

MIT Open Access Articles

*Role for Interactive Tradespace Exploration
in Multi-Stakeholder Negotiations*

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

Citation: Ross, Adam, Hugh McManus, Donna Rhodes, and Daniel Hastings. "Role for Interactive Tradespace Exploration in Multi-Stakeholder Negotiations." In AIAA SPACE 2010 Conference & Exposition. American Institute of Aeronautics and Astronautics, 2010.

As Published: <http://dx.doi.org/10.2514/6.2010-8664>

Publisher: American Institute of Aeronautics and Astronautics

Persistent URL: <http://hdl.handle.net/1721.1/81979>

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



Role for Interactive Tradespace Exploration in Multi-Stakeholder Negotiations

Adam M. Ross*

Massachusetts Institute of Technology, Cambridge, MA, 02139

Hugh L. McManus†

Metis Design, Cambridge MA 02139

and

Donna H. Rhodes‡ and Daniel E. Hastings§

Massachusetts Institute of Technology, Cambridge, MA 02139

The significant time, effort, and resource expenditures needed to design and develop aerospace systems motivate on-going research into developing methods for generating, evaluating, and selecting candidate system solutions that can deliver more benefit for a given cost. Compounding the problem is the multiplicity of perspectives of the many stakeholders for such systems, altering the meaning of “benefit” and “cost” depending on the stakeholder considered. Tradespace exploration techniques have been used in the past to generate large datasets in order to gain insights into design-value, cost-benefit tradeoffs for complex aerospace systems. Using interactive tradespace exploration to support multi-stakeholder negotiations can reveal these tradeoffs not only for individuals, but also across a group. A method is introduced and applied to two aerospace cases in order to explore the potential for interactive tradespace exploration to support stakeholder negotiations. Preliminary results indicate the method to be a rapid and beneficial technique, which generated compromise alternatives, guided the elicitation of previously unarticulated information, and resulted in increased confidence and solution buy-in of participating stakeholders.

I. Introduction

TYPICALLY space systems require large amounts of resources and political capital in order to be developed, motivated by meeting the needs of many potential stakeholders. In fact, the justification for the large expenses incurred in the development of these systems is the ultimate promise of providing essential benefit (e.g., national defense, or invaluable scientific insights) to many individuals and organizations. Unfortunately for the engineers designing and developing these systems, the problem of meeting these promises is an asymmetric one: it is relatively easy to incur large expenses, while it is relatively difficult to incur large benefits (i.e., there are many more possible “bad” potential systems than “good” ones.) As the number of stakeholders, and therefore likely diverse expectations, increases, the more difficult it may become to satisfy everyone. The problem of satisfying multiple stakeholders is not a new one,¹ however, advances in computation and visualization capabilities may improve the prospects for doing so. At its core, the problem relates to incomplete information (e.g., what does each person want, how do those wants relate to one another, and what is actually possible to accomplish with a system), as well as allocation of resources (e.g., what to build, who pays, and when should it get built). In absence of perfect information, two example approaches to this problem are to compromise on preferences, or to compromise on design.

In the first approach, key stakeholders are gathered to discuss and aggregate their preferences in order to drive the design decisions. The benefit of this approach is that it is a human process and through interaction, stakeholders

* 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

may discover unarticulated desires. Additionally a straightforward set of requirements may result, which provides a single set of criteria for determining a “successful” design. Additionally, the process itself is relatively transparent, helping to build trust among the parties involved. The downside of this approach is that it prematurely constrains the possible set of alternatives, where stakeholders, unaware of potential “win-win” systems, may “compromise away” preferences that could have been easily met “for free.” The end result may be the development of a system that only partially satisfies everyone, even if it meets the aggregated set of requirements. In the second approach, each stakeholder is treated separately, for example distinct mission stakeholders, and the best solution for each is developed and then merged together in one of two ways: the best features from each solution are combined to form a “mixed-breed” system, or a new system is designed to have the performance that is a superset of all individual systems. The first combination solution can result in a system that is no good to anyone (especially if the parts conflict), and the second combination solution may have excellent performance, but may not succeed due to cost, scheduling, or technology development issues.

This paper demonstrates an emerging use of interactive tradespace exploration as a means for communicating multiple individual stakeholder preferences and “best” system designs in support of discovering superior group “compromise” designs. A method using question-driven interactive tradespace exploration is introduced, along with a process for three individual stakeholders to discover their own “best” designs, the tradeoffs among these best designs, and a structured approach for identifying and developing a group “best” design. The method is illustrated using two case studies: X-TOS,² a single satellite-based science mission in support of neutral atmospheric density modeling, and SRS,³ a constellation of satellites-based defense mission for real-time radar imaging and tracking of targets of interest. The method is shown to hold promise in reducing incomplete information, as well as generating trust among the negotiating parties.

A. Prior Work in Design Negotiation

Research in negotiation in engineering design has pointed out that design negotiation differs from other types of negotiation, such as labor dispute⁴ or house-buying,⁵ in that design negotiation tends to be a cooperative endeavor, where all parties tend to actively seek a mutually acceptable solution and may be willing to give up individual levels of satisfaction in order to ensure the ultimate development of the design.⁶ Given the cooperative nature of design negotiation, work using game theory has been done to assist negotiation agent formulation, whereby the utilities of individual agents are jointly considered and traded to develop compromise solutions.⁷ Various formulations have been used to represent the preferences of multiple agents, including both utility and fuzzy-based approaches.⁸ Complex, real design problems require human intervention, however, especially when one considers the prevalence of incomplete or ambiguous information, both in terms of stated preferences of stakeholders, as well as information about the design problem itself. This recognition has led to the development of automated multi-attribute negotiation, which can be used to assist autonomous agents, as well as human decision makers, in order to streamline and accelerate the negotiation process.⁹ Ref. 9 discusses the additional difficulties facing a multi-attribute, multi-stakeholder negotiation through an illustration of a single attribute, simple decision problem. The difficulties in the simplified case include the following dilemmas facing a participant: assessing an individual’s utility function is multi-dimensional, assessing other participants’ utility functions is more difficult, the context may vary over time, and discovering Pareto optimal solutions is more difficult. Coincidentally, each of these issues has arisen in the context of tradespace exploration for individual decision makers, suggesting a potential may exist for using interactive tradespace exploration to support multi-attribute, multi-stakeholder negotiation.¹⁰ For example, Interactive Decision Maps have been used to help discover and gain intuition over tradeoffs on complex multi-criteria Pareto Surfaces.¹¹ Additionally, prior work has demonstrated the power of interactive multi-dimensional data exploration tools,^{12,13} as well as the emergent benefits of discovering unexpressed desires when confronted with an array of options during the design process.¹⁴ Some work has even applied this to interactive design using fuzzy sets, with humans in the loop, but the application has tended to be to smaller, less complex (than aerospace) systems.¹⁵

Looking at current engineering design approaches, a problem that emerges is that many methods inadvertently or arbitrarily aggregate across stakeholder preferences and therefore may result in system designs that either fail to meet expectations, or are far less capable or more expensive than if properly developed. Worse, in practice, most attempts to create mathematical decision criteria, even for a single stakeholder, are prone to imperfections and omissions.^{16,17} These imperfections are both multiplied by having multiple stakeholders, and stressed by the negotiation process, which may cause stakeholders to modify their criteria, or discover the importance of previously unexpressed ones.

B. Tradespace Exploration

This paper addresses the use of tradespace exploration with multiple decision makers represented both mathematically *and* as humans in the loop. A tradespace, for the purposes of this paper, is represented by a data set (computed ahead of time) that includes a large number of possible designs, and their analyzed performances.^{18,19} The designs are expressed parametrically by varying the elements of a *design vector*. The performance is expressed in terms of *attributes*, metrics of interest to, and chosen by, the stakeholders. The upper and lower bounds of these metrics (expressed as worst acceptable performance, beyond which the system is useless to the stakeholder, and best needed performance, beyond which there is no extra value) are specified *for each stakeholder*. Preferences for each of these attributes for each stakeholder are explicitly captured as weights and utility curves. A computational model is used to calculate the attribute values for some or all of the possible designs in the design vector, their single (SAU) and multi-attribute utilities (MAU) for one or more stakeholders, and their costs. These values, along with key intermediate variables that are helpful for describing and further analyzing the designs, are stored in a database, and are shown in Fig. 1 (left).

