

MIT Open Access Articles

Responsive systems comparison method: Dynamic insights into designing a satellite radar system

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Ross, Adam M. et al. "Responsive systems comparison method: Dynamic insights into designing a satellite radar system." in Proceedings of the AIAA SPACE 2009 Conference & Exposition, Pasadena, CA, 14-17 September, 2009, AIAA.

As Published: <https://www.aiaa.org/ProceedingsDetail.aspx?id=3895>

Publisher: American Institute of Aeronautics and Astronautics

Persistent URL: <http://hdl.handle.net/1721.1/72318>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of use: Creative Commons Attribution-Noncommercial-Share Alike 3.0



Responsive Systems Comparison Method: Dynamic Insights into Designing a Satellite Radar System

Adam M. Ross*

Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

Hugh L. McManus†

Metis Design, Cambridge, MA, 02141, USA

Donna H. Rhodes‡ and Daniel E. Hastings§

Massachusetts Institute of Technology, Cambridge, MA, 02139, USA

and

Andrew Long**

Chantilly, VA, 20152, USA

Often shifts in context, such as changes in budgets, administrations, and warfighter needs, occur more frequently than high-cost space-based system development timelines. In order to ensure the successful development and operation of such systems, designers must balance between anticipating future needs and meeting current constraints and expectations. This paper describes the application of Multi-Epoch Analysis on a previously introduced satellite radar system program case study, quantitatively analyzing the impact of changing contexts and preferences on “best” system designs for the program. Each epoch characterizes a fixed set of context parameters, such as available technology, infrastructure, environment, and mission priorities. For each epoch, several thousand design alternatives are parametrically assessed in terms of their ability to meet imaging, tracking, and programmatic expectations using Multi-Attribute Tradespace Exploration. While insights on tradeoffs are discovered within a particular epoch, further dynamic insights become apparent when comparing tradespaces across multiple epochs. The Multi-Epoch Analysis reveals three key insights: 1) the ability to quantitatively investigate the impact of “requirements” across many systems and contexts, 2) the ability to quantitatively identify value “robust” systems, including both passively robust and changeable systems, and 3) the ability to quantitatively identify key system tradeoffs and compromises across stakeholders and missions.

I. Introduction

Often shifts in context, such as changes in budgets, administrations, and warfighter needs, occur more frequently than high-cost space-based system development timelines.¹ In order to ensure the successful development and operation of such systems, designers and system architects must balance between anticipating future needs and meeting current constraints and expectations. Scenario analysis and system modeling techniques have been developed in order to quantify and gain insight into the design ramifications of such changes,² however, current quantitative methods in decision analysis and forecasting tend to focus on probabilistic formulations of the future.³

* Research Scientist, Engineering Systems Division, E38-574, and AIAA Senior Member.

† Senior Special Projects Engineer, 10 Canal Park, AIAA Associate Fellow.

‡ Principal Research Scientist, Engineering Systems Division, E38-572, and AIAA Member.

§ Professor of Aeronautics and Astronautics and Engineering Systems, Dean for Undergraduate Education, 4-110, AIAA Fellow.

** Affiliate, 43727 Scarlet Sq, AIAA Member.

Probabilistic treatments of the future are hampered by a real difficulty in estimating probabilities, as well as the inevitable impact of “unknown unknown” futures.⁴ Instead of pursuing a probabilistic characterization of the future, an approach is used that enumerates possible futures through parametric variation of sources of uncertainty representing snapshots of potential contexts and needs facing a system. This paper describes the application of Multi-Epoch Analysis on a previously introduced satellite radar system program case study⁵, quantitatively analyzing the impact of changing contexts and preferences on “best” system designs for the program. Each epoch characterizes a fixed set of context parameters, such as available technology, infrastructure, environment, and mission priorities. During a fixed context, several thousand design alternatives are parametrically assessed in terms of their ability to meet imaging, tracking, and programmatic expectations using Multi-Attribute Tradespace Exploration.^{6,7}

II. Method Used

The case study described in this paper was used to help develop and test the development of a multi-epoch tradespace exploration methodology. The following section describes this methodology, along with precursor methods and case studies that provided the foundation for its development.

A. Responsive Systems Comparison method overview

The Responsive Systems Comparison (RSC) methodology utilizes Dynamic Multi-Attribute Tradespace Exploration (MATE)⁷ coupled with Epoch-Era Analysis¹ to take a designer or system analyst (RSC practitioner) through a step-by-step process of designing and evaluating dynamically relevant system concepts. These processes are designed to guide the RSC practitioner through the steps of determining how a system will deliver value, brainstorming solution concepts, identifying variances in contexts and needs (Epochs) that may alter the perceived value delivered by the system concepts, evaluating key system trade-offs across varying epochs, assessing the changeability of the various system concepts, developing possible sequences of epochs (Eras) to be encountered by the system, and lastly developing strategies for how a designer might develop and transition a particular system concept through and in response to these varying epochs. At the end of the RSC method the RSC practitioner will have a characterization of the desired value-space, identification of system concepts comprising the design-space, the means to identify and assess dynamically relevant systems concepts (via RSC-specific metrics), insights into key design trades, and an overall development and deployment strategy with which to develop their system in response to varying contexts and needs.

The RSC method consists of seven (7) processes – as shown in Figure 1 below. These processes, with a brief description of purpose, are as follows:

1. **Value-Driving Context Definition** - Identify overall problem / needs statement.
2. **Value-Driven Design Formulation** - Elicit stakeholder needs statements (attributes) and formulate system solution concepts (design variables).
3. **Epoch Characterization** - Parameterize the range of contextual uncertainties (epochs) under consideration.
4. **Design Tradespace Evaluation** - Gain an understanding, via modeling and simulation, of how key system concepts and trades (design variables) fulfill the overall value-space (attributes) in response to contextual uncertainties (epochs).
5. **Multi-Epoch Analysis** - Identify value robust system designs across changing contexts and needs.
6. **Era Construction** - Develop era timelines from the set of enumerated epochs.
7. **Lifecycle Path Analysis** – Develop near- and long-term system value delivery strategies in response to time-dependent contextual uncertainties (described via era timelines).

Overall, the seven RSC processes occur in sequence, but some feedback can occur between process 2, 3, and 4. The primary purpose of the feedback is to refine the system concepts, various contexts, value propositions, and assumptions under evaluation. This refinement can occur for several reasons including desiring higher model fidelity, re-scoping due to new information concerning the problem scope or problem statement, or availability of new system concepts and/or models as insights are gained from previous iterations.

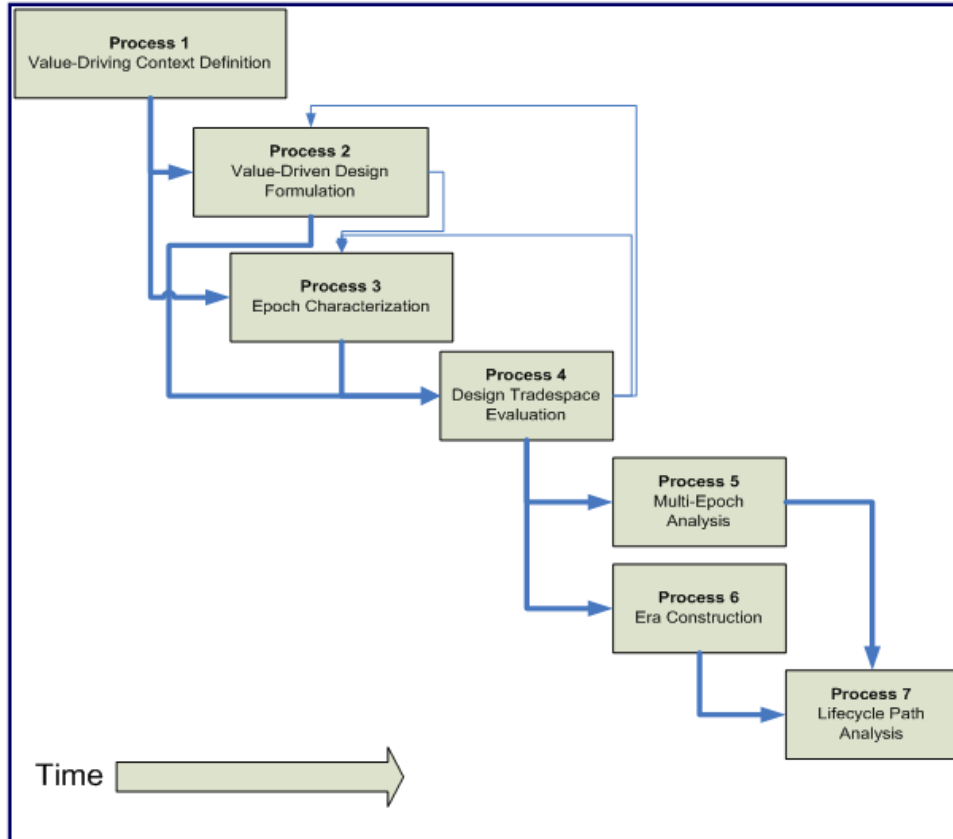


Figure 1. Responsive Systems Comparison (RSC) method process flowchart.

Value-Driving Context Definition (Process 1) and part of Value-Driven Design Formulation (Process 2), as applied to this case study, are described in Ref. 5, illustrating the early implementation of RSC. What follows is a brief summary of the early processes (1-3) with some introduction to Process 4, and a detailed focus on Multi-Epoch Analysis (Process 5). Era Construction (Process 6) is discussed in Ref. 8, and Lifecycle Path Analysis (Process 7) is still being applied to this case study, to be published at a future date.

III. SRS Case Summary

The following sections describe the application of the various RSC processes to the Satellite Radar System (SRS) case study.

A. Value-Driving Context Definition

The first process of RSC is intended to capture the overall problem statement (what is the problem, why is it important, and who cares about the problem and solution). Additionally, the key contextual factors that may affect the problem and/or its solution should be identified, including past, present, and possible future instances of these factors. For this study, a satellite radar system (SRS) was proposed as a solution to the problem of providing 24-hour, all-weather imaging and tracking of targets of interest. The purpose of the case study was defined as assessing the ability of potential satellite radar system architectures to satisfy potential stakeholders over a large range of possible future situations (epochs). The overall value proposition for the case study was “to determine which SRS architecture a notional SRS Program Manager should select to maximize the chances that stakeholders will remain satisfied throughout the system lifecycle” (i.e., which system will provide the highest degree of *value robustness*⁹).

B. Value-Driven Design Formulation

The second process is an essential part of RSC where key decision makers are interviewed to determine their needs statements as expressed by objectives. A set of attributes, metrics that reflect how well a decision maker-defined objective is met, are elicited from each decision maker for each mission of interest for the system.^{10,11} These

attributes are the criteria used for distinguishing the “goodness” of each design alternative considered. The attributes elicited through interviews with proxy users in the community and derived from public data regarding satellite radar¹² are listed in Table 1.

Table 1. Attribute definitions for three missions: tracking, imaging, and “program.”

	Attribute Name	units	range (U=0 to U=1)
Tracking	Minimum Target RCS	m ²	1000 --> 0
	Min. Discernable Velocity	m/s	50 --> 5
	Number of Target Boxes	#	1 --> 10
	Target Acquisition Time	min	120 --> 10
	Target Track Life	min	0 --> 60
	Tracking Latency	min	240 --> 0
Imaging	Resolution	m	5 --> 0.01
	Targets per Pass	#	1 --> 10 ⁵
	Field of Regard	km ²	1000 --> 10 ⁶
	Revisit interval	min	300 --> 10
	Imaging Latency	min	720 --> 60
	Geo-Location Accuracy	m	500 --> 50
Program	Baseline Cost	\$B	
	Deviation from Cost	%	
	Baseline Schedule	years	
	Deviation from Schedule	years	

After the attributes are defined, the next step is for the designer to develop a set of concepts for meeting those attributes. For this case study, the concept of a satellite radar constellation was assumed and the design variables represent aspects of this concept that can be traded (i.e., are within the control of the designer). The final list of design variables and their defined ranges are listed in Table 2. Fully enumerating across each of the defined ranges would result in 124,416 different design configurations. The grayed out values represent enumeration levels that were ultimately not assessed in the case study, reducing the tradespace size to 23,328 design configurations.

Table 2. Satellite radar design variables and ranges.

