IDA | RESEARCH NOTES

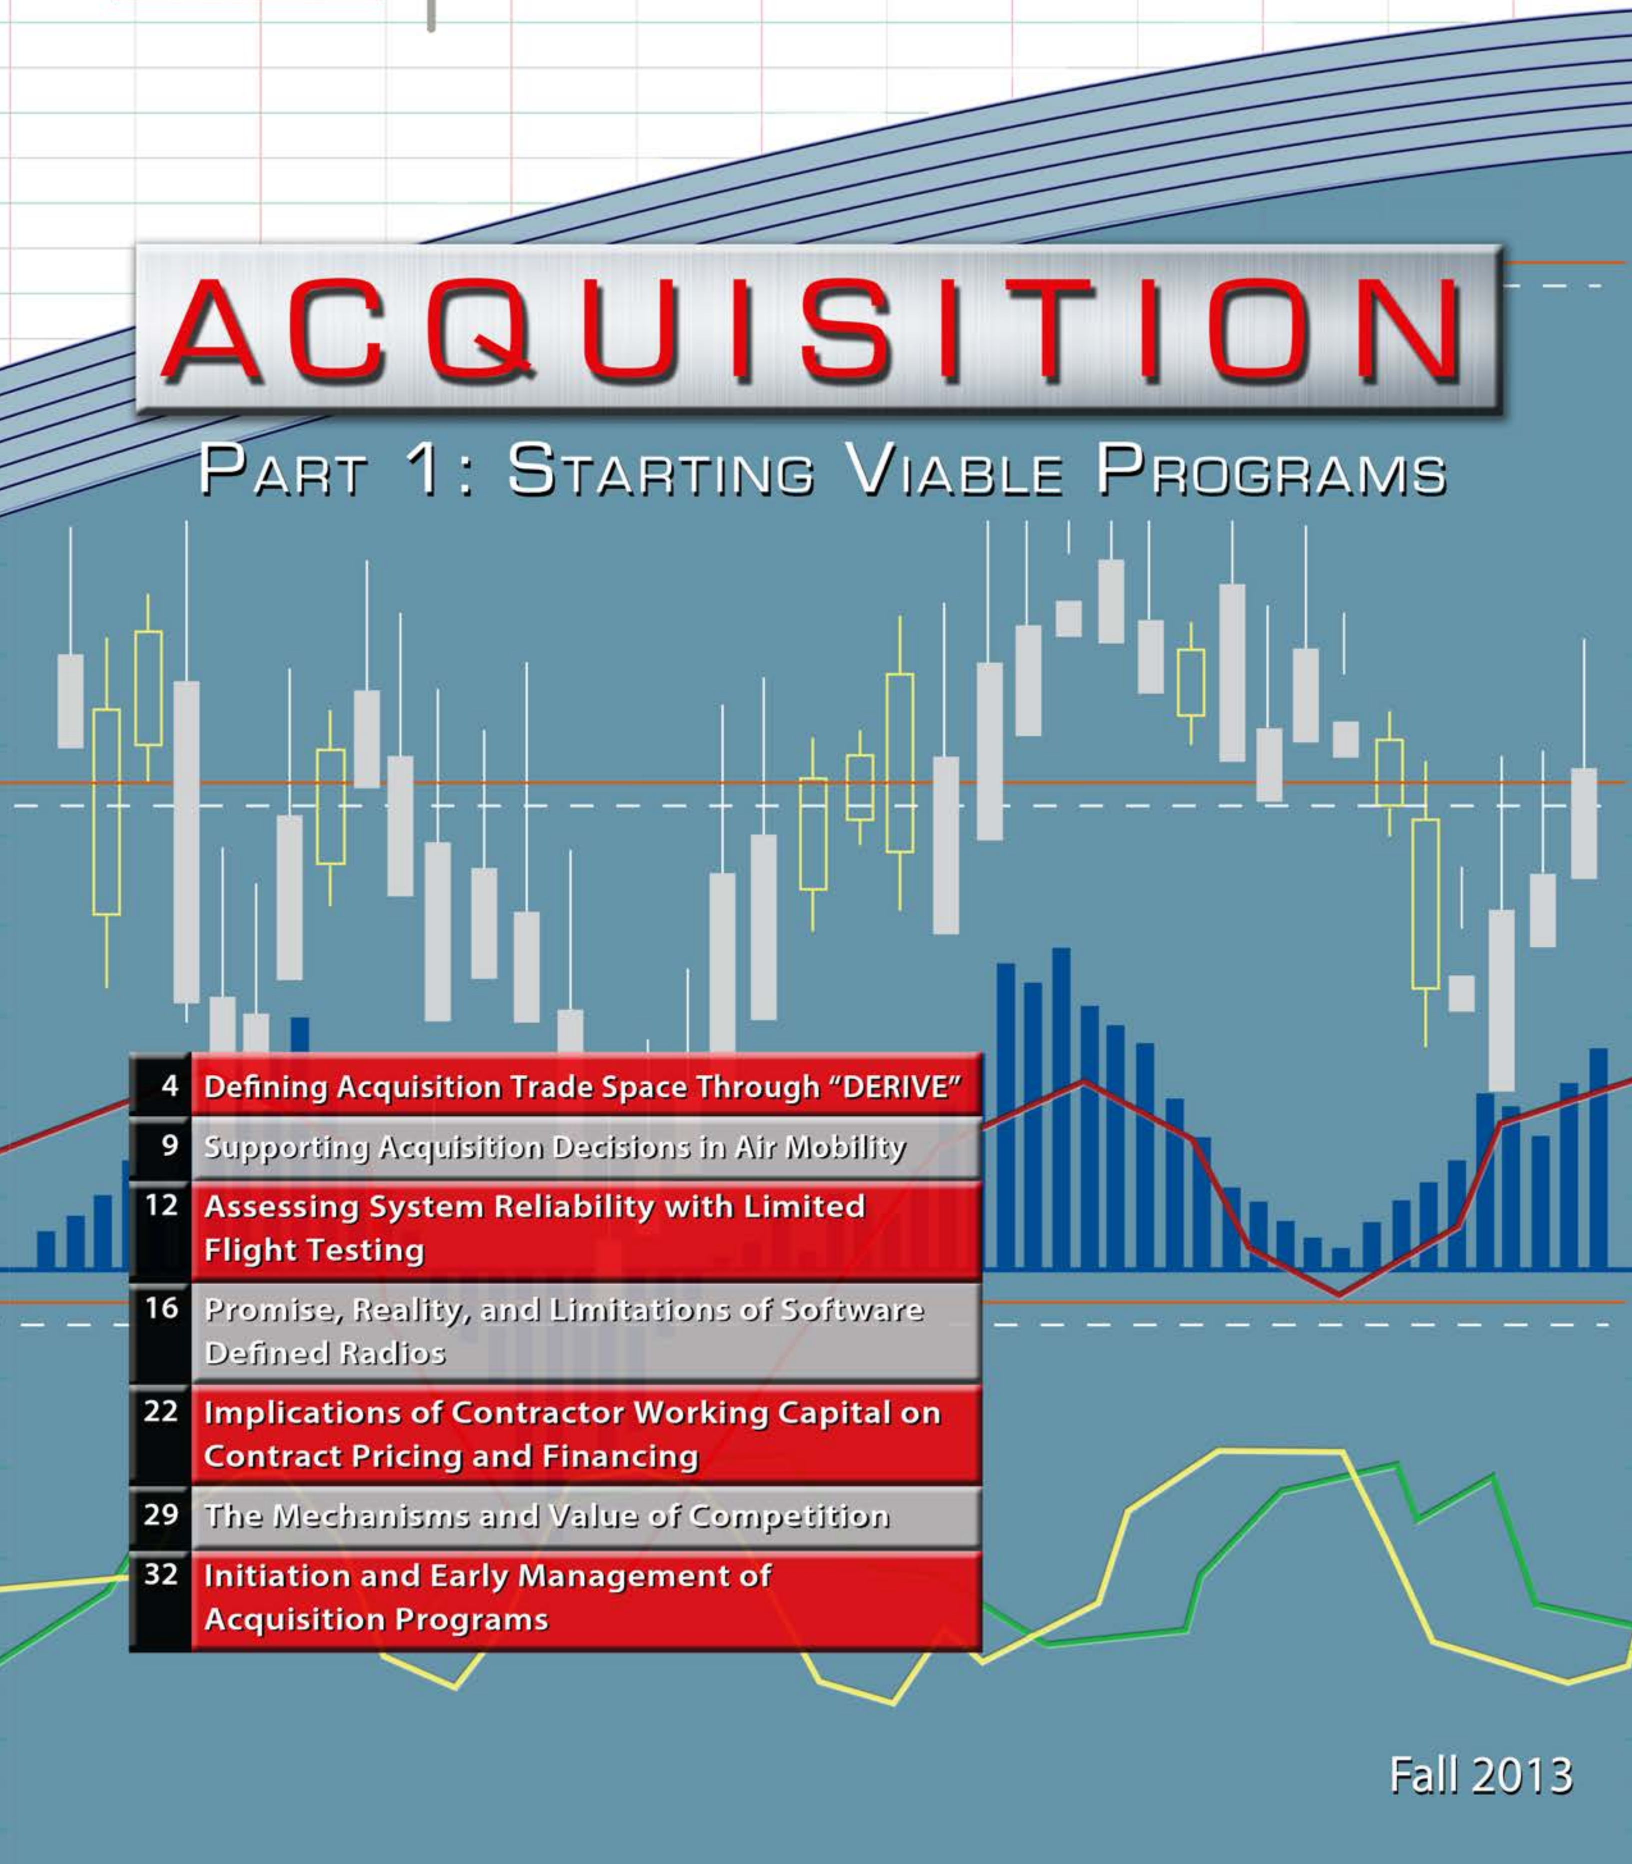

ACQUISITION

Part 1: Starting Viable Programs

4 Defining Acquisition Trade Space Through "DERIVE"
9 Supporting Acquisition Decisions in Air Mobility
12 Assessing System Reliability with Limited Flight Testing
16 Promise, Reality, and Limitations of Software Defined Radios
22 Implications of Contractor Working Capital on Contract Pricing and Financing
29 The Mechanisms and Value of Competition
32 Initiation and Early Management of Acquisition Programs

Fall 2013

IDA is the Institute for Defense Analyses, a non-profit corporation operating in the public interest.

IDA's three federally funded research and development centers provide objective analyses of national security issues and related national challenges, particularly those requiring extraordinary scientific and technical expertise.

The articles in this edition of IDA Research Notes were written by researchers within the following four IDA divisions. The directors of those divisions would be glad to respond to questions about the specific research topics or related issues.

Cost Analysis and Research Division (CARD)

Dr. David J. Nicholls, *Director* (703-575-4991, dnicholl@ida.org)

Science and Technology Division (STD)

Dr. Leonard J. Buckley, *Director* (703-578-2800, lbuckley@ida.org)

Strategy, Forces and Resources Division (SFRD)

Mr. Michael Dominguez, Director (703-845-2527, mdomingu@ida.org)

System Evaluation Division (SED)

Dr. Steve Warner, *Director* (703-845-2096, swarner@ida.org)

Institute for Defense Analyses
4850 Mark Center Drive
Alexandria, Virginia 22311
www.ida.org

ACQUISITION, PART 1: STARTING VIABLE PROGRAMS

This is the first of two issues of IDA's *Research Notes* focused on acquisition in the Department of Defense. These two issues will cover topics related to program inception, program execution, and managing a portfolio of acquisition programs. In this issue, we present articles related to program inception—specifically, starting viable programs.

Establishing feasible programs at Milestone B is truly a critical step. Programs that enter development with flawed assumptions or concepts must address the consequent problems during execution—sometimes resulting in Nunn-McCurdy breaches. This issue contains articles on three challenges that programs must address at inception: selecting the system's capability and its specific attributes, accurately estimating cost and schedule, and choosing an appropriate acquisition strategy.

