 4). While the techniques developed to rank the lifecycle value of alternative designs are not strictly applicable when epoch shifts are reversible or when user satisfaction is continually experienced rather than delivered over time, other complex systems that operate in dynamic environments where operational opportunity can be captured can still be evaluated using this framework.

9.1.3.3 Comparisons to Current Methods for Evaluating Space Science Missions

By answering the two research questions together, a method was developed to assess space science missions that leverage DSSs. As a result, the modified RSC method also addresses and improves upon many of the weaknesses of previous methods. To illustrate and compare the strengths of the new method to previous methods, Table 9-1 tabulates the properties of typical space science missions and shows which methods and value-centric design methodologies (VCDMs) are appropriate for which properties.

It is clear from this table that the modified RSC method presented in this work addresses most of the identifiable weaknesses of previous methods in space science missions. Although the modifications made during the course of this research are new, the research leverages and builds upon much of the previous 15 years of research at SEArI and other institutes. The modifications have firmly cemented the various pieces of research together to strengthen those previous weaknesses, especially when examining lifecycle value.

Table 9-1: Summary of the properties of typical space science missions and how well some VCDMs can be used to assess their value. (This table is a combination of Table 4-1 and Table 4-3).

Properties of Typical Space Science Missions	NPV	CBA	TOPSIS	MAUT	Original RSC	Modified RSC
Value can be derived from multiple sources	X	✓	✓	✓	✓	✓
Value is NOT discounted cash flow	X	X	✓	✓	✓	✓
Mutual Utility Independence	X	X	X	✓	✓	✓
Stakeholder decision is under uncertainty	X	X	✓	✓	✓	✓
Value is an ordered, non-ratio comparison	X	X	✓	✓	✓	✓
Context may change over time	X	X	X	X	✓	✓
Applicable with multiple stakeholders	✓	✓	X	X	O	O
Value is delivered over time	✓	✓	X	X	X	✓
Value can be delivered synergistically	X	X	O	O	O	✓

When Prof. Daniel Hastings stepped down as Chief Scientist of the Air Force in 1999, he knew there needed to be a better way to handle making decisions under uncertainty, which is what started much of the research that has built up to this work. It is now clear that the time has come for over a decade of research to be brought out of the academic world and into the design world. The methods developed in this work are a great improvement not only over what the Air Force was using in the 1990's but what NASA is using today and should be applied to all future mission designs, not just the ones that leverage the emergent capabilities of DSS to achieve their goals.

9.1.4 Linking Emergent Capabilities to Specific Analytical Techniques

For engineers and scientists attempting to leverage the emergent capabilities of DSS in their own concept exploration programs, it important to know which aspects of the modified RSC method are especially important for accurately evaluating and comparing alternative designs. The findings described in this section are the result of applying the second research question to the answer to the first question to see which techniques are applicable to the specific emergent capability being leveraged.

When leveraging *shared sampling*, the value of DSS may be self-evident, especially when a single asset cannot be used to provide minimally satisfactory value, but it is also evident when attributes such as coverage are of high value to the stakeholders compared to spatial resolution or other attributes that normally have better levels in a monolithic system. Understanding how a stakeholder perceives the benefits of shared sampling comes from careful value elicitation. Additionally, added value through risk mitigation or opportunistic changeability can be shown through EEA.

When leveraging *simultaneous sampling*, the value of DSS may also be self-evident, but careful elicitation and construction of coverage attributes are also important to explicitly reveal the added value of additional assets. Identifying and constructing these attributes may require some ingenuity and familiarity with the scientific literature, but the value model elicitation will primarily show that single assets are unacceptable, by definition, in these cases. Additionally, added value through risk mitigation or opportunistic changeability can be shown through EEA.

When leveraging *self-sampling*, the value of DSS should be evident mostly by comparisons to similar monolithic systems that can achieve the same goal but at far greater costs. The data products from these systems may seem complicated to obtain compared to the products from monolithic systems, but the

monolithic method of obtaining a data product that can be obtained through self-sampling is usually completely infeasible or prohibitively expensive. Performance modeling and TSE can show these alternative concepts on the same tradespace compared using the same value model.

When leveraging *census sampling*, the value of DSS should be evident when different statistical metrics are applied to show the value of multiple measurements. In this type of mission, there is a tension between the breadth and the depth of the data; more capable missions cannot sample as many targets. Furthermore, time plays a crucial role for gathering the data, and if completion time is important, the value of multiple assets will be self-evident as monolithic assets fail to meet the minimum satisfaction levels for statistical confidence in that time period (or if it is infeasible for a single asset to sample more than one target to begin with). Additionally, if the stakeholder's perceptions of satisfaction can change as a result of data gained through the mission, EEA can show how the effects can change over the course of the system lifecycle so added value through opportunistic changeability can be shown.

When leveraging *stacked sampling*, the value of DSS should be evident through the use of the MAUT hierarchy tree that specifically aggregates synergistic sets of data. This comes from proper value model elicitation, where the interviewer specifically asks how the stakeholder perceives the value of sets of data in combination. Whether these sets are redundant or synergistic depends on stakeholder perceptions.

When leveraging *staged sampling*, the value of DSS should be evident through Epoch-Era Analysis using Value at Opportunity to quantify the uncertain benefits that may be captured by changing course during the mission to capture operational opportunity. It only makes sense to pursue operational opportunity if the added value of that opportunity is more than the initial value of the deployed asset, so this capability is most valuable when the cost per asset is low or the potential value is high.

When leveraging *sacrifice sampling*, the value of DSS should be self-evident by the much lower cost of the assets that are being sacrificed relative to the cost of a system that would survive. It can also become evident in the analysis of the environment those assets are being sent to, but it is important to understand the costs of designing for survivability in those environments versus the cost of intentionally destroying assets relative to the scientific value, especially when this capability is leveraged in tandem with another emergent capability.

9.1.5 Summary of Case-Specific Insights

The four case studies in Chapter 5 through 8 expounded on methodology and successfully showed how the emergent capabilities defined in Chapter 3 could be used to answer currently unanswered scientific questions or scientific questions that had not been thoroughly answered, despite over half a century of space exploration under a monolithic systems paradigm (although other factors such as science politics and limited budgets are key reasons as well). The proposed methods showed how a scientist or engineer could analyze a system that leverages any of the seven emergent capabilities, and why those techniques add potential value to scientific progress. Some of the specific findings and insights from the four case studies are summarized here so that scientists and engineers can see how a mission concept selection process may be different under a DSS paradigm compared to previous selection processes.

9.1.5.1 Case Study #1: ExoplanetSat

The primary purpose of the ExoplanetSat mission is to detect transiting, Earth-like exoplanets around nearby, Sun-sized stars. This is accomplished by using assets that leverage *shared sampling* to conduct

observations over periods of time greater than one year that a single satellite could not conduct alone. The assets take turns observing targets as their orbits precess, so that by the time one asset can no longer observe one target, another asset is in a position to continue the observation. With many assets working together, sharing observation responsibilities, a larger area of the sky can be observed.

The results from this case study showed a stark contrast between this implementation of the RSC method and the methods employed by NASA. After a two-day concurrent engineering session in a world-class design studio, JPL's Team X generated two solutions that fit within the budget, were technically feasible, and fit the description of the mission profile. However, they did not fundamentally understand how value was generated in this mission, nor did they explore a wealth of orbital options. As a result, the system they proposed would not have achieved the science goal. If they had derived a more complete set of requirements through better value elicitation, it is possible they would have been able to deliver a satisfactory mission, but it may not have been optimal for a given number of assets.

Not only did the RSC method help explicitly define the stakeholder's perception of value using a constructed value model, but it also showed how to deliver the most value given a limited number of satellites. Affordability is a major concern, and the number of available assets is dependent on the maximum budget. The most important takeaway of this case study, from the stakeholder's perspective, was to understand the maximum capability of a mission, given a certain number of satellites.

When designing a mission that leverages shared sampling, it is important to understand how failure modes impact the value delivery. If value comes from having a completely unbroken chain of measurements, redundancy or reconfigurability in the system is critical to sustaining value in the event of a disturbance that causes an asset to fail. If value comes from multiple repeated measurements, the loss of a single asset may not be as detrimental, but a repeating gap in the measurements may introduce some bias in the data products that could be more detrimental to the stakeholder than a designer may initially assume. This is why understanding what drives value is important when defining the attributes for coverage.

An example of how value delivery from shared sampling can be easily mischaracterized is ground coverage for a constellation of satellites. Suppose a coverage attribute for a satellite constellation was "revisit time." One design has a "revisit time" level of 30 minutes; another has a "revisit time" level of 45 minutes. If a satellite in the first constellation fails, the "average revisit time" may only increase to 33 minutes, but the "maximum revisit time" may now be one hour. If "average revisit time" is the attribute on which a stakeholder is mentally judging the constellation, the first constellation would still be acceptable. However, if "maximum revisit time" was how the stakeholder perceived value, the second constellation would now be perceived as better. Good value elicitation will help designers understand how value is derived from the attributes to reduce ambiguity in their interpretation, and structured value modeling combined with era analysis will help designers properly express these risks to stakeholders.

9.1.5.2 Case Study #2: HOBOCOP

The primary purpose of the HOBOCOP mission is to study the solar magnetosphere and other phenomena that are influenced by it. This is accomplished by using assets that leverage *simultaneous sampling* to conduct in-situ measurements of an extended field that give greater spatial and time resolution than any single asset could. A single satellite working alone would be incapable of providing any new information

that could not already be gained from existing missions, so no deeper scientific understanding could be achieved without multiple satellites.

Not only did using the RSC method show how different technology paradigms affected the design options and stakeholder opinions on the age of data, but it also showed how the heterogeneity of assets was important for preventing single assets from ballooning in size and cost. Only the in-situ payloads needed to be placed close to the Sun; the heavier observational payload would have required much more mass to shield itself from the increased heat flux, which would in turn require that the propulsion system be much more massive as well. Meanwhile, the different systems still required similar propulsion capabilities and were able to share those commonalities.

When designing a mission that leverages simultaneous sampling, it is important to understand how simultaneous measurements are generating value. If remote observations of an object are driving the value of the mission, then the number of assets required is generally low. The STEREO mission leverages two spacecraft along with terrestrial assets in an effective way, and having more spacecraft would add little value other than increased redundancy. Even in HOBOCOP's case, the stakeholder saw no added *scientific* value (as opposed to lifecycle value from redundancy) from having more than four assets.

In contrast, if in-situ measurements of a spatially distributed field are driving the value of the mission, there may be *no upper limit* to the feasible number of assets that would add value. As more assets are added, the spatial resolution of the measurements increases, creating finer and finer detail that may be important for science goals. Higher resolution may be especially valuable if it captures some previously unknown emergent property of a target of investigation (for instance, the properties of the entire ocean cannot be defined by observing a single water molecule). In cases like these, it is important to have some coverage metric to describe how “good” one design is compared to another, especially when limited, intelligent asset distribution is preferable to the arbitrary distribution of many assets.

The fact that there may be no upper limit to the number of value-adding assets provides a concrete example of the weakness of MAUT applied to multiple stakeholders, or even multiple science goals as perceived by a single stakeholder. Suppose the maximum perceived level of an attribute affected by the number of simultaneous samples is high. It may then be impossible to completely satisfy that attribute under a given budget. A design may have the ability to fundamentally revolutionize our scientific understanding of a phenomenon while still having a science goal $MAU < 0.05$ if the maximum level is too high. Meanwhile, another design may satisfy a different goal with much more achievable expectations with a science goal $MAU = 1$ but only deliver $MAU = 0.001$ for the first goal. A designer unfamiliar with the analytic techniques presented in this thesis may erroneously conclude that the second design is better, but the stakeholder would clearly want the first one.

Validating the value model through multiple iterations with stakeholder feedback is critical, especially when inputs from multiple stakeholders are being considered. Just because one design scores $MAU = 0.1$ from one stakeholder, and another design scores $MAU = 0.5$ from a different stakeholder, does not mean that the second design is more valuable than the first design to either stakeholder. This is also why recent research in multi-stakeholder negotiations is helpful for not only finding compromise solutions in competing stakeholder needs, but also for finding the most value-driving designs among missions with multiple goals.

9.1.5.3 Case Study #3: ÆGIR

The primary purpose of the ÆGIR mission is to characterize the water content of many asteroids for future in-situ resource extraction and utilization. This is accomplished by using assets that leverage *census sampling* to conduct many measurements of many targets much more quickly and feasibly than any single spacecraft could. The assets could also leverage *self-sampling* with the right payload to conduct radio observations to gather even more data on asteroid interiors that cannot be gathered as effectively with any single asset.

Value modeling using the RSC method showed that the time sensitivity of the data returned was critical for mission success in a census. Additionally, it showed how self-sampling can be a much less expensive way to gather data that would otherwise require far more equipment that may obtain higher-quality results but on a much smaller target area. Compared to expendable penetrators that could only be used once or landers with massive drills, radar penetration is a fast, high-TRL method for gathering data over a large area that could take months for a single surface asset to gather.

When designing a mission that leverages census sampling, it is important to understand and articulate the fundamental goal to help define the relationship between the number of assets that would be needed and the capability that is required of each asset. As with simultaneous sampling, there is a balance between the number of measurements that can be taken and the quality of those measurements, depending on the instrumentation required. If an individual asset's capability is too low, the measurements may have too much uncertainty, which defeats the purpose of a census to statistically characterize a large population.

Understanding both the time value of the sampling and the statistical characterization is critical to delivering maximum stakeholder value. A mission with three assets, each with tremendous capability, may be able to visit half as many asteroids as a mission with six assets, each with the capability to achieve just the bare minimum necessary of the goals, or visit the same number of asteroids but in twice the amount of time. Whether quality outweighs quantity is dependent on what defines quality and how quantity delivers value; the RSC method can show the best way to balance the total number of assets with the capabilities of each asset to find Pareto-optimal designs for any given stakeholder value model.

When designing a mission that leverages self-sampling, it may be valuable to gather information on the alternative options available to complete the mission. Doing so may show that it is unwise to attempt a monolithic option for that mission. GRAIL and GRACE were much more affordable than having ground assets travel to make gravimetric measurements, even if ground measurements would have been more accurate. ASI is a tremendous technical challenge, especially with DSS in orbit compared to conducting interferometry with ground assets, but compared to launching a telescope or antenna hundreds of meters across, the technical challenges seem minor.

Understanding the tradeoff between some metric of resolution relative to some metric of coverage is critical to the value proposition. If the large-scale, emergent properties of an individual target of investigation are more important than detailed investigations of individual points on the target, and if a self-sampling technique can deliver data over that scale, a self-sampling DSS may have the same value as hundreds of those satellites that make point measurements.

9.1.5.4 Case Study #4: GANGMIR

The primary purpose of the GANGMIR mission is to study many of the highest priority goals related to human exploration of Mars, which involved studying both Mars' atmosphere and regolith. The mission's goals are accomplished by using assets that leverage *stacked sampling* to conduct both in-situ and remote measurements observations of the weather and atmospheric temperature that a single asset would typically be designed for. Value is also added by leveraging *staged sampling* to capitalize on opportunity to conduct follow-up observations in the event of either a small find that the stakeholders may desire better confirmation of or a major discovery that would require a higher burden of proof to accept.

Value was added to this mission not only by assets operating in high-risk, high-value environments considered too risky for a single asset, but also by deploying those assets in heterogeneous environments. These assets can collect synergistic or complementary data from environments in which a single satellite would be ill-equipped to operate.

Not only did the RSC method find tensions between design variable options that would satisfy one goal but hurt another, it also helped express the added value of carrying additional assets and the inherent flexibility of being able to change operating locations over the mission lifecycle. Scientists will always desire "just one more," but RSC showed what the returns on "just one more" truly are compared to static evaluation methods that do not consider alternative scenarios in the future, when stakeholder perceptions of value may change considerably.

When designing a mission that leverages sacrifice sampling, it is helpful to understand the relationship between cost and operational risk. Reducing risk may cause significant design cost increases that make additional assets far less affordable, but accepting too much risk could result in dissatisfied stakeholders in the statistically possible event that every asset fails. Likewise, as in the case of census sampling, there is a balance between the added value of exploring a site in more detail, such as with a rover, and the added value of exploring many sites but in less detail. Visiting more sites may please more scientists, but full satisfaction may not be achieved by any of them if the assets do not have enough individual capability. If spatial distribution of those sites is important, then a single rover is far less likely to be as scientifically valuable as many penetrators.

From a broad perspective, the GANGMIR mission has the potential to leverage nearly every emergent capability defined in Chapter 3, depending on what assets and instruments it uses. The penetrators are being sent to locations from which they will never return, and the journey itself is risky, so it is a case of *sacrifice sampling*. The fact that they are being deployed in a time-delayed fashion rather than all at once shows that they are leveraging *staged sampling*. The small satellites are working together to lower the revisit time of measurements of the atmospheric temperature, which is a case of *shared sampling*. The weather stations on the penetrators are sampling points on the planet at the same time, which is a limited case of *simultaneous sampling*. The combination of data from the ground and from orbit shows synergy, which is a case of *stacked sampling*. Finally, the penetrators are sampling a number of scientifically interesting locations around the planet from a list of candidate sites, which could be seen as a limited case of *census sampling*. The only emergent capability that is missing from this mission is *self-sampling*.

The RSC method and the modifications and assumptions that have been made to define value over the lifecycle show how *all* of the emergent capabilities of DSS can be leveraged to achieve new scientific vistas in space science missions.

9.2 Potential Impact on Systems Engineering

This work is meant to serve as a framework or blueprint for scientists, engineers, or teams consisting of one or both working together to design a mission to achieve fundamental space science goals. The methods that have been developed are applicable to a wide range of applications, not exclusively DSS space science missions. This section serves as a message specifically to the systems engineering community.

9.2.1 Consequences of the Proposed MAUT Hierarchy

The MAUT hierarchy shown in Figure 4-5 was developed specifically for this work on DSS space science missions. While the hierarchical structure was only intended to break down a complex problem into smaller subcomponents, an unintended consequence of this formulation is that it naturally fits within the scope of how science missions are already developed with different levels of mission managers, instrument PIs, and instrument technicians.

In complex space science missions, no single person may have enough expertise to have expert opinions at every level of the design that are valid for eliciting. There are mission PIs who are in charge of the whole mission, but there are also instrument PIs who may be in charge of a single science goal. Working with the PI is an instrument scientist or technician that may have greater expertise on the lowest level of the attribute hierarchy. When value models are elicited from these stakeholders using the proposed MAUT hierarchy, they naturally feed upward into ranking full mission satisfaction.

Some space science missions, especially ones with a narrow focus or a single goal, may not fit as rigorously within the hierarchy, but another consequence of how the hierarchy is formulated is that it has the ability to collapse on itself when the value model complexity is low. For instance, in ExoplanetSat, the value model collapsed from a series of attributes into a single mathematical function that was then mapped to a single utility. By casting a wide net from the beginning with any complex design, the value model can be at least bounded with this approach.

9.2.2 The Nature of Performance Modeling

Performance modeling can be an open-ended process that constantly improves the quality and certainty of how well designs are assigned attribute levels and ranked. However, at some point, a level of fidelity will be reached where improving the performance model only increases the computation time and modeling expense without showing any worthwhile changes in the utility of the designs that are being compared. The main question modelers need to ask themselves is “when is this good enough to stop?”

The performance model should be driven by the value model, meaning the outputs of the performance model need to be high-fidelity enough to capture changes in the attributes that the stakeholders care about. For instance, an EDL model for a spacecraft would be very different depending on whether the stakeholder cared about 100 km or 10 m landing accuracy. Sometimes basic approximations are fine at this stage of the design process, especially given the large amount of uncertainty associated with this stage, but determining how accurate the performance model needs to be is really dependent on the case.

On the other hand, modelers should not assume that the initial bounds of the attribute levels as they were solicited during stakeholder interviews are necessarily the final levels that the performance model would need to consider. Any epoch shift can change the minimum acceptable and maximum perceivable levels

of the attributes, so the performance model should be accurate over a range of feasible levels that the stakeholder might care about.

In the case studies explored in this thesis, there are many instances where a performance model was created with unnecessarily high fidelity, such as the GANGMIR EDL model and the HOBOCOP heat shield model. One way to speed computation time or reduce the calculations required for the full tradespace is to decrease a module's fidelity. In these cases, the lower-fidelity model was validated against the higher-fidelity model to ensure that it was correct enough before being put into general use when evaluating the full tradespace. Another way to reduce computation time was to use batches of data generated from high-fidelity models and then interpolate the results to other designs.

There are also many instances where a performance model was initially too low fidelity, such as the HOBOCOP propulsion and astrodynamics module and the ExoplanetSat orbit propagator. In these cases, the potential for error was caught after studying how second-order effects affected the final outcome on a few examples, rather than building an entirely new system to analyze every possible variable. In these cases it was necessary to understand what the next level of fidelity would require so it could be implemented at least once.

Decisions about whether to pursue higher fidelity in other cases can be made with careful sensitivity analysis during the construction of the full performance model. The best practice is for the model to be as modular as possible so that each module can be upgraded in parallel. High-fidelity models take too long to develop and run compared to what is gained as a result of higher fidelity at such an early stage in the design process when programmatic uncertainty is larger than model uncertainty.

9.2.3 Modified RSC Method in Other Applications

While the framework developed in Chapter 4 was specifically designed to compare DSSs, the assumptions that were needed to develop a value model and articulate uncertain benefits over a system lifecycle are applicable to a number of other fields, some of which have already benefited from being examined through the lens of the value-driven design paradigm.

The benefit of the proposed methods is that new metrics can be applied to a range of other complex design problems that evolve in time. In contrast, the general case of EEA allows for systems to revert back to previous epochs over the course of their lifecycle. This makes it difficult to compare system lifecycle values for different designs across different eras because it may not be fair to analyze all eras equally. In the general case, the proposed value metrics are not accurate. However, for systems wherein the stakeholder perceives the change in value as a march in time across epochs rather than a random walk among all possible epochs, VaO and the expected value of the system can be used to objectively compare alternative designs.

The military systems that EEA was developed for may switch back and forth between epochs like "peacetime," "wartime," "search and rescue," and "budget shortfall" quite randomly. For many other systems, the construction of an era may never result in revisiting the initial epoch for satisfaction. For instance, a car buyer looking at options for a BMW may learn about additional luxuries that would only raise their minimum expectations for the car that they buy, and those expectations would never go back down once they learn about them. Likewise, the expectations of a home buyer looking into local school districts would only increase over time, which would create epoch shifts in the value model of the buyer.

The assumptions that were made in Chapter 4 for DSS space science missions are also applicable to a number of space missions, whether they are monolithic or distributed. The method may also benefit design work involving prospective opportunities, such as mining, oil drilling, and other Earth-bound exploration design.

More importantly, in any design project, it is imperative to understand what the customer wants before the architect can develop a design for the customer. The value modeling techniques described herein can be applied to almost any product acquisition process. Although in some instances it may be deemed as overkill, it may also uncover previously unarticulated stakeholder expectations that would never have been uncovered otherwise. Rather than waiting until the end of the systems engineering engine to validate the product against the stakeholder's expectations, as is shown in NASA's Systems Engineering Handbook, the methods described in this work show how product validation and sensitivity analysis can be achieved up front, rather than as an afterthought that may accidentally be forgotten, especially under programmatic uncertainty like budget cuts.

9.2.4 Recommendations for Concurrent Design Studios

The work in this dissertation is especially valuable as a set of examples for concurrent design engineering studios that quickly develop mission concepts to estimate feasibility and costs, such as JPL's Team X. There are a number of recommendations that could be made to such studios engaging in rapid concurrent design work, not only for space science missions but also for any other complex systems.

Among the top priorities for these studios is to revisit how they interact with customers. In the design studios, there are a number of stations dealing with individual engineering or science specialties. By adding one person to the team to serve as a value modeler and stakeholder interviewer, these studios can drastically increase the value of the output from concurrent engineering sessions so that the stakeholder actually receives what he or she wants in the first place.

Most of the work that this extra role would need to accomplish can be done up front without the entire team, which saves costs. The more that the stakeholder's needs can be communicated to members of the team through this new position, the more that time is saved so that the cost of these sessions can be lowered. If the value modeler can communicate an "engineering" version of the scientists' needs before the work begins, the entire team will have a clearer objective in mind during the design session.

One of the major flaws in rapid prototyping and concurrent engineering studios is the limitations that are unintentionally applied by quickly spiraling towards point solutions. During Pre-Phase A concept exploration studies, it is important to cast as wide a net as possible to capture many solutions to a problem, so that the stakeholders can eventually decide for themselves what solution is best for them.

Concurrent engineering studios are already well-equipped for handling the massive amount of data that can be generated in a concept selection, but current systems engineering practices are geared toward point solutions rather than tradespace exploration. By using value-driven design practices up front, studios can identify what variables and variables levels drive value the most and understand the synergies between different options as the stakeholder perceives them. If concurrent engineering studios continue spiraling toward a point design, they will prematurely limit the design space and therefore the customers' potential options for achieving mission success as they perceive it.

Finally, even though it is clearly better to explore many design options, it is still important to help the customer sift through the data that has been generated as a result of these concurrent engineering sessions. The value modeler should continue working with the customers after the concurrent engineering sessions to help them decide on a final solution; after all, only one mission will be designed in the end, and this mission must suit the stakeholder's needs within their proposed budget. It is difficult to examine the intricate data sets that can be generated using the RSC method without any previous experience, and more tools can be developed to communicate the options to the customers with more practice using value-driven design and the modified RSC method presented herein.

9.3 Potential Impact on Space Science Missions

This work is meant to serve as a framework or blueprint for scientists, engineers, or teams consisting of one or both working together to design a mission to achieve fundamental space science goals. The methods that have been developed are applicable to a wide range of applications, not exclusively DSS space science missions. This section serves as a message specifically to the scientific community.

9.3.1 Mission Concept Development

The most important message to scientists is that in order to use the methods developed in this dissertation, missions should be question-oriented instead of destination-oriented. These four case studies have shown that there are scientific opportunities that are made possible by leveraging the emergent capabilities of distributed satellites that have not yet been achieved in previous space science missions.

One way to avoid a destination-oriented design approach is to consider all of the unanswered research questions and prioritize them by scientific field rather than by solar system destination. This is partially accomplished by the Decadal Survey, but when concept missions are proposed, the destination tends to take the higher priority. In addition to the planetary formation question posed by Castillo-Rogez et al. involving visits to both Phobos and Deimos [198], there may be other broad questions that have been left unanswered because it is impossible to answer them with a monolithic system. The current concept selection and systems engineering paradigms make distributed options less attractive when other sets of goals or destinations can be chosen, but if the value of the goals is reconsidered in a broader context, distributed systems will be more likely to be launched to answer bigger, more fundamental questions.

Another important message to scientists is that bigger is not always better. There is a balance between fewer (or one) high-performance assets and many, low-performance assets. The point of this research is not to state unequivocally which one is better, but rather to give scientists the tools to decide which design choice is best for their mission and goals. For example, in HOBOCOP, a large number of very small magnetometers would map the 3D structure of the solar wind better than any single high-performance magnetometer could.

A third message (that is somewhat related to the second message) is that heterogeneous design options should be considered in the concept selection process despite the fact that the redundant subsystems required by multiple free-flying spacecraft potentially add cost. Some science goals may be achieved by sending assets to locations where it is too expensive or risky to design other instrument packages to survive in the same environment. Separating delicate instruments from ones that travel to high-risk environments can reduce overall risk. For example, the HOBOCOP VFMs were kept safe from overheating by not being allowed in the closest orbits, and the Galileo atmospheric probe leveraged

sacrifice sampling to gather valuable data while the main spacecraft continued on its mission. Additionally, some instruments work better when separated from others or sent to alternate locations. For example, in GANGMIR, it would be pointless to have a sounding radiometer underground, so it was not placed on the penetrators. Even a small number of heterogeneous assets can drastically increase the value proposition of a space science mission.

9.3.2 Articulating Science Returns

The methods developed in this research may seem to put an additional burden on scientists while they are trying to articulate mission goals and requirements. That additional burden is related to the extra considerations in value-driven design processes. Scientists have to be able to express quantitative levels of satisfaction for different levels of different attributes. Coming up with this list of attributes, their minima and maxima, and the utility curves is not a fast or easy process; it requires the scientists to spend more time thinking about what specifically delivers satisfaction (e.g. the characteristics of data packages that come from missions and how and how well they achieve science goals). In addition to identifying the goals, it is crucial to identify what about those goals will define success.

A major difference between typical design processes and value-driven design in tradespace exploration is that “requirements” should not enter the lexicon in the concept development stage. Divorcing the concept exploration process from requirements is crucial because requirements inhibit creativity and prematurely place unnecessary constraints on the design space. These constraints are detrimental to the tradespace exploration process because it prevents objective comparisons from one design to another – if a design meets all the requirements, it is exactly equal in perceived satisfaction to every other design that also meets the requirements.

It may be unusual at first to have so many possibilities and so few boundaries and to have to express an *opinion* when attempting to research a scientific *fact*, but with well-articulated science goals, feasible and value-delivering missions can be designed. The value-defining methods developed in this thesis allow scientists to focus less on the “requirements” and more on the drivers that deliver satisfaction to the scientists and answer the fundamental science questions.

The value elicitation process helps scientists articulate to the engineers what the requirements might be once a design is chosen. As a result, the validation stage in the systems engineering engine is emphasized up front and is always at the forefront of the requirements definitions process rather than

Value elicitation using the methods described herein is similar to the techniques that are currently in use, such as articulating goals, measurements, and acceptable levels; a key difference is that scientists have to emphasize where value is derived, which can become complex with complex goals. Exactly how value is perceived is not always obvious from the beginning, but articulating value is the cornerstone of the RSC method. More will be discussed on this important subject in Section 9.4.2.

9.3.3 Complementary and Competing Science Goals

The case studies presented herein show how missions can benefit from synergies between science goals. Stakeholders perceived higher satisfaction by partially accomplishing multiple goals compared to fully satisfying only one goal in both HOBOCOP and ÆGIR. The complementarity between the goals leads to synergy in designs that raises their rankings in the tradespace.

However, the RSC method can also show when goals compete with each other to the point where trying to satisfy more than one goal is detrimental to overall mission satisfaction. The best solution for some cases when this happens may be to split the mission into multiple independent missions with separate stakeholder value models.

In GANGMIR, the goal to observe and characterize dust was at such a low priority that increasing the radiometer's aperture and maximum wavelength *or* adding a dust camera to the penetrators to accomplish this goal removed those designs from the Pareto front. The competing needs of these science goals caused the dust goal to only be partially satisfied in any Pareto-efficient design.

Another example comes from Ross's master's thesis in the first Terrestrial Observer Swarm (A-TOS) case study [25]. This mission had competing science goals, and it made sense to split the mission into either two separate ones or use heterogeneous spacecraft in different orbits instead of shoehorning instruments for both science goals onto the same spacecraft. This is because the best orbits for each goal were very different, and trying to satisfy both goals from the same orbit was more than double the cost of launching two separate missions (see the figure on page 80 of [25]); as a result, compromising made no one happy. This discrepancy was even more obvious than it was in HOBOCOP, where the designs that tried to accomplish everything with a single spacecraft weren't even considered because they were technically infeasible, not just expensive.

The RSC method is intended to be a prescriptive method to help identify sufficiently satisficing designs early on in the system design lifecycle. If competing priorities and design variables are seen up front, the prescriptive recommendation from RSC may be to remove some goals from the scope of a mission or have two or more dedicated missions. Such decisions may require inputs from programmatic leaders and thus were not considered within the scope of this research, but the RSC method can help determine a good course of action up front.

9.4 Additional Research Observations and Findings

In addition to the answers to the primary research questions posed in Chapter 1, there were a number of other findings and observations that came from the case studies. These observations provide more heuristic guidance for a scientist or engineer engaging in a Pre-Phase A study that follows the methods developed in this work. This section will discuss some of those observations and make specific recommendations for others to follow.

9.4.1 The Truth of Phenomena

Although the breadth of science goals explored in the case studies was intentional, one *unintended* variable among the science goals was the difference in the knowledge available before each mission about the scientific phenomena that was being explored. All four case studies explored science goals with different levels of uncertainty; all science achievements can be thought of as reducing the uncertainty in the collective ability to express, model, or predict phenomena.

There are many taxonomies for expressing different levels of uncertainty. In systems engineering, McManus and Hastings [37] developed a taxonomy to characterize uncertainty in a broader sense and mitigate it in complex systems (although there are other taxonomies of uncertainty in the research

community [520]). According to McManus and Hastings's taxonomy, uncertainty can have "one of several flavors":

- Statistically Characterized (Random) Variables/Phenomena: Examples include the fatigue properties of specific materials, seasonal weather patterns, and population distributions. Anything that can be modeled with a statistical distribution and a confidence level falls into this category.
- Known Unknowns: Examples include Hawking radiation and the properties of quasars. These are phenomena that may exist but are not well characterized. Their properties are at best bounded, and they are usually handled qualitatively.
- Unknown Unknowns: Examples include the San Andreas Fault and the Tunguska event before the year 1900. These phenomena cannot be planned for, only speculated about. However, a large part of the spirit of exploration is to discover Unknown Unknowns to ask research questions about them.