The use of the tradespace in multi-stakeholder negotiation rests on the fact that new sets of stakeholder preferences, (new limits, weights, and utilities for each attribute) can be quickly created and calculated “real time” for even a large number of designs, assuming the existence of a database with a reasonable selection of designs and an appropriate set of calculated attributes. The interactive process of stakeholders viewing the impact of their preferences on their own design alternatives, as well as on other stakeholders’ alternatives provides powerful insight into how discovery of “better” solutions can occur. The feedback of changing preferences is shown in Fig. 1 (right).

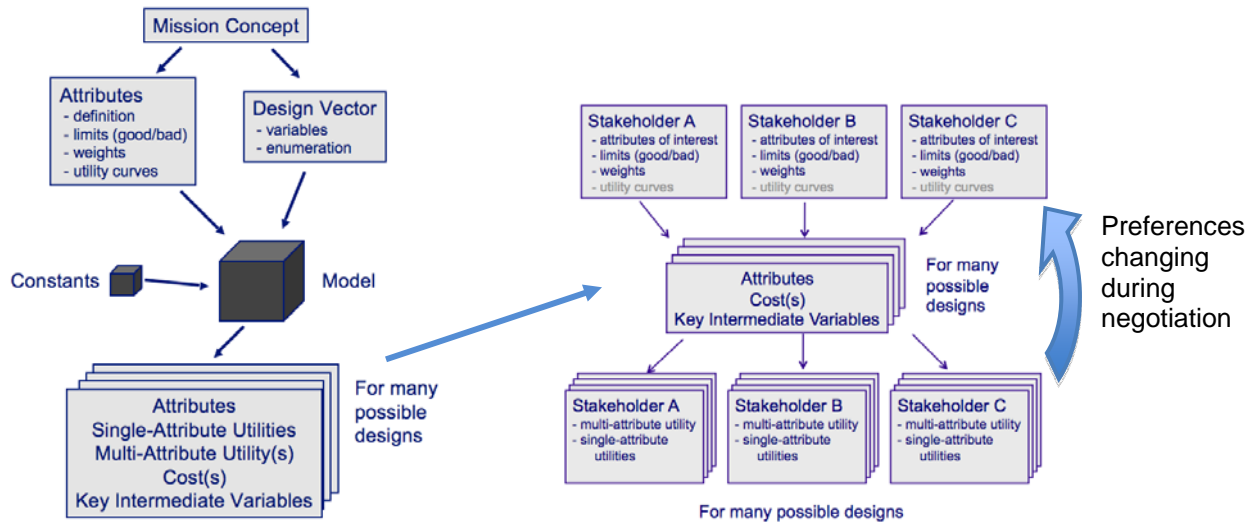


Figure 1. Calculation of tradespace (for one or more initial stakeholders (left)); Use of tradespace to quickly assess designs for multiple stakeholders with changing preferences (right)

II. Method and Results

The proposed method for a multi-stakeholder tradespace exploration includes the following steps:

- A. Each stakeholder individually finds a few “good” designs. Typically this involves examining a utility/cost tradeoff, selecting several efficient designs, and examining and comparing their physical characteristics.
- B. The stakeholders share their “favorite” designs, and evaluate the other stakeholders’ favored designs using their own criteria. Negotiations and compromises commence, driven by the information gained from the tradespace, including both computationally-derived multi-dimensional Pareto sets, as well as human-expressed alterations of stated preferences. The human negotiation has the advantage of both instilling confidence in the stakeholders and making explicit what is being traded; the tools often find solutions that may never be found through “simple” trades. In either case, the “compromise” designs can be instantly displayed in the tradespace of each stakeholder, so that it is clear how much is being gained or lost in the negotiations.

- C. Advanced compromises are developed through interactive relaxation of constraints and proposals of altered preferences or design alternatives.
- D. Once acceptable compromise designs have been discovered, stakeholders jointly examine the tradespace near acceptable compromise designs to see if “better” compromises can be found. Note that in this case, “better” is defined through negotiation rather than in any mathematical sense.
- E. Finally, design(s) that are acceptable to all stakeholders may be explored in more depth for robustness, looking at sensitivities and time-based context effects (e.g., using Epoch-Era analysis^{3,20}).

For the remainder of this section, the method is applied in two case studies, X-TOS and SRS. The models were basic physics, textbook equation type models, without operational detail, and were intended to generate low fidelity insights for complex problems, not to provide accurate design details for the proposed solution systems. For each case study, three “decision makers” role-played a session devoted to finding acceptable “compromise” designs, faithfully representing their own preferences for a system solution. In the SRS case, the decision makers* specified their preferences (attributes, acceptability ranges, and weights) in advance, but were then free to alter their preferences during the session. In the X-TOS case, the decision makers specified their preferences at the start of the session and likewise were able to alter these preferences during their session. For each of the cases, the three decision makers, along with a facilitator, participated in the session using a visualization software-supported venue. Each decision maker was given “control” over their own interactive tradespaces, with the ability to specify and alter their preferences, as well as explore design alternatives. The venue allowed for the simultaneous observation of each decision maker’s distinct preferences, as well as preferred design alternatives. The facilitator guided the participants through the method described above, tracked the group “compromise” decisions, and served as a third-party mediator for negotiations among the decision makers.

A. Each Stakeholder Finds “Good Value” Designs

Each stakeholder needs to quantify their needs through the selection of attributes, and the specification of worst acceptable and best desired values, utility curves, and weights for each attribute. Ideally, all stakeholder would have had input into the definitions of the attributes used to create the tradespace model and datasets. The problem addressed in this method is that each stakeholder may have different attributes, each with different weights (e.g., levels of importance), and different worst and best values and utility curves.

After the preference sets are specified, each stakeholder needs to use the methods of basic tradespace exploration¹⁹ to find a set of favored designs. In addition to forming the starting point for negotiations, this step allows the stakeholders to understand the nature of the tradespace, the characteristics of their favored solutions, and some of the trades that they might have to make, even acting alone without the complication of considering the other stakeholders.

1. X-TOS Example

Table 1 below shows the best and worst acceptable values, and weights, specified for the five attributes of the X-TOS system by each of the three stakeholders. The stakeholders represented scientists interested in different aspects of the X-TOS mission, which collects data from the neutral atmosphere using a variety of sensors. Stakeholder “G” is interested in disruptive “storms” in the atmosphere, characterized by data taken at low altitude, near the equator, over relatively short periods of time (months). Stakeholder “H” is interested in the slower evolution of the atmosphere over the middle latitudes. Stakeholder “F” is interested in the evolution of the atmosphere over the course of the 11-year solar cycle, and so would like to observe a non-trivial fraction of this cycle. These preferences are compared below. The values in bold are the ones where the stakeholder differs from the original stakeholder used to create the original tradespace.²¹

* A decision maker is a type of stakeholder with significant control or influence over defining needs and/or allocation of resources for a system. For the purposes of this paper, since multi-stakeholder negotiation is intended to find compromises affecting the driving needs and allocated resources, the terms “decision makers” and “stakeholders” are used interchangeably.

Table 1. Attribute limits and weights for X-TOS stakeholders

Attribute	"G" – low alt "storms"			"H" – mid-latitude events			"F" – solar cycle effects		
	Worst	Best	Weight	Worst	Best	Weight	Worst	Best	Weight
Lowest Data Altitude (km)	230	150	0.4	1000	150	0.2	1000	150	0.425
Data Life Span (years)	1	11	0.1	1	11	0.35	5.7	11	0.5
Time near Equator (hrs/day)	7	24	0.3	0	7	0.05	0	24	0.175
Latitude Diversity (deg)	0	180	0.05	120	180	0.4	0	180	0.125
Data latency (hrs)	120	1	0.1	120	1	0.1	120	1	0.1

The tradespace representations to follow will typically illustrate the design alternative tradeoffs in terms of multi-attribute utility (y-axis, indicating benefit to a stakeholder) versus lifecycle cost (x-axis, indicating costs incurred for system). In the top right corner is the number of feasible designs for that tradespace (N, which is out of a possible 9930 for the X-TOS original database).

