

Tradespace exploration of in-space communications network architectures

Peter Davison^a, Bruce Cameron^{a*} and Edward F. Crawley^b

^aMassachusetts Institute of Technology, Cambridge, MA 02139

^bSkolkovo Institute of Science and Technology, Skolkovo 143025, Russia

Making architectural decisions in long lifecycle systems is challenging because the time between system definition and end of operations can span multiple decades, resulting in shifts in stakeholder needs and major advances in technologies. Space based communications using relay satellite constellations is one such example, requiring substantial up-front planning to define capabilities and size capacity due to the large investment of time and resources. Additionally, there are numerous viable system architectures. In this paper, we develop a graph-based decision method to assess and explore architectural flexibility in the future evolution of long lifecycle systems. The tradespace graph defines edges between similar architectures, quantifies the switching cost between architectures, using graphs to analyze the potential system evolution pathways. In a test case on NASA communication satellites, we find that hosting government communications payloads, in particular optical payloads, on commercial satellites could reduce cost and increase flexibility of the NASA network.

Keywords: space communications; communications architectures; tradespace exploration; decision-making; NASA TDRS network

1. Introduction

Selecting an architecture for a space system or augmenting the architecture of a system already in operation are decisions that have a large impact on system cost, performance, and flexibility. Space systems that consist of multiple satellites and span a long operational lifetime present a unique set of challenges for architecture decision makers as the architecture must be planned so that it may evolve over time. This architecture change over the lifecycle of the system can be driven by the the finite lifetime of in-space assets, the development of new technologies, and changes in stakeholder needs that affect the utility of the system. One subset of systems affected by such change is space-based communications relay networks such as the US MILSATCOM networks, Europe's planned EDRS network, and NASA's Space Network (SN). The SN, composed of a constellation of Tracking Data Relay Satellites (TDRS), in particular presents unique challenges as it serves a variety of stakeholders and currently has a stable architecture. The system in its current form satisfies a set of user needs, however as existing assets reach end of life, NASA has begun to question how the system should be renewed and evolved. This system evolution must be selected from a large number of viable pathways in the face of substantial uncertainty; moreover it must satisfy long-term stakeholder demand as communications technologies change and new on-orbit users are introduced that require relay capacity on the network.

Architecture-level changes to space-based relay systems, which can be realized through the decommissioning of expiring satellites and the development and launch of additional satellites, allow the decision maker to improve system performance and recurring cost over time. In cases where new technologies are to be incorporated into the architecture, technology investment decisions must be made well before the architecture has been defined. As an example, if NASA wishes to incorporate optical inter-satellite links in its next generation of TDRS satellites, the development of optical communications payloads would require

*Corresponding author. Email: bcameron@mit.edu

significant up-front investment before the configurations of the satellites themselves are defined. During this multi-year development period, costs and stakeholder demands would likely change, resulting in uncertainty in the performance and cost of the architecture. Therefore it is essential that both technology investment and architecture decisions are made in order to maintain flexibility in the future architecture.

The problem of architecture flexibility studied here can be recast as a problem of analyzing architectural decisions to determine how they affect the cost and performance of changing from one system architecture to another. As such, the idea of a ‘development pathway’, or a sequence of architectures along a decision timeline, is central to how we quantify flexibility. If we can define an architecture by four independent decisions, each with three options, then we have a tradespace of 81 unique architectures. This tradespace is almost small enough for a human decision maker to fully analyze, however when we consider development paths we see that there are over half a million possible sequences of three architectures that make up a path. This explosion of possible options to analyze necessitates the development of a computational tool to organize this vast amount of information and reduce it to a form that can be understood and used by a decision maker.

The approach taken in this paper to tackling this problem is to encode the tradespace in a graph, with architectures modeled as vertices and feasible decisions to change the architecture as edges. A path through a tradespace graph of communications relay architectures represents a feasible sequence of satellite launch and decommission decisions to evolve the system from a known current state to a future architecture. The performance, or benefit, of an architecture is a property of the vertex itself that can be used to characterize the value of a particular pathway, while the cost of architectural transformation is encoded in the edge weight. The tradespace graph, outlined and applied in Davison et al. (2014b,a) has been shown to enable decision makers to examine the structure of the tradespace and pinpoint key decisions that drive performance, cost, and flexibility. Such a graph-based tradespace model provides a foundation to deploy efficient domain-independent tools such as Dijkstra’s Shortest Path algorithm to exhaustively search potential development paths through the tradespace.

In this paper, we apply this method to the future evolution of the TDRS architecture, specifically focusing our analysis on the impact of hosting a relay payload on a commercial satellite bus on the cost, performance, and flexibility of the future system. In Section 2 we discuss the background of the TDRS network and present a short literature review of existing tradespace analysis methodologies. Section 3 discusses the underlying architecture model used to represent the tradespace studied in this paper, while Section 4 develops the graph-based model that defines the core of our decision support tool. In Section 5 we present the results of our analysis on the flexibility and cost impact of hosted payload decisions, and conclude with an overview of this tool’s impact and points of future work in Section 6.