Unfortunately, this classification does not discern with fine enough detail the differences between the levels of uncertainty in the science goals described in this work. Under this taxonomy, the ground science goals of GANGMIR and the goals of ÆGIR are both "Known Unknowns," while the remaining goals can all be classified as "Statistically Characterized Phenomena" despite huge differences in describing how to accomplish those goals within each of those categories (corresponding to the three "flavors" listed above). The goals of ÆGIR and ExoplanetSat both involve statistical satisfaction criteria based on *mathematically constructed* value models despite being in separate categories, while the ground science goals ("Known Unknowns") and atmospheric science goals ("Statistically Characterized Phenomena") were both well-suited for MAUT-based design rankings despite being in different categories.

Based on the experience gained through these case studies, it is the opinion of the author that the amount of knowledge known before a mission is launched lies somewhere on the spectrum between "Unknown Unknowns" and "Known Knowns." The level of uncertainty associated with a science goal has a *direct impact on how the value model for the science goal should be elicited* and whether or not a MAUT-based value model or mathematically constructed value model would be better suited for ranking design alternatives.

A summary of these findings is shown in Figure 9-1. On the far right of the spectrum are the "Unobserved Phenomena," or things that we only have enough knowledge about to ask the question itself. An example of this is "what is the distribution of alien lifeforms in the galaxy." There is enough knowledge to ask the question because the concept of "aliens" has been proposed, therefore it is no longer an "Unknown Unknown"; however, no credible evidence for the existence of aliens has been discovered, let alone what the distribution might be. Once a phenomenon has been observed, more questions about it can be raised that can help focus a scientific inquiry and the uncertainty moves into the second category, "Observed Phenomena."

The ground science goals in GANGMIR lie somewhere within these first two categories. Some could say that there is evidence for biohazards due to the discovery of perchlorates from the Phoenix lander, though no concrete evidence exists yet for biohazards beneath the surface. A value model based in MAUT was adequate for describing scientific stakeholder satisfaction because the elements of characterizing the phenomena were based on perceptions of certain attributes, including how many attempted observations

that could be made. A successful mission would move the science goal from either of the first two categories into the next category, “Bound Phenomena.”

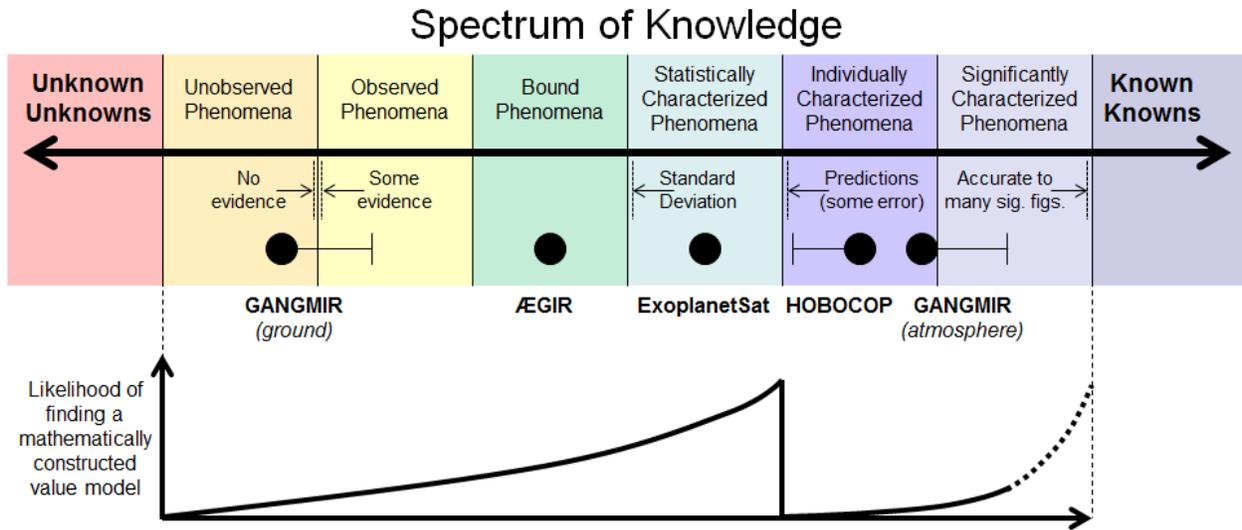


Figure 9-1: Spectrum of science knowledge for a science goal.

“Bound Phenomena” are phenomena that have been observed several times but are still poorly understood and have too much uncertainty to be “Statistically Characterized” within McManus and Hastings’s taxonomy. ÆGIR’s science goals related to understanding the statistics of water distribution in asteroids, which has been observed many times from remote measurements and meteorite samples, but there is not enough data to statistically classify or predict with relative certainty how much water is present on a particular asteroid. While there were enough descriptive attributes to partially construct a value model based in MAUT, statistical metrics did have some influence on how satisfaction was perceived, even if the stakeholder did not perceive the statistics directly. Both ÆGIR and the Kepler missions’ goals are to move a “Bound Phenomenon” into the “Statistically Characterized Phenomenon” category.

“Statistically Characterized Phenomena” are phenomena that can be described with standard deviation and more certainty, but they cannot necessarily be predicted with accuracy. Thanks to the efforts of the Kepler mission to move exoplanet knowledge into this category, ExoplanetSat hopes to discover exoplanets so they can be characterized further with follow-up studies. This mission’s value model was entirely described by a mathematically constructed value model, a fact that cannot be ignored in the bigger picture of designing space science missions. Once a particular exoplanet is found, that particular planet can then be “Individually Characterized.”

“Individually Characterized Phenomena” can be predicted with some degree of accuracy. For instance, the heliospheric current sheet is known to vary as a Parker spiral, and HOBOCOP intends to map it with greater detail. Likewise, the temperature on Mars is known and has been modeled, but higher accuracy is needed for longer periods of time in order to accomplish the goals of landing humans on the surface using aerobraking. Describing these phenomena in more detail suddenly changes from a strict statistical approach and back to an attribute-based approach, so once this threshold has been crossed, MAUT becomes very relevant again. The switch from characterizing phenomena using a statistical approach to an individualized one is illustrated by the sudden drop in the probability of finding a mathematically constructed value model at the bottom of Figure 9-2.

“Significantly Characterized Phenomena” were not particularly well explored in this work, but examples on the far right end of the spectrum include natural constants that are known to very high degrees of certainty. One could say that the number π is one of the closest things there is to a “Known Known” because it can be calculated to billions of digits and mathematically defined. Although the author has no experience in this area, it is his opinion that the value model for describing these phenomena in more detail may be better described by specific mathematical objectives rather than a collection of attributes.

Scientists or engineers planning future space science missions leveraging DSS must be aware of the uncertainty inherent in the mission’s science goals so they can be aware that instead of using a general value modeling technique, like MAU or AHP, one could instead use a more explicit value model based on the known relationships between value metrics of interest. By casting a wide net with these four very different case studies, this dissertation captured unexpected insights that were not intentionally sought through the research approach. Other scientists or engineers can avoid the pitfalls of choosing the wrong type of model if they know that a constructed model based on scientific literature may better describe the science goal.

This is not meant to imply that eliciting a value model is “simpler” for statistically characterized and significantly characterized phenomena than it is for other phenomena. The attributes for coverage involved in some cases can be complex and require great care to be validated against the stakeholder’s expectations. It is very easy to calculate the wrong coverage metric, especially if the stakeholder isn’t communicating their idea of coverage with the appropriate terminology. The application of this spectrum of knowledge taxonomy to non-science inquiries such as military systems or other consumer products is an open area for future research.

9.4.2 The Nature of Modeling Uncertainty in RSC

The Responsive Systems Comparison method was developed specifically to handle uncertain assumptions in the models that are used to conduct concept exploration case studies and to identify sensitivities that can quickly change the value of a given design. EEA is set up to allow for minor variations in the parameters that are critical for ranking designs, such as performance model (design) assumptions, exogenous (context) parameters, and value model (expectations) inputs; the framework can be seen as a set of broad sensitivity analyses across uncertain (possible) futures. These analyses are done in the beginning of the design process rather than at the end when there may be limited resources to conduct a proper sensitivity analysis or when it’s too late to incorporate any knowledge gained from this exercise.

Additionally, MATE can be used in tradespace exploration studies for uncertainty representation and propagation. Statistically characterized uncertain input parameters can be used to represent error ellipses and compare the relative uncertainty of single designs or compare the variability among many designs and design families [28]. Such analyses are done even in older methods like the ones shown in NASA’s systems engineering handbook [12], but the framework of MATE allows for easier understanding of how uncertainty propagates across the entire tradespace at once.

For cost models, minor variations in the cost coefficients tend to affect all designs in the same direction, though not necessarily with the same magnitude or ratio. In ExoplanetSat, only one type of asset was modeled, and all assets were the same mass, so minor variations in cost between cost model epochs were mostly due to variations in orbital parameters. Although the Pareto front experienced minor variations due

to changes in dominance, all designs that were Pareto-efficient in the initial cost epoch were Pareto-efficient in the others.

For other missions with very different concept designs and a wide range of potential orbital masses, models with mass-based CERs tend to have larger error cost ellipses on high-mass designs compared to low-mass designs [29]. This is evident in the Space Tug case study, where the high-mass nuclear designs were spread out on the cost axis, whereas the other concept designs with lower mass were spread more along the utility axis [50]. In a mission like GANGMIR, where costs can change due to a variation in the cost of one asset but not another, or one instrument but not another, only the designs that use that instrument or asset would change position on the tradespace.

For performance models, the assumption that minor variations affect all designs equally is not necessarily accurate, depending on what is being varied and whether it affects all designs or only a subset. In HOBOCOP, an alternate epoch was modeled to account for possible changes in mass and performance coefficients. Many designs that were Pareto-efficient in the optimistic technology epoch were infeasible in the pessimistic epoch. Additionally, with more complicated missions where a change in an assumption affects a subset of designs far more than another, such as one with multiple potential assets, the designs that do not carry the asset examined by one variable's sensitivity will not experience a change in utility compared to designs that do carry that asset.

For value models, changes in the value model can have drastic effects on the ranking of designs, especially when redundancy or substitution effects are present. Some sensitivities that could be examined include the k -weights at every level of the MAUT hierarchy, the minimum acceptable and maximum perceivable attribute levels, and the baseline points on the attribute utility curves. Research on value robustness has explored this in order to determine possible attractive designs that are insensitive to various assumptions (or mistakes) in value model elicitation and formulation [41], [192].

Sensitivity analyses in the value model are relatively easy to perform because they do not require rerunning the performance models, which can be complicated and take hours or days. Each k -weight can be varied and the whole tradespace's utilities can be recalculated. Tools like IVTea can quickly help visualize these changes and identify Pareto-efficient designs or designs with high NPT, fNPT, eNPT, or efNPT. Because these metrics have already been explored in previous research, delving into them with these case studies was considered beyond the scope of the research presented herein.

9.4.3 The Difficulties of Stakeholder Value Modeling

Stakeholder value modeling is not an easy process; it is more of an art than a science. As expected, there was a great deal of difficulty extracting a value model from the stakeholders in each case study for a number of different reasons. Some of this came from fundamentally improper understandings of the phenomena being observed. This required a reexamination of the goals from an entirely different angle and extensive research in the field of study so that the designers understood why certain aspects were valuable, while the original assumptions about value were not correct or otherwise inconsequential. Other difficulties arose as a result of communications issues between the designer and the stakeholders, or because it was difficult to communicate the purpose and importance of MAUT as a VCDM. Stakeholders

who were unfamiliar with the process were challenged to concretely define quantitative levels of satisfaction and choose between example alternatives under uncertainty.

This section describes some of the insights and difficulties that arose during the interview processes. Any person acting in the interviewing role and eliciting a value model from someone else should understand how their skills in value elicitation can be improved.

9.4.3.1 Method Acting As a Scientist

In many ways, value elicitors must metaphorically put themselves in the shoes of the stakeholders they are interviewing to truly grasp how the stakeholders perceive value, in a similar way that method actors prepare for their roles in film. Method actors take their roles to an extreme level by transforming their lives into the lives of the characters they play. Even the most mundane details of the character, whether or not those details are even shown through the screen, are studied and practiced before the cameras roll. Nicolas Cage learned to speak the Navajo language fluently just to understand his character in *Windtalkers*, even though he never spoke the language on screen. Robert De Niro worked 12 hour shifts as a cab driver to prepare for his role in *Taxi Driver*. Christian Bale lost 60 lbs. for his role in *The Machinist*.

If the interviewer in a case study lacks a fundamental understanding of the science goals of a mission or does not have the expertise to understand related papers written by the scientist stakeholder, then it is nearly impossible to understand the source of the stakeholder's perception of value. A simple crash course on Wikipedia is not enough to bring the modeler up to speed; diving into the scientific literature is usually necessary, which is why such great care was taken in each case study chapter to review the literature relevant to the science goals and previous missions. The interviewer cannot simply ask a scientific stakeholder what is important and expect them to articulate it in the proper technical language of MAUT; they must be able to speak using the language of their scientific field. (Scientists designing their own value model by conducting a self-interview would not have this problem.) Using techniques from empathy in design [521]–[523] may alert the interviewer to some cognitive dissonance between a constructed value model and their mental value model (as it did in ExoplanetSat), which will signify that the interviewer is missing a key component somewhere, but this alone will not create the value model if the interviewer is not familiar enough with the literature.

Improper assumptions about the Sun's magnetic field led HOBOCOP down the wrong path for several months before being corrected in an eye-opening conversation with an expert. The ExoplanetSat value model would not have been articulated as it had been without reading published literature. The spectral signature of pure ice versus ice mixed with other clays fundamentally alters the value of the asteroid as a target for future mining operations. A value modeler needs to be able to understand the scientific nuances involved in the mission so that they can ask the right questions to allow the stakeholder to better expressing their perception of value.

9.4.3.2 Stakeholder Interview Difficulties

Even with a more fundamental understanding of a scientific objective, value models are difficult to elicit from stakeholders who are unfamiliar with MAUT. Simply explaining how MAUT works up front is a time-consuming process, and the stakeholders will still most likely question what is the purpose of playing a game of chance with the ASSESS software.

There is value in being in the same room as the person whose perceptions of value are being elicited via a stakeholder interview. Sometimes the non-verbal cues an interviewer notices can be more important than the words the interviewee says. Paying attention to these cues helps sort out what the stakeholder feels. If the constructed value model does not represent the needs of the stakeholder, there could be some non-verbal tension in their demeanor.

During an interview with Dr. Seager, it was clear she was uncomfortable with the value model as it existed, meaning something was fundamentally wrong with how her perceptions were being represented. Rather than pressing forward and obtaining incorrect numbers, other options were explored, which is how the statistically derived constructed value model was discovered. Once that representation was presented to her, it was clear she was much more confident in expressing her needs.

Additionally, if constructed value models do exist and are more appropriate metrics for satisfaction than a simple aggregation of attributes, the stakeholder may not know about them. If they do, they may not even mention them during the interview. Other times, the stakeholder may have an idea on how to measure the value, but it is not fully developed. This also happened in ExoplanetSat, where a paper coauthored by the PSS almost fully articulated the constructed value model but stopped one step shy of expressing how to measure that mission's satisfaction. A simple change in the highest formula of that analysis was all that was needed.

9.4.3.3 Stakeholder Risk Attitudes

Understanding the stakeholder's views on risk can inform the choice between a cautious, traditional, monolithic system and a DSS mission. Thus it is important not only to make sure the attributes are independent, but also to ensure that the stakeholder will accept a design that only meets the minimum standards. This is especially important when dealing with the multi-tiered MAU structure.

An example from this research was in the series of interviews to elicit the value model for HOBOCOP. During one of the later interviews with the PSS, the articulated minimum acceptable attribute levels of the VFM instrument were possibly *too* low. When it came time to weigh the instrument MAU against the other constellation properties in HOBOCOP, this caused her to have a low risk attitude toward the instrument MAU compared to the other science goal SAUs. As a consequence, the weights of the VFM were more than double the weights on the other attributes, which would mean that a great deal of satisfaction could be achieved with a single satellite compared to one with multiple viewing angles.

The same MAU weighting interview was conducted with the graduate student, but without telling her what the minimum values on the attributes were. She only knew that her choices were "poor, but acceptable" and "the most advanced VFM to date". Knowing this, her attribute weightings varied wildly compared to the PSS's weights. With attributes of "VFM MAU," "Inclination," and "Simultaneous Viewing Angles," the PSS's weights were [0.20, 0.07, 0.03] while the graduate student's were [0.20, 0.42, 0.62]. The graduate student's point of view is that multiple angles and inclinations have not been accomplished before in space science missions, whereas high-quality VFMs have been launched on previous missions and even work on ground telescopes. Delphi methods and other techniques to achieve stakeholder consensus may potentially help alleviate this problem [524], [525].

9.4.3.4 Stakeholder Availability Issues

Stakeholder availability proved to be a recurring challenge. Whether or not the value models truly represent the accurate value or even address all the relevant attributes of a mission can only be known through appropriate follow-up interviews to validate the tradespace. This can be done by making comparisons of designs that are very different in their individual attributes but have similar MAUs, and seeing if the stakeholders have a preference for one over the other. This was difficult to do due to the limited time and availability of the stakeholders. This was not just a limitation simply because a graduate student was asking for more time with research scientists with limited availability; this is also a limitation in real missions when the attention of stakeholders may be diverted by other projects they are leading.

Additionally, the PSSs for several missions changed over the course of the case studies, resulting in a completely different value model from additional stakeholder interviews. If the value model changed significantly, the performance model also needed to be changed significantly if the attributes were not already accounted for or the minimum acceptable and maximum perceived values were outside the modeling scope of the modules. A solid foundation when beginning a study will cut down on the number of iterations required to reach validated conclusions. One can never predict when stakeholders will change their minds, but based on the experience gained through the case studies presented in this work, stakeholders' perceptions for the baseline scenario change little compared how they would change if a new stakeholder became involved with the project.

9.4.3.5 Stakeholder Minimum Satisfaction Levels and Attribute Independence

One final observation from stakeholder interviews was that the definition of “minimally acceptable” for the attributes may not be perceived as entirely independent. During discussions with several different stakeholders to identify the minimum acceptable levels of different attributes after the value-driving attributes had been identified, the stakeholders agreed that they would accept designs with the minimum level attribute for each attribute but would not be satisfied with the previously selected minimum acceptable levels of *all* attributes. Such a situation is shown in contrast to the assumed levels of independent attributes in Figure 9-2.

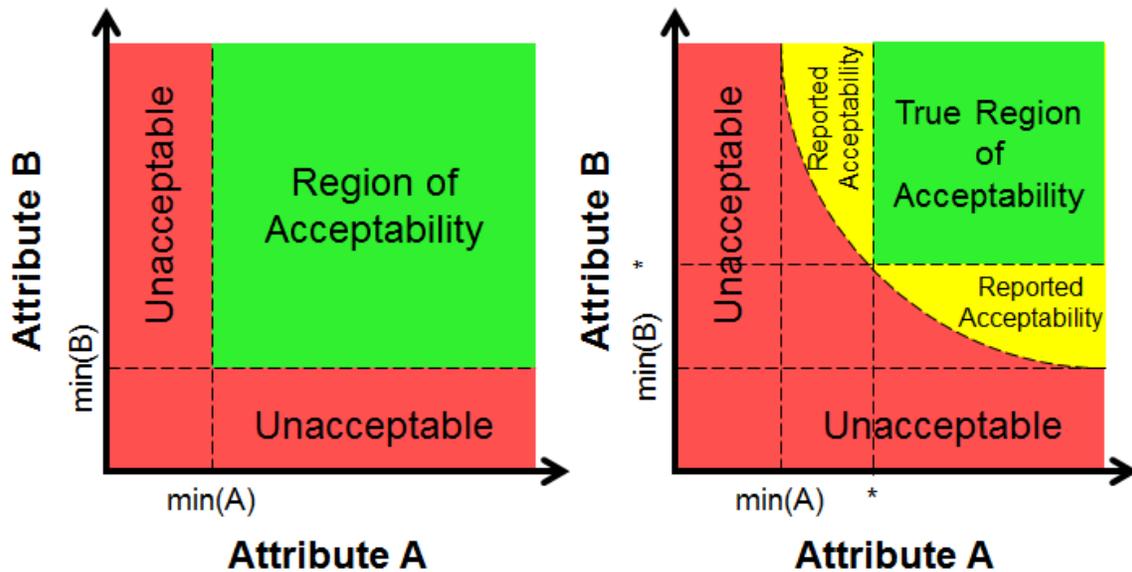


Figure 9-2: Comparison between the reported minimum acceptable levels of attributes for (left) independent attributes, and (right) dependent attributes or incorrectly reported acceptable minima.

The situation on the left is ideal; the stakeholder is still willing to accept (with the utility $MAU \equiv 0$) the design in the bottom left corner of the green region. There is no question that the attributes are independent. The situation on the right is less than ideal; it opens up the possibility that the attributes that have been chosen are not actually independent of each other, and that the stakeholder perceives a relationship between these attributes when mentally evaluating the system.

If the situation on the right side of Figure 9-2 occurs during a stakeholder interview, it is wise to reexamine the attributes to ensure they are independent; otherwise, one of the fundamental assumptions of MAUT value models is broken. Iterations on attributes are expected early in the RSC method, and attributes are allowed to change or be retailored. However, if the attributes are still considered correct representations of the stakeholder's mental model of the system, it may be a case where the stakeholder has not accurately described their acceptable minimum values for the attributes. It may also be that the stakeholder is willing to compromise on their reported minima for one attribute if they have high satisfaction in others, meaning the attributes only meet the definition of independence at higher levels than the reported minimum levels.

One way to fix this assumption is to explore sets of attribute levels along the lines of the region where the reported minimum acceptable values are. The stakeholder may be willing to accept designs with minimal satisfaction along the boundary of the yellow region on the right side of the figure because other attribute levels are high enough to achieve this minimum satisfaction. A designer could then find a set of levels (denoted on the axes by stars) that represents the best boundaries on the limits of the true region of acceptability. Defining the levels in this manner may prematurely eliminate some designs from the tradespace of acceptable designs, but these designs should already be ranked low in the tradespace (depending on the levels and weights of the other attributes) given that MAUT uses multiplicative weights rather than additive weights. It is best to limit the difference between the reported minimum and the new lower bound as much as possible so fewer designs are prematurely discarded.

9.4.3.6 Quantifying Scientific Benefits Versus Ranking Design Alternatives

One of the recurring problems when communicating the process and results from the modified RSC method is how the term “zero” is used. When defining the utility for instrument attributes on the third tier of the proposed MAUT structure, it was usually clear to the stakeholders what $U(x) = 0$ meant. However, once other goal-level and primary mission attributes on the second and first levels entered the equation, it became more difficult for them to understand why something that was below the acceptable level still delivered “zero” value to the mission.

It is important to communicate and reiterate that “value” in EEA is an ordinal ranking based on the integration of utility, not a cardinal aggregation of imaginary monetary science dollars, and that zero utility does not necessarily mean zero scientific benefit. A sketch of the difference between these two is shown in Figure 9-3. The utility function that is elicited from a stakeholder is bounded by the minimum and maximum acceptable attribute levels regardless of how far away the minimum acceptable level is from the level equal to zero.

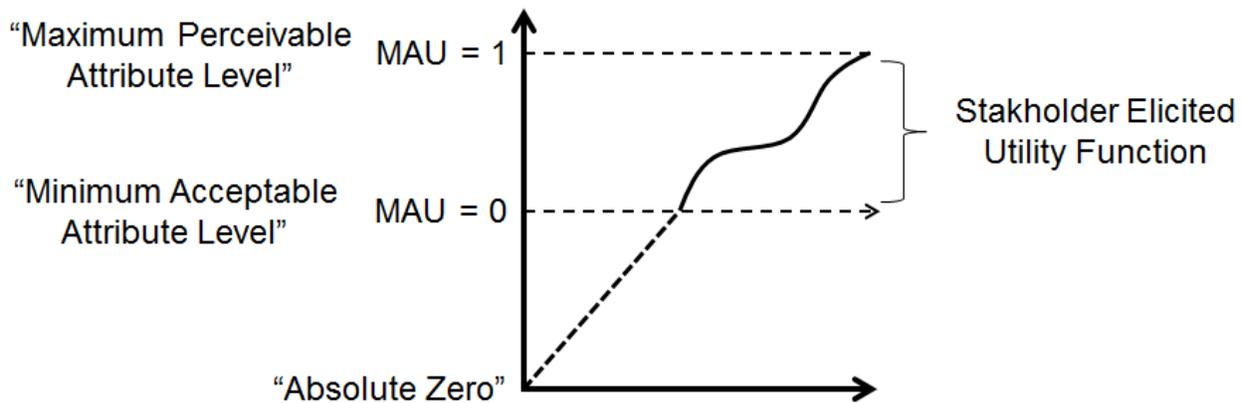


Figure 9-3: Comparison of different interpretations of "zero value" when examining systems with multiple discrete measurements.

For illustrative purposes, imagine the attribute level in Figure 9-3 is “observation lifetime,” and the minimum acceptable attribute level is two years. A satellite network that observes its target for 1.99 years is by definition unacceptable to the stakeholder because it does not meet the minimum satisfaction criteria; however, 1.99 years of observations are probably not “worthless” to the stakeholder. Since the value metric in EEA is defined by the integration under the attribute curve, it does not count the area under the diagonal dotted line or the horizontal arrow where $MAU = 0$.

However, if the minimum acceptable levels of attributes that are not directly related to the instruments are all set to zero, it may actually be appropriate to equate “zero cumulative value over the lifecycle” and “worthless.” If the stakeholder prefers to think of the quantity of science returns, this may be one way to objectively compare systems near the bottom end of acceptability.

The difficulty of communicating the true lower limits of acceptability also relates to the previous section where scientists may have different definitions of “acceptable” under different circumstances. A high-performance constellation that falls slightly short of its minimum acceptable observation lifetime attribute level (say, 1.99 years) will be perceived as having greater lifecycle value compared to a low-performance constellation that meets the minimum. More research and validation with test case studies is needed to say

with certainty that attributes like observation time and number of samples or visits work better if the minimum acceptable level is set to zero.

9.4.4 Adding Instruments or Assets to a DSS

One of the defining characteristics of monolithic systems is that adding instruments to the spacecraft or increasing the fidelity of the existing instruments nearly always adds additional benefit. This is because if there is room and budget for it, more or better instrumentation means more return on investment. If one instrument provides synergistic gains with other instruments, there is even more motivation to add a new instrument (such as what Selva showed in his dissertation [85]), especially if that instrument is small enough that the added cost is low.

However, using the monolithic approach to a DSS mission design can be detrimental. First, the addition of instruments can cause per-asset costs to rise, threatening the affordability of the original number of assets, especially in cases where there is a strict cost limit. Whether the addition of that instrument at the cost of an asset increases or decreases the value of the entire mission is dependent on how valuable individual assets are and what attributes are being driven primarily by leveraging the emergent capabilities of DSS. If the number of assets is already large (such as with many small satellites leveraging simultaneous sampling) or there are spare assets (such as with staged and sacrifice sampling), the utility penalty of losing an asset may be offset by the addition of the instrument. On the other hand, if assets are working together in series (such as with shared sampling), losing a single asset may decrease other value metrics and increase the risk of the mission becoming unacceptable if an unbroken chain of operating assets are required to meet the minimum satisfaction criteria. The RSC method displays this tradeoff in an objective way so designers can make this choice.

Second, the extra mass can change the value proposition in other ways, even possibly making the overall value less as a result. For instance, in HOBOCOP, the desired orbits for the HOBO satellites maximized eccentricity so that the satellite could make closer approaches to the Sun. Adding mass caused the final orbits to be more circular because of the nature of electric propulsion systems and the decreased acceleration with the same thrust. Decreasing I_{SP} to compensate for this drastically increases the fuel mass and therefore lowers the total number of assets that can be deployed.

By contrast, in GANGMIR, the extra mass added to the penetrators by an additional instrument makes their mass go up, which means the depth of penetration goes up. A small penetrator optimized to carry only one instrument may not penetrate deep enough into the surface, depending on how deep the stakeholders want to go. If ballast needed to be added to reach the desired minimum depth, that ballast provides far less value than another instrument or suite of instruments that can piggyback on the first instrument's investigations. On the other hand, if multiple instruments were required to divide a sample taken by a penetrator, a bigger sample or multiple drilling points may be required.

9.5 Future Work and Open Research Questions

Each case study presented in this work is meant to serve as an example to a variety of audiences. First, these studies show how distributed satellites *can* address fundamental science goals. Second, these studies show how to *define* science goals so that the stakeholders are actually getting what they want out of the mission. With improper stakeholder value models, missions may achieve the wrong goals, or fail to be attractive enough to the scientific community to get to the next phase of development.

There are also a number of areas of research that have been identified that can strengthen the modified RSC method proposed in this work.

9.5.1 Future Work on Developing Methods

Modifications to the RSC method that have been made also open up ways to revisit past research in fractionated satellites. Looking back at old research with fresh tools is not uncommon; after several satellite telecommunications programs filed for bankruptcy, de Weck et al. examined how those companies could have reduced the financial risks by designing more flexible satellites with staged deployment [78]. The value propositions can be reexamined with VaO to show whether, or under what conditions, fractionated satellite systems have a solid business case.

9.5.2 Future Work on Case Studies

Future work on these case studies can always improve on the value, performance, and cost models. Declaring a “stopping point” for when the problem setup is “good enough” was difficult and motivated almost entirely by deadlines. Given enough time, systems engineers can tinker with code, research more information, and build complexity into these models in a never-ending spiral if they are left alone with no well-defined stopping points.

The methods presented herein show not only how utility-driving design variables can be explored, but also how cost- and utility-driving subsystems options compare with previously assumed subsystems architectures. Future work can revisit the subsystems models, especially the communications subsystems, as higher-fidelity performance and cost models are developed for the different concept architectures that can be used in space systems. Having a tool like this will help express the added benefit of certain concepts to mission managers so that they can make an informed decision about the risks associated with untested subsystems rather than strictly choose off-the-shelf components wherever possible.

Ideally, each case study could be further developed into a major planetary science mission. However, if this happened for any case study, the value models may change significantly with more focused stakeholder input depending on the politics of scientific research, and the final design choice may be drastically different from the recommended design solutions discussed at the end of each chapter.

9.5.3 Performance Modeling Research Questions

There are certain levels of fidelity these models cannot achieve without future research. The most critical relationship that has been glossed over in this work is the relationship among radiation protection, cost, and spacecraft lifetime. Is there a way to estimate or place bounds on the added mass and cost of increasing the satellite’s lifetime through radiation protection (as opposed to increases only due to operations costs) during a Pre-Phase A study without building a bottom-up model of the system? A model that could answer this would be useful in EEA to show the added value of longer lifetimes by capturing opportunity or the risk mitigation by planning launches in multiple waves as assets fail over time, if they are cheap to replace and upgrade.

It would be easy to assume an exponential model on the thickness of a protection layer and some extinction coefficient depending on the material and assume a failure rate through a Poisson model. However, even this requires assumptions about the electronics board layouts, the other structures of the spacecraft, and information about how these would differ from one instrument selection to the next. Not every component of the spacecraft is susceptible to radiation damage, so radiation protection is placed

strategically on a satellite to minimize the necessary protection mass. Although uncertainty in Pre-Phase A is high, such a model may cause unfair penalizations toward either end of the lifetime length spectrum. Future studies on the operational lifetime of miniaturized satellites like CubeSats will help inform such models, but more time is needed to gather more data for a parametric model.

Another important area of future research for DSS is the miniaturization of communications technologies, including optical communication and their effect on the value proposition for all DSS missions, not just science missions. ExoplanetSat demonstrates a level of pointing precision that is unprecedented for satellites of this size, and the innovative use of a piezostage in conjunction with standard reaction control wheels allows this technology to also be used for optical communications. By reducing the contact time or total number of contacts necessary, ground operations costs with DSS can be substantially curtailed when the number of assets that need to communicate directly with the ground increases. This can also be alleviated by path diversity and using satellites as relays, but if all options can be compared, the value proposition is much clearer.

9.5.4 Value Modeling Research Questions

Developing value models led to some interesting possibilities for redefining levels of acceptability within MAUT. Section 9.4.3.6 showed an alternative method of ranking designs for quantifying scientific return that makes more sense to stakeholders who are unfamiliar with the definitions in MAUT. Future research on quantifying scientific returns can determine whether or not this alternative method will work and be more accessible to stakeholders reviewing the results from RSC.

