Communications Satellites: Time Expanded Graph Exploration of a Tradespace of Architectures*†

Peter Davison, Demetrios Kellari‡, Edward F. Crawley, Bruce G. Cameron

Abstract

The choice of architecture in the design of large-scale complex systems has an enormous bearing on the cost, performance and adaptability of the system. It is often the case that the operational environment of the system changes over time, necessitating evolution of the architecture. In this paper, tradespace graphs of candidate space-based relay networks for NASA's Space Communication and Navigation (SCaN) program, focusing on the Tracking and Data Relay Satellite (TDRS) system, are analyzed in order to examine the potential pathways through the tradespace for a specified planning horizon. We utilize a method that enumerates architectures as vertices in a tradespace graph, taking into account both the fixed properties of the architecture and the variable ages of its assets. We then examine what real options can be put in place to enable future flexibility given uncertainty about the lifespan of existing TDRS assets. In a separate set of analysis, these time-expanded graphs are used to price the potential inclusion of hosted payloads against a range of possible architectural decisions.

1. Introduction

The selection of architecture of a complex, large-scale, engineering system, such as space systems, has large implications on lifecycle cost, performance and flexibility of the system [1,2]. It has been well documented that the majority of costs are committed in the conceptual stages of design [2,3], placing emphasis on the need to choose an architecture wisely [1]. It is arguably equally important however to enable architectural flexibility in order to mitigate the uncertainty of future operating scenarios [4,5,6,7].

In the case of satellite constellations with long operational lifespans, architectural changes may be required to accommodate new disruptive technologies [8,9,10] to account for the limited lifecycle of in-space assets or to react to unexpected changes in stakeholder needs [11,12,13]. Additionally it is often the case that technology development must occur before the architecture is fixed [14,15]; therefore prioritization of technology development, given limited investment resources, is an essential decision [18,19,20,21].

To facilitate decision-making in the face of uncertainty and continue delivering value to stakeholders over the entire lifecycle of the system, a method is required to simplify the complex architecture transformation space into a simpler tradespace. Consequently, the main goal of this work is to turn complex and potentially ambiguous architectural information about the flexibility of architectures into information of value to a system architect.

---

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

† Suggested Reviewers: Prof. Alessandro Golkar (golkar@skolkovotech.ru), Dr. Matthew Silver (silver@cambrianinnovation.com)

‡ Corresponding Author. Tel: +18572726991. Email: dkellari@mit.edu.

NASA's Space Communications and Navigation (SCaN) program aims to address current challenges with the Tracking and Data Relay Satellite (TDRS) network [22,23]. Currently this system satisfies the existing operating requirements including the stakeholder needs, however these requirements are not fixed over time. Changing stakeholder needs, development of new optical technologies and current assets reaching the end of their lifecycle offers an opportunity to augment the architecture of the system. The planning horizon of the SCaN tradespace is 20 to 30 years – the current TDRS network is a mixture of new assets with a planned life of 10 to 15 years and older assets working beyond their intended lifetimes [16,23].

Davison et al. [24,25,26] modeled architecture tradespaces as graphs with vertices representing architectures and edges representing feasible transformations from one architecture to another. This method was applied to the tradespace of potential future SCaN space segment architectures in [27]. Encoding the tradespace in this way has provided a method to scour a vast tradespace of potential architectures for feasible paths, through the exploitation of domain-independent algorithms developed in graph theory. However, the methodology presented in [27] ignored the passage of time between each architecture transformation decision, and as such cannot capture the effects that limited asset lifetimes have on decision-making.

In this paper the concept of tradespace graph is built upon to allow the aging of in-space assets with limited lifespans to be modeled along the evolution of the architecture. The modeling of these graphs, their generation, and their analysis are significantly more complex than the time-independent (static) graph methodology, resulting in a richer set of results that allow several non-intuitive and valuable conclusions to be drawn from the raw tradespace data.

This paper begins, in Section 2, with a review of past research on representing and evaluating architecture tradespaces. In Section 3, the time-expanded tradespace graph is presented and two problems are formulated, in the context of SCaN. The graphs are analyzed for each of the problems and the main findings are presented in Section 4. Section 4 is organized in two main parts corresponding to the two problems and finally, the implications of the analysis with respect to decision-making are presented, followed by a discussion of future work.

2. Literature Review

There is a large body of work on the modeling of system architectures. One of the most widely used tools is the design structure matrix (DSM) [1,28], which models the relationships between system functions and objects of form in a matrix. Another modeling tool developed by Dori et al. [28,29] comprises a graphic modeling language, Object Process Methodology, which represents architectures using the interconnections between system objects, operands and processes. While these are useful representations of architectures, the DSM is limited by the size of the architecture and cannot model constraints, whereas OPM is not an efficient method to generate architecture tradespaces. In addition to these, architectures can be modeled as a set of decisions, featuring a choice among various options, using decision trees [31]. In [32] the concept of an architecture decision graph is proposed for representing architectural options, which employs a more scalable approach than a decision tree. Arguable the most commonly used method for architecture modeling is

the design vector model [32,33], which provides a transformation from input design parameters to output system evaluation metrics by creating a model for the represented system.

The above methods are advantageous for representing single architectures in different cases, but only a few, such as the design vector model, can be used to enumerate an architecture tradespace. A tradespace is a multi-dimensional design space populated with many architectures, where the dimensions consist of architecture evaluation metrics, such as can be seen in [35]. After a tradespace has been defined, it is common to engage in a tradespace exploration to identify the 'optimal' architecture. A variety of methods have been employed in previous work for this purpose including incremental changes to a baseline architecture using an Ant Colony Optimization [36] and Genetic Algorithms [37]. Identifying the optimal architecture can prove difficult since there are multiple architecture evaluation metrics; therefore usually the set of non-dominated architectures is found using Pareto front analysis [39,40].

In large-scale, complex engineering systems it is often found that external forces may disrupt the architectural assumptions, making the chosen architecture suboptimal. In [26] Davison et al. use a weighted graph analysis method to demonstrate the low performance of non-flexible architectures when subjected to externally driven change. In [40] and [41] Genetic Algorithms are used to allow the architecture to evolve when faced with exogenous changes. In the context of decision-making, an architect may wish to understand how an architecture can be extended, to improve its performance or cost in reaction to environmental changes, after it has already been selected. The design of architectures for flexibility or extensibility is distinct from searching for a shifting optimum design point, since now there is assumed to be a cost for architectural changes, such as an increase in functionality.

The staged development and deployment of a constellation of satellites, reacting to uncertainty in demand, is studied in [35]. An initial suboptimal architecture is selected allowing it to be reconfigured with additional assets in response to stakeholder needs. The lack of optimality in initial architecture is offset by the system upgrade options as more information becomes available. In [42] Arney examines the reuse of systems within a specific architecture for different missions in order to evaluate the flexibility of feasible paths through the tradespace.

