

# Multi-Epoch Analysis of a Satellite Constellation to Identify Value Robust Deployment across Uncertain Futures

Andrew A. Rader<sup>1</sup>  
*SpaceX, Hawthorne, CA, 90250*

*and*

Adam M. Ross<sup>2</sup> and Matthew E. Fitzgerald<sup>3</sup>  
*SEArI, Massachusetts Institute of Technology, Cambridge, MA, 02139*

The value of a system depends heavily on the future contexts it will encounter. For complex space systems with multi-year design and deployment phases, it is useful to design a system so that it delivers value to stakeholders over a wide range of future contexts. Epoch-Era Analysis, a computational scenario planning approach, decomposes the lifecycle of a system into sequential epochs that each have fixed contexts and value expectations. This paper applies Multi-Epoch Analysis (a subset of Epoch-Era Analysis) along with Multi-Attribute Tradespace Exploration (MATE) to the design of a satellite constellation, with the aim of maximizing value across a range of end-user subscription and geographic distribution contexts. The system level tradespace is assembled using a bottom-up iterative approach based on expert knowledge, and accounts for performance attributes metrics such as revisit times, data latencies, observation times, and data downlink volumes. Competing designs consisting of alternative orbital, ground station location, and deployment configurations are evaluated in terms of their fuzzy Normalized Pareto Trace (fNPT) across epochs. The resulting staged deployment strategy delivers robust value based on stakeholder preference across a wide range of future contexts.

## Nomenclature

$D_Z$	=	deployment design variable index $Z$
$fNPT$	=	fuzzy normalized pareto trace
$G_Y$	=	ground station design variable index $Y$
$K$	=	fuzziness factor
$MAU$	=	multi-attribute utility
$NPT$	=	normalized pareto trace
$O_X$	=	orbit design variable index $X$

## I. Introduction

Multi-attribute tradespace exploration (MATE) examines the tradeoffs amongst diverse early design choices and their impacts on lifecycle cost and performance. The result is an improved understanding of the relative impacts of decisions on expected utility (a measure of perceived benefit to a stakeholder) and cost. In the real world, however, neither mission context, nor stakeholder preference, are expected to remain fixed throughout the lifecycle of a project. Since the definition of utility is expected to change, sustained value delivery to stakeholders depends on

---

<sup>1</sup> Mission Integrator, 1 Rocket Rd.

<sup>2</sup> Research Scientist, Engineering Systems Division, 77 Massachusetts Ave., E38-574, AIAA Senior Member

<sup>3</sup> Research Assistant, Aeronautics and Astronautics, 77 Massachusetts Ave., E38-555, AIAA Student Member

the ability of a design to be robust to variations in a wide range of exogenous uncertainties. This problem is analogous to trying to optimize a system while faced with objective functions that change over time: any optimized design is only optimum for a particular set of conditions.

Designs can take advantage of several means of ensuring that value delivered remains high regardless of future context. *Changeability* is the ability to efficiently alter either the physical design parameters or operations of the system at any point during the system's design, development, or operational lifecycle<sup>1</sup>. For example, the ability to efficiently redesign a subsystem to meet a requirements update is an example of design-phase changeability, whereas the ability to efficiently change orbit during operations is an example of operational-phase changeability. There are several types of changeability; for example Fricke and Schulz (Ref. 2) identified four key types: *flexibility*, *adaptability*, *agility*, and *robustness*. Changeability can be incorporated into the design process and/or the design itself<sup>3</sup>. Flexibility and adaptability both refer to the ability to alter a system in response to exogenous change; however, while flexibility relies on change agents external to the system, adaptability requires the system itself to act as the change agent. Agility refers to the degree to which a system can change in a short time span. Other forms of changeability include *modifiability* and *scalability*, which refer to the ability to add or remove design parameters, and to change the magnitude of parameters, respectively.

A system is value robust if it continues to deliver high value to stakeholders across a wide range of changing operational environments and dynamic contexts<sup>4</sup>. Passive value robustness delivers value while maintaining its design parameters, while dynamic value robustness permits changes in the system<sup>5</sup>. Value is here defined as a measure of benefit per unit cost for established stakeholder preferences, which may be uncertain for future contexts. Myriad exogenous factors may impact the value delivered by a system, such as the current set of stakeholders and their individual preferences, future market demand, resources available to deploy the system, the competitive environment, technological advances, and the regulatory environment.

For satellite systems, some elements of changeability can be incorporated into the design process, for example, design flexibility in the hardware elements of certain subsystems. Additionally, design changeability lends itself especially well to software, firmware, ground segment, and some operational elements. However, all changeability comes at a cost, and maintaining significant degrees of changeability in the hardware elements of an operational satellite system can be prohibitively expensive in terms of mass, volume, power, or design cost.

This study considers a constellation with 1-9 microsattellites in low-Earth orbit performing a data relay communications RF mission. The satellites are in contact with globally distributed ground stations for operational telecommanding and data downlink, and are tasked with observing in two predefined geographic regions (one low-latitude, one high latitude), in addition to targets which are globally distributed. The focus is on low-cost modular satellite platforms based on existing bus and payload designs. The satellites will be in contact with a ground-based operations team on a regular basis (generally once per orbit, and once per day at minimum), so system changes can be implemented on a regular basis as part of normally scheduled operations. The requirement to use existing technology with few modifications places restrictions on the changeability that can be incorporated into the system itself. Apart from the ability to enter a "safehold" mode where a satellite would await commands from the ground, adaptability is not required for the system. Additionally, the exogenous factors we will consider in this study (such as geographic and application usage contexts) change slowly compared with system changes that can be directly implemented via telecommand, so agility is not a major concern in this system.

The design variables we consider for the satellite constellation are related to deployment of the system (orbit, deployment schedule) and ground station architecture. Moreover, the restriction on using existing hardware means that once the satellites are deployed, only minor changes can be implemented (e.g., software, hardware, RF bands, firmware), or changes must be planned for in advance of launch. However, especially because the system would provide a new service, the exogenous environment in terms of usage context application, geographic location, and volume is highly uncertain. Thus, we are primarily concerned with designing the system for passive value robustness across uncertain usage contexts, with some consideration of flexibility where it can be included at a reasonable cost.