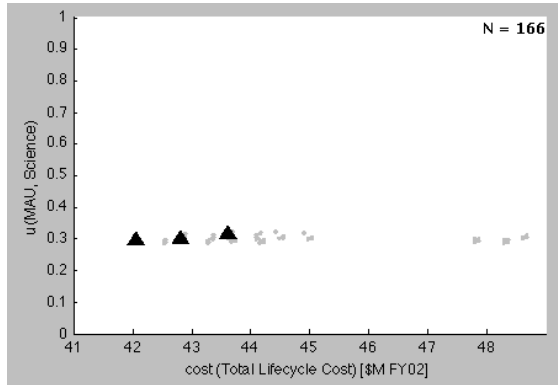


Figure 2. Initial tradespace for stakeholder G




Figure 3. Initial tradespace for stakeholder H

Stakeholder G finds that the restrictions he placed on the worst acceptable values for the attributes severely restricts his tradespace. Note the very low number of acceptable designs in Fig 2. Stakeholder H places only one meaningful restriction on the tradespace, the requirement for at least 120 degrees of latitude diversity. This eliminates about one-third of the original tradespace, but still leaves plenty of designs for consideration (Fig. 3). The stakeholder H picks two favorites, which differ only by amount of on-board delta V. Stakeholder F demands a longer lifetime mission (at least 5.7 years). This has the effect of eliminating about two thirds of the tradespace, including all of the cheaper designs. This stakeholder selects three designs for further examination (Fig. 4).

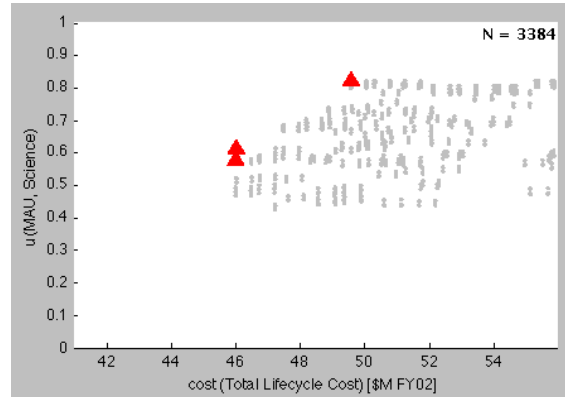


Figure 4. Initial tradespace for stakeholder F

The selected "favorite" designs for the three stakeholders are show in Table 2. The first three are stakeholder G's favorites, the next two stakeholder H's, and the third set stakeholder F's. Stakeholder G has selected designs similar to those selected by the original stakeholder. Stakeholder H's designs are distinguished by the polar (90 deg. inclination) orbits to achieve the coverage he wants. Stakeholder F's designs are longer lifetime; this necessitates higher perigees (less drag) and larger delta V (i.e., more on-board fuel).

Table 2. Favored initial designs for the three X-TOS stakeholders

Design ID	Design Variables					Key attributes				"G"	"H"	"F"	Cost M\$
	Inclin. (deg)	Perogee (km)	Apogee (km)	Delta V (m/s)	Power Type	Lifespan (yrs)	Sample Alt (km)	Lat. (deg)	Dev T @ eq (hr/dy)	MAU	MAU	MAU	
4576	30	200	350	1000	1	1.26	200	60	11.5	0.30	X	X	42.05
2689	30	200	950	650	0	1.44	200	60	11.5	0.30	X	X	42.81
4775	30	200	650	1000	1	2.17	200	60	11.5	0.31	X	X	43.60
2577	90	200	950	700	0	1.43	200	180	5.3	X	0.70	X	42.07
3936	90	200	950	900	0	2.17	200	180	5.3	X	0.74	X	42.89
3886	30	250	950	900	0	6.47	250	60	11.6	X	X	0.61	45.98
4585	30	250	650	1000	0	6.20	250	60	11.5	X	X	0.58	42.81
3209	30	300	950	800	0	11.00	300	60	11.6	X	X	0.82	49.54

2. *SRS Example*

In this case, three stakeholders picked three to five attributes each, which they use to assess potential system alternatives for performing a radar-based imaging and/or tracking mission, defined worst acceptable and best desired values for each attribute, and assigned weights.* Stakeholder “E” was interested in observing and monitoring large areas in a timely fashion (e.g., for maritime surveillance or strategic verification). Attributes for stakeholder E in Table 3 include revisit interval, targets per pass, and data latency. Stakeholder “N” was interested in the spotting and tracking of small objects of interest, as well as the rapid reporting of data (e.g., for tactical observations). This stakeholder had the largest number of attributes showing in Table 4, including the technically challenging ability to define the global location of the spotted objects (i.e., geo-location accuracy). Stakeholder “R” was interested in wide area surveillance for specific, smaller objects, shown in Table 5.

Table 3. SRS attributes for stakeholder “E”

Attribute Definition	Worst	Best	Weight
Revisit Interval (min) = average time between observations on the same target	500	30	0.50
Targets Per Pass = number of targets of a given size that can be detected per pass	1	100000	0.35
Image Latency (min) = time from target view to user receipt of analyzed image data	720	60	0.15

Table 4. SRS attributes for stakeholder “N”

Attribute Definition	Worst	Best	Weight
Geo-Location Accuracy (m) = error in absolute position of target on the ground. Assumes no references	100	3	0.50
Target Acquisition Time (min) = 95% worst case time to acquire target	120	10	0.30
Tracking Latency (min) = time from target view to user receipt of data	10	1	0.30
Minimum RCS (m2) = Theoretical minimum RCS necessary for target to be distinguished	1000	0.01	0.10
Minimum Detect. Velocity (m/s) = Theoretical minimum velocity necessary for moving target to be detected	50	5	0.10

* The “weights” are not importance weights per se, but rather represent the “level of satisfaction” on a 0 to 1 scale, when that attribute is at its best value and all of the other attributes are at their worst level. This type of weight is called a “swing weight” and requires the simultaneous consideration of all attributes for a given stakeholder. It allows for more nuanced understanding of benefits, including complementarities and substitution effects among the attributes. As such, a multiplicative multi-attribute utility function is used, where the weights do not have to add up to 1, but must fall between 0 and 1. In the special case where the weights add up to 1, the function degenerates into a linear weighted sum and the attributes can be considered as independently contributing to the overall utility.

Table 5. SRS attributes for stakeholder “R”

Attribute Definition	Worst	Best	Weight
Image Resolution (m) = theoretical minimum separation to distinguish different targets.	5	0.01	0.45
Number of Target Boxes = number of (200km) ² boxes scanned at required resolution in one pass	1	10	0.35
Revisit Interval (min) = average time between observations on the same target	500	50	0.20

The SRS stakeholders examined their tradespaces, and picked favored designs. In addition to manually picked favorites, the Pareto set was calculated for each stakeholder, containing non-dominated designs that had the highest utility for given cost, or the lowest cost for a given utility. Both sets of designs for each stakeholder are shown in Fig. 5, Fig. 6, and Fig. 7 below. The favorite designs for each stakeholder are listed in Table 6 below. For these tradespaces, a total of 23,328 design alternatives were evaluated. Stakeholder E had a large tradespace of acceptable designs -- very few designs were eliminated from the tradespace by his attribute requirements. He favored high altitude vehicles with high bandwidth. He picked only a few representative favorites; the Pareto front had many more designs. Stakeholder N, by contrast, had attribute limits that eliminated most of the designs from the tradespace. Only a small number of favored designs delivered utility at a “reasonable” price; the calculated Pareto set added only a few more designs. Note that low cost designs (less than \$10B) are absent from this tradespace. Stakeholder R also has a large tradespace, but with a sparse Pareto front. His “favorite” designs actually include several non-Pareto designs; this stakeholder was curious and wished to explore designs near but not on the front. Stakeholder R’s favored designs are high altitude, with larger antennas, higher power antennas, but with smaller numbers of vehicles.