The first three articles address two aspects of selecting a capability—its specific attributes, and how to assess whether the capability has been achieved. In the opening paper, **Prashant Patel** *et al* describe the IDA-developed Deducing Economically Realistic Implications Via Engineering (DERIVE) framework. This physics-based assessment tool fuses a variety of information sources to provide managers and decision makers early in the process with new insights into costs, likely risks, and potential trades. It also helps define major assumptions, the validity of which should be tracked as the program evolves. In the second paper, **William Greer** summarizes the specific results and overall lessons learned from selected airlift cost-effectiveness analyses conducted at IDA. One lesson learned is the importance of assessing all options, including upgrades, new procurements, and, if necessary, termination. **Joseph Buontempo** discusses the challenges in assessing reliability when the expense of testing means there will be very few tests. Specifically, he discusses how Bayesian approaches can significantly reduce the uncertainty in such assessments.

Inaccurate cost and schedule estimates are a major source of future program problems. Such inaccuracies are not generally due to methodological or computational errors, but rather to the invalidity of fundamental assumptions that define the nature of the program. **Lawrence Goeller** and **Patricia Bronson** discuss the assumption, critical for software-defined radios, that hardware could be both sufficiently powerful for signal processing in real-time and yet flexible enough to support future undetermined waveforms on the Joint Tactical Radio System (JTRS) programs.

The next two articles address issues associated with the selection of an effective acquisition strategy that aligns contractor incentives with government goals. **Scot Arnold** examines the implications of contractor working capital on contract pricing and financing, and the difficult problem facing the government in negotiating cost-efficient sole-source procurement with fixed-price contracts. **James Dominy***et al* look at the mechanisms and value of competition, which is often considered a cornerstone of the Federal government's acquisition process. In particular, he shows how the competition for a development program is really about the potential profits in production and discusses the implications of that for the government.

The final article looks broadly across the issues related to program inception. **Royce Kneece** and **Gene Porter** address issues in the initiation and early management of major defense acquisition programs. Topics covered include the difficulties often faced in defining requirements and an IDA proposal for a more analytically based process.

DEFINING ACQUISITION TRADE SPACE THROUGH "DERIVE"

Prashant R. Patel, David Gillingham, and David Sparrow

The Problem

A lack of analytical rigor and poor communication between the program developers and the acquisition oversight community often lead to "false-starts" and "do-overs" during program initiation.

IDA's trade space framework—Deducing Economically Realistic Implications Via Engineering (DERIVE)—links engineering and physics analysis, operational constraints, and semi-parametric cost estimates.

IDA's trade space framework—Deducing Economically Realistic Implications Via Engineering (DERIVE)—links engineering and physics analysis, operational constraints, and semi-parametric cost estimates. The goal is to increase the efficiency of the acquisition process by reducing friction between the program office, the Services, the Joint Staff, and the Office of the Secretary of Defense (OSD), especially at program initiation and during the early stages of development.

IDA designed the DERIVE framework to link important technical inputs to programmatic and operational outputs in a straightforward, traceable, and transparent manner. The framework provides an analytic structure that could be used to build understanding and communicate intent. It could be especially helpful for programs whose complex interactions between requirements, operational restrictions, and technology—rather than any individual issue—drive acquisition outcomes.

Trade Space

The use of trade studies in engineering is not new. It has a long history in the technical community and has now been formally adopted into the DoD acquisition decision-making process. Recent experiences suggest that the Services' trade-space tools are being used to inform their internal deliberations. However, several recent new-start proposals have been the subject of follow-on trade studies and amended Analysis of Alternatives efforts, suggesting room for improvement.

Schedule delays associated with follow-on analyses can be avoided if the trade study processes and analytical outputs are structured to support both user and oversight objectives. The outputs of IDA's DERIVE framework are constructed to achieve this goal by enhancing traceability and transparency of inputs, outputs, and decision making.

Traceability

Traceability is used by systems engineers to manage technically complex endeavors by flowing down program objectives into discrete technical goals. Alternatively, students employ traceability to demonstrate to professors that they have a firm grasp of the nature of problems even if small errors are present in the analysis. Traceability can also be leveraged by the Services and program offices to demonstrate that they have rigorously analyzed the operational environment and have a firm understanding of the technical issues and programmatic consequences for a new program.

DoD asked IDA to develop and demonstrate DERIVE on a generic infantry fighting vehicle (IFV). The results of that effort will be used below to illustrate how DERIVE's outputs are designed to foster traceability.

Creating traceability requires exposing objectives of the program, how they relate to technical assumptions, and how the various elements interact to drive results. An output of the DERIVE process traces the desired capabilities to the commensurate technical inputs. Table 1 shows how key performance and programmatic attributes can be mapped to specific technical requirements for an IFV.

Cross-referencing the technical assumptions and desired capabilities in a single, compact form provides two benefits. First, it allows the program developers to clearly articulate the user's goals and the technical requirements necessary to achieve those goals. Second, it allows the oversight community to understand the potential loss of capability if there are technical shortfalls during development.

Similarly, Table 2 shows how cost traceability can be achieved. Various cost categories are mapped to the data sources and assumptions used in generating the cost estimate. This trace-

Table 1. Performance and Technical Traceability Matrix

Capability Area		Specifications (Desires)	Analytical Implication
Force Protection	**Ballistic**	Trade space	Integral ballistic armor must be able to passively defeat ballistic threats.
	Explosive	Survive an X class of IED and a Y RPG	Supports 45 pounds/square foot (psf) of integral underbody armor and 95 psf or add-on EFP armor.
Passenger Capacity		Trade space	Interior volume scales based on human factors and number of passengers (32 cubic ft/person and 450lbs/person).
Full Spectrum	**Weight**	Desire system to reliable	Structure, engine, transmission, etc. must be sized to support add-on EFP armor.
	Power	Increased exportable power	Has a 50-horsepower generator for electrical power.
Timing		Field System Quickly	Uses currently producible armor materials, engines, etc.
Transportability		Transportable by C-17	IDA-defined combat weight limited to 130,000 lbs and must fit inside compartment E of C-17.
Mobility		Speed of X up a grade of Y	Uses an Abrams-like track and has 20 horsepower/ton of engine power.
Lethality		Lethal to similar class of vehicles	Has a manned turret. Reserved 2.1 tons for non-armored turret weight and 120 cubic feet of volume. Also, 2.5 tons for ammunition and fuel.
Electronics and Sensors			Has sensors/electronics similar to Abrams and Bradley.
General			Includes other fixed vehicle components (e.g., wiring, bolts, weld material). Weight allocated to these types of items is 2.5 tons.

ability matrix allows oversight organizations to qualitatively assess the riskiness and fidelity of the estimate.