9.6 Summary of Unique Thesis Contributions

The broad scope of the case studies touched on a variety of topics in the fields of science, spacecraft engineering, systems engineering, value elicitation, and mission planning. The work presented in this thesis makes significant contributions to a number of areas that are summarized here.

9.6.1 Definitions of Emergent Capabilities

The most basic and fundamental contribution from this work is the set of definitions of the emergent capabilities of DSS. By articulating these capabilities, others can examine ways to leverage them in space science missions to achieve unique science goals.

9.6.2 Case Study Data Sets

The data sets from all four case studies are publicly available, and all relevant code and data is included in the companion DVD submitted with this thesis. The only proprietary information that cannot be included is the launch cost estimates in the alternate cost models of ExoplanetSat. These data sets will be added to the SEARi case study database so that students can use them in future research to test new methods and tools or simply use as models for how to design RSC case studies.

Additionally, since each case study is open source, future researchers can modify the case studies by adding new modules, increasing the fidelity of the current models, updating the assumptions on critical engineering parameters, and applying new methods developed in future TSE research. Each case study also can be considered the first steps of a Pre-Phase A study that could be expanded upon and proposed to a space agency for consideration as a real mission, not just a concept study.

9.6.3 Epoch-Era Analysis Value Modeling

The four case studies are among the few so far that have included multi-era analysis, though they are not the first to use EEA to show lifecycle benefits. Modifications made to the RSC method to apply to space science missions may also apply to other fields, meaning that some of the inherent weaknesses of MAUT can be overcome in a variety of studies that can be used to show true lifecycle value of designs.

9.6.4 Constellations for Astrophysics

The coverage model developed for ExoplanetSat was the first examination into year-round coverage of targets by a constellation of satellites looking up rather than down. Although the pointing requirements are strict by most standards, the model does show how other constellations could plan their pointing scheduling and understand the capabilities of any given constellation design. This work is especially relevant to instruments that have high solar exclusion angle requirements such as other photometry studies, but it is also extremely relevant for gamma-ray burst and other phenomena that require long-duration access to targets in the sky.

9.6.5 Interplanetary Federated Communications Network

Significant effort went into examining communications architectures for deep space communications, although it was not presented in complete detail in Chapter 6. A minor contribution was made regarding the development of an interplanetary federated satellite network. The RSC framework with a monetized VCDM like Net Present Value can be used to identify the conditions under which such a network would be affordable and valuable to other missions, including human missions to Mars.

9.7 Final Thoughts

This work has shown that the near future of space science can be exciting if the paradigms of the past can be cast away to open up a new age of distributed satellite systems. There is much exploring to do, and with the tools and examples presented herein, scientists and engineers can work together to expand our scientific and technological horizons beyond the scope of what was possible yesterday.

*So for now wave good-bye, leave your hands held high
Hear this song of courage long into the night
And the wind will bear my cry to all who hope to fly
Lift your wings up high my friend **fearless to the end...***

MANOWAR, "COURAGE"

References

- [1] A. L. Weigel, “Bringing policy into space systems conceptual design: Qualitative and quantitative methods,” Massachusetts Institute of Technology, 2002.
- [2] A. Brown, “Mars Science Laboratory: the budgetary reasons behind its delay,” *The Space Review*, 02-Mar-2009. [Online]. Available: <http://www.thespacereview.com/article/1318/1>. [Accessed: 25-Jan-2014].
- [3] A. Svitak, “NASA’s Overbudget Mars Rover in Need of Another Cash Infusion,” *SpaceNews.com*, 28-Jan-2011. [Online]. Available: <http://www.spacenews.com/article/nasas-overbudget-mars-rover-need-another-cash-infusion>. [Accessed: 25-Jan-2014].
- [4] I. Klotz, “NASA budget plan saves telescope, cuts space taxis | Reuters,” *Reuters*, 16-Nov-2011. [Online]. Available: <http://www.reuters.com/article/2011/11/16/us-usa-space-budget-idUSTRE7AF06320111116>. [Accessed: 25-Jan-2014].
- [5] O. Brown and P. Eremenko, “The Value Proposition for Fractionated Space Architectures,” 2006.
- [6] D. H. Rhodes and A. M. Ross, “Summary of a New Method,” in *Value-Driven Tradespace Exploration for System Design*, SEArI, 2010.
- [7] T. P. Hughes, *Rescuing Prometheus*. 1998. New York: Pantheon books, 2007.
- [8] M. D. Fagen, A. E. Joel, and G. E. Schindler, “A history of engineering and science in the Bell system,” *Hist. Eng. Sci. Bell Syst. Fagen MD Joel Amos E Schindler GE N. Y. Lab. 1975- Bell Teleph. Lab. Inc.*, vol. 1, 1975.
- [9] Y. Monden, *Toyota production system: an integrated approach to just-in-time*. CRC Press, 2011.
- [10] “The House of Quality,” *Harvard Business Review*. [Online]. Available: <http://hbr.org/1988/05/the-house-of-quality/ar/1>. [Accessed: 15-Jan-2014].
- [11] Y. Akao, “Development history of quality function deployment,” *Cust. Driven Approach Qual. Plan. Deploy.*, pp. 339–351, 1994.
- [12] S. J. Kapurch, *NASA Systems Engineering Handbook*. DIANE Publishing, 2010.
- [13] A. M. Ross and D. E. Hastings, “The tradespace exploration paradigm,” in *INCOSE International Symposium 2005*, 2005, p. 13.
- [14] N. P. Diller, “Utilizing multiple attribute tradespace exploration with concurrent design for creating aerospace systems requirements,” Massachusetts Institute of Technology, 2002.
- [15] D. Sturtevant and Z. Szajnfarber, “Systems Engineering and MATE: Historical Roots.”
- [16] G. B. Shaw, D. W. Miller, and D. E. Hastings, “The generalized information network analysis methodology for distributed satellite systems,” Citeseer, 1999.
- [17] G. B. Shaw, D. W. Miller, and D. E. Hastings, “Development of the quantitative generalized information network analysis methodology for satellite systems,” *J. Spacecr. Rockets*, vol. 38, no. 2, pp. 257–269, 2001.
- [18] C. D. Jilla, D. W. Miller, and R. J. Sedwick, “Application of multidisciplinary design optimization techniques to distributed satellite systems,” *J. Spacecr. Rockets*, vol. 37, no. 4, pp. 481–490, 2000.
- [19] C. D. Jilla and D. W. Miller, “Assessing the performance of a heuristic simulated annealing algorithm for the design of distributed satellite systems,” *Acta Astronaut.*, vol. 48, no. 5, pp. 529–543, 2001.
- [20] C. D. Jilla, “A multiobjective, multidisciplinary design optimization methodology for the conceptual design of distributed satellite systems,” Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2002.
- [21] J. Von Neumann and O. Morgenstern, “The theory of games and economic behavior,” 1947.
- [22] H. Raiffa and R. Keeney, “Decisions with multiple objectives: Preferences and value tradeoffs,” *Decis. Mult. Object. Prefer. Value Tradeoffs*, 1976.
- [23] A. M. Ross, D. E. Hastings, J. M. Warmkessel, and N. P. Diller, “Multi-attribute tradespace exploration as front end for effective space system design,” *J. Spacecr. Rockets*, vol. 41, no. 1, pp. 20–28, 2004.
- [24] H. L. McManus and T. E. Schuman, “Understanding the Orbital Transfer Vehicle Trade Space,” *AIAA Pap.*, vol. 6370, 2003.
- [25] A. M. Ross, “Multi-attribute tradespace exploration with concurrent design as a value-centric framework for space system architecture and design,” Massachusetts Institute of Technology, 2003.
- [26] F. Black and M. Scholes, “The pricing of options and corporate liabilities,” *J. Polit. Econ.*, pp. 637–654, 1973.

- [27] M. Walton and D. Hastings, "Quantifying embedded uncertainty of space systems architectures in conceptual design," in *Proc. of AIAA Space 2001 Conference and Exposition, Albuquerque, NM, USA*, 2001.
- [28] M. A. Walton, "Managing uncertainty in space systems conceptual design using portfolio theory," Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2002.
- [29] M. A. Walton and D. Hastings, "Applications of Uncertainty Analysis Applied to Architecture Selection of Satellite Systems," *J. Spacecr. Rockets*, vol. 41, no. 1, pp. 75–84, 2004.
- [30] J. E. Derleth, "Multi-Attribute Tradespace Exploration and its Application to Evolutionary Acquisition," Massachusetts Institute of Technology, 2003.
- [31] T. J. Spaulding, "Tools for Evolutionary Acquisition: A Study of Multi-Attribute Tradespace Exploration (MATE) Applied to the Space-Based Radar (SBR)," DTIC Document, 2003.
- [32] S. Seshasai and A. Gupta, "Knowledge-Based Approach to Facilitate Engineering Design," *J. Spacecr. Rockets*, vol. 41, no. 1, pp. 29–38, 2004.
- [33] S. Seshasai, A. Gupta, and A. Kumar, "An integrated and collaborative framework for business design: A knowledge engineering approach," *Data Knowl. Eng.*, vol. 52, no. 1, pp. 157–179, Jan. 2005.
- [34] A. M. Ross, "Managing unarticulated value: changeability in multi-attribute tradespace exploration," Citeseer, 2006.
- [35] A. M. Ross and D. E. Hastings, "Assessing changeability in aerospace systems architecting and design using dynamic multi-attribute tradespace exploration," in *AIAA Space*, 2006.
- [36] A. M. Ross, D. H. Rhodes, and D. E. Hastings, "Defining changeability: Reconciling flexibility, adaptability, scalability, modifiability, and robustness for maintaining system lifecycle value," *Syst. Eng.*, vol. 11, no. 3, pp. 246–262, 2008.
- [37] H. McManus and D. Hastings, "A framework for understanding uncertainty and its mitigation and exploitation in complex systems," *IEEE Eng. Manag. Rev.*, vol. 34, no. 3, pp. 81–94, 2006.
- [38] H. McManus, M. Richards, A. Ross, and D. Hastings, "A Framework for Incorporating 'ilities' in Tradespace Studies," in *AIAA Space*, 2007.
- [39] A. M. Ross and D. H. Rhodes, "Using natural value-centric time scales for conceptualizing system timelines through epoch-era analysis," in *INCOSE International Symposium*, 2008.
- [40] A. M. Ross and D. H. Rhodes, "Architecting Systems for Value Robustness: Research Motivations and Progress," in *Systems Conference, 2008 2nd Annual IEEE*, 2008, pp. 1–8.
- [41] A. M. Ross, D. H. Rhodes, and D. E. Hastings, "Using pareto trace to determine system passive value robustness," in *Systems Conference, 2009 3rd Annual IEEE*, 2009, pp. 285–290.
- [42] C. J. Roberts, M. G. Richards, A. M. Ross, D. H. Rhodes, and D. E. Hastings, "Scenario planning in dynamic multi-attribute tradespace exploration," in *Systems Conference, 2009 3rd Annual IEEE*, 2009, pp. 366–371.
- [43] L. Viscito and A. M. Ross, "Quantifying flexibility in tradespace exploration: Value weighted filtered outdegree," *AIAA Space 2009*, pp. 1–9, 2009.
- [44] L. Viscito, D. Chattopadhyay, and A. M. Ross, "Combining Pareto Trace with Filtered Outdegree as a Metric for Identifying Valuably Flexible Systems," *CSE 2009*, vol. 2009, pp. 1–9, 2009.
- [45] M. G. Richards, A. M. Ross, N. B. Shah, and D. E. Hastings, "Metrics for evaluating survivability in dynamic multi-attribute tradespace exploration," *J. Spacecr. Rockets*, vol. 46, no. 5, pp. 1049–1064, 2009.
- [46] M. G. Richards, A. M. Ross, D. E. Hastings, and D. H. Rhodes, "Multi-attribute tradespace exploration for survivability," Massachusetts Institute of Technology, Engineering Systems Division, 2009.
- [47] M. E. Fitzgerald and A. M. Ross, "Sustaining lifecycle value: Valuable changeability analysis with era simulation," in *Systems Conference (SysCon), 2012 IEEE International*, 2012, pp. 1–7.
- [48] M. E. Fitzgerald and A. M. Ross, "Mitigating contextual uncertainties with valuable changeability analysis in the multi-epoch domain," in *Systems Conference (SysCon), 2012 IEEE International*, 2012, pp. 1–8.
- [49] M. E. Fitzgerald, A. M. Ross, and D. H. Rhodes, "Assessing Uncertain Benefits: a Valuation Approach for Strategic Changeability (VASC)," in *INCOSE International Symposium*, 2012.
- [50] A. A. Rader, A. M. Ross, and D. H. Rhodes, "A methodological comparison of Monte Carlo Simulation and Epoch-Era Analysis for tradespace exploration in an uncertain environment," in *Systems Conference, 2010 4th Annual IEEE*, 2010, pp. 409–414.
- [51] A. A. Rader, F. T. Newland, and A. M. Ross, "An Iterative Subsystem-Generated Approach to Populating a Satellite Constellation Tradespace," in *Proceedings from AIAA SPACE 2011 Conference & Exposition*, Long Beach, CA, 2011.
- [52] A. M. Ross, H. L. McManus, D. H. Rhodes, and D. E. Hastings, "A Role for Interactive Tradespace Exploration in Multi-Stakeholder Negotiations," *AIAA Space 2010*, 2010.

- [53] A. M. Ross, H. L. McManus, D. H. Rhodes, and D. E. Hastings, "Revisiting the Tradespace Exploration Paradigm: Structuring the Exploration Process," *AIAA Space 2010*, 2010.
- [54] J. C. Beesemyer, A. M. Ross, and D. H. Rhodes, "Case Studies of Historical Epoch Shifts: Impacts on Space Systems and their Responses," 2012.
- [55] A. M. Ross, M. G. O'neill, D. E. Hastings, and D. H. Rhodes, "Aligning Perspectives and Methods for Value-Driven Design," in *American Institute of Aeronautics and Astronautics: Space 2010 Conference and Exposition*, 2010.
- [56] D. M. (Douglas M. Schofield, "A framework and methodology for enhancing operational requirements development: United States Coast Guard cutter project case study," Thesis, Massachusetts Institute of Technology, 2010.
- [57] B. Mekdeci, "Managing the impact of change through survivability and pliability to achieve viable systems of systems," Massachusetts Institute of Technology, 2013.
- [58] M. A. Schaffner, "Designing systems for many possible futures: the RSC-based method for affordable concept selection (RMACS), with multi-era analysis," Massachusetts Institute of Technology, 2014.
- [59] M. A. Schaffner, A. M. Ross, and D. H. Rhodes, "A Method for Selecting Affordable System Concepts: A Case Application to Naval Ship Design," *Procedia Comput. Sci.*, vol. 28, pp. 304–313, 2014.
- [60] A. L. Piña, "Applying epoch-era analysis for homeowner selection of distributed generation power systems," Thesis, Massachusetts Institute of Technology, 2014.
- [61] M. A. Schaffner, M. W. Shihong, A. M. Ross, and D. H. Rhodes, "Enabling Design for Affordability: An Epoch-Era Analysis Approach," DTIC Document, 2013.
- [62] M. S. Wu, A. M. Ross, and D. H. Rhodes, "Design for Affordability in Complex Systems and Programs Using Tradespace-based Affordability Analysis," *Procedia Comput. Sci.*, vol. 28, pp. 828–837, 2014.
- [63] N. Ricci, A. M. Ross, and D. H. Rhodes, "A Generalized Options-based Approach to Mitigate Perturbations in a Maritime Security System-of-Systems," *Procedia Comput. Sci.*, vol. 16, pp. 718–727, 2013.
- [64] N. Ricci, D. H. Rhodes, and A. M. Ross, "Evolvability-Related Options in Military Systems of Systems," *Procedia Comput. Sci.*, vol. 28, pp. 314–321, 2014.
- [65] N. Ricci, M. A. Schaffner, A. M. Ross, D. H. Rhodes, and M. E. Fitzgerald, "Exploring Stakeholder Value Models via Interactive Visualization," *Procedia Comput. Sci.*, vol. 28, pp. 294–303, 2014.
- [66] M. E. Fitzgerald and A. M. Ross, "Controlling for Framing Effects in Multi-Stakeholder Tradespace Exploration," *Procedia Comput. Sci.*, vol. 28, pp. 412–421, 2014.
- [67] S. R. Goerger, A. M. Madni, and O. J. Eslinger, "Engineered Resilient Systems: A DoD Perspective," *Procedia Comput. Sci.*, vol. 28, pp. 865–872, 2014.
- [68] J. H. Saleh, "Weaving time into system architecture: new perspectives on flexibility, spacecraft design lifetime, and on-orbit servicing," Massachusetts Institute of Technology, 2001.
- [69] J. H. Saleh, D. E. Hastings, and D. J. Newman, "Flexibility in system design and implications for aerospace systems," *Acta Astronaut.*, vol. 53, no. 12, pp. 927–944, 2003.
- [70] J. H. Saleh, E. S. Lamassoure, D. E. Hastings, and D. J. Newman, "Flexibility and the value of on-orbit servicing: New customer-centric perspective," *J. Spacecr. Rockets*, vol. 40, no. 2, pp. 279–291, 2003.
- [71] J. H. Saleh, G. Mark, and N. C. Jordan, "Flexibility: a multi-disciplinary literature review and a research agenda for designing flexible engineering systems," *J. Eng. Des.*, vol. 20, no. 3, pp. 307–323, 2009.
- [72] R. Nilchiani, "Measuring space systems flexibility: a comprehensive six-element framework," Massachusetts Institute of Technology, 2005.
- [73] R. Nilchiani, C. Joppin, and D. Hastings, "Calculations of Flexibility in Space systems," *Summ. Rep.*, 2005.
- [74] R. Nilchiani and D. E. Hastings, "Measuring the Value of Flexibility in Space Systems: A Six-Element Framework," *Syst. Eng.*, vol. 10, no. 1, pp. 26–44, 2007.
- [75] C. Joppin, "On-orbit servicing for satellite upgrades," Thesis, Massachusetts Institute of Technology, 2004.
- [76] A. M. Long, "Framework for evaluating customer value and the feasibility of servicing architectures for on-orbit satellite servicing," Thesis, Massachusetts Institute of Technology, 2005.
- [77] M. G. Richards, "On-orbit serviceability of space system architectures," Massachusetts Institute of Technology, 2006.
- [78] O. L. De Weck, R. D. Neufville, and M. Chaize, "Staged deployment of communications satellite constellations in low earth orbit," *J. Aerosp. Comput. Inf. Commun.*, vol. 1, no. 3, pp. 119–136, 2004.
- [79] M. R. Silver and O. L. de Weck, "Time-expanded decision networks: A framework for designing evolvable complex systems," *Syst. Eng.*, vol. 10, no. 2, pp. 167–188, 2007.
- [80] B. G. Cameron, E. F. Crawley, G. Loureiro, and E. S. Rebentisch, "Value flow mapping: Using networks to inform stakeholder analysis," *Acta Astronaut.*, vol. 62, no. 4, pp. 324–333, 2008.

- [81] O. L. de Weck, A. M. Ross, and D. H. Rhodes, "Investigating Relationships and Semantic Sets amongst System Lifecycle Properties (Ilities)," in *3rd International Engineering Systems Symposium*, 2012.
- [82] Z. J. Bailey, A. Babuscia, L. Alkalai, and D. W. Miller, "A Trade Space Model for Distributed Lunar Surface Robotic Exploration."
- [83] F. Alibay, "Evaluation of Multi-Vehicle Architectures for the Exploration of Planetary Bodies in the Solar System," Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2014.
- [84] F. Alibay, V. R. Desaraju, J. E. Duda, and J. A. Hoffman, "Fractionated robotic architectures for planetary surface mobility systems," *Acta Astronaut.*, vol. 95, pp. 15–29, Feb. 2014.
- [85] D. Selva, B. G. Cameron, and E. F. Crawley, "Rule-based System Architecting of Earth Observing Systems: The Earth Science Decadal Survey."
- [86] D. Selva and E. F. Crawley, "VASSAR: Value assessment of system architectures using rules," in *Aerospace Conference, 2013 IEEE*, 2013, pp. 1–21.
- [87] D. Selva, "Executable Science Traceability Matrices." (unpublished), 2014.
- [88] A. Aliakbargolkar, "A framework for space systems architecting under stakeholder objectives ambiguity," Thesis, Massachusetts Institute of Technology, 2012.
- [89] C. Paredis, J. Aughenbaugh, R. Malak, and S. ReKuc, "Set-based design: a decisiontheoretic perspective," in *Proc. Frontiers in Design & Simulation Research & e006 Workshop*, 2006, pp. 1–25.
- [90] A. C. Ward, "A Theory of Quantitative Inference Applied to a Mechanical Design Compiler," Jan. 1989.
- [91] D. K. Sobek, A. C. Ward, and J. K. Liker, "Toyota's principles of set-based concurrent engineering," *Sloan Manage. Rev.*, vol. 40, no. 2, pp. 67–84, 1999.
- [92] R. J. Malak, J. M. Aughenbaugh, and C. J. Paredis, "Multi-attribute utility analysis in set-based conceptual design," *Comput.-Aided Des.*, vol. 41, no. 3, pp. 214–227, 2009.
- [93] R. Malak, "Using Parameterized Pareto Sets to Model Design Concepts," *J. Mech. Des.*, vol. 132, p. 041007, 2010.
- [94] R. Malak, "Using Support Vector Machines to Formalize the Valid Input Domain of Predictive Models for Systems Design Problems," *J. Mech. Des.*, vol. 132, p. 101001, 2010.
- [95] G. M. Stump, "Multi-Criteria Design Optimization Using Multi-Dimensional Data Visualization," The Pennsylvania State University, 2004.
- [96] D. E. Carlsen, "Assessment of User-Guided Visual Steering Commands During Trade Space Exploration," The Pennsylvania State University, 2008.
- [97] G. Stump, M. Yukish, J. D. Martin, and T. W. Simpson, "The ARL trade space visualizer: An engineering decision-making tool," in *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, 2004, vol. 30.
- [98] G. M. Stump, M. Yukish, T. W. Simpson, and J. J. O'Hara, "Trade space exploration of satellite datasets using a design by shopping paradigm," in *Aerospace Conference, 2004. Proceedings. 2004 IEEE*, 2004, vol. 6, pp. 3885–3895.
- [99] G. M. Stump, M. Yukish, T. W. Simpson, and L. Bennett, "Multidimensional visualization and its application to a design by shopping paradigm," in *9th AIAA/ISSMO Symposium on Multidisciplinary Analysis and Optimization*, 2002, pp. 4–6.
- [100] G. Stump, S. Lego, M. Yukish, T. W. Simpson, and J. A. Donndelinger, "Visual steering commands for trade space exploration: User-guided sampling with example," *J. Comput. Inf. Sci. Eng.*, vol. 9, no. 4, p. 044501, 2009.
- [101] D. D. Jordan, D. B. Spencer, T. W. Simpson, M. A. Yukish, and G. M. Stump, "Optimal Continuous-Thrust Trajectories Via Visual Trade Space Exploration," *J. Astronaut. Sci.*, vol. 57, no. 4, pp. 755–766, 2009.
- [102] A. C. Clarke, "Extra-terrestrial relays," *Wirel. World*, vol. 51, no. 10, pp. 305–308, 1945.
- [103] J. G. Walker, "Circular orbit patterns providing continuous whole earth coverage," DTIC Document, 1970.
- [104] J. G. Walker, "Continuous whole-earth coverage by circular-orbit satellite patterns," DTIC Document, 1977.
- [105] B. Evans, "Sirius Rising: Proton-M Ready to Launch Digital Radio Satellite Into Orbit," *AmericaSpace*. [Online]. Available: <http://www.americaspace.com/?p=43791>. [Accessed: 17-Apr-2015].
- [106] J. R. Wertz, "Orbit and Constellation Design - Selecting the Right Orbit," in *Space Mission Engineering: The New SMAD*, 1st ed., Hawthorne, CA: Microcosm Press, 2011, pp. 235–288.
- [107] D. C. Beste, "Design of satellite constellations for optimal continuous coverage," *Aerosp. Electron. Syst. IEEE Trans. On*, no. 3, pp. 466–473, 1978.
- [108] L. Rider, "Optimal Orbital Constellations For Global Viewing Of Targets Against A Space Background," *Opt. Eng.*, vol. 19, no. 2, pp. 192219–192219–, 1980.

- [109] L. Rider, "Analytic design of satellite constellations for zonal earth coverage using inclined circular orbits," *J. Astronaut. Sci.*, vol. 34, pp. 31–64, 1986.
- [110] A. T. Takano and B. G. Marchand, "Optimal Constellation Design for Space Based Situational Awareness Applications," in *AIAA/AAS Astrodynamics Specialists Conference, Girdwood, AK*, 2011.
- [111] A. T. Takano, "Numerical analysis and design of satellite constellations for above the horizon coverage," 2011.
- [112] B. G. Marchand and C. J. Kobel, "Above the Horizon Satellite Coverage with Dual-Altitude Band Constraints," *J. Spacecr. Rockets*, vol. 46, no. 4, pp. 845–857, 2009.
- [113] A. D. Biria, "Analytical approach to the design of optimal satellite constellations for space-based space situational awareness applications," 2012.
- [114] N. C. Deschamps, C. C. Grant, D. G. Foisy, R. E. Zee, A. F. Moffat, and W. W. Weiss, "The BRITE space telescope: Using a nanosatellite constellation to measure stellar variability in the most luminous stars," *Acta Astronaut.*, vol. 65, no. 5, pp. 643–650, 2009.
- [115] "BRITE-Constellation." [Online]. Available: <http://www.brite-constellation.at/>. [Accessed: 04-Feb-2015].
- [116] W. W. Weiss, A. F. J. Moffat, A. Schwarzenberg-Czerny, O. F. Koudelka, C. C. Grant, R. E. Zee, R. Kuschnig, S. Mochnecki, S. M. Rucinski, J. M. Matthews, P. Orleanski, A. Pamyatnykh, A. Pigulski, J. Alves, M. Guedel, G. Handler, G. A. Wade, and A. L. Scholtz, "BRITE-Constellation: Nanosatellites for Precision Photometry of Bright Stars," *ArXiv13095531 Astro-Ph*, Sep. 2013.
- [117] "BRITE Austria - eoPortal Directory - Satellite Missions." [Online]. Available: <https://directory.eoportal.org/web/eoportal/satellite-missions/b/brite-austria>. [Accessed: 17-Jan-2014].
- [118] T. L. Cline, "Developments in High-Precision Gamma-Ray Burst Source Studies," *Adv. Space Res.*, vol. 3, no. 4, pp. 175–183, 1983.
- [119] O. Brown and P. Eremenko, "Fractionated space architectures: a vision for responsive space," DTIC Document, 2006.
- [120] O. Brown and P. Eremenko, "Application of value-centric design to space architectures: the case of fractionated spacecraft," *AIAA Pap.*, vol. 7869, p. 2008, 2008.
- [121] O. Brown, P. Eremenko, and P. Collopy, "Value-Centric Design Methodologies for Fractionated Spacecraft: Progress Summary from Phase I of the DARPA System F6 Program," 2009.
- [122] J. M. Lafleur and J. H. Saleh, "GT-FAST: a point design tool for rapid fractionated spacecraft sizing and synthesis," in *AIAA SPACE 2009 Conference & Exposition*, 2009.
- [123] D. Maciuca, J. K. Chow, A. Siddiqi, O. de Weck, S. Alban, L. D. Dewell, A. S. Howell, J. M. Lieb, B. P. Mottinger, and J. Pandya, "A modular, high-fidelity tool to model the utility of fractionated space systems," *AIAA Pap.*, vol. 6765, p. 2009, 2009.
- [124] D. McCormick, B. Barrett, and M. Burnside-Clapp, "Analyzing fractionated satellite architectures using RAFTIMATE-a Boeing tool for value-centric design," *AIAA Pap.*, vol. 6767, p. 2009, 2009.
- [125] M. G. O'Neill, "Assessing the Impacts of Fractionation on Pointing-Intensive Spacecraft," SM Thesis," *Aeronaut. Astronaut. Dept Mass. Inst. Technol. Camb. MA*, 2009.
- [126] M. G. O'Neill and A. L. Weigel, "Assessing the impacts of fractionation on pointing-intensive spacecraft," *SM Thesis Aeronaut. Astronaut. Mass. Inst. Technol.*, 2009.
- [127] M. G. O'Neill, "Spacecraft Evaluation Tool Verification and Validation." SEARi, 2010.
- [128] M. G. O'Neill and A. L. Weigel, "Assessing Fractionated Spacecraft Value Propositions for Earth Imaging Space Missions," *J. Spacecr. Rockets*, vol. 48, no. 6, pp. 974–986, 2011.
- [129] M. G. O'Neill, H. Yue, S. Nag, P. Grogan, and O. de Weck, "Comparing and optimizing the DARPA system F6 program value-centric design methodologies," in *AIAA SPACE 2010 conference & exposition, Anaheim, CA*, 2010, vol. 30, pp. 391–412.
- [130] W. Ferster, "DARPA Cancels Formation-flying Satellite Demo | SpaceNews.com," *Space News*, 17-May-2013.
- [131] A. Aliakbargolkar, "Architecting Federated Satellite Systems for Successful Commercial Implementation," 2013.
- [132] A. Golkar, "Federated Satellite Systems: A Case Study on Sustainability Enhancement of Space Exploration Systems Architectures," presented at the 64th International Astronautical Congress, Beijing, China, vol. IAC-13-D3.4.
- [133] A. Golkar, "Federated Satellite Systems (FSS): A Vision Towards an Innovation in Space Systems Design." 2013.
- [134] A. Golkar, O. de Weck, M. Z. Win, and S. Seager, "Federated Satellite Systems (FSS): A Novel Approach to Space Systems Design." 2013.

- [135] A. Golkar and I. L. i Cruz, "The Federated Satellite Systems paradigm: Concept and business case evaluation," *Acta Astronaut.*, vol. 111, pp. 230–248, 2015.
- [136] A. Fox, R. Griffith, A. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, and I. Stoica, "Above the clouds: A Berkeley view of cloud computing," *Dept Electr. Eng Comput Sci. Univ. Calif. Berkeley Rep UCBECS*, vol. 28, 2009.
- [137] I. Lluch and A. Golkar, "Satellite-to-satellite coverage optimization approach for opportunistic inter-satellite links," in *Aerospace Conference, 2014 IEEE*, 2014, pp. 1–13.
- [138] I. Lluch and A. Golkar, "Resource Balancing for Federated Satellite Systems," presented at the AIAA SPACE 2014 Conference and Exposition, San Diego, CA, 2014.
- [139] S. Kirkpatrick, "Optimization by simulated annealing: Quantitative studies," *J. Stat. Phys.*, vol. 34, no. 5–6, pp. 975–986, 1984.
- [140] P. J. Van Laarhoven and E. H. Aarts, *Simulated annealing*. Springer, 1987.
- [141] E. Aarts and J. Korst, "Simulated annealing and Boltzmann machines," 1988.
- [142] R. C. Eberhart and J. Kennedy, "A new optimizer using particle swarm theory," in *Proceedings of the sixth international symposium on micro machine and human science*, 1995, vol. 1, pp. 39–43.
- [143] J. D. Schaffer, "Multiple objective optimization with vector evaluated genetic algorithms.," in *Proceedings of the 1st International Conference on Genetic Algorithms, Pittsburgh, PA, USA, July 1985*, 1985, pp. 93–100.
- [144] D. E. Goldberg and J. H. Holland, "Genetic algorithms and machine learning," *Mach. Learn.*, vol. 3, no. 2, pp. 95–99, 1988.
- [145] L. Davis and others, *Handbook of genetic algorithms*, vol. 115. Van Nostrand Reinhold New York, 1991.
- [146] D. M. Olsson and L. S. Nelson, "The Nelder-Mead simplex procedure for function minimization," *Technometrics*, vol. 17, no. 1, pp. 45–51, 1975.
- [147] J. C. Lagarias, J. A. Reeds, M. H. Wright, and P. E. Wright, "Convergence properties of the Nelder–Mead simplex method in low dimensions," *SIAM J. Optim.*, vol. 9, no. 1, pp. 112–147, 1998.
- [148] B. Corbin and T. Steiner, "Multidisciplinary System Design Optimization for a Distributed Solar Observation Constellation," presented at the Federated Satellite Systems Workshop, Skolkovo, Russia, 2015.
- [149] I. Das and J. Dennis, "Normal-Boundary Intersection: An Alternate Method for Generating Pareto Optimal Points in Multicriteria Optimization Problems.," DTIC Document, 1996.
- [150] I. Das and J. E. Dennis, "Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems," *SIAM J. Optim.*, vol. 8, no. 3, pp. 631–657, 1998.
- [151] I. Das, "On characterizing the 'knee' of the Pareto curve based on normal-boundary intersection," *Struct. Optim.*, vol. 18, no. 2–3, pp. 107–115, 1999.
- [152] A. Messac, A. Ismail-Yahaya, and C. A. Mattson, "The normalized normal constraint method for generating the Pareto frontier," *Struct. Multidiscip. Optim.*, vol. 25, no. 2, pp. 86–98, 2003.
- [153] I. Y. Kim and O. L. De Weck, "Adaptive weighted-sum method for bi-objective optimization: Pareto front generation," *Struct. Multidiscip. Optim.*, vol. 29, no. 2, pp. 149–158, 2005.
- [154] C. E. Gilchrist, "Spacecraft mass trade-offs versus radio-frequency power and antenna size at 8 GHz and 32 GHz," 1987.
- [155] K. Bedingfield, R. D. Leach, and M. B. Alexander, "Spacecraft System Failures and Anomalies Attributed to the Natural Space Environment," Aug. 1996.
- [156] "NASA - NSSDC - Spacecraft - PDMP Details." [Online]. Available: <http://nssdc.gsfc.nasa.gov/nmc/spacecraftTelemetry.do?id=1989-084B>. [Accessed: 18-Apr-2015].
- [157] "2014.B.1.4 Ultra-Compact Ka-Band Parabolic Deployable Antenna for CubeSats," *iCubeSat*. .
- [158] R. E. Freeland, G. D. Bilyeu, and G. R. Veal, "Development of flight hardware for a large, inflatable-deployable antenna experiment," *Acta Astronaut.*, vol. 38, no. 4–8, pp. 251–260, Feb. 1996.
- [159] R. E. Freeland, G. D. Bilyeu, G. R. Veal, M. D. Steiner, and D. E. Carson, "Large inflatable deployable antenna flight experiment results," *Acta Astronaut.*, vol. 41, no. 4–10, pp. 267–277, Aug. 1997.
- [160] A. Babuscia, B. Corbin, M. Knapp, R. Jensen-Clem, M. Van de Loo, and S. Seager, "Inflatable antenna for cubesats: Motivation for development and antenna design," *Acta Astronaut.*, vol. 91, pp. 322–332, Oct. 2013.
- [161] A. Babuscia, B. Corbin, R. Jensen-Clem, M. Knapp, I. Sergeev, M. Van de Loo, and S. Seager, "CommCube 1 and 2: A CubeSat series of missions to enhance communication capabilities for CubeSat," in *2013 IEEE Aerospace Conference*, 2013, pp. 1–19.
- [162] S. J. Dolinar and Yuen, "A Note on Deep Space Optical Communications Link Parameters," TDA Progress Report 42-70.