2. Background and Literature Review

2.1. Background

NASA’s Space Communication and Navigation (SCaN) program was established in 2006 to centralize the management and systems engineering responsibilities of NASA’s communications networks. The three component networks of the SCaN program - the Near-Earth Network (NEN), the Space Network (SN), and the Deep Space Network (DSN) - have evolved from a long series of ground-based and space-based networks as communications technologies and user needs have developed over time (Tsiao (2008)). The NEN is a collection of ground stations which communicate directly with users on orbit, the DSN is a network of three ground stations used to communicate with mission beyond Earth orbit, and the SN consists of a constellation of geosynchronous Tracking Data Relay Satellites (TDRS) and two dedicated ground stations.

The primary goal of the SCaN program is to develop these three communications networks to provide the highest feasible data rates with modern technologies and protocols to all current and future users (Sanchez et al. (2014); NASA (2011)). One study currently underway is a cooperative effort between Goddard Space Flight Center and MIT to develop a tradespace enumeration and evaluation tool for future Space Network architectures (Sanchez et al. (2013, 2014)). As the SN is currently being upgraded with the launch of the third generation of TDRS satellites (TDRS K, L and M), the scope of this tool is concerned with the mid-to-long-term architecture of the network to reduce cost, improve performance, and maintain flexibility for continued development. This tradespace enumeration tool, which will be discussed further in Section 3,

generates the input architecture data to the tradespace graph analysis which is the focus of this paper.

2.2. Literature Review

The methodology presented in this paper can be grouped with a broad set of Architecture Tradespace Exploration tools whose goal is to find the optimal architecture or architectures among a potentially large and complex set. Both Koo et al. (2009) and Arney Arney (2012) find the optimal architecture by making incremental changes to a starting baseline architecture. Genetic Algorithms (GA), a group of optimization tool modeled after natural evolution processes, are used in Brown and Thomas (1998) and Singh and Dagli (2009) to search the tradespace for optimal architectures in a similar manner. A variety of nonlinear programming algorithms are implemented to iteratively search the tradespace for global optima in many MDO methods (Martins and Lambe (2012)). Once an optimal architecture is found, a number of studies have examined how changing requirements and environmental factors perturb the solution space (see, Maher and Poon (1996); Haris and Dagli (2011)).

When evaluating an architecture along multiple dimensions, such as cost and performance, a Pareto frontier can be used to gain insight into the tradespace by pinpointing the non-dominated architectures (McManus and Warmkessel (2004); Battat et al. (2013)). Several tradespace exploration methods are presented in Ross and Hastings (2005) that expand this basic Pareto front analysis to take into account performance or cost uncertainty, changing system requirements, and emergent system properties such as flexibility and robustness.

In addition to searching a tradespace for a optimal standalone architectures, several methods have been proposed to study the evolution of an architecture through multiple configurations. Real options analysis has been used to assess the value of different deployment sequences of a communications constellations (De Weck et al. (2004)), with a similar approach used to quantify the value gained from asset reuse for space exploration systems (Arney (2012)). A retrospective analysis of architecture development is performed in Nakamura and Basili (2005), charting the evolution of a piece of software from initial deployment to its final release, to analyze how different development pathways affect the cost of the final system. A dynamic technology strategy framework is developed to facilitate decision making in Chiesa and Mazini (1998) and the concept of path dependence and its implications on technological and firm evolution are discussed in Araujo and Harrison (2002).

There is an abundance of work done in the realm of space communications; traditionally the approaches that have been taken utilize network simulators and point designs. In Jennings and Heckman (2008); Barritt et al. (2010); Baranyai et al. (2005) space network simulators, created using COTS software, are used to mimic the performance of the Space Network in LEO support operations. Another similar concept is presented in Spangelo and Cutler (2013), in which a framework is developed to analyze small satellite communications operations and maximize performance for important stakeholder metrics, given a particular mission. In the point design method, a limited number of architectures are examined in terms of cost and performance. Bhasin et al. (2008) analyze a point design in the context of a commercial broadband mobile service system while Bhasin et al. (2006) presents a small number of point designs for a Lunar-based relay network. The work in the realm of communications is often done at a component or sub-system level, in contrast to architecture studies which are carried out at the system or the system of systems level. The broader perspective of an architecture-level study enables comparisons of a multitude of options to be made, albeit at the cost of a lower fidelity, more qualitative analysis. Nevertheless, in the presence of uncertainty and a large tradespace of options, architecture level analysis is essential to support decision-makers. NASA has identified the needs which need to be addressed by a future communications network (NASA (2011)), providing specific requirements and network protocols for such a system in Bhasin and Hayden (2001) and Schier et al. (2005).

The exploration of architecture tradespaces has been previously carried out in the context of LEO commercial communication networks (see, de Weck et al. (2003); De Weck et al. (2004); Siddiqi et al. (2005)). In the formulation by De Weck et al. (2004), architectures are encoded as a set of design variables with alternative options, which account for the major architectural decisions. De Weck et al. identify the 5 design decisions to characterize a communications network architecture; from the various alternatives these decisions can take, a tradespace of 600 architectures is created and analyzed in terms of cost and performance.