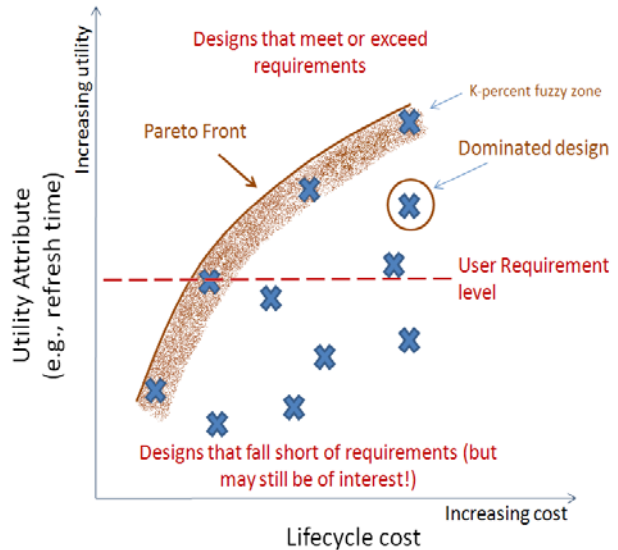
## II. Populating the Tradespace

Populating the tradespace for a satellite constellation system involves balancing the performance of competing design alternatives, each of which has implications for the mission that are difficult to quantify independently. For an earth observation mission with requirements for two-way communications, global coverage, high observation times, frequent revisits, and large downlink volumes, it is not clear how these attributes should be traded off versus each other or versus costs. Moreover, most of the requirements derive from multiple subsystems in the design of the constellation, from hardware mounted on each individual satellite, to satellite spacing and ground communication infrastructure. It is not straightforward to span the tradespace by simply selecting design variables at only the system or subsystem level. In order to mobilize domain experts and their tools, we can perform the individual subsystem trades in parallel and assess their interactions at the system level at key nodal checkpoints<sup>6</sup>. This bottom-up approach allows for an independent partitioning of the design space by expertise (e.g., orbit, ground, launch, payload, etc.) in order to assemble a system level tradespace that captures the individual subsystem level trades.

Moreover, it is not clear how attributes should be valued. Since the mission level attributes are time varying, how should peak performance be valued versus mean, median, or worst case performance? How should global coverage be valued versus coverage over a particular area of interest? How much value is derived from additional ground stations as compared with additional satellites? We can examine these factors and account for stakeholder preference by employing tradespace exploration.

Tradespace exploration provides a mathematical construct by which the subjective utility of competing designs is assessed on a common basis (e.g., stakeholder utility vs. lifecycle cost). Since tradespaces are constructed using concept neutral criteria (perceived benefits and costs), they allow for a comparison of vastly different concepts on the same basis<sup>7</sup>. Several features are highlighted by the example tradespace shown in Fig. 1:

- Increasing *value* is comprised of both increased utility (increasing along the y-axis) and lower cost (decreasing along the x-axis).
- The Pareto Front represents designs that have the highest utility for a given cost, or the lowest cost for a given utility.
- Designs falling closer to the Pareto Front are higher in value. Designs falling farther below the Pareto Front are dominated by higher value designs, yielding less utility for the same cost.
- Designs falling in the K-percent fuzzy zone are considered K-percent fuzzy efficient.
- If the objective is simply to meet a requirement (and there is no desire to exceed it), a location near the intersection of the Pareto Front and the user requirement level yields the most efficient implementation.
- Since multi-attribute utility (MAU) captures contributions from all attributes, there may be cases where a design that performs worse for one attribute is preferred based on its overall utility.



**Figure 1. Example tradespace (each x indicates a unique point design)**

The key factors represented in a tradespace are vectors of *design variables* (i.e., the tradable features of alternate potential designs) and estimated performance (*utility*) *attributes* (i.e., design performance with respect to stakeholder preference). Design variables at the system level (e.g., orbit) can be decomposed into *subsystem design variables* specific to one or more subsystems (e.g., ‘orbit’ can be decomposed into ‘plane’, ‘inclination’, ‘spacing, etc.). Analysis of the coupling between design variables in the tradespace is achieved using iterations, during which each subsystem is analyzed in relative isolation by domain experts, but at the completion of which, the next iteration is seeded by analyses from all subsystems during the prior iteration. The impacts of the design variables on the utility attributes are assessed at the system level.

Utility attributes are quantitative criteria, which are representations of stakeholder preference; each one, derived from needs identification and preference elicitation, captures the magnitude of one aspect of perceived utility<sup>8, 9</sup>. Attributes are bounded by a stakeholder-defined range from the least acceptable performance (beyond which the system is useless to the stakeholder) to the most desirable performance (beyond which there is no extra benefit). A set of attributes should be complete, non-redundant, operational, decomposable, minimal, and perceived independently. Satisfaction for performance at different levels of attributes can be assessed using single attribute utility functions that typically vary from 0 (minimally acceptable) to 1 (most desirable)<sup>10</sup>.

For the satellite constellation application, the first step was the definition of utility-driving attributes based on user requirements (Fig. 2, top left). Next, the design team, composed of five expert teams divided by subsystem (orbit, payload, bus, ground, and launch), defined relevant design variables at their subsystem level, and the relationship between the subsystem and system was quantified (Fig. 2, bottom left). The third step was to examine the resulting tradespace in terms of the relative utility attribute vs. cost of the resulting potential designs (Fig. 2, middle left). Finally, design performance in terms of stakeholder preferences is explored and synthesized using Multi-Attribute Tradespace Exploration (MATE)<sup>11</sup>.

Once the tradespace has been constructed, we can apply Epoch-Era Analysis (EEA) to examine the impact of exogenous factors on design performance for alternative future contexts.

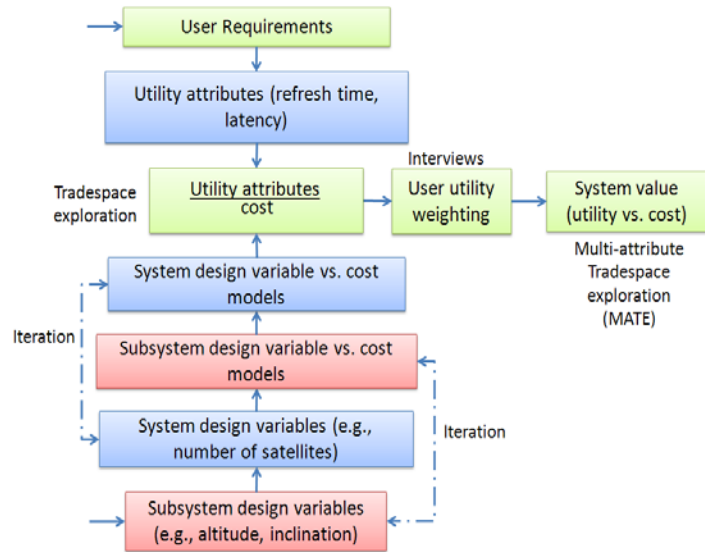


Figure 2. Tradespace exploration flow<sup>6</sup>

### A. Satellite Constellation Utility Attributes

For the satellite constellation application, utility attributes were derived from a decomposition of user requirements into quantifiable metrics (one for each major quantifiable requirement, e.g., revisit time, space to ground data latency, observation time, data downlink time). All of these attributes are impacted by more than one subsystem, and thus their performance can only be quantified at the system level. For example, data downlink time depends on characteristics of the orbit, spacecraft, and ground station infrastructure, such as minimum elevation angle, frequency band, gain, noise, and ground station location distribution. Additionally, some performance metrics are time varying: for example, revisit time varies by individual satellite, and by orbital parameters, orbital epoch, location of interest on the Earth, and other factors such as orbital maneuvers and equipment duty cycle. For such cases, it is useful to define an attribute by weighting two or more performance relevant metrics (e.g., 'low end ~ 90<sup>th</sup> percentile' and 'worst case' performance). Also note that the satellite constellation provides service for 3 distinct geographic regions (region 1, region 2, and global), and performance for each attribute is expected to vary by region.

