Application of Value-Centric Design to Space Architectures: The Case of Fractionated Spacecraft*

Owen Brown†
Defense Advanced Research Projects Agency

Paul Eremenko‡
Booz Allen Hamilton

Recently, the superiority of value-centric (in contrast to cost-centric) design methodologies has been much touted in the literature, including by the present authors. Here, we describe a specific application of value-based methods to the design of space architectures. We suggest that a family of architectures which we term "fractionated," i.e., with mission and support functionality distributed across multiple wirelessly-interacting spacecraft modules flown in cluster orbits, offers a superior value proposition over conventional "monolithic" satellites. We describe the cost and value drivers that differentiate the potential architectural solutions and provide two notional approaches to quantifying the net value and risk (which we treat as the variance in net value) in a manner that supports performing architectural trades. We conclude with some thoughts on the incorporation of such value-based methods into the existing systems engineering and government procurement processes.

INTRODUCTION TO FRACTIONATION

Several years ago—precipitated by the question: if the connectivity between various components of a satellite is made wireless, do all components need to be housed in the same structure?—we proposed the concept of fractionated spacecraft.[1] Taken to its logical extreme, a spacecraft could be decomposed—fractionated, in our parlance—into a plurality of modules flown in cluster orbits, each corresponding to a particular subsystem or function needed for the mission.

So, for instance, one could envision a computation module, a data handling module, and a ground telemetry link interconnected among themselves

---

† Tactical Technology Office. E-mail: owen.brown@darpa.mil.

‡ Science & Technology Practice. E-mail: eremenko@alum.mit.edu.

[1] Brown, O. & Eremenko, P., "The Value Proposition for Fractionated Space Architectures," AIAA-2006-7506, *AIAA Space 2006*, San Jose, CA (2006).

AIAA-2008-7869

and with the rest of the spacecraft through a wireless data network. Power generation can likewise be concentrated on a single spacecraft module and disseminated wirelessly to the rest of the cluster through a variety of power beaming technologies. One can also imagine a single navigation module responsible for inertial determination of its attitude and position; the rest of the modules could derive their inertial state by sensing only their relative state with respect to the navigation module, or conversely, the navigation module could determine the attitude of all other modules relative to itself. Similarly, inertial stationkeeping can be centralized on a single module equipped with thrusters. The rest of the cluster could perform relative stationkeeping through, for instance, electromagnetic force and torque transfer. In addition to all these free-flying bus elements, or resource modules, it is not difficult to conceive flying each payload, e.g., a sensor, transponder, etc., on an independent payload module relying on the rest of the cluster for its resource needs. We term this entire construct heterogeneous fractionation: the decomposition of a space system into wirelessly-interacting dissimilar functional elements.

Homogeneous fractionation, i.e., decomposition into identical or functionally similar modules, is also possible, of course. This could be a means, for instance, for effecting subsystem redundancy by flying more than one module corresponding to a particular function in the cluster. Or, it could simply be a means of decoupling multiple payloads. In principle, many satellite constellations today, e.g., Iridium, Globalstar, GPS, etc. can be construed as homogeneously fractionated space systems.

Both heterogeneous and homogeneous fractionation offer compelling vision for the future of space systems. One can envision a cosmic future where a virtual global "space bus" infrastructure supports an assortment of individual payloads—ranging from bare sensors that receive their data, power, navigation, and stationkeeping externally, to clouds of "pixie dust" acting as distributed sensor arrays or communications relays—that need only be hurled skyward. The marginal cost of space missions could be driven to levels of widespread affordability. Realistically, a near-term "optimal" fractionated architecture likely consists of a mix of heterogeneous and homogeneous fractionation. We return to the question of what we mean by an "optimal" fractionated architecture shortly.

SYSTEM F6

In early 2008, the Defense Advanced Research Project Agency (DARPA) embarked on a program, entitled System F6, to develop, mature, and demonstrate on orbit the key enablers of spacecraft fractionation.[2] Central to the F6 program is *the network*. As envisioned in the F6 demonstrator, a fractionated

[2] "DARPA Awards Contracts for Fractionated Spacecraft Program," Press Release, Defense Advanced Research Projects Agency, Feb. 26, 2008, available at http://www.darpa.mil/body/news/2008/F6.pdf

spacecraft[3] consists of a multitude of network nodes distributed across a cluster of spacecraft modules. One of the key enabling technologies being developed under this effort is a component appliqué—a universal adapter of sorts—that provides an interface for existing off-the-shelf spacecraft components and enables each one to become a network-addressable and network-accessible device.

The F6 effort also seeks to develop a set of physical and protocol network layers that are appropriate for a spacecraft cluster environment. The problem is analogous to DARPA's effort to develop ARPANET over two decades ago, but with notable complications. Many spacecraft devices (henceforth referred to as nodes) require real-time or at least near real-time communications between each other since they are frequently interrogated or actuated by closed-loop control systems with stringent bandwidth requirements. The changing cluster geometry, however, makes it exceedingly difficult to guarantee link latency. Other problems specific to the space environment also arise, including Doppler effects, occlusion, and of course stringent security and information assurance requirements.

One of the key aspects of the F6 network development effort is the desire to develop and promulgate open network interface standards and specifications to enable future generations of government and commercial systems to use them without encountering proprietary or other restrictions. Some of the F6 performers, in fact, are using an open-source model for their software development efforts.

Beyond the physical and protocol network layers, the F6 program is developing a real-time middleware layer that enables distributed computing and resource sharing across the nodes of the fractionated spacecraft network. Each node exposes its resources to the network, and any node can, in principle, utilize those resources subject to a prioritization scheme. Thus, for instance, a spacecraft module whose computing node fails, can utilize spare capacity of a computing node resident elsewhere on the network (e.g., on another module, or even on the ground if link latency and availability permits).

The F6 effort is exploring a variety of other potentially enabling technologies and technology off-ramps. Wireless power transmission, for instance, can be a key enabler if reasonable efficiencies can be attained. Electromagnetic formation flying can also potentially yield great payoff for certain mission classes. In addition to re-architecting monolithic satellites as fractionated clusters from the outset, DARPA is exploring the possibility of equipping near-term traditional monolithic spacecraft with "fractionation

[3] Note that in light of the preceding discussion the notion of a single fractionated "spacecraft" can get fuzzy. We prefer to talk about systems rather than individual spacecraft. This is because a single set of infrastructure modules can support multiple payloads and multiple missions. Furthermore, these can change over time, merge with other clusters, and otherwise be reconfigured. Where a single spacecraft ends and another begins ceases to be a meaningful inquiry.

technologies" that would enable the replacement or augmentation of certain functions by deploying additional modules into cluster formations throughout the life of the monolithic spacecraft. This would provide an incremental, evolutionary approach to effecting fractionated capabilities on orbit.

The architecture being developed under the F6 program is fundamentally mission agnostic. It is being designed with the explicit flexibility to support a wide array of potential mission concepts and payloads (perhaps even simultaneously). We believe that the architectural paradigm will be an appealing option for a variety of mission areas. The ultimate objective of the F6 program is to demonstate on orbit one or more such missions with realistic warfighter utility.

Perhaps the most critical architectural problem being addressed by the F6 program is the question of the optimal allocation of nodes to modules; this, after all, is what differentiates fractionated from monolithic systems in the first place. A methodology for answering questions such as "how many modules should a particular spacecraft be broken down into?" and "how should capability and components be distributed across those modules?" must be developed to bring fractionation from the realm of being an ad hoc exercise into the fold of rigorous systems engineering. So what are the appropriate criteria for accomplishing this allocation? Minimum cost? Minimum mass? We maintain—in a later section of this essay—that both of these are misleading metrics for evaluating the merits of a design.[4] Instead, we answer this question by introducing net lifecycle value and net lifecycle risk as the objective metrics in this architectural optimization problem. But before we discuss how these value and risk metrics are computed, we consider the attributes of fractionated systems and how they differ from their monolithic counterparts. These attributes are the ultimate source of the enhanced value and reduced risk offered by fractionated architectures.

ATTRIBUTES OF FRACTIONATED SPACECRAFT

In order to define what we mean by an "optimal" fractionated architecture, it behooves us to discuss some of the attributes of fractionated systems generally. Certainly fractionated space systems enable a whole array of new missions unattainable with monolithic satellites. These include, for instance, distributed aperture sensing and observation, interferometry missions, and missions requiring satellites beyond the capacity of the largest single existing launch vehicle. This, however, is not the *raison d'être* for fractionation. We posit that fractionated architectures are an improved paradigm for the implementation of many, if not most, existing space missions.

---

[4] These arguments are also explored at length in Saleh, J., "Flawed Metrics: Satellite Cost per Transponder and Cost per Day," *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 44, No. 1, pp. 147-156 (2008).

Cost

In demonstrating the basis for this assertion, we first turn to the cost differences between traditional monolithic spacecraft and comparable fractionated versions thereof. Table 1, below, summarizes many of the elements of lifecycle costs that are likely to differ between fractionated and monolithic architectures, given the same mission. The most notable cost penalty for fractionated systems is, of course, the "overhead" associated with wireless connectivity between modules (i.e., the transmitter and receiver equipment for data, power, etc.) and the necessary duplication of some components (e.g., thermal management, structure, etc.) across all modules.