	Variable Name	Definition Range
DESIGN VARIABLES	Peak Transmit Power	1.5, 10, 20 [kW]
	Radar Bandwidth	0.5, 1, 2 [GHz]
	Physical Antenna Area	10, 40, 100, 200 [m ²]
	Antenna Type	Mechanical, AESA
	Altitude	800, 1200, 1500 [km]
	Constellation Configuration	8 Walker IDs
	Communication Downlink	Relay, Direct downlink
	Tactical Communication Able	Yes, No
	Maneuver Capability	1x, 2x, 4x base fuel
	Tugable	Yes, No
	Constellation Option	None, long-lead, spares

C. Epoch Characterization

After defining the value proposition and possible design responses in order to develop a potentially successful system, this process seeks to characterize the contextual uncertainties that may prevent system success in spite of the designer’s intentions. The process, described in Ref 5 and Ref 8, results in a set of context and need uncertainty categories (Figure 2) from which a set of parameters can be derived that capture possible combinations of unfolding uncertainties that may arise (Table 3). The parameters used in this way are called *epoch variables* and the span of these enumerated epoch variables is the epoch space, representing the possible future contexts and needs for the system.

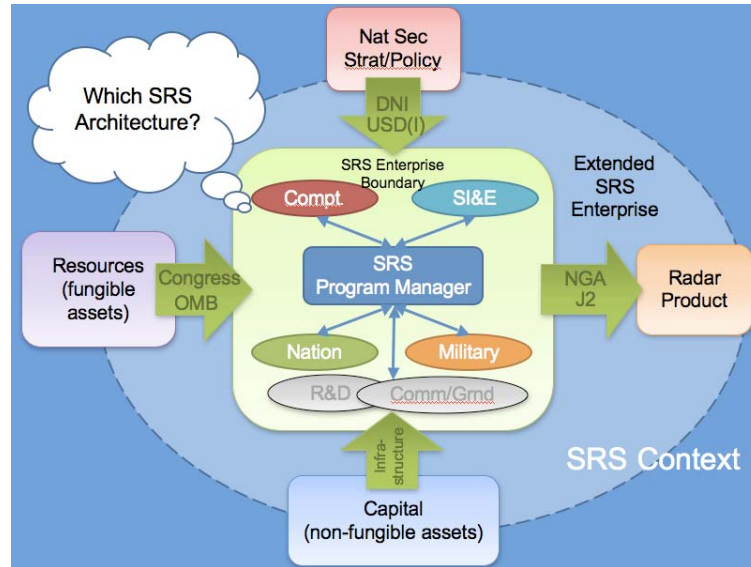


Figure 2. Satellite Radar System scope definition with exogenous uncertainty categories.⁵

Table 3 Epoch variable set definition⁸

Exogenous Uncertainty Category	Epoch Variables	Number of Steps	Units/Notes
Strategy/Policy	Imaging vs. Tracking Mission Priorities	3	1=SAR<GMTI 2=SAR=GMTI 3=SAR>GMTI
Resources	Budget Constraint	na	Use tradespace to vary "costs"
Capital	Radar Technology	3	1=Mature 2=Medium, 3=Advanced
	Communication Infrastructure	2	1=AFSCN 2=WGS+AFSCN
	Collaborative AISR Assets	2	1=Available 2=Not available
Radar Product	Operations Plans	9	Lookup table of geographic region & target op. plans
	Threat Environment	2	1=No jamming 2=Hostile jamming

Enumerating across the range of epoch variables results in 648 unique epochs (contexts and needs) for the Satellite Radar System. These unique epochs will form the basis for evaluating the potential system tradespace across changing futures in order to identify passively value robust designs (those designs that continue to deliver value, quantified as high utility and low cost), or designs that are changeable and able to actively pursue sustained high value.

IV. Results

The primary data for analysis from an application of RSC is generated during Processes 4 and 5. What follows are the results from the Design Tradespace Evaluation (Process 4) and Multi-Epoch Analysis (Process 5) processes.

A. Design Tradespace Evaluation

For each of the enumerated epochs, the tradespace (the span of enumerated design variables) is evaluated using models and simulations, outputting attributes, utilities, and costs of each design for each mission (Figure 3). A high level representation of the tradespace data is usually a Utility-Cost scatterplot with each point representing a feasible design solution for the epoch (i.e. the design is technically feasible and falls within the acceptable bounds for each attribute for each decision maker).

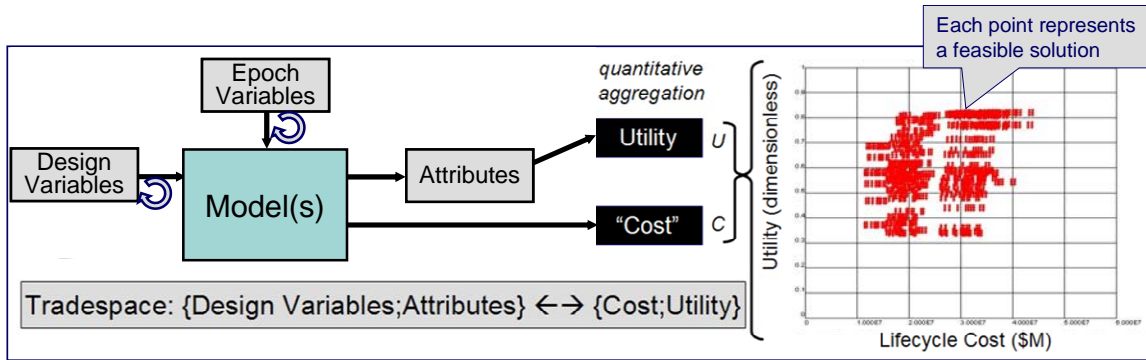


Figure 3. Data flow for generating tradespaces for each epoch using design variables and epoch variables as inputs.

Figure 4 shows a typical graphic representation of a static tradespace, here generated by ATSV software.¹³ The vertical axis shows total utility in a given epoch, the horizontal axis shows total lifecycle cost with each point representing an evaluated and feasible design. In interpreting the figure, the designs at the top (higher performance) and to the left (less cost) are the preferred (Pareto set) designs. A simple use of such a plot would be to locate promising designs (such as the ones on the “knee”) for further evaluation.

Additional information exists in the data beyond lifecycle cost and utility. As a first step, color is used to plot a third dimension of data, in this case antenna area. Larger antennas (100 m², red points) are higher performance but more expensive; in this epoch, the medium sized antenna (40 m², blue points) can provide good performance at lower cost. Designs with small antennas (10 m²) do not appear; further investigation of the data set reveals that they do not meet the minimum requirements for spotting the targets in this epoch and were therefore infeasible in this tradespace. Each of the 648 epochs can be evaluated in this way, coloring by various design variables to glean patterns, correlations, and insights across the tradespace; however the process can be very time-consuming and possibly intractable due to the sheer size of the data set.

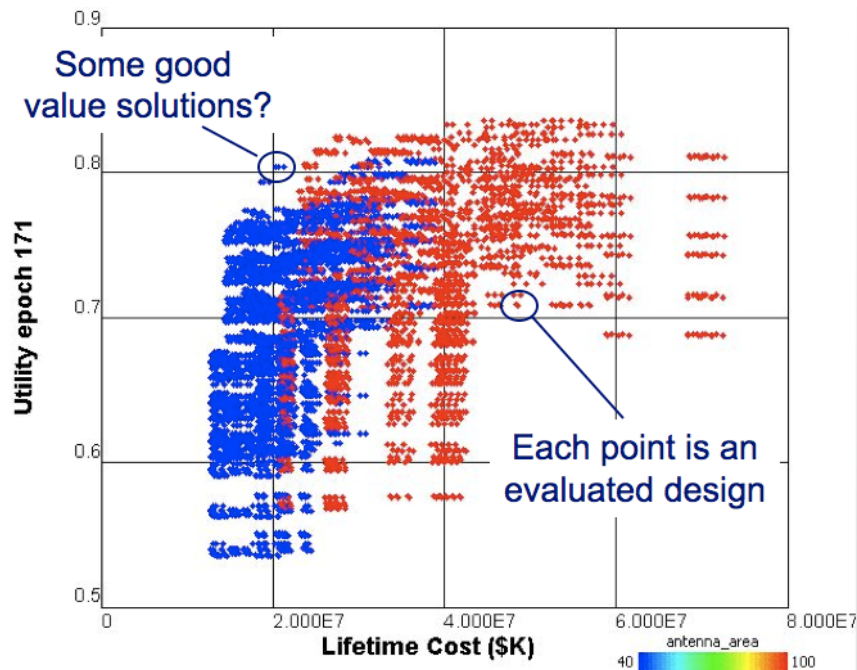


Figure 4. Static tradespace example.

In order to more quickly gain insights into the tradespace across epochs, a number of analysis techniques are applied within Multi-Epoch Analysis (Process 5), which follows.

B. Multi-Epoch Analysis

1. Cross-epoch: Tradespace Yield

As discussed in Section II, there are 23,328 total possible designs in the Satellite Radar System tradespace. However, for one or more of several reasons, a design may be deemed ‘infeasible.’ A design may become infeasible if it is physically impossible, it cannot fit on a launch vehicle, or it fails to achieve the minimum acceptable level of one or more of the attributes identified by the stakeholder(s). When that happens, an infeasible design is not plotted or analyzed in later steps of the process. This means that a tradespace may display considerably fewer than the 23,328 possible designs. The yield percentage for each epoch’s tradespace was calculated, and is shown below in Figure 5. This only shows data for those epochs that were evaluated (245 epochs), so blank spaces indicate an epoch that was not evaluated.^{††}

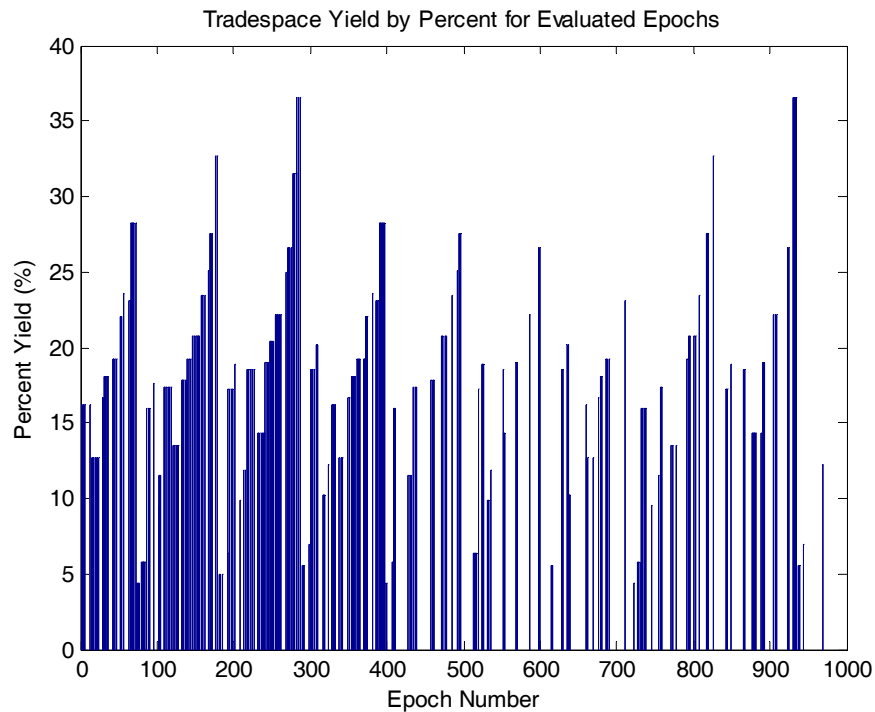


Figure 5. Tradespace yield by epoch number.

The tradespace yields were also plotted on a histogram, as shown in Figure 6. The yields are distributed in an approximate bell curve. The low yield epochs (bins to the left) are epochs that tend to have very demanding target sets. The high yield epochs (bins to the right) are those epochs that incorporate airborne intelligence, surveillance, and reconnaissance (AISR) assets. Most epochs have a yield of 18-20%.

When selecting epochs for further analysis, it is important to understand the yield statistics. How typical is a selected “baseline” epoch? Is it more demanding than most epochs (which might be useful for setting requirements)? Is the epoch an outlier (which might result in unrealistic and/or very expensive requirements not needed in most user environments)?

^{††} Enumerating the full range of the original epoch variables resulted in a total of 972 distinct epochs. Further model development revealed a flaw in the model, resulting in a reduction of distinct epochs to 648. The original numbering scheme was retained, however, so the epoch identifiers range up to 972. Of the 648 possible epochs, 245 were evaluated at the time of writing. On-going research is evaluating the remaining epochs and will verify stability of insights.