Studies have shown that having too many attributes (e.g., more than 5 to 7) per stakeholder makes it difficult to accurately elicit stakeholder preferences when all attributes are considered simultaneously<sup>5</sup>. Thus, it is probably a good idea to limit the number of system level attributes considered per stakeholder to fewer than seven, and ideally no more than five. For this application, four primary system level attributes were considered (see Table 1). Each of the utility attributes consists of the worst-case and 90<sup>th</sup> percentile worst-case performance weighted evenly.

**Table 1. Utility attributes and contributing subsystems for the satellite constellation**

Utility attribute	Contributing subsystems
Revisit time	<ul style="list-style-type: none"> <li>• Orbit (time to revisit anywhere on the globe)</li> <li>• Launch (only as an orbit driver)</li> <li>• Distribution of targets on Earth</li> </ul>
Data latency	<ul style="list-style-type: none"> <li>• Orbit (time to revisit the selected area)</li> <li>• Launch (only as an orbit driver)</li> <li>• Ground infrastructure (time to revisit a ground station after collecting data)</li> <li>• Downlink rate</li> <li>• Ground-based data distribution networks</li> </ul>
Observation time	<ul style="list-style-type: none"> <li>• Orbit</li> <li>• Launch (only as an orbit driver)</li> <li>• Payload (footprint, duty cycle)</li> </ul>
Data downlink time	<ul style="list-style-type: none"> <li>• Orbit (ground station contact time)</li> <li>• Launch (only as an orbit driver)</li> <li>• Ground infrastructure (location, elevation angle)</li> </ul>

**B. Satellite Constellation Design Variables and their Associated Cost Drivers**

Many design variables are composed of contributions from multiple subsystem design variables. Most of the design variables have associated cost drivers. Some cost drivers can be determined at the subsystem level, while others can only be quantified at the system level. Table 2 shows the design variables considered for the satellite constellation application and their associated cost drivers.

**Table 2. Satellite Constellation design variables and their associated cost drivers**

System design variables	Subsystem design variables	Associated cost drivers
Spacecraft bus	<ul style="list-style-type: none"> <li>• Bus selection</li> <li>• Controlled/uncontrolled</li> </ul>	<ul style="list-style-type: none"> <li>• Design &amp; build</li> </ul>
Orbital characteristics	<ul style="list-style-type: none"> <li>• Orbit planes</li> <li>• Eccentricities</li> <li>• Separation</li> <li>• Special cases (e.g., sun-synchronous)</li> </ul>	<ul style="list-style-type: none"> <li>• Launch</li> <li>• Propulsion &amp; delta-v</li> </ul>
Deployment	<ul style="list-style-type: none"> <li>• Number of satellites deployed</li> <li>• Fixed deployment schedule vs. flexible deployment path</li> </ul>	<ul style="list-style-type: none"> <li>• Design &amp; build</li> <li>• Launch</li> <li>• Operations</li> </ul>
Launch vehicle	<ul style="list-style-type: none"> <li>• Number of satellites per launch</li> <li>• Primary vs. secondary payload</li> <li>• Replacement availability</li> </ul>	<ul style="list-style-type: none"> <li>• Launch vehicles</li> <li>• Launch operations</li> </ul>
Payload	<ul style="list-style-type: none"> <li>• Selection of payload for mission</li> </ul>	<ul style="list-style-type: none"> <li>• Equipment design &amp; build cost</li> <li>• Operations</li> </ul>
Downlink	<ul style="list-style-type: none"> <li>• Frequency</li> <li>• Bandwidth</li> <li>• Geographic availability</li> </ul>	<ul style="list-style-type: none"> <li>• Equipment design &amp; build cost</li> <li>• Regulatory issues</li> </ul>
Redundancy, reliability, & replacement strategy	<ul style="list-style-type: none"> <li>• On-orbit (hot or cold) spares vs. replacement</li> <li>• Spacecraft reliability (expected lifespan)</li> </ul>	<ul style="list-style-type: none"> <li>• Initial cost vs. replacement cost</li> </ul>
Ground station	<ul style="list-style-type: none"> <li>• Number and placement of stations</li> <li>• Build vs. buy</li> <li>• Stationary vs. mobile</li> </ul>	<ul style="list-style-type: none"> <li>• Equipment design &amp; build cost</li> <li>• Start-up vs. operational costs</li> </ul>
Data processing & handling	<ul style="list-style-type: none"> <li>• Storage</li> <li>• Security</li> </ul>	<ul style="list-style-type: none"> <li>• Start-up costs</li> <li>• Operational costs</li> </ul>

Note that while many of the design variables considered for this application allow for changeability during the design process, very little changeability is possible once the system is deployed except for certain mission elements

(e.g., ground station infrastructure, re-launch of replacement or additional satellites, and software/firmware upgrades). Thus, in our following study, we are primarily concerned with passive value robustness across alternate future contexts.

For clarity of illustration, this paper will examine only a subset of the design variables considered in the full satellite constellation study. Tables 3 to 5 show the primary orbit, ground segment, and deployment trades considered for this application. The designs considered in this example derive from the full factorial expansion of these subsystem design variables (i.e.,  $\mathbf{O}_X \times \mathbf{G}_Y \times \mathbf{D}_Z$ ), except that orbit index **O4** was always used with the ‘flexible deployment’ deployment (**D’**) cases, and never with the ‘optimized constellation’ deployment (**D**) cases due to restrictions in launching fewer satellites in a single launch by using secondary rideshare launches. This strategy results in lower overall launch costs, but less optimized orbital locations.

Based on the enumerated orbital, ground station, and deployment design variables, we get a total of 22 total designs ( $(3 \times 2 \times 3) + (1 \times 2 \times 2)$ ). Note that since deployment context is time dependent, for the ‘single satellite’ and ‘partial constellation’ deployments, the option exists to upgrade the design into ‘full constellation’ designs within the ‘optimized constellation’ or ‘flexible deployment’ path.

**Table 3. Orbit design variables**

Orbit index	Descriptor	Altitude	Inclination
<b>O1</b>	Near equatorial (+ polar satellites for deployed constellation)	600km	5°
<b>O2</b>	Low inclination (+ polar satellites for deployed constellation)	600km	63°
<b>O3</b>	All polar satellites	600km	97.8°
<b>O4</b>	Mixed – flexible deployment	mixed	Mixed

**Table 4. Ground station design variables**

Ground station index	Descriptor	Ground Station Locations
<b>G1</b>	Minimized ground stations	<ul style="list-style-type: none"> <li>Equatorial case: 1 equatorial ground station</li> <li>Mid-latitude case: 3 mid-latitude ground stations</li> <li>Polar case: 2 polar ground stations</li> </ul>
<b>G2</b>	Maximized downlink	<ul style="list-style-type: none"> <li>7 globally distributed ground stations</li> </ul>

**Table 5. Deployment design variables**

Deployment index	Descriptor	No. sat	No. planes	Launches required
<b>D1</b>	Single Satellite	1	1	1
<b>D2</b>	Partial constellation – optimized constellation	5	3	3
<b>D2’</b>	Partial constellation deployment – flexible deployment	5	5	5 (secondary as available)
<b>D3</b>	Full constellation – optimized constellation	9	5	5
<b>D3’</b>	Full constellation deployment – flexible deployment	9	9	9 (secondary as available)

### C. Example Single Attribute Tradespace

Now that we have defined the utility attributes and design variables, we can plot the tradespace for each individual utility attribute (i.e., revisit time, data latency, observation time, and downlink time). Fig. 3 shows an example single attribute tradespace, in this case for accumulated observation time of the targets of interest. Note that, in addition to each utility attribute, there is a resulting tradespace for each of the selected attribute quantifying metrics (90<sup>th</sup> percentile worst case and worst case performance), and each of the 3 geographic regions (low latitude, high latitude, and global, resulting in  $4 \times 3 \times 3 = 36$  total tradespaces). Additionally, for presentation here, costs have been normalized by the expected cost of the most expensive design.