Table 1

<table>
<tr><th colspan="4">LIFECYCLE COSTS OF FRACTIONATED ARCHITECTURES</th></tr>
<tr><th><u>Phase</u></th><th><u>Term</u></th><th><u>Definition</u></th><th><u>Favors</u></th></tr>
<tr><td>Development</td><td rowspan="2">Requirements decoupling</td><td rowspan="2">Pointing and jitter requirements for sensors/payloads are isolated to individual modules reducing flywheel mass and IA&T costs
Systems engineering is at the individual module level reducing IA&T costs
Design life and reliability can be treated as an independent variable that can be optimized separately for each module based on cost of reliability for various components</td><td rowspan="2">Fractionated</td></tr>
<tr><td rowspan="4">Fabrication</td></tr>
<tr><td>Fractionation overhead</td><td>Mass penalty due to wireless transceivers and some component duplication on each module</td><td>Monolith</td></tr>
<tr><td>Production learning</td><td>Learning curve effects due to the fabrication of multiple modules of the same type</td><td>Fractionated</td></tr>
<tr><td>Commoditization</td><td>Ability to transition to assembly line or other mass production fabrication techniques</td><td>Fractionated</td></tr>
<tr><td>Launch</td><td>Small launch vehicles</td><td>Fractionated spacecraft favor (but do not require) the use of small, responsive launch vehicles; the economic efficiency of these vehicles is yet unknown</td><td>Ambiguous</td></tr>
<tr><td>Operations</td><td>Operating complexity</td><td>Operating a cluster of spacecraft versus a single spacecraft may increase complexity; this may be moot with autonomous cluster ops</td><td>Ambiguous</td></tr>
<tr><td>All</td><td>Decoupling of classification</td><td>Stringent classification requirements can be imposed on individual modules – which would communicate to the ground via VPN – and not the entire architecture</td><td>Fractionated</td></tr>
</table>

Several cost elements clearly favor fractionated systems. Likely most significant among these is the decoupling of individual module requirements. The effect is threefold. First, stringent payload pointing, jitter, isolation, and other requirements need only be imposed on the module containing that particular payload—and not on the rest of the spacecraft (including on other payloads). This can lead to a significant downsizing of attitude control systems; one can even conceive that a computation or data storage module could be gravity- or spin-stabilized (assuming omnidirectional wireless links and full-body solar cell coverage of course).

Second, the integration, assembly, and test (IA&T) problem for several smaller modules interconnected only through a strictly-defined wireless interface is a significantly simpler one than for a single larger spacecraft; this has both technical and programmatic implications. Technically, the larger spacecraft would include most of the contents of the individual smaller modules, but their interactions—in addition to the desirable data and power interfaces—would also include undesirable ones such as vibrations, thermal effects, and electromagnetic interference. Programmatically, the implication is that a single module or payload element does not lie in the critical path of the program. Thus, should a payload module be delayed for one reason or another, the entire program is not delayed.

And third, but certainly not least, is the decoupling of design reliability or design lifetime of the mission from that of individual modules. In other words, each module can be designed for some optimal design lifetime; multiple generations may be deployed throughout the mission life of the system, and this may differ according to the appropriate obsolescence timescale across modules. The ability to tailor design lifetime across the system can potentially yield significant cost savings. The cost of attaining a given level of reliability varies dramatically between different subsystems and components. In a fractionated spacecraft, each subsystem or group of components can be designed for some optimal design lifetime specific to that subsystem or component. Hardware with similar design lifetimes can be aggregated on modules thus creating some short- and some long-lived modules. Each can be replaced on its own time frame. Such a design—effectively making reliability/design lifetime an independent, tradeable design variable—is likely to cost significantly less than an integrated monolithic system built to some exogenously-specified (and frequently arbitrary) lifetime requirement.

Another facet of decoupling is the separation of security environments between modules. In principle, for instance, one can conceive of an architecture where a sensitive payload is developed and launched separately from the other modules (which may include other, non-sensitive payloads). Once on orbit, the sensitive payload could downlink to an appropriate ground facility over a shared link utilizing a properly-accredited virtual private network (VPN) or comparable technology. Confining security and classification requirements to a single module can result in a tremendous cost savings in the development of other payload and resource modules.

The standardization of resource modules across multiple fractionated systems opens the door to (hitherto largely elusive in the space industry) volume-based cost reductions. One such effect is the learning curve.[5] The learning curve is predicated on the notion that individual and organizational learning in the course of fabricating multiple instances of the same item reduces recurring labor costs. Somewhat distinct from the learning curve, which is an artifact of experience, are the potential cost savings due to the commoditization of production, which is based on a fundamental change in fabrication processes.

While learning curve effects can be significant over the production of just a handful of items, commoditization requires volumes that have never hitherto been attained in the space industry.[6] Commoditization typically involves a significant reduction in the recurring labor content of production at the expense of capital investment in significant automation and a reduction in customization. The automotive and personal computer industries are excellent examples of commoditized manufacturing. Whether fractionation will provide sufficient volumes to warrant a transition to commoditized production is largely an unresolved question, predicated on the breadth of its adoption. Should widespread adoption take place, however, it would be as profound and transformative event for the affordability of space as the Ford Model T was for the affordability of automobiles.

Also somewhat ambiguous in terms of its overall impact on the cost-based merits of fractionation is the situation with operating costs. Although conventional wisdom suggests that operating a cluster of spacecraft would be a more complex and costly undertaking than operating a single satellite, it is precisely the complexity of operating a cluster and the need to ensure collision-free trajectory planning that is likely to force a significantly increased degree of autonomy in fractionated spacecraft operations. Autonomous cluster operations is an explicit objective of the F6 development effort. Furthermore, the construct of treating ground assets not as an entirely separate "ground segment," but simply as another node or set of nodes on the network promises to simplify them and could reduce cost.

Finally, it is worthy of note that launch and operating costs are two areas where fractionated architectures can either suffer in comparison to their monolithic counterparts, or—alternatively—are areas of great opportunity for them. Fractionated modules certainly lend themselves to launches using

---

[5] Wright, T.P., "Factors Affecting the Cost of Airplanes," *Journal of Aeronautical Sciences*, Vol. 3, No. 4, pp. 122-128 (1936).

[6] The Iridium, Globalstar, and Orbcomm constellations are, perhaps, the highest-volume space systems to date. And while the learning curve effects there were significant, anecdotal evidence suggests that no fundamental transformation of Lockheed Martin's or Loral's production processes was warranted by quantities in the dozens.

emerging small, responsive vehicles such as the SpaceX Falcon 1 (to LEO) and the Falcon 9 (to GEO). The responsiveness and flexibility of such vehicles is certainly a potential advantage for fractionated satellites. Their promise of reduced cost per kilogram to orbit, however, is widely regarded with skepticism. This skepticism is frequently premised on optimistic launch volume assumptions, which fractionated architectures could go a long way toward realizing. Should these cost savings materialize, they will significantly strengthen the case for fractionated systems based on cost alone. If they do not, however, it is important to note that fractionated modules can be launched in bundles on conventional launch vehicles or, if the orbitology is favorable, can be individually launched as secondary payloads on unrelated missions.

So where does that leave us? There are some certain cost penalties associated with fractionation. There are also undisputable advantages. The overall cost proposition is ambiguous until further development and operational testing can be accomplished. But we can say, with a reasonable degree of certainty, that if the fractionated paradigm gains widespread acceptance, it is likely to be transformative even on a cost basis alone.[7] In order to gain widespread acceptance in the first place, of course, it must have other demonstrable advantages. Thus, we turn our attention to the flexibility and robustness of fractionated spacecraft.

Flexibility

We define flexibility as the ability of a system to change on demand. A flexible system is one that offers its owner, developer, operator, or other interested party—whom we shall term the stakeholder—options to alter the system in some way or leave it unchanged. Why is flexibility a good thing? In a perfect and static—albeit rather boring—world, flexibility would be worthless. The world we live in, however, is uncertain. It is impossible to design a system from the outset that anticipates every possible occurrence throughout its lifetime. So we design systems that possess attributes such as scalability, upgradeability, maintainability, and adaptability—all manifestations of flexibility and all are a way of delaying architectural decisions about the system into the future. A list of these "-ilities" applicable to fractionated spacecraft is provided in Table 2, below, with our working definitions for each.

[7] One subject which we touched upon in the introduction, but have not explicitly enumerated in Table 1 or the subsequent discussion is that of amortizing costs of fractionated systems across multiple missions through shared infrastructure modules. This fundamentally alters the economics of space systems and makes it virtually impossible to compare fractionated to monolithic systems. For that reason, the subject warrants a separate discussion in a future essay and is not addressed in any detail here.