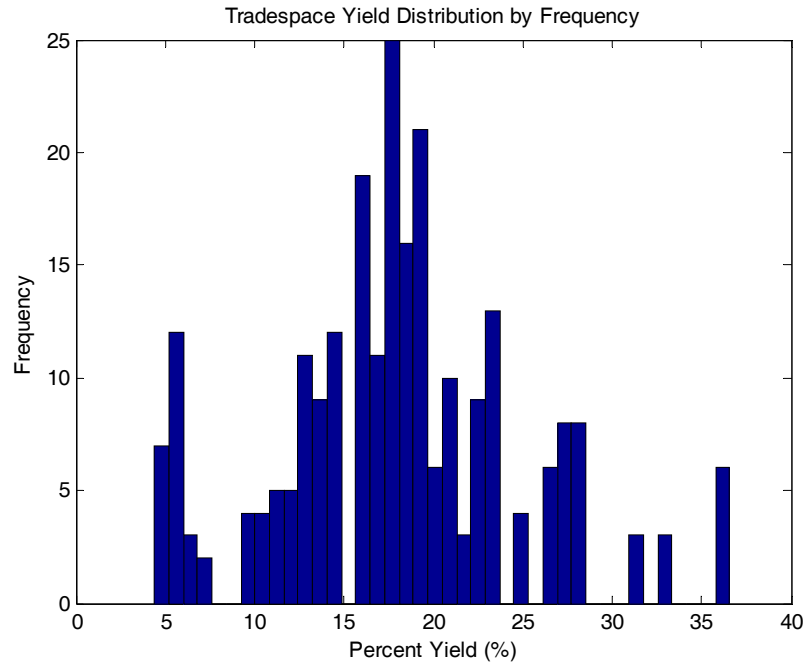


Figure 6. Tradespace yield frequencies across 245 epochs.

2. *Within-epoch: Distribution of attribute performance vs. limits*

The role that the individual attributes played in the aggregated tradespace was explored by a different graphical approach. Figure 7 shows plots for the six attributes of interest to the imaging user. Each plot has the attribute value on the x axis and the frequency at which a given range of values occurs in the designs on the tradespace on the y axis. The plots are for Epoch 171, with ALL designs (even those that were disqualified for violating minimum performance criteria) shown. Also noted on the plots are attribute values corresponding to some value(s) of utility.^{‡‡}

These plots present the effect of the attribute acceptability ranges in the overall tradespace results, and hence offer insight into how demanding expectations are for the stakeholders. For example, the imaging latency has a large number of designs (a group that uses a communication backbone) with utilities just above 0.25. The arrow indicates the “good” direction for this attribute. Other designs have a range of latencies depending on constellation and communication systems, but these designs all have low utility. A few designs are disqualified for having unacceptable latency. Imaging latency is a lightly weighted attribute ($k=1$) as well, so the overall effect on the utility of the designs in the tradespace will be modest (except for the designs rendered infeasible).

Number of targets per pass, on the other hand, has values distributed throughout the utility range (except for very low utility systems, which turn out to be excluded from the tradespace for failing other criteria). Number of targets per pass is a heavily weighted attribute, so it will have a large effect on the utility of the designs in the overall tradespace. As noted previously, this attribute is driven by peak power and antenna area, reinforcing the well-known importance of the power-aperture product for a radar system.

The other heavily weighted attribute, theoretical maximum resolution, does not affect the tradespace as strongly due to the CONOPs assumption in the model: that a system will dwell on a target long enough to achieve its theoretical maximum resolution based on the bandwidth and dimensions of the system. All systems can get reasonable performance on this attribute (utility between 0.6 and 0.8). The distribution is banded because, given unlimited viewing times, the three levels of bandwidth dominates the attribute.

^{‡‡} Sets of attributes are defined per decision maker per mission. For each attribute, an acceptability range is defined corresponding to utility=0 (minimally acceptable) to utility=1 (maximally desirable). Designs are assessed in terms of their performance in these attributes, which may or may not achieve acceptable levels according to these ranges.

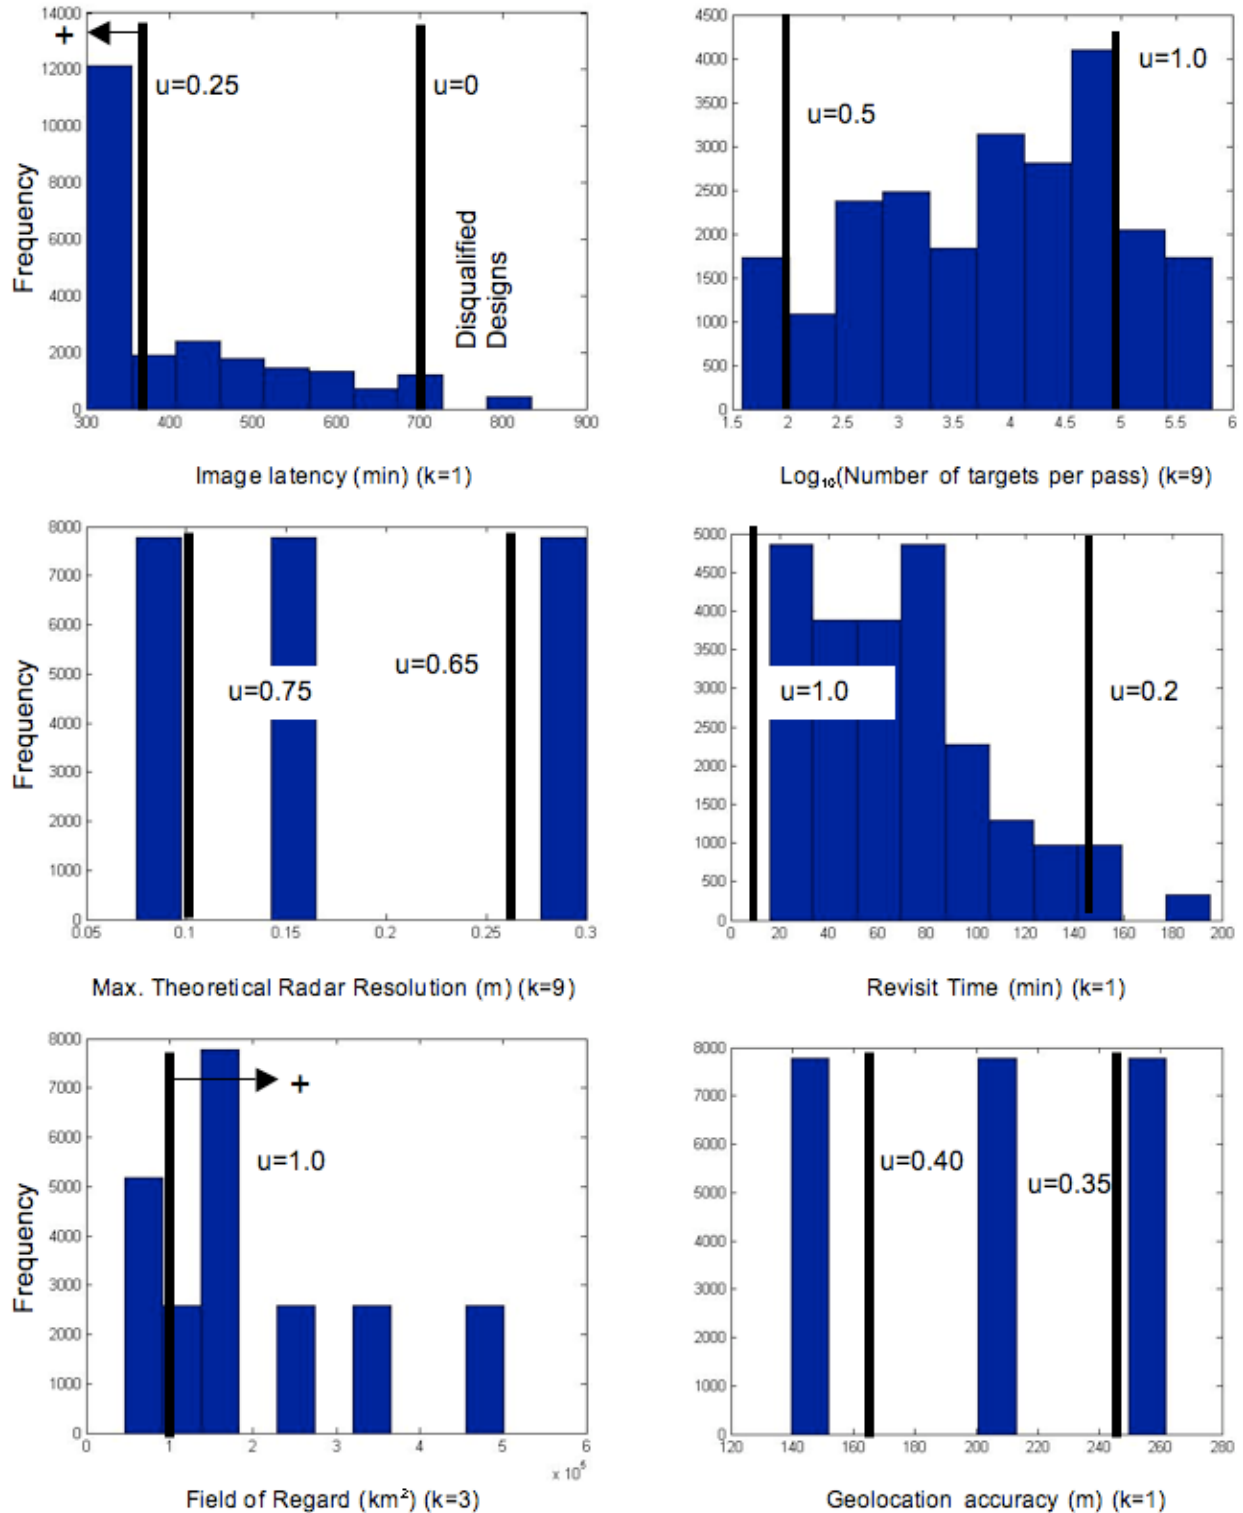


Figure 7. Distributions of image mission attributes (Epoch 171).

Revisit time has values throughout the utility range, but is only modestly effective due to its low weighting. Almost all systems score a “1” (maximum utility) on field of regard, making it almost non-discriminating (despite its moderate weighting); and all systems have lightly weighted geolocation utility in a very narrow range, making it unimportant to the decision maker as a distinguisher among design alternatives.