- [163] W. K. Marshall and B. D. Burk, "Received optical power calculations for optical communications link performance analysis," *TDA Prog. Rep.* 42, vol. 87, pp. 32–40, 1986.
- [164] M. Sans, Z. Sodnik, I. Zayer, and R. Daddato, "Design of the ESA Optical Ground Station for Participation in LLCD," 2012.
- [165] "LLCD | Goddard Space Flight Center." [Online]. Available: <http://esc.gsfc.nasa.gov/267/271.html>. [Accessed: 16-Jan-2014].
- [166] F. Dietrich and R. Davies, "Communications Architecture," in *Spacecraft Mission Analysis and Design*, 3rd ed., Microcosm Press and Kluwer Academic Publishers, 1999, pp. 533–586.
- [167] C.-C. Chen, "Link and System Design," in *Deep Space Optical Communications*, 2006, pp. 83–119.
- [168] A. Biswas and S. Piazzolla, "Deep-space optical communications downlink budget from Mars: System parameters," *IPN Prog. Rep.*, vol. 42, no. 154, pp. 0–1, 2003.
- [169] B. Moision and J. Hamkins, "Deep-space optical communications downlink budget: modulation and coding," *IPN Prog. Rep.*, vol. 42, no. 154, pp. 1–28, 2003.
- [170] H. Hemmati, *Deep space optical communications*, vol. 11. Wiley. com, 2006.
- [171] A. T. Ronny, "Link Performance Analysis of a Ship-to-Ship Laser Communication System," DTIC Document, 2012.
- [172] "Deep Space Network," NASA. [Online]. Available: <http://www.nasa.gov/content/deep-space-network/#.Utgx6bSqTYQ>. [Accessed: 16-Jan-2014].
- [173] NASA, "DSN Aperture Fee Tool." Jan-2010.
- [174] Z. Szajnarfarber and A. L. Weigel, "Towards an empirical measure of spacecraft innovation: The case of communication satellites," *Acta Astronaut.*, vol. 66, no. 7–8, pp. 1266–1279, Apr. 2010.
- [175] J. Mukherjee and B. Ramamurthy, "Communication Technologies and Architectures for Space Network and Interplanetary Internet," 2013.
- [176] S. Squyres and others, "Vision and voyages for planetary science in the decade 2013–2022," *Natl. Res. Counc. Publ.*, 2011.
- [177] R. A. Brealey, S. C. Myers, and F. Allen, *Principles of corporate finance*. McGraw-Hill/Irwin, 2008.
- [178] E. J. Mishan and E. Quah, *Cost benefit analysis*. Routledge, 2007.
- [179] M. K. Macauley, "The value of information: Measuring the contribution of space-derived earth science data to resource management," *Space Policy*, vol. 22, no. 4, pp. 274–282, 2006.
- [180] R. L. Keeney and H. Raiffa, *Decisions with multiple objectives: preferences and value trade-offs*. Cambridge university press, 1993.
- [181] D. Kahneman and A. Tversky, *Choices, values, and frames*. Cambridge University Press, 2000.
- [182] T. L. Saaty, "Decision making—the analytic hierarchy and network processes (AHP/ANP)," *J. Syst. Sci. Syst. Eng.*, vol. 13, no. 1, pp. 1–35, 2004.
- [183] T. L. Saaty, "The analytic hierarchy and analytic network processes for the measurement of intangible criteria and for decision-making," in *Multiple Criteria Decision Analysis: State of the Art Surveys*, Springer, 2005, pp. 345–405.
- [184] Y. Wang and Z. He, "A TOPSIS based robust optimization methodology for multivariable quality characteristics," in *IEEE International Conference on Service Operations and Logistics, and Informatics, 2008. IEEE/SOLI 2008*, 2008, vol. 2, pp. 2558–2561.
- [185] D. H. Rhodes and A. M. Ross, *Value-Driven Tradespace Exploration for System Design*. 2010.
- [186] M. Köppen, R. Vicente-Garcia, and B. Nickolay, "Fuzzy-pareto-dominance and its application in evolutionary multi-objective optimization," in *Evolutionary Multi-Criterion Optimization*, 2005, pp. 399–412.
- [187] *Trucklandia*. SEArI, 2013.
- [188] P. N. Johnson-Laird, *Mental models*. Harvard University Press, 1986.
- [189] P. N. Johnson-Laird, "Mental models and thought," *Camb. Handb. Think. Reason.*, pp. 185–208, 2005.
- [190] B. Fischhoff, "Value elicitation: is there anything in there?," *Am. Psychol.*, vol. 46, no. 8, p. 835, 1991.
- [191] K. Pedersen, J. Emblemsvag, R. Bailey, J. K. Allen, and F. Mistree, "Validating design methods and research: the validation square," in *Proceedings of the ASME Design Theory and Methodology Conference*, 2000.
- [192] A. M. Ross, D. H. Rhodes, and M. E. Fitzgerald, "Interactive Value Model Trading for Resilient Systems Decisions," *Procedia Comput. Sci.*, vol. 44, pp. 639–648, 2015.
- [193] J. H. Saleh, N. C. Jordan, and D. J. Newman, "Shifting the emphasis: From cost models to satellite utility or revenue models: The case for a value-centric mindset in space system design," *Acta Astronaut.*, vol. 61, no. 10, pp. 889–900, 2007.

- [194] P. D. Collopy and P. M. Hollingsworth, "Value-driven design," *J. Aircr.*, vol. 48, no. 3, pp. 749–759, 2011.
- [195] J. Brathwaite and J. H. Saleh, "On the concept of value and its importance to space systems design and acquisition," *AIAA Pap.*, vol. 7866, 2008.
- [196] P. Jorion, *Value at Risk, 3rd Ed.: The New Benchmark for Managing Financial Risk*. McGraw Hill Professional, 2006.
- [197] N. N. Das, D. Entekhabi, and E. G. Njoku, "An algorithm for merging SMAP radiometer and radar data for high-resolution soil-moisture retrieval," *Geosci. Remote Sens. IEEE Trans. On*, vol. 49, no. 5, pp. 1504–1512, 2011.
- [198] J. C. Castillo-Rogez, M. Pavone, I. A. Nesnas, and J. A. Hoffman, "Expected science return of spatially-extended in-situ exploration at small Solar system bodies," in *Aerospace Conference, 2012 IEEE*, 2012, pp. 1–15.
- [199] K. P. Yoon and C.-L. Hwang, *Multiple attribute decision making: an introduction*, vol. 104. Sage Publications, 1995.
- [200] S. Pugh, *Total design: integrated methods for successful product engineering*. Addison-Wesley, 1991.
- [201] R. Eggert, *Engineering design*. Pearson/Prentice Hall, 2005.
- [202] P. C. Fishburn, "Exceptional paper-lexicographic orders, utilities and decision rules: A survey," *Manag. Sci.*, vol. 20, no. 11, pp. 1442–1471, 1974.
- [203] R. W. Janke, "A lexicographic combination model of decision making by superintendents among multiattribute instructional program alternatives," *J. Sch. Psychol.*, vol. 12, no. 4, pp. 318–325, 1974.
- [204] G. A. Hazelrigg, "Bad design decisions: Why do we make them," *New Des. Paradig. Pasadena CA*, 2001.
- [205] G. A. Hazelrigg, *Systems engineering: an approach to information-based design*. Pearson College Division, 1996.
- [206] W. J. Larson and J. R. Wertz, "Space mission analysis and design," Microcosm, Inc., Torrance, CA (US), 1992.
- [207] J. R. Wertz, D. F. Everett, and J. J. Puschell, *Space mission engineering: the new SMAD*. Microcosm Press, 2011.
- [208] Analytical Graphics Inc., *Systems Toolkit (STK)*. AGI, 2012.
- [209] R. H. Battin, *An introduction to the mathematics and methods of astrodynamics*. Aiaa, 1999.
- [210] D. A. Vallado, *Fundamentals of astrodynamics and applications*, vol. 12. Springer Science & Business Media, 2001.
- [211] G. P. Sutton and O. Biblarz, *Rocket propulsion elements*. John Wiley & Sons, 2010.
- [212] R. G. Jahn, *Physics of electric propulsion*. Courier Corporation, 2012.
- [213] H. P. Hitchcock, "On Estimating the Drag Coefficient of Missiles," DTIC Document, 1951.
- [214] A. J. Eggbbbs, "Bodies of revolution having minimum drag at high supersonic airspeeds," 1957.
- [215] H. G. Heinrich, S. R. Hess, and G. Stumbris, "Drag Coefficients of Several Bodies of Revolution at Transonic and Supersonic Velocity," DTIC Document, 1964.
- [216] R. L. McCoy, "MC Drag-A Computer Program for Estimating the Drag Coefficients of Projectiles," DTIC Document, 1981.
- [217] I. Celmins, "Projectile Supersonic Drag Characteristics," DTIC Document, 1990.
- [218] T. J. McCabe and C. W. Butler, "Design complexity measurement and testing," *Commun. ACM*, vol. 32, no. 12, pp. 1415–1425, 1989.
- [219] B. Edmonds, "What is Complexity?-The philosophy of complexity per se with application to some examples in evolution," *Evol. Complex.*, 1995.
- [220] D. A. Bearden, "A complexity-based risk assessment of low-cost planetary missions: when is a mission too fast and too cheap?," *Acta Astronaut.*, vol. 52, no. 2, pp. 371–379, 2003.
- [221] P. Nguyen, N. Lozzi, W. Galang, S. Hu, A. Sjovald, P. Young, S. Book, and J. Schmilz, "Unmanned Space Vehicle Cost Model," *US Air Force Space Missile Syst. Cent. SMCFCM*, vol. 2430, pp. 90245–4687, 1994.
- [222] H. Habib-Agahi, G. Ball, and G. Fox, "NICM schedule & cost rules of thumb," in *AIAA space 2009 Conference & Exposition. AIAA, Pasadena*, 2009, pp. 6512–6512.
- [223] H. Habib-Agahi, J. Mrozinski, and G. Fox, "NASA instrument cost/schedule model Hamid Habib-Agahi," in *Aerospace Conference, 2011 IEEE*, 2011, pp. 1–19.
- [224] B. W. Boehm, R. Madachy, B. Steece, and others, *Software cost estimation with Cocomo II with Cdrom*. Prentice Hall PTR, 2000.
- [225] H. Apgar, "Cost Estimating," in *Space Mission Engineering: The New SMAD*, 2011.
- [226] Goddard Space Flight Center, "Near Earth Network (NEN) Users' Guide Revision 1." 15-Jan-2010.

- [227] B. Fox, K. Brancato, and B. Alkire, “Guidelines and metrics for assessing space system cost estimates,” DTIC Document, 2008.
- [228] W. Marshall and C. Boshuizen, “Planet Labs’ Remote Sensing Satellite System,” 2013.
- [229] P. La Tour, “Cost Improvement Estimates on Small Satellites.”
- [230] C. Foerste, F. Flechtner, R. Schmidt, U. Meyer, R. Stubenvoll, F. Barthelmes, M. Rothacher, R. Biancale, S. Bruinsma, and J. M. Lemoine, “A new high resolution global gravity field model from the combination of GRACE satellite mission and altimetry/gravimetry surface gravity data,” in *Geophysical Research Abstracts*, 2005, vol. 7, p. 04561.
- [231] M. T. Zuber, D. E. Smith, D. H. Lehman, T. L. Hoffman, S. W. Asmar, and M. M. Watkins, “Gravity Recovery and Interior Laboratory (GRAIL): Mapping the Lunar Interior from Crust to Core,” *Space Sci. Rev.*, vol. 178, no. 1, pp. 3–24, Sep. 2013.
- [232] S. A. Lane, T. W. Murphey, and M. Zatman, “Overview of the innovative space-based radar antenna technology program,” *J. Spacecr. Rockets*, vol. 48, no. 1, pp. 135–145, 2011.
- [233] D. O. Fulcoly, N. B. Shah, A. M. Ross, and D. H. Rhodes, “Exploiting Multidimensional Design of Experiments and Kriging Methods: An Application to a Satellite Radar System Tradespace and Orbital Transfer Vehicle Tradespace,” in *AIAA SPACE 2012 Conference & Exposition*, 2012.
- [234] N. Duffo, F. Torres, I. Corbella, V. González, S. Blanch, A. Camps, M. Vall-Ilossera, J. L. Alvarez, S. Ribo, and M. Martín-Neira, “Some results of the MIRAS-SMOS demonstrator campaigns,” in *Geoscience and Remote Sensing Symposium, 2007. IGARSS 2007. IEEE International*, 2007, pp. 3639–3642.
- [235] G. Fjeldbo, A. J. Kliore, and V. R. Eshleman, “The neutral atmosphere of Venus as studied with the Mariner V radio occultation experiments,” *Astron. J.*, vol. 76, p. 123, 1971.
- [236] G. F. Lindal, G. E. Wood, G. S. Levy, J. D. Anderson, D. N. Sweetnam, H. B. Hotz, B. J. Buckles, D. P. Holmes, P. E. Doms, V. R. Eshleman, and others, “The atmosphere of Jupiter: An analysis of the Voyager radio occultation measurements,” *J. Geophys. Res. Space Phys.* 1978–2012, vol. 86, no. A10, pp. 8721–8727, 1981.
- [237] G. F. Lindal, J. R. Lyons, D. N. Sweetnam, V. R. Eshleman, D. P. Hinson, and G. L. Tyler, “The atmosphere of Uranus: Results of radio occultation measurements with Voyager 2,” *J. Geophys. Res. Space Phys.* 1978–2012, vol. 92, no. A13, pp. 14987–15001, 1987.
- [238] E. R. Kursinski, G. A. Hajj, J. T. Schofield, R. P. Linfield, and K. R. Hardy, “Observing Earth’s atmosphere with radio occultation measurements using the Global Positioning System,” *J. Geophys. Res. Atmospheres* 1984–2012, vol. 102, no. D19, pp. 23429–23465, 1997.
- [239] M. B. McElroy and D. M. Hunten, “The ratio of deuterium to hydrogen in the Venus atmosphere,” *J. Geophys. Res.*, vol. 74, no. 7, pp. 1720–1739, 1969.
- [240] T. M. Donahue, J. H. Hoffman, R. R. Hodges, and A. J. Watson, “Venus was wet: a measurement of the ratio of deuterium to hydrogen,” *Science*, vol. 216, no. 4546, pp. 630–633, 1982.
- [241] J.-L. Bertaux and J. T. Clarke, “Deuterium content of the Venus atmosphere,” 1989.
- [242] J.-L. Bertaux, D. Nevejans, O. Korablev, E. Villard, E. Quémerais, E. Neefs, F. Montmessin, F. Leblanc, J.-P. Dubois, E. Dimarellis, and others, “SPICAV on Venus Express: Three spectrometers to study the global structure and composition of the Venus atmosphere,” *Planet. Space Sci.*, vol. 55, no. 12, pp. 1673–1700, 2007.
- [243] J. T. Clarke, J. Bertaux, C. J. Carveth, J. Chaufray, R. Gladstone, and N. Darling, “Coordinated Sounding Rocket, HST, and SPICAV Observations of Venus in Nov. 2013,” in *AGU Fall Meeting Abstracts*, 2013, vol. 1, p. 08.
- [244] J. W. Harbaugh, J. H. Doveton, and J. C. Davis, “Probability methods in oil exploration.[Book with glossary],” 1977.
- [245] A. A. Golkar, R. Keller, B. Robinson, O. L. de Weck, E. F. Crawley, and others, “A methodology for system architecting of offshore oil production systems,” in *Proceedings of the international design structure matrix conference (DSM’09). Greenville, South Carolina, USA*, 2009, pp. 343–356.
- [246] “Mars Exploration Rover Mission: All Spirit Updates.” [Online]. Available: http://mars.nasa.gov/mer/mission/status_spiritAll_2010.html#sol2478. [Accessed: 15-Jan-2015].
- [247] “Mars Exploration Rover Mission: All Opportunity Updates.” [Online]. Available: http://mars.nasa.gov/mer/mission/status_opportunityAll.html. [Accessed: 15-Jan-2015].
- [248] “Curiosity Rover Drive Log.” [Online]. Available: <http://curiositylog.com/>. [Accessed: 15-Jan-2015].
- [249] O. S. Card, *Ender’s Game*. Macmillan, 2006.
- [250] R. E. Young, “The Galileo probe: how it has changed our understanding of Jupiter,” *New Astron. Rev.*, vol. 47, no. 1, pp. 1–51, 2003.

- [251] M. H. Wong, P. R. Mahaffy, S. K. Atreya, H. B. Niemann, and T. C. Owen, “Updated Galileo probe mass spectrometer measurements of carbon, oxygen, nitrogen, and sulfur on Jupiter,” *Icarus*, vol. 171, no. 1, pp. 153–170, 2004.
- [252] Z. Manchester and M. A. Peck, “Stochastic Space Exploration with Microscale Spacecraft,” in *AIAA Guidance, Navigation and Control Conference*, 2011, pp. 8–11.
- [253] Z. Manchester, M. Peck, and A. Filo, “KickSat: A Crowd-Funded Mission to Demonstrate the World’s Smallest Spacecraft,” 2013.
- [254] C. Haskins, K. Forsberg, M. Krueger, D. Walden, and D. Hamelin, “Systems engineering handbook,” in *INCOSE*, 2006.
- [255] M. Zeleny and J. L. Cochrane, *Multiple criteria decision making*, vol. 25. McGraw-Hill New York, 1982.
- [256] R. L. Keeney, “Creativity in decision making with value-focused thinking,” *Sloan Manage. Rev.*, vol. 35, pp. 33–33, 1994.
- [257] R. L. Keeney, “Value-focused thinking: Identifying decision opportunities and creating alternatives,” *Eur. J. Oper. Res.*, vol. 92, no. 3, pp. 537–549, 1996.
- [258] R. L. Keeney and R. L. Keeney, *Value-focused thinking: A path to creative decisionmaking*. Harvard University Press, 2009.
- [259] M. E. Fitzgerald and A. M. Ross, “Guiding Cooperative Stakeholders to Compromise Solutions Using an Interactive Tradespace Exploration Process,” *Procedia Comput. Sci.*, vol. 16, pp. 343–352, 2013.
- [260] A. Ross, “Interactive Model-Centric Systems Engineering (IMCSE),” Massachusetts Institute of Technology, 2014.
- [261] P. Delquié, *ASSESS - An Interactive Program for Utility Assessment*. Fontainebleau, France: INSEAD, 2007.
- [262] P. Delquié, “Interpretation of the risk tolerance coefficient in terms of maximum acceptable loss,” *Decis. Anal.*, vol. 5, no. 1, pp. 5–9, 2008.
- [263] R. de Neufville and P. Delquie, “A model of the influence of certainty and probability ‘effects’ on the measurement of utility,” in *Risk, decision and rationality*, Springer, 1988, pp. 189–205.
- [264] R. Cohen, *The Fast and the Furious*. 2001.
- [265] L. Shaffer, “Are fossil fuels about to become extinct?,” *CNBC*. [Online]. Available: <http://www.cnbc.com/id/102591064>. [Accessed: 23-Apr-2015].
- [266] M. D. Curry and A. M. Ross, “Considerations for an Extended Framework for Interactive Epoch-Era Analysis,” *Procedia Comput. Sci.*, vol. 44, pp. 454–465, 2015.
- [267] C. M. Pong, M. W. Smith, M. W. Knutson, S. Lim, D. W. Miller, S. Seager, J. N. S. Villasenor, and S. D. Murphy, “One-arcsecond line-of-sight pointing control on exoplanetsat, a three-unit CubeSat,” 2011.
- [268] C. M. Pong, S. Lim, M. W. Smith, D. W. Miller, J. S. Villaseñor, and S. Seager, “Achieving high-precision pointing on ExoplanetSat: initial feasibility analysis,” in *SPIE Astronomical Telescopes and Instrumentation: Observational Frontiers of Astronomy for the New Decade*, 2010, p. 77311V–77311V.
- [269] M. Smith, S. Seager, C. Pong, M. Knutson, D. Miller, T. Henderson, S. Lim, T. Brady, M. Matranga, and S. Murphy, “The ExoplanetSat Mission to Detect Transiting Exoplanets with a CubeSat Space Telescope,” 2011.
- [270] M. W. Smith, S. Seager, C. M. Pong, M. W. Knutson, D. W. Miller, T. C. Henderson, S. Lim, T. M. Brady, M. J. Matranga, and S. D. Murphy, “ExoplanetSat Mission,” *Draper Technol. Dig.*, p. 20, 2012.
- [271] M. Knapp, M. Smith, C. Pong, J. Nash, and S. Seager, “ExoplanetSat: High Precision Photometry for Exoplanet Transit Detections in a 3U Cubesat,” in *Proceedings from the 63rd International Astronautical Congress*, Naples, Italy, 2012.
- [272] M. Guelman, A. Kogan, A. Kazarian, A. Livne, M. Orenstein, and H. Michalik, “Acquisition and pointing control for inter-satellite laser communications,” *Aerosp. Electron. Syst. IEEE Trans. On*, vol. 40, no. 4, pp. 1239–1248, 2004.
- [273] K. Warfield, “Exoplanet Satellite 2018-08 Study Final Report,” Jet Propulsion Laboratory, 02-Aug-2013.
- [274] S. Seager, *Exoplanets*. University of Arizona Press, 2010.
- [275] L. Gizon and S. K. Solanki, “Determining the inclination of the rotation axis of a Sun-like star,” *Astrophys. J.*, vol. 589, no. 2, p. 1009, 2003.
- [276] T. Arentoft, H. Kjeldsen, T. R. Bedding, M. Bazot, J. Christensen-Dalsgaard, T. H. Dall, C. Karoff, F. Carrier, P. Eggenberger, and D. Sosnowska, “A Multisite Campaign to Measure Solar-like Oscillations in Procyon. I. Observations, Data Reduction, and Slow Variations,” *Astrophys. J.*, vol. 687, no. 2, p. 1180, 2008.

- [277] G. Mallén-Ornelas, S. Seager, H. K. C. Yee, D. Minniti, M. D. Gladders, G. M. Mallén-Fullerton, and T. M. Brown, “The EXPLORE Project. I. A deep search for transiting extrasolar planets,” *Astrophys. J.*, vol. 582, no. 2, p. 1123, 2003.
- [278] T. G. Beatty and S. Seager, “Transit Probabilities for Stars with Stellar Inclination Constraints,” *Astrophys. J.*, vol. 712, no. 2, p. 1433, 2010.
- [279] S. Rucinski, K. Carroll, R. Kuschnig, J. Matthews, and P. Stibrany, “Most (microvariability & oscillations of stars) Canadian astronomical micro-satellite,” *Adv. Space Res.*, vol. 31, no. 2, pp. 371–373, 2003.
- [280] G. Walker, J. Matthews, R. Kuschnig, R. Johnson, S. Rucinski, J. Pazder, G. Burley, A. Walker, K. Skaret, R. Zee, S. Grocott, K. Carroll, P. Sinclair, D. Sturgeon, and J. Harron, “The MOST Asteroseismology Mission: Ultraprecise Photometry from Space,” *Publ. Astron. Soc. Pac.*, vol. 115, pp. 1023–1035, Sep. 2003.
- [281] J. F. Rowe, J. M. Matthews, S. Seager, E. Miller-Ricci, D. Sasselov, R. Kuschnig, D. B. Guenther, A. F. J. Moffat, S. M. Rucinski, G. A. H. Walker, and W. W. Weiss, “The Very Low Albedo of an Extrasolar Planet: MOST Space-based Photometry of HD 209458,” *Astrophys. J.*, vol. 689, pp. 1345–1353, Dec. 2008.
- [282] B. Croll, J. M. Matthews, J. F. Rowe, B. Gladman, E. Miller-Ricci, D. Sasselov, G. A. H. Walker, R. Kuschnig, D. N. C. Lin, D. B. Guenther, A. F. J. Moffat, S. M. Rucinski, and W. W. Weiss, “Looking for Super-Earths in the HD 189733 System: A Search for Transits in MOST Space-based Photometry,” *Astrophys. J.*, vol. 671, pp. 2129–2138, Dec. 2007.
- [283] A. Baglin, E. Michel, and M. Auvergne, “the CoRoT team 2006,” in *Proc. SOHO*, 2006, vol. 18.
- [284] T. Appourchaux, E. Michel, M. Auvergne, A. Baglin, T. Toutain, F. Baudin, O. Benomar, W. Chaplin, S. Deheuvels, and R. Samadi, “CoRoT sounds the stars: p-mode parameters of Sun-like oscillations on HD 49933,” *Astron. Astrophys.*, vol. 488, no. 2, p. 705, 2008.
- [285] R. A. García, S. Mathur, D. Salabert, J. Ballot, C. Régulo, T. S. Metcalfe, and A. Baglin, “CoRoT reveals a magnetic activity cycle in a Sun-like star,” *Science*, vol. 329, no. 5995, pp. 1032–1032, 2010.
- [286] J. Ballot, T. Appourchaux, T. Toutain, and M. Guittet, “On deriving p-mode parameters for inclined solar-like stars,” *ArXiv Prepr. ArXiv08030885*, 2008.
- [287] “Kepler: Target Stars.” [Online]. Available: <http://kepler.nasa.gov/Science/about/targetFieldOfView/>. [Accessed: 28-Feb-2015].
- [288] D. Huber, W. J. Chaplin, J. Christensen-Dalsgaard, R. L. Gilliland, H. Kjeldsen, L. A. Buchhave, D. A. Fischer, J. J. Lissauer, J. F. Rowe, and R. Sanchis-Ojeda, “Fundamental Properties of Kepler Planet-candidate Host Stars using Asteroseismology,” *Astrophys. J.*, vol. 767, no. 2, p. 127, 2013.
- [289] T. L. Campante, R. Handberg, S. Mathur, T. Appourchaux, T. R. Bedding, W. J. Chaplin, R. A. García, B. Mosser, O. Benomar, and A. Bonanno, “Asteroseismology from multi-month Kepler photometry: the evolved Sun-like stars KIC 10273246 and KIC 10920273,” *ArXiv Prepr. ArXiv11083807*, 2011.
- [290] T. Hirano, R. Sanchis-Ojeda, Y. Takeda, N. Narita, J. N. Winn, A. Taruya, and Y. Suto, “Measurements of stellar inclinations for kepler planet candidates,” *Astrophys. J.*, vol. 756, no. 1, p. 66, 2012.
- [291] J. H. J. De Bruijne, “Science performance of Gaia, ESA’s space-astrometry mission,” *Astrophys. Space Sci.*, vol. 341, no. 1, pp. 31–41, 2012.
- [292] G. R. Ricker, D. W. Latham, R. K. Vanderspek, K. A. Ennico, G. Bakos, T. M. Brown, A. J. Burgasser, D. Charbonneau, M. Clampin, L. D. Deming, and others, “Transiting Exoplanet Survey Satellite (TESS),” in *Bulletin of the American Astronomical Society*, 2010, vol. 42, p. 459.
- [293] G. R. Ricker, D. W. Latham, R. K. Vanderspek, K. A. Ennico, G. Bakos, T. M. Brown, A. J. Burgasser, D. Charbonneau, L. D. Deming, J. P. Doty, and others, “The Transiting Exoplanet Survey Satellite (TESS),” in *Bulletin of the American Astronomical Society*, 2009, vol. 41, p. 193.
- [294] “TESS Overview.” [Online]. Available: <http://tess.gsfc.nasa.gov/overview.html>. [Accessed: 28-Feb-2015].
- [295] C. Broeg, A. Fortier, D. Ehrenreich, Y. Alibert, W. Baumjohann, W. Benz, M. Deleuil, M. Gillon, A. Ivanov, R. Liseau, and others, “CHEOPS: A transit photometry mission for ESA’s small mission programme,” in *EPJ Web of Conferences*, 2013, vol. 47, p. 03005.
- [296] C. Catala, P. Consortium, and others, “PLATO: PLANetary Transits and Oscillations of stars,” *Exp. Astron.*, vol. 23, no. 1, pp. 329–356, 2009.
- [297] H. Rauer and C. Catala, “The PLATO mission,” *Proc. Int. Astron. Union*, vol. 6, no. S276, pp. 354–358, 2010.
- [298] J. A. S. correspondent and J. Amos, “European Space Agency picks Plato planet-hunting mission,” *BBC News*. [Online]. Available: <http://www.bbc.com/news/science-environment-26267918>. [Accessed: 09-Mar-2015].