As expected, observation time increases with increased numbers of satellites deployed. Although the increased observation time appears almost linear, there is in fact a somewhat diminishing return because additional satellites occupy less optimized orbits and sometimes provide redundant coverage that is not included in the accumulated observation time metric (also the cost increase is not quite linear with number of satellites). Increased ground station infrastructure has no effect on accumulated observation time, but does have an effect on some of the other attributes, and also increases the cost of the system. The flexible deployment path generally results in lower performance as a result of the reduced orbital optimization, but this comes at a lower cost due to the opportunity to make use of secondary launch opportunities as they arise. Furthermore, note that both **O1D1** designs have been eliminated from the analysis because region 1 spans into higher latitudes than the satellite is able to access.




Figure 3. Single attribute tradespace of 90<sup>th</sup> percentile worst case observation time for geographic region 1

### D. Example Multi-Attribute Tradespace

Once we have defined the individual attribute tradespaces, we can convert performance in terms of the individual utility attributes into non-dimensional utility based on user preference in order to construct utility vs. cost

tradespaces for each attribute. (This allows for capturing potential perceived nonlinear benefits of different levels of the performance attributes.) Next, we are in a position to construct a multi-attribute tradespace by assigning utility values to the selected attributes and weighting each attribute based on user preference<sup>7</sup>. There are multiple techniques for mapping attributes to utility, including direct elicitation, indirect elicitation, assignment by proxy, and assignment by “educated” justification. Due to the desire to illustrate a representative hypothetical user, this example does not use elicited single and multi-attribute utility functions, but rather assigns single-attribute utilities based on the following assumptions for the satellite constellation problem<sup>6</sup>:

- All attributes are continuous.
- Meeting a requirement achieves a single-attribute utility of 0.5.
- Doubling or exceeding double a requirement is assigned a utility of 1.
- Falling short of a requirement by half or more is assigned a utility of 0.
- Utilities between 0 and 1 are interpolated on a linear scale, giving a utility value of between 0 and 1.

Thus a utility of 0.5 can be interpreted as meeting the requirements “on average”. A utility higher than 0.5 performs better than the requirements “on average”, and a utility less than 0.5 fails to meet the requirements “on average”. However, this does not mean that a utility of higher than 0.5 necessarily achieves all the requirements using this definition of utility.

Following a computation of single attribute utility, the single attribute utilities are assigned swing weights that determine their relative contribution to overall utility *with respect to each other, and to the utility they provide* in order to construct multi-attribute utility. For example, it might be much more utility-generating to achieve a data refresh time requirement than a latency requirement, or vice-versa. Note that this is a separate process from the computation of single attribute utility. This application computed utility weightings based on a simplified conjoint analysis, where users are presented with a menu of options that all have *equal nominal value* (i.e., assuming all attributes should be weighted equally). However, users are unlikely to actually rank nominally equal alternatives equally, and this difference between nominal rankings and user-assigned rankings allows “true” weightings to be elucidated. The resulting utility weightings are shown in table 6. This utility formulation assumes that each attribute contributes utility independently of one another (therefore weights add to 1 and the multi-attribute utility function becomes a simple linear weighted sum of single attribute utilities).

**Table 6. Single attribute utility weightings to construct multi-attribute utility (90<sup>th</sup> = 90<sup>th</sup> percentile worst case; WC = worst case)**

Attribute	Revisit time	Data latency	Observation time	Data downlink time
<b>Weighting</b>	0.5	0.2	0.15	0.15
<b>Breakdown</b>	90 <sup>th</sup> (0.25); WC (0.25)	90 <sup>th</sup> (0.1); WC (0.1)	90 <sup>th</sup> (0.075); WC (0.075)	90 <sup>th</sup> (0.075); WC (0.075)

Now that we have derived utility definitions and weightings, we can assemble the multi-attribute tradespace for each of the 3 geographic regions. Examples are shown in Figs. 4 and 5. Note that for each of the geographic regions, the relationship between cost and performance does not follow the same overall trend, and designs that are favored in one geographic region are not necessarily the same designs that are favored in another. For example, region 1 is located much closer to the equator at latitudes of 15-35° compared with region 2 at 35-60°, so low-inclination orbits (**O2** designs) are preferred for region 1 whereas polar orbits (**O3** designs) perform better for region 2. Additionally, the flexible path deployment provides a slightly better utility vs. cost tradeoff for region 2, because there are more options for polar launches and the region is less sensitive to orbital optimization.