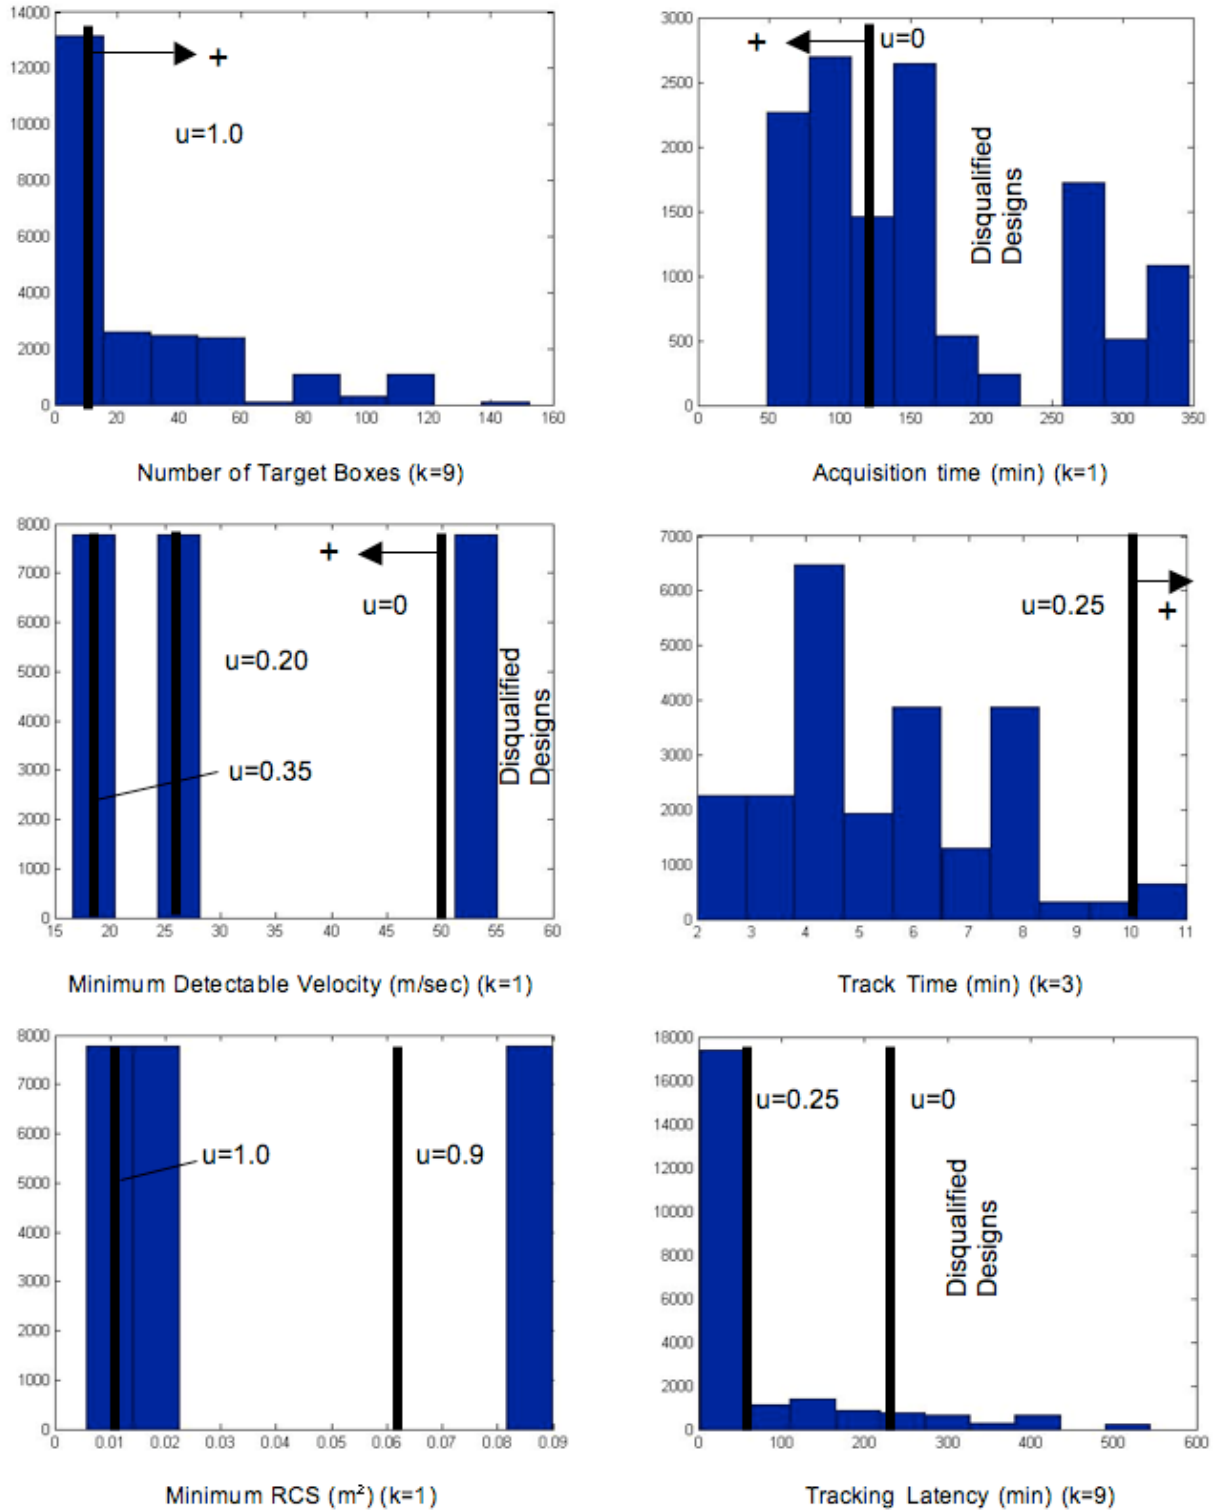


Figure 8. Distribution of tracking mission attributes (Epoch 171).

A similar study was done for the tracking user, with rather different results, shown in Figure 8. The number of target boxes attribute (analogous to the number of targets per pass attribute for the imaging user) was maxed (utility = 1) for many designs, making it only lightly discriminating despite a high weighting. The acquisition time attribute

was only lightly discriminating *except* that a fair number of systems were disqualified for (often just barely) failing the minimum requirement.

In both of these cases, in a “real” SRS study, the decision maker might want to revisit the utility functions. For example, it is possible that number of target boxes, an attribute that the user said was very important, is not discriminating between systems with widely varying performance because the user did not imagine that such high levels of performance were possible, and hence assigned them all a “best possible” utility of 1. In the acquisition time case, whole classes of solutions are eliminated from the solution set due to their inability to meet the stated minimum requirements for a relatively unimportant attribute. Of course, these results may be correct – the extra target boxes may in fact have no utility above a certain threshold, and the acquisition time may indeed be unacceptable at a sharply defined maximum; the point is the user needs to consider these possibilities and decide.

Minimum detectable velocity, despite a low weighting, was quite important as systems fell in three bands, one with a moderate utility, one with a low one, and one which was disqualified. Track life performance was mediocre (low utility) for all systems, making it not much of a discriminator despite an intermediate weighting; minimum RCS (analogous to resolution) was good for all systems. The most important discriminator was the tracking latency, which had a distribution very similar to that of the image latency in Figure 7, but was highly weighted, thus strongly favoring systems with communications backbone and tactical communications capability.

3. Calculating Pareto Trace

Understanding the impact of context changes allows one to quantify system properties known as ‘ilities’. Examples of such ‘ilities’ include robustness, flexibility, adaptability and survivability.

A system is said to be robust when it is relatively insensitive to changes in the system context. Value robustness, the ability to be robust in “value delivery” to stakeholders, can be of two types – active and passive. Actively value robust designs are those that can be changed by a system internal or external change agent below acceptable cost and time thresholds. Passively value robust designs are system designs that maintain high value delivery in a changing system context. System context changes may include changes in the system environment, or changes in stakeholder preferences.

One approach for quantifying the “best” value solutions in a tradespace is to find the non-dominated solutions in the multi-objective space: utility and cost. Non-dominated solutions are those designs at highest utility for a given cost or lowest cost at a given utility. In order to find a solution that does better in one objective, one must do worse in the other. Figure 9 below illustrates the Pareto Set for Epoch 171 tradespace for the Image stakeholder. These “best” designs are context dependent, however, since the tradespace will shift from epoch to epoch as designs move in and out of feasibility and the costs and benefits of alternatives are reassessed. In order to generalize the Pareto Set metric, a new metric, called *Pareto Trace*, was developed to capture a sense of this non-dominated, “best” value across epochs.¹⁴

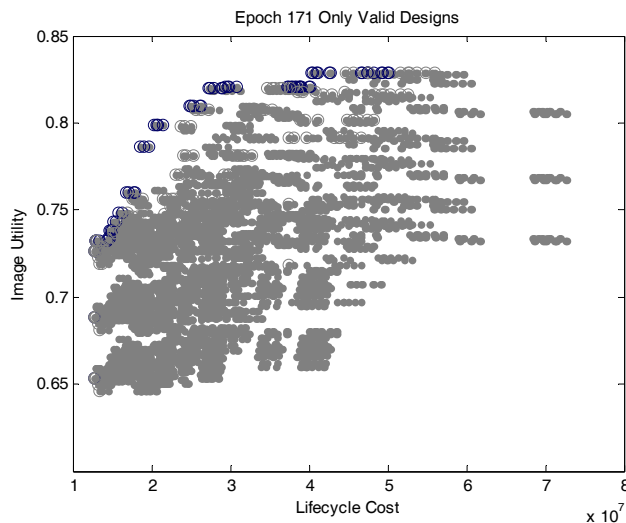


Figure 9. Pareto Set for a single epoch (171).

Pareto Trace can be used as a metric for quantification of value robustness of systems. The Pareto Trace of a system over a number of different changing contexts, i.e. a number of epochs, can inform the decision maker about

the value robustness of that system in comparison to the others in the design space. To determine the Pareto Trace number, the Pareto efficient set of designs for each epoch in the test set is calculated, and the relative frequency of occurrence of the designs in the superset constructed of all the Pareto sets is determined. A high relative frequency, i.e. a high Pareto Trace, indicates that a design is value robust over many changes in the system context. Calculating the Pareto Trace over a large number of epochs will generate a list of passive value robust designs.

In this case study, a sample of 245 epochs was used to demonstrate the Pareto Trace metric and generate a list of value robust designs. In Figure 10, the Pareto Trace for the designs in the design space is shown. The subplots indicate the different objective functions under which the Pareto sets for each trace calculation was done. From this plot, the value robust designs can be identified easily by noting the designs numbers which have a high Pareto Trace. For example, these plots indicate that Design 3435 in the design space has a high Pareto trace under multiple objectives, and may thus be a highly passively robust design, and a possible candidate for detailed design studies. Once highly passively value robust designs have been identified, focused tradespace exploration can be done around the chosen designs to provide more detailed system tradeoff information.

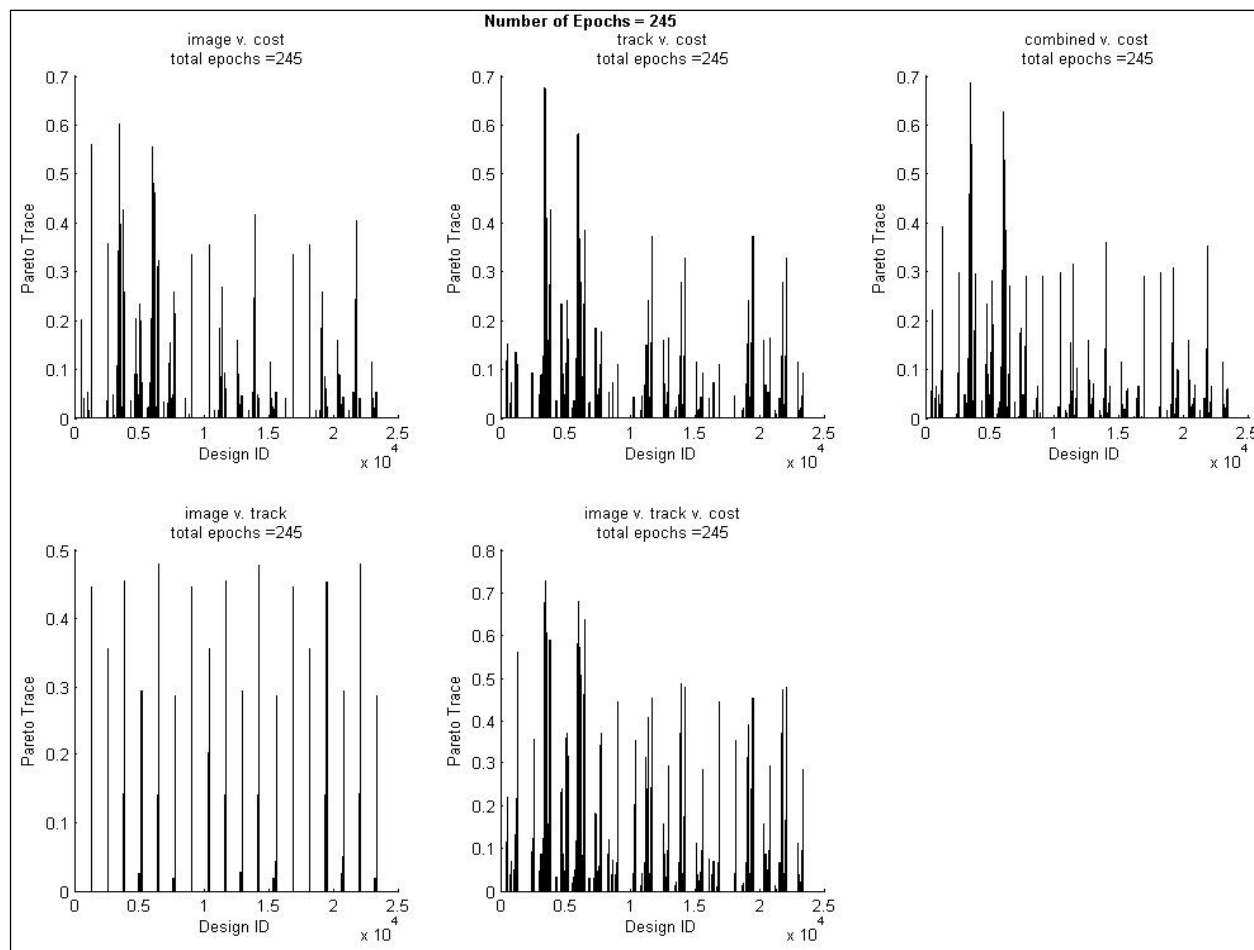


Figure 10. Pareto Trace across 245 epochs.¹⁴

4. Identifying Passive Value Robust Designs

Passive value robust designs are those that retain high value over many different epochs. These designs tend to be those that are on the Pareto Front, but are not close to physical constraints or utility limits. Given the two missions, there is a Pareto Set for the SAR stakeholder and the GMTI stakeholder for each epoch, as well as a combined Pareto Surface for both.