Table 2. Cost Elements and Costing Assumptions and Data Sources

Cost Element	Description/Sources/Methodology
Hull/Frame	Cost estimating relationship depends on material type and weight. Assumed a buy-to-fly of 1.
Suspension, Engine, Transmission, Auxiliary Automotive, Integration, Assembly, Tests, and Evaluation	Army Ground Vehicle Systems Bluebook (2006)
Add-on EFP armor	Estimated as cost per ton from budget data and publicly reported contract values.
Electronics/sensors	Estimated from President's Budget submissions for ground vehicle upgrade programs. Focused on sensors and electronic upgrades
Contractor non-prime mission product cost elements	Estimated using historical contractor cost data reports. Applied as a multiplication factor on the prime mission product.
Support	Estimated using Selected Acquisition Reports. Applied as a factor on contractor costs.
Deflation/inflation rates and conversions	Joint Inflation Calculator (http://www.asafm.army.mil/offices/office.aspx?officecode=1400) .

Finally, the logic used to estimate the costs and performance of the IFV trade space is described in Figure 1.

- **Determine size of the box (volume under armor)**
  - Number of dismounts and crew; soldier space claim
  - Interior mission equipment and auxiliary automotive space claim
- **Determine weight of the box**
  - Front, side, rear, ballistic force protection; underbody and EFP protection
  - Areal density of protection technologies
  - Other - radios, seats, steering, soldiers, etc.
- **Determine weight and size of subsystems that move the box**
  - Drivetrain, suspension, support structure
  - Engine track/tires based on mobility requirements – hp/ton, ground pressure, etc.
- **Cost the system based on identified materials and components**
  - Scale contractor and program costs
- **Prune infeasible solutions**
  - Impose constraints such as transportability weight restrictions

Figure 1. Outline of Process Used in Creating Infantry Fighting Vehicle Trade Space

In sum, the DERIVE framework helps program developers and the acquisition oversight community build a common understanding of the key technical, operational, and cost drivers of new capabilities being sought by the Department.

Transparency

The DERIVE framework improves the transparency of the analyses supporting acquisition decisions. Figure 2 shows an output of the DERIVE framework for the IFV example. It enhances transparency by illustrating the entire trade space rather than a few point designs. Showcasing the full trade space demonstrates the thoroughness of the investigation and reduces the possibility of having to include additional cases. Also, instead of using a value function, the analysis simply highlights the desired point solutions and lists the rationale for the decision and the relevant trade-offs that were considered and accepted as part of the decision-making process. Showing trade space data, the rationale, and the resulting decision together serves to enhance trust, convey thoroughness, and reduce institutional friction.

Conclusion

DERIVE and similar approaches provide a framework that can be used to engage and improve acquisition outcomes. DERIVE fuses a variety of information sources (capabilities,

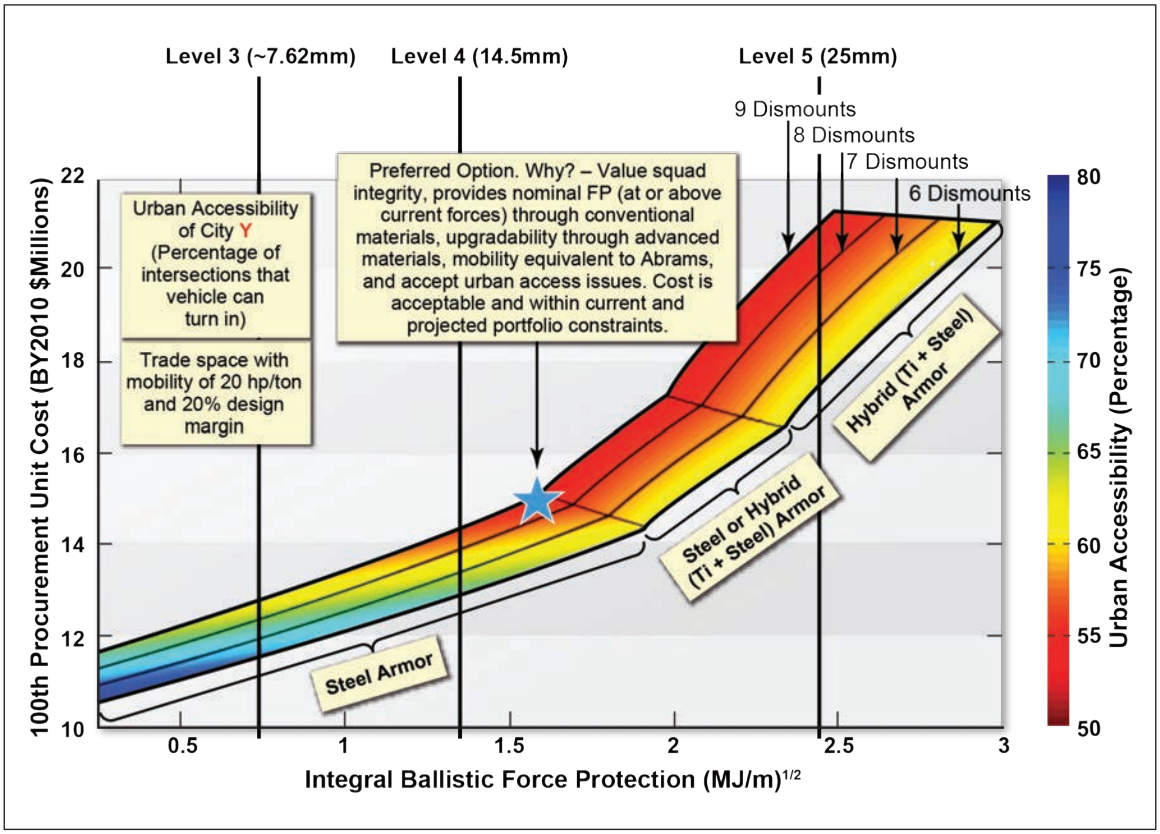

Figure 2. Infantry Fighting Vehicle Trade Space with Logic for Decision

operational, technical, and cost) to enable more thorough analyses in support of decision making and to reduce friction between program developers and the acquisition oversight community.