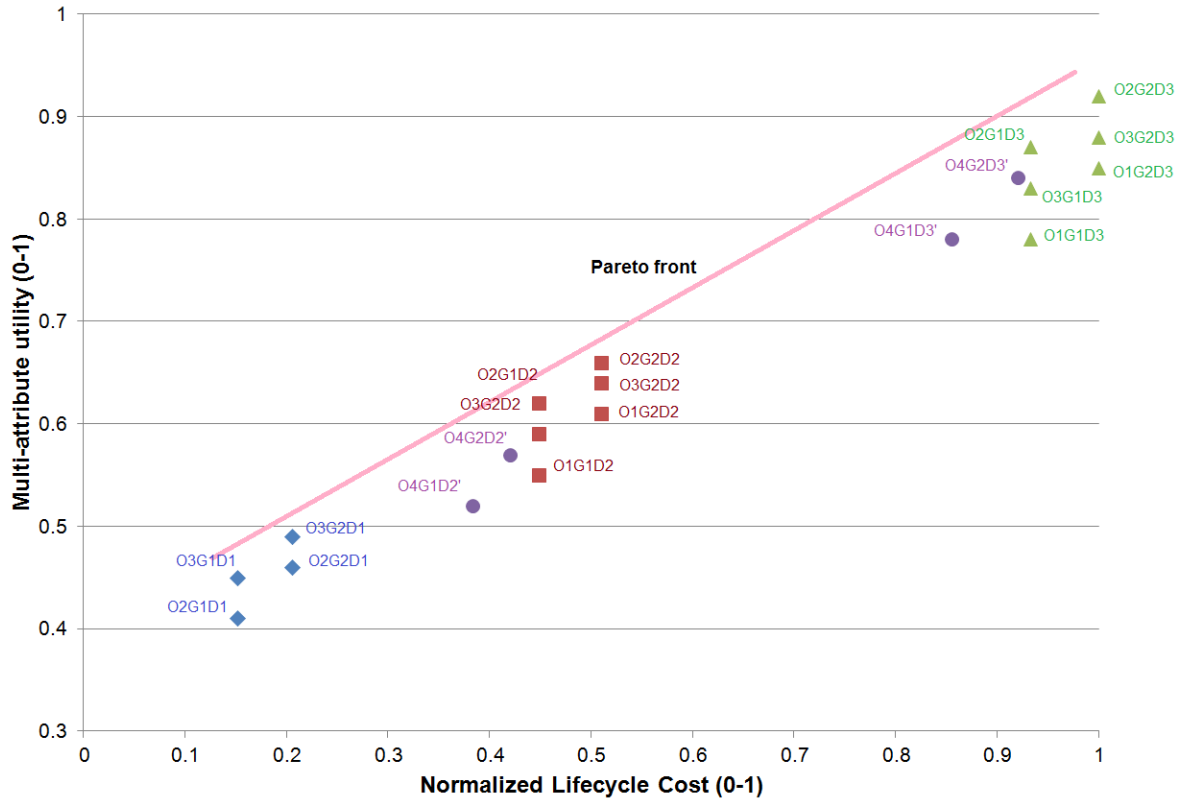


Figure 4. Multi-attribute utility tradespace for geographic region 1

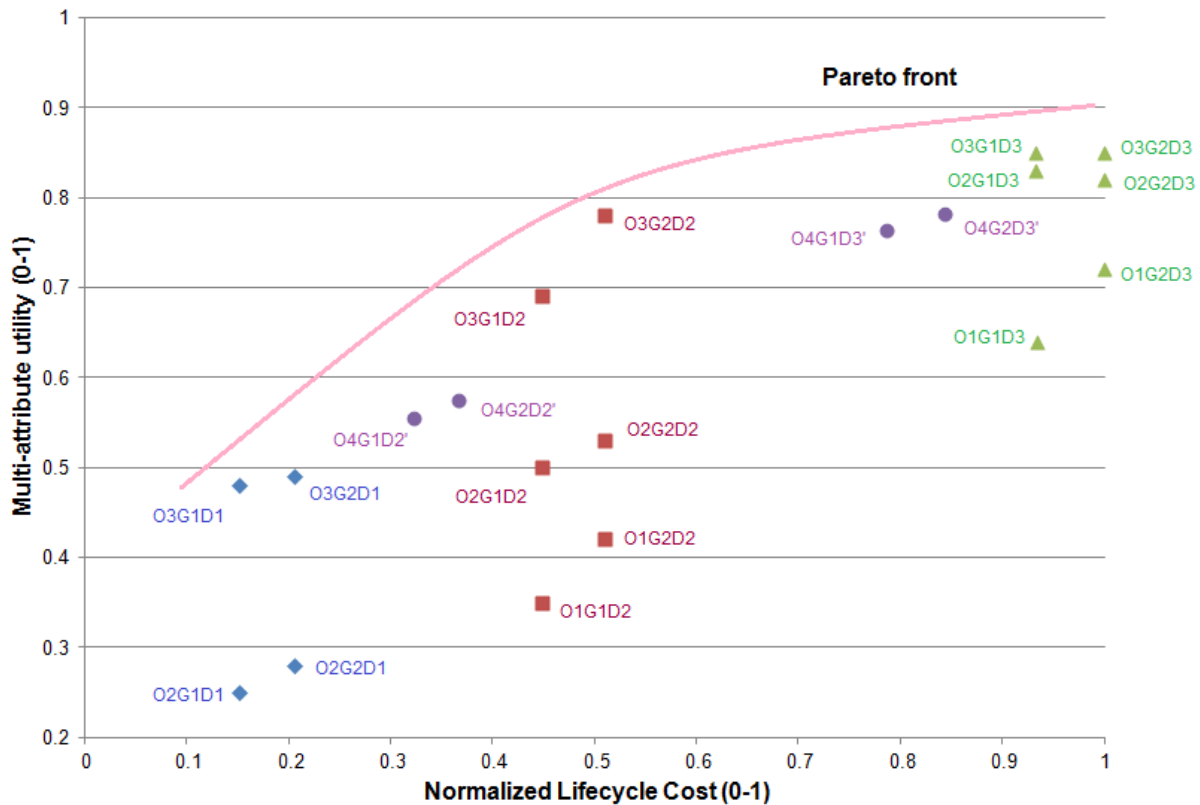


Figure 5. Multi-attribute utility tradespace for geographic region 2

### III. Multi-Epoch Analysis

Epoch-Era Analysis (EEA) is a computational scenario planning framework that is particularly suited to consider the changeability and robustness of a design across many alternative potential futures. EEA provides a structured way to analyze the temporal system value environment in order to determine the expected value distribution given uncertain future contexts<sup>12</sup>. EEA’s temporally ordered structure is designed to account primarily for exogenous sources of uncertainty, particularly those arising from future conditions over the dynamic lifecycle of a project. EEA decomposes the lifecycle of a system (comprising an “era”) into sequential epochs that each have fixed context and value expectations (see Fig. 6). Each epoch consists of a set of value expectations and the key exogenous factors describing the context (see Fig. 7). The epochs comprising a given era can either be pre-defined, “hand-picked” to fit imagined future scenarios of interest, or the process can be automated to construct eras from sets of possible epochs. The structured ordering of epochs creates an intuitive framework that considers full project lifecycles. Comparing system performance (in utility and cost) between eras consisting of different epochs provides insight into value delivery over uncertain, developing futures; however, we can also use a straightforward multi-epoch analysis to gain insight about the point impact on value of each of many possible future contexts.

While EEA often considers time dependence in the ordering of epochs into eras, we can invoke multi-epoch analysis to generate many of the same insights as EEA at a lower level of effort by employing a simple time independent analysis of the utility across selected epochs without combining them into time-ordered eras<sup>1, 13</sup>. This provides a lower effort alternative in absence of sufficient information to build up eras reliably. The lower fidelity associated with multi-epoch analysis as opposed to full EEA means that we will have less information about time-evolution strategies such as the phased deployment case. For example, although we will examine the performance of the system for the three different deployment configurations (single satellite, partial deployment, and full deployment) the staged deployment strategy might not be the “best” path given that path dependency was not considered. When additional information is available about epoch path dependence or ordering, a full EEA can be conducted.