The Pareto Sets for an example epoch for both missions are shown below in Figure 11.

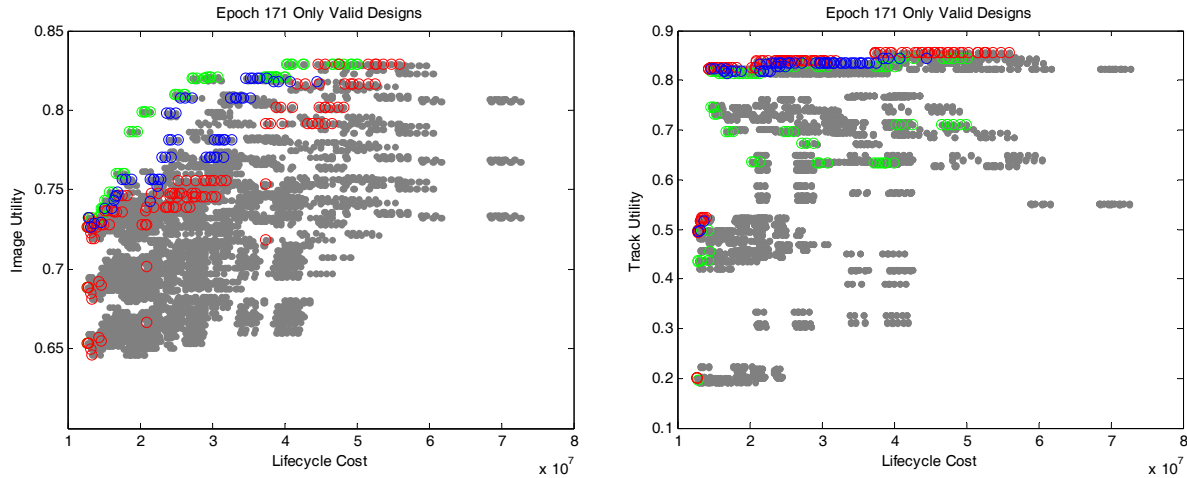


Figure 11. Pareto Sets for imaging and tracking missions (green=imaging, red=tracking, blue=compromise).

The green points are the Pareto efficient designs for Image Utility and Cost. The red points are the Pareto efficient designs for Track Utility and Cost. The blue points are those additional designs that appear on the three dimensional Pareto surface of Image Utility - Track Utility - Cost. These additional designs do not show up in the Pareto Sets for the individual missions, but rather represent efficient compromise across the missions. As a starting point, it would seem that these designs are a good place to start looking for compromise designs.

Of the 23,328 designs in the full tradespace, across epochs, Pareto Sets tend to contain several hundred designs, ranging from 500 to 1000 in any given set. In order to identify the designs that perform well across many tradespaces, Pareto Trace was used to identify those designs that occur in the Pareto Set for many epochs.

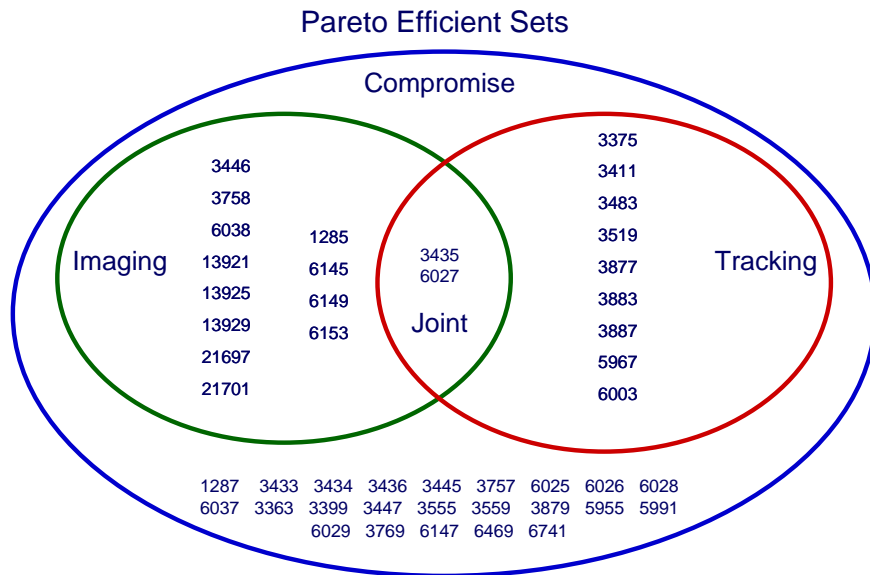


Figure 12. Pareto efficient designs across missions.

As shown in Figure 12, there are designs that work for both the SAR stakeholder as well as the GMTI stakeholder. These designs tend to be rather unremarkable beyond the fact that they are non-dominated in both utility sets. There are also many designs that appear on the compromise Pareto surface, but not in the distinct stakeholder sets.

The set of designs listed above was obtained from the data by taking the Pareto Trace for each stakeholder across each of the epochs. The designs are a subset of that set, and are the designs with Pareto Trace three sigma above the average. In general, the designs listed here have a Pareto Trace above 100, or they appear in the Pareto Sets of over one hundred epochs, out of the total 245 evaluated.

The ‘Joint’ designs are the ones that appear in both the Image only and Track only Pareto Sets. In general, the designs in these two sets are very different, however, Design 3435 and Design 6027 are Pareto Efficient for both stakeholders in many epochs.

The compromise designs are the ones that appear on the Pareto Surface for both stakeholders together with cost, but not on the individual stakeholder’s Pareto Sets. These designs are then the ones that neither stakeholder would consider Pareto Efficient, but do maximize the combined utility of the two stakeholders.

Design 3435 has the highest Pareto Trace of the design space. It also appears in the Joint set above. This indicates it may be a highly passive value robust design in that it will satisfy the stakeholders over many epochs. Design 3435 is, all in all, a fairly unremarkable design, somewhere in the middle of the design space, but in the utility cost space it consistently rises to the top. The reason for this is not immediately clear, but appears to be a combination of the coverage statistics from the orbit, and the number of satellites in this design, as well as the relatively flat area of the utility curves in the region this design inhabits. It is also a valid design for most epochs, so it is not near any of the invalidating constraints.

From the Pareto Trace analysis, Design 3435 was identified as a passively value robust design. The design variables of Design 3435 are shown in Table 4.

Table 4. Design variable values for Design 3435.

Design ID Number		3435
Pareto Trace #	Image vs. Cost	99
	Track vs. Cost:	109
	Combined vs. Cost:	109
	Image vs. Track:	0
	Image vs. Track vs. Cost:	109
Orbit Altitude (km)		1500
Walker ID		53deg. 10sat 5planes
TX Freq (GHz)		10
Antenna Area (m ²)		40
Antenna Type		AESA
Radar Bandwidth (GHz)		2
Peak Power (kW)		10
Communication Architecture		No
Communication Tactical		Yes
Tugable (<i>Era Analysis</i>)		Yes
Maneuver (<i>Era Analysis</i>)		Baseline
Procurement Strategy (<i>Era Analysis</i>)		No Spares

To obtain more detailed system tradeoff information, a focused tradespace exploration was performed centered on that design. A finer mesh of approximately 9000 points in the design space around Design 3435 was enumerated and plotted. Figure 13 and Figure 14 show the imaging utility and tracking utility of the designs versus lifetime cost.

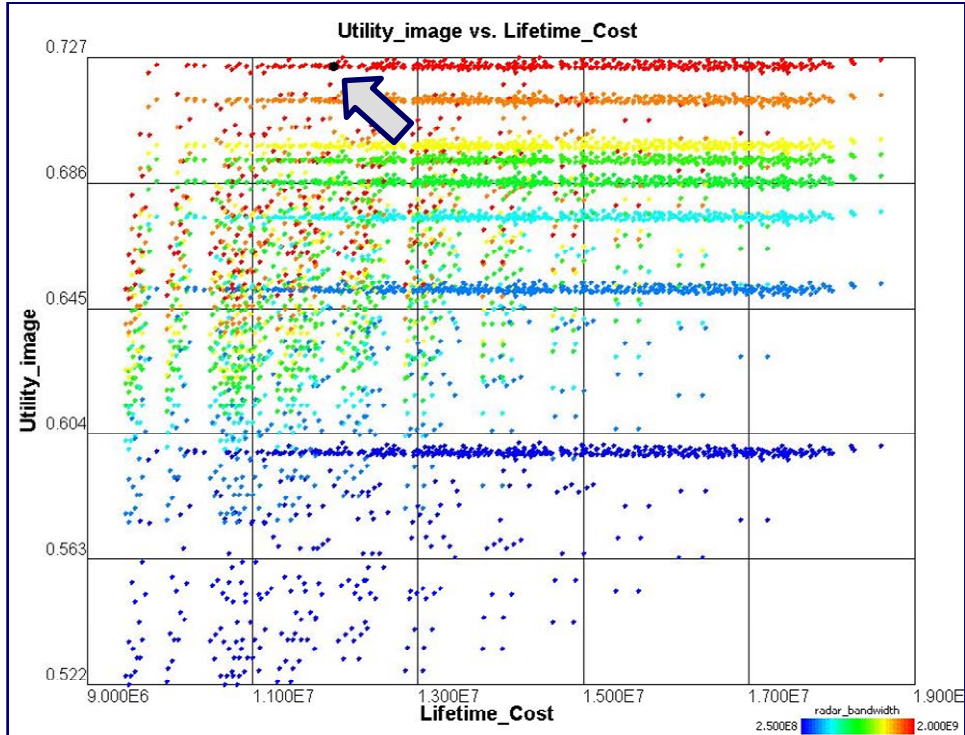


Figure 13. Imaging utility v. cost of focused tradespace around design 3435 (arrow indicating 3435), colored by increasing radar bandwidth.

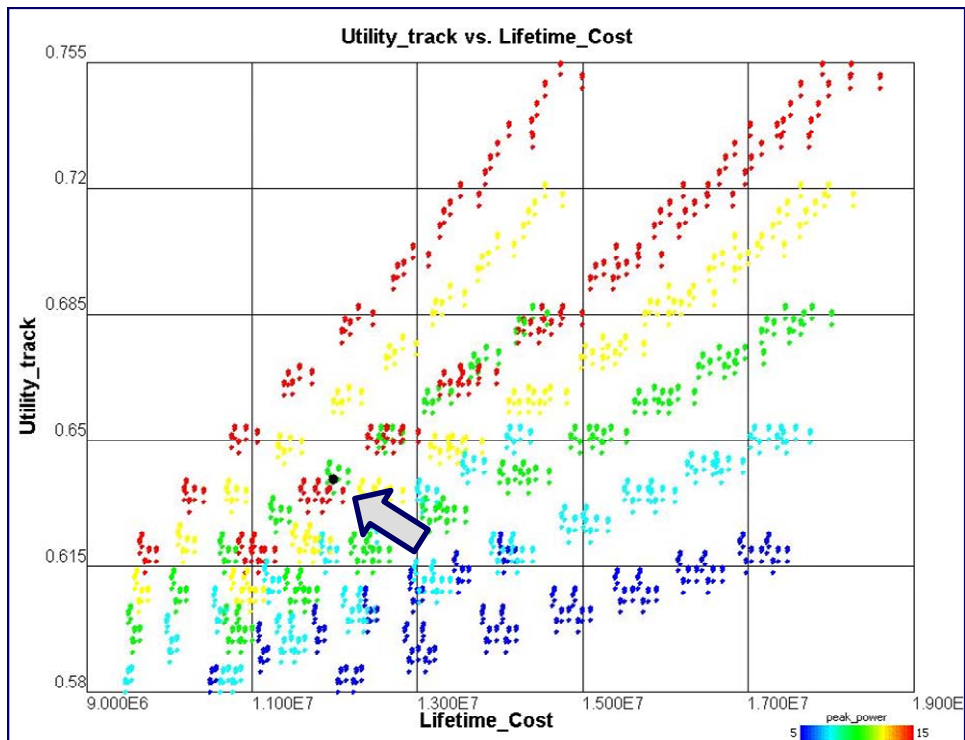


Figure 14. Tracking utility v. cost of focused tradespace around Design 3435 (arrow indicating 3435), colored by increasing peak power.