- [299] P. B. Selding, “Planet-hunting Plato Wins Backing of European Scientists for \$1B Medium-class Mission,” *SpaceNews.com*. [Online]. Available: <http://spacenews.com/39335planet-hunting-plato-wins-backing-of-european-scientists-for-1b-medium/>. [Accessed: 09-Mar-2015].
- [300] C. M. Pong, “High-precision pointing and attitude estimation and control algorithms for hardware-constrained spacecraft,” Massachusetts Institute of Technology, 2014.
- [301] M. Knapp, R. Jensen-Clem, S. Seager, D. Miller, and M. W. Smith, “ExoplanetSat Constellation,” in *Bulletin of the American Astronomical Society*, 2011, vol. 1, p. 40601.
- [302] A. Cassan, D. Kubas, J.-P. Beaulieu, M. Dominik, K. Horne, J. Greenhill, J. Wambsganss, J. Menzies, A. Williams, U. G. ae Jørgensen, and others, “One or more bound planets per Milky Way star from microlensing observations,” *Nature*, vol. 481, no. 7380, pp. 167–169, 2012.
- [303] B. Corbin, *Movie 01 Slow*. <https://www.youtube.com/watch?v=3pkbzaOZXYM>, 2014.
- [304] B. Corbin, *Movie 02 Fast*. <https://www.youtube.com/watch?v=9FmPR0kCM6s>, 2014.
- [305] “GridSphere - File Exchange - MATLAB Central,” 26-Sep-2010. [Online]. Available: http://www.mathworks.com/matlabcentral/fileexchange/file_infos/28842-gridsphere. [Accessed: 02-Mar-2015].
- [306] M. W. Smith, C. M. Pong, and L. A. Hamlin, “Kickoff Meeting for ASTERIA: Arcsecond Space Telescope Enabling Research in Astrophysics (Phaeton Project),” Jet Propulsion Laboratory, 24-Oct-2014.
- [307] “Onset HOBO Data Logger.” [Online]. Available: <http://www.onsetcomp.com/>. [Accessed: 22-Mar-2015].
- [308] M. L. Kaiser, T. A. Kucera, J. M. Davila, O. S. Cyr, M. Guhathakurta, and E. Christian, “The STEREO mission: An introduction,” in *The STEREO Mission*, Springer, 2008, pp. 5–16.
- [309] Corbin, Benjamin, *HOBOCOP Network Goes Metal*. <https://www.youtube.com/watch?v=XEPyPrTxd5g>, 2013.
- [310] Corbin, Benjamin, *HOBOCOP Version 2.0*. <https://www.youtube.com/watch?v=zn19XOs91Tc>, 2014.
- [311] Corbin, Benjamin, *HOBOCOP Version 3.0*. <https://www.youtube.com/watch?v=bPjMo4cojh8>, 2015.
- [312] H. W. Babcock, “The Topology of the Sun’s Magnetic Field and the 22-YEAR Cycle,” *Astrophys. J.*, vol. 133, p. 572, Mar. 1961.
- [313] P. Zeeman and M. Bôcher, “Zeeman effect,” *Nature*, vol. 55, no. 1424, p. 347, 1897.
- [314] H. W. Babcock, “Zeeman Effect in Stellar Spectra,” *Astrophys. J.*, vol. 105, p. 105, 1947.
- [315] J. W. Belcher, “Alfvénic wave pressures and the solar wind,” *Astrophys. J.*, vol. 168, p. 509, 1971.
- [316] A. J. Hundhausen, “Coronal expansion and solar wind,” *Coronal Expans. Sol. Wind XII 238 Pp 101 Figs Springer-Verl. Berl. Heidelb. N. Y. Also Phys. Chem. Space Vol. 5*, vol. 1, 1972.
- [317] J. W. Belcher, A. J. Lazarus, R. L. McNutt, and G. S. Gordon, “Solar wind conditions in the outer heliosphere and the distance to the termination shock,” *J. Geophys. Res. Space Phys. 1978–2012*, vol. 98, no. A9, pp. 15177–15183, 1993.
- [318] N. V. Pogorelov, S. Borovikov, J. Heerikhuisen, I. Kryukov, G. P. Zank, M. C. Bedford, and T. K. Kim, “Multi-scale Modeling of the Solar Wind Interaction with the Interstellar Medium and Its Connection to Voyager Results,” in *AGU Fall Meeting Abstracts*, 2013, vol. 1, p. 1969.
- [319] E. N. Parker, “Interaction of the solar wind with the geomagnetic field,” *Phys. Fluids 1958-1988*, vol. 1, no. 3, pp. 171–187, 1958.
- [320] M. J. Owens, H. E. Spence, S. McGregor, W. J. Hughes, J. M. Quinn, C. N. Arge, P. Riley, J. Linker, and D. Odstrcil, “Metrics for solar wind prediction models: Comparison of empirical, hybrid, and physics-based schemes with 8 years of L1 observations,” *Space Weather*, vol. 6, no. 8, 2008.
- [321] D. N. Baker, G. Poh, D. Odstrcil, C. N. Arge, M. Benna, C. L. Johnson, H. Korth, D. J. Gershman, G. C. Ho, W. E. McClintock, and others, “Solar wind forcing at Mercury: WSA-ENLIL model results,” *J. Geophys. Res. Space Phys.*, vol. 118, no. 1, pp. 45–57, 2013.
- [322] “Heliospheric current sheet,” *Wikipedia, the free encyclopedia*. 05-Feb-2015.
- [323] E. J. Smith, “The heliospheric current sheet,” *J. Geophys. Res. Space Phys.*, vol. 106, no. A8, pp. 15819–15831, Aug. 2001.
- [324] “Missions to study the Sun.” [Online]. Available: <http://www.planetary.org/explore/space-topics/space-missions/missions-to-study-the-sun.html>. [Accessed: 07-Mar-2015].
- [325] K.-P. Wenzel, R. G. Marsden, D. E. Page, and E. J. Smith, “Ulysses: The first high-latitude heliospheric mission,” *Adv. Space Res.*, vol. 9, no. 4, pp. 25–29, 1989.
- [326] A. Balogh, T. J. Beek, R. J. Forsyth, P. C. Hedgecock, R. J. Marquedant, E. J. Smith, D. J. Southwood, and B. T. Tsurutani, “The magnetic field investigation on the ULYSSES mission-Instrumentation and preliminary scientific results,” *Astron. Astrophys. Suppl. Ser.*, vol. 92, pp. 221–236, 1992.

- [327] K. P. Wenzel, R. G. Marsden, D. E. Page, and E. J. Smith, “The ULYSSES mission,” *Astron. Astrophys. Suppl. Ser.*, vol. 92, p. 207, 1992.
- [328] A. Balogh, R. G. Marsden, and E. J. Smith, *The heliosphere near solar minimum: The Ulysses perspective*. Springer Science & Business Media, 2001.
- [329] V. Domingo, B. Fleck, and A. I. Poland, “The SOHO mission: an overview,” *Sol. Phys.*, vol. 162, no. 1–2, pp. 1–37, 1995.
- [330] B. Fleck, *The SOHO mission*. Springer, 1995.
- [331] J.-P. Delaboudiniere, G. E. Artzner, J. Brunaud, A. H. Gabriel, J. F. Hochedez, F. Millier, X. Y. Song, B. Au, K. P. Dere, R. A. Howard, and others, *EIT: extreme-ultraviolet imaging telescope for the SOHO mission*. Springer, 1995.
- [332] E. C. Stone, A. M. Frandsen, R. A. Mewaldt, E. R. Christian, D. Margolies, J. F. Ormes, and F. Snow, “The advanced composition explorer,” in *The Advanced Composition Explorer Mission*, Springer, 1998, pp. 1–22.
- [333] D. J. McComas, S. J. Bame, P. Barker, W. C. Feldman, J. L. Phillips, P. Riley, and J. W. Griffiee, “Solar wind electron proton alpha monitor (SWEPAM) for the Advanced Composition Explorer,” in *The Advanced Composition Explorer Mission*, Springer, 1998, pp. 563–612.
- [334] R. P. Lin, G. J. Hurford, N. W. Madden, B. R. Dennis, C. J. Crannell, G. D. Holman, R. Ramaty, T. T. von Roseninge, A. Zehnder, H. F. van Beek, and others, “High-energy solar spectroscopic imager (HESSI) small explorer mission for the next (2000) solar maximum,” in *SPIE’s International Symposium on Optical Science, Engineering, and Instrumentation*, 1998, pp. 2–12.
- [335] R. P. Lin, B. R. Dennis, G. J. Hurford, D. M. Smith, A. Zehnder, P. R. Harvey, D. W. Curtis, D. Pankow, P. Turin, M. Bester, and others, “The Reuven Ramaty high-energy solar spectroscopic imager (RHESSI),” in *The Reuven Ramaty High-Energy Solar Spectroscopic Imager (RHESSI)*, Springer, 2003, pp. 3–32.
- [336] T. Kosugi, K. Matsuzaki, T. Sakao, T. Shimizu, Y. Sone, S. Tachikawa, T. Hashimoto, K. Minesugi, A. Ohnishi, T. Yamada, and others, “The Hinode (Solar-B) mission: an overview,” *Sol. Phys.*, vol. 243, no. 1, pp. 3–17, 2007.
- [337] J.-L. Bougeret, K. Goetz, M. L. Kaiser, S. D. Bale, P. J. Kellogg, M. Maksimovic, N. Monge, S. J. Monson, P. L. Astier, S. Davy, and others, “S/WAVES: The radio and plasma wave investigation on the STEREO mission,” *Space Sci. Rev.*, vol. 136, no. 1–4, pp. 487–528, 2008.
- [338] B. E. Wood, R. A. Howard, S. P. Plunkett, and D. G. Socker, “Comprehensive Observations of a Solar Minimum Coronal Mass Ejection with the Solar Terrestrial Relations Observatory,” *Astrophys. J.*, vol. 694, no. 2, p. 707, Apr. 2009.
- [339] C. J. Eyles, R. A. Harrison, C. J. Davis, N. R. Waltham, B. M. Shaughnessy, H. C. A. Mapson-Menard, D. Bewsher, S. R. Crothers, J. A. Davies, G. M. Simnett, and others, “The heliospheric imagers onboard the STEREO mission,” *Sol. Phys.*, vol. 254, no. 2, pp. 387–445, 2009.
- [340] I. G. Richardson, T. T. von Roseninge, H. V. Cane, E. R. Christian, C. M. S. Cohen, A. W. Labrador, R. A. Leske, R. A. Mewaldt, M. E. Wiedenbeck, and E. C. Stone, “> 25 MeV Proton Events Observed by the High Energy Telescopes on the STEREO A and B Spacecraft and/or at Earth During the First Seven Years of the STEREO Mission,” *Sol. Phys.*, vol. 289, no. 8, pp. 3059–3107, 2014.
- [341] D. J. McComas, H. O. Funsten, S. A. Fuselier, W. S. Lewis, E. Möbius, and N. A. Schwadron, “IBEX observations of heliospheric energetic neutral atoms: Current understanding and future directions,” *Geophys. Res. Lett.*, vol. 38, no. 18, p. L18101, Sep. 2011.
- [342] D. J. McComas, D. Alexashov, M. Bzowski, H. Fahr, J. Heerikhuisen, V. Izmodenov, M. A. Lee, E. Möbius, N. Pogorelov, N. A. Schwadron, and G. P. Zank, “The Heliosphere’s Interstellar Interaction: No Bow Shock,” *Science*, vol. 336, no. 6086, pp. 1291–1293, Jun. 2012.
- [343] K. Scherer and H. Fichtner, “The Return of the Bow Shock,” *Astrophys. J.*, vol. 782, no. 1, p. 25, Feb. 2014.
- [344] D. J. McComas, M. A. Dayeh, H. O. Funsten, G. Livadiotis, and N. A. Schwadron, “The heliotail revealed by the Interstellar Boundary Explorer,” *Astrophys. J.*, vol. 771, no. 2, p. 77, 2013.
- [345] J. Schou, P. H. Scherrer, R. I. Bush, R. Wachter, S. Couvidat, M. C. Rabello-Soares, R. S. Bogart, J. T. Hoeksema, Y. Liu, T. L. D. Jr, D. J. Akin, B. A. Allard, J. W. Miles, R. Rairden, R. A. Shine, T. D. Tarbell, A. M. Title, C. J. Wolfson, D. F. Elmore, A. A. Norton, and S. Tomczyk, “Design and Ground Calibration of the Helioseismic and Magnetic Imager (HMI) Instrument on the Solar Dynamics Observatory (SDO),” in *The Solar Dynamics Observatory*, P. Chamberlin, W. D. Pesnell, and B. Thompson, Eds. Springer US, 2011, pp. 229–259.
- [346] W. D. Pesnell, B. J. Thompson, and P. C. Chamberlin, *The solar dynamics observatory (SDO)*. Springer, 2012.

- [347] P. H. Scherrer, J. Schou, R. I. Bush, A. G. Kosovichev, R. S. Bogart, J. T. Hoeksema, Y. Liu, T. L. Duvall Jr, J. Zhao, C. J. Schrijver, and others, “The helioseismic and magnetic imager (HMI) investigation for the solar dynamics observatory (SDO),” in *The Solar Dynamics Observatory*, Springer, 2012, pp. 207–227.
- [348] G. Thuillier, S. Dewitte, and W. Schmutz, “The PICARD mission: scientific objectives and status of development,” in *36th COSPAR Scientific Assembly*, 2006, vol. 36, p. 170.
- [349] G. Thuillier, S. Dewitte, W. Schmutz, and others, “Simultaneous measurement of the total solar irradiance and solar diameter by the PICARD mission,” *Adv. Space Res.*, vol. 38, no. 8, pp. 1792–1806, 2006.
- [350] J. Burt and B. Smith, “Deep Space Climate Observatory: The DSCOVR mission,” in *Aerospace Conference, 2012 IEEE*, 2012, pp. 1–13.
- [351] M. Cash, “The DSCOVR Solar Wind Mission: Algorithm Development to Enhance Space Weather Forecasting,” in *2014 AGU Chapman Conference on Low-Frequency Waves in Space Plasmas*, 2014.
- [352] A. Gandorfer, S. K. Solanki, J. Woch, V. M. Pillet, A. Á. Herrero, and T. Appourchaux, “The Solar Orbiter Mission and its Polarimetric and Helioseismic Imager (SO/PHI),” *J. Phys. Conf. Ser.*, vol. 271, no. 1, p. 012086, Jan. 2011.
- [353] S. K. Solanki, J. C. del T. Iniesta, J. Woch, A. Gandorfer, J. Hirzberger, W. Schmidt, T. Appourchaux, A. Alvarez-Herrero, and for the S. team, “The Polarimetric and Helioseismic Imager for Solar Orbiter: SO/PHI,” *ArXiv150203368 Astro-Ph*, Feb. 2015.
- [354] Y. Guo, “Solar Probe Plus: Mission design challenges and trades,” *Acta Astronaut.*, vol. 67, no. 9, pp. 1063–1072, 2010.
- [355] S. P. Plus, “Report of the Science and Technology definition Team,” *PRE-Publ. VERSION Febr*, vol. 14, 2008.
- [356] S. Probe, “Mission Engineering Study Report, prepared by The Johns Hopkins University Applied Physics Laboratory under NASA Contract NNN06AA01C,” *Laurel MD March 10 2008*.
- [357] C. P. Escoubet, M. Fehringer, and M. Goldstein, “The Cluster mission,” *Ann. Geophys.*, vol. 19, no. 10/12, pp. 1197–1200, 2001.
- [358] J. M. Weygand, M. G. Kivelson, K. K. Khurana, H. K. Schwarzl, S. M. Thompson, R. L. McPherron, A. Balogh, L. M. Kistler, M. L. Goldstein, J. Borovsky, and D. A. Roberts, “Plasma sheet turbulence observed by Cluster II,” *J. Geophys. Res. Space Phys.*, vol. 110, no. A1, p. A01205, Jan. 2005.
- [359] V. Angelopoulos, *The THEMIS mission*. Springer, 2009.
- [360] E. Friis-Christensen, H. Lühr, and G. Hulot, “Swarm: A constellation to study the Earth’s magnetic field,” *Earth Planets Space*, vol. 58, no. 4, pp. 351–358, 2006.
- [361] E. Friis-Christensen, H. Lühr, D. Knudsen, and R. Haagsmans, “Swarm—an Earth observation mission investigating geospace,” *Adv. Space Res.*, vol. 41, no. 1, pp. 210–216, 2008.
- [362] S. Macmillan and N. Olsen, “Observatory data and the Swarm mission,” *Earth Planets Space*, vol. 65, no. 11, pp. 1355–1362, 2013.
- [363] A. S. Sharma and S. A. Curtis, “Magnetospheric multiscale mission,” in *Nonequilibrium Phenomena in Plasmas*, Springer, 2005, pp. 179–195.
- [364] C. Zeitlin, D. M. Hassler, F. A. Cucinotta, B. Ehresmann, R. F. Wimmer-Schweingruber, D. E. Brinza, S. Kang, G. Weigle, S. Böttcher, E. Böhm, S. Burmeister, J. Guo, J. Köhler, C. Martin, A. Posner, S. Rafkin, and G. Reitz, “Measurements of Energetic Particle Radiation in Transit to Mars on the Mars Science Laboratory,” *Science*, vol. 340, no. 6136, pp. 1080–1084, May 2013.
- [365] A. J. Tylka, J. H. Adams, P. R. Boberg, B. Brownstein, W. F. Dietrich, E. O. Flueckiger, E. L. Petersen, M. A. Shea, D. F. Smart, and E. C. Smith, “CREME96: A revision of the cosmic ray effects on microelectronics code,” *Nucl. Sci. IEEE Trans. On*, vol. 44, no. 6, pp. 2150–2160, 1997.
- [366] R. A. Weller, M. H. Mendenhall, R. A. Reed, R. D. Schrimpf, K. M. Warren, B. D. Sierawski, and L. W. Massengill, “Monte Carlo simulation of single event effects,” *Nucl. Sci. IEEE Trans. On*, vol. 57, no. 4, pp. 1726–1746, 2010.
- [367] M. H. Mendenhall and R. A. Weller, “A probability-conserving cross-section biasing mechanism for variance reduction in Monte Carlo particle transport calculations,” *Nucl. Instrum. Methods Phys. Res. Sect. Accel. Spectrometers Detect. Assoc. Equip.*, vol. 667, pp. 38–43, 2012.
- [368] S. Agostinelli, J. Allison, K. al Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, and others, “GEANT4—a simulation toolkit,” *Nucl. Instrum. Methods Phys. Res. Sect. Accel. Spectrometers Detect. Assoc. Equip.*, vol. 506, no. 3, pp. 250–303, 2003.
- [369] D. Heynderickx, B. Quaghebeur, E. Speelman, and E. J. Daly, “ESA’s Space Environment Information System (SPENVIS): a WWW interface to models of the space environment and its effects,” *Proc AIAA*, vol. 371, 2000.

- [370] A. Panetti, "Thermal," in *Spacecraft Mission Analysis and Design*, 3rd ed., Microcosm Press and Kluwer Academic Publishers, 1999.
- [371] J. E. Oberright, "Spacecraft Thermal Control," in *Space Mission Engineering: The New SMAD*, 2011.
- [372] R. J. Strangeway, C. T. Russell, V. Angelopoulos, P. J. Chi, D. R. Pierce, K. M. Rowe, and M.-C. Vassal, "Digital Fluxgate Magnetometer Optimized for Lunar Networks," *LPI Contrib.*, vol. 1530, p. 3026, 2010.
- [373] B. Weiss, "Magnetic Measurements of Primitive Bodies," presented at the International Primitive Body Exploration Working Group, Nice, France, 2013.
- [374] "Magnetometer." [Online]. Available: <http://www.ssbv.com/ProductDatashets/page39/page31/index.html>. [Accessed: 21-Mar-2015].
- [375] S. R. Burr, "The design and implementation of the Dynamic Ionosphere Cubesat Experiment (DICE) science instruments," Citeseer, 2013.
- [376] A. W. Case, J. C. Kasper, P. S. Daigneau, D. Caldwell, M. Freeman, T. Gauron, B. A. Maruca, J. Bookbinder, K. E. Korreck, J. W. Cirtain, and others, "Designing a sun-pointing Faraday cup for solar probe plus," in *SOLAR WIND 13: Proceedings of the Thirteenth International Solar Wind Conference*, 2013, vol. 1539, pp. 458–461.
- [377] Q. Schiller and A. Mahendrakumar, "REPTile: A miniaturized detector for a CubeSat mission to measure relativistic particles in near-earth space," 2010.
- [378] L. Blum and Q. Schiller, "Characterization and testing of an energetic particle telescope for a CubeSat platform," 2012.
- [379] D. L. Glaser, J. S. Halekas, P. Turin, D. W. Curtis, D. E. Larson, S. E. McBride, and R. P. Lin, "STEIN (SupraThermal Electrons, Ions and Neutrals), A New Particle Detection Instrument for Space Weather Research with CubeSats," 2009.
- [380] R. P. Lin, G. K. Parks, J. S. Halekas, D. E. Larson, J. P. Eastwood, L. Wang, J. G. Sample, T. S. Horbury, E. C. Roelof, D. Lee, and others, "Cinema (cubesat for ion, neutral, electron, magnetic fields)," in *AGU Fall Meeting Abstracts*, 2009, vol. 1, p. 09.
- [381] G. Crowley, C. Fish, C. Swenson, R. Burt, E. Stromberg, T. Neilsen, S. Burr, A. Barjatya, G. Bust, and M. Larsen, "Dynamic ionosphere cubesat experiment (dice)," 2011.
- [382] C. Fish, C. Swenson, T. Neilsen, B. Bingham, J. Gunther, E. Stromberg, S. Burr, R. Burt, M. Whitely, G. Crowley, and others, "DICE Mission Design, Development, and Implementation: Success and Challenges," 2012.
- [383] A. M. Amelushkin, V. V. Bogomolov, V. V. Benghin, G. K. Garipov, E. S. Gorbvskoy, B. Grossan, P. A. Klimov, B. A. Khrenov, J. Lee, V. M. Lipunov, and others, "Space experiments on-board of Lomonosov mission to study gamma-ray bursts and UHECRs," *EAS Publ. Ser.*, vol. 61, pp. 545–552, 2013.
- [384] A. M. Amelushkin, V. V. Benghin, V. V. Bogomolov, G. K. Garipov, E. S. Gorbvskoy, B. Grossan, A. F. Iyudin, B. A. Khrenov, P. A. Kilmov, J. Lee, and others, "Space Mission Lomonosov on Study Gamma-Ray Bursts and UHECR." 2014. [Online]. Available: <http://elfin.igpp.ucla.edu/wp-content/uploads/2013/08/002-ELFINOverview2014-00-Q.pdf>. [Accessed: 21-Mar-2015].
- [385] A. A. Norton, J. P. Graham, R. K. Ulrich, J. Schou, S. Tomczyk, Y. Liu, B. W. Lites, A. L. Ariste, R. I. Bush, H. Socas-Navarro, and others, "Spectral line selection for HMI: a comparison of Fe I 6173 \AA and Ni I 6768 \AA," *Sol. Phys.*, vol. 239, no. 1–2, pp. 69–91, 2006.
- [386] B. Fleck, S. Couvidat, and T. Straus, "On the formation height of the SDO/HMI Fe 6173 \AA Doppler signal," *Sol. Phys.*, vol. 271, no. 1–2, pp. 27–40, 2011.
- [387] ESA, "Solar Orbiter: Exploring the Sun-heliosphere connection." European Space Agency, Jul-2011.
- [388] J. Hirzberger, "SO/PHI: The Solar Orbiter Polarimetric and Helioseismic Imager," presented at the MPS Solar Group Seminar, 28-Oct-2014.
- [389] V. M. Pillet, J. C. del T. Iniesta, A. Alvarez-Herrero, V. Domingo, J. A. Bonet, L. G. Fernandez, A. L. Jimenez, C. Pastor, J. L. G. Blesa, P. Mellado, J. Piqueras, B. Aparicio, M. Balaguer, E. Ballesteros, T. Belenguier, L. R. B. Rubio, T. Berkefeld, M. Collados, W. Deutsch, A. Feller, F. Girela, B. Grauf, R. L. Heredero, M. Herranz, J. M. Jeronimo, H. Laguna, R. Meller, M. Menendez, R. Morales, D. O. Suarez, G. Ramos, M. Reina, J. L. Ramos, P. Rodriguez, A. Sanchez, N. Uribe-Patarroyo, P. Barthol, A. Gandorfer, M. Knoelker, W. Schmidt, S. K. Solanki, and S. V. Dominguez, "The Imaging Magnetograph eXperiment (IMaX) for the Sunrise balloon-borne solar observatory," *ArXiv10091095 Astro-Ph*, Sep. 2010.
- [390] M. H. Hecht, "Reliability," in *Space Mission Engineering: The New SMAD*, 1st ed., Microcosm Press, 2011, pp. 753–767.

- [392] Corbin, Benjamin, *Heliospheric Current Sheet Model Mk2*. https://www.youtube.com/watch?v=3b_iS3JrKw, 2015.
- [393] Corbin, Benjamin, *HOBOCOP Version 4.3*. <https://www.youtube.com/watch?v=hRZgwp0CuXw>, 2015.
- [394] R. P. Binzel, “The Torino impact hazard scale,” *Planet. Space Sci.*, vol. 48, no. 4, pp. 297–303, 2000.
- [395] C. to R. N.-E. O. Surveys and H. M. Strategies, *Defending Planet Earth: Near-Earth Object Surveys and Hazard Mitigation Strategies: Final Report*. National Academies Press, 2010.
- [396] B. Corbin, “Implementing Advanced Technologies and Models to Reduce Uncertainty in a Global, Cost-Effective Asteroid Mitigation System,” in *Proceedings from the 61st International Astronautical Congress*, Prague, CZ, 2010, vol. IAC-10–E3.1B.5.
- [397] J. S. Lewis, “Mining the sky: untold riches from the asteroids, comets, and planets,” *Read. Mass Addison-Wesley Pub Co C1996*, vol. 1, 1996.
- [398] G. K. O’neill and others, *The high frontier*. Corgi Books London, 1978.
- [399] B. Hays, “NASA contracts two firms to work on asteroid mining,” *UPI*, 24-Nov-2014. [Online]. Available: http://www.upi.com/Science_News/2014/11/24/NASA-contracts-two-firms-to-work-on-asteroid-mining/5301416856690/. [Accessed: 07-May-2015].
- [400] H. Y. McSween, *Meteorites and their parent planets*. Cambridge University Press, 1999.
- [401] D. Vokrouhlický, A. Milani, and S. R. Chesley, “Yarkovsky effect on small near-Earth asteroids: Mathematical formulation and examples,” *Icarus*, vol. 148, no. 1, pp. 118–138, 2000.
- [402] W. F. Bottke, D. Vokrouhlický, M. Brož, D. Nesvorný, and A. Morbidelli, “Dynamical spreading of asteroid families by the Yarkovsky effect,” *Science*, vol. 294, no. 5547, pp. 1693–1696, 2001.
- [403] S. R. Chesley, S. J. Ostro, D. Vokrouhlický, D. Čapek, J. D. Giorgini, M. C. Nolan, J.-L. Margot, A. A. Hine, L. A. Benner, and A. B. Chamberlin, “Direct detection of the Yarkovsky effect by radar ranging to asteroid 6489 Golevka,” *Science*, vol. 302, no. 5651, pp. 1739–1742, 2003.
- [404] W. F. Bottke Jr, D. Vokrouhlický, D. P. Rubincam, and D. Nesvorný, “The Yarkovsky and YORP effects: Implications for asteroid dynamics,” *Annu Rev Earth Planet Sci*, vol. 34, pp. 157–191, 2006.
- [405] “The International Astronomical Union Minor Planet Center Home Page,” *The International Astronomical Union Minor Planet Center*. [Online]. Available: <http://www.minorplanetcenter.net/>. [Accessed: 09-Jan-2014].
- [406] M. L. Nelson, D. T. Britt, and L. A. Lebofsky, *Review of asteroid compositions*. the University of Arizona Press, Tucson and London, 1993.
- [407] O. R. Norton, *The Cambridge encyclopedia of meteorites*, vol. 1. 2002.
- [408] P. J. Hevey and I. S. Sanders, “A model for planetesimal meltdown by ^{26}Al and its implications for meteorite parent bodies,” *Meteorit. Planet. Sci.*, vol. 41, no. 1, pp. 95–106, 2006.
- [409] J. E. Chambers, “Planetary accretion in the inner Solar System,” *Earth Planet. Sci. Lett.*, vol. 223, no. 3, pp. 241–252, 2004.
- [410] L. T. Elkins-Tanton, “Linked magma ocean solidification and atmospheric growth for Earth and Mars,” *Earth Planet. Sci. Lett.*, vol. 271, no. 1, pp. 181–191, 2008.
- [411] K. Lodders, “Solar system abundances and condensation temperatures of the elements,” *Astrophys. J.*, vol. 591, no. 2, p. 1220, 2003.
- [412] T. H. Prettyman, “Chapter 41 - Remote Chemical Sensing Using Nuclear Spectroscopy,” in *Encyclopedia of the Solar System (Second Edition)*, L.-A. M. R. W. V. Johnson, Ed. San Diego: Academic Press, 2007, pp. 765–786.
- [413] W. C. Feldman, S. Maurice, D. J. Lawrence, R. C. Little, S. L. Lawson, O. Gasnault, R. C. Wiens, B. L. Barraclough, R. C. Elphic, T. H. Prettyman, and others, “Evidence for water ice near the lunar poles,” *J. Geophys. Res. Planets 1991–2012*, vol. 106, no. E10, pp. 23231–23251, 2001.
- [414] D. J. Lawrence, W. C. Feldman, J. O. Goldsten, S. Maurice, P. N. Peplowski, B. J. Anderson, D. Bazell, R. L. McNutt, L. R. Nittler, T. H. Prettyman, and others, “Evidence for water ice near Mercury’s north pole from MESSENGER Neutron Spectrometer measurements,” *Science*, vol. 339, no. 6117, pp. 292–296, 2013.
- [415] W. V. Boynton, W. C. Feldman, S. W. Squyres, T. H. Prettyman, J. Brückner, L. G. Evans, R. C. Reedy, R. Starr, J. R. Arnold, D. M. Drake, and others, “Distribution of hydrogen in the near surface of Mars: Evidence for subsurface ice deposits,” *Science*, vol. 297, no. 5578, pp. 81–85, 2002.
- [416] T. H. Prettyman, D. W. Mittlefehldt, N. Yamashita, D. J. Lawrence, A. W. Beck, W. C. Feldman, T. J. McCoy, H. Y. McSween, M. J. Toplis, T. N. Titus, P. Tricarico, R. C. Reedy, J. S. Hendricks, O. Forni, L. L. Corre, J.-Y. Li, H. Mizzon, V. Reddy, C. A. Raymond, and C. T. Russell, “Elemental Mapping by Dawn Reveals Exogenic H in Vesta’s Regolith,” *Science*, vol. 338, no. 6104, pp. 242–246, Oct. 2012.

- [417] T. H. Prettyman, W. C. Feldman, F. P. Ameduri, B. L. Barraclough, E. W. Cascio, K. R. Fuller, H. O. Funsten, D. J. Lawrence, G. W. McKinney, C. T. Russell, and others, “Gamma-ray and neutron spectrometer for the Dawn mission to 1 Ceres and 4 Vesta,” *Nucl. Sci. IEEE Trans. On*, vol. 50, no. 4, pp. 1190–1197, 2003.
- [418] T. H. Prettyman, W. C. Feldman, H. Y. McSween Jr, R. D. Dingler, D. C. Enemark, D. E. Patrick, S. A. Storms, J. S. Hendricks, J. P. Morgenthaler, K. M. Pitman, and others, “Dawn’s gamma ray and neutron detector,” in *The Dawn Mission to Minor Planets 4 Vesta and 1 Ceres*, Springer, 2012, pp. 371–459.
- [419] I. G. Mitrofanov, M. L. Litvak, A. B. Varenikov, Y. N. Barmakov, A. Behar, Y. I. Bobrovitsky, E. P. Bogolubov, W. V. Boynton, K. Harshman, E. Kan, and others, “Dynamic Albedo of Neutrons (DAN) Experiment Onboard NASA’s Mars Science Laboratory,” *Space Sci. Rev.*, vol. 170, no. 1–4, pp. 559–582, 2012.
- [420] A. S. Rivkin, E. S. Howell, F. Vilas, and L. A. Lebofsky, “Hydrated Minerals on Asteroids: The Astronomical Record,” 2003.
- [421] L. A. Lebofsky, M. A. Feierberg, A. T. Tokunaga, H. P. Larson, and J. R. Johnson, “The 1.7-to 4.2- μm spectrum of asteroid 1 Ceres: Evidence for structural water in clay minerals,” *Icarus*, vol. 48, no. 3, pp. 453–459, 1981.
- [422] E. M. Palmer, E. Heggy, M. T. Capria, F. Tosi, and C. T. Russell, “Dielectric Properties of the Surface of Asteroid Vesta from Dawn VIR Thermal Observations,” in *Lunar and Planetary Institute Science Conference Abstracts*, 2013, vol. 44, p. 2476.
- [423] D. J. Daniels, *Ground penetrating radar*. Wiley Online Library, 2005.
- [424] A. Martinez and A. P. Byrnes, *Modeling dielectric-constant values of geologic materials: An aid to ground-penetrating radar data collection and interpretation*. Kansas Geological Survey, University of Kansas, 2001.
- [425] W. Kofman, A. Hérique, J.-P. Goutail, T. Hagfors, I. P. Williams, E. Nielsen, J.-P. Barriot, Y. Barbin, C. Elachi, P. Edenhofer, and others, “The comet nucleus sounding experiment by radiowave transmission (CONCERT): a short description of the instrument and of the commissioning stages,” *Space Sci. Rev.*, vol. 128, no. 1–4, pp. 413–432, 2007.
- [426] H. Devlin and S. Correspondent, “Rise and shine: Rosetta’s Philae probe could be awake within weeks,” *the Guardian*. [Online]. Available: <http://www.theguardian.com/science/2015/apr/14/rise-and-shine-rosetta-philae-probe-could-be-awake-within-weeks>. [Accessed: 01-May-2015].
- [427] R. E. Gold, S. C. Solomon, R. L. McNutt, A. G. Santo, J. B. Abshire, M. H. Acuña, R. S. Afzal, B. J. Anderson, G. B. Andrews, P. D. Bedini, and others, “The MESSENGER mission to Mercury: scientific payload,” *Planet. Space Sci.*, vol. 49, no. 14, pp. 1467–1479, 2001.
- [428] N. Maeno, T. Araki, J. Moore, and M. Fukuda, “Dielectric response of water and ice in frozen soils,” *Phys. Chem. Ice Ed. Maeno N Hondoh*, pp. 381–386, 1992.
- [429] P. D. Kaviratna, T. J. Pinnavaia, and P. A. Schroeder, “Dielectric properties of smectite clays,” *J. Phys. Chem. Solids*, vol. 57, no. 12, pp. 1897–1906, 1996.
- [430] D. J. Tholen, “Asteroid Taxonomy from Cluster Analysis of Photometry,” 1984.
- [431] B. Zellner, D. J. Tholen, and E. F. Tedesco, “The eight-color asteroid survey: Results for 589 minor planets,” *Icarus*, vol. 61, no. 3, pp. 355–416, 1985.
- [432] S. Xu, R. P. Binzel, T. H. Burbine, and S. J. Bus, “Small main-belt asteroid spectroscopic survey,” in *Bulletin of the American Astronomical Society*, 1993, vol. 25, p. 1135.
- [433] S. J. Bus and R. P. Binzel, “Phase II of the Small Main-Belt Asteroid Spectroscopic Survey: A Feature-Based Taxonomy,” *Icarus*, vol. 158, no. 1, pp. 146–177, Jul. 2002.
- [434] F. E. DeMeo, R. P. Binzel, S. M. Slivan, and S. J. Bus, “An extension of the Bus asteroid taxonomy into the near-infrared,” *Icarus*, vol. 202, no. 1, pp. 160–180, 2009.
- [435] S. J. Bus, F. E. DeMeo, R. P. Binzel, and S. M. Slivan, “Bus-DeMeo taxonomy: Extending asteroid taxonomy into the near-infrared,” in *Bulletin of the American Astronomical Society*, 2008, vol. 40, p. 440.
- [436] C. Magri, S. J. Ostro, K. D. Rosema, M. L. Thomas, D. L. Mitchell, D. B. Campbell, J. F. Chandler, I. I. Shapiro, J. D. Giorgini, and D. K. Yeomans, “Mainbelt asteroids: Results of Arecibo and Goldstone radar observations of 37 objects during 1980–1995,” *Icarus*, vol. 140, no. 2, pp. 379–407, 1999.
- [437] C. Magri, M. C. Nolan, S. J. Ostro, and J. D. Giorgini, “A radar survey of main-belt asteroids: Arecibo observations of 55 objects during 1999–2003,” *Icarus*, vol. 186, no. 1, pp. 126–151, 2007.
- [438] E. T. Lu, H. Reitsema, J. Troeltzsch, and S. Hubbard, “The B612 Foundation Sentinel Space Telescope,” *New Space*, vol. 1, no. 1, pp. 42–45, 2013.