The concept of a Time-Expanded Decision Network is introduced by Silver and de Weck in [8] for the exploration of development pathways through a design space. A comprehensive cost model, including system development, operation and retirement is used to formulate a dynamic network, containing the possible architecture sequences over a period of time, from an architectural tradespace. The development pathway of the architecture among the many available options is then considered to be the shortest possible path through this network. In this method the architecture performance and cost at each time step is evaluated based on dynamic stakeholder demand profiles.

It will be seen in Section 3 that the notion of distance between architectures is essential to the method proposed in this paper. The distance, in this context, represents the difficulty or cost of switching from one architecture to another. In [43] the number of common systems between architectures is used as a proxy for evolving the system. The delta DSM, proposed in [44], is a matrix construct of the difference between two architectures, which can be used

to quantify the distance between architectures. De Weck and Silver [8] use a switching cost metric, involving a detailed calculation of system costs. In this method the architectures are enumerated, switching costs are calculated and a static tradespace graph is created. In order to carry out a 10-year analysis with time steps of 1 year, a time-expanded vertex is generated for each static vertex at each time-step, followed by tradespace exploration utilizing graph search algorithms. This labor-intensive, bottom-up approach is limited by the size of the tradespace. The method presented in this paper differs in that a static graph is built followed by a single time-expansion step corresponding to the possible ages of the assets; therefore the aging of the architecture is accounted for using a lower ratio of time-expanded vertices to static vertices.

In summary, three broad uses of tradespace exploration have been identified, which motivate much of the research in this domain: to enable the inclusion of architectural flexibility [6] in conceptual design to respond to changes in stakeholder needs or changes in operating scenario [11,12,13]; to incorporate new disruptive technologies without a complete system overhaul [8,9,10,53]; and lastly to develop new iterations of the system over time to improve performance or reduce cost [46].

3. Problem Formulation

3.1. Tradespace graph and Time Expanded Graph

Time-independent tradespace graphs have been developed and analyzed in [26]. This approach transforms the space of feasible architectures into a weighted graph, consisting of vertices and edges, where the architecture vertices are connected by edges representing feasible transformations and weighted according to the high-level cost of this transformation.

In order to incorporate the passage of time into these analyses thus taking into consideration the effect of planning horizons on architecture development, our attention is turned to time-expanded graphs. In [8], Time Expanded Decision Networks are used so as to convert a dynamic graph, with dynamic vertex characteristics, into a time-invariant graph, with fixed vertices and edges. The main advantage of this tradespace approach, over examining architectures in isolation, is it incorporates the relationships between architectures to allow the implementation of methods used in graph theory, such as Dijkstra's Shortest Path Algorithm [47].

In Time Expanded Decision Networks, exogenous factors including different stakeholder needs, drive the changes in time and determine the performance and cost of the architectures. The Time-Expanded Graph method, developed in this paper, models endogenous changes such as the aging of architectural assets. This provides the opportunity to optimize architectures for a fixed planning horizon considering the finite lifecycle of the system components. Hence moving from vertex to vertex along an architecture pathway in the time-expanded tradespace represents the evolution of an architecture and its asset portfolio over one time step, which is specified by the user.

3.1.1. Vertex Expansion

The global time step is the number of years by which the age of each constellation will be incremented at each step between vertices. The size of the time step determines the level of fidelity in which to model the decision-making process as a function of time; thus the

optimal time step is one which approximates the time period of major SCaN architecture decisions. Considering that NASA's current upgrade cycle is approximately 5-7 years and constellation ages very from 10 to 20 years [22], a time step of 5 years is used in this analysis.

In this formulation, the constellation ages of a vertex represent the number of years that the constellation has been in orbit at the beginning of the time step. To generate the time-expanded vertices corresponding to an architecture, all the possible combinations of feasible constellation ages are found for the given lifetime of each constellation. A vertex is then created in the graph for each of these age sets, storing information on both the architecture and constellation ages. The rules for vertex expansion are best illustrated graphically in *Figure 1* for an architecture with two constellations, C1 and C2, each with lifetime equal to 15.

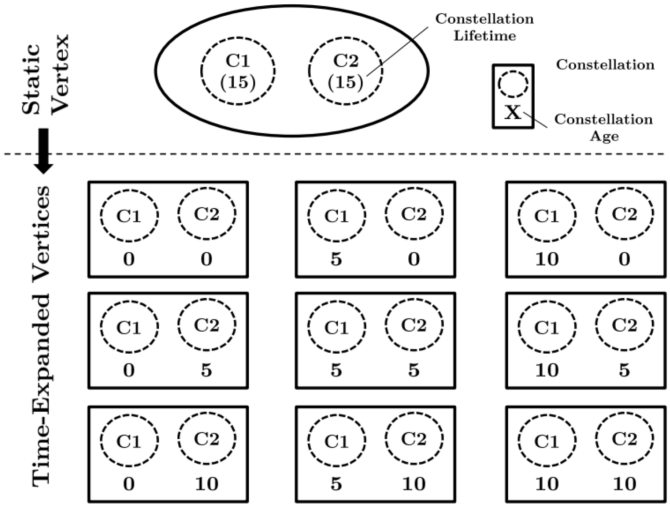

Figure 1: Example vertex expansion for the SCaN Time-Expanded Tradespace Graph. The static vertex is mapped to 9 time-expanded vertices.

3.1.2. Edge Existence

The key to this method is in the modeling of the relationships between vertices, since the presence of an edge represents a feasible step along the evolutionary pathway of an architecture. In the context of SCaN, a time-expanded graph edge represents both the addition/removal of a constellation and the passage of time. Due to the incorporation of time the edge is directional and contains a 'source' and a 'sink' vertex.

Each edge in the time-independent graph, described in [26], is used to form the time-expanded graph edges. These original edges are defined using the relationship between asset portfolios of the two architectures across an edge; the condition for edge existence is that at most a single asset can be added across it. In the time-expanded case, the vertices are 'expanded', to form new vertices incorporating all the possible ages of the constellations (see *Figure 1*). A time-expanded edge is placed between vertices, which satisfy three criteria:

1. Each constellation that is common to both architectures has its age incremented by the time step.
2. Any constellation added to the sink architecture has an age equal to zero.