Tradespace exploration focused around identified passively value robust designs such as Design 3435 allows the modeling and analysis efforts to be efficiently used. Enumerating the design space to this level of detail in the initial

stages of analysis would require prohibitive computing power. However, using an initial rough enumeration to identify value robust candidates for further study reduces the computation power drastically. Designs that are identified through focused tradespace exploration around selected value robust designs might be good candidates for detailed design.

5. Calculating Filtered Outdegree

The “ilities” such as flexibility, adaptability, scalability, and modifiability are related to the question of whether a system can be changed, which can be quantified as the changeability of the system. In RSC, a tradespace is used to characterize design alternatives. Whether these alternatives can be changed into one another is the subject of Tradespace Network Analysis.⁷ A tradespace network is formed by proposing transition rules that can be followed by system designs during different phases of their lifecycle. For the RSC SRS case study, the lifecycle was partitioned into four distinct phases: Design (D), Build (B), Test and Integrate (T), and Operate (O). Given these phases, the design vector was inspected and possible mechanisms for transitioning from one design vector state to another design vector state were brainstormed. Eight example transition rules were proposed and used for the case study. These rules were:

1. Redesign (Design phase)
2. Redesign (Build phase)
3. Redesign (Test and Integrate phase)
4. Alter constellation through adding satellites (Operate phase)
5. Alter constellation altitude through burn with onboard fuel (Operate phase)
6. Alter constellation altitude through tug by external vehicle (Operate phase)
7. Alter constellation inclination through burn with onboard fuel (Operate phase)
8. Alter constellation inclination through tug by external vehicle (Operate phase)

A Rule-Effects Matrix, shown in Figure 15 below, lists the eight transition rules and the qualitative effect that executing each one would have on the design variables. The effect can be to increase, decrease, or increase/decrease a design variable value. *Path Enablers* are extra design variables in the design vector whose purpose is not to drive attribute performance, but rather to allow for the execution of transition rules at lower cost (e.g. buying spare satellites in advance lowers the cost of Rule 4: add satellites to constellation). The *Change Origin* column indicates whether the change mechanism can be considered to be a flexible-type or adaptable-type change, which is determined by the instigator of the change: if done by the system, it is adaptable-type; if done by a system-external force, it is flexible-type. The *Phase Allowed* column indicates that these rules are only available during specific phases in the system lifecycle.

Rule-Effects Matrix SRS		Design Variables											Path Enablers			Change		Phase allowed			
		OA	WN	TF	AS	AT	RB	PP	CA	CT	T	M	CO	Origin	Flex	Adp	D	B	T	O	
		DV1	DV2	DV3	DV4	DV5	DV6	DV7	DV8	DV9	IV1	IV2	IV3				1	2	3	4	
Rule	Redesign (Design)	R1	↕↕	↕↕		↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕				
	Redesign (Build)	R2	↕↕	↕↕		↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕				
	Redesign (Test)	R3	↕↕	↕↕		↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕	↕↕				
	Alter Const (add sats)	R4	↕↕	↑											↕↕	↕↕					
	Alter altitude with onboard fuel	R5	↕↕											↕↕	↕↕	↕↕					
	Alter altitude with tugging	R6	↕↕										↕↕	↕↕	↕↕	↕↕					
	Alter const (change inclination) with onboard fuel	R7		↕↕									↕↕	↕↕	↕↕	↕↕					
	Alter const (change inclination) with tug	R8		↕↕									↕↕	↕↕	↕↕	↕↕					

Figure 15. Rule-Effects Matrix for Satellite Radar System case.

These transition rules were coded into a logical MATLAB routine to determine whether a particular design is reachable from another design through following a particular rule, and if so, calculate the cost and schedule impact of executing a particular transition. Three transition matrices were created: T^{Acost} , T^{Atime} , and T^{AV} . The entries of these matrix indicate whether a transition is allowed (non-zero value) and for how much of the appropriate cost type (dollars, days, delta-v). Each matrix was $N \times N \times K$, where N is the tradespace size (23,328) and K is the number of rules (8). Figure 16 below illustrates the relationship between the transition matrix and the tradespace network.

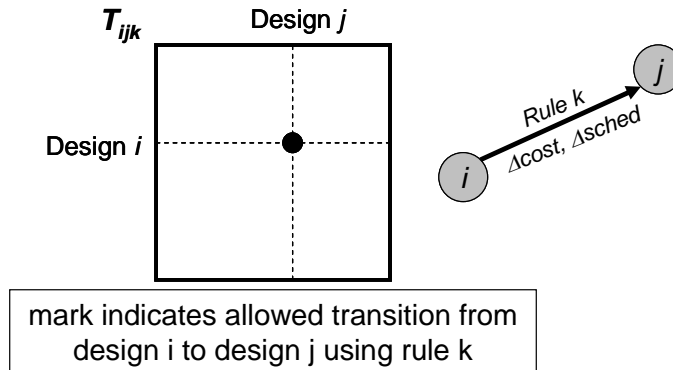


Figure 16. Illustration of transition matrix used for tradespace networks.

The transition matrices must be recalculated for each epoch of interest, as only feasible designs are considered in the creation of the matrix, so the allowable origin and destination designs will vary by epoch, as will the execution cost and time for following particular rules.

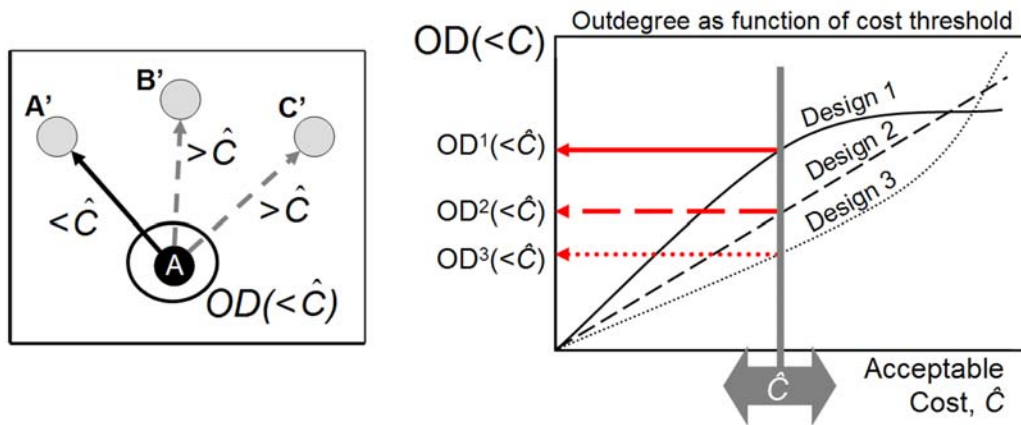


Figure 17. Filtered Outdegree and the Filtered Outdegree Function¹⁵

Once the transition matrices are created, the filtered outdegree function of designs can be calculated. As illustrated in Figure 17, the Filtered Outdegree is the number of outgoing arcs, or transitions, from a given design to all other designs, at acceptable cost. In theory, the maximum outdegree for a given tradespace is $(N-1) \times K$, where N is the size of the tradespace and K is the number of rules. The “filter” represents the “acceptable cost” criterion, meaning that transitions that cost more dollars/time/delta-v than is acceptable will not be followed or counted in the metric. The filtered outdegree metric quantifies the perceived changeability of a particular design for a given acceptable “cost” threshold. (Note “cost” can be in dollars, time, or delta-v for this case study. For the following examples in this section, only dollars are used for the purpose of clarity.) As the acceptable cost threshold varies, so too does the perceived changeability (Filtered Outdegree) of a given design. It has been seen that different designs have different shapes in the filtered outdegree function, providing insights into key cost thresholds for changeability and also as a focus point for designers to motivate efforts to lower the cost of transitioning, through creating new path enablers or transition rules.¹⁵

Since Design 3435 was identified to be of high interest in the Pareto Trace analysis, it will be used here to illustrate some of the changeability data derived from the case study.

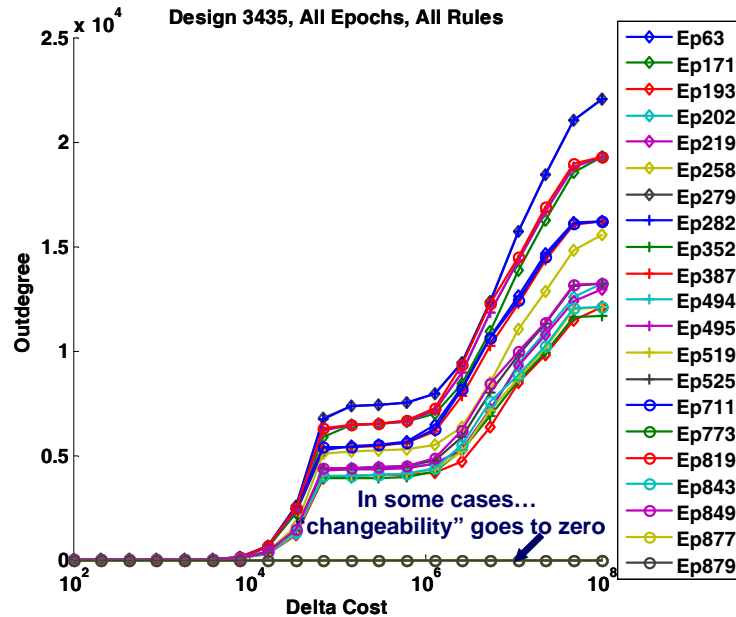


Figure 18. Filtered Outdegree function for Design 3435 across many epochs.

Figure 18 illustrates several key points related to changeability across the SRS tradespace. First, the outdegree—the metric of changeability (number of other states to which the design could change)—is a function of acceptable change cost with several critical value thresholds. Second, the Filtered Outdegree not only varies by willingness to pay, but also by context, which affects the cost of transition, as well as the feasibility of prospective end states for the design. The shapes of these curves is typical of most designs in the tradespace and suggest several cost levels that “activate” new mechanisms, akin to being able to pay for a new “option” to change the system to a different state. Since 10^5 \$K appeared to be one of the interesting regions of the function, it was selected as the cost threshold for the following analysis.

The rules that are available in a given lifecycle phase play a critical role in determining the changeability of a design alternative. Some designs are more changeable than others at different phases of the lifecycle. Figure 19 below illustrates the Filtered Outdegree for the feasible designs in Epoch 171 across the four phases of the lifecycle. The gaps in the figure do not necessarily reflect designs with no-changeability, but can also correspond to infeasible (unacceptable) designs in the epoch. Figure 20 below lists the feasible designs in this epoch (tradespace size=23328, feasible designs=6426). It turns out that all designs have maximum outdegree during the Design phase for this cost threshold, meaning that any design can be redesigned into any other design during the Design phase. As the system progresses from Design to Build to Test to Operate, the filtered outdegree drops, partly as a result of the increased expense of changing the system later in its development, and also partly due to the difficulty in accessing the system in order to change it.