Table 6. Initial favorite designs of SRS stakeholders

Stakeholder	Design ID	Design Variables							"E" MAU	"N" MAU	"R" MAU	Cost B\$
		Altitude (km)	Num. Sats	Ant. Area (m ²)	Bandwidth (GHz)	Power (kw)	Tac Comm					
E	3328	1500	10	10	2	10	no	0.72	X	0.62	9.6	
	3436	1500	10	40	2	10	no	0.81	X	0.92	14.2	
	3760	1500	20	40	2	10	no	0.86	X	0.92	20.8	
N	471	800	10	40	1	1.5	yes	0.66	0.53	0.58	13.3	
	3387	1500	10	40	1	1.5	yes	0.75	0.47	0.58	12.8	
	6303	1200	20	40	1	1.5	yes	0.81	0.57	0.58	18.4	
R	*424	800	10	10	2	20	no	0.56	X	0.78	12.0	
	2692	1500	5	10	2	20	no	0.59	X	0.66	9.1	
	*2752	1500	5	40	1	10	no	0.60	X	0.80	10.8	
	2788	1500	5	40	2	10	no	0.60	X	0.85	10.8	
	3436	1500	10	40	2	10	no	0.81	X	0.92	14.2	




Figure 5. Initial tradespace for stakeholder E, picked favorites and Pareto set




Figure 6. Initial tradespace for stakeholder N, picked favorites and Pareto set

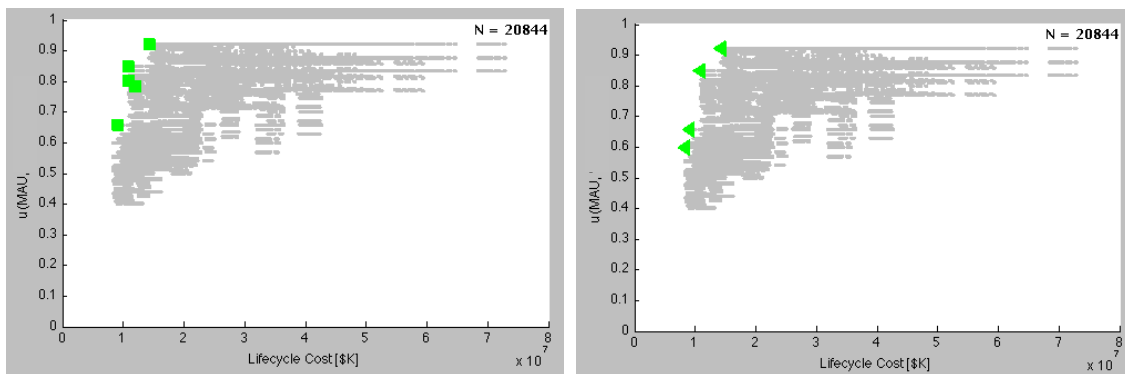


Figure 7. Initial tradespace for stakeholder R, picked favorites and Pareto set

B. Stakeholders Compare Designs and Make “Easy” Compromises

Once the stakeholders have expressed their needs and have found physically realizable designs in a common tradespace, they are ready to begin productive negotiations on selecting a group agreed upon design. Both mathematical techniques and humans-in-the-loop negotiations can be used. Quick compromises fall into a number of categories:

- For some multi-stakeholder problems, there may solutions that satisfy everyone. If one or more designs exist on everyone’s “favorites” list, they are clear winners, easy for everyone to accept. (These solutions that do not require compromise are considered “joint” solutions.)
- Expanding the search to include the entire MAU-versus-cost Pareto front of each stakeholder may find solutions that are on each stakeholder’s Pareto front, even if not on their favored list. These solutions are also often “easy” compromises. (These solutions that mathematically do not require compromise are considered “joint best” solutions.)

- Relaxing the requirement that the solutions be on the Pareto front (i.e., accepting a “fuzzy” Pareto set) will expand the solution space while still requiring stakeholders to make only small compromises.
- Examining the multi-dimensional Pareto surface of cost and all stakeholders’ MAUs may find optimal “compromise” solutions. These solutions are not on the individual stakeholders’ Pareto fronts, but are on a more difficult to visualize higher dimensional surface of optimal tradeoffs between stakeholders’ MAUs and cost. These extra mathematically generated compromises are the most efficient MAU-cost tradeoff solutions that are more difficult to determine through human-in-the-loop approaches since it requires making the “right” trades to give up utility from one stakeholder in order to most efficiently give utility to a different stakeholder.

In all cases (except, usually, the last) the solutions can be found either by mathematical tools, or by humans-in-the-loop trading of either attributes or design features. The favorite lists are short enough to be compared directly. The full Pareto fronts can be compared with tools, but the stakeholders can often find joint Pareto designs by negotiation and joint exploration. Once joint exploration begins, the “fuzzy” Pareto front is often explored intuitively through stakeholders making small compromises on design features that other stakeholders want, but which push the design slightly off their own Pareto front. Often, these small compromises are in the compromise solution multi-dimensional Pareto set, but it is easier for the stakeholders to understand why they work if they arrive at them by negotiation.

A common problem with “easy” negotiations is the so-called “gold-plated” system. Solutions that simply include all the features the stakeholders want will often appear on everyone’s Pareto front, but at the high end of the cost range. Both negotiation (through adding of design features) and mathematical techniques (the highest-cost, highest-utility designs will tend to appear on or near everyone’s Pareto front) can yield these proposed solutions.

1. X-TOS Example

The X-TOS system has no “easy” compromises that all three stakeholders can agree upon. Each stakeholders’ favorites are not only not favored by the other stakeholders, but in fact do not appear in their tradespaces at all, having been disqualified by the restrictions the stakeholders placed on worst acceptable value of certain attributes. This is indicated in Table 2 by the “X”s under the other stakeholders’ MAUs. Other stakeholders’ favorites fail to satisfy stakeholder G either due to not enough equator time or not a low enough perigee; stakeholder H cannot accept the others’ favorites lack of latitude diversity, and stakeholder F requires longer lifetime missions that the others have dismissed due to cost. Evidently this group is not starting in a good place for negotiating a joint mission.

In fact, stakeholder G and F share no designs at all due to the conflicting demands of a low altitude and a long lifetime. Some easy compromise is possible between stakeholders H and F. Stakeholder F does not care about latitude one way or another, so it is an “easy” compromise for stakeholder F to look at high inclination missions; the compromise H needs to make is that these missions will be higher cost.

2. SRS Example

In contrast with the X-TOS example, multiple “easy compromises” were necessary to get the SRS stakeholders aligned. Entering this step of the negotiation process, design number 3436 was on both stakeholder E and stakeholder R’s favorites lists. On the other hand, stakeholder N excluded all of the other stakeholders’ favored designs from his tradespace. Comparing the designs directly, it appears that N requires a tactical communication system; the other stakeholders desire higher power and bandwidth.

The first step beyond simple comparison of favorites is to compare the entire Pareto fronts. Stakeholder E has the most agreeable tradespace, with several of the other stakeholders’ Pareto designs on or near his Pareto front. Stakeholder N, on the other hand, rejects all of the other stakeholders’ Pareto designs. Stakeholder R does not eliminate the other stakeholders’ Pareto sets from his tradespace, but finds them to be non-Pareto optimal.

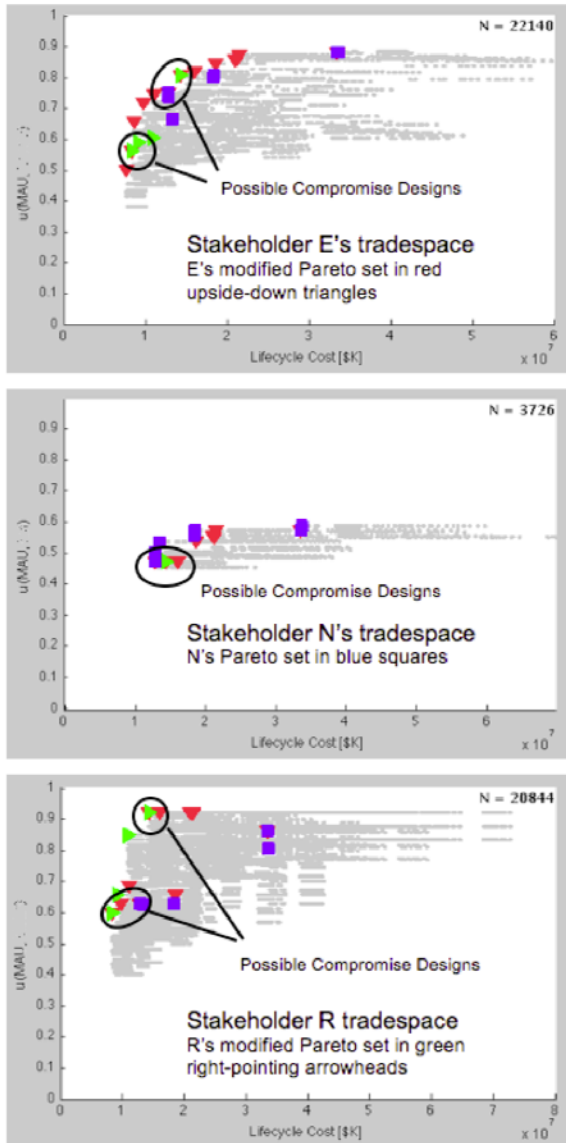


Figure 8. Modified Pareto sets include mutually acceptable solutions

designs are acceptable to all stakeholders. They are remarkably similar, differing only in a few particulars, despite having been collected by differing techniques. The last two designs are lower cost designs that were identified on Figure 8 but are not acceptable to stakeholder N; these designs will be reviewed in the next subsection.