3. Any constellation that is removed from the source architecture has an age equal to its lifetime minus the time step.

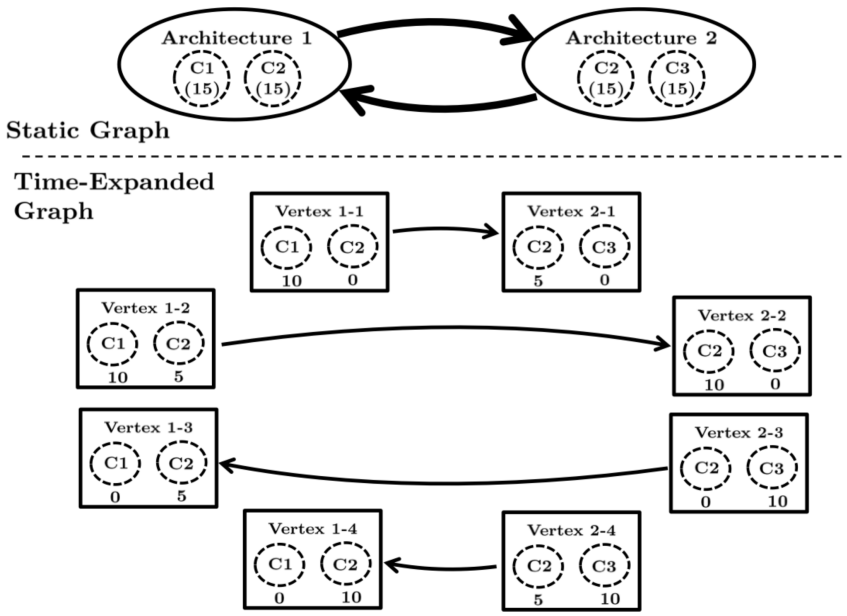

Figure 2: Example edge expansion for the SCaN Time-Expanded Tradespace Graph.

A representation of the expansion of edges for two architectures with two constellations each can be seen in *Figure 2*. Taking the lifetime of the constellations to be 15 years and the time step to be 5 years, the new time-expanded edges are generated using the time-independent relations.

In addition to deciding to add and/or remove constellations at each time step, a decision maker might also wish to retain the architecture in its current configuration. This decision is desirable to model since, in reality, NASA has neither the resources nor the need to add assets to the Space Network every five years. Hence, allowing the architecture to go unchanged over a time step enables a great deal of modeling flexibility in terms of analyzing decisions in response to changing demand or constraints on non-recurring costs.

In the implementation of this graph it must be ensured that, across an edge, all constellations increase in age by the time step, all constellations assigned the maximum age are removed, and all added constellations are assigned an age equal to zero. The time-expansion step of the tradespace generation is suited to a rule-based implementation, due to the well-defined logic required to perform this task [56,57,58].

3.1.3 Edge Weight

Finally, with the structure of the time-expanded graph defined, it is possible to calculate the weights for the newly created edges, based on recurring and non-recurring cost. Recurring cost includes the cost of maintaining and operating the constellations over a given period of time (see *Equation 1*) and non-recurring cost comprises the addition of a new constellation. The fundamental idea behind the allocation of an edge weight is to indicate the cost of architecture transformation from one vertex to another.

$$C_{rec} = \begin{cases} operations\text{-}cost & \text{for procurement} \\ service\text{-}fee & \text{for hosted-payload} \\ service\text{-}fee \times \frac{ops}{ops+ant+bus+payl+launch} & \text{for commercial} \end{cases} \quad (1)$$

The edge weight is equal to the sum of the recurring cost of all constellations in the source architecture per time-period, and the non-recurring cost of any constellation added across the edge. Hence, the edge weight calculation can be represented as,

$$W_{time-expanded} = C^*_{nrec} + T \sum_i {C^i_{rec}}/{L^i}, \quad (2)$$

where L^i is the lifetime in years of the ith constellation in the source vertex, C^i_{rec} is the total recurring cost of the ith constellation, T is the time step in years and C^*_{nrec} is the non-recurring cost of the constellation added across the edge.

3.2. Problem 1 Formulation: Extension of current TDRS by 5 years

Firstly the problem related to the expected lifetime of TDRS will be analyzed; that is, given a specific planning horizon and the possibility that the current TDRS constellation will outlive its design lifetime, can a plan be formulated which minimizes the development costs while maintianing the system functionality.

3.2.1. TDRS Age Assumptions

This analysis uses a reference year of 2020 to begin the development path, since this is likely the earliest point in which this type of analysis could be applied to SCaN architecture decisions, with the planning window stretching to 2050.

The current TDRS system consists of seven operational GEO satellites grouped into three generations (or constellations) [51]. IT is assumed that the first and second generations of satellites will no longer be operational in 2020. That leaves the third generation of TDRS satellites, of which two are currently in orbit, and a third is planned for launch in 2015 [51]. With a planned retirment around 2030, all current TDRS assets will not be operational 10 years into our development path.

However, like many in-space assets there is a significant probability that the operational lifetime of the current constellation will be extended beyond its intended lifetime. Consequently a decision maker may be interested in analyzing how an evolution plan can be implemented that accounts for these future operating scenarios [52].

3.2.2. Tradespace Description

For our analysis, we make the assumption that there is a known probability, p, that the constellation lifetime is extended by five years beyond the planned retirement date, i.e. shifting it from 2030 to 2035. Starting our analysis in 2020 it is assumed that the actual decision to decommission the TDRS is not made until 2025. This implies that during the period from 2020 to 2025 the architectural decisions must be sufficiently flexible to account for the uncertainty of the retirement date. Ideally, the initial decision should reflect the estimated probability of TDRS being extended the extra fives years. The tradespace to be

used in this analysis can be seen in *Figure 3* plotted along the lifecycle cost and architecture benefit.

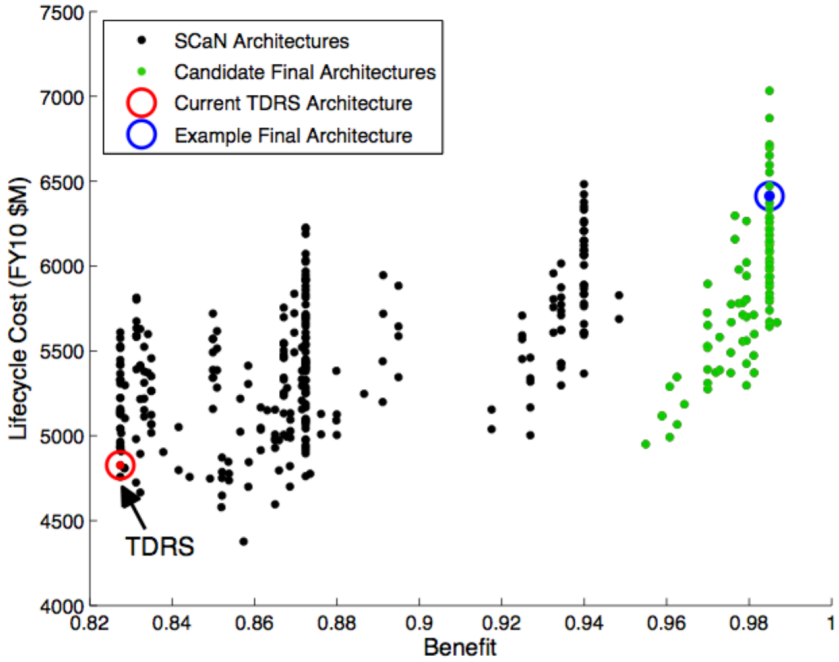

Figure 3: The SCaN tradespace showing initial architecture, candidate final architectures with benefit greater than 0.95, and an example final architecture.