Graph-based methods have been proposed as a method to model relationships between architectures for




Figure 1: Hierarchy of elements in a SCaN Model architecture.

the purpose of exploring development pathways through a large tradespace (see, Silver and De Weck (2007); Davison et al. (2014b,a)). A network model is proposed in Concho et al. (2011) utilizing an optimization approach to consider future system alternatives. In Silver and De Weck (2007), Silver and de Weck introduce the concept of time-expanded decision networks, essentially static representations of dynamic architecture relationships, to find the progression of launch vehicle configurations that maximize stakeholder value. The problem of finding the optimal sequence of configurations is recast as a shortest path problem through the graph. In Davison et al. (2014b,a) a similar approach is applied at larger scale to the tradespace of architectures for NASA’s future in-space transportation infrastructure.

3. SCaN Architecture Model

The starting point for the application of the tradespace graph framework is a fully enumerated and evaluated set of architectures which span the set of decisions that are of interest to the system architect (Davison et al. (2014b)). The tradespace that forms the core data set for this study was developed in conjunction with NASA’s Space Communication and Navigation (SCaN) program. The tradespace is encoded and generated using the SCaN Tradespace Model (Sanchez et al. (2013, 2014)), which is a computational tool based on the VASSAR framework (Value Assessment of System Architectures using Rules, see Selva and Crawley (2013); Selva (2012)). The core of this framework is the implementation of a rule-based expert system (RBES) to assess the value of an architecture based on the matching of capabilities with stakeholder requirements.

Although an in-depth discussion of architecture enumeration and evaluation is beyond the scope of this paper (we refer the reader to Sanchez et al. (2013, 2014); Selva and Crawley (2013)), it is important to develop an understanding of the SCaN Architecture model, as this model will be closely tied to our development of the tradespace graph. At the highest level, an architecture is represented as a collection of space assets, ground assets, and stakeholders that together determine the capabilities, cost, and benefit of the architecture. The graphical decomposition of elements that define an architecture is shown in Figure 1.

The space assets of an architecture consists of one or more constellations, which are sets of one or more

satellites that relay data between in-space users and ground stations. An important assumption made in this architecture model is that all satellites in a given constellation are identical, both in design and orbit. The constellation is the fundamental building block of our tradespace exploration framework.

At the next level of decomposition, a constellation itself is characterized by a multitude of properties that are used within the architecture evaluator (Sanchez et al. (2014)) to calculate architecture performance and cost. The constellation properties relevant to our tradespace graph framework are:

- *Communications Payloads* - The payload(s) carried by each satellite in the constellation, characterized by the frequency band of the transponder. Options: S, Ka, Ku, optical, Ka+Ku, Tri-band (Ka+Ku+S).
- *Contract Modality* - Characterizes the relationship between the owner of the payloads, the owner of bus, and the consumer of the payload capacity. Options: procurement (NASA owned and operated bus and payloads), hosted-payload (NASA-owned payloads hosted on commercial bus), commercial (commercial bus and payload, capacity purchased by NASA).
- *Lifetime* - The planned time length of time for which satellites in the constellation can perform their communications functions.
- *Satellite Bus* - The underlying infrastructure of each satellite in the constellation. The bus includes non-payload subsystems, such as power and attitude control, that provide a platform for the communications payloads and antennae.
- *Orbit* - All satellites in a constellation share the same orbit, characterized by the semi-major axis, inclination, and eccentricity.
- *Number of Planes* - The number of different orbital planes in which a constellation has satellites.
- *Satellites per Plane* - The number of satellites a constellation has in each of the orbital planes.

On the second major branch of the architecture decomposition, ground assets consist of ground stations which provide the means through which data is passed from the space segment to terrestrial users. Ground stations are characterized by their payloads, location, antennae, and contract modality. In this study ground stations will not be considered an architecture variable and will be fixed across the tradespace. For both space and ground assets, a rule-based model (Sanchez et al. (2013)) is used to estimate the recurring and non-recurring costs of the assets to compare development and operations costs across multiple architectures.

Finally, the set of stakeholders of a particular architecture is what drives the performance, or benefit, of an architecture. Stakeholders such as NASA, NOAA, and USGS are given relative weights in the architecture, with their individual satisfactions computed in the evaluation stage and combined to determine the architecture benefit. Benefit is evaluated by computing stakeholder satisfaction based on the coverage the network provides to a given set of users.

4. SCaN Tradespace Graph Generation

With the SCaN Architecture model described above, we now move to the generation of the tradespace graph. In this implementation of the tradespace graph framework, each vertex has a one-to-one correspondence with an architecture in the SCaN Tradespace such that a hop from one vertex to another represents a decision to change the architecture. The existence of an edge between two vertices represents a feasible architecture change, while the weight of the edge encodes the cost of the change resulting from the development and deployment of new architecture assets.