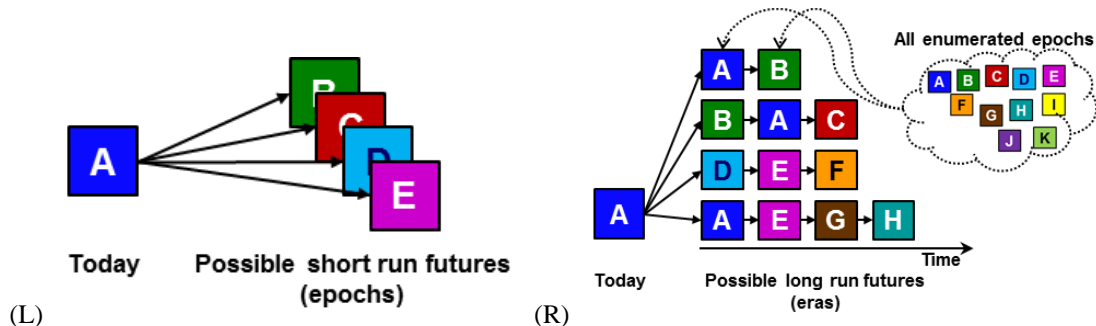
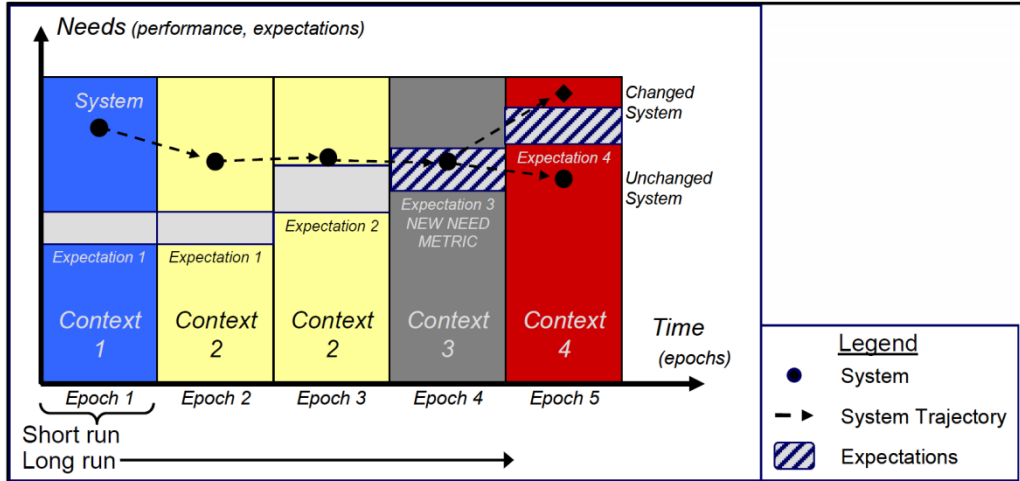


Figure 6. (L) Epochs as possible short run futures, while (R) an era spans a system lifecycle and is subdivided into time sequence of epochs that define alternative future value expectations and contexts.<sup>6</sup>



**Figure 7. A system path can be traced across an era, which is comprised of an ordered sequence of epochs, each of which has a context and set of value expectations.<sup>7</sup>**

Several methods have been developed to examine design efficiency across eras and epochs. The simplest method is to examine the impact of epoch variables on the value as determined directly from resulting tradespaces. A more analytical method uses the metric *Pareto Trace*<sup>1,14</sup>. Pareto trace is a multi-epoch metric defined as the number of epochs in which a design is located on the Pareto front. Since Pareto Trace operates across epochs, it provides a method to identify designs that are passively value robust. Comparing the characteristics of designs with a high Pareto Trace gives insight into which design variables most impact system value across alternative future contexts<sup>15</sup>. In order to cast a wider net by including not only the designs on the Pareto front, but also those near the Pareto front, we can consider “fuzziness” to the Pareto set. A *K-percent fuzzy Pareto Trace* extends the classic notion of Pareto optimality to include those solutions within K fraction of the Pareto Frontier (i.e., those just within the dominated solution space). This allows designers to examine solutions that perform well across many conditions, even if they are not the dominant solution for any particular condition<sup>15</sup>. In addition, variation in the K level can give insight into the effect of uncertainty on design variables and resulting value robustness. The Pareto Trace for a design is dependent on the number of epochs that are considered. We can divide the Pareto Trace by the number of epochs considered in order to yield a bounded solution in the range 0-1. This *Normalized Pareto Trace* (NPT) and the corresponding “fuzzy” version, *Fuzzy Normalized Pareto Trace*, (fNPT), make it easier to compare designs independent of the number of epochs considered<sup>5</sup>.

### A. Epoch Enumeration

Now that we have defined the satellite constellation tradespace, we can examine the impact that future end-user involvement and market conditions play by including active geographic region(s) and user subscription type (none, trial, partial adoption, full adoption) as epoch variables. The satellite service has multiple potential applications in widespread geographic regions, and it is not fully clear which applications will be widely subscribed, and which geographical locations will be the most active. For this example, thirteen applications were examined in the case study (enumerated here as A through M). Additionally, while it is possible that some users would fully adopt the satellite service with only a single satellite deployed, many applications would require the performance afforded by a partial or full constellation deployment (e.g., some users would likely delay full service implementation until revisit times are reduced and service redundancy is assured).

For this example, the satellite offers two distinct types of service: a 1-way data transmission service (service type 1; ST1), and a 2-way bi-directional communication service (service type 2; ST2). Each application has an associated geographic region and primary service type, although based on discussion with stakeholders, it is assumed that applications may occasionally make use of the other service type. Service can be adopted on a trial basis, or a full use of service basis. Thus each epoch is defined by a usage level for each application varying from no use of service to full use of both service types 1 and 2. Each usage case corresponds to a utility weighting modifier assigned as shown in Table 7. A full list of epochs to be considered (I through X) with corresponding application usage index is

shown in Table 8 (e.g., in epoch **V** application **A** adopts usage index 3, corresponding to ST1 full use, and no use of ST2). The enumerated epochs have been hand-picked to represent a likely development path based on discussion with the application user group, but it is not strictly necessary for the epochs to be assembled in order for multi-epoch analysis (as opposed to the ordered sets which would be generated for full EEA analysis). For example, because we don't consider path dependence, this multi-epoch analysis would still be relevant if some users adopt the service on a trial basis and then return to no use; alternatively, it is possible that some users might skip the trial use segment and adopt the full service from the start.

**Table 7. Usage cases and corresponding utility weighting modifier**

Usage case	Usage index	Utility modifier (if ST1 is the primary service for the application)	Utility modifier (if ST2 is the primary service for the application)
No use of service	0	0	0
ST1 trial	1	0.25	0.05
Full trial	2	0.33	0.33
ST1 full use, no use of ST2	3	0.8	0.33
ST1 full use, ST2 trial	4	0.9	0.6
Full use of service	5	1	1

**Table 8. Application and usage epoch enumeration**

Application index	Primary Service Type	Area of Interest	Epoch & usage index based on Table 7									
			I	II	III	IV	V	VI	VII	VIII	IX	X
<b>A</b>	ST1	Global	3	3	3	3	3	4	4	4	5	5
<b>B</b>	ST1	Global	1	3	3	3	3	4	4	4	5	5
<b>C</b>	ST1	Region 1	0	0	1	1	2	2	4	4	5	5
<b>D</b>	ST1	Region 1	0	0	1	1	2	2	4	4	5	5
<b>E</b>	ST1	Region 1	0	0	1	1	2	2	4	4	5	5
<b>F</b>	ST2	Global	0	0	1	1	2	2	4	4	5	5
<b>G</b>	ST2	Region 1	0	0	1	1	2	2	4	4	5	5
<b>H</b>	ST2	Region 1	0	0	1	1	2	2	4	4	5	5
<b>I</b>	ST2	Region 1	0	0	1	1	2	2	4	4	5	5
<b>J</b>	ST1	Global	0	0	0	1	0	2	2	4	2	5
<b>K</b>	ST2	Region 2	0	0	0	1	0	2	2	4	2	5
<b>L</b>	ST1	Region 2	0	0	0	1	0	2	2	4	2	5
<b>M</b>	ST2	Region 2	0	0	0	1	0	2	2	4	2	5