Table 2

FLEXIBILITY – ability of a system to change on demand		
Term	*Synonyms*	*Definition*
▸ Scalability	▸ Incremental deployment	Ability to add components or capability to a system throughout its lifetime
▸ Evolvability	▸ Upgradeability	Ability to replace components due to technology obsolescence
▸ Maintainability	▸ Sustainability	Ability to replace components that have failed or are near end of life
▸ Adaptability	▸ Reconfigurability ▸ Versatility	Ability to reconfigure existing system functionality to meet new needs or circumstances

A useful construct for thinking about flexibility and its various manifestations is options. A flexible system is one that gives its stakeholder an assortment of options that he can choose to exercise throughout the life of the system. Scalability is the option to add modules to a system. Evolvability and maintainability are options to remove modules and add others (or vice versa). And adaptability is the option to somehow modify the system. There is, of course, always the fifth option at any given time, which is to do nothing.

We suggest that fractionated systems are significantly more flexible than monolithic ones because they give the stakeholder a wider array of options throughout the lifetime of the system, and these options are easier for the stakeholder to exercise. So, for instance, a fractionated system can be easily scaled up in capability throughout its lifetime through the incremental addition of individual modules in response to changing demand. Similarly, the capability of the system could be scaled down during the acquisition, design, and build cycle simply by completing fewer payload elements (and potentially deploying fewer resource/infrastructure modules) than originally planned. The only way to scale up the capability of a traditional monolithic satellite is to launch another one. Scaling down a monolithic system is also a difficult proposition, which would have huge technical implications if such a decision were made beyond the initial design stages of the program. Likewise, as particular technologies become obsolete throughout the lifetime of the system, individual modules housing those components in a fractionated architecture can be replaced with newer ones. In a monolithic system, the only option is to launch a new one. With maintainability, again, failed modules in a fractionated cluster can be individually replaced; a monolith, in the absence of on-orbit servicing, goes out of service and necessitates the deployment of a new one. And fractionated systems are easily adaptable to unforeseen circumstances. So, for instance, if a fractionated spacecraft module loses its on-board processor, it would be able to utilize one on another module. A traditional monolithic spacecraft in this situation—as countless real life examples have shown over the years—would be doomed.

Robustness

While flexibility refers to the range of options that a stakeholder has to alter the system in response to unforeseen circumstances, robustness is the intrinsic ability of a system to maintain functionality in response to unforeseen circumstances. The circumstances in this latter case tend to be internal or external perturbations that impact the system's ability to perform its nominal mission. It is important to note the distinction. Flexibility requires a stakeholder to take some action to alter the system in response to an event. Robustness refers to the system's own intrinsic response to events.

Table 3

ROBUSTNESS – retention of functionality in response to an internal or external stimulus		
Term	*Synonyms*	*Definition*
‣ Reliability	--	Ability of a system to function under nominal conditions
‣ Survivability	--	Ability of a system to function off-nominal or unanticipated conditions
‣ Resilience to fragility	‣ Robustness to fragility	The (in)frequency with which a system succumbs to unmodeled, catastrophic, cascading failures
‣ Fault tolerance	‣ Graceful degradation	Gradual loss of system functionality due to one or more failures

We posit that fractionated systems will be significantly more robust than traditional monolithic satellites. To see why, it is helpful to construct a taxonomy of failure types to further elucidate the forms that robustness can take. Table 3, above, provides a summary of our taxonomic scheme. The most familiar form of robustness is reliability. Reliability reflects the expected design life of a system, or, for a fixed design life, the probability with which a system will last that length of time. As a design variable, reliability reflects the designer's certainty that a system will withstand some nominal operating environment and nominal set of perturbations throughout its life. There are two principal tools that a designer has for effecting reliability in systems: redundancy and qualification. Although redundancy is accomplished somewhat differently in fractionated architectures—by duplicating a module delivering a particular functionality—this should not impact overall reliability. Neither is the qualification of fractionated components significantly different than that for monolithic spacecraft. Thus, we see no fundamental reason why the reliability of fractionated systems would differ significantly from that of monolithic ones.[8]

More interesting in the present context is another facet of robustness—survivability. In contrast to reliability, survivability reflects the system's response to future events or conditions that were unanticipated (or at least not specifically planned for) in the system design. These events or conditions are typically ones that are severe and relatively improbable, so that it is impractical to design the system to be reliable to them. For instance, while a spacecraft is typically designed to be reliable to temperature fluctuations and radiation, accommodation of a meteorite impact or an anti-satellite attack is typically not an explicit design constraint; a spacecraft that does not fail in these circumstances would be termed survivable.

Fractionated spacecraft are likely to offer enhanced survivability. We proffer two main reasons to justify this assumption. First, fractionated systems have a spatial distribution of components over a larger volume—target spreading, in essence. Second, the aggregation of several objects which can autonomously relocate themselves into new geometries would be certain to confuse any autonomous tracking system.

A third form of robustness is associated with fragility. Fragility is the tendency of complex systems to experience infrequent and improbable, but highly catastrophic failures resulting from multiple cascading malfunctions. These have been termed "black swans,"[9] and have been attributed to transitions of dynamical systems to chaotic regimes. Most major catastrophes afflicting engineering systems are results of fragilities—interactions of multiple unmodeled failure modes. Modern examples include Apollo, Challenger, Chernobyl, and Columbia.

The propensity of a system to experience such "black swan" failure events appears to be related to the complexity of the system.[10] Fractionated spacecraft, by decomposing what might be a highly complex space system into smaller individual modules linked only via strictly enforced wireless interfaces, reduce the effective complexity of the overall system to that of any individual module. Fractionation is, in essence, analogous to the principal mechanism for managing complexity in software systems: abstraction, i.e., defining and enforcing interfaces and otherwise separating the whole into discrete, modular parts which need not know anything about each other except at the interfaces. (Interfaces in software systems are, of course, not electromechanical in nature.)

And finally, one might note that fault tolerance, which appears in Table 3, is not part of our tripartite taxonomy for the forms or types of robustness. It is

---

[8] Though we certainly do not dispute that the addition of wireless links and other fractionation-related subsystems will introduce new failure modes into the analysis.

[9] Taleb, N., *The Black Swan: Impact of the Highly Improbable*, Random House (2007).

[10] Carlson, J.M. & Doyle, J., "Complexity and Robustness," *Proceedings of the National Academy of Sciences*, Vol. 99, Suppl. 1, pp. 2538-2545 (2002).

simply an attribute that describes how a system behaves in response to the various stimuli or conditions if the systems turns out not be robust to that particular stimulus or condition. As a gross generalization, monolithic systems are prone to step changes in functionality, while appropriately-designed fractionated systems should degrade capability gracefully or incrementally.

Industrial Base Effects

While not attributes of specific fractionated architectures, we believe that the paradigm as a whole would have some interesting—and largely positive—effects on the structure of the aerospace industry and the defense industrial base. We identify several specific effects that are likely to result from the proliferation of fractionated architectures; these are summarized in Table 4, below.

Table 4

INDUSTRIAL BASE EFFECTS OF FRACTIONATED ARCHITECTURES	
Term	*Definition*
Reduced barrier to entry	The splitting of a single large spacecraft into smaller pieces permits players outside of the large traditional primes to participate in development and fabrication
Increased number of competitive opportunities	With a standard open interface among modules, a spacecraft mission or program can be split up among multiple contractors raising today's very small number of competitive opportunities
Volume for responsive spacelift	Launching fractionated spacecraft modules may be a significant market for small launch vehicle payloads
Improved operator NPV	If fractionated spacecraft really offer an enhanced value proposition over traditional monolithic ones, this will enhance the NPV and shareholder returns for commercial operators

Note that most of the industrial based effects discussed here are positive externalities of fractionated systems. They have no direct cost, value, or risk impact from architecture to architecture or from mission to mission. In the aggregate, however, we suggest that the proliferation of the fractionated paradigm would significantly enhance the economics and market structure of the aerospace industry.

THE VALUE PARADIGM

We have noted a variety of attributes which might differentiate fractionated architectures from monolithic ones and may, in fact, differentiate one fractionated architecture from another. But how does one make decisions on the basis of these characteristics? How much should one be willing to pay for a more flexible or more robust system? Is it possible to compare to systems on the basis of their relative flexibility or robustness?

The answer to these questions lies in the notion of value. Value is a measure—wholly apart from cost—that reflects the utility of a particular system to its owner or operator (whom we term the stakeholder, as before). We,[11] and many others,[12] have previously extolled the virtues of value-based metrics for system design and there is little need to rehash all of the arguments here. It will suffice to note that outside the aerospace industry, the value metric (in the form of net present value, NPV) is universally applied to evaluating the merits of various systems, projects, products, etc. The history of why value is not used as a criterion for the design and acquisition of aerospace systems is reserved for discussion later in this essay. For the purposes of our discussion here, we use the term value to refer to the total lifecycle value delivered by a system. Net value, then, is the total lifecycle value minus total lifecycle cost.[13]

The primary source of value for any system—fractionated, monolithic, or otherwise—is its mission; this is, after all, why we bother developing the system in the first place. While the most significant source, however, it is not the one that is of greatest interest to us here. With the exception of several new missions enabled by spatial distribution of payload components discussed *supra*, we postulate that one fractionated architecture will deliver the same mission capability as another, differently fractionated architecture, which is the same as the mission capability of a corresponding monolithic system. The baseline mission value is, therefore, identical for all of them.[14]

What differentiates fractionated systems from one another and from monolithic ones is the value of the derivative attributes—flexibility, robustness, etc.—painstakingly enumerated in the preceding section. These attributes are derivative because their value is derived from the value of the underlying mission. In the absence of an underlying value stream, the value of flexibility and

---

[11] See, e.g., note 1.