Two approaches were taken to find mutually acceptable solutions. First, design features were negotiated. On examining the favorite solutions and the larger Pareto sets, it was found that stakeholder N required the tactical communication system in all acceptable designs, and stakeholder R had a strong preference for high bandwidth systems. These design features had minimal effects on the other stakeholder's utilities, but (compared with the very high cost of SRS systems) could be included relatively cheaply. With the stakeholder's consent, the Pareto sets of all stakeholders were modified so that high bandwidth radar and tactical communication systems were included in all designs. The comparison between stakeholders was repeated. Figure 8 shows each stakeholder's tradespace with all stakeholders' modified Pareto optimal designs on them. Stakeholder E's Pareto set is shown with red downward-pointing triangles; stakeholder N with blue squares, and stakeholder R with green arrowheads (sideways-pointing triangles). Many stakeholders' Pareto set designs now appear on, or at least very near, the Pareto front of other stakeholders, indicating potential mutually acceptable designs.

The other approach to identify potential compromises was to perform an automated search on the multi-dimensional space of cost and the three stakeholders' MAUs to find a Pareto surface. This technique found large numbers of designs that are reasonable compromises in the mathematical sense. Interestingly, many of the designs found by the negotiation process described above were on this surface. Two sets of designs were selected from this search: some "gold-plated" designs that performed well for all stakeholders at a price, and some "compromise" solutions that were not favorites of any individual stakeholder but were on the mutual Pareto front, were not too far from the individual stakeholders Pareto front, and were reasonably priced.

The compromise designs collected by these methods are shown in Table 7 below. Designs that are on or very near a stakeholder's Pareto front are indicated with a bold MAU. A non-bold MAU indicates an acceptable but, for that stakeholder, dominated design. An "X" indicates the design is not acceptable to that stakeholder. Five of the

Table 7. Designs that satisfy multiple stakeholders

Stakeholder	Design ID	Design Variables						"E"	"N"	"R"	Cost B\$
		Altitude (km)	Num. Sats	Ant. Area (m ²)	Bandwidth (GHz)	Power (kw)	Tac Comm	MAU	MAU	MAU	
All	3423	1500	10	40	2	1.5	yes	0.75	0.47	0.63	12.8
	3435	1500	10	40	2	10	yes	0.81	0.47	0.92	14.3
	*6027	1200	10	40	2	10	yes	0.80	0.50	0.92	14.6
	*519	800	10	40	2	10	yes	0.68	0.53	0.88	15.8
	**6349	1500	20	40	2	10	yes	0.87	0.57	0.92	22.3
E+R	2679	1500	5	10	2	10	yes	0.56	X	0.60	8.3
	2787	1500	5	40	2	10	yes	0.60	X	0.85	10.9

no star – arrived at by design feature compromise
 * selected from analytically identified multi-stakeholder Pareto set
 ** gold-plate design, Pareto for everyone at high cost

C. Negotiations and Relaxation of Attribute Requirements

If, as is often the case, there are no easy compromises that can satisfy all stakeholders at a reasonable price, more detailed negotiations need to commence based on the idea of relaxing the individual stakeholder’s attribute limits. The basic tradespace procedure is straightforward, although the humans-in-the-loop negotiations may not be.

The overall yield of each stakeholder’s tradespace can be used to identify difficult to satisfy stakeholders who are eliminating much of the tradespace due to strict attribute limits. These limits may be justified, but if they prevent multi-stakeholder agreement they should be examined. If the stakeholders can be persuaded to relax them, the tradespace will include more designs. If the expanded tradespace yields new designs for consideration, they can be examined and compared by each stakeholder on tradespace plots and by tabulating design vector elements and attributes. This forms a common basis for negotiation, with explicit acknowledgement of what each stakeholder is getting (e.g., attributes and overall utility) and giving up (e.g., relaxed limits, cost, lower attributes than “favored” designs).

1. X-TOS Example

The yield for stakeholder G indicates that this stakeholder has eliminated from consideration most of the possible designs due to strict limits on the attributes he will accept. Figures 9 and 10 are plots of individual attributes versus cost, with all designs shown. In these figures, the lowest sample altitude and time near the equator attributes are shown, with a red line indicating the worst acceptable limit for stakeholder G. Designs on the wrong side of this line are insufficient in the attribute and do not appear on stakeholder G’s tradespace. Color is used to highlight what is eliminated by each restriction. The sample altitude restriction eliminates all longer-lifetime missions (which are key to stakeholder F) and the time near the equator restriction eliminates all higher-inclination missions (key to stakeholder H). If this stakeholder wishes to participate in a joint mission, these restrictions will need to be relaxed.

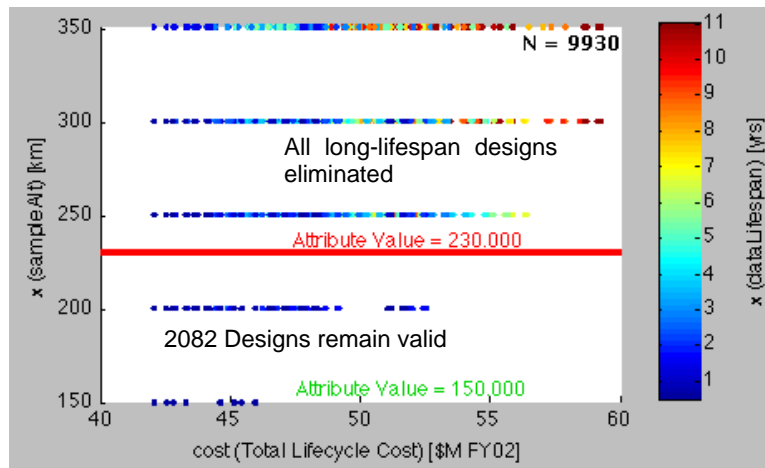


Figure 9. Stakeholder G limit on lowest sample altitude eliminates longer lifespan designs

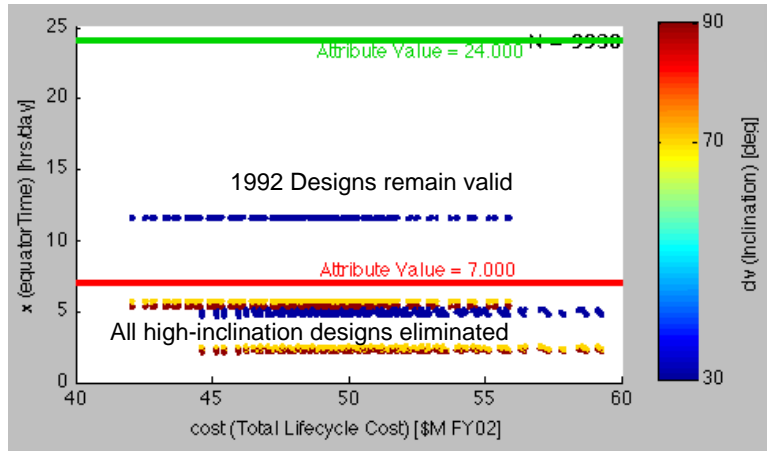


Figure 10. Stakeholder G limit on equator time eliminates all higher-inclination designs

Fortunately, it can be observed directly from the figures that even a modest relaxation of the restrictions will allow much more of the design space to be included. A relaxation from 230 to 270 km in data altitude, and from 7 to 4 hrs/day in equator time was negotiated. Stakeholders H and F correctly stuck to their less restrictive limits, as short lifetime systems are useless to stakeholder F, as are low-latitude missions to stakeholder H.