The problem of selecting an initial architecture is given, since it must represent the TDRS architecture with three satellites carrying two Tri-band (Ka, Ku, and S band) [53] antennae and the respective constellation ages. Meanwhile, the set of candidate final architectures, shown in *Figure 3*, consists of all architectures with a benefit greater than 0.95, therefore satisfying the stakeholders' high performance criteria. The architecture evolution pathways are calculated to be the shortest path through the graph from the TDRS vertex to each of vertices in the set of candidate final architectures.

3.2.3. Definition and Computation of Path Sets

Two sets of paths to each of these candidate final vertices are computed. The first set, known as the *projected lifetime set*, assumes that TDRS is retired in 2030 after 10 years of operation along the development path, as was the initial design. The second set, known as the *extended lifetime set*, assumes that TDRS is extended five additional years until 2035 for 15 years of operation along the development path; furthermore it is assumed that knowledge of this extension comes after the first time period.

In order to accomplish the latter point, the edge connecting the initial TDRS architecture to itself is removed from the analysis. This is done to incorporate flexibility in this decision, that is, the decision to maintain the architecture 'as-is' does not account for multiple possible future scenarios. Both sets of 30-year paths are shown in *Figure 4* with the projected lifetime set in blue and the extended lifetime set in red.

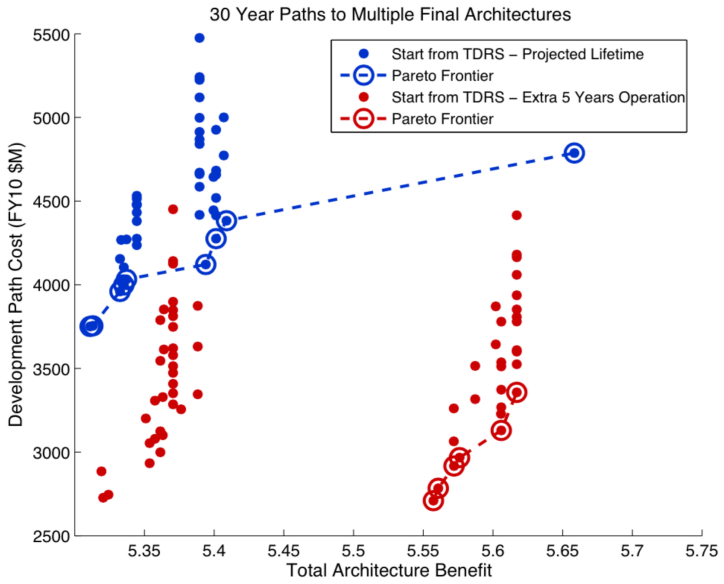

Figure 4: The development paths from TDRS to each candidate time-expanded vertex with a 30 year development path. Paths shown for both the planned retirement of the TDRS constellation in 2030 and for the scenario in which TDRS operations are extended five years to 2035. Extended lifetime is shown in red and projected lifetime in blue.

3.2.4. Pairing of Extended and Projected Paths

A valuable result for decision makers involves evaluation of the choice of architecture in the initial time step to maximize the benefits, considering the need for flexibility of the TDRS architecture to account for the two operating scenarios. To assess the performance and cost of the first evolution decision, i.e. the first edge in the graph, paths from the projected lifetime set are paired with paths from the extended lifetime set. These are then compared in terms of the expected development cost and aggregate benefit based on the probability of the TDRS lifetime being extended. The pairs are found by searching for paths with the same architecture after the first time step.

3.2.5. Expected Costs and Benefits

The expected development cost of the path pair is defined as,

$$E[C] = p \times C_{Extended} + (1-p) \times C_{Projected}, \quad (3)$$

and the expected total benefit is defined as,

$$E[B] = p \times B_{Extended} + (1-p) \times B_{Projected}, \quad (4)$$

where p is equal to the estimated probability that TDRS will be extended for an additional five years, C is the total path cost and B is the aggregate benefit along the 30 year development path. This formulation enables uncertainty in the future scenario to be considered in the model.

3.3. Problem 2: Effect of Hosted Payloads

The use of commercial or governmental satellite buses to host scientific, military, or commercial payloads has been proven to reduce costs over traditional procurement strategy [54]; therefore NASA has shown interest in deploying hosted payloads for the SCaN

architecture [55]. The second result of this paper aims at trying to understand the benefits of introducing hosted payloads into a 30 year evolution of the SCaN architecture, in particular what is the effect of varying the type of payload.

3.3.1. Constraints

It is reasonable to assume that NASA stakeholders would not wish to make drastic changes to the current architecture to mitigate the associated risks such as delayed deployment. Therefore hosted payloads are gradually incorporated into the model, with a limit of one per constellation and only one asset assigned to it; hence this enables the network to continue functioning in it's absence. Finally, the candidate final architectures are constrained to be procurement-only; this is done to model the realistic scenario in which NASA may want to examine the available options, in terms of hosted payloads, until a procured satellite can be developed and deployed.

3.3.2. Tradespace Description

The shortest paths between the TDRS vertex and each of the candidate final vertices are computed for two distinct cases: one including hosted payloads and one excluding hosted payloads in order to compare the two. Note that there is a one-to-one mapping between the procurement-only paths and the paths which incorporate hosted payloads.

4. RESULTS

The first result examines the case in which the lifetime of TDRS is extended by 5 years and the second result examines the effect of introducing hosted payloads into the development path of the system.

4.1. Results: Extension of Current TDRS by Five Years

The expected cost and benefit for all of these pairings are plotted in *Figure 5* for two values of p: p = 0.25 on the left and p = 0.75 on the right.