[12] See, e.g., Saleh, J. et al., "Flexibility and the Value of On-Orbit Servicing: New Customer-Centric Perspective," *Journal of Spacecraft and Rockets*, Vol. 40, No. 2, pp. 279-291 (2003); Nilchiani, R. & Hastings, D., "Measuring the Value of Flexibility in Space Systems: A Six-Element Framework," *Systems Engineering*, Vol. 10, No. 1, pp. 26-44 (2007).

[13] Whether we treat the value and cost as discounted to the present is immaterial to our discussion here. A quantitative treatment would require the determination of an appropriate discount rate. Interestingly, while economically-savvy branches of the federal government routinely rely on Treasury Note rates for discounting future cash flows (see, e.g., Nussle, J., "2008 Discount Rates for OMB Circular No. A-94," OMB Memo M-08-08, White House Office of Management and Budget, Jan. 14, 2008), Congress and the Department of Defense use an effective discount rate of zero in procurement planning.

[14] Note that this assumption lies in great contrast to some contemporary thought that small spacecraft can not provide capability on par with large ones." Here we suggest otherwise, i.e., that small spacecraft together can create value (based solely on static capability; irrespective of the flexibility and robustness that are the major focus of the rest of this essay) at least equal to that of large monoliths. Of course, we maintain that clever combinations of small spacecraft—what we call fractionated clusters—offer *more* value than traditional monoliths.

robustness is precisely zero. But, as we postulate some nominal mission capability for all systems, it becomes meaningful to compare the architectures on the basis of how flexible, robust, etc. they are. Furthermore, if the value of these attributes is quantified in units that are commensurable with cost, meaningful cost-value tradeoffs for various architectural options can be made, and programmatic risk for various architectures can be traded against their expected net value.

LIFECYCLE UNCERTAINTIES GIVING RISE TO THE VALUE OF FLEXIBILITY & ROBUSTNESS

As we noted in our taxonomic discussion of flexibility and robustness, both of these attributes become important only in the presence of perturbing events or uncertainties. Flexibility and robustness are both worthless in a world that is static, certain, and *a priori* deterministic. For better or worse, we do not live in such a (rather dull) world. In Table 5, below, we enumerate some of the lifecycle uncertainties or perturbations that affect space systems. [15]

Table 5

LIFECYCLE PERTURBATIONS		
Perturbation	*Definition*	*Example(s)*
▸ Development problem	A problem arising during the development phase leading to a workstream delay or budget overrun	▸ Failure to mature a technology ▸ Payload damaged during test
▸ Funding fluctuation	Volatility (usually, but not necessarily, a reduction) in the available budget	▸ Funding cut ▸ Funding increase
▸ Requirements change	A change in program objectives subsequent to their initial definition and commencement of development	▸ New threat identified during design ▸ New stakeholder enters program
▸ Launch failure	Failure of launch vehicle payload to reach desired orbit	▸ Explosion of launch vehicle
▸ On-orbit failure	Component failure on orbit due to internal or external event	▸ Processor failure ▸ Micrometeorite impact
▸ Demand increase	Demand for mission service exceeds original expectations during development	▸ Increased bandwidth needed in theater ▸ New TV technology requires additional bandwidth
▸ Obsolescence	Emergence of new technologies favor newly-designed components over existing assets	▸ Emergence of novel antenna or sensor technology ▸ Moore's Law

[15] A more extensive discussion of the uncertainty environment over the life of fractionated—or any other—satellites can be found in Brown, O. & Eremenko, P., "Fractionated Architectures: A Vision for Responsive Space," AIAA-RS-2006-1002, 4th *AIAA Responsive Space Conference*, Los Angeles, CA (2006).

Each of these perturbations can be modeled as a probabilistic event with an appropriate distribution. The distribution can be based on historical data (e.g., for launch vehicle failures[16] and development overruns[17]) or on first-principles or quasi-empirical analyses for the mechanism underlying the perturbation (e.g., for component operational failures[18] and obsolescence[19]).

METHODOLOGIES FOR CALCULATION OF VALUE, COST, AND RISK

Equipped with a lengthy list of cost drivers (e.g., module size, mass, power, design lifetime) and value attributes (e.g., mission performance, flexibility, robustness) that are of interest in the design of fractionated architectures, a nominal assumed mission utility, and a set of perturbations or uncertainties that the system might experience over its mission life, we are now ready to explore the trade space of fractionated systems on the basis of their value, cost, and risk.

In order to conduct meaningful exploration of the architectural trade space for fractionated spacecraft (of which the conventional monolithic satellite is, of course, a special case), a design methodology must be able to take an architecture defined by a relatively small set of tradable design parameters and compute a metric on the basis of which the architecture may be compared to others. Such a tool can be used to compare the relative merits of architectures that are outputs of other design methods, or it can be used iteratively to find optimal top-level architectural parameters for specific missions. We revisit this question of how a methodology of this type might be incorporated into the systems engineering process in a subsequent section.

[16] Guikema, S. & Paté-Cornell, N., "Bayesian Analysis of Launch Vehicle Success Rates," *Journal of Spacecraft and Rockets*, Vol. 41, No. 1, pp. 93-102 (2004).

[17] Arena, M. et al., "Historical Cost Growth of Completed Weapon System Programs," TR-343-AF, RAND Corporation, Santa Monica, CA (2006).

[18] Sullivan, B. & Akin, D., "A Survey of Serviceable Spacecraft Failures," AIAA-2001-4540, *AIAA Space 2001*, Albuquerque, NM (2001).

[19] Moore, G., "Cramming More Components onto Integrated Circuits," *Electronics*, Vol. 38, No. 8 (1965).

A design parameter space that is useful for defining a fractionated spacecraft configuration might include the following (Table 6):

Table 6

ARCHITECTURAL DESIGN PARAMETERS	
Term	*Definition*
Degree of fractionation	The number of physical spacecraft modules comprising a fractionated spacecraft cluster.
Distribution of nodes across modules	The distribution of functionality across each fractionated module. In a network-centric view, this would be the assignment of network nodes to physical spacecraft modules.
Reliability/design lifetime of each module	The design lifetime/reliability of each module, indepent of the overall mission lifetime. In a maintainable architecture a uniform reliability distribution across all components is unlikely to be the optimal solution since the "cost of reliability" and the rate of obsolescence varies widely from subsystem to subsystem
Modes of connectivity between modules	The types of links for data, power, and force/torque transfer between modules and between components and nodes within a single module.
Choice of launch opportunity for each module	Launch vehicle and launch event in which a particular module will be deployed to orbit. All modules could be launched together, each one individually, or most likely, a hybrid solution is optimal.

It is noteworthy that this list of design parameters is rather short. We have focused on selecting those that are either unique to fractionated systems or ones that are likely to differentiate fractionated systems amongst each other or from monolithic ones. A much more extensive set of parameters is needed, of course, to define a design. The goal of the methodological discussion on comparing the value, cost, and risk of various architectures which follows, however, is not to develop a multi-disciplinary optimization tool that yields optimized detailed system designs. Such a tool would be theoretically feasible, of course, with value, cost, and risk serving as "fitness functions" for evaluating the optimality of a particular design. The problems concomitant to the development of such an ambitious optimization methodology, however, would obfuscate the essence, novelty, and elegance of the fundamental value-based architecting paradigm.[20] And, as will become apparent in our subsequent discussion of systems engineering, traditional (cost- and mass-based) design methods can be applied to fleshing out most of the design of individual fractionated modules without

[20] Unfortunately, it is difficult to develop design rules-of-thumb based on value that are useful at the subsystem or component level. While it may be easy to use rough relationships for estimating cost on the basis of size, weight, and power of subsystems and components (it scales roughly linearly or as a power law with an exponent that is not too far above 1.0), the relationship between value and most concrete design variables that are of interest during detailed design trades is obscure. Hence our focus on architectural design variables that have clear relationships to the lifecycle value streams.

significantly (or at all) compromising the value proposition of the overall architecture. In the meantime, we turn our attention to the goal of demonstrating two techniques for mapping each architecture in a tradespace spanned by a simple design vector to value, cost, and risk metrics that can be used to rank them.

First Sample Methodology

The first approach may be described as static, forward-looking valuation: the total cost and value streams for the architecture are predicted for its entire lifetime, given a particular perturbation environment that may lead to varying degrees of system degradation. Here, lifecycle cost is relatively easy to estimate using traditional cost modeling approaches. Parametric cost estimating relationships (CERs) are particularly well-suited to use in a near-real-time tradespace exploration tool. Unfortunately, traditional CERs are not well calibrated for estimating the costs of fractionated spacecraft modules, which may differ significantly from simply small satellites of comparable size. Newly-developed or adapted CERs may be appropriate. Bottom-up cost build-up can, of course, offer improved fidelity and the only realistic way (at least until a statistically-significant body of fractionated systems exists) to model all the requisite cost effects.