- [439] M. J. S. Belton, J. Veverka, P. Thomas, P. Helfenstein, D. Simonelli, C. Chapman, M. E. Davies, R. Greeley, R. Greenberg, J. Head, and others, “Galileo encounter with 951 Gaspra: First pictures of an asteroid,” *Science*, vol. 257, no. 5077, pp. 1647–1652, 1992.
- [440] C. R. Chapman, J. Veverka, P. C. Thomas, K. Klaasen, M. J. S. Belton, A. Harch, A. McEwen, T. V. Johnson, P. Helfenstein, M. E. Davies, and others, “Discovery and physical properties of Dactyl, a satellite of asteroid 243 Ida,” *Nature*, vol. 374, no. 6525, pp. 783–784, 1995.
- [441] T. J. Ahrens, T. Takata, J. D. O’Keefe, and G. S. Orton, “Impact of comet Shoemaker-Levy 9 on Jupiter,” *Geophys. Res. Lett.*, vol. 21, no. 11, pp. 1087–1090, 1994.
- [442] R. W. Carlson, P. R. Weissman, M. Segura, J. Hui, W. D. Smythe, T. V. Johnson, K. H. Baines, P. Drossart, T. Encrenaz, and F. E. Leader, “Galileo infrared observations of the Shoemaker-Levy 9 G Impact Fireball: A Preliminary report,” *Geophys. Res. Lett.*, vol. 22, no. 12, pp. 1557–1560, 1995.
- [443] D. L. Matson, L. J. Spilker, and J.-P. Lebreton, “The Cassini/Huygens mission to the Saturnian system,” in *The Cassini-Huygens Mission*, Springer, 2003, pp. 1–58.
- [444] S. Kempf, R. Srama, N. Altobelli, S. Auer, V. Tschernjawski, J. Bradley, M. E. Burton, S. Helfert, T. V. Johnson, H. Krüger, and others, “Cassini between Earth and asteroid belt: first in-situ charge measurements of interplanetary grains,” *Icarus*, vol. 171, no. 2, pp. 317–335, 2004.
- [445] J. Veverka, B. Farquhar, M. Robinson, P. Thomas, S. Murchie, A. Harch, P. G. Antreasian, S. R. Chesley, J. K. Miller, W. M. Owen, and others, “The landing of the NEAR-Shoemaker spacecraft on asteroid 433 Eros,” *Nature*, vol. 413, no. 6854, pp. 390–393, 2001.
- [446] M. T. Zuber, D. E. Smith, A. F. Cheng, J. B. Garvin, O. Aharonson, T. D. Cole, P. J. Dunn, Y. Guo, F. G. Lemoine, G. A. Neumann, and others, “The shape of 433 Eros from the NEAR-Shoemaker laser rangefinder,” *Science*, vol. 289, no. 5487, pp. 2097–2101, 2000.
- [447] S. A. Sandford, J. Aléon, C. M. Alexander, T. Araki, S. Bajt, G. A. Baratta, J. Borg, J. P. Bradley, D. E. Brownlee, J. R. Brucato, and others, “Organics captured from comet 81P/Wild 2 by the Stardust spacecraft,” *Science*, vol. 314, no. 5806, pp. 1720–1724, 2006.
- [448] S. Abe, T. Mukai, N. Hirata, O. S. Barnouin-Jha, A. F. Cheng, H. Demura, R. W. Gaskell, T. Hashimoto, K. Hiraoka, T. Honda, and others, “Mass and local topography measurements of Itokawa by Hayabusa,” *Science*, vol. 312, no. 5778, pp. 1344–1347, 2006.
- [449] A. Fujiwara, J. Kawaguchi, D. K. Yeomans, M. Abe, T. Mukai, T. Okada, J. Saito, H. Yano, M. Yoshikawa, D. J. Scheeres, and others, “The rubble-pile asteroid Itokawa as observed by Hayabusa,” *Science*, vol. 312, no. 5778, pp. 1330–1334, 2006.
- [450] T. Okada, K. Shirai, Y. Yamamoto, T. Arai, K. Ogawa, K. Hosono, and M. Kato, “X-ray fluorescence spectrometry of asteroid Itokawa by Hayabusa,” *Science*, vol. 312, no. 5778, pp. 1338–1341, 2006.
- [451] M. F. A’Hearn, M. J. S. Belton, W. A. Delamere, J. Kissel, K. P. Klaasen, L. A. McFadden, K. J. Meech, H. J. Melosh, P. H. Schultz, J. M. Sunshine, and others, “Deep impact: excavating comet Tempel 1,” *science*, vol. 310, no. 5746, pp. 258–264, 2005.
- [452] K. J. Meech, N. Ageorges, M. F. A’Hearn, C. Arpigny, A. Ates, J. Aycock, S. Bagnulo, J. Bailey, R. Barber, L. Barrera, and others, “Deep Impact: Observations from a worldwide Earth-based campaign,” *Science*, vol. 310, no. 5746, pp. 265–269, 2005.
- [453] C. M. Lisse, J. VanCleve, A. C. Adams, M. F. A’hearn, Y. R. Fernández, T. L. Farnham, L. Armus, C. J. Grillmair, J. Ingalls, M. J. S. Belton, and others, “Spitzer spectral observations of the Deep Impact ejecta,” *Science*, vol. 313, no. 5787, pp. 635–640, 2006.
- [454] P. H. Schultz, B. Hermalyn, and J. Veverka, “The Deep Impact crater on 9P/Tempel-1 from Stardust-NExT,” *Icarus*, vol. 222, no. 2, pp. 502–515, 2013.
- [455] J. Veverka, K. Klaasen, M. A’Hearn, M. Belton, D. Brownlee, S. Chesley, B. Clark, T. Economou, R. Farquhar, S. F. Green, and others, “Return to comet Tempel 1: overview of Stardust-NExT results,” *Icarus*, vol. 222, no. 2, pp. 424–435, 2013.
- [456] K.-H. Glassmeier, H. Boehnhardt, D. Koschny, E. Kührt, and I. Richter, “The Rosetta mission: flying towards the origin of the solar system,” *Space Sci. Rev.*, vol. 128, no. 1–4, pp. 1–21, 2007.
- [457] J. Biele and S. Ulamec, “Capabilities of Philae, the Rosetta lander,” *Space Sci. Rev.*, vol. 138, no. 1–4, pp. 275–289, 2008.
- [458] C. T. Russell, F. Capaccioni, A. Coradini, M. C. De Sanctis, W. C. Feldman, R. Jaumann, H. U. Keller, T. B. McCord, L. A. McFadden, S. Mottola, and others, “Dawn mission to Vesta and Ceres,” *Earth Moon Planets*, vol. 101, no. 1–2, pp. 65–91, 2007.

- [459] C. T. Russell, C. A. Raymond, A. Coradini, H. Y. McSween, M. T. Zuber, A. Nathues, M. C. De Sanctis, R. Jaumann, A. S. Konopliv, F. Preusker, and others, “Dawn at Vesta: Testing the protoplanetary paradigm,” *Science*, vol. 336, no. 6082, pp. 684–686, 2012.
- [460] J. Kawaguchi, A. Fujiwara, and T. Uesugi, “Hayabusa—Its technology and science accomplishment summary and Hayabusa 2,” *Acta Astronaut.*, vol. 62, no. 10, pp. 639–647, 2008.
- [461] D. S. Lauretta and O.-R. Team, “An Overview of the OSIRIS-REx Asteroid Sample Return Mission,” in *Lunar and Planetary Institute Science Conference Abstracts*, 2012, vol. 43, p. 2491.
- [462] E. L. Lawler, J. K. Lenstra, A. R. Kan, and D. B. Shmoys, *The traveling salesman problem: a guided tour of combinatorial optimization*, vol. 3. Wiley New York, 1985.
- [463] J. T. Betts, “Survey of Numerical Methods for Trajectory Optimization,” *J. Guid. Control Dyn.*, vol. 21, no. 2, pp. 193–207, 1998.
- [464] C. G. Sauer Jr, “Optimization of multiple target electric propulsion trajectories,” 1973.
- [465] T. T. McConaghy, T. J. Debban, A. E. Petropoulos, and J. M. Longuski, “Design and optimization of low-thrust trajectories with gravity assists,” *J. Spacecr. Rockets*, vol. 40, no. 3, pp. 380–387, 2003.
- [466] A. E. Petropoulos and J. M. Longuski, “Shape-based algorithm for the automated design of low-thrust, gravity assist trajectories,” *J. Spacecr. Rockets*, vol. 41, no. 5, pp. 787–796, 2004.
- [467] J. R. Stuart and K. C. Howell, “An Automated Search Procedure to Generate Optimal Low-Thrust Rendezvous Tours of the Sun-Jupiter Trojan Asteroids.”
- [468] “GTOC 4 – Asteroids billiard | GTOC Portal.” .
- [469] “GTOC 7 – Multi-spacecraft exploration of the asteroid belt | GTOC Portal.” .
- [470] S. Nozette, P. Rustan, L. P. Pleasance, J. F. Kordas, I. T. Lewis, H. S. Park, R. E. Priest, D. M. Horan, P. Regeon, C. L. Lichtenberg, and others, “The Clementine mission to the Moon: Scientific overview,” *Science*, vol. 266, no. 5192, pp. 1835–1839, 1994.
- [471] J. Mueller, M. Paluszczek, S. Thomas, A. Knutson, D. Klein, and M. Tam, “Asteroid Prospector,” 2013.
- [472] M. Carr, D. Beaty, and et al., “Analysis of Strategic Knowledge Gaps Associated with Potential Human Missions to the Martian System.” MEPAG/SBAG, 30-Jun-2012.
- [473] Corbin, Benjamin, *Mars Precession Demo*. <https://www.youtube.com/watch?v=au00YJfwsQA>, 2014.
- [474] S. R.-S. A. G. MEPAG, “Findings of the Mars special regions science analysis group.,” *Astrobiology*, vol. 6, no. 5, p. 677, 2006.
- [475] J. D. Rummel, M. S. Race, D. L. DeVinzi, P. J. Schad, P. D. Stabekis, M. Viso, and S. E. Acevedo, *A Draft Test Protocol for Detecting Possible Biohazards in Martian Samples Returned to Earth*. 2002.
- [476] W. V. Boynton, W. C. Feldman, I. G. Mitrofanov, L. G. Evans, R. C. Reedy, S. W. Squyres, R. Starr, J. I. Trombka, C. D’uston, J. R. Arnold, P. a. J. Englert, A. E. Metzger, H. Wänke, J. Brückner, D. M. Drake, C. Shinohara, C. Fellows, D. K. Hamara, K. Harshman, K. Kerry, C. Turner, M. Ward, H. Barthe, K. R. Fuller, S. A. Storms, G. W. Thornton, J. L. Longmire, M. L. Litvak, and A. K. Ton’chev, “The Mars Odyssey Gamma-Ray Spectrometer Instrument Suite,” in *2001 Mars Odyssey*, C. T. Russell, Ed. Springer Netherlands, 2004, pp. 37–83.
- [477] W. C. Feldman, W. V. Boynton, R. L. Tokar, T. H. Prettyman, O. Gasnault, S. W. Squyres, R. C. Elphic, D. J. Lawrence, S. L. Lawson, S. Maurice, and others, “Global distribution of neutrons from Mars: Results from Mars Odyssey,” *Science*, vol. 297, no. 5578, pp. 75–78, 2002.
- [478] I. Mitrofanov, D. Anfimov, A. Kozyrev, M. Litvak, A. Sanin, V. Tret’yakov, A. Krylov, V. Shvetsov, W. Boynton, C. Shinohara, D. Hamara, and R. S. Saunders, “Maps of Subsurface Hydrogen from the High Energy Neutron Detector, Mars Odyssey,” *Science*, vol. 297, no. 5578, pp. 78–81, Jul. 2002.
- [479] “Catalog Page for PIA04907.” [Online]. Available: <http://photojournal.jpl.nasa.gov/catalog/PIA04907>. [Accessed: 25-Mar-2015].
- [480] S. Karunatillake, J. J. Wray, O. Gasnault, S. M. McLennan, A. D. Rogers, S. W. Squyres, W. V. Boynton, J. R. Skok, L. Ojha, and N. Olsen, “Sulfates hydrating bulk soil in the Martian low and middle latitudes,” *Geophys. Res. Lett.*, vol. 41, no. 22, p. 2014GL061136, Nov. 2014.
- [481] M. H. Hecht, S. P. Kounaves, R. C. Quinn, S. J. West, S. M. M. Young, D. W. Ming, D. C. Catling, B. C. Clark, W. V. Boynton, J. Hoffman, L. P. DeFlores, K. Gospodinova, J. Kapit, and P. H. Smith, “Detection of Perchlorate and the Soluble Chemistry of Martian Soil at the Phoenix Lander Site,” *Science*, vol. 325, no. 5936, pp. 64–67, Jul. 2009.
- [482] S. P. Kounaves, M. H. Hecht, J. Kapit, K. Gospodinova, L. DeFlores, R. C. Quinn, W. V. Boynton, B. C. Clark, D. C. Catling, P. Hredzak, D. W. Ming, Q. Moore, J. Shusterman, S. Stroble, S. J. West, and S. M. M. Young, “Wet Chemistry experiments on the 2007 Phoenix Mars Scout Lander mission: Data analysis and results,” *J. Geophys. Res. Planets*, vol. 115, no. E1, p. E00E10, Jan. 2010.

- [483] P. R. Christensen, B. M. Jakosky, H. H. Kieffer, M. C. Malin, H. Y. McSween Jr, K. Nealson, G. L. Mehall, S. H. Silverman, S. Ferry, M. Caplinger, and others, “The thermal emission imaging system (THEMIS) for the Mars 2001 Odyssey Mission,” *Space Sci. Rev.*, vol. 110, no. 1–2, pp. 85–130, 2004.
- [484] M. P. Golombek, R. A. Cook, H. J. Moore, and T. J. Parker, “Selection of the Mars Pathfinder landing site,” *J. Geophys. Res. Planets 1991–2012*, vol. 102, no. E2, pp. 3967–3988, 1997.
- [485] M. P. Golombek, J. A. Grant, T. J. Parker, D. M. Kass, J. A. Crisp, S. W. Squyres, A. F. C. Haldemann, M. Adler, W. J. Lee, N. T. Bridges, and others, “Selection of the Mars Exploration Rover landing sites,” *J. Geophys. Res. Planets 1991–2012*, vol. 108, no. E12, 2003.
- [486] R. Arvidson, D. Adams, G. Bonfiglio, P. Christensen, S. Cull, M. Golombek, J. Guinn, E. Guinness, T. Heet, R. Kirk, and others, “Mars Exploration Program 2007 Phoenix landing site selection and characteristics,” *J. Geophys. Res. Planets 1991–2012*, vol. 113, no. E3, 2008.
- [487] R. L. Kirk, E. Howington-Kraus, M. R. Rosiek, J. A. Anderson, B. A. Archinal, K. J. Becker, D. A. Cook, D. M. Galuszka, P. E. Geissler, T. M. Hare, I. M. Holmberg, L. P. Keszthelyi, B. L. Redding, W. A. Delamere, D. Gallagher, J. D. Chapel, E. M. Eliason, R. King, and A. S. McEwen, “Ultra-high resolution topographic mapping of Mars with MRO HiRISE stereo images: Meter-scale slopes of candidate Phoenix landing sites,” *J. Geophys. Res. Planets*, vol. 113, no. E3, p. E00A24, Mar. 2008.
- [488] M. Golombek, J. Grant, D. Kipp, A. Vasavada, R. Kirk, R. Fergason, P. Bellutta, F. Calef, K. Larsen, Y. Katayama, and others, “Selection of the Mars Science Laboratory landing site,” *Space Sci. Rev.*, vol. 170, no. 1–4, pp. 641–737, 2012.
- [489] P. Mahaffy, “Exploration of the Habitability of Mars: Development of Analytical Protocols for Measurement of Organic Carbon on the 2009 Mars Science Laboratory,” *Space Sci. Rev.*, vol. 135, no. 1–4, pp. 255–268, Jul. 2007.
- [490] P. R. Mahaffy, C. R. Webster, M. Cabane, P. G. Conrad, P. Coll, S. K. Atreya, R. Arvey, M. Barciniak, M. Benna, L. Bleacher, and others, “The sample analysis at Mars investigation and instrument suite,” *Space Sci. Rev.*, vol. 170, no. 1–4, pp. 401–478, 2012.
- [491] “Leak in Curiosity’s Wet Chemistry Test Finds Organics,” *DNews*. [Online]. Available: <http://news.discovery.com/space/leak-in-curiositys-wet-chemistry-test-finds-mars-organics-150317.htm>. [Accessed: 22-Mar-2015].
- [492] C. Freissinet, D. P. Glavin, A. Buch, C. Szopa, S. Kashyap, H. B. Franz, J. L. Eigenbrode, W. B. Brinkerhoff, R. Navarro-Gonzalez, S. Teinturier, C. A. Malespin, B. D. Prats, P. R. Mahaffy, and SAM and MSL Science Teams, “First In Situ Wet Chemistry Experiment on Mars Using the SAM Instrument: MTBSTFA Derivatization on a Martian Mudstone,” presented at the 46th Lunar and Planetary Science Conference, The Woodlands, Texas, 2015.
- [493] R. Gowen, A. Smith, A. Griffiths, A. Coates, A. Ball, S. Barber, A. Hagermann, S. Sheridan, I. Crawford, P. Church, and others, “Kinetic penetrators for exploration of solar system bodies,” in *Fifth International Planetary Probes Workshop (IPPW-5)*, Bordeaux, France, 2007, pp. 23–29.
- [494] A. J. Ball, J. R. Garry, R. D. Lorenz, and V. Kerzhanovich, *Planetary landers and entry probes*, vol. 1. Cambridge University Press Cambridge, UK, 2007.
- [495] S. A. Gavit and G. Powell, “The new millennium program’s Mars microprobe mission,” *Acta Astronaut.*, vol. 39, no. 1, pp. 273–280, 1996.
- [496] R. C. Blue, “Mars microprobe project instrumentation package,” *Acta Astronaut.*, vol. 45, no. 4, pp. 585–595, 1999.
- [497] S. Smrekar, D. Catling, R. Lorenz, J. Magalhães, J. Moersch, P. Morgan, B. Murray, M. Presley, A. Yen, A. Zent, and others, “Deep Space 2: the Mars microprobe mission,” *J. Geophys. Res. Planets 1991–2012*, vol. 104, no. E11, pp. 27013–27030, 1999.
- [498] R. D. Lorenz, J. E. Moersch, J. A. Stone, A. Ron Morgan Jr, and S. E. Smrekar, “Penetration tests on the DS-2 Mars microprobes: penetration depth and impact accelerometry,” *Planet. Space Sci.*, vol. 48, no. 5, pp. 419–436, 2000.
- [499] R. A. Gowen, A. Smith, A. D. Fortes, S. Barber, P. Brown, P. Church, G. Collinson, A. J. Coates, G. Collins, I. A. Crawford, and others, “Penetrators for in situ subsurface investigations of Europa,” *Adv. Space Res.*, vol. 48, no. 4, pp. 725–742, 2011.
- [500] P. Weiss, K. L. Yung, N. Kömle, S. M. Ko, E. Kaufmann, and G. Kargl, “Thermal drill sampling system onboard high-velocity impactors for exploring the subsurface of Europa,” *Adv. Space Res.*, vol. 48, no. 4, pp. 743–754, 2011.
- [501] C. W. Young, “Penetration equations,” Sandia National Labs., Albuquerque, NM (United States), 1997.

- [502] “Deep Space 2 Technology/Probe: Forebody,” *NASA JPL New Millenium Program*. [Online]. Available: <http://nmp.jpl.nasa.gov/ds2/tech/forebody.html>. [Accessed: 17-Mar-2015].
- [503] G. Collinson, U. P. Consortium, and others, “Planetary Penetrators-The Vanguard for the Future Exploration of the Solar System,” *J. Br. Interplanet. Soc.*, vol. 61, pp. 198–202, 2008.
- [504] A. Smith, I. A. Crawford, R. A. Gowen, R. Ambrosi, M. Anand, B. Banerdt, N. Bannister, N. Bowles, C. Braithwaite, P. Brown, and others, “Lunar Net—a proposal in response to an ESA M3 call in 2010 for a medium sized mission,” *Exp. Astron.*, vol. 33, no. 2–3, pp. 587–644, 2012.
- [505] A. Smith, I. A. Crawford, R. A. Gowen, A. J. Ball, S. J. Barber, P. Church, A. J. Coates, Y. Gao, A. D. Griffiths, A. Hagermann, and others, “LunarEX—a proposal to cosmic vision,” *Exp. Astron.*, vol. 23, no. 3, pp. 711–740, 2009.
- [506] J. Burns, “Missile practice for Moon mission,” *BBC*, 06-Jun-2008.
- [507] A. M. Harri, W. Schmidt, V. Linkin, S. Alexashkin, and L. Vázquez, “MetNet Mars Mission—New Lander Generation for Martian in situ Surface Observations,” in *European Planetary Science Congress 2012*, 2012, vol. 1, p. 671.
- [508] M. F. Mora, F. Greer, A. M. Stockton, S. Bryant, and P. A. Willis, “Toward total automation of microfluidics for extraterrestrial in situ analysis,” *Anal. Chem.*, vol. 83, no. 22, pp. 8636–8641, 2011.
- [509] F. W. Taylor, J. T. Houghton, G. D. Peskett, C. D. Rodgers, and E. J. Williamson, “Radiometer for remote sounding of the upper atmosphere,” *Appl. Opt.*, vol. 11, no. 1, pp. 135–141, 1972.
- [510] C. M. R. Platt, “Remote sounding of high clouds: I. Calculation of visible and infrared optical properties from lidar and radiometer measurements,” *J. Appl. Meteorol.*, vol. 18, no. 9, pp. 1130–1143, 1979.
- [511] S. T. Brown, B. Lambrigtsen, R. F. Denning, T. Gaier, P. Kangaslahti, B. H. Lim, J. M. Tanabe, and A. B. Tanner, “The High-Altitude MMIC sounding radiometer for the Global Hawk unmanned aerial vehicle: instrument description and performance,” *Geosci. Remote Sens. IEEE Trans. On*, vol. 49, no. 9, pp. 3291–3301, 2011.
- [512] D. J. McCleese, J. T. Schofield, F. W. Taylor, S. B. Calcutt, M. C. Foote, D. M. Kass, C. B. Leovy, D. A. Paige, P. L. Read, and R. W. Zurek, “Mars Climate Sounder: An investigation of thermal and water vapor structure, dust and condensate distributions in the atmosphere, and energy balance of the polar regions,” *J. Geophys. Res. Planets 1991–2012*, vol. 112, no. E5, 2007.
- [513] “GOMSpace Nanocam CIU Datasheet.”
- [514] Students of the Space Missions and Systems Design Class and S. Corpino, “Space Missions & Systems Design - Project Work 2014: Team C (Ground).” Politecnico di Torino, 16-Jun-2014.
- [515] K. Zacny, Y. Bar-Cohen, M. Brennan, G. Briggs, G. Cooper, K. Davis, B. Dolgin, D. Glaser, B. Glass, S. Gorevan, and others, “Drilling systems for extraterrestrial subsurface exploration,” *Astrobiology*, vol. 8, no. 3, pp. 665–706, 2008.
- [516] Y. Bar-Cohen and K. Zacny, *Drilling in extreme environments: penetration and sampling on earth and other planets*. John Wiley & Sons, 2009.
- [517] K. A. Zacny and G. A. Cooper, “Friction of drill bits under Martian pressure,” *J. Geophys. Res. Planets 1991–2012*, vol. 112, no. E3, 2007.
- [518] K. Zacny and G. Cooper, “Coring basalt rock under simulated martian atmospheric conditions,” *Mars*, vol. 3, 2007.
- [519] K. Zacny, P. Chu, G. Paulsen, J. Spring, M. Hedlund, B. Mellerowicz, A. Garcia, S. Indyk, and J. Craft, “Sample Acquisition and Caching Architectures for the Mars 2020 Rover Mission,” *AIAA Space 2013*, 2013.
- [520] S. Blanchemanche, A. Rona-Tas, A. Cornuéjols, A. Duroy, and C. Martin, “An Ontology of Scientific Uncertainty: Methodological Lessons from Analyzing Expressions of Uncertainty in Food Risk Assessment,” in *ISA/ESA Amsterdam conference, Amsterdam, Nederland*, 2013.
- [521] M. Kouprie and F. S. Visser, “A framework for empathy in design: stepping into and out of the user’s life,” *J. Eng. Des.*, vol. 20, no. 5, pp. 437–448, 2009.
- [522] C. Leyva, “Empathy in design,” University of Cincinnati, 2013.
- [523] C. L. Marino, “Empathy in Design,” UNIVERSITY OF CINCINNATI, 2013.
- [524] N. Dalkey and O. Helmer, “An experimental application of the Delphi method to the use of experts,” *Manag. Sci.*, vol. 9, no. 3, pp. 458–467, 1963.
- [525] H. A. Linstone, M. Turoff, and others, *The Delphi method: Techniques and applications*, vol. 29. Addison-Wesley Reading, MA, 1975.

A. Appendix – ExoplanetSat

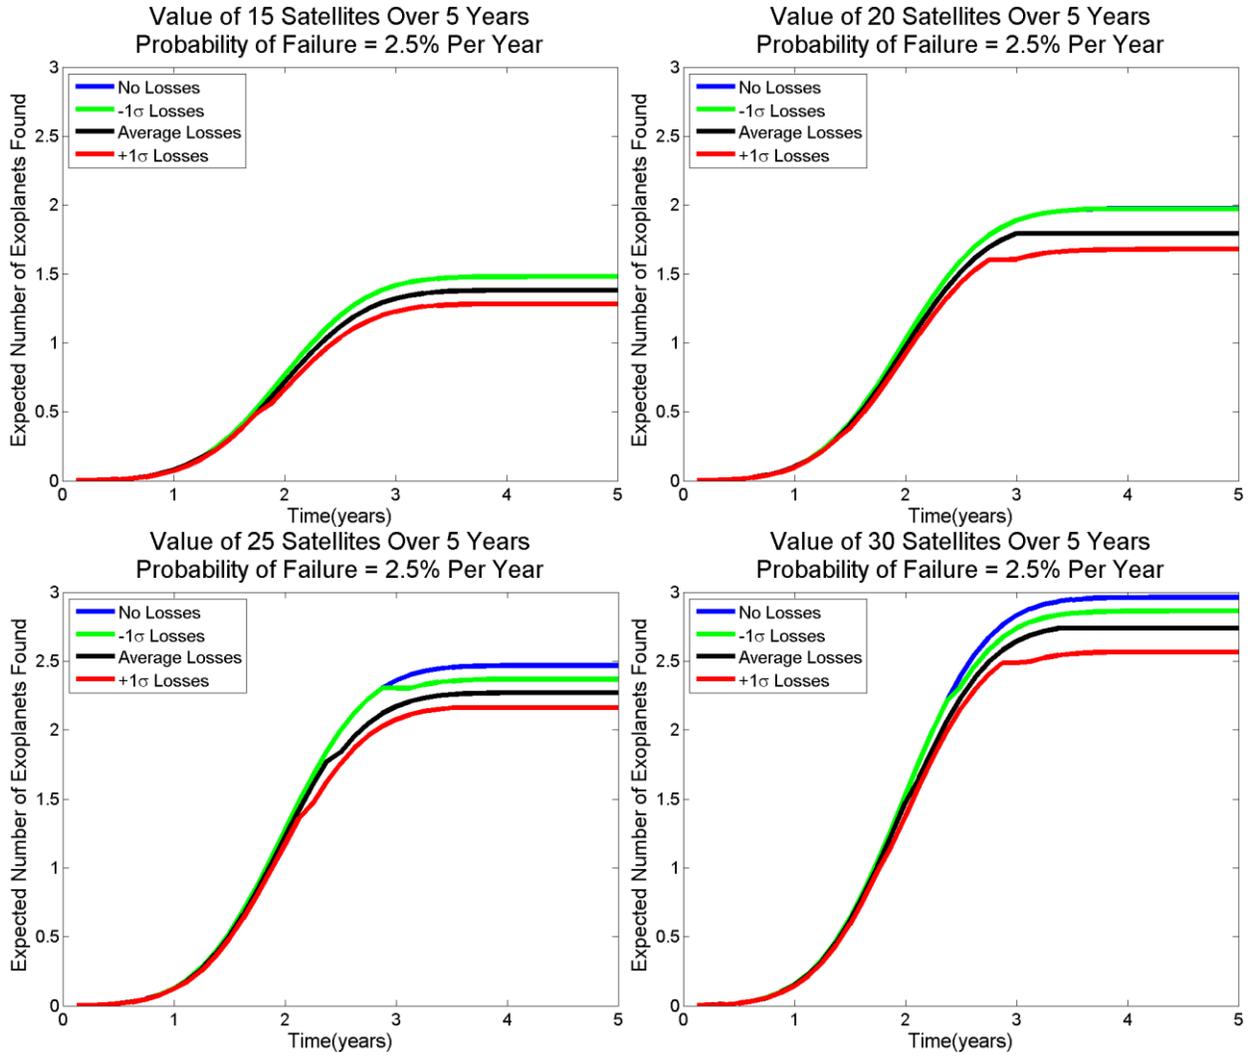


Figure A-1: Comparison of the Expected Number vs. mission lifetime for 15 (top left), 20 (top right), 25 (bottom left), and 30 (bottom right) satellites per constellation with 2.5% failure rate.