Therefore the structure of the graph models the relationships between architectures in the tradespace and the feasibility of decisions to transform the architecture from one form to another. A path from one vertex to another, potentially through a number of intermediate vertices, represents an architecture development pathway. It is important to note that in this ‘time-independent’ implementation of the tradespace graph, the passage of time between each sequential decision is not modeled. Dynamic aging of assets in the network will be presented in future work.




Figure 2: A conceptual example of the static graph for the SCaN tradespace.

4.1. Edge Existence

The fundamental idea behind the generation of edges in the graph is that the addition and/or removal of a constellation is the primary means of transforming one architecture into another. In other words, a decision to change the architecture is equivalent to adding and/or removing constellations of satellites, with any constellations common to the old and new architectures retained across the change. This decision model mirrors the process by which NASA has developed the Space Network architecture over its lifetime with the launch of batches of newly procured satellites and the decommissioning of existing ones that have reached the end of their lifetimes (see, Tsiao (2008)).

It should be clear that if we define an architecture change decision only by the constellations that are added and removed, we could jump between any two architectures by deleting and adding an arbitrary number of constellations. However, this type of transformation does not characterize a feasible decision available to NASA since the SCaN program has a limited annual budget with which to augment and operate the network. For this reason, we introduce the constraint that in order for an architecture transformation to be feasible, exactly one constellation must be added. This ‘single added constellation’ rule defines our basic edge existence criterion, with the direction of the edge pointing to the architecture with the added constellation.

It is important to note that this edge existence rule does not take into account the number of constellations that might be removed from the architecture across the jump since we assume that the cost of decommissioning a constellation is already taken into account in the evaluation of non-recurring cost. Figure 2 shows a visualization of a small conceptual tradespace graph to illustrate this simple edge rule.

Looking at the constellations shared between pairs of architectures connected by an edge, we see that exactly one constellation is ‘added’ in the direction of the arrow. For edge 1-3 (i.e. the directed edge from Architecture 1 to Architecture 3), constellation C2 is added with no constellations removed. The same is true for edge 3-4 with C4 added. Architectures 2 and 3 share constellation C2, with C1 added across edge 2-3 and C3 added across edge 3-2. Finally, edge 4-2 is an edge across which two constellations (C1 and C4) are removed, and a single constellation (C3) is added.

In order to implement the edge existence rule above, we must first determine which constellations are shared between two architectures. The SCaN Tradespace Model is built such that each architecture possesses

its own unique constellation objects and identifiers, so to determine whether two constellations are identical we must explicitly match constellation details. To match constellations we will use the architecturally distinguishing features identified in Section 3.

4.2. *Edge Weight*

Once all edges between vertices have been enumerated, the final step in the construction of the tradespace graph is the assignment of edge weights. In this application of the framework, edge weight reflects the cost of switching the architecture from one vertex to another. This switching cost is assumed to be the cost of adding the new constellation to the architecture.

In turn we assume that the cost of adding the constellation to the architecture is well modeled by the non-recurring cost of the constellation itself, which includes the costs associated with designing, manufacturing, and launching each of the satellites. Therefore, the edge weight is assigned to be the non-recurring cost of the constellation added across the edge, which itself is a function of the contract modality of the constellation, the payloads selected, and the satellite bus used.

4.3. *Graph Analysis*

With the structure of the tradespace modeled in the structure of the graph, we can now abstract the multi-faceted tradespace exploration problem into a much simpler set of graph search problems. To find a development path through the tradespace, we must first define the endpoint vertices that represent the initial architecture and final architecture along the pathway. Since we are interested in analyzing the evolution of the SCaN architecture from its current state, and the existing TDRS architecture is well-modeled in the tradespaces we use below, the initial architecture for all development paths is well defined.

The selection of the final architecture is the primary lever by which we vary the analysis cases to produce useful conclusions from the graph. We can select the final architecture based on cost and/or benefit thresholds, constraints on the architecture, and future architectures currently under investigation by SCaN system architects. The simplicity of the model and the efficiency of graph search algorithms allows us to pick many different final architectures and compare them to one another across simple metrics.

To find the development path from the fixed initial architecture to any final architecture, we use Dijkstra's Algorithm (Dijkstra (1959)) to find the shortest path from the initial architecture vertex to the final architecture vertex. The path that is returned is the minimum development cost path from initial to final that obeys the constraint that a single constellation is added to the architecture at each step.

5. SCaN Tradespace Graph Results

In this section we present two sets of results to demonstrate the utility of the tradespace graph and quantitatively analyze several decisions of direct relevance to SCaN stakeholders.

5.1. *Development Cost vs. Benefit Analysis*

The architectures that define the tradespace under analysis are shown in Figure 3 as the black points plotted according to evaluated architecture benefit and lifecycle cost. This tradespace spans architectures over all architecture parameters defined in Section 3 with the exception of contract modality, which is restricted to procurement-only constellations. The tradespace graph generated from these architectures consists of 365 vertices and 6,092 edges.