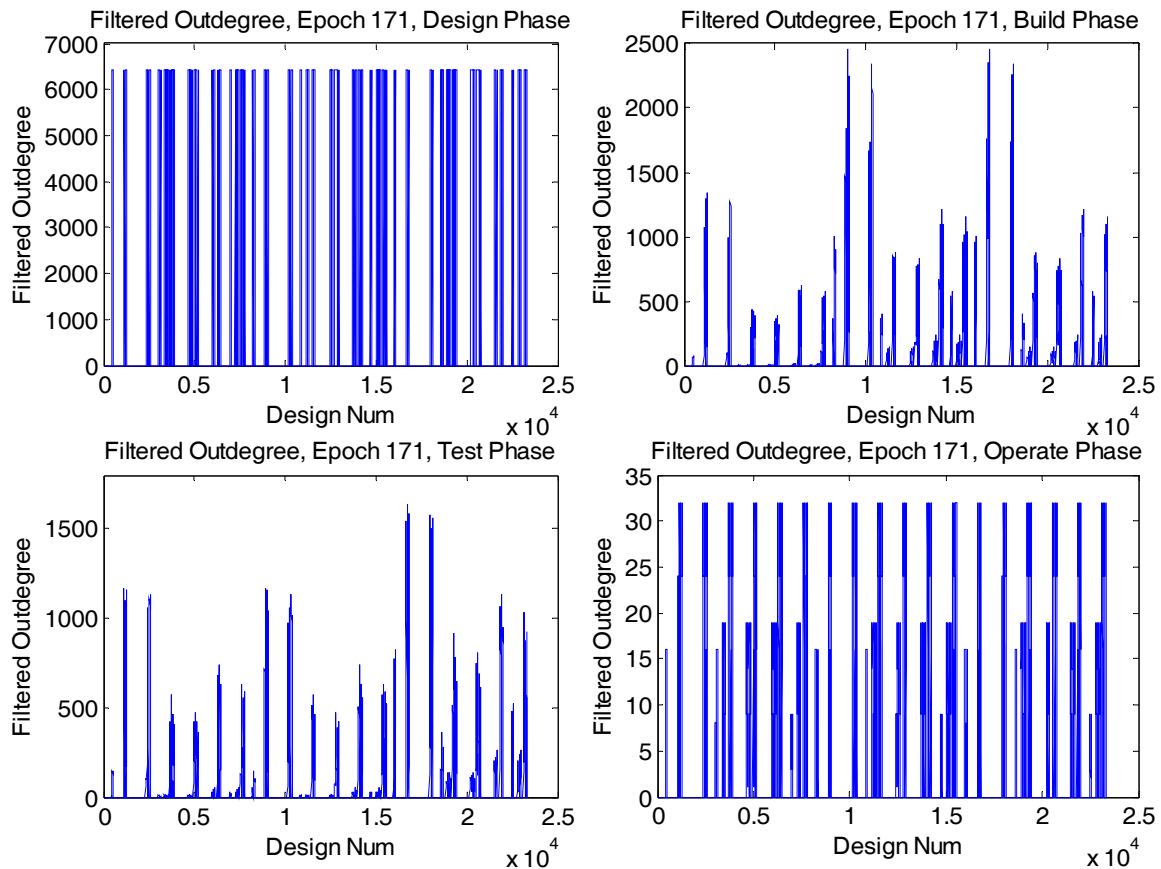


Figure 19. Filtered Outdegree (cost threshold= 10^5 \$K) by phase for Epoch 171.

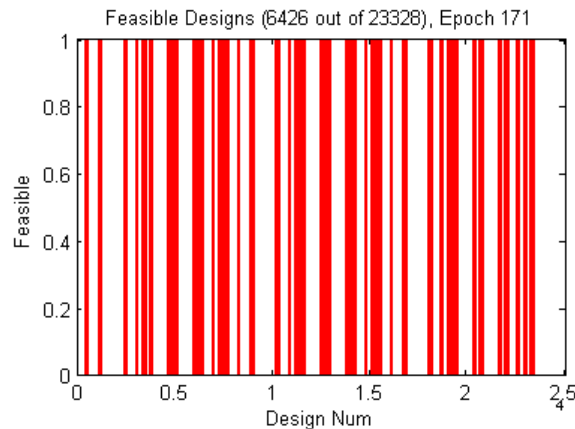


Figure 20. Feasible designs for Epoch 171.

Other epochs had similar results, with the tradespace fully connected during the Design phase through the “redesign” transition rule, and lesser connection during later phases. In general there was an extremely large drop off in outdegree when the system entered the Operate phase. Even though there were more transition rules available during this phase (Rules 4-8), they were rarely executable at acceptable cost.

6. Identifying Changeable Designs

After generating the transition matrices for the tradespace in each epoch of interest, and calculating the Filtered Outdegree functions of the designs, the next step was to identify designs that were most changeable per phase in the epochs of interest. In order to do so, four epochs are highlighted in Table 5 below: Epochs 63, 171, 193, and 202.

These correspond to the unique epochs that formed the basis for Era 1.⁸ The highest Filtered Outdegrees were selected and normalized by the size of the feasible tradespace for that epoch. A score of 100 means that the entire tradespace was accessible from the indicated design. A score of 36.8 means that 36.8% of the tradespace was accessible from the most changeable design in the tradespace during that epoch and indicated phase. It is important to note that the “best across” design number means that the scores are not necessarily from the same design, but rather indicate the best possible score for each phase in a given epoch. Design 3435’s score is indicated per epoch as a point of comparison. It is immediately apparent that the highly passive value robust design is not very changeable when compared to other designs in the tradespace.

Table 5. Highest Filtered Outdegree as percent of feasible designs across four epochs and phases.

Epoch	Design #	Phase			
		Design	Build	Test	Operate
63	Best across	100	36.8	45.5	0.6
63	3435	100	0.04	0.3	0.2
171	Best across	100	38.2	25.5	0.5
171	3435	100	0.02	0.1	0.1
193	Best across	100	16.0	14.8	0.8
193	3435	100	0	0.05	0.2
202	Best across	100	14.6	13.5	0.7
202	3435	100	0	0.05	0.2

Since we know that these are the best scores for changeability, the next question is whether single designs can score well in all phases and whether particular designs appear best across multiple epochs.

Table 6. Percent Filtered Outdegree across four phases and epochs for most changeable designs.

Epoch	Design #	Phase			
		Design	Build	Test	Operate
63	8921	100	36.8	30.2	0.4
63	16697	100	36.8	37.3	0.4
63	16701	100	36.2	45.5	0.6
63	1089	100	3.5	10.8	0.6
171	9029	100	38.2	17.6	0.4
171	16805	100	38.2	23.6	0.4
171	16701	100	27.0	25.5	0.5
171	1089	100	1.8	3.4	0.5
193	9029	100	16.0	8.2	0.6
193	16805	100	16.0	13.8	0.6
193	16809	100	15.0	14.8	0.8
193	1101	100	2.0	2.9	0.8
202	9029	100	14.6	7.4	0.6
202	16805	100	14.6	12.7	0.6
202	16809	100	13.6	13.5	0.7
202	1101	100	1.8	2.7	0.7

Designs that were able to score best in each phase across the four epochs considered are shown in Table 6. In particular, Designs 16701, 16805, and 16809 appear as highly changeable across the four epochs, scoring at or close to the highest in a given phase for a given epoch. (Bold indicates highest score for that tradespace epoch in that phase.)

The designs of interest are listed in Table 7 below, with comparisons to follow. First, the three designs are virtually the same, with only two design variables changing across them: antenna area and maneuver package. Two of the three have the largest antenna area in the tradespace (100 m²) and two have the highest maneuver package (6=baseline fuel x4). Second, the three designs all have high scores in the path enabler values. In fact all but one path enabler is at its maximum value. The result confirms the expectation that path enablers exist to lower transition costs and enable additional transitions. By contrast, Design 3435 has the lowest values of the path enablers allowed in the tradespace, as from a passive value perspective, adding these path enablers do not contribute to attribute performance and only increase cost. Since changeability is not valued through utility, it is not “rewarded” for the additional cost of adding a path enabler.

Table 7. Most changeable designs in the four epochs, along with Design 3435.

		16701	16805	16809	3435
Design Variables	Orbit Alt (km)	800	800	800	1500
	Walker ID (53deg, 4=20 sats, 5 planes; 3=10 sats, 5 planes)	4	4	4	3
	Tx Freq (GHz)	10	10	10	10
	Ant Area (m ²)	40	100	100	40
	Ant Type (0=AESA)	0	0	0	0
	Radar Bandwidth (GHz)	1	1	1	2
	Peak Power (kW)	20	20	20	10
	Comm Arch (0=backbone, 1=ground)	0	0	0	1
	Tactical Comm (0=yes)	0	0	0	0
Path Enablers	Tugable (1=yes)	1	1	1	1
	Maneuver package (4=baseline, 5=x2, 6=x4)	6	5	6	4
	Constellation Option (1=nothing, 3=sparcs built)	3	3	3	1

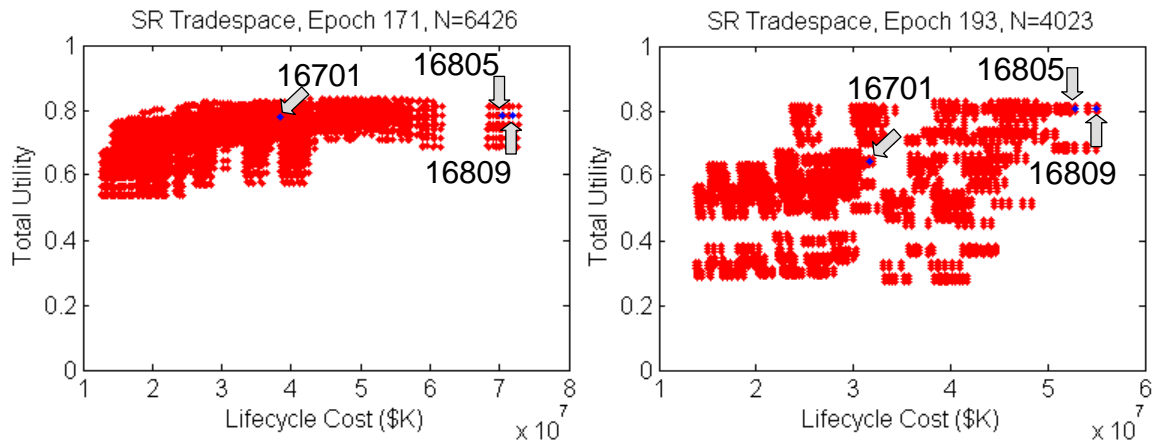


Figure 21. Tradespaces for Epochs 171 and 193 with three highly changeable designs.

Now that highly changeable designs have been identified, the next step is to determine how they fit into the overall tradespace. Figure 21 illustrates where these three designs fall into Epoch 171 and Epoch 193. Two of the three are very high cost options, however one of them, Design 16701, is an intermediate compromise, and it is clearly not Pareto Efficient in value delivery, even though it is highly changeable.

Since the most highly changeable designs may not be attractive from a mission utility perspective, it may be instructive to inspect how changeable the entire tradespace is as a function of lifecycle phase.

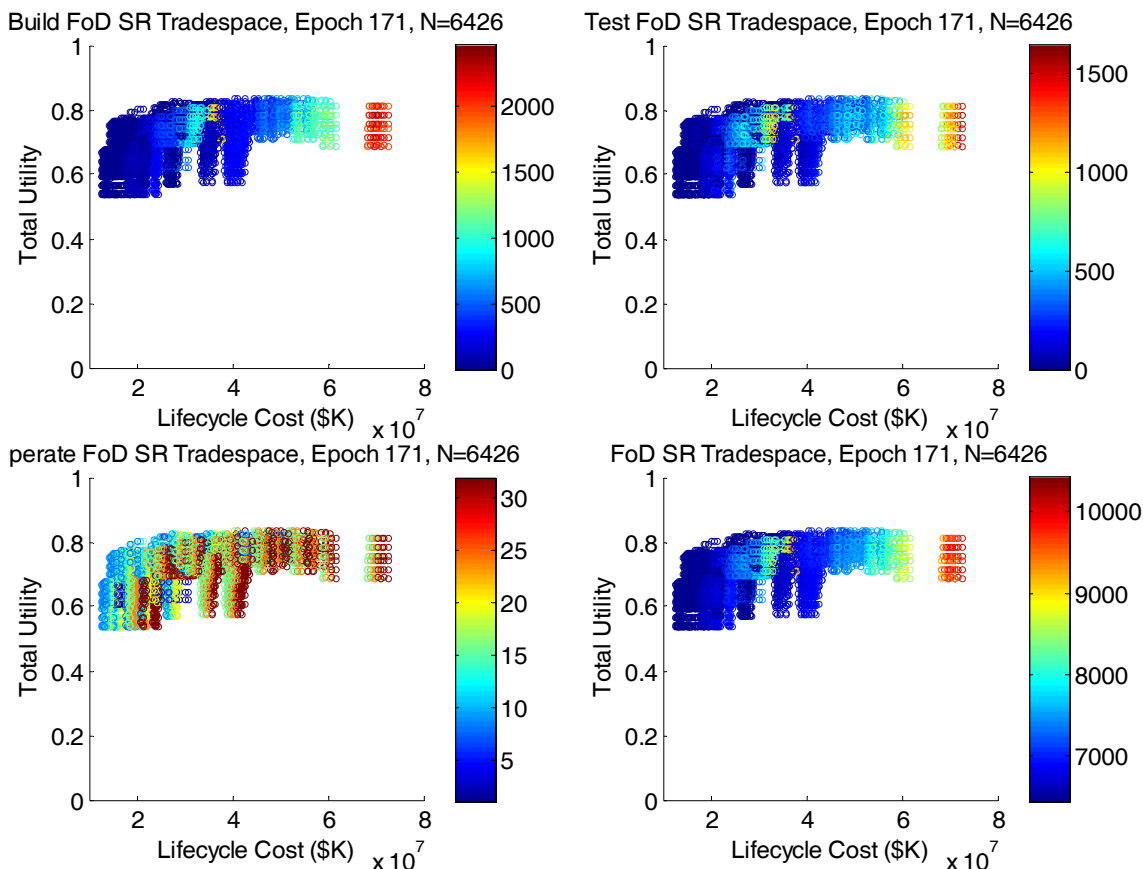


Figure 22. Changeability tradespaces for Epoch 171 across lifecycle phases.