Using the modified attributes, the stakeholders' tradespaces were re-calculated and compromise designs found. In this case, the process was essentially manual, based on starting with existing favored designs and modifying them to satisfy the other stakeholders. For example, an inclination of 70 degrees was not included on anyone's favorite list, but was a balance between G's need for time near the equator and H's need for higher-latitude data. The other major trade was between low altitude and the need for lifetime; this was explored by lowering the perigee of F's favored designs until the lifetime became too short. It was found that a perigee of 250 km was possible. This led quickly to a single compromise design.

2. SRS Example

There are solutions acceptable to all SRS stakeholders in the design space. However, SRS is a very expensive system, and none of the lower-cost solutions are included in the mutually acceptable designs. We can see from Fig. 8 that stakeholder N is excluding a large portion of the tradespace, including lower cost designs acceptable to the other stakeholders. Examining the attributes of stakeholder N shows that the worst acceptable value of three attributes, time to acquire target, tracking latency, and minimum detectable velocity, are excluding large numbers of designs. The tracking latency limit requires a tactical communication system an issue that has already been dealt with through negotiations with other stakeholders. The other two attribute limits are excluding low cost designs, as can be seen in Fig. 11. In this figure, blue 5-point stars are selections from the multi-stakeholder Pareto set, the yellow diamond is the gold-plated design, the negotiated compromises based on stakeholder favorites are shown as black stars, and the lighter (pink) stars are the proposed lower-cost designs.

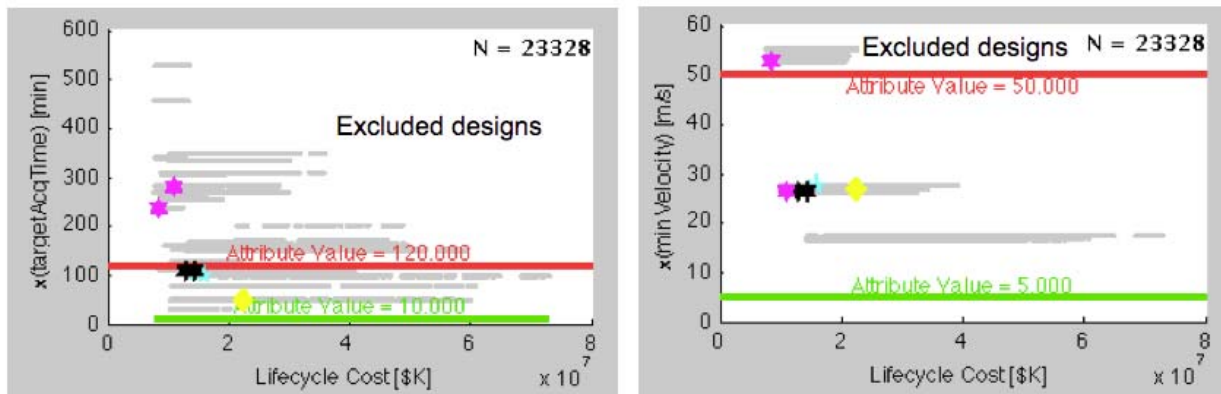


Figure 11. Stakeholder N attribute limits exclude lower cost designs

Relaxing these limits for stakeholder N is not a trivial matter, however. The target acquisition time attribute in particular is important to this stakeholder. However, as a “what-if” exercise, the stakeholder was willing to investigate the effects of relaxing these attributes. The results are shown in Fig. 12. Relaxing the target acquisition time attribute limit (to 300 min) adds a limited number of lower-utility, lower-cost solutions to the tradespace, including one of the designs favored by the other stakeholders. Relaxing the minimum velocity attribute (to 60 m/s) on its own is not effective, but relaxing both limits adds a large number of designs, some of them quite attractive to this stakeholder. With both limits relaxed, the utilities of all proposed compromise designs are recalculated. The results can be seen in Table 8.

Table 8. Candidate designs after stakeholder N attribute modification

Symbol	Design ID	Design Variables						Tac Comm	"E" MAU	"N" MAU	"R" MAU	Cost B\$
		Altitude (km)	Num. Sats	Ant. Area (m ²)	Bandwidth (GHz)	Power (kw)						
Black Star	3423 3435	1500 1500	10 10	40 40	2 2	1.5 10	yes yes	0.75 0.81	0.56 0.56	0.63 0.92	12.8 14.3	
Blue 5pt star	6027 519	1200 800	10 10	40 40	2 2	10 10	yes yes	0.80 0.68	0.58 0.60	0.92 0.88	14.6 15.8	
Diamond	6349	1500	20	40	2	10	yes	0.87	0.62	0.92	22.3	
Pink Star	2679 2787	1500 1500	5 5	10 40	2 2	10 10	yes yes	0.56 0.60	0.49 0.46	0.60 0.85	8.3 10.9	

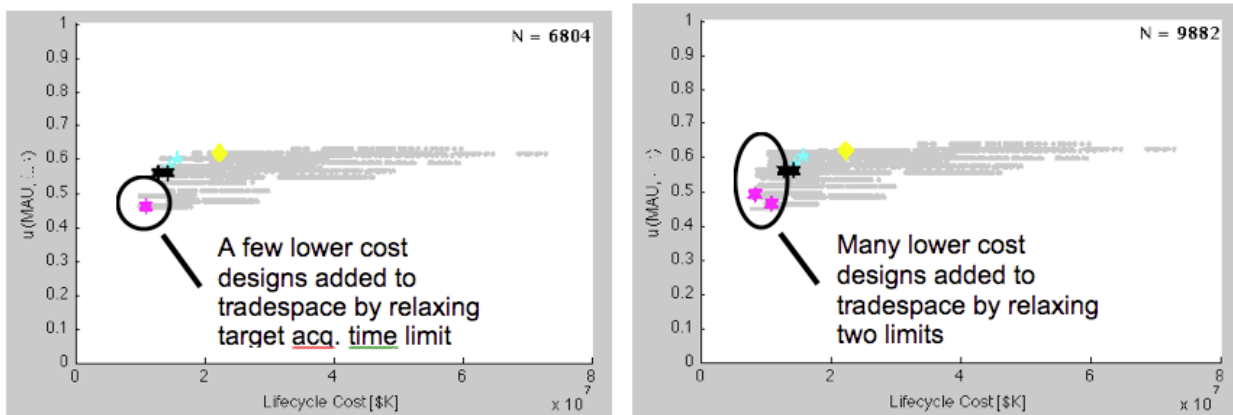


Figure 12. Relaxing attribute limits expands stakeholder N’s tradespace

D. Finding Better Compromises

Once a favorite design has been selected, the question remains – can better compromise designs be found in the multi-stakeholder tradespace? The two examples given here both involve attempts to improve the solution for the stakeholders that have had to make the most compromises. This is NOT a question of mathematical optimization, but is much more an exercise in getting confidence that there are no better solutions for the affected stakeholders, and generating buy-in from those stakeholders. The basic method is to see if perturbations in design variables (choices) can “tweak” the design in a way that helps assure stakeholder buy-in. These permutations may involve improvements for all stakeholders (better designs in an absolute sense), or they may require tradeoffs (improvements for some stakeholders at the expense of others). In the illustrated examples in this paper, the “tweaks” can only be as fine as the enumeration of the design variables. The issue of locally recalculating the tradespace around good designs is relevant to this question, but not covered here.