The current TDRS architecture, which defines the initial architecture for all development paths, is shown on the far left of the plot. For this set of results, we assume that the final architecture is fixed to be the *maximum benefit* architecture in the tradespace, shown on the far right of the plot. Using Dijkstra's algorithm, we can find the minimum cost path between these two architectures, which is plotted as the solid red line in Figure 3.

In the first step along this path a constellation consisting of one GEO satellite with an optical payload is added to the initial TDRS constellation, substantially increasing benefit. In the second step, the TDRS constellation is removed and a two-satellite GEO constellation with two KaKu payloads and one S-band




Figure 3: Two development paths through the tradespace from TDRS to TDRS- Final.

payload is added to the architecture. The final step removes the single optical satellite and adds a two-satellite GEO constellation carrying KaKu, S-band, and optical payloads. The total projected development cost of this (i.e. the total non-recurring cost for all added constellations) is projected to be approximately \$2.8 billion.

If instead of adding the single satellite optical constellation in the first step, we add a constellation of two satellites carrying optical payloads, we change the progression of architectures through the graph. This alternate path, which is no longer the minimum cost path, is shown in Figure 3 as the dashed line. The addition of the second satellite results in an increase in the benefit of the intermediate architectures along the path due to the additional capacity provided by the extra satellite, which can be seen graphically as a shift to the right of the intermediate path points. However, this additional path comes with an increase in development cost of \$210 million.

This trade between development cost and architecture benefit can be expanded by finding other paths that are close to the minimum cost path. We can use another graph search tool, Yen’s Algorithm (Yen (1971)), to efficiently find other paths between the initial and final architectures by finding the K-shortest paths between the two vertices.

Applying this algorithm to the same tradespace and endpoints above, the 100 shortest paths are shown in Figure 4. In this plot, each point represents a unique path from the TDRS architecture to the defined final architecture. Each path is plotted according to the sum of development costs along it and the average benefit of all the architectures along the path. By noting the Pareto Frontier of this data set, we enumerate the paths that represent the non-dominated choices for a decision maker and quantify the options available to trade development cost for stakeholder benefit.

5.2. Static Graph Result: Hosted Payload Decision Analysis

Since using commercial or government satellite buses to host scientific, military, or commercial payloads has already been proven as a viable and cost-efficient capability, NASA is interested in deploying hosted payloads in future versions of the SCaN architecture. In this second set of results we analyze a specific question of direct interest to SCaN stakeholders: what is the benefit of introducing hosted payload constellations into the development path, and how does this benefit vary depending on what type of payload is hosted?




Figure 4: The 100 lowest cost paths through the tradespace from TDRS to TDRS- Final.

5.3. Hosted Payload Constraints

In this section, we include architectures with hosted payload contract modalities in the graph. Although hosting payloads on commercial buses is a proven capability, it is reasonable to assume that NASA stakeholders would be hesitant to make an immediate and drastic change to the current architecture by switching a large portion of the SCaN architecture to a hosted payload contract structure. To account for this preference, we introduce two constraints on the tradespace graph: first only one hosted payload constellation can be present in an architecture, and second this constellation can only be allocated one payload. In other words, if all hosted payloads were to fail the architecture could still maintain an acceptable level of performance.

5.4. Graph Preliminaries

The tradespace used for this analysis is shown in Figure 5 with TDRS highlighted as the initial architecture. The tradespace static graph contains 538 vertices and 10,408 edges. We define a set of candidate final architectures for this result as the collection of 111 candidate final architecture with benefit greater than 0.95, which are shown in green in Figure 5.

5.5. Path Breakdown - Type of Hosted Payload

We find the shortest path from TDRS to each of these 111 architectures through the tradespace graph. These paths contain a mixture of procurement-only and partial hosted payload architectures that obey the two constraints above. We can categorize these paths through the tradespace based on what type of payload (if any) is hosted in the constellations along the development path.

However, each of the 111 development paths (i.e. one path for each final architecture) are not equivalent in terms of their relative cost and the benefit they deliver. One way to compare paths of equal merit to one another is to plot the paths according to cost and average benefit, and then to take samples of the whole set based on proximity to the Pareto Frontier. The concept of the *Fuzzy Pareto set* (?) will be used here compare development paths to one another.

Figure 6 shows the breakdown of paths based on the type of payload hosted along the path for four sizes of the Fuzzy Pareto set. The 10% set, which consists of the 11 ‘best’ paths through the tradespace is




Figure 5: The tradespace used for the static graph hosted payload analysis.

dominated by optical payloads, indicating that investments in optical hosted payload technologies have the highest likelihood of paying off in terms of a favorable development path.

Further, as we include more paths into the Fuzzy Pareto set, we see that the percentage of optical payload paths decreases, indicating that the majority of these paths are close to the Pareto Frontier. By the same argument, the percentage of KaKu and Tri-band payloads increases as we expand the Fuzzy Pareto set, indicating that the distribution of these paths is skewed towards the region of the path space that is further from the Pareto Frontier.