## B. Comparing Designs across Epochs

Now that we have defined the epochs of interest, we can compute the Fuzzy Normalized Pareto Trace (fNPT) for the designs enumerated in tables 3 to 5. In each epoch, one or more application is 'activated' to be included in the analysis. For each application, a utility is contributed based on the application's region-specific multi-attribute tradespace (e.g., see Fig. 6), and its usage case for the epoch in question (Table 8). The resulting fNPT across all 10 Epochs (including all applications and geographic regions) are shown in Figs. 7 to 9 for K-fuzziness levels of 5%, 10%, and 25%. Recall that the 5% fuzziness level identifies the number of epochs for which a design is within 5% of Pareto optimality both in terms of utility and cost, relative to all other designs considered. Since the fNPT is normalized by the total number of epochs considered (10), the fNPT must be a multiple of 0.1 in the range 0 to 1. Additionally, the 10% and 25% K fuzziness levels increase the zone of inclusion about the Pareto efficiency zone (up to 25% off Pareto optimality in terms of utility and cost), so include progressively more designs in the Pareto trace (i.e., the fNPT for a design can go up with K-value, but cannot go down). In general, the higher the fNPT, particularly at a lower K level, the more passively robust the design for the epochs considered.




Figure 7. Fuzzy Normalized Pareto Trace across 10 Epochs for K = 5%

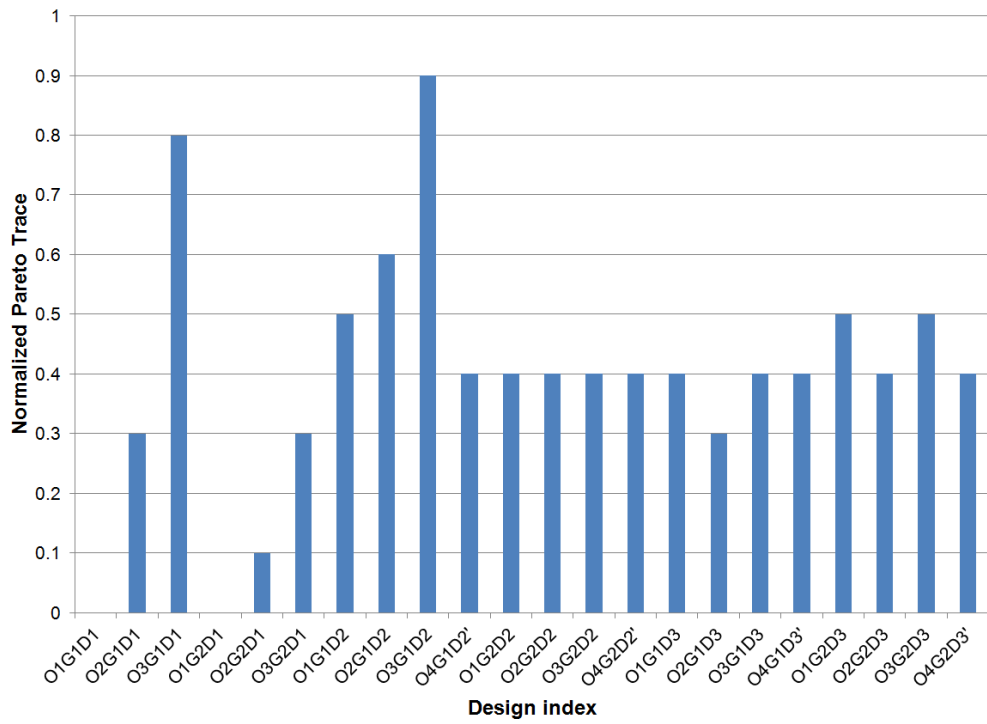


Figure 8. Fuzzy Normalized Pareto Trace across 10 Epochs for K = 10%




Figure 9. Fuzzy Normalized Pareto Trace across 10 Epochs for K = 25%

#### IV. Discussion

Based on the fNPT results (Fig. 7 to 9), we find that designs **O3G1D1**, **O3G1D2**, and **O2G1D2** are the most passively robust across the epochs considered (with design **O1G2D3** being the most passively robust for the full constellation deployment case). Based on the multi-attribute utilities shown in Figs. 4 and 5, some of these designs gain a performance advantage for particular regions, but none of these designs stand out as having a performance advantage across the tradespaces for all geographic regions. Indeed, rather than performing optimality for a single region and set of end user subscription, these designs instead demonstrate their increased robustness when we consider a wide range of conditions and user adoption contexts.

Several trends are evident in the fNPT results. Most of the designs have relatively high fNPT values at the 25% level for many epochs. This is largely due to the fact that this analysis only considered designs that have been determined by the domain experts to be relatively efficient; that is, here we are considering pre-filtered tradespaces<sup>6</sup>. Increasing the K-fuzziness level includes more designs in those considered fuzzy Pareto optimal, however, no designs achieve fuzzy Pareto optimality in all epochs even at the K = 25% level. Additionally, while the single satellite and partial deployment designs have more variance in fNPT, they also tend to have higher fNPT values. This is because a smaller number of satellites can achieve a better performance to cost ratio because the location of the first few satellites is the most optimized, with additional satellites providing increasingly redundant coverage for the same extra cost; however, a small number of satellites are more likely to perform especially poorly by failing to cover a particular region adequately or at all.

Based on a more detailed examination of the individual regional tradespaces, we can identify the inherent advantages of designs **O3G1D1**, **O3G1D2**. Since the **O3** designs represent constellations comprised of polar satellites, they are more versatile in providing good service to a wide range of geographic locations without being optimized for any particular location. Although region 1 is located at a relatively low latitude (15-35°) so it is serviced well by low-inclination and even somewhat by equatorial orbits, the higher latitudes of region 2 (35-60°) require polar satellites to achieve comprehensive coverage with a small number of satellites. Additionally, applications that are distributed globally benefit from the coverage provided by polar satellites, with less advantage from equatorial satellites. This relative advantage declines for the partially and especially fully deployed constellation cases, where polar satellites are added in all designs. In fact, in the fully deployed constellation case,

the design containing the near-equatorial satellite (**O1G2D3**) is the most passively robust, because it provides global coverage based on a polar constellation with an equatorial coverage supplement, improving its performance for applications in region 1 and distributed globally.

The upgraded ground station configurations (**G2** designs) are generally not worth the additional cost, except in the fully deployed constellation cases, where they yield a general increase in fNPT compared with **G1** designs. This is because until a full constellation is deployed, the ground stations are significantly under-utilized and fail to contribute their full potential value. Additionally, in the case of a full constellation deployment, the ground station infrastructure contributes a much lower fraction of the overall system cost. However, note that although this analysis considered the construction of new ground station infrastructure at a fixed cost, the relative value afforded by more ground stations may be worth the cost in the case of additional time purchased on existing ground stations.

In general, the ‘flexible deployment’ path (**O4**) designs failed to provide as much passive robustness as the planned deployment path designs. However, this analysis considered only the relative performance impacts and cost saving of these paths. There may be other unconsidered factors such as capital investment cash flow timelines, or additional uncertainties in user adoption or the competitive environment which could tilt the balance in favor of these designs. Further analysis could consider the advantages yielded by designing the system for changeability within the context of robustness, for example by invoking a Valuation Approach for Strategic Changeability analysis (VASC)<sup>13</sup>.