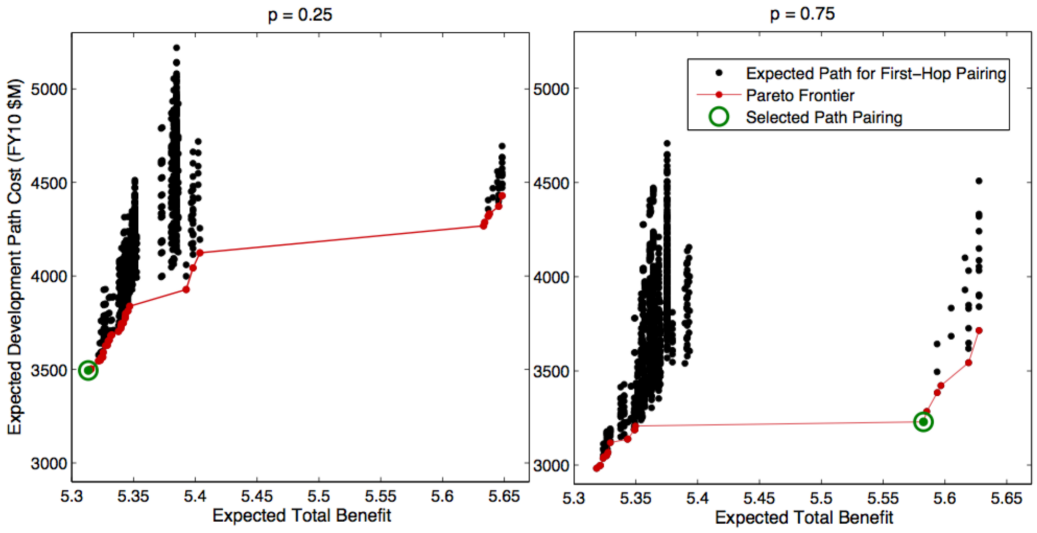

Figure 5: Expected development path cost and average architecture benefit for each pair of pathways that share the same initial step. Two values of p, the probability of TDRS constellation extension by 5 years, are shown.

In the next sections, a Pareto optimal pairing from each of the probability assignments is selected for further exporation.

4.1.1. Exploring Paired Paths - p = 0.25

In the case where p = 0.25 the minimum expected cost pairing is chosen, which can be seen highlighted on the left plot in *Figure 5*. A detailed schematic decomposition of the stages of the evolution of both paths in the pairing can be observed in *Figure 6*. The first step, from 2020 to 2025, along both paths is by definition the same; the new architecture consists of the addition of a single satellite carrying a KaKu and S-band payload to the already present TDRS constellation [53]. In the next time step the path diverges into two distinct trajectories, with the top branch representing the extended lifetime case and the bottom representing the projected lifetime case.

It can be seen that total development cost for the extended lifetime path is approximately $1 billion less than the projected lifetime path. Naturally, the cause of this cost difference can be traced to the non-recurring cost associated with the added constellations along each branch; in the extended lifetime path three constellations are developed and launched over the 30 year model, while four constellations are added along the projected lifetime path in this same period.

Although there are cost savings due to the TDRS extension, the value of the assets at the end of this period has not been accounted for. That is, at the end of the period, in the extended lifetime path there are 15 'satellite-years' remaining among all the constellations and in the projected lifetime path there are 25 satellite-years remaining.

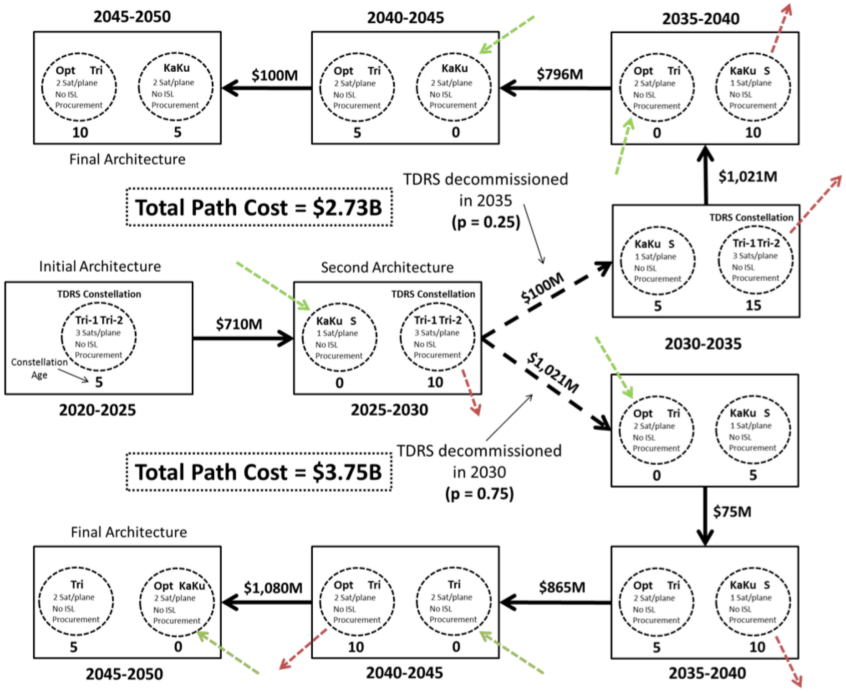

Figure 6: Diagram of the two possible evolution branches that are available given the selection of the first hop in the tradespace to minimize expected development cost with p=0.25. This selection is shown as the green circle in the left plot of Figure 5.

4.1.2. Exploring Paired Paths - p = 0.75

A pair of archtecture development paths from the Pareto frontier is selected at random for the case when p=0.75, as highlighted in *Figure 5*.

A schematic breakdown of the steps along these two evolution pathways are displayed in *Figure 7*. In this case an optical payload is added to the existing TDRS assets in the first step, which agrees with the earlier observation that high values of p favour optical payloads. Comparing the projected lifetime trajectory of this path with that from the p=0.25 path, it is noticeable that it is significantly more costly; however due to the low value of (1-p) in *Equation 3* in this scenario, the higher cost has a lower impact on the calculation of expected cost (and expected benefit). It is then intuitive that this first decision favors the extended lifetime over the projected lifetime branch, thus making it a poor fit for the scenario in which TDRS is to be decommissioned in 2030. Note that by not including this probability and simply taking an average, the initial decision taken in the previous result is more attractive due to lower costs for a similar benefit.