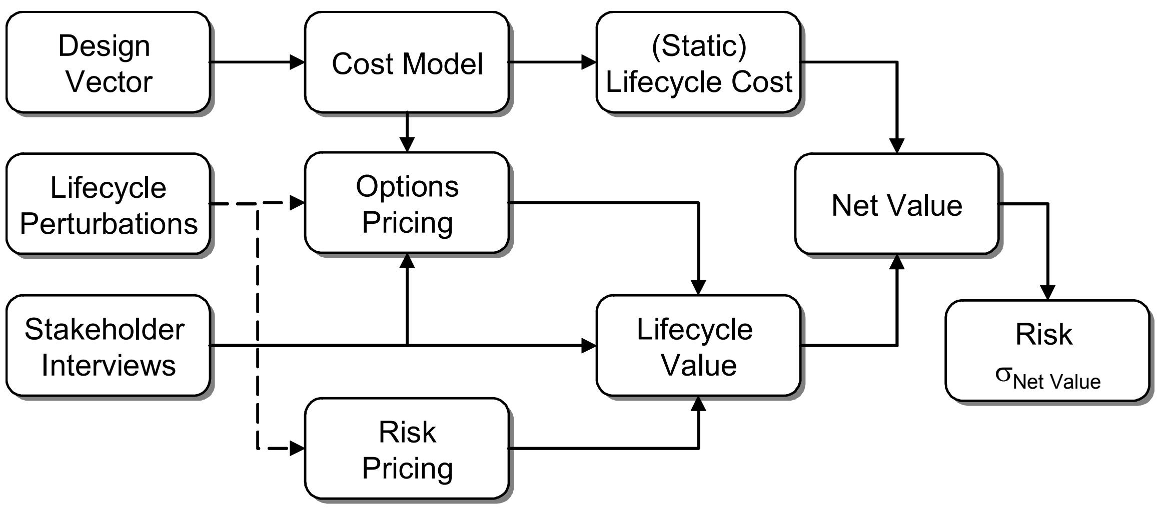

Figure 1: Forward-looking modeling of cost and value streams in a predicted risk environment

Figure 1 depicts, in detail, one hypothetical set of analyses that could constitute a forward-looking valuation given a particular risk environment. The top branch of the process takes a particular design vector (which, of course, can be iterated upon to map an entire design space) and coverts it into an estimate of static lifecycle cost for the architecture represented by this design vector. Additionally, the cost of exercising various options is estimated on the basis of the design vector and cost model. The valuation of these various options—representative of flexibility (see below)—is based on the cost increment associated with option exercise, the value of the underlying mission (as derived, e.g., from stakeholder interviews), and the assumed volatility of the value of the underlying mission resulting from lifecycle perturbations. The value of the "bundle" of options, combined with the risk premium associated with the lifecycle perturbations, are added to the value of the underlying mission to compute the lifecycle value of the architecture. Net value and the variance of net value are the ultimate output metrics.

The value of the mission, as described above, is the value of the service provided by the mission payload and is a function of the various mission performance attributes. For a commercial mission, such as communications, the value may be derived from market pricing based on bandwidth, frequency, coverage, availability, etc. For government missions, where no market exists in which the users' valuation of the service may be reflected in prices, a stakeholder interview approach can be applied. A family of techniques called multi-attribute utility theory may be used to elicit a stakeholder's relative valuation of multiple performance attributes.[21] If one of the attributes can also somehow be valued absolutely (e.g., cost in dollars), then the relative valuations of the other attributes can also be converted to absolute valuations. The selection of representative stakeholders is a significant challenge (e.g., is a user, operator, procurement official, or agency head the "right" stakeholder for a government satellite system?). It must also be noted that the technique is only effective over attributes which are readily quantified and perceived by the stakeholder. Thus, for instance, asking the interviewee to trade off bandwidth and coverage area is reasonable; asking him or her to do the same between bandwidth and flexibility is not. This is an important point. It is typical in the conduct of Analyses of Alternatives and other exercises that perform conventional cost-benefit analyses to (at least qualitatively) use mission capability, flexibility, and robustness (or some of the attributes subsumed under these labels) as independent measures of benefit. This practice is fallacious since, as discussed earlier, flexibility and robustness derive their value (benefit) from the mission capability

It is for this reason that the valuation of flexibility and robustness (and all the constituent attributes encompassed under these rubrics, as previously discussed) requires a different approach. The ability to add, remove, or reconfigure capability on orbit can be modeled as a real option and valued, analogously to financial options, using the Black-Scholes formula.[22] In its

21 The classic overview of multi-attribute utility theory and related techniques is Keeney, R. & Raiffa, H., *Decisions with Multiple Objectives: Preferences and Value Trade-Offs*, Cambridge University Press, Cambridge, UK (1993).

22 An excellent treatise on real options generally is Trigeorgis, L., *Real Options: Managerial Flexibility and Strategy in Resource Allocation*, MIT Press, Cambridge, MA (1996).

simplest form, the Black-Scholes model expresses the value of an option as a function of the price of the underlying asset, the exercise price of the option, and an estimate of the future volatility of the underlying asset. In the real option analogue as applied to fractionated spacecraft, the underlying asset is the present value of future cash flows if the option (to add, remove, or modify modules) is exercised; the exercise price is the cost of adding or removing modules or executing the modification (this may be the opportunity cost of fuel used for an orbit change maneuver, for instance); and the future volatility of the underlying asset is the variance associated with future cash flows contingent on the exercise of the option. This latter volatility measure is an amalgamation of the relevant risk environment. Thus, for instance, the option to replace a failed module on orbit would be valued based on a volatility estimate that incorporates the risks associated with all possible failure causes: component malfunction, environmental factors, antisatellite action, etc. The option to remove a module from the design late in the development process (but before launch), for instance, would be valued on the basis of a volatility estimate that includes the probability of a drop in user demand, a reduction in funding, requirements changes, and technology development problems.

Somewhat analogously, the valuation of robustness can be accomplished through the calculation of a risk premium to be included in the net value calculation. This risk premium is a quantitative reflection of the fact that an architecture which has a lower cost and/or value stream variance over its lifetime is more valuable than one with a higher variance. A useful way of thinking about this effect is through the insurance analogy. Distributing a monolithic architecture across multiple modules—where a single failure has only a small effect on capability and/or replacement cost—is effectively equivalent (in a purely economic sense, of course) to self-insuring, and thereby saving the insurance premium. In the case of the government, which is theoretically a risk-neutral actor, the premium can be viewed as an assurance premium. Thus, if a particular level of assured capability is required, the risk premium is the difference in cost between procuring two different architectures to meet this level of assured capability. In a fractionated system, one might have to buy, say, 20% more modules; in a monolithic system, one might have to procure two whole satellites (100% more) to meet a given assurance level.[23] Risk premiums can be computed using any number of standard actuarial or financial tools.

To summarize this first sample approach, it is essentially a quantitative cost-benefit analysis (where the benefit is quantified in units commensurable with cost) of an architecture that encompasses some non-traditional attributes on top of the usual measures of performance. A tally of the projected lifecycle cost

---

[23] This idea is explored in greater depth in Brown, O., "Reducing Risk of Large Scale Space Systems Using a Modular Architecture," *Space Systems Engineering & Risk Management Symposium*, The Aerospace Corporation, Los Angeles, CA (2004).

and lifecycle value streams in light of a predicted risk environment are taken to present value. Although the various lifecycle options (e.g., to upgrade or replace modules) are included in the valuation, no assumptions are made about actual sustainment of the system over its life. Thus, the system may be permitted to degrade or otherwise change from its initial state over its lifetime.

Second Sample Methodology

The second approach differs from the first primarily in the fact that system capability is not permitted to degrade permanently due to the risk environment. Instead, an operator behavior model responds (or anticipates) events such as failures, and acts to maintain a minimum objective level of capability throughout the system's life. This necessitates a dynamic simulation of the risk environment (rather than *a priori* assumptions about future volatility associated with various perturbations) and resultant impact on the system, but does not require individual accounting of value streams. Instead, the total lifecycle cost needed to maintain capability under uncertainty is a metric that encompasses the value of flexibility and robustness (and all the other "-ilities" subsumed thereunder).[24] Perhaps the most significant advantage of this approach is that an explicit valuation of the mission/payload service is not needed, as it is assumed that all architectures are required to provide the same assured level of capability—thereby eliminating the need for cumbersome stakeholder interviews.

The primary disadvantage is that differences in service availability between architectures are not fully captured. By assuming that the level of mission services is constant, differences between architectures in the rapidity of initial deployment, the mission value derived from partial on-orbit capability derived from incremental deployment, and the rapidity with which replenishment (with either terrestrial or on-orbit backups) in response to on-orbit failures can be accomplished are all ignored. In essence, the difference in the responsiveness of various architectures is not captured by this family of methods.

[24] An illustrative attempt at capturing the value of fractionated spacecraft in comparison to analogous monolithic systems using the lifecycle cost under uncertainty method is described in Brown, O. et al., "System Lifecycle Cost Under Uncertainty as a Design Metric Encompassing the Value of Architectural Flexibility," AIAA-2007-6023, *AIAA Space 2007*, Long Beach, CA (2007).