Dr. Patel is a Research Staff Member in IDA's Cost Analysis and Research Division. He holds a doctorate in aerospace engineering from the University of Michigan.

Dr. Gillingham is a Research Staff Member in IDA's Science and Technology Division. He holds a doctorate in physics from the University of Maryland, College Park.

Dr. Sparrow, a Research Staff Member in IDA's Science and Technology Division, received a doctorate in physics from the Massachusetts Institute of Technology.

SUPPORTING ACQUISITION DECISIONS IN AIR MOBILITY

William L. Greer

The Problem

In recent years, DoD has faced several difficult decisions regarding the modernization and recapitalization of U.S. airlift forces, and, in each case, analysis of the effectiveness and costs of available options provided key insights to inform the decisions, while offering lessons for analysts going forward.

Air mobility forces—airlifters and aerial tankers—serve a crucial role, both in peacetime and in wartime military operations. We focus in this article on airlift, the rapid movement of cargo and passengers to, from, or within a theater. The cargo can include a diverse range of materiel, including mail, spare parts, and combat vehicles and ammunition. Passenger airlift can include rapid medical evacuation as well as the airborne movement of troops.

IDA has conducted a number of airlift cost-effectiveness analyses over the last 20 years. These assessments were variously called *Cost and Operational Effectiveness Analyses (COEAs), Analyses of Alternatives (AoAs)*, or simply cost-benefit trade studies. All served the same end: to inform decision makers about the desirability of major acquisition programs and their alternatives.

IDA has conducted a number of airlift cost-effectiveness analyses over the last 20 years. All serve the same end: to inform decision makers about the desirability of major acquisition programs and their alternatives.

Case Studies

We discuss here three examples of program designs that have been informed by the IDA studies on airlift.

C-17 COEA: What Kind of Airlifters Should DoD Buy?

Congress mandated an IDA analysis, *Cost and Operational Effectiveness Analysis of the C-17 Program* (1993), which was intended initially to investigate whether the aging C-141 airlifters should be given an extended life or whether a new airlifter—the C-17—should be bought instead. If C-17s were to be bought, the C-141s would be retired. The IDA analysis included C-141s and C-17s as alternatives, but added to the list of alternatives several military-modified commercial cargo airlifters. Our analysis showed that requirements could be met in the most cost-effective way not by buying just C-17s or by extending C-141 life, but by a mix of some C-17s and the less costly modified commercial cargo aircraft. This mixed fleet solution was influential in subsequent DoD decisions and led to a competition between the manufacturer of the C-17 and manufacturers of large commercial cargo planes. In the end, although DoD decided to buy only C-17s, the intense competition forced improved performance and led to lower costs for that choice.

C-5M: Should the C-5 Fleet Be Upgraded with New Engines and Improved Reliability?

In the late 1990s, DoD was faced with a decision about improving the C-5 airlifter, the single largest airlifter in the U.S. inventory: invest in replacing the older C-5 engines and low-reliability parts with new ones, or maintain the existing systems. Which would be more cost effective? At that time, 126 C-5A/Bs were in the fleet. In the 1997 *Independent Analysis of C-5 Modernization Study*, IDA researchers estimated that the cost of replacing all engines and low-reliability parts plus the costs for maintaining the upgraded fleet over 25 years would be lower than the costs to maintain existing systems over the same time period. Despite the large initial expenses for new parts and engines, the ultimate savings expected from lower maintenance costs for the improved C-5, dubbed the C-5M, offset the initial new parts acquisition costs. And the added reliability of these improvements was projected to immediately improve mission capable rates and departure reliabilities. Table 1 shows this comparison. DoD decided to proceed with the C-5M program, made a request for proposals to industry, and shared our analytical results with industry through a redacted version of the report that removed proprietary data provided to IDA by major contractors.

Size and Mix: Under What Conditions Should the C-17 Production Line Be Continued or Terminated?

When do you stop a program, particularly a program such as the C-17, which had proven to be effective operationally? Because successful programs often develop strong proponents, the arguments for termination need to be solid. In the congressionally mandated 2009 *Study on Size and Mix of Airlift Force*, we assessed the pros and cons for terminating the C-17 acquisition program. Terminating U.S. procurement didn't necessarily mean the production line would be shut down; the C-17 was still being built and delivered to foreign governments. But eliminating U.S. procurement would certainly limit the number of aircraft being built per year, leading to employee layoffs and increasing the unit cost of the smaller numbers of aircraft that were produced each year. Obviously, termination would not be warranted if U.S. military needs were not being met. But that was not the

Table 1. Comparison of Cost and Effectiveness of C-5 Alternatives

126 Aircraft Fleet	25-Year Cumulative Net Present Value (Billions of Dollars)	Operational Effectiveness	
		Mission Capable Rate (%)	Departure Reliability (%)
Full Modernization to C-5M	15.6	74	91
Base C-5A/B	16.1	64	85

case. Prior to IDA's research, the main argument against termination was that more U.S. C-17s should be procured as a hedge against emerging needs that demanded a larger-than-planned C-17 fleet. However, the IDA analyses showed that there were more than enough C-17s available for anticipated needs out to 10 years. Second, we showed that the cost of stopping and then restarting the C-17 line would be lower under reasonable discounting assumptions than the cost of retaining an open line and producing aircraft at a low sustaining rate, for the purpose of keeping the line open. We estimated that if the line were to be terminated and then reopened in 10 years, DoD would have saved money by waiting, in spite of large restart costs. Plus, in that period of time, it is likely that an entirely new airlifter design could be under serious consideration as a competitor to the C-17.