5.6. Path Breakdown - Hosted Payload Savings

To assess the impact of each hosted payload technology in terms of the cost savings realized relative to a procurement-only development, we have to extract additional information from the graph. We start by matching each of the candidate final architectures in the full space (i.e. with hosted payloads allowed) to ones in the space of procurement-only architectures. We do this by finding the architecture which has the same constellations, but with procurement contracts in place of any hosted payload entries.

We then find the shortest path from TDRS to each of these matching final architectures through the tradespace graph *with all architectures with hosted payloads removed*. Now for each path through the full tradespace, we have a matching ‘procurement-only’ path that can be compared directly. Figure 7 plots the paths through the full tradespace in red, with the procurement-only paths shown in black.

Based on the Pareto Frontiers of these two sets of paths it is clear that allowing hosted payloads in the tradespace, even in the limited manner demonstrated here, results in cost savings across a range of final architecture options. Interestingly, the minimum cost path for both the hosted payload and procurement-only tradespaces overlaps. In other words, if the decision maker is only concerned with minimizing cost without taking into account future benefit above the minimum 0.95 threshold, then allowing hosted payloads in the tradespace would not affect the sequence of decisions. Upon closer examination we find that this is due to the fact that this particular final architecture requires the addition of only a single constellation from the TDRS starting point, and therefore can be reached across one edge.

We can go one step further and assess the average cost savings introduced by each type of hosted payload. Since for every hosted payload path we have a matching procurement-only path, we can take the difference between the development cost of the two paths to determine the savings realized by the hosted




Figure 6: The distribution of paths that contain each type of hosted payload for different sizes of the Fuzzy Pareto set.

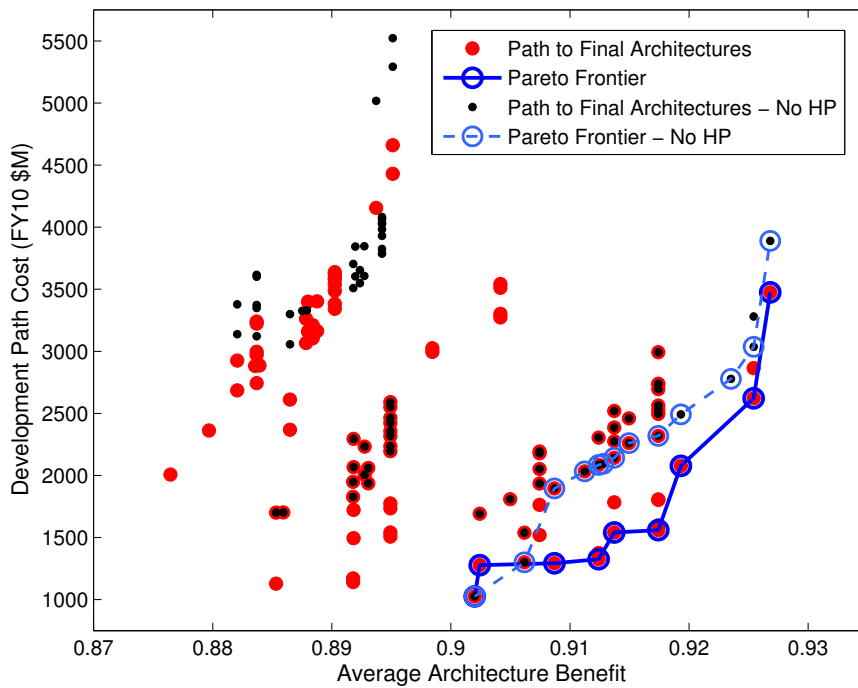


Figure 7: A plot of the development paths from TDRS to 111 final architectures in the hosted payload tradespace and the corresponding procurement-only paths.

payload technology. By using the same Fuzzy Pareto set approach, grouping paths according to hosted payload type, and taking the average savings across all paths in each group, we generate the cost savings data shown in Figure 8.

An important caveat to this data is that the savings calculated are highly sensitive to the hosted payload pricing model used to generate the underlying architecture cost data, so uncertainties in the cost model must be carried forward to the savings analysis. However, using the given model, we can see that average savings across the board is in the \$500-\$700 million range, which, considering that the majority of total development cost is in the \$1-\$3 billion range, tells us that hosted payloads provide significant savings. We also note that in the best paths (small Fuzzy Pareto set size) KaKu payloads provide greater cost savings when hosted versus Optical payloads.

This observation must be taken side by side to the one in Figure 6 that the majority of ‘good’ paths host Optical payloads. Together these points indicate that hosting KaKu payloads provides a greater return on investment with a smaller number of favorable development path options, while hosting Optical payloads provides a marginally smaller return but with a much high degree of flexibility in the number of development options available.