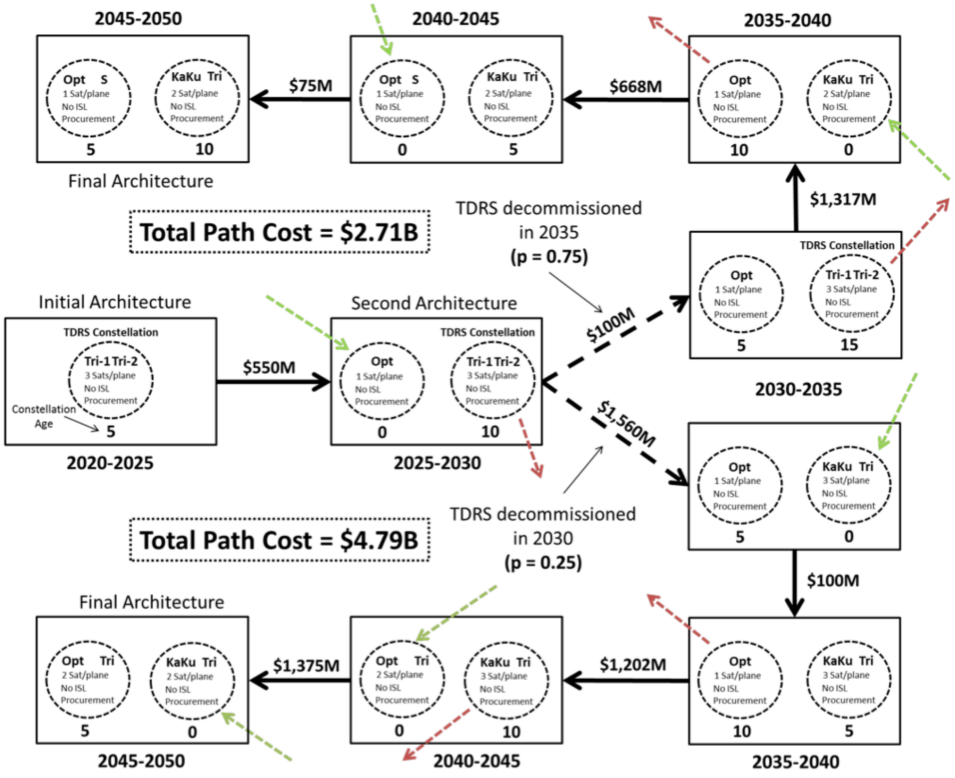

Figure 7: Diagram of the two possible evolution branches that are available given the selection of the first hop in the tradespace to lay on the expected development cost - expected benefit Pareto front with p=0.75. This selection is shown as the green circle in the right plot of Figure 5.

It is immediately clear that the extra five years of TDRS lifetime results in significant cost savings and benefit improvements for almost all final vertices. These cost savings emerge from the fact that extending the lifetime of TDRS by five years reduces the number of constellations to be developed, thus dramatically reducing non-recurring cost.

4.1.3. Remarks

Apart from the difference in the development costs between the two graphs in *Figure 5*, it is noticeable that the structure of the plot changes. The cluster of architectures on the right of the plot are predominately those in which an optical payload is added after the first time step. This suggests that an optical installation is much better suited to a high probability of TDRS extension, perhaps since it operates in synergy with the large Tri-band TDRS constellation. Such observations emerge from this formulation of a tradespace of architecture development pathways. Analysis of these graphs, including changing probabilities of future scanrios, enable the architect to select the decision, which maximizes the performance, while maintaining flexibility due to uncertainty.

4.2. Results: Effect of Hosted Payloads

The shortest paths between the TDRS vertex and each of the candidate final vertices are computed for two distinct cases: one including hosted payloads and one excluding hosted payloads, in order to compare the two.

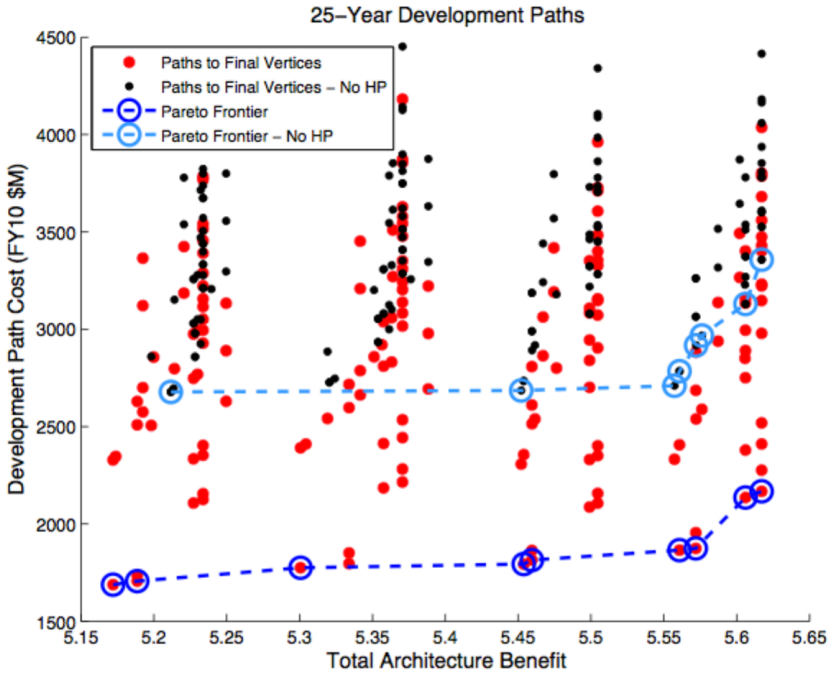

Figure 8: Plot of all the development paths with length equal to 30 years including all the time-expanded vertices corresponding to candidate final architectures. Red points show paths through the tradespace with single hosted payload constellations permitted, while black points enforce the constraint that all constellations are procured.

It is clear from the plot in *Figure 8* that there is a large gap between the two sets of paths, with a vertical shift of approximately $1 billion between the two Pareto frontiers. This observation tells us that when a 30 year development period is taken into consideration, the effect of introducing hosted payloads into the tradespace, is considerable. Note that the architecture benefit along the development pathway remains the same because only the contract modality is changing.

4.2.1. Example Path Comparison - Maximum Savings

This result, compares the details of two paths which start and end at the same vertex, one utilizing hosted payloads and the other procurement-only. The detailed steps of the evolution can be seen in the schematic in *Figure 9*, with a red outline representing a constellation containing a hosted payload.