1. X-TOS Example

The stakeholders, having found a good design, examined the sensitivities on all design variables by direct examination. It was found that the selected compromise was quite optimal, given the existing enumeration of the design space. Changes in inclination caused the latitude diversity or the time at the equator to become unacceptable to some stakeholders. Changes in perigee altitude were similarly precluded due to needs for lifetime and low altitude data. Likewise, changes in apogee altitude changed either lifespan or equator time in unacceptable ways. One change was found that produced a slightly better compromise: reducing the delta V of the system produced a slight, but acceptable reduction of the lifetime of the system, and lowered its cost somewhat. This could be perceived as beneficial to stakeholders G and H, who were having to settle for a more expensive system to keep

Table 9. Final design set for X-TOS

Design ID	Design Variables					Key attributes				"G" MAU	"H" MAU	"F" MAU	Cost M\$
	Inclin. (deg)	Perogee (km)	Apogee (km)	Delta V (m/s)	Power Type	Lifespan (yrs)	Sample Alt (km)	Lat. Dev (deg)	T @ eq (hr/dy)				
4576	30	200	350	1000	1	1.26	200	60	11.5	0.43	X	X	42.05
2689	30	200	950	650	0	1.44	200	60	11.5	0.43	X	X	42.81
4775	30	200	650	1000	1	2.17	200	60	11.5	0.45	X	X	43.60
2577	90	200	950	700	0	1.43	200	180	5.3	0.34	0.70	X	42.07
3936	90	200	950	900	0	2.17	200	180	5.3	0.35	0.74	X	42.89
3886	30	250	950	900	0	6.47	250	60	11.6	0.34	X	0.61	45.98
4585	30	250	650	1000	0	6.20	250	60	11.5	0.34	X	0.58	42.81
3209	30	300	950	800	0	11.00	300	60	11.6	X	X	0.82	49.54
4619	70	250	950	1000	0	7.67	250	140	5.74	0.25	0.74	0.66	47.45
3915	70	250	950	900	0	6.47	250	140	5.73	0.25	0.72	0.60	45.98

stakeholder F involved. Stakeholder F had to settle for a system a little further off his Pareto front, but still very close. Table 9 and Fig. 13 show both the original compromise design, and the improved compromise.

2. *SRS Example*

These techniques applied to the SRS tradespace do not result in further candidate designs. The primary reason can be seen by examining Table 7. The candidate designs are all quite similar, differentiated by the variation of only a few design variables. These design variations tend to please one stakeholder at the expense of reducing a different stakeholder’s utility or increasing cost (i.e., these are Pareto trade-offs). The SRS design space is very coarsely enumerated, so these design variations are large. It is likely that a finer enumeration of the design space in the vicinity of the candidate designs would reveal better compromises between the different stakeholders’ utilities and cost.

E. Sensitivity to Time and Change

Epoch-Era Analysis^{3,22} was used to analyze the effects of a changing environment and stakeholder needs on the compromise solutions in the SRS tradespace. An epoch is a time period of fixed needs, constraints, and design alternatives. A time ordered series of epochs, or era, is a long run scenario that includes changes in use, technology, and supporting infrastructure.²³ In these analyzed eras, the first epoch is considered the “baseline” context. In the next epoch, a supporting communications infrastructure becomes available. The next epoch sees a change in use such that target objects get smaller. The fourth epoch includes more supporting infrastructure; the last epoch is a return to the old target objects. This era was explored by the three stakeholders. The results are shown in Fig. 14, with the rows showing the evolution of the tradespaces across the era (from top to bottom row) for stakeholder E, N, and R respectively. The 3rd and 4th epochs involve the need to spot smaller objects and have the most effect on the tradespace. Stakeholder E sees the relative worth of the more expensive “gold plate” design increase, and some even more expensive designs appear attractive (have higher utility) in these epochs. Stakeholder N is relatively

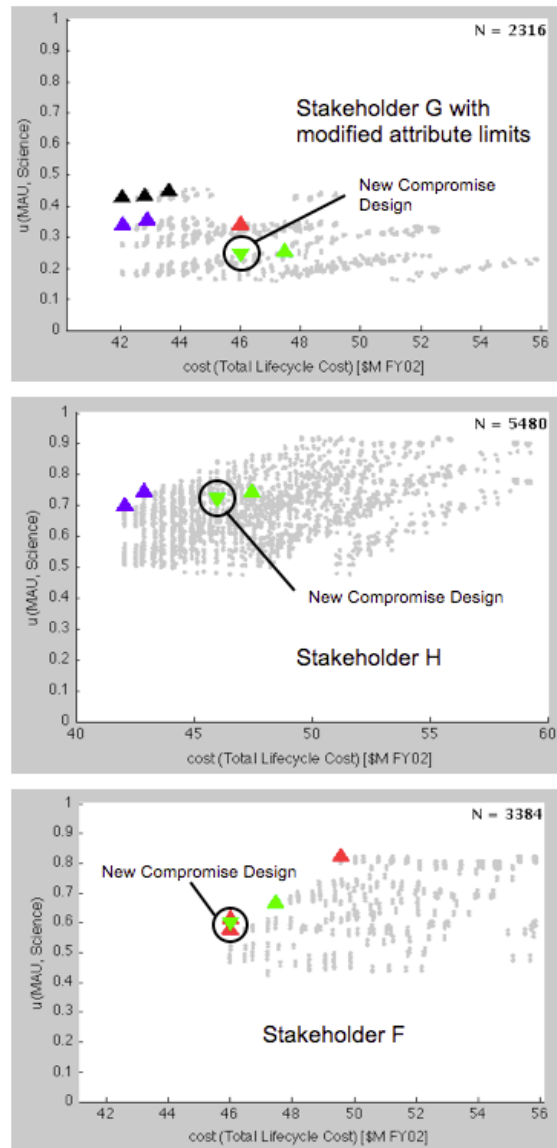


Figure 13. X-TOS stakeholder tradespaces showing compromise designs

unconcerned with the changes. Stakeholder R sees the biggest changes, with even the “gold plate” design becoming dominated by large numbers of other designs in the 3rd and 4th epochs.

A quick exploration of which designs might be more attractive to stakeholders E and R in the third epoch reveals difficulty for low cost compromises: stakeholder E wants many satellites in the system, while stakeholder R favors large, powerful antennas on each satellite. These desires are expensive and preclude lower cost compromises. Further negotiations are needed to determine if the possibility of this future epoch justifies high costs in order to have higher performing systems during that time period, or if the existing compromise solutions are adequate. Some of the existing compromise solutions remain near-Pareto for all stakeholders in all epochs.

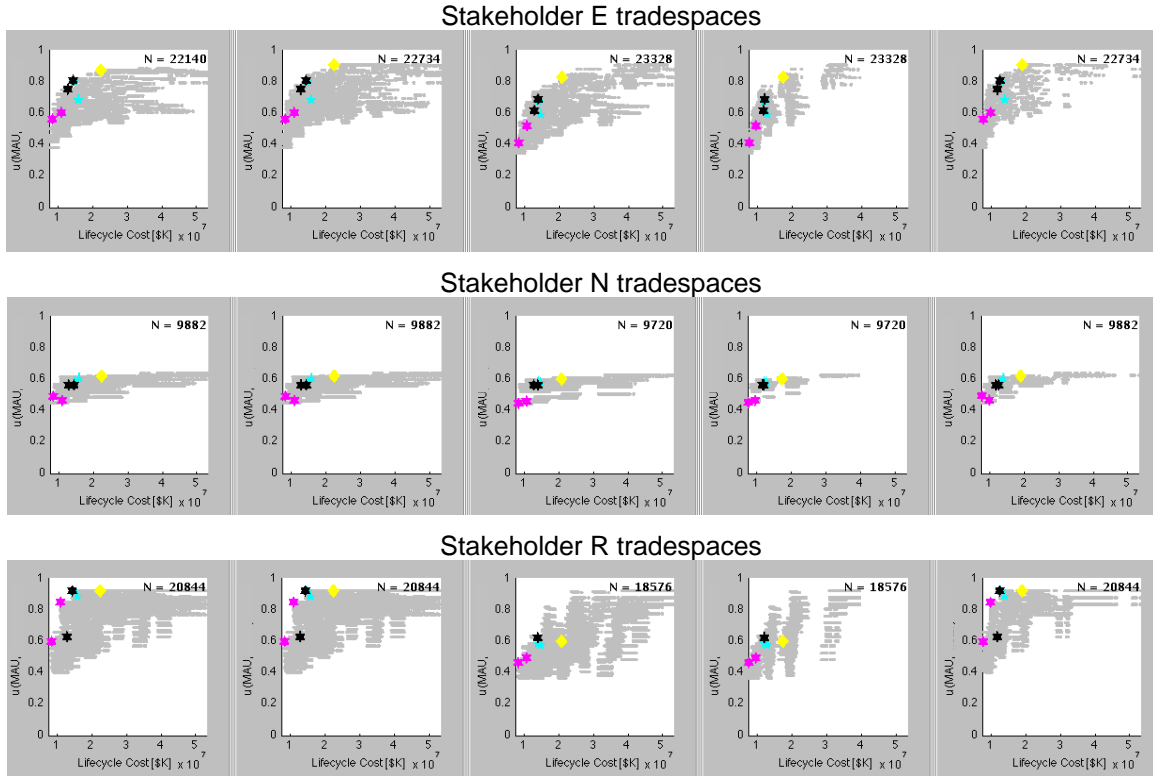


Figure 14. Evolution of three stakeholder tradespaces through an era (top to bottom: E, N, R)