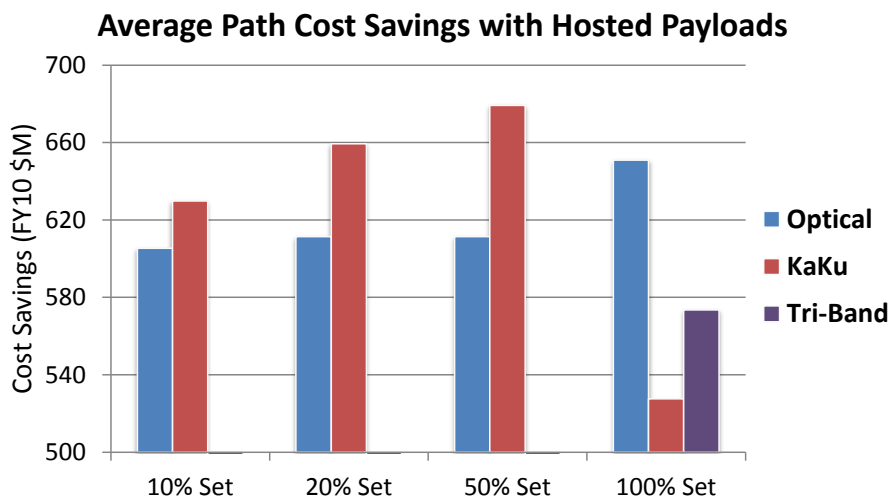


Figure 8: The average cost savings along the development path broken down by payload type and Fuzzy Pareto set size.

6. Conclusions

The uncertainty in future NASA Space Network stakeholder needs, coupled with the large tradespace of available options for the system architecture, produce a problem that fits well within the computational graph based framework proposed in this paper. This graph-based methodology helped to identify architectures with a high level of flexibility and highlight suitable evolution pathways through the tradespace, distilling valuable decision-making information from a complex and uncertain tradespace.

By comparing pairs of paths from the current TDRS configuration to potential future architectures, one path allowing architectures hosted payloads and the other excluding them, we were able to assess the cost savings realized by the introduction of hosted payload technology given a high degree of uncertainty in the future evolution of the architecture. It was found that the hosting of an optical payload appeared in the most efficient paths through the tradespace, however the highest cost savings was realized when hosting KaKu payloads. Furthermore, the tradespace graph approach allowed us to quantify the expected savings from introducing hosted payload technologies despite ambiguity in the exact development of the system over time, with this savings on average in the range of \$500 million to \$700 million.

Considering the high level of architectural ambiguity associated with complex systems, this methodology enables decision makers to account for uncertainty by incorporating flexibility into the architectural decisions.

Architectural change over the lifetime of a space communications system can be driven by many factors; the static tradespace graph analysis allows for the inclusion of multiple future scenarios which translates into a lower system cost and higher system performance, therefore satisfying stakeholder needs in the long-term.

Acknowledgments

This research has been supported by the Skolkovo Institute of Science and Technology.

References

- Araujo, L. and D. Harrison (2002). Path dependence, agency and technological evolution. *Technology Analysis & Strategic Management* 14(1), 5–19.
- Arney, D. (2012). *Rule-Based Graph Theory to Enable Exploration of the Space System Architecture Design Space*. Ph. D. thesis, Georgia Institute of Technology.
- Baranyai, L., E. G. Cuevas, S. Davidow, C. Demaree, and P. DiCaprio (2005). End-to-end network modeling and simulation of integrated terrestrial, airborne and space environments. In *Aerospace Conference, 2005 IEEE*, pp. 1348–1353. IEEE.
- Barritt, B., K. Bhasin, W. Eddy, and S. Matthews (2010, April). Unified approach to modeling and simulation of space communication networks and systems. In *Systems Conference, 2010 4th Annual IEEE*, pp. 133–136.
- Battat, J. A., B. Cameron, A. Rudat, and E. F. Crawley (2013). Technology Decisions Under Architectural Uncertainty. *Journal of Spacecraft and Rockets* 5(5).
- Bhasin, K., A. W. Hackenberg, R. Slywczak, P. Bose, M. Bergamo, and J. Hayden (2006). Lunar relay satellite network for space exploration: architecture, technologies and challenges. In *Proceedings 24th AIAA International Communications Satellite Systems Conference*, pp. 12–14.
- Bhasin, K. and J. L. Hayden (2001). Space Internet architectures and technologies for NASA enterprises. In *Aerospace Conference, 2001, IEEE Proceedings.*, Volume 2, pp. 2–931. IEEE.
- Bhasin, K. B., J. L. Hayden, T. Sartwell, R. A. Miller, and J. J. Hudiburg (2008). Integrated network architecture for nasa’s orion missions. In *SpaceOps 2008 Conference*.
- Brown, A. and M. Thomas (1998). Reengineering the Naval Ship Concept Design Process. In *Ship Systems Engineering Symposium, ASNE*.
- Chiesa, V. and R. Mazini (1998). Towards a framework for dynamic technology strategy. *Technology Analysis & Strategic Management* 10(1), 111–129.
- Concho, A., J. E. Ramirez-Marquez, T. Herald, and B. Sauser (2011). Optimal component substitution within system evolution planning considering multiple-vendor, functionally equivalent commercial products. *Technology Analysis & Strategic Management* 23(5), 509–526.
- Davison, P., B. Cameron, and E. F. Crawley (2014a, March). In-space transportation infrastructure architecture decisions using a weighted graph approach. In *Aerospace Conference, 2014 IEEE*, pp. 1–9.
- Davison, P., B. Cameron, and E. F. Crawley (2014b). Technology Portfolio Planning by Weighted Graph Analysis of Space System Architectures. *Systems Engineering (Accepted for Publication)*.
- de Weck, O., R. de Neufville, and M. Chaize (2003). Enhancing the economics of communications satellites via orbital reconfigurations and staged deployment. *American Institute of Aeronautics and Astronautics*.
- De Weck, O. L., R. D. Neufville, and M. Chaize (2004). Staged Deployment of Communications Satellite Constellations in Low Earth Orbit.
- Dijkstra, E. W. (1959). A note on two problems in connexion with graphs. *Numerische Mathematik* 1, 269–271.
- Haris, K. and C. H. Dagli (2011). Adaptive reconfiguration of complex system architecture. In *Procedia Computer Science*, Volume 6, pp. 147–152.
- Jennings, E. and D. Heckman (2008, March). Performance Characterization of Space Communications and Navigation (SCaN) Network by Simulation. In *Aerospace Conference, 2008 IEEE*, pp. 1–9.