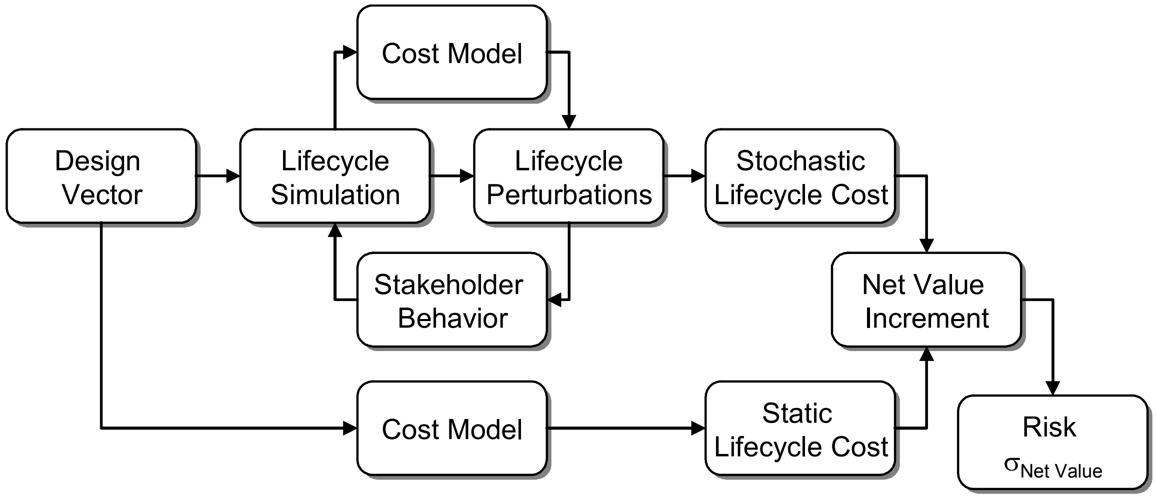

Figure 2: Dynamic modeling of total lifecycle cost to maintain an assured level of capability in a simulated risk environment

Figure 2, above, depicts a functional diagram of a hypothetical implementation of this modeling approach. The top branch of the diagram essentially computes the lifecycle cost of the system accounting for responses (based on the stakeholder behavior model) to various lifecycle perturbations. This is the stochastic lifecycle cost. The bottom branch of the diagram computes the unperturbed (static) cost of the system corresponding to the same lifecycle. The difference between these two cost metrics corresponds to the net value increment associated with this particular architecture. This net value increment corresponds to the value of flexibility and robustness (but not the value of the underlying mission).

To illustrate this approach with an example, let us consider a warfighter need for the assured (to nearly 100%) presence of ten communications transponders in a particular geographic area. We wish to assess a monolithic and a fractionated architecture (with each of ten transponders on an independent spacecraft module) based on their net lifecycle costs. Let's assume a simple lifecycle perturbation environment which includes funding fluctuations, launch failure risk, and on-orbit failure risk. Realistically, stochastic models should be utilized to model all three as random variables. For our purposes, let us suppose simply that funding is halved for both the monolithic and fractionated architectures at some point in the design cycle. The monolithic spacecraft must be re-designed significantly and the cost and schedule impact of this is modeled through non-recurring (NRE) cost estimating relations (CERs). The fractionated architecture simply halves the number of modules (this, of course, assumes homogenous distribution of both payload functionality and bus support functionality) with minimal or no NRE cost impact.

Since no launch vehicle is 100% reliable, and supposing, again, that a launch failure event happens to occur to both the monolithic and fractionated architecture, two monoliths must be built and readied for launch on the assumption that the first launch will fail (but recall that we need ~100% assured capability). While the fractionated architecture will require perhaps one or two additional spare modules to compensate for potential launch failures. This assumes, of course, independent launches for each of the fractionated modules. The launch vehicles are likely to be smaller and more expensive (on a per-unit-mass basis), thereby incurring greater total launch and launch integration costs, though lower total spacecraft fabrication costs.

Now, finally, assuming that the monolithic spacecraft is also not ~100% reliable, our requirement for mission assurance will necessitate a spare (i.e., 100% duplication). The fractionated spacecraft will likewise require a few spare modules, though probably not complete duplication of each one (i.e., somewhat less than 100% duplication). And so the model goes on simulating the lifecycle of each architecture given a particular risk environment. A somewhat comprehensive list of potential risks were listed in a preceding section on lifecycle perturbations. All of these can be modeled as random variables on the basis of various empirical datasets or theoretical models (normal or log-normal distributions are a good start for many).

In these simple examples, the stakeholder's appropriate response to a perturbation is obvious. However, once the possibility of anticipatory strategies is allowed, modeling stakeholder behavior in response to perturbations becomes a rather complex task. Consider, for instance, the question of whether the stakeholder should be assumed to have full knowledge of the statistics of the perturbation environment? Partial knowledge? No knowledge? So, for instance, if nearly 100% mission assurance is required, the operator must likely maintain a fleet of on-orbit spare modules (how big depends on the stakeholder's knowledge and knowledge of the uncertainty in his knowledge of the probability of failures), but perhaps also a fleet of terrestrial spares—depending on fabrication and launch vehicle timelines (and how well, one might ask, should the stakeholder be assumed to be able to predict these timelines?). Developing an operator response strategy to a particular type of risk is likely to be an optimization problem that analyzes a variety of response scenarios under different informational assumptions and seeks to minimize their overall cost. Alternatively, to improve computational performance, a simple rule set—which may yield near-optimal operator response strategies—can be utilized instead.

The net lifecycle cost under uncertainty approach is a neater and more flexible technique, albeit a more computationally intensive one, than an explicit forward-looking build-up of lifecycle cost and value streams. By assuming a constant level of mission capability for all architectures, it also omits some aspects of architectural responsiveness which may be important design discriminators in certain cases. However, most of the other critical attributes of the system, including flexibility and robustness, are "automatically" accounted for in assessing the system's ability to respond to different types of perturbations on the basis of some notional stakeholder behaviors.

Validation

It must be emphasized that the twin methodologies described in the previous section are purely illustrative—countless other approach can, no doubt, be devised to model the attributes enumerated and discussed here. Whatever approach is selected, it is of paramount importance that the resultant framework or tool be exercised against realistic architectures to assess its potential for real-world applicability to the spacecraft procurement and design processes.

In a perfect world, the development and validation of a design tool precedes the actual design of the system. Lamentably, however, the world we live in is sorely imperfect, and it appears that the development of the value-centric design methodology will be largely concurrent with the design of the System F6 demonstrator platform. A point of consolation, perhaps, is that since fractionated architectures have the inherent flexibility to adapt to uncertainty, any modifications that may emerge as a result of the value-centric design framework may be incorporated into the system rather late in the design phase. It is critical for its ultimate acceptance, and likely somewhat disruptive in the interim, that this novel design methodology be incorporated into the satellite builders' processes, and that their spacecraft development framework be made to comport with the unconventional design guidance that may result.

As primarily a tradespace exploration or top-level architectural design tool, the methodology described here must necessarily rely on parametrics or other simple models in its estimation of cost (and value, if appropriate). To lend credence to the hypothesis that fractionated architectures provide a superior solution for certain classes of missions, a more detailed econometric analysis of a specific point design must be undertaken. Such an analysis would start with a specific program of record with which a satellite builder is intimately familiar. A notional fractionated design with identical mission functionality would be developed using a methodology analogous to the one described above. Then, a detailed ground-up estimate of an appropriate cost/value metric would be developed to enable a high-confidence comparison between the monolithic and fractionated designs. The purpose of such a comparison would be to provide compelling evidence of the utility (or dis-utility) of fractionated architectures in a specific, realistic mission scenario that may be used to inform future satellite procurement decisions. DARPA's System F6 program is seeking to validate the tools being developed against a series of point designs derived from programs of record for which detailed technical and costing information is available.

Traditional Systems Engineering

One manifestation of the traditional systems engineering process used for most aerospace and defense systems is depicted in Figure 3, below. Stakeholder mission requirements are the primary input to the process. These are converted to architectural requirements through the requirements loop which iterates between a requirements flowdown (analysis) activity which converts mission requirements into their architectural implications (better known as design requirements) and a functional allocation activity which decomposes the system into lower-level functions and accordingly allocates design requirements. The output of the requirements loops is a functional decomposition of the system and flowdown of high-level mission requirements to lower-level design requirements at all functional levels.