Based on the fNPT, the most efficient deployment path to achieve robust value across a wide range of geographic and user subscription contexts starts with designs **O3G1D1**, **O3G1D2**. However, once a partial polar constellation has been deployed, near-equatorial satellites offer the largest increase in passive robustness. Thus, we can imagine a most efficient deployment path that is somewhat reversed from the originally envisioned order, whereby the last launch (instead of the first) would consist of one or more near-equatorial satellites. Additionally, once the partial constellation has been deployed, we might then consider upgrading the ground station infrastructure. For example, instead of the planned polar constellation deployment path of **O3G1D1** → **O3G1D2** → **O3G1D3**, we might consider a deployment path such as **O3G1D1** → **O3G1D2** → **O1G2D3**. Since additional ground stations could be added at any time, and satellites could be deployed in any order, there should be no physical barrier to this alternate deployment path. Such a strategy could be considered in the context of a full time-order dependent EEA that examines constellation buildup, deployment timeline, and changeability in the design for eras that account for a realistic sequence of exogenous variables.

## V. Conclusions

Tradespace exploration provides a good starting point to examine the impact of design choices on the ability of competing designs to deliver value to stakeholders at an efficient performance vs. cost tradeoff. The first step is to define the attributes that provide utility to stakeholders and enumerate the design variables that can be traded at the subsystem and system level. Once that is accomplished, a tradespace can be assembled for each individual attribute that is relevant to stakeholders, either directly based on user requirements, or based on another form of preference elucidation. The tradespace can either be a full set of all potential designs, or we can leverage existing domain expert knowledge to generate pre-filtered tradespaces contain only designs of relatively high value<sup>6</sup> (variations from these designs can be included in the analysis at a future stage if desired). By assigning utility to attributes, and weighting them based on stakeholder preference, we can examine the performance of a design for a fixed set of conditions, taking into account nonlinear perceptions of benefit and representing them on a common basis.

For complex satellite systems with uncertain futures, we can employ multi-epoch analysis or Epoch-Era Analysis in an early design stage to provide insight on how competing designs will be poised to yield passive value robustness across a wide range of alternative future contexts. Fuzzy Normalized Pareto Trace (fNPT) is a good metric to evaluate the passive value robustness of designs across many epochs representing alternative future conditions. Using this metric, we can identify designs that provide high value robustness, and gain insight into other factors, such as which deployment paths will likely achieve a large degree of robustness over the system lifecycle. By modifying the K-fuzziness parameter, we can adjust the analysis to be more or less inclusive, providing insight on how close each design is to the Pareto optimal condition across the projected future epochs. The insight yielded by this analysis can be used to drive a re-examination of user requirements, stakeholder utility valuation, and

assumptions about future contexts. Additionally, it identifies designs that warrant further investigation, focusing the designer's attention on the design variables that most impact performance across a wide range of future contexts. The results from the illustrated case study suggested an initially non-intuitive deployment strategy that was the reverse of original plans. Upon further consideration, the deployment strategy suggested by the results did make sense over the traditional strategy, illustrating the promise of using a structured framework such as Epoch-Era Analysis for proposing temporal strategies to achieve value robustness.

## Acknowledgments

This work was carried out by the lead author while in the Mission Development Group at COM DEV International, Ltd.

## References

- <sup>1</sup> Fitzgerald, M.E. and Ross, A.M., "Sustaining Lifecycle Value: Valuable Changeability Analysis with Era Simulation," *SysCon2012 – IEEE International Systems Conference 2012*. Vancouver, Canada, March 2012.
- <sup>2</sup> Fricke, E. and Schulz, A., "Design for Changeability (DfC): Principles to Enable Changes in Systems Throughout their Entire Lifecycle," *Systems Engineering* **8**(4): 432-359, 2005.
- <sup>3</sup> Saleh, J. Mark, G., and Jordan, N., "Flexibility: a multi-disciplinary literature review and a research agenda for designing flexible engineering systems," *Journal of Engineering Design* **20**: 307-323, 2009.
- <sup>4</sup> Ross, A.M. and Hastings, D.E., "Assessing Changeability in Aerospace Systems Architecting and Design Using Multi-Attribute Tradespace Exploration," *AIAA Space 2006* San Jose, CA, September 2006.
- <sup>5</sup> Roberts, C.J., Richards, M.G., Ross, A.M., and Rhodes, D.H., "Scenario Planning in Dynamic Multi-Attribute Tradespace Exploration," in *SysCon2009 – IEEE International Systems Conference*, Vancouver, BC, Canada 2009.
- <sup>6</sup> Rader, A.A., Newland, F.T., and Ross, A.M., "An Iterative Subsystem-Generated Approach to Populating a Satellite Constellation Tradespace," *AIAA Space 2011*, Long Beach, CA, September 2011.
- <sup>7</sup> Ross, A.M. and Hastings, D.E., "The Tradespace Exploration Paradigm," *INCOSE International Symposium 2005*, Rochester, NY, July 2005.
- <sup>8</sup> Ross, A. M., Diller, N. P., Hastings, D. E., and Warmkessel, J. M. "Multi-Attribute Tradespace Exploration with Concurrent Design as a Front-End for Effective Space System Design," *Journal of Spacecraft and Rockets*, Vol. 41, No. 1, 2004, pp. 20-28.
- <sup>9</sup> Keeney, R. L., and Raiffa, H., *Decisions with Multiple Objectives--Preferences and Value Tradeoffs*. 2nd ed. Cambridge: Cambridge University Press, 1993.
- <sup>10</sup> Ross, A.M., McManus, H.L., Rhodes, D.H., and Hastings, D.E., "A Role for Interactive Tradespace Exploration in Multi-Stakeholder Negotiations," *AIAA Space 2010*, Anaheim, CA, September 2010.
- <sup>11</sup> Ross, A.M. and Rhodes, D.H., "Using Natural Value-Centric Time Scales for Conceptualizing System Timelines through Epoch-Era Analysis" *INCOSE International Symposium 2008*, Utrecht, the Netherlands, June 2008.
- <sup>12</sup> Rader, A.A., Ross, A.M., and Rhodes, D.H., "A Methodological Comparison of Monte Carlo Methods and Epoch-Era Analysis for System Assessment in Uncertain Environments," 4th Annual IEEE Systems Conference, San Diego, CA, April 2010.
- <sup>13</sup> Fitzgerald, M.E. and Ross, A.M., "Mitigating Contextual Uncertainties with Valuable Changeability Analysis in the Multi-Epoch Domain," *SysCon2012 – IEEE International Systems Conference 2012*. Vancouver, Canada, March 2012.
- <sup>14</sup> Ross, A.M., Rhodes, D.H., and Hastings, D.E., "Using Pareto Trade to Determine System Passive Value Robustness," *SysCon2009 – IEEE International Systems Conference 2009*. Vancouver, Canada, March 2009.
- <sup>15</sup> Ross, A.M., McManus, H.L., Rhodes, D.H., and Hastings, D.E., "Revisiting the Tradespace Exploration Paradigm: Structuring the Exploration Process," *AIAA Space 2010*, Anaheim, CA, September 2010.