Since during the design phase all designs can maximally transition, all other phases are depicted in Figure 22 for Epoch 171. Overall there is a clear trend that changeability comes at a price, though there appear to be some lower cost options with high changeability. If highly changeable systems are prized, this type of analysis can be used to identify design types and tradespace regions that may provide good tradeoff of value and changeability. In particular, if one does not want to pursue a strategy of seeking passively value robust designs (which require some foreknowledge of possible epochs), one can seek actively changeable designs, which are ones that can change over time to move into high regions of the tradespace, as real epochs and eras unfold over time.

V. Discussion

Through the application of RSC to the SRS design problem, a number of insights were discovered, though no definitive prescriptive conclusion can be made, as the models and data used were more “illustrative” than technically in-depth and rigorous.

A. Insight 1: quantitative investigation of “requirement” effect (e.g. on yield)

As discussed in section IV.B.1 and IV.B.2, since application of RSC results in the assessment of a large number of design alternatives on a common basis (e.g. decision maker-defined attributes), basic statistics can be applied to

understand the effects of expectation constraints on a design space. This means that given a design space—the span of the design variable enumerations—the accessibility of an attribute space can be assessed through models and simulations. The display of the distribution of frequency versus levels of attributes (Figure 7 and Figure 8) gives rapid first order insight into how difficult (or easily) achieved different levels of attribute performance can be attained. For attributes whose performance are below “acceptable” levels, a designer can either generate new levels of design variables (i.e. new design alternatives), or have a discussion with the decision maker whose needs are not being met to better understand the “rigidity” the acceptability cut-off for that attribute. Requirements in this context can be thought of as the band of minimally acceptable attribute values in this distribution. Especially when looking at these distributions across decision makers, and seeing the yield (number of “valid” designs versus the total number of designs investigated), one can see how strong a driver different attribute cut-offs may be, helping to focus the attention of the designer on potential areas of concern for solving stakeholder needs. For the SRS case study this issue arose for the tracking mission attributes of *target acquisition time* and *minimum detectable velocity*. Many designs failed to reach minimum acceptable levels of these attributes and therefore were deemed “infeasible” for both missions (since the SRS must satisfy both missions, to fail in one is to fail for both). Designs “close” to cut-offs may be most sensitive to changes in expectations, especially as epochs change and new contexts arise.

B. Insight 2: quantify “value robust” and correlate to design features

As the number of designs evaluated increases (23,328 per tradespace) and the number of epochs evaluated increases (245 of 648 tradespaces), the amount of data relating the behavior of various designs under various contexts and expectations grows tremendously. Inspecting a series of Utility-Cost plots for each of these tradespaces would be time-consuming and difficult to generate aggregate insights. When thinking about the dynamic value delivery problem (i.e. ‘how can SRS continue to be useful across changing contexts and expectations?’), an aggregate measure of value sustainment across epochs becomes useful. The Pareto Trace metric, defined as a count of the number of times a design appears in a Utility-Cost Pareto Set across epochs, was shown to be a useful screening metrics to quickly identify designs that stay “valuable” across changing contexts and needs. Applying this metric to include not only Utility-Cost per decision maker (image and track), but also to include $Utility_{image} - Utility_{track} - Cost$, resulted in the generation of supersets of designs that contained win-win solutions as well as efficient compromises across the decision makers. Looking at this larger set across epochs allowed for the identification of design alternatives that may not have been selected when optimizing for one mission alone. The “compromise” designs appear to be dominated on a per decision maker basis, but in fact are the “best” that can be done for the cost across the decision makers (the blue designs in Figure 11).

The Pareto Trace metric only captures the value robustness of a design that does not have to change to continue to deliver value, and only across epochs that have been enumerated and evaluated. Value robustness across unanticipated epochs is not guaranteed unless similar to epochs already assessed. The Filtered Outdegree metric attempts to hedge against the “unknown” uncertainties by indicating the changeability of various designs. The transition rules, which specified the allowed mechanisms for designs to change states, were mostly related to “redesign” and to orbit changes. Due to cost and schedule constraints, captured as a “willingness to pay dollars and time” to enact a change, the number of transitions open to a given design decreased rapidly as the program progressed from design to build to integrate and test to operations. Even though this case application was mostly done at an “illustrative” level, the results reflect the conditions affecting most space programs: high cost to change designs, and the inaccessibility to change the system once in operation. Investigating the designs that were most changeable showed that they had the highest levels of the “path enablers” set, design features chosen to enhance a system’s ability to change, rather than perform its mission. For SRS these involved extra fuel, grapple points for a space “tug”, and building and deploying extra satellites to allow for constellation expansion or replacement. The most changeable designs appeared to deliver the same utility, but at a much higher cost, as alternative Pareto efficient designs, however if the epochs changed, these designs may be more likely to be able to retain high value through executing change. In the case of unanticipated epochs, the heuristic of “more changeable is better” may be an approach to take, however, the costs of acquiring this changeability is readily quantified (epoch-independent), but the benefits are not readily quantified (epoch-dependent, so one must at least guess).¹⁶

Another question is whether one should not just care about whether a design is changeable, but whether a design can change to a higher value state. Ongoing research is investigating developing metrics that capture this benefit aspect of change¹⁷, however the context-dependent nature of benefit requires some ability to anticipate the epoch of the destination design, including expectations on the system. The inherent bias towards readily accounting for the cost, but not the benefit, of changeable designs makes this type of quantification difficult.

C. Insight 3: quantify key system tradeoffs

As discussed in prior work using Multi-Attribute Tradespace Exploration,^{18,19} knowledge generated through investigating the mapping of design alternatives to attributes allows designers to not only better understand the “easy” and “difficult” performance levels for potential system designs, but also to better communicate the impact on system feasibility by the expectation levels of attributes set by decision makers. The SRS case study reinforced this lesson through clearly illustrating a design heuristic for radar (power-aperture product drives much of performance), as well as illustrating the tensions between mission areas, as discussed above on the impacts of “requirements” across missions. When demonstrating the case application to government practitioners, it was noted that this capability does not replace the wisdom of trained engineers, however it does ensure that wisdom becomes predictably accessible to many engineers, including those with less experience. Asking questions about the existence of tradeoffs leads to learning more about the system technologies and design, as well as expectations. Some tradeoffs are inherent in the solution to a problem (e.g. using Keplerian orbits requires trading coverage versus dwell time), while others are inherent to assumptions about a problem (e.g. necessity of using Keplerian orbits in the first place). Visual and statistical exploration of tradeoffs during tradespace exploration result in artifacts (tradespace plots) and discussions that create traceability to help justify the ultimate selection of a design for further development, as well as contingency plans if the need for design change arises.

VI. Conclusion

Responsive Systems Comparison Method is a set of seven processes aimed at formally guiding a tradespace exploration study to assess a large number of design alternatives on a common value-centric basis across changing contexts and needs. The ultimate goal is to generate knowledge about tradeoffs, compromises, and risks to a system development project, and allow for ready communication among designers and stakeholders. Ongoing research continues to refine and apply the seven processes to both aerospace and other complex system development projects in order to produce a methodology that can prescriptively generate system evolutionary strategies and justifications to encourage the acquisition of value robust system architectures that can sustain value delivery to stakeholders across dynamic futures.

Acknowledgments

The authors are grateful to the RSC development and satellite radar case study teams for their hard work and dedication. Funding for this work was generously provided by the U.S. Government and the Systems Engineering Advancement Research Initiative (seari.mit.edu), a consortium of systems engineering leaders from industry, government, and academia.

References

- ¹ 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, 2008, pp. 15.
- ² Lempert, R. J., Popper, S. W. and Bankes, S. C., *Shaping the Next One Hundred Years--New Methods for Quantitative Long-Term Policy Analysis*, RAND, Santa Monica, CA, 2003, pp. 187.
- ³ Harries, C., "Correspondence to What? Coherence to What? What is Good Scenario-based Decision Making?," *Technological Forecasting and Social Change*, Vol. 70, No. 2003, pp. 797-817.
- ⁴ McManus, H. and Hastings, D., "A Framework for Understanding Uncertainty and Its Mitigation and Exploitation in Complex Systems," *IEEE Engineering Management Review*, Vol. 34, No. 3, Third Quarter 2006, pp. 81-94.
- ⁵ Ross, A. M., McManus, H., Long, A., Richards, M. G., Rhodes, D. H. and Hastings, D., "Responsive Systems Comparison Method: Case Study in Assessing Future Designs in the Presence of Change," *AIAA Space 2008*, San Diego, CA, 2008, pp. 9.
- ⁶ 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, pp. 20-28.
- ⁷ Ross, A. M. and Hastings, D. E., "Assessing Changeability in Aerospace Systems Architecting and Design Using Dynamic Multi-Attribute Tradespace Exploration," *AIAA Space 2006*, 2006-7255, AIAA, San Jose, CA, 2006, pp. 18.
- ⁸ Roberts, C. J., Richards, M. G., Ross, A. M., Rhodes, D. H. and Hastings, D. E., "Scenario Planning in Dynamic Multi-Attribute Tradespace Exploration," *3rd Annual IEEE Systems Conference*, Vancouver, Canada, 2009, pp. 7.
- ⁹ Ross, A. M. and Rhodes, D. H., "Architecting Systems for Value Robustness: Research Motivations and Progress," *2nd Annual IEEE Systems Conference*, Montreal, Canada, 2008, pp. 8.
- ¹⁰ Keeney, R. L. and Raiffa, H., *Decisions with Multiple Objectives--Preferences and Value Tradeoffs*, 2nd, Cambridge University Press, Cambridge, UK, 1993, pp. 569.

¹¹ Ross, A. M., "Multi-Attribute Tradespace Exploration with Concurrent Design as a Value-centric Framework for Space System Architecture and Design," Dual-SM Thesis, Aeronautics and Astronautics and Technology & Policy Program, Massachusetts Institute of Technology, Cambridge, MA, 2003.

¹² "Alternatives for Military Space Radar," Congressional Budget Office, Pub. 1609, Washington, DC, Jan 2007.

¹³ ATSV, ARL Trade Space Visualizer, Software Package, Ver. 3.3.5, Penn State University Applied Research Laboratory, <http://www.tradespaceexploration.psu.edu/>, State College, PA, 2008.

¹⁴ Ross, A. M., Rhodes, D. H. and Hastings, D. E., "Using Pareto Trace to Determine System Passive Value Robustness," *3rd Annual IEEE Systems Conference*, Vancouver, Canada, 2009, pp. 7.

¹⁵ Ross, A. M., Rhodes, D. H. and Hastings, D., "Defining Changeability: Reconciling Flexibility, Adaptability, Scalability, Modifiability, and Robustness for Maintaining Lifecycle Value," *Systems Engineering*, Vol. 11, No. 3, Fall 2008, pp. 246-262.

¹⁶ Shah, N. B., Viscito, L., Wilds, J., Ross, A. M. and Hastings, D., "Quantifying Flexibility for Architecting Changeable Systems," *6th Conference on Systems Engineering Research*, Los Angeles, CA, 2008,

¹⁷ Viscito, L., "Quantifying Flexibility in the Operationally Responsive Space Paradigm," SM Thesis, Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, MA, 2009.

¹⁸ Ross, A. M. and Hastings, D. E., "The Tradespace Exploration Paradigm," *INCOSE International Symposium 2005*, Rochester, NY, 2005,

¹⁹ Ross, A. M., "Managing Unarticulated Value: Changeability in Multi-Attribute Tradespace Exploration," PhD Dissertation, Engineering System Division, Massachusetts Institute of Technology, Cambridge, 2006.