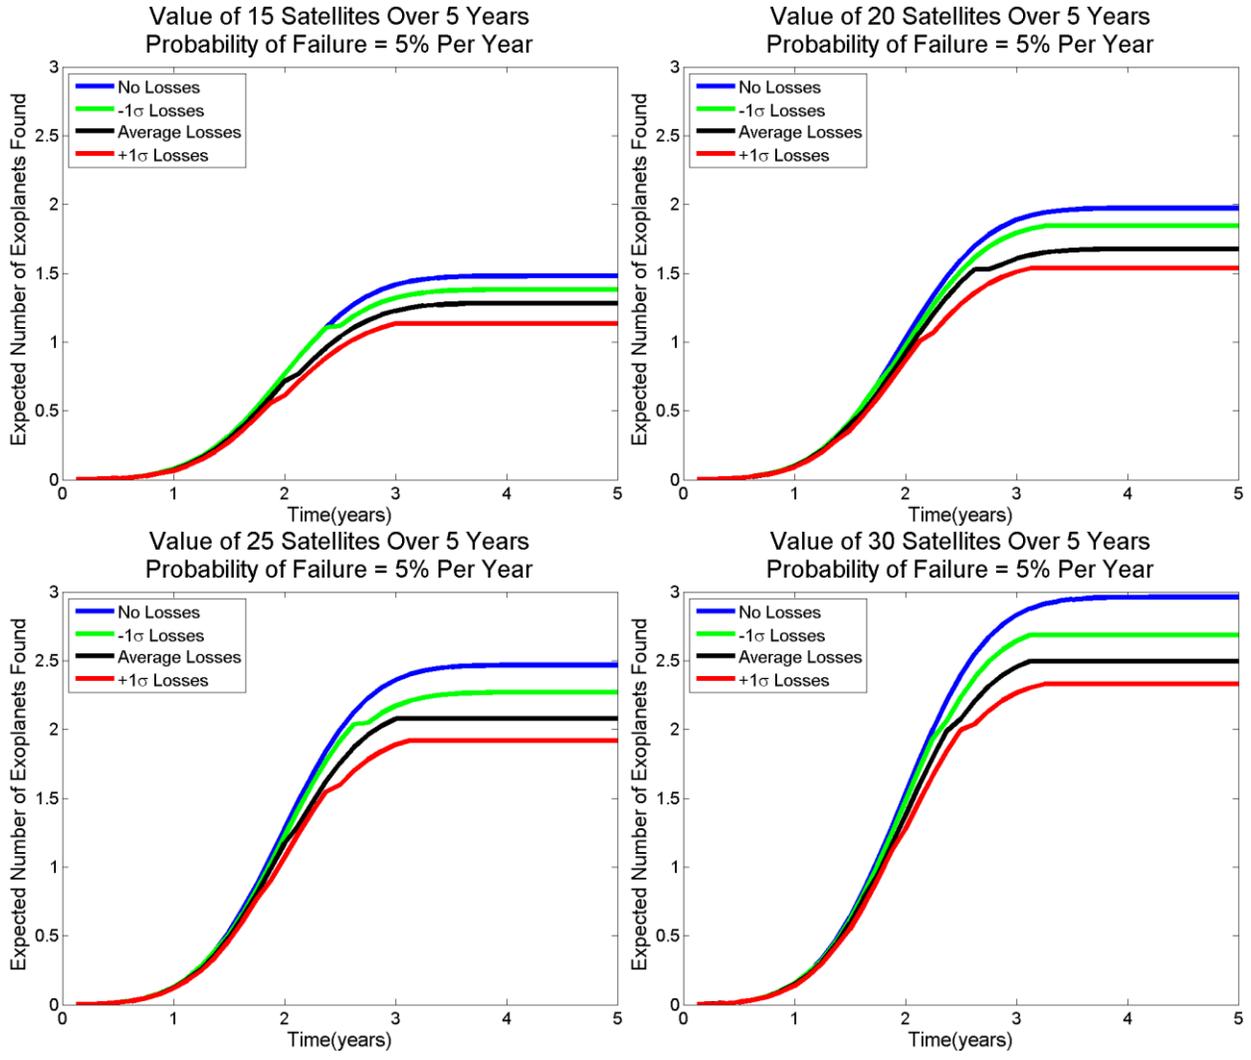


Figure A-2: Comparison of the Expected Number vs. mission lifetime for 15 (top left), 20 (top right), 25 (bottom left), and 30 (bottom right) satellites per constellation with 5% failure rate.

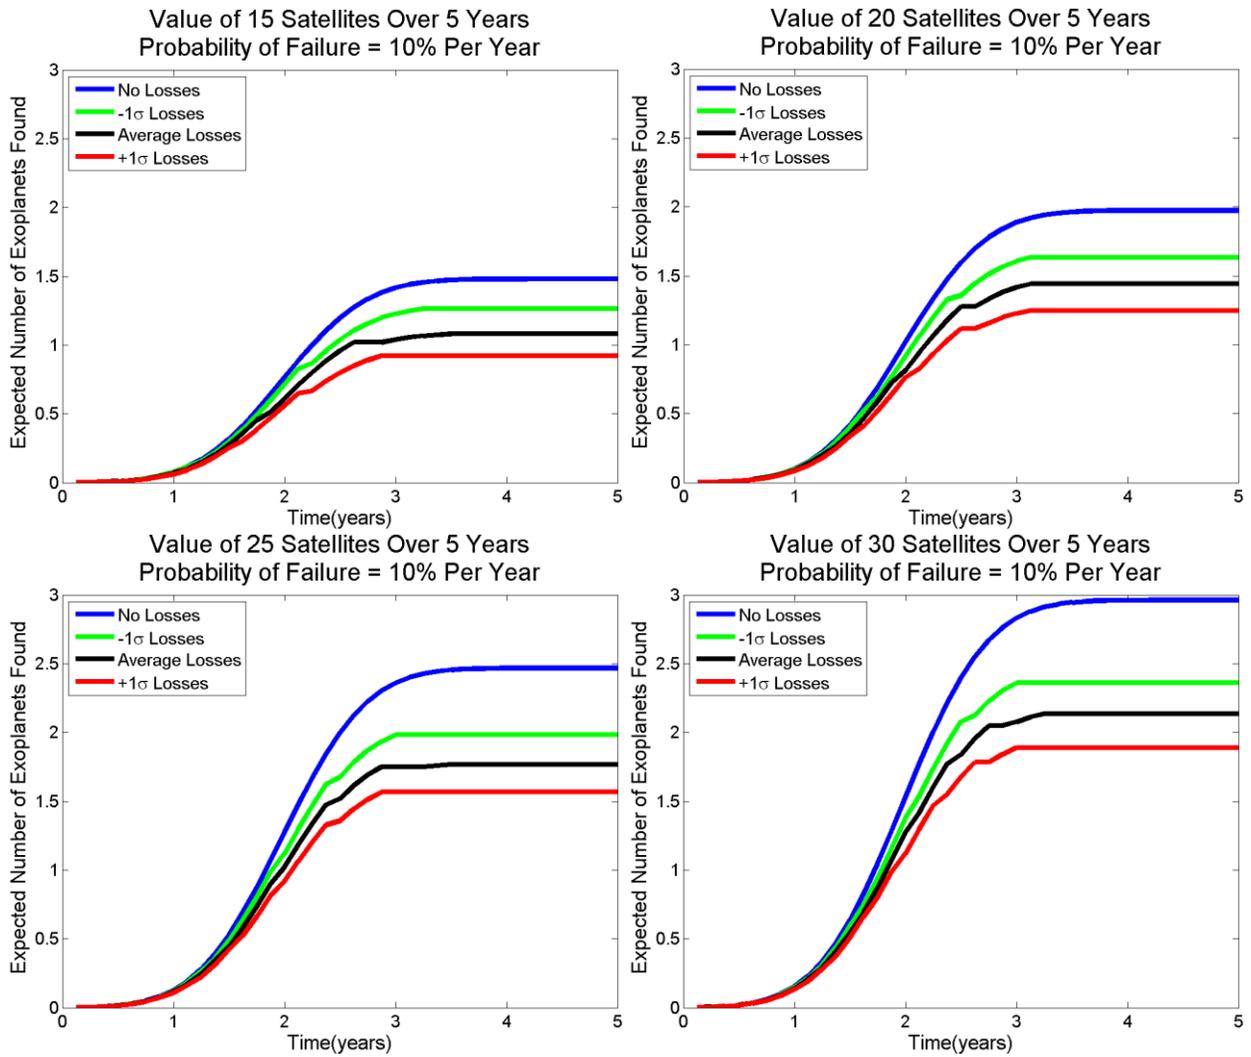


Figure A-3: Comparison of the Expected Number vs. mission lifetime for 15 (top left), 20 (top right), 25 (bottom left), and 30 (bottom right) satellites per constellation with 10% failure rate.

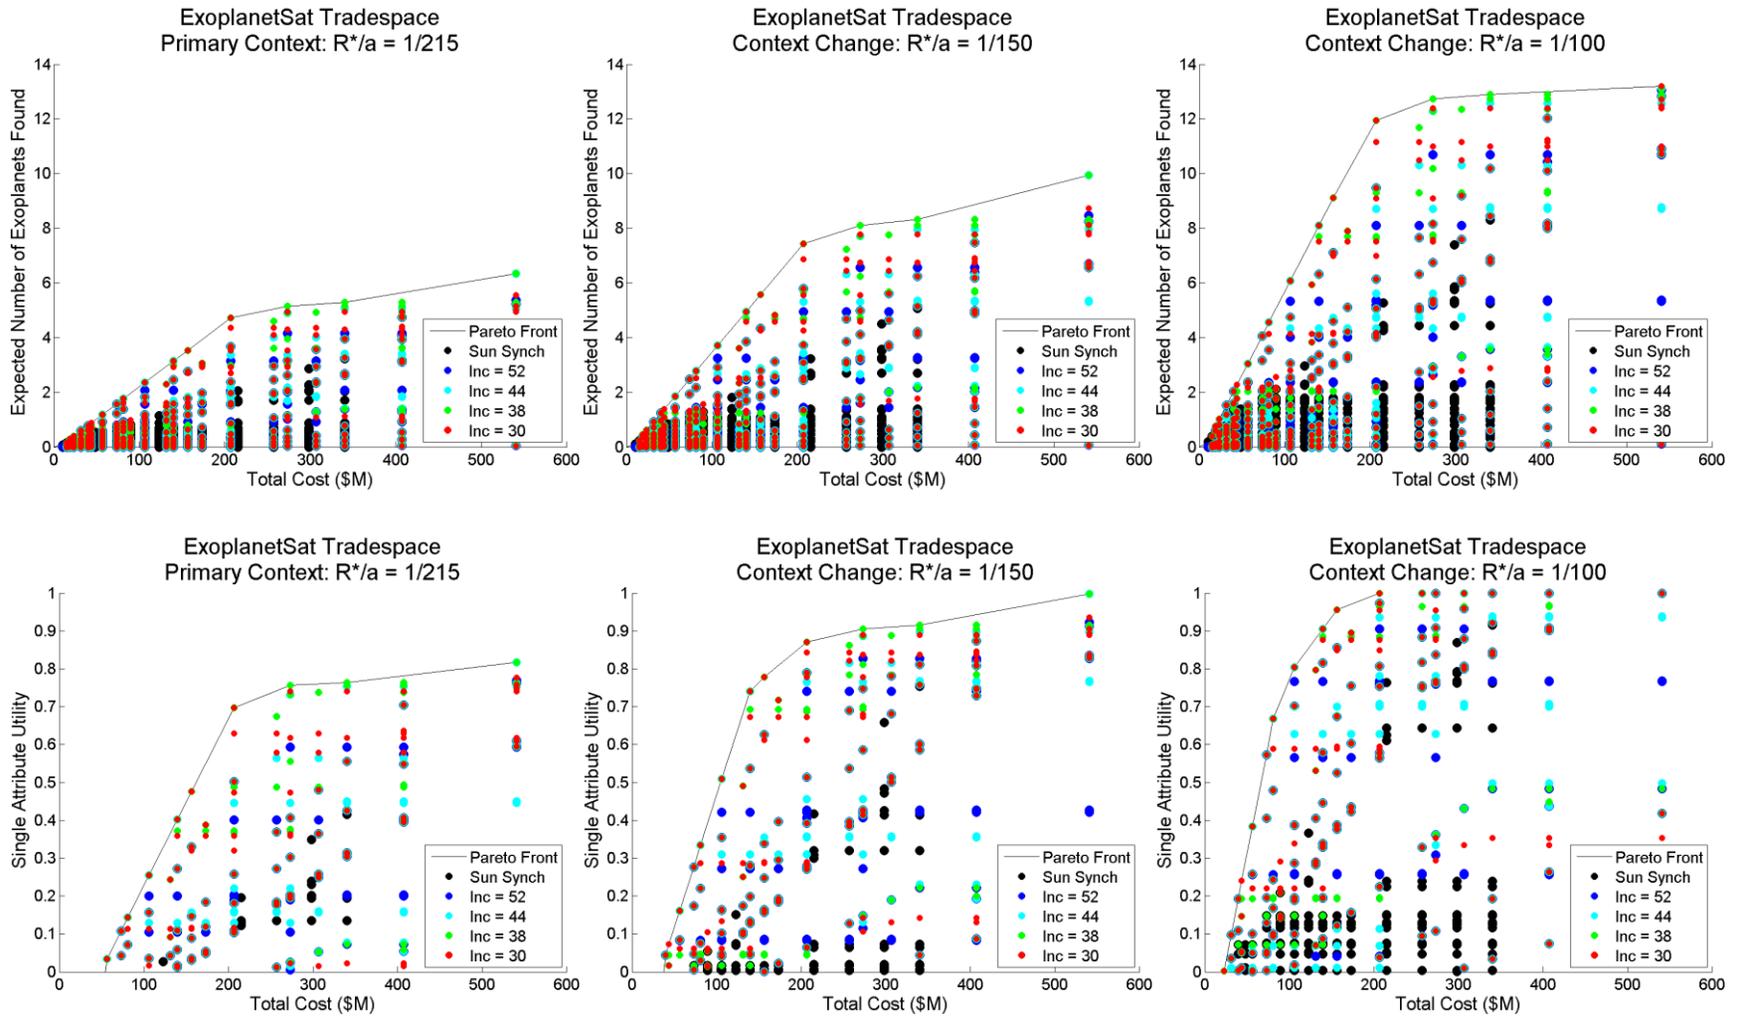


Figure A-4: Tradespaces showing variations in the Expected Number of Exoplanets Discovered (top) and Single-Attribute Utility (bottom) in contexts ranging from R^*/a equals 1/215 (left), 1/150 (center), and 1/100 (right). As the ratio gets higher, utility tends to rise for all designs, though the Pareto front does not change significantly on the lower end of the curve.

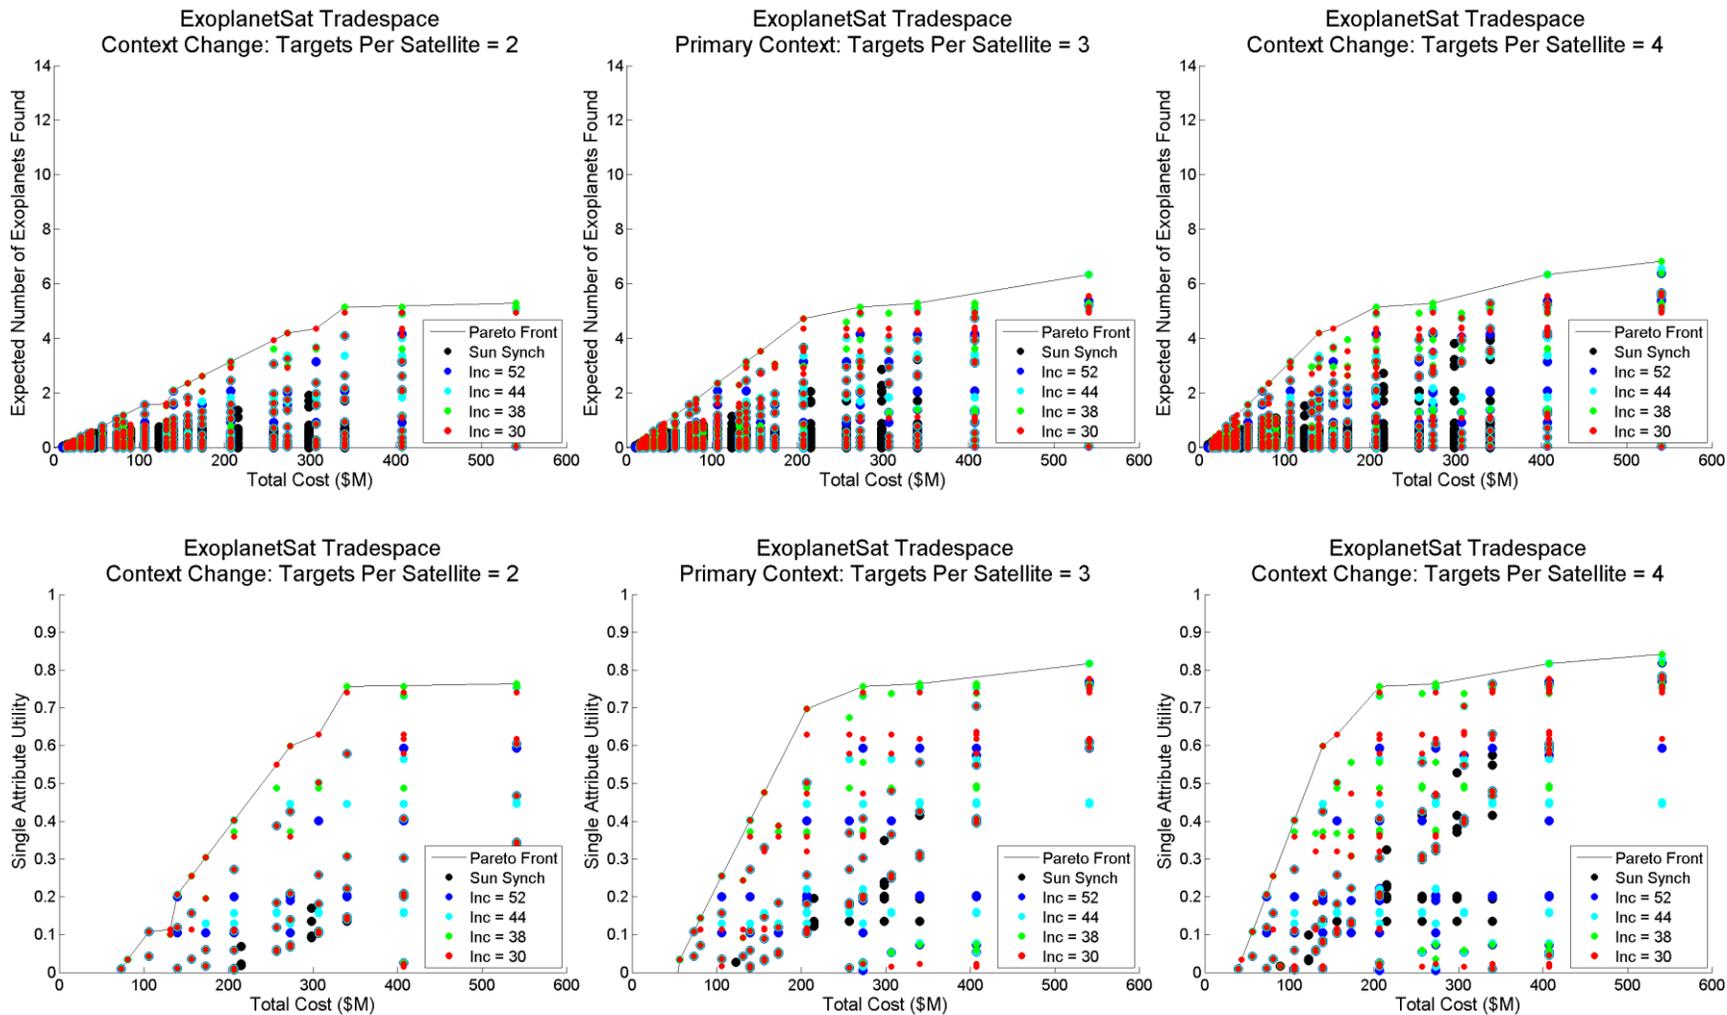


Figure A-5: Tradespaces showing variations in the Expected Number of Exoplanets Discovered (top) and Single-Attribute Utility (bottom) in contexts ranging from Number of Targets Per Satellite equals 2 (left), 3 (center), and 4 (right).

B. Appendix – HOBOCOP

Table B-1: List of natural constants.

Parameter	Value	Units
Astronomical Unit (AU)	149,597,871	km
Solar Flux at 1 AU	1,365	W/m ²
Solar Mass	1.989e30	kg
Stefan-Boltzmann Constant	5.67e-8	W/m ² -K ⁴

Table B-2: List of parameter assumptions in the thermal submodule.

Parameter	Value	Units
White Epoxy Absorptivity	0.248	--
White Epoxy Emissivity	0.888	--
Torlon 5030 Coefficient of Thermal Conductivity	0.37	W/m-K
Torlon 5030 Density	1,610	kg/m ³
CFOAM Density	270	kg/m ³
CFOAM Heat Capacity	700	J/kg-K
CFOAM Coefficient of Thermal Conductivity	0.334	W/m-K
Solar Panel Absorptivity	0.805	--
Solar Panel Emissivity	0.825	--
Solar Panel Maximum Temperature	125	C
Satellite Structure Density	2,700	kg/m ³
Satellite Structure Heat Capacity	920	J/kg-K
Heat Shield Diameter	25	cm
Effective Temperature, Sun	600	K
Effective Temperature, CMB	5	K

Table B-3: List of satellite component parameters and assumptions.

Parameter	Value	Units
Star Tracker Mass	0.164	kg
Star Tracker Power	0.5	W
CPU Mass	0.100	kg
CPU Power	0.231	W
Rigid Solar Panel Mass	0.270	kg
Rigid Solar Panel Efficiency	30	%
Reaction Control Wheel Mass	0.360	kg
Reaction Control Wheel Power	0.5	W
Battery Energy Density	125	W-hr/kg
Battery Depth of Discharge	99	%
Battery Transmission Efficiency	98	%
Mass Margin	20	%

Table B-4: List of propulsion system component parameters and assumptions.

Parameter	Value	Units
Propellant Tank Starting Temperature	300	K
Propellant Tank Safety Factor	2	--
Xenon Molar Mass	131.3	kg/mol
Universal Gas Constant	8.314	J/mol-K
Kevlar Density	1,440	kg/m ³
Kevlar Maximum Stress	3.620	GPa
Specific Impulse	3,000	s

Table B-5: List of assumed instrument data rates.

Parameter	Value	Units
Magnetometer Data Rate	25	bps
Fast Particle Analyzer Data Rate	1,000	bps
Vector Field Magnetograph Data Rate	25,000	bps

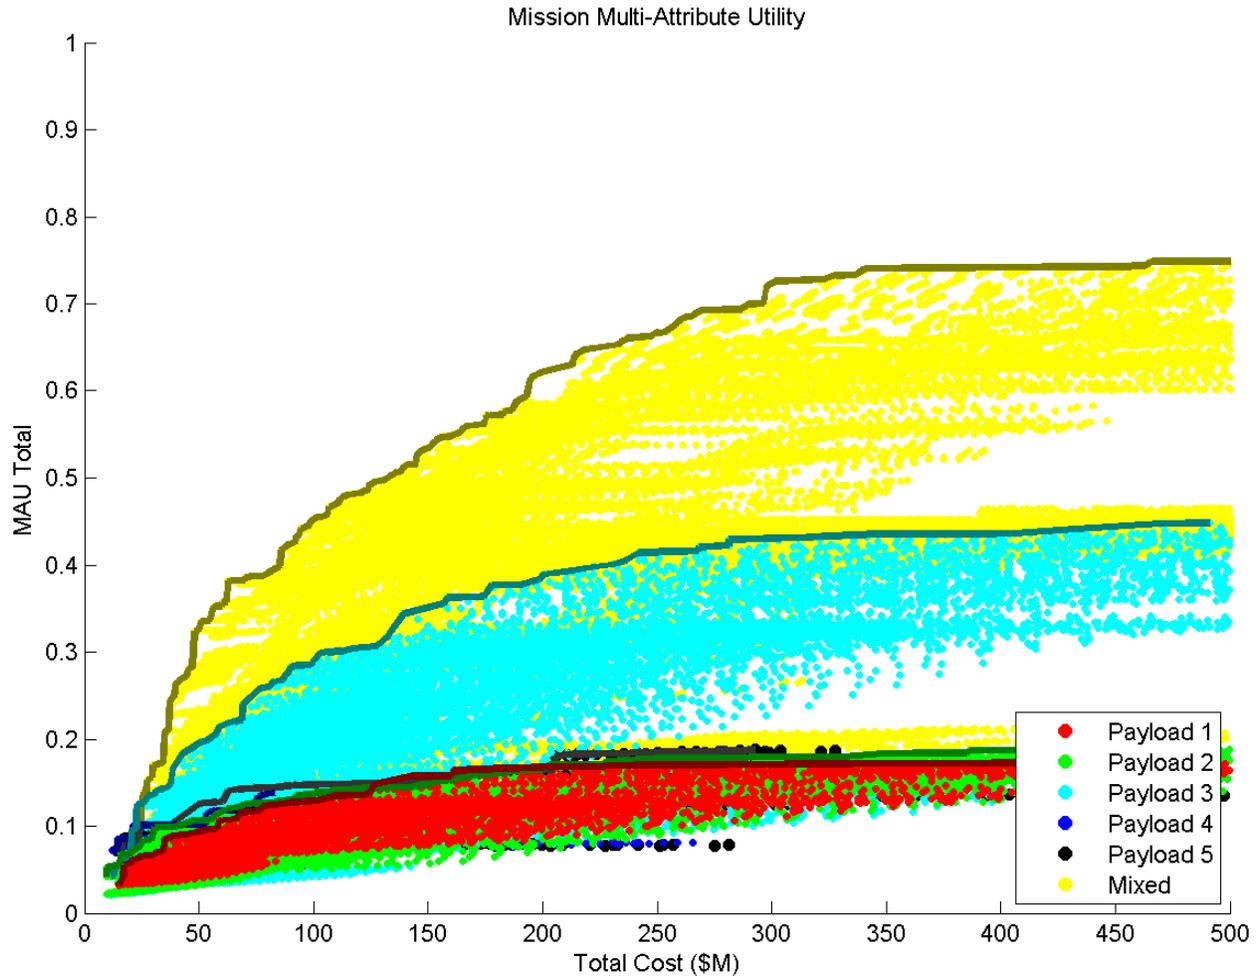


Figure B-1: Enlarged tradespace of HOBOCOP.

Design: 153081 Value Model: 1 Probability of Failure: 0.1 Contacts Per Year: 1

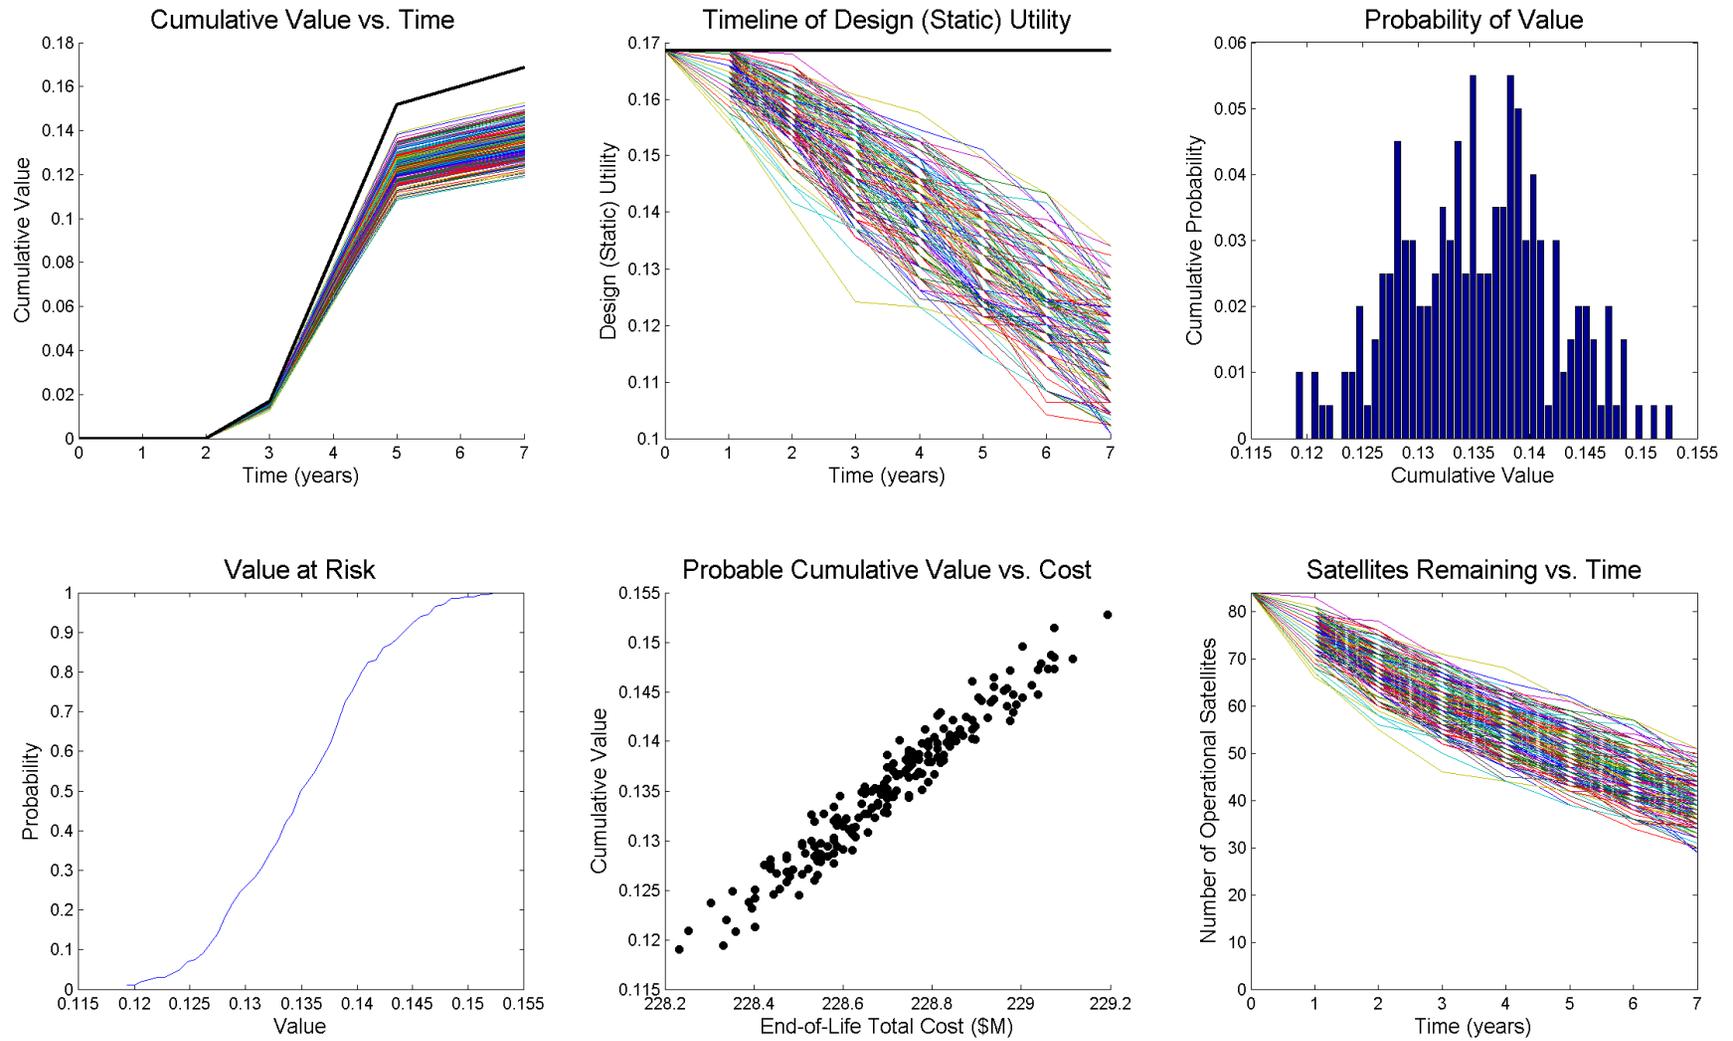


Figure B-2: Era analysis results for a particular design (designation 153081) with one contact per year.