- Koo, B. H. Y., W. L. Simmons, and E. F. Crawley (2009). Algebra of systems: A metalanguage for model synthesis and evaluation. *IEEE Transactions on Systems, Man, and Cybernetics Part A:Systems and Humans* 39, 501–513.
- Maher, M. L. and J. Poon (1996). Modeling Design Exploration as Co-Evolution. *Computer-Aided Civil and Infrastructure Engineering* 11, 195–209.
- Martins, J. R. R. A. and A. B. Lambe (2012). Multidisciplinary Design Optimization: A Survey of Architectures. *AIAA Journal (Submitted for Publication)*.
- McManus, H. and J. M. Warmkessel (2004). Creating Advanced Architecture for Space Systems: Emergent Lessons from New Processes.
- Nakamura, T. and V. R. Basili (2005). Metrics of software architecture changes based on structural distance. In *Proceedings - International Software Metrics Symposium*, Volume 2005, pp. 55–64.
- NASA (2011). Space Communications and Navigation (SCaN) Network Architecture Definition Document (ADD) Volume 1: Executive Summary. Technical report.
- Ross, A. M. and D. E. Hastings (2005). The Tradespace Exploration Paradigm. In *INCOSE International Symposium 2005*, pp. 13–25.
- Sanchez, M., D. Selva, B. Cameron, E. Crawley, A. Seas, and B. Seery (2013). Exploring the architectural trade space of NASAs Space Communication and Navigation Program. In *Aerospace Conference, 2013 IEEE*, pp. 1–16. IEEE.
- Sanchez, M., D. Selva, B. Cameron, E. Crawley, A. Seas, and B. Seery (2014). Results of the MIT Space Communication and Navigation Architecture Study. In *2014 IEEE Aerospace Conference, Big Sky, Montana*.
- Schier, J. S., J. J. Rush, W. D. Williams, and P. Vrotsos (2005). Space communication architecture supporting exploration and science: Plans and studies for 2010-2030. In *1st Space Exploration Conference: Continuing the Voyage of Discovery*, Volume 30.
- Selva, D. (2012). *Rule-Based System Architecting of Earth Observation Satellite Systems*. Ph. D. thesis, Massachusetts Institute of Technology.
- Selva, D. and E. F. Crawley (2013). VASSAR: Value assessment of system architectures using rules. In *IEEE Aerospace Conference Proceedings*.
- Siddiqi, A., J. Mellein, and O. L. De Weck (2005). Optimal Reconfigurations for Increasing Capacity of Communication Satellite Constellations. In *Proceedings of 46th AIAA Structures, Structural Dynamics & Materials Conference, Austin, TX*, pp. 2005–2065.
- Silver, M. R. and O. L. De Weck (2007). Time-expanded decision networks: A framework for designing evolvable complex systems. *Systems Engineering* 10, 167–186.
- Singh, A. and C. H. Dagli (2009). Multi-objective stochastic heuristic methodology for tradespace exploration of a network centric system of systems. In *Systems Conference, 2009 3rd Annual IEEE*, pp. 218–223. IEEE.
- Spangelo, S. and J. Cutler (2013). Analytical Modeling Framework and Applications for Space Communication Networks. *Journal of Aerospace Information Systems* 10, 452–466.
- Tsiao, S. (2008). *Read You Loud and Clear: The Story of NASA’s Spaceflight Tracking and Data Network*. Washington D.C.: NASA History Division.
- Yen, J. Y. (1971). Finding the k Shortest Loopless Paths in a Network. *Management Science* 17(11), 712–716.