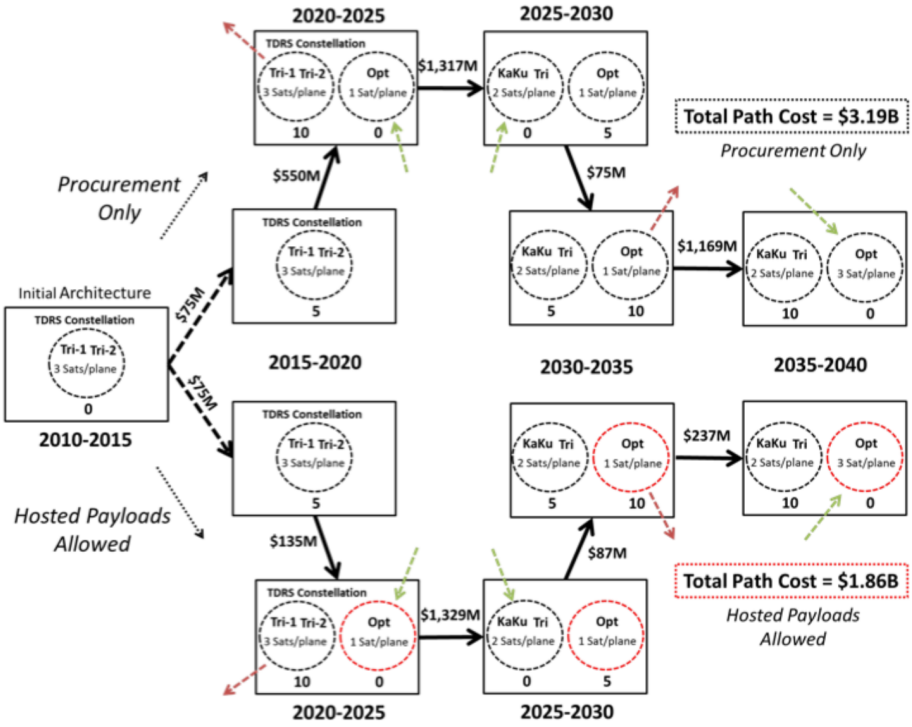

Figure 9: Diagrams of the matching hosted payload and non-hosted payload paths through the time-expanded graph that together represent a development path savings of $1.33 billion with the introduction of single hosted payload constellations into the tradespace. The red constellation outline indicates hosted payload contract modality.

These paths are located on the Pareto frontier of their respective graphs, and represent the architecture evolution with the maximum hosted payload savings; over the 30 year development interval, this is calculated to be approximately $1.33 billion.

It can be seen that the difference between the two paths in intermediate decisions, between the initial and final architecture, is the contract modality of the constellation added across the edge. Note that this result is not general, since the added constellations differ in many other pairings. The cost difference, as expected, stem from the non-recurring costs of the procured constellations which amount to being a factor of five more expensive than hosted payload constellations. In fact, the recurring costs of the hosted payload constellations are higher than the corresponding procurement constellations, due to high service fees, as is evident from the weight along the 2030 edge.

5. Conclusion

The concept of a time-expanded graph has been demonstrated, which builds on the tradespace graph model of the system architecture development, to account for the aging of in-space assets and constraints on the length of the development time line. The main new concept applied in this paper was the inclusion of the lifecycle of satellites, to account for time in the architecture tradespace analysis including launch, operation and decommission. Posing the architectural tradespace exploration in this way yields more realistic results to support in decision-making when considering a specific planning horizon, and offers the opportunity to augment the system architecture.

In the first set of results, it was found that extending the lifetime of the TDRS constellation by an extra five years results in a cost and benefit improvement on average over the 30 year development cycle. Based on the expected probability of TDRS being decommissioned or extended an extra five years, the initial Pareto-optimal architecture decisions are determined.

In the next set of results the effect of utilizing hosted payload, over a 30-year development period, was analyzed and compared with previous results which used a time-independent approach. It was found that the introduction of a time constraint magnifies the savings achieved due to hosted payloads, and changes the deployment breakdown of hosted payloads along the development timeline.

The methods presented in the paper account for a variety of future operating scenarios, evaluating the performance and cost of architecture evolution based on a probability of each scenario. The outcome of the analysis is to pinpoint key decisions or architectural changes which ensure a Pareto-optimal architecture, while taking into account flexibility. Identifying common decisions among the best performing architectures is an indication that this architectural change will introduce flexibility. This can be taken one step further by isolating many architectures with the same decisions given a range of probabilities of future scenarios, and performing tradeoffs. In this way the opportunities, which arise from the limited lifecycle of in-space assets, can be realized while accounting for uncertainty.

In this work architecture evolution with time is quantitively encoded in the time-expanded tradespace graph. This enables the development of a multitude of tools for analysis of information pertaining to the relationships of architectures, considering the passage of time. The ultimate goal is to assist decision makers with a tool that mitigates the uncertainty associated with complex, large-scale projects.

6. References

[1] E. Crawley, B. Cameron, D. Selva, System Architecture, Pearson, 2015.

[2] E. Crawley, O. Weck de, S. Eppinger, C. Magee, J. Moses, W. Seering, et al., The Influence of Architecture in Engineering Systems, Eng. Syst. Monogr. (2004) 30.

[3] M. Krueger, C.D. Walden, E.R.D. Hamelin, Systems Enginering Handbook, International Council on Systems Engineering, 2011.

[4] E. Rechtin, The art of systems architecting, IEEE Spectr. 29 (1992) 66–69. doi:10.1109/6.158642.

[5] H. McManus, D. Hastings, A framework for understanding uncertainty and its mitigation and exploitation in complex systems, IEEE Eng. Manag. Rev. 34 (2006) 81–94.

[6] S.H. Thomke, The role of flexibility in the development of new products: An empirical study, Res. Policy. 26 (1997) 105–119.

[7] A.F. Mehr, I.Y. Tumer, Risk-Based Decision-Making for Managing Resources During the Design of Complex Space Exploration Systems, J. Mech. Des. 128 (2006) 1014.

[8] M.R. Silver, O.L. De Weck, Time-expanded decision networks: A framework for designing evolvable complex systems, Syst. Eng. 10 (2007) 167–186.

[9] C.M. Christensen, J.L. Bower, Customer Power, Strategic Investment, and the Failure of Leading Firms, Strateg. Manag. J. 17 (1996) 197–218.

[10] C.M. Christensen, The Innovator's Dilemma, Harvard Business Review Press, 1st edition, 1997.

[11] R.M. Henderson, K.B. Clark, Architectural Innovation: The Reconfiguration of Existing Product Technologies and the Failure of Established Firms, Adm. Sci. Q. 35 (1990) 9–30.

[12] P.M.E. Porter, Competitive Strategy, Meas. Bus. Excell. 1 (1997) 12–17.

[13] B.G. Cameron, E.F. Crawley, G. Loureiro, E.S. Rebentisch, Value flow mapping: Using networks to inform stakeholder analysis, Acta Astronaut. 62 (2008) 324–333.

[14] R.E. Freeman, Strategic Management: A Stakeholder Approach, 1984.

[15] A. Yassine, D. Braha, Complex Concurrent Engineering and the Design Structure Matrix Method, Concurr. Eng. 11 (2003) 165–176.

[16] G. Sohlenius, Concurrent Engineering, CIRP Ann. - Manuf. Technol. 41 (1992) 645–655.

[17] W.J.A. Dahm, Technology Horizons a Vision for Air Force Science & Technology During 2010-2030, US Air Force, Tech. Rep. 2 (2010) 237.