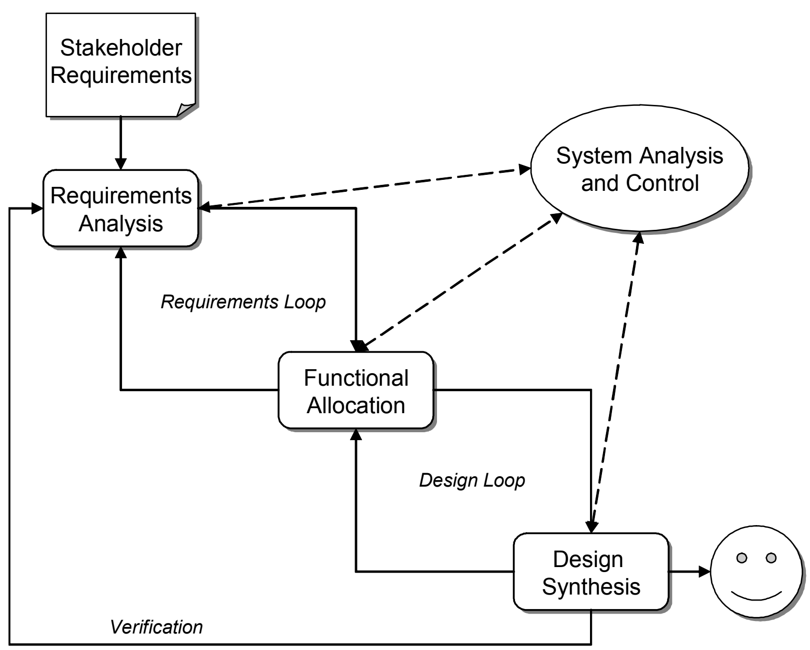

Figure 3: Traditional systems engineering process[25]

The design loop, in turn, is responsible for the genesis of an optimal design that conforms to the design requirements at all functional levels. The design synthesis activity generates and verifies closure of alternative design concepts, while iterating with the functional allocation activity to ensure that the functional decomposition is consistent with the physical manifestation of the system. The verification loop confirms the integrity of the overall process by

25 Adapted from *Systems Engineering Fundamentals*, Defense Acquisition University Press (2001).

verifying the performance attributes of each synthesized design against stakeholder mission requirements.

The system analysis and control activity performs a variety of "housekeeping" functions such as project, schedule, configuration, and interface management. Most interestingly for our purposes, however, it is also where responsibility for performing trade-off studies among alternative design concepts resides. Although modern manifestations of the systems engineering process indicate that "[p]erformance, cost, safety, reliability, risk, and other effectiveness measures must be traded against each other and against physical characteristics,"[26] the more routine practical implementation of this guidance is to select the minimum-cost alternative that meets the stakeholder-specified performance objectives. This drive to minimize cost is arguably driven by fiscal responsibility. More likely, however, it is driven by the practical demands of executing functional allocation and design synthesis.

Cost is a relatively simple proxy variable to use as a figure of merit for particular designs. Cost is allocable and additive across components of the overall system. It also has the convenient property of correlating—nearly linearly in many cases—with system mass. Mass, in turn, is one of the foremost design parameters for aerospace systems; engineers are comfortable using it as a *de facto* "currency" in system design and trading it across subsystems and integrated product teams (IPTs).[27] Imagine trying to trade flexibility, robustness, or any of the other non-traditional attributes discussed here to inform a decision about what type of weight, thrust, specific impulse, and reliability for a thruster to use. It appears nigh impossible. But making such a decision on the basis of cost or mass—for a given minimum level of performance required to meet mission requirements—would be second nature to any engineer.

The problem with this conventional approach is two-fold. First, it largely makes unrealistic the possibility of trade-offs between performance, cost, risk, etc. since these are—in the traditional paradigm—incommensurable metrics. We do not deny, of course, that such trade-offs are sometimes made. But they are typically done on an ad hoc and qualitative basis at senior-most levels of the program. These are not routine trades that are part of the "inner loops" of the iterative systems engineering process. Second, this approach altogether ignores design flexibility, robustness, and the other derivative attributes. On occasion, these may be stakeholder requirements or engineering "best practices." Without a rigorous definition, much less quantification, however, it is unclear how these attributes can be systematically built into a design if a cost or mass penalty is incurred as a consequence.

---

[26] *Ibid.* at p. 112.

[27] So much so, in fact, that gross take-off weight for aircraft and launch mass for spacecraft is frequently a stakeholder-specified requirement—in stark violation of best systems engineering practices which seek to confine stakeholder requirements to architecture-independent mission performance attributes.

It should not come as a surprise to the reader, at this point, that we maintain that net value and the variance thereof are the appropriate metrics on which system trades should be based. The use of the net value metric addresses both of the problems we noted with the traditional approach: it makes the quantities of performance, cost, and risk commensurable and therefore tradable, and it can be made to encompass flexibility, robustness, and other derivative attributes in a quantitative and largely transparent manner. So, it would seem, the simple solution is to decree net value as the metric by which system trade-off studies within the system analysis and control function are to be performed, provide a tool for estimating net value for a particular architecture, and we are done. Unfortunately, net value fails to address the original reason we cited for the popularity of cost- and mass-based figures of merit for system design. Value is not an easily allocable or tradable parameter below the system level (i.e., at the subsystem and below). It makes little sense to construct a "value budget" or trade value between subsystems. Thus, it is likely neither a convenient nor useful design metric for design synthesis below the architectural level.

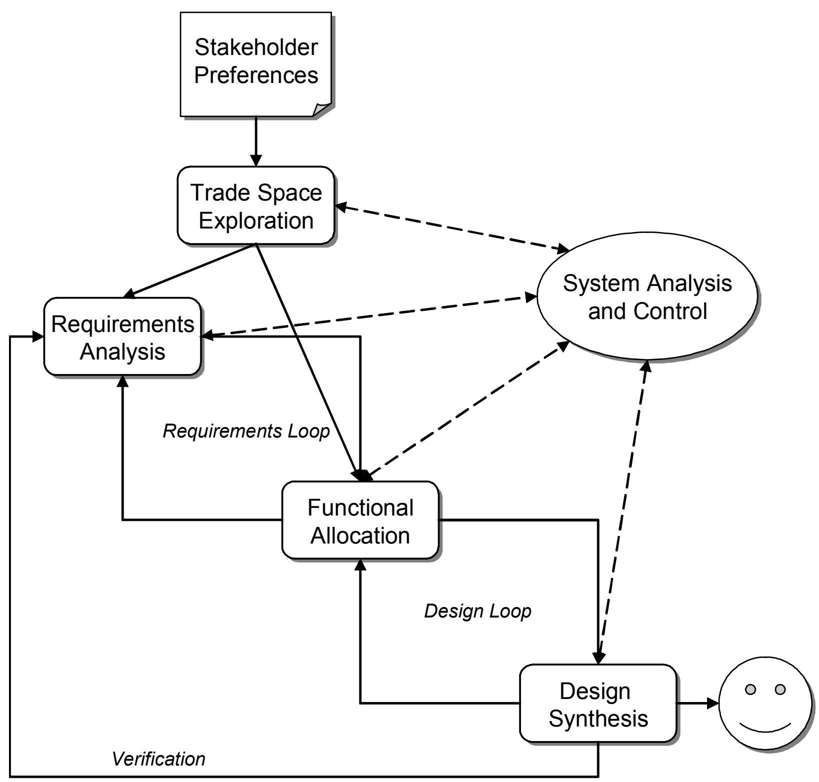

Figure 4: Systems engineering with value-centric tradespace exploration

We argue that the types of value estimating methodologies described here are best applied as a sort of requirements generation tool or a vehicle for preliminary tradespace exploration. In utilizing the tool in this manner, it supplants—to some extent—stakeholder requirements. A stakeholder would, instead, supply (either through interviews or other means) a preference function across the various mission performance attributes. "Hard" requirements, i.e., absolute constraints on some mission parameter, can be modeled by zeroing the value of all non-conforming architectures. The output of the tool would be a family of mean-variance efficient architectures that maximize net value and minimize the variance thereof.[28] These architectures would be expressed in terms of the rather brief vector of design parameters, an example of which was provided in Table 6. These architectural design parameters are certainly insufficient to specify a design at the subsystem or component level (as one might expect the output of the systems engineering process to be). Instead, they are best construed as top-level design requirements, and the value estimating methodology as a tool for converting stakeholder preferences into architectural design requirements. Put another way, the value tool effectively incorporates internally the first iteration of the requirements loop. Its output is a set of initial functional requirements and functional allocations at the system or architecture level. Once these requirements are generated, the conventional systems engineering process commences without further interference from value metrics.[29]

Incorporating Value Metrics into Risk Management

The linkage, in Figure 4, between the value-based tradespace exploration activity and the system analysis and control function is solely intended to refer to the risk management capabilities of the value estimation tools. Since, as we argued earlier in this essay, the calculation of net value also enables the quantification of risk as the variance in net value, this provides a quantitative framework for management of top-level program risks. We think that such an approach to risk management is a dramatic improvement over the traditional risk management approach which we describe as a "stop light" approach because of the qualitative color-coding of risk probabilities and impacts that lies at its heart.

---

[28] This is not to discount the importance of ensuring cost conformance with fiscal constraints. The overall cost contribution to net value can, of course, be an explicit output and parameter for the elimination of architectures that do not conform to budgetary limitations.

[29] We are aware of the potential objection that this approach leaves open the possibility that a suboptimal design decision (from a net value perspective) would be made at a subsystem or component level since the traditional process would likely rely on cost or mass minimization. Short of turning the value tool into a complete multi-disciplinary optimization engine, we do not know of a practical solution to this problem. Our suggestion is that the key drivers of architectural flexibility and robustness be identified at the outset and included in the architectural design vector which the value methodology then optimizes at the outset of the systems engineering process.