Lessons Learned

Although each study ended with its own set of specific insights and recommendations, several overarching factors emerged from our findings. The following lessons relate to airlift acquisitions but could be extended to other large military systems.

Competition

Competition forces all manufacturers to provide the most cost-effective aircraft they can, a pressure that would likely be missing if there were no alternative choices. The C-17 COEA provides a good illustration of that. The government actually ended up with the airlifter they were hoping to buy initially, but at lower cost and with greater capability when competition was introduced.

Comparisons

Instead of comparing airlifters one-on-one, the appropriate comparison should be fleet-on-fleet, and the comparison medium might be mission accomplishment rather than a specific performance factor. This gives all airlifters in the fleet the opportunity to carry what they carry best, allowing for a potentially lower cost *fleet* alternative than would be achieved by a fleet with only one kind of airlifter.

Upgrades

Sometimes upgrades (rather than recapitalization) can be a cost-effective way to improve the airlifter fleet. And sometimes recapitalization is the more cost-effective route. There is no set answer; the answer depends on the details (e.g., service life remaining after upgrades are installed, relative costs of acquiring and maintaining different types of aircraft).

Termination

The point at which DoD should terminate acquisition is informed by the peacetime and wartime airlift requirements as well as affordability considerations. It will be a different answer for any specific sets of cases.

Dr. Greer is an Assistant Director in IDA's System Evaluation Division. He holds a doctorate in chemical physics from the University of Chicago and a BA in chemistry from Vanderbilt University.

ASSESSING SYSTEM RELIABILITY WITH LIMITED FLIGHT TESTING

Joseph T. Buontempo

The Problem

Expensive flight tests often cannot be conducted a sufficient number of times to yield estimates of system reliability with low uncertainty.

One of the challenges for DoD in developing and fielding the Ground-Based Midcourse Defense (GMD) system is testing the performance and reliability of the Ground-Based Interceptors (GBIs). (The GMD system is described below.) Because GBI flight tests are expensive (on the order of $250 million each), the Missile Defense Agency (MDA) can conduct only a limited number of them—typically one flight test per year. Such infrequent testing would yield considerable uncertainty in estimates of GBI reliability—and therefore in assessments of whether the interceptors are meeting requirements. In the past, similar challenges have arisen in assessments of the reliability of nuclear power plants, nuclear weapons, and some other weapon systems due to limited system-level testing. These situations helped drive the development of the Bayesian methodology for estimating reliability.

IDA examined the use of the Bayesian methodology as a way to estimate GBI reliability given limited flight testing and to reduce the uncertainty and risks associated with these estimates.

IDA examined the use of the Bayesian methodology as a way to estimate GBI reliability given limited flight testing and to reduce the uncertainty and risks associated with these estimates. A Bayesian approach quantifies a starting state of knowledge using a probability distribution (called the "prior distribution"), uses data as they become available to modify the state of knowledge (using the formalism of Bayes' Theorem), and summarizes the resulting state of knowledge with a refined probability distribution (the "posterior distribution").

In determining the starting state of knowledge for estimating GBI reliability, the GBIs can be divided into components (or subsystems), and the reliability of each component can be modeled

The GMD system is intended to engage limited intermediate- and long-range ballistic missile threats in the midcourse phase of flight to protect the United States. The GMD system, which employs GBIs, is supported by multiple sensors that detect and track the ballistic missile threats. The GBI is a three-stage, solid-fuel rocket carrying a 230-pound Exo-atmospheric Kill Vehicle (EKV) toward the target's predicted location in space. Once released from the booster, the EKV uses data received in-flight from ground-based radars and its own on-board sensors to attempt to close with and, using the kinetic energy from a direct hit, destroy the target outside Earth's atmosphere.

with its own probability distribution. For illustration, one can make the initial simplifying assumptions that the starting state of knowledge of each component is identical; the initial reliability of each component is unknown; and the components have independent reliabilities. The product of the component reliabilities, therefore, is the system reliability. Given these assumptions, the prior distribution can be produced for GBI reliability, and as new data become available, the reliability distribution for each component—and thus the GBI system—can be updated using Bayes' Theorem (see the callout box for details).

In general, the Bayesian methodology is flexible enough to capture other (more useful) starting states of knowledge and to incorporate many different kinds of data. Regarding the former, different prior distributions could be specified for each component, for example, rather than assuming that all components start with the same uncertainty about reliability. Regarding the incorporation of different kinds of data, appropriate models could allow the use of data obtained from modeling and simulation, ground tests, bench tests, and correlated failures. Through these methods, an initial prior distribution can be constructed that would enable more certain reliability estimates given a limited number of system-level tests (see Figure 1).

Thus, the use of the Bayesian methodology and other sources of data in addition to infrequent flight tests can reduce uncertainty in the estimates of GBI reliability and help identify and fix failure modes more rapidly. This analytic process provides a pathway to help reduce risk and enable assessments of whether the

The reliability of each GBI component can be modeled with a two-parameter (α, β) distribution called a beta distribution. The beta distribution gives the probability density of a value x (here reliability) on the interval [0, 1]:

$$Prob\ (x) = \frac{x^{\alpha-1}(1-x)^{\beta-1}}{B(\alpha,\beta)},$$

where $\alpha > 0$, $\beta > 0$, and $B(\alpha, \beta)$ is the beta function (which serves as a normalizing constant to ensure that the total area under the density curve equals 1).

The choice of parameter values for the GBI beta prior distribution depends on the assumptions regarding the starting state of knowledge (see main text) and the number of GBI components (n). If the initial reliability of each component is assumed to be unknown, the initial values of the two parameters would be chosen to fit a uniform distribution between 0 and 1 (roughly, all values between 0 and 1 are equally likely). The actual parameter values, as derived by Redd and Reese [1], following Goodman [2], are:

$$\alpha = \frac{(2/3)^{1/n}-1}{1-(4/3)^{1/n}}, \beta = \alpha \cdot (2^{1/n}-1).$$

This starting state of knowledge is updated with new data using Bayes' Theorem, which shows that the posterior distribution is equal to the prior times a factor that is dependent on the data. If we make the simplifying assumption that each trial is successful with the same probability, this implies that the number of successful trials for a particular component can be described by a binomial distribution. Using the beta distribution to describe the prior distribution and a binomial distribution to describe the data leads to a posterior distribution that is also a beta distribution.