[18] A. Eder, M. Linde, Efficient and Dynamic--The BMW Group Roadmap for the Application of Thermoelectric Generators., in: Second Thermoelectr. Appl. Work. San Diego, 2011.

[19] W.M. Hawes, M.R. Duffey, Formulation of Financial Valuation Methodologies for NASA's Human Spaceflight Projects., Proj. Manag. J. 39 (2008) 85–94.

[20] Z. Szajnfarber, M.G. Richards, A.L. Weigel, Challenges to Innovation in the Government Space Sector, DTIC Document, 2011.

[21] M. Sanchez, D. Selva, B. Cameron, E. Crawley, A. Seas, B. Seery, Results of the MIT Space Communication and Navigation Architecture Study, in: 2014 IEEE Aerosp. Conf. Big Sky, Mont., 2014.

[22] J.S. Shrier, B. Younes, NASA. Space Communications and Navigation (SCaN) Network Architecture Definition Document (ADD) Volume 1: Executive Summary, 2011.

[23] M.J. Liberatore, An extension of the analytic hierarchy process for industrial R&D project selection and resource allocation, IEEE Trans. Eng. Manag. EM-34 (1987) 12–18.

[24] P. Davison, B. Cameron, E.F. Crawley, Technology Portfolio Planning by Weighted Graph Analysis of System Architectures, Syst. Eng. (2014).

[25] P.L. Davison, Tradespace Exploration for Space System Architectures : A Weighted Graph Framework, 2014.

[26] P. Davison, B. Cameron, E.F. Crawley, In-space transportation infrastructure architecture decisions using a weighted graph approach, in: Aerosp. Conf. 2014 IEEE, 2014: pp. 1–9.

[27] P.L. Davison, B.G. Cameron, E.F. Crawley, Tradespace Exploration of In-Space Communications Network Architectures, (Pending Review), (2014).

[28] T.R. Browning, Applying the design structure matrix to system decomposition and integration problems: a review and new directions, IEEE Trans. Eng. Manag. 48 (2001).

[29] D. Dori, Object-Process Methodology: A Holistic Systems Paradigm; with CD-ROM, Springer, 2002.

[30] I. Reinhartz-Berger, D. Dori, S. Katz, OPM/Web – Object-Process Methodology for Developing Web Applications, Ann. Softw. Eng. 13 (2002) 141–161.

[31] Z. Covaliu, R.M. Oliver, Representation and solution of decision problems using sequential decision diagrams, Manage. Sci. 41 (1995) 1860.

[32] W. Simmons, A framework for decision support in systems architecting, Ph.D. Diss. Dep. Aeronaut. Astronaut. Massachusetts Inst. Technol. (2008).

[33] A.M. Ross, D.E. Hastings, The tradespace exploration paradigm, in: INCOSE Int. Symp. 2005, 2005: p. 13.

[34] T.W. Simpson, T.M. Mauery, J. Korte, F. Mistree, Kriging models for global approximation in simulation-based multidisciplinary design optimization, AIAA J. 39 (2001) 2233–2241.

[35] O.L. De Weck, R. De Neufville, M. Chaize, Staged Deployment of Communications Satellite Constellations in Low Earth Orbit, J. Aerosp. Comput. Information, Commun. 1 (2004) 119–136.

[36] B.H.Y. Koo, W.L. Simmons, E.F. Crawley, Algebra of systems: A metalanguage for model synthesis and evaluation, IEEE Trans. Syst. Man, Cybern. Part ASystems Humans. 39 (2009) 501–513.

[37] A. Singh, C.H. Dagli, Multi-objective stochastic heuristic methodology for tradespace exploration of a network centric system of systems, in: Syst. Conf. 2009 3rd Annu. IEEE, IEEE, 2009: pp. 218–223.

[38] A.A. Rudat, HEXANE: architecting manned space exploration missions beyond low-earth orbit, (2013).

[39] H. McManus, J.M. Warmkessel, Creating Advanced Architecture for Space Systems: Emergent Lessons from New Processes, J. Spacecr. Rockets. 41 (2004) 69–74.

[40] K. Haris, C.H. Dagli, Adaptive reconfiguration of complex system architecture, in: Procedia Comput. Sci., 2011: pp. 147–152.

[41] M. Lou Maher, J. Poon, Modeling Design Exploration as Co-Evolution, Comput. Civ. Infrastruct. Eng. 11 (1996) 195–209.

[42] D.C. Arney, Rule-based graph theory to enable exploration of the space system architecture design space, Georgia Institute of Technology, 2012.

[43] D.C. Arney, A.W. Wilhite, Modeling Space System Architectures with Graph Theory, J. Spacecr. Rockets. 51 (2014) 1413–1429.

[44] R. Smaling, O. De Weck, Assessing risks and opportunities of technology infusion in system design, Syst. Eng. 10 (2007) 1–25.

[45] J.A. Battat, B. Cameron, A. Rudat, E.F. Crawley, Technology Decisions Under Architectural Uncertainty: Informing Investment Decisions Through Tradespace Exploration, J. Spacecr. Rockets. 51 (2014) 523–532.

[46] T.R. Browning, S.D. Eppinger, Modeling impacts of process architecture on cost and schedule risk in product development, IEEE Trans. Eng. Manag. 49 (2002) 428–442.

[47] E.W. Dijkstra, A note on two problems in connexion with graphs, Numer. Math. 1 (1959) 269–271.

[48] E.F. Hill, Jess in Action: Java Rule-Based Systems, 2003.

[49] D. Selva, E.F. Crawley, VASSAR: Value assessment of system architectures using rules, in: IEEE Aerosp. Conf. Proc., 2013.

[50] D. Selva Valero, Rule-based system architecting of Earth observation satellite systems, (2012).

[51] NASA Goddard Space Flight Center: Tracking and Data Relay Satellite, (n.d.). http://tdrs.gsfc.nasa.gov/ (accessed November 03, 2014).

[52] M.A. Walton, D. Hastings, Applications of Uncertainty Analysis Applied to Architecture Selection of Satellite Systems, J. Spacecr. Rockets. 41 (2004) 75–84.

[53] J. Teles, M. V Samii, C.E. Doll, Overview of TDRSS, Adv. Sp. Res. 16 (1995) 67–76.

[54] M. Andraschko, J. Antol, S. Horan, D. Neil, Commercially hosted government payloads: Lessons from recent programs, in: IEEE Aerosp. Conf. Proc., 2011.

[55] M. Andraschko, J. Antol, R. Baize, S. Horan, D. Neil, P. Rinsland, et al., The potential for hosted payloads at NASA, in: IEEE Aerosp. Conf. Proc., 2012.