First, it is worthy of note that our aspiration here is not to throw out this traditional "stop light" risk management framework which has arguably served the systems engineering community rather well over many decades, but rather to see if it can be improved, informed, or enhanced in light of the availability of a quantitative cost and value estimating tool, and as part of a broader re-thinking of the systems engineering process which is obviously being precipitated by the value-based approach to system architecting.

Second, it appears that there may be some room for commonality in the representation of specific risks between the risk management and value modeling frameworks. As we noted previously, in order for flexibility and robustness to have any value in the first place, the value modeling framework will have to model an assortment of lifecycle perturbations, including those occurring during the development phase. Furthermore, the value model will likely represent these risks as random variables, and therefore assume some probability distribution and empirically-derived characteristics thereof. The "stop light" method is essentially a qualitative approach to doing precisely that. But if a quantitative approach will have to be undertaken in the value model, why not leverage it to represent the risks the same way in the risk management process, i.e., as probability distributions of some relevant program or system parameters?

Third, the mental paradigm precipitated by the traditional risk management approach is to "burn down" risk as quickly as possible, typically by reducing uncertainty through the maturation, early demonstration, etc. of specific technologies or components. This is not always a value-maximizing strategy, however. Instead of selecting a specific design solution and seeking to mature it as quickly as possible, it may instead be advantageous to postpone a particular design decision, and instead incorporate an option that allows the decision to be made later. In essence, the toolset of risk mitigation strategies ought to be significantly broader; more importantly, some principled approach needs to exist for deciding whether it is better—as a risk mitigation strategy—to pick a specific solution and mature it rapidly, or to design in an option instead.

Fourthly, and finally, in selecting a risk mitigation strategy, there needs to exist a methodology by which the net value of the system can be traded off against its variance, i.e., risk. Put another way, one can easily conceive of scenarios where a stakeholder might choose a system with lower net value in exchange for lower risk. This tradeoff must be made on the basis of the risk aversion profile of the stakeholder, i.e., what is the value premium necessary to justify a given level of risk incurred. So in selecting a particular risk mitigation strategy, the impact on system value, cost, and the anticipated reduction in the variance of value and cost due to the risk mitigation strategy must be considered. We believe that the value modeling methodology is a tool that could effectuate such tradeoffs.

It is easy to get overly ambitious with this theoretically elegant vision. Being able to do cost and value tradeoffs for multiple risk mitigation approaches for a myriad of component- or even subsystem-level risks is likely infeasible if we are to avoid developing a gargantuan multi-disciplinary optimization tool based on value metrics. But it is essential to be able to do this for certain major subsystem-, system-, and program-level risks, as the value impact of the different mitigation strategies might, indeed, be significant. Our proposed approach, therefore, is that each risk in the traditional risk tracking and management process should be tested against the value tool. That is to say, if the risk transpires, what effect on net lifecycle value and variance will there be? This is a quantitative measure of risk impact. The probability of occurrence ought to be harmonized with the data used to create the perturbation models that drive the value methodology. And finally, each potential mitigation methodology should also be vetted with value tool to assess its impact on net lifecycle value and its variance. Thus, the value methodology brings an element of rigor to assessing risk severity and selecting value-optimal mitigation strategies. Many risks, of course, will not have an impact on the estimate of net value or risk. This is because the type of model that we propose here is based on a small handful of high-level architectural design parameters. Those risks must, of course, continue to be managed conventionally, while perhaps qualitatively mindful of flexibility and robustness considerations.

INCORPORATING VALUE METRICS INTO ACQUISITION & PROCUREMENT

The procurement framework used by the Pentagon today has its roots in the 1961 innovation spearheaded by Defense Secretary Robert McNamara and Pentagon Comptroller Charles Hitch known as Planning, Programming, and Budgeting Systems (PPBS). PPBS was a significant improvement over the blind allocative approach employed since the American Revolution, and remains alive and well at the core of the Pentagon's contemporary acquisition framework (it has been re-titled to PPBE, with the "E" standing for execution). The theory behind PPBS was and remains sound. Rooted in Hitch's profound RAND treatise, *The Economics of Defense in the Nuclear Age*, whose publication preceded his appointment to the Pentagon post, it sought to tie the budgeting process to military strategy through a series of systems analyses. In this work, Hitch acknowledges that (what we call) net value is the appropriate criterion for procurement planning: "If gains and costs can be measured in the same unit, then to maximize gains-minus-costs is certainly an acceptable criterion-form—the equivalent of making the most out of whatever actions can be taken."[30] Hitch,

[30] Hitch, C. & McKean, R., *The Economics of Defense in the Nuclear Age*, RAND Corporation Report R-346, p. 175 (1960).

however, can be seen to struggle with the so-called problem of "incommensurables;" he writes:

> "Nor is it to say that it is hopeless to try to value other effects which appear at first glance to be incommensurables. In specific analyses ingenuity can often go a long way toward measuring such effects in terms of the common unit. We cannot say just where to draw the line between effects that should be measured in terms of the common denominator and those that should not…. A Congressional committee was probably justified in concluding its review of an evaluation of federal resource development projects: 'Some of the effort to place monetary values on indirect benefits is nothing short of ludicrous.'"[31]

In implementing PPBS, McNamara and Hitch eschewed attempting to remedy this problem of incommensurables altogether. In his retrospective on PPBS, Hitch writes:

> "Thus, the problem of allocating resources within the Department of Defense itself involves the choosing of doctrines, weapons, equipment, and so forth, so as to get the most defense out any given level of available resources or, what is logically equivalent, to achieve a given level of defense at the least cost…. Approaching the problem from the second point of view—achieving a given level of defense at the least cost, which is the way Secretary McNamara prefers to look at the problem—we work in terms of marginal products and marginal costs in order to help the top decision-maker choose the appropriate level of resources."[32]

By adopting a minimum cost criterion—achieving a given level of defense at the least cost—McNamara and Hitch effectively capitulated on the question of quantifying benefit or value.[33] Theoretically, of course, the minimum cost formulation can be construed as equivalent. For a fixed value or benefit level, the cheapest system will be the one with the highest net value. The difficulty—which we discussed at length in the preceding discussion on systems engineering—stems in the practical implementation. This approach levies the burden of

---

[31] *Ibid.* at p. 185.

[32] Hitch, C., *Decision-Making for Defense*, University of California Press, p. 52 (1965).

[33] In the past few decades, one of the much-touted reforms to the procurement process has been that of "best-value contracting," whereby more than just the cost of a system is considered. Specifically, the system that offers the "best value to the government" is selected. Unfortunately, value is defined purely as a function of cost, performance, quality, and schedule (and not other non-performance attributes of the architecture), and provides no method for weighting or otherwise combining these so-called key performance indicators (KPIs) into a single measure of value. Flexibility and robustness are ignored altogether, and the statutory requirement is only to provide a ranked list of these KPIs when used as source selection criteria (and not their actual weighting) in competitive procurements.

formulating exactly what the value or benefit of the system should consist of upon the stakeholder (i.e., procurer or user). The stakeholder, by necessity, focuses on his mission performance, not the architectural nuances of the system. Thus, he tends to formulate a narrow set of mission requirements which set the baseline value or benefit of the system. The system designer or developer is then left with devising a minimum-cost solution to those mission requirements. Attributes like flexibility and robustness are never levied as requirements—both because the stakeholder has little insight into the architectural options that would effectuate these qualities in a system, and because no apparent quantitative metrics for levying objective flexibility and robustness requirements are in common use.

By contrast, the value-centric procurement approach which we advocate requires us to tackle the problem of incommensurables head on. Though open to criticism, debate, and undoubtedly significant improvements, a set of credible techniques for quantifying value, as we have shown here, does exist. By making net system value the criterion for guiding procurement decisions, the stakeholder is reduced to specifying—in general terms—the "level of defense" that he seeks from the system. It is then up to the system designer or developer to determine how to maximize the net value of his proffered architecture. The designer's toolbox then includes improving performance, reducing cost, or incorporating non-performance value-enhancing attributes such as flexibility and robustness into the system design. He is incentivized, of course, to do all three—and not just minimize cost as under the *status quo*.

CONCLUSION

The purpose of this essay has been to revisit the concept of fractionated spacecraft (particularly in light of the recently-commenced DARPA System F6 program), describe their merits and demerits relative to the traditional paradigm of monolithic satellites, and to introduce a family of value-based techniques for their design and procurement. The detailed results of the DARPA technology maturation effort, development of open-source value-based space architecture design tools, and their application to actual space missions will, hopefully, be the topics of numerous publications forthcoming from the various performers of the DARPA program,[34] government stakeholders, and other interested parties. Additionally, the topics of value-based systems engineering and acquisition practices are the subject of ongoing additional work by the authors and will be discussed in detail in future publications.

---

[34] The performers for the preliminary design phase of the program include teams led by Orbital Sciences, Northrop Grumman, Lockheed Martin, and Boeing.