Sources:

[1] Redd, T. and S. Reese, Brigham Young University, personal communication.

[2] Goodman, L., "The Variance of the Product of K Random Variables," *Journal of the American Statistical Association*, Vol. 57 No. 297, March 1962.

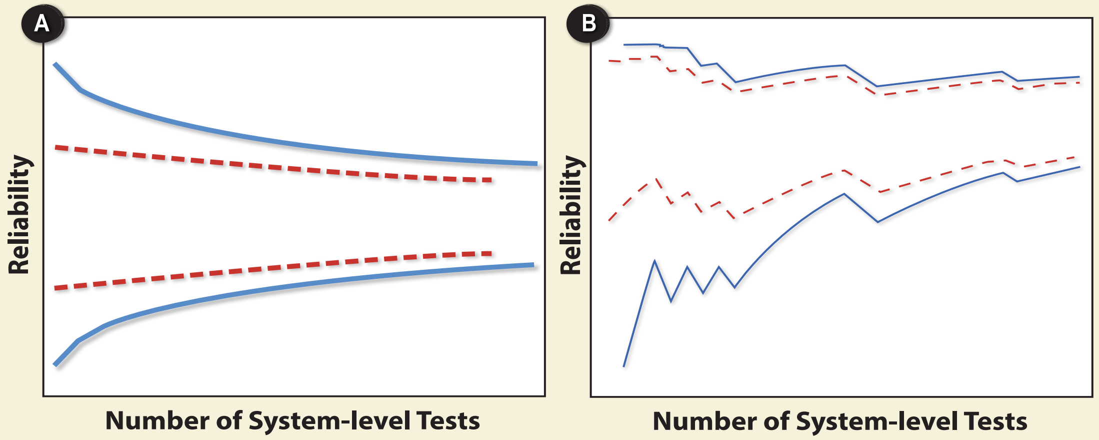

Ⓐ provides a notional depiction of the reduction in uncertainty. The solid lines represent estimates based only on data from the system-level tests, and the dashed lines represent estimates that also include other suitable sources of data using the Bayesian methodology. Ⓑ shows a similar set of results obtained for a simple problem* and contains noise and asymmetric behavior (the lower bound increases more rapidly than the upper bound decreases due to the nature of the problem). In these illustrations, the uncertainty estimates for a limited number of system-level tests are smaller when one incorporates a starting state of knowledge (based, for example, on component-level testing and/or modeling and simulation) prior to the onset of system-level testing. Over a sufficient number of system-level tests, the solid and dashed lines approach similar estimates of uncertainty. However, in practice, conducting a sufficient number of tests might not be possible, and thus the inclusion of other suitable sources of data can provide smaller uncertainty for a limited number of system-level tests.

* For Ⓑ, results were randomly generated for 30 "tests" (observation of 1 for "success" and 0 for "fail") with probability of success equal to 0.75. For the estimates using the Bayesian methodology, the starting prior had 90 percent probability that the probability of success was between 0.5 and 0.95. The solid lines are the upper and lower classical 90 percent confidence interval. The dashed lines are the upper and lower bound of a Bayesian credible interval. These intervals, confidence and credible, correspond to our "uncertainty estimates." *[Acknowledgment: IDA Research Staff Member Alyson Wilson]*

Figure 1. Reduction in Uncertainty of Reliability Estimates with Increasing Number of System-level Tests

GBIs are meeting reliability requirements. In addition, projections of future reliability can also be made by employing additional assumptions involving, for example, the degradation of GBI reliability due to age, the occurrence of unforeseen failure modes, and the effectiveness of repairs.

IDA researchers are also employing Bayesian methods to estimate reliability for other military systems. For example, since there are many similarities among the vehicles in the Stryker family, we used Bayesian hierarchical models to estimate the reliability of the family based on data obtained for the different vehicles and across both developmental and operational tests. These models let the data determine the appropriate weighting of information across vehicle variants and test phases. The combination of information improves estimation and reduces uncertainty.

IDA is also examining the use of Bayesian methods for improving the estimation of reliability for the Ground Combat Vehicle. These methods are particularly well suited to integrating heterogeneous information sources. We are starting by considering information obtained from modeling and simulation to develop prior distributions that summarize what we anticipate seeing in subsequent testing. This will help with test design and potentially reduce uncertainty in the reliability estimates.

For many military systems, conducting a sufficient amount of system-level testing in operational environments to provide reliability estimates with low uncertainty can be quite expensive. The use of Bayesian methods can allow researchers to design test strategies involving a limited number of tests and to incorporate other sources of data—such as modeling and simulation, ground tests, and bench tests—that could help reduce uncertainty and risk.

Dr. Buontempo is an Assistant Director in IDA's System Evaluation Division. He holds a doctorate in physical chemistry from the University of Chicago.

PROMISE, REALITY, AND LIMITATIONS OF SOFTWARE-DEFINED RADIOS

Lawrence Goeller and Patricia Bronson

The Problem

The goal of using software-defined radios to provide interoperable communications equipment across U.S. military forces has proven to be technologically challenging to implement.

A long-standing Department of Defense goal has been to achieve interoperability. Some of the most challenging problems in this area involve the acquisition of interoperable communications equipment. Since the Services depend on radios that have been acquired over decades, it is infeasible to start over and replace all of them with new, more interoperable systems; the transition would have to be gradual.

The vision of an SDR—a communications device that can synthesize any waveform, including new waveforms, solely by changes to the software—has been, and remains, highly attractive... but, it seems evident, we are not there yet.