Design: 153081 Value Model: 1 Probability of Failure: 0.1 Contacts Per Year: 2

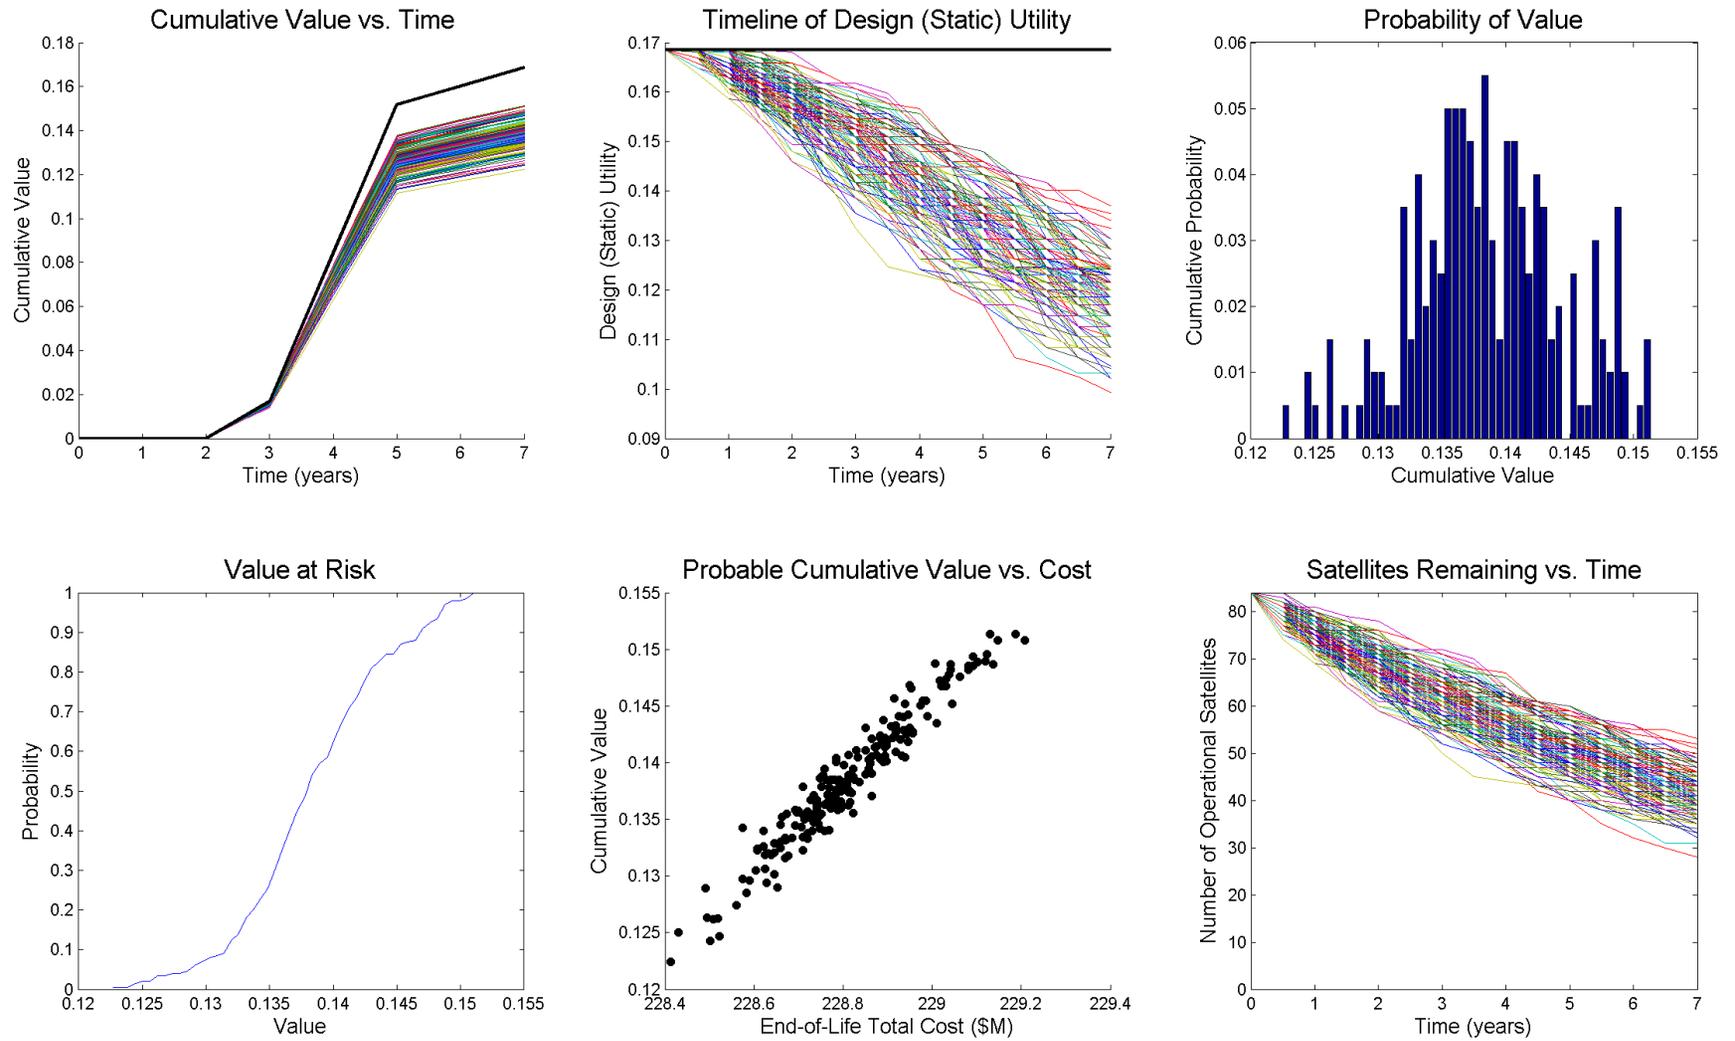


Figure B-3: Era analysis results for a particular design (designation 153081) with two contacts per year.

Design: 153081 Value Model: 1 Probability of Failure: 0.1 Contacts Per Year: 3

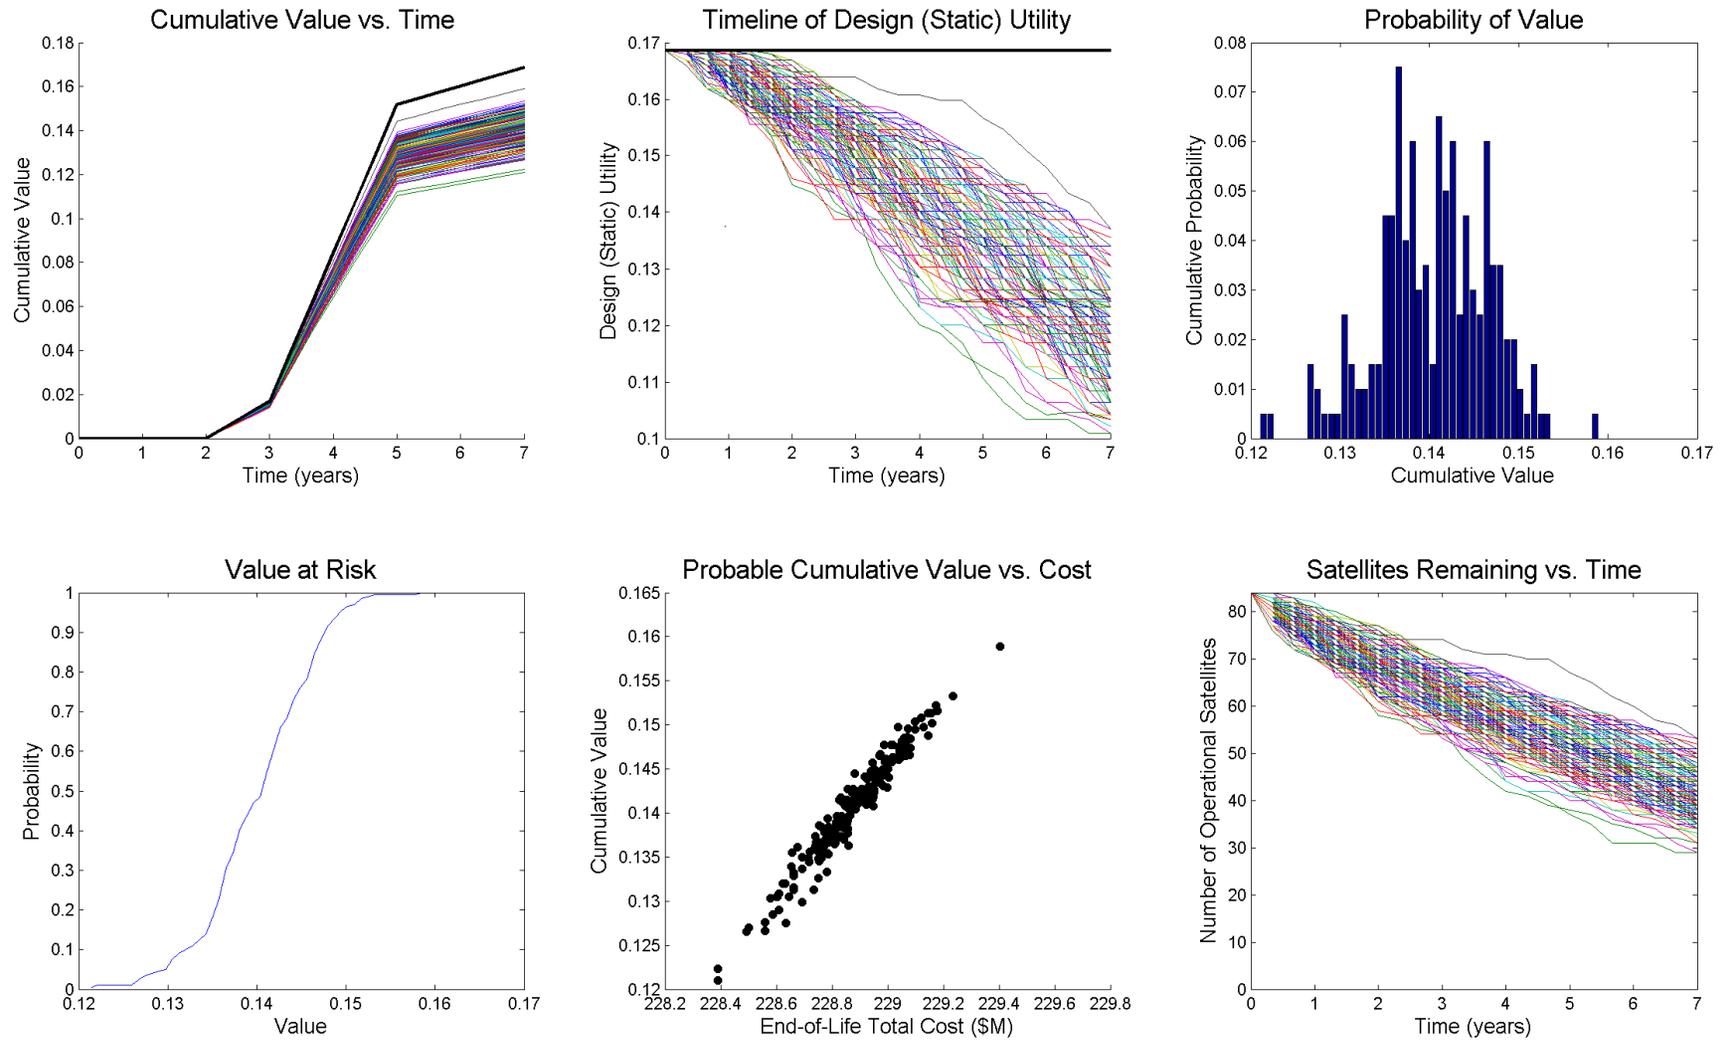


Figure B-4: Era analysis results for a particular design (designation 153081) with three contacts per year.

Design: 153081 Value Model: 1 Probability of Failure: 0.1 Contacts Per Year: 4

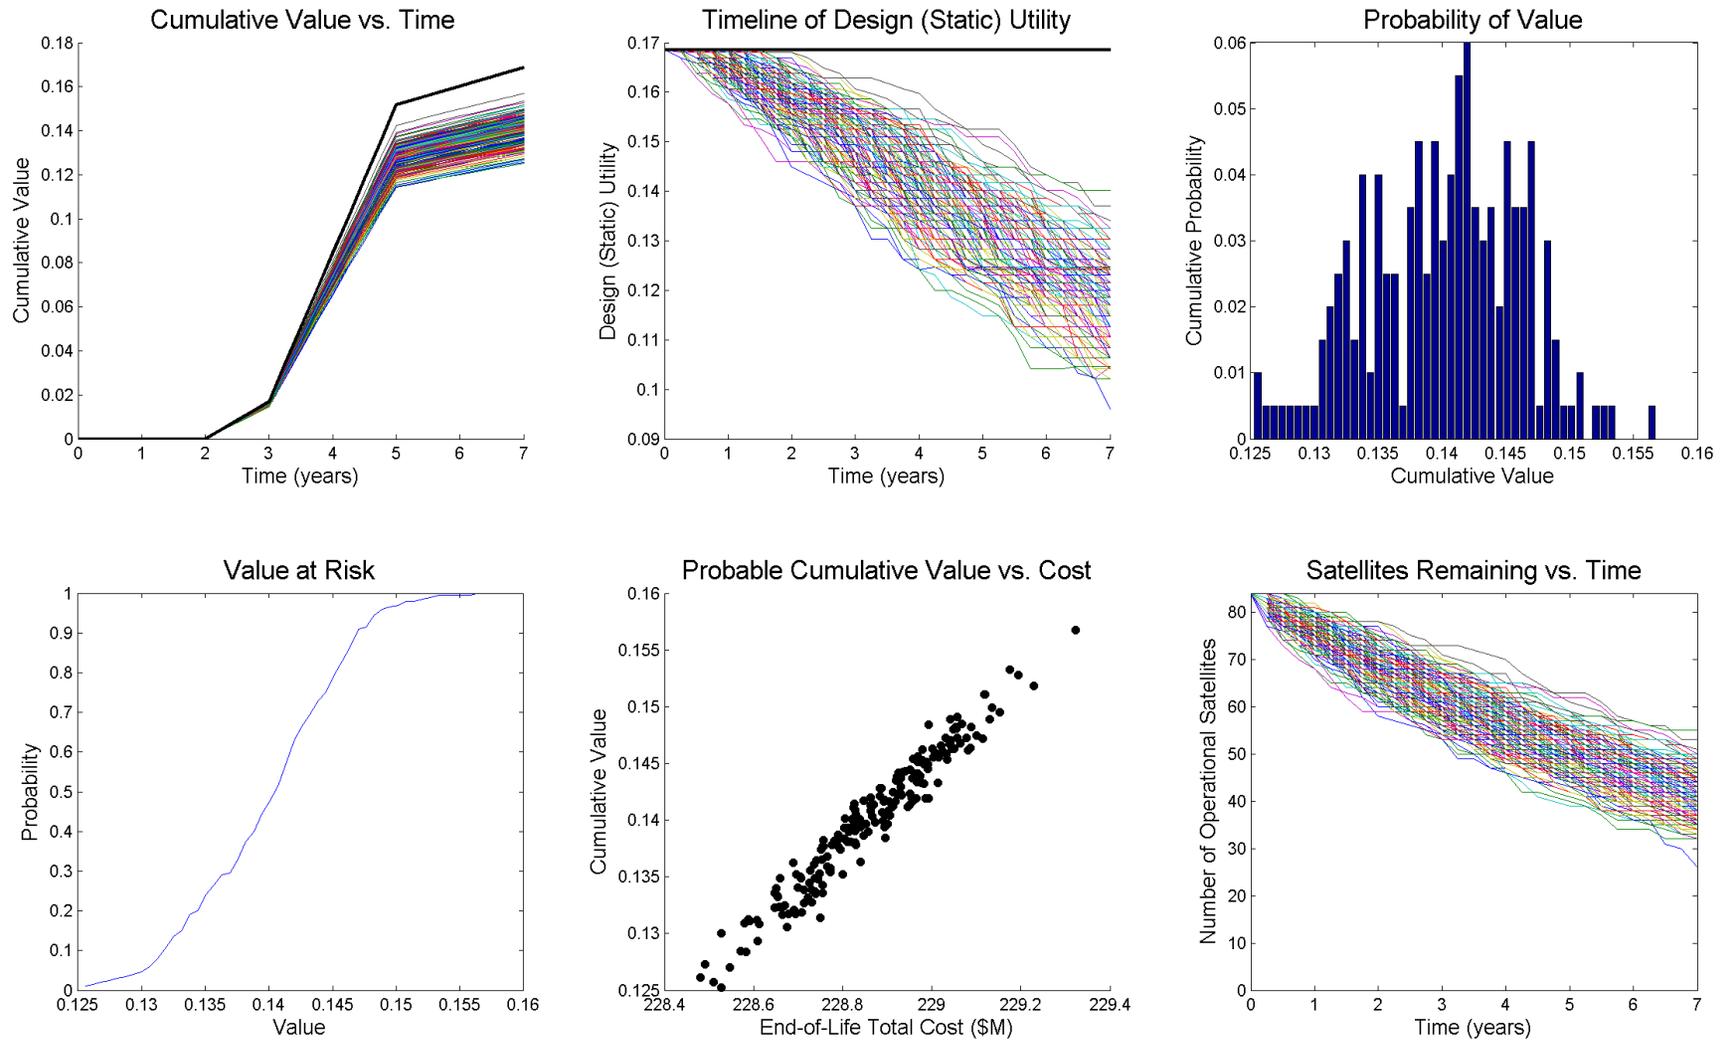


Figure B-5: Era analysis results for a particular design (designation 153081) with four contacts per year.

C. Appendix – GANGMIR

Table C-1: List of natural constants.

Natural Constant	Value	Units
Radius of Mars	3396.2	km
Mass of Mars	6.4174e23	kg
Mars Day Length	24.6597	hours
Mars Year Length	686.973	days
Solar Flux at Mars	593	W/m ²
Mars J ₂ Term	1.96045e-3	--

Table C-2: List of parameters of payloads used to build parametric payload models.

Satellite Payload Parametric Model Inputs	Value	Units
GOMSpace Camera Mass	0.166	kg
GOMSpace Camera Power	0.6	W
GOMSpace Camera Aperture Diameter	3.5	cm
GOMSpace Camera Focal Ratio	1.9	--
ExoplanetSat Camera Mass	0.86	kg
ExoplanetSat Camera Power	2.3	W
ExoplanetSat Camera Aperture Diameter	8.5	cm
ExoplanetSat Camera Focal Ratio	1.4	--
Planet Labs Camera Mass	2.3	kg
Planet Labs Camera Power	10	W
Planet Labs Camera Aperture Diameter	9	cm
Planet Labs Camera Focal Ratio	12.66	--
Mars Climate Sounder Mass	9	kg
Mars Climate Sounder Power	11	W
Mars Climate Sounder Aperture Diameter	4	cm
Mars Climate Sounder Number of Bands	9	--

Table C-3: List of parameters for payload performance

Satellite Payload Parameters	Value	Units
Camera Pixels Per Sweep	2048	--
Camera Operational Wavelength	600	nm
Camera Image Quality	0.4	--
Camera Pixel Size Limit	5	μm
Radiometer Focal Ratio	1.7	--
Radiometer Field of Interest	120	Km
Radiometer Image Quality	1.1	--
Radiometer Integration Time	2	s
Radiometer Bits Per Sample	16	bits/sample

Table C-4: List of satellite components and assumptions.

Satellite Components	Value	Units
Bus Power	5	W
CPU Mass	0.5	kg
Reaction Control Wheel Power	2.8	W
Reaction Control Wheel Mass	0.36 (+0.12/6 kg of total mass)	kg
Star Tracker Power	0.5	W
Star Tracker Mass	0.158	kg
Solar Panel Efficiency	0.3	--
Solar Panel Mass Per Unit Area	5.4	kg/m ²
Antenna Diameter (Gain ~ 3dB)	0.05	m
Antenna Mass Per Unit Area	2.94	kg/m ²
Transmitter Mass Per Unit Power	0.217	kg/W
Power Converter Mass Per Unit $\sqrt{\text{Power}}$	0.38	kg/W ^{1/2}
Transmission Efficiency	60	%
Time Per Orbit in Sunlight	50	%
Operational Duty Cycle	100	%
Battery Depth of Discharge	50	%
Battery Energy Density	125	W-hr/kg
Battery Transmission Efficiency	98	%
Engine Mass	0.5	kg
Engine Specific Impulse	280	S
Structural Mass Fraction (< 5kg subtotal)	7	%
Structural Mass Fraction (> 5kg subtotal)	13	%
Thermal Mass Fraction	3	%

Table C-5: List of mothership components and assumptions.

Mothership Components	Value	Units
Additional Payload Mass	20	kg
Bus and Additional Payload Power	300	W
Transmitter Power	10	W
Transmitter Diameter	2	m
Reaction Control Wheel Mass	20	kg
Structural Mass Fraction	12	%
Thermal Mass Fraction	5	%
Heat Shield Mass Fraction	5	%
Engine Mass	20	kg
Engine ISP	300	s

Table C-6: List of mothership arrival parameters.

Mothership Arrival Parameters	Value	Units
Orbital Inclination	70	deg
Hohmann Transfer Insertion Velocity	2.649	km/s
Aerobraking Velocity Change	50	%
Solar Flux at Mars	593	W/m ²

Table C-7: List of penetrator components and assumptions.

Penetrator Components	Value	Units
Drill Mass	0.1	kg
Drill Energy	5	W-hr
Mass Spectrometer Mass	0.75	kg
Mass Spectrometer Energy	30	W-hr
X-Ray Spectrometer Mass	0.26	kg
X-Ray Spectrometer Energy	20	W-hr
Volatile Spectrometer Mass	0.75	kg
Volatile Spectrometer Energy	15	W-hr
Thermometer/Barometer/Hygrometer Mass	0.13	kg
Thermometer/Barometer/Hygrometer Power	0.045	W
Magnetometer/Seismometer/Camera Mass	0.53	kg
Magnetometer/Seismometer/Camera Power	0.37	W
Dust Camera Mass	0.25	kg
Dust Camera Power	0.1	W
Accelerometer Mass	0.07	kg
Accelerometer Energy	0.5	W-hr
Transmitter Mass	0.1	kg
Transmitter Energy (Short-term)	5	W-hr
Transmitter Power (Long-term) (Average)	0.0243	W
CPU Mass	0.05	kg
CPU Energy (Short-term)	5	W-hr
CPU Power (Long-term)	0.05	W
Primary Battery Energy Density	430	W-hr/kg
Secondary Battery Energy Density	125	W-hr/kg
Secondary Battery Depth of Discharge	75	%
Solar Panel Efficiency	20	%
Solar Panel Mass Per Unit Area	5.4	kg/m ²
Battery Transmission Efficiency	98	%
Solar Panel Duty Cycle (% of Time Charging)	25	%
Penetrator Diameter	12	cm
Penetrator Shroud Mass	3	kg
Penetrator Shroud Diameter	60	cm
Penetrator Nosecone Length	24	cm

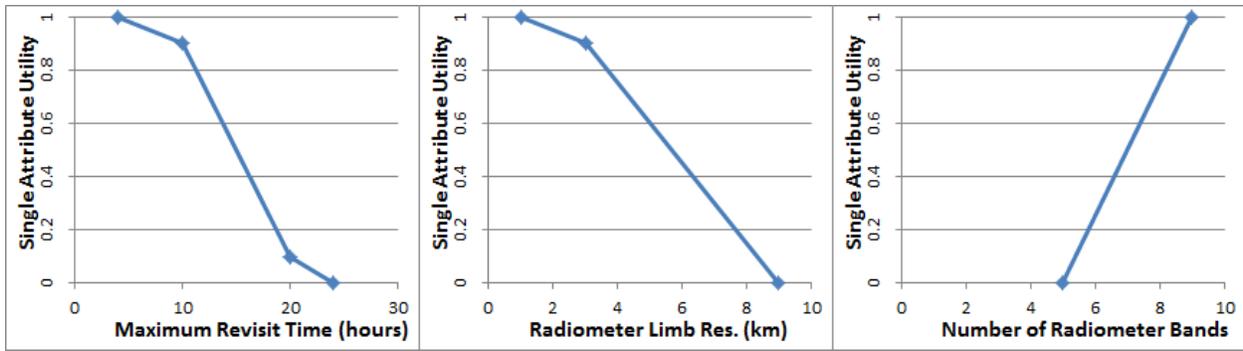


Figure C-1: SAU functions for Science Goal #3 GANGMIR (Observe Global Aerosol Composition)

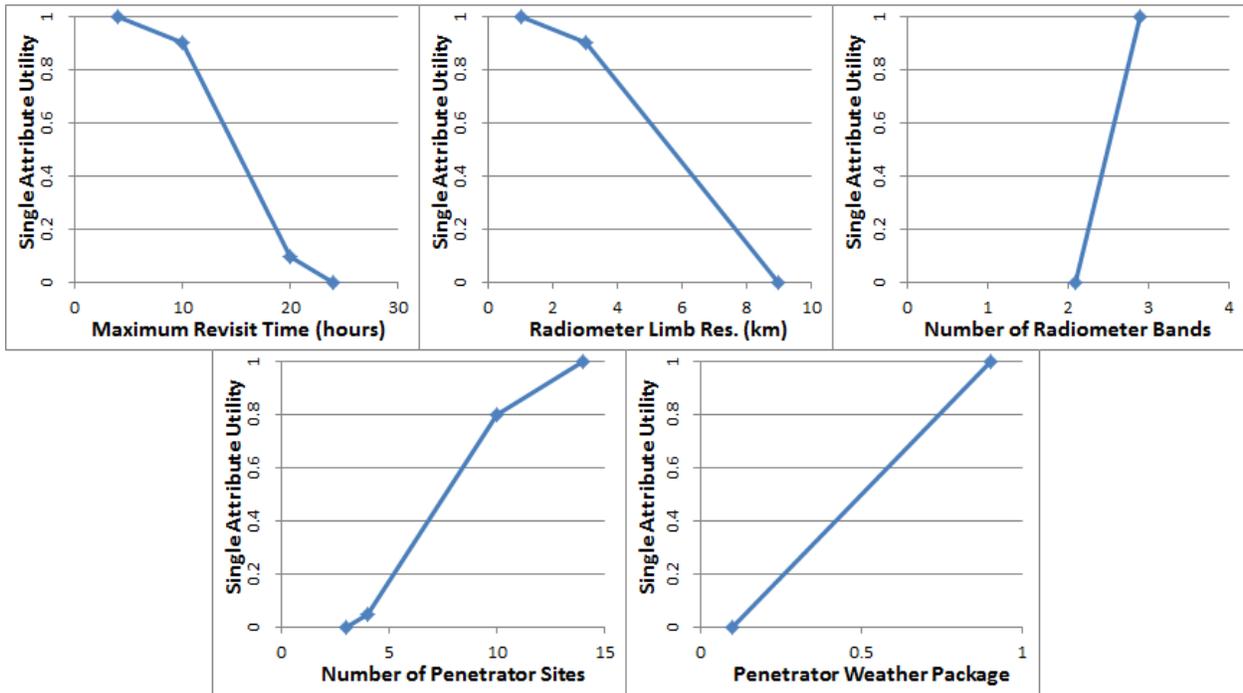


Figure C-2: SAU functions for Science Goal #4 in GANGMIR (Measure Global Surface Pressure).

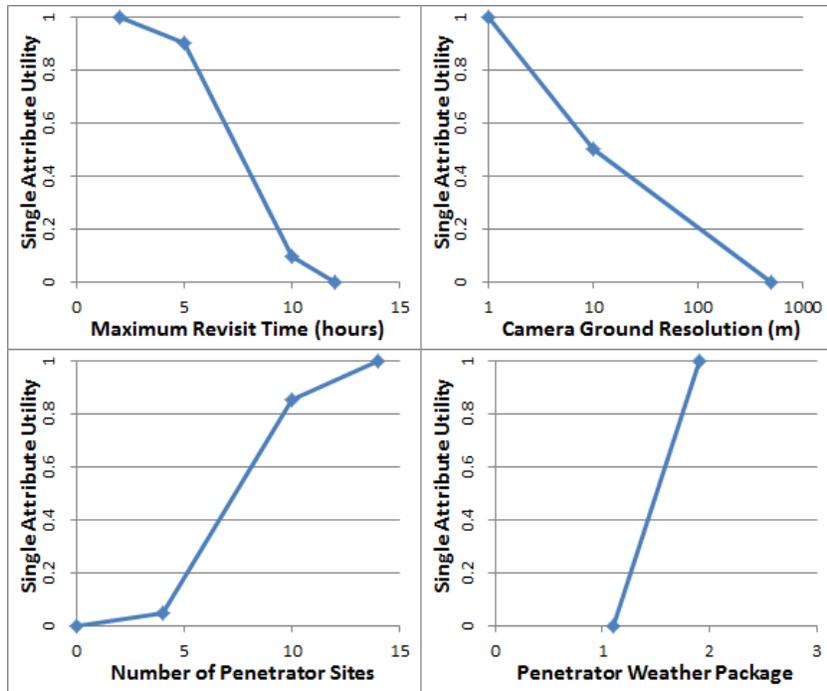


Figure C-3: SAU functions for Science Goal #5 GANGMIR (Observe Local/Regional Weather).

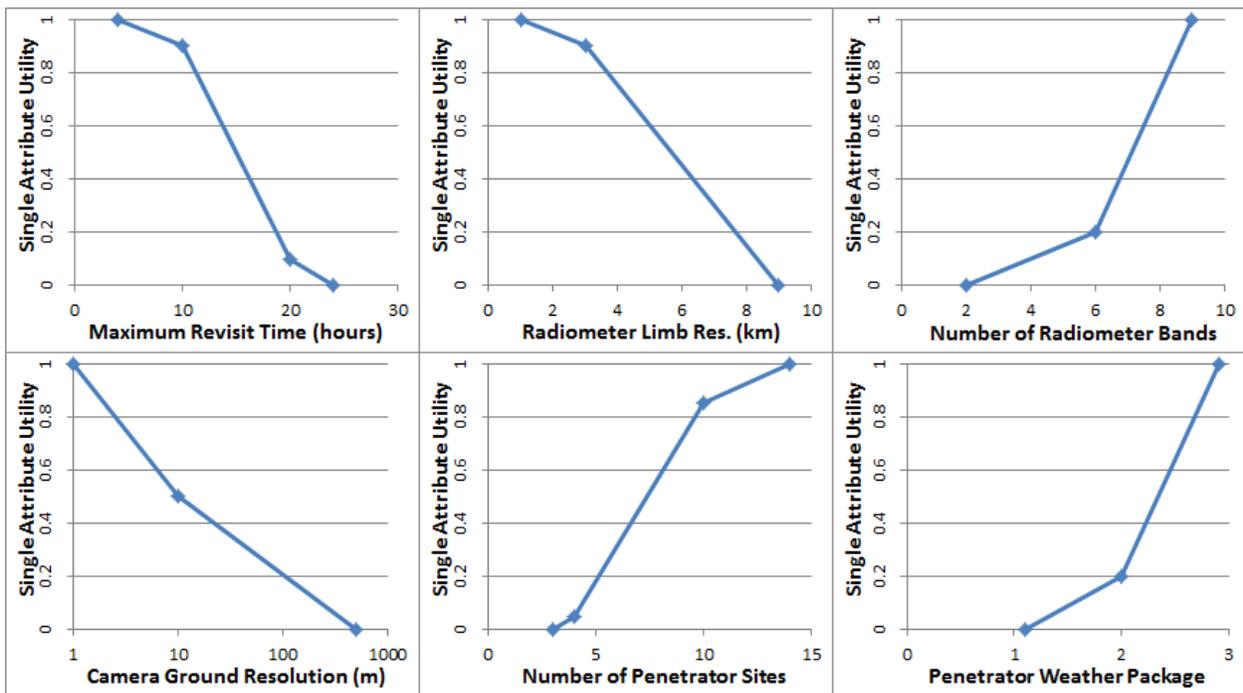


Figure C-4: SAU functions for Science Goal #6 in GANGMIR (Observe Dust and Aerosol Activity).

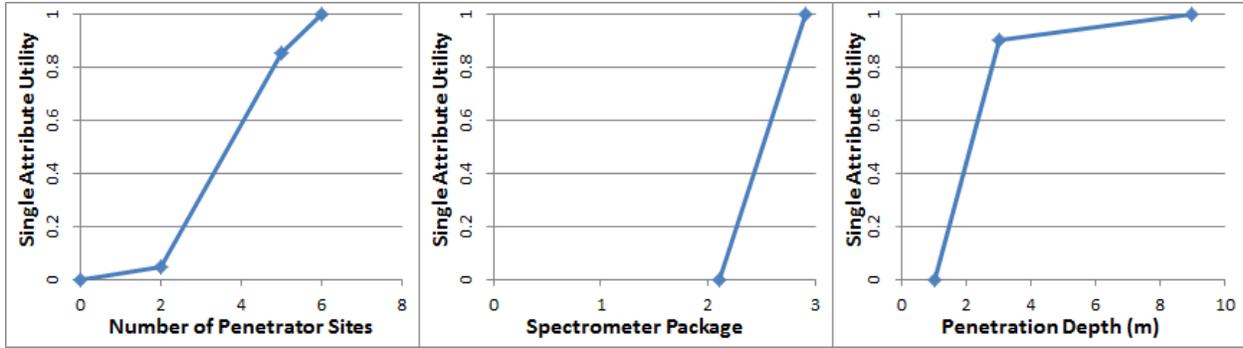


Figure C-5: SAU functions for Science Goal #7 in GANGMIR (Measure Presence of Ground Ice).

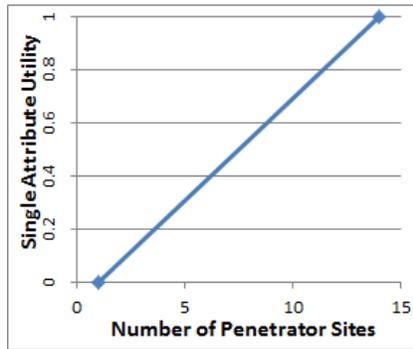


Figure C-6: SAU functions for Science Goal #8 in GANGMIR (Obtain Atmospheric EDL Profiles).