III. Discussion

The structured approach described in this paper combines traditional negotiation techniques with computationally-aided, interactive tradespace exploration in a unique and promising way, illustrating its applicability for rapid, transparent consideration of conflicting stakeholder needs, for identifying win-win compromise solutions, as well for understanding the technical and cost implications for such compromises. The approach also has a number of intrinsic advantages:

- The use of interactive tradespace exploration analyses and representation techniques allow negotiation processes to proceed very quickly. Proposed compromises can be assessed by each stakeholder in real time, and what the stakeholder is gaining or losing in the compromise is immediately visible.
- The tradespace representations provide a common “boundary object” on which the stakeholders can focus their discussions.²⁴
- Techniques such as multi-dimension Pareto analyses can help the stakeholders find non-intuitive compromises. Other techniques can also be used to help suggest productive directions for negotiations, such as identifying design variables key to each stakeholder, or attribute limits that are constraining the potential solution set.
- On the other hand, having the negotiations driven by the stakeholders “in person,” rather than “in proxy,” allowing imperfectly expressed stakeholder needs to be corrected in real time; “trivial” compromises can be

made very quickly, and the trading of satisfying stakeholders (one stakeholder's utility for another's) can be done by the stakeholders, instead of resorting to methods with known problems, such as imposed assumptions about the relative importance of stakeholders..

The work presented in this paper is still preliminary in that the use of interactive tradespace exploration for multi-stakeholder negotiations was not the original intended use of interactive tradespace exploration. This usage, instead, emerged as a stakeholder-identified promising approach over the course of maturing tradespace exploration techniques in general.¹⁹ A more targeted research effort that develops tools, representations, and techniques to streamline the steps suggested in this paper, including the ability to readily track impacts of negotiations, such as "lost utility" and "extra cost incurred" due to compromises can help to make explicit fundamental tensions and tradeoffs that inherently exist in multi-attribute, multi-stakeholder design problems.

IV. Conclusion

The approach illustrated in this paper, to use interactive tradespace exploration to assist and accelerate multi-stakeholder negotiations around cooperatively-developed complex systems, is still in its early stages but already shows promise in accelerating negotiations and group knowledge sharing. The approach addresses known problems in both the current state of negotiation in practice, and in theoretical multi-stakeholder decision making. Next steps in maturing the approach should involve the development of more metrics and tools in order to efficiently characterize the tradeoffs performed in negotiation (both in terms of design options and attributes, utilities, and cost), as well as demonstration of various negotiation techniques developed in the economic literature of cooperative games¹. Agent-based approaches could be used to help suggest negotiation pathways, as well as the generation of new alternatives, which could then be re-inserted back into interactive tradespace exploration to reveal superior "win-win" solutions. Preliminary use of the tradespace-assisted negotiations approach on several different problems has provided confidence in the approach's practical usefulness as well as insights into the nature of the multi-stakeholder problem, possibly allowing for the identification of better compromise solutions than may be proposed through more traditional approaches.

Acknowledgments

Funding for this research was provided by the Systems Engineering Advancement Research Initiative (SEArI), a research initiative within the Engineering Systems Division at the Massachusetts Institute of Technology. SEArI (<http://seari.mit.edu>) brings together a set of sponsored research projects and a consortium of systems engineering leaders from industry, government, and academia. SEArI gratefully acknowledges the funding support of the U.S. government for this work.

References

-
- ¹ Lai, G., Li, C., Sycara, K., and Giampapa, J., "Literature Review on Multi-Attribute Negotiations," *Technical Report CMU-RI-TR-04-66*, Carnegie Mellon University Robotics Institute, December 2004, pp. 37.
 - ² Ross, A.M., Hastings, D.E., Warmkessel, J.M., and Diller, N.P., "Multi-Attribute Tradespace Exploration as a Front-End for Effective Space System Design," *AIAA Journal of Spacecraft and Rockets*, Jan/Feb 2004.
 - ³ Ross, A.M., McManus, H.L., Rhodes, D.H., Hastings, D.E., and Long, A.M., "Responsive Systems Comparison Method: Dynamic Insights into Designing a Satellite Radar System," *AIAA Space 2009*, Pasadena, CA, September 2009.
 - ⁴ Sycara, K., "Problem Restructuring in Negotiation," *Management Science*, Vol. 37, No. 10, October 1991, pp. 1248-1268.
 - ⁵ Kersten, G. E., Michalowski, W., Szpakowicz, S. and Koperczak, Z., "Restructurable Representations in Negotiation," *Management Science*, Vol. 37, No. 10, October 1991, pp. 1269-1290.
 - ⁶ Kusiak, A. and Wang, J., "Negotiation in Engineering Design," *Group Decision and Negotiation*, Vol. 3, No. 1, April 1994, pp. 69-91.
 - ⁷ Kalai, E. and Smorodinsky, M., "Other Solutions to the Nash Bargaining Problem," *Econometrica*, Vol. 43, 1975, pp. 513-518.
 - ⁸ Scott, M. and Antonsson, E., "Formalisms for Negotiation in Engineering Design," *ASME Design Engineering Technical Conference (DETC)*, Irvine, CA, August 1996.

-
- ⁹ Lai, G. and Sycara, K., "A Generic Framework for Automated Multi-attribute Negotiation," *Group Decision and Negotiation*, Vol. 18, No. 2, March 2009, pp. 169-187.
- ¹⁰ Wolf, D., Simpson, T. W. and Zhang, X. L., "A Preliminary Study of Novice and Expert Users' Decision-Making Procedures during Visual Trade Space Exploration," *IDETC/CIE 2009*, ASME, San Diego, CA, August-September 2009, pp. 11.
- ¹¹ Lotov, A. V., Bushenkov, V. A. and Kamenev, G. K., *Interactive Decision Maps: Approximation and Visualization of Pareto Frontier*, Applied Optimization, Kluwer Academic Publishers, Boston, 2004, pp. 310.
- ¹² 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.
- ¹³ Stump, G. M., Yukish, M. A. and Martin, J. D., "The ARL Trade Space Visualizer: An Engineering Decision-Making Tool," *10th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, AIAA 2004-4568, AIAA, Albany, NY, 2004, pp. 11.
- ¹⁴ Stump, G., Yukish, M., Simpson, T. W. and O'Hara, J. J., "Trade Space Exploration of Satellite Datasets Using a Design by Shopping Paradigm," *2004 IEEE Aerospace Conference*, Big Sky, MT, March 2004.
- ¹⁵ Wang, J. and J. Terpenney, "Interactive Evolutionary Solution Synthesis in Fuzzy Set-based Preliminary Engineering Design," *Journal of Intelligent Manufacturing*, Vol. 14, 2003, pp. 153-167.
- ¹⁶ Arrow, K. J., *Social Choice and Individual Values*, Yale University Press, New Haven, 1963.
- ¹⁷ Hazelrigg, G. A., "Bad design decisions: Why do we make them?" *New Design Paradigms*, Pasadena, CA. June 2001.
- ¹⁸ Ross, A.M. and Hastings, D.E., "The Tradespace Exploration Paradigm," *INCOSE International Symposium 2005*, Rochester, NY, July 2005.
- ¹⁹ 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.
- ²⁰ 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.
- ²¹ Ross, A.M., Multi-Attribute Tradespace Exploration with Concurrent Design as a Value-centric Framework for Space System Architecture and Design, Dual Master of Science Thesis, Aeronautics and Astronautics and Technology and Policy Program, MIT, June 2003.
- ²² Ross, A.M., "Managing Unarticulated Value: Changeability in Multi-Attribute Tradespace Exploration," Doctor of Philosophy Dissertation, *Engineering Systems Division*, MIT, June 2006.
- ²³ 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, March 2009.
- ²⁴ Carlile, P.R., "A Pragmatic View of Knowledge and Boundaries: Boundary Objects in New Product Development," *Organization Science*, Vol. 13, No. 4 (Jul. - Aug., 2002), pp. 442-455.