The Ups and Downs of the Joint Tactical Radio System (JTRS)

Considering the constraints, and given the rapid progress in the 1980s and 1990s in powerful, inexpensive Personal Computer (PC) technology, one way forward seemed especially promising: if radios could migrate from the traditional mixer-based modulation technique to one based on a programmable microprocessor, the *waveforms* could be loaded and run on the underlying hardware much like word processing and spreadsheet programs were run on a PC (Figure 1). Such devices were referred to as *software-defined radios* (SDRs). These new SDRs needed to be, first and foremost, backward-compatible with the existing, fielded radios that each Service had previously purchased. Second, they had to

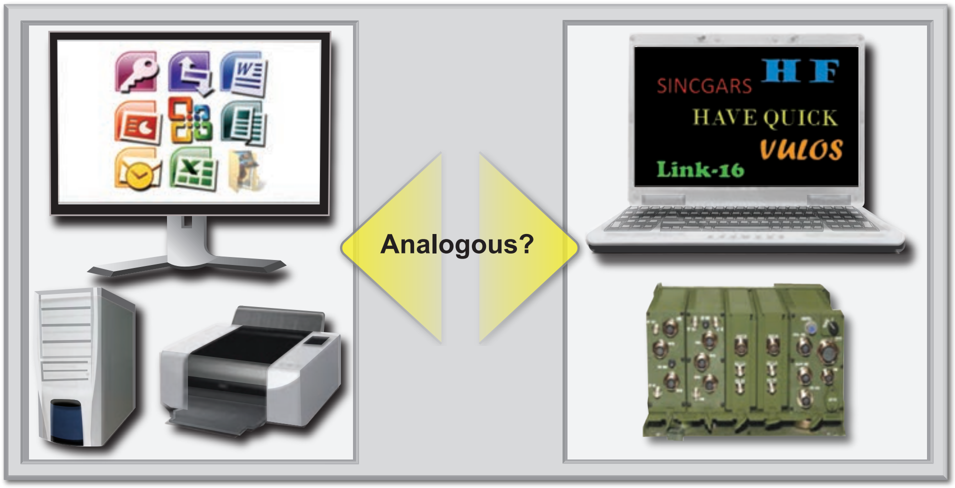

Figure 1. SDR Envisioned as a PC Running "Waveforms" Instead of Applications

be programmable to "run" at least some of the waveforms of different Services, allowing interoperability without the presumed excessive cost for additional dedicated circuitry for each additional waveform. Third, this approach offered the promise that new and better waveforms could be readily ported to SDR equipment after it had been acquired and fielded. In 1997, the Joint Tactical Radio System (JTRS) SDR program was born, and the Office of the Secretary of Defense directed that, after a certain date, the Services would not be permitted to acquire any non-JTRS radio without a waiver.

Unfortunately, the JTRS program has struggled. Despite a number of program changes, requirements easing, management restructures, and the investment of more than a billion dollars, the Ground Mobile Radio (GMR) version of the JTRS radio was terminated in 2011 after declaring a Nunn-McCurdy breach.[1] The Airborne Maritime Fixed (AMF) JTRS prime contract was ended a year later, with no fielded hardware produced. The Handheld, Manpack, and Small Form Fit (HMS) JTRS and Multifunctional Information Distribution System (MIDS) JTRS programs have produced some hardware as of this writing, and other parts of the program are continuing, but, on the whole, the program has been a disappointment. The question is: are there underlying technical issues that fundamentally preclude success, or can the promise of the SDR still be met with some combination of newer technology and different management? This is the question that the Office of the Under Secretary of Defense for Acquisition, Technology and Logistics' Performance Assessments and Root Cause Analyses (PARCA) organization asked IDA to investigate.

The Basics of Radios and Waveforms

With few exceptions, the function of any radio is to use a baseband signal to modulate a carrier wave and then amplify and transmit the resulting signal—and then reverse the process at the receiving end. The baseband signal could be analog or digital, and the carrier wave can be at almost any frequency. Since the modulation/demodulation process has traditionally been performed via non-linear mixing and amplification of the baseband and carrier signals, such devices are sometimes called *mixers*. Other operations may be performed on the baseband signal before modulation (such as channel encoding, if the signal is digital) and after modulation (such as multiplexing, or combining, this signal with others transmitted at the same time). Collectively, the combination of channel encoding, modulation, multiplexing, and other processes is referred to as the *waveform.* The waveform may be thought of as the information the receiver needs to know about the transmitted signal in order to correctly recover the baseband information.

At the heart of most modern radios is the replacement of the mixer, and often other components, with some sort of digital processing device.

---

[1] The Nunn–McCurdy Amendment requires the Secretary of Defense to notify the Congress if the cost per unit of a program grows more than 15 percent beyond what was originally estimated and calls for the termination of programs with total cost growth greater than 25 percent.

This is a revolutionary change; the task of modulating a carrier has been changed to "synthesizing" what the modulated signal *would have looked like* if conventional hardware had been used via what is, in practice, a digital signal processor. (Both mixer-based and microprocessor-based radios are capable of transmitting and receiving either analog or digital waveforms.) These devices are not necessarily SDRs, however; in many designs, the digital logic circuit synthesizes only one or a few waveforms, and these cannot be modified or added to, once fabricated. For a logic-based radio to be considered an SDR, the waveforms it synthesizes must be controlled externally by a software program; hence, the name.

In the 1980s and early 1990s, it was envisioned that a so-called general purpose processor (GPP), such as an Intel Pentium, would be able to perform the waveform synthesis. It rapidly became apparent, however, that it could not, in practice, support such tasks. To make a long story short, GPPs were found to be unsuitable for the real-time and multiplication-intensive requirements associated with signal processing. Even in PCs, such operations are performed by digital signal processors (DSPs) located in math coprocessors or graphics cards. Although many military radio systems were developed that synthesize waveforms via DSP-based circuits, DSPs are not nearly as flexible as GPPs. They are designed to perform a specific, limited set of operations very quickly, but it is in general difficult to add new algorithms to them after fabrication. This constrains their utility to the SDR vision. This gap between the flexibility of a GPP and the operational performance of a DSP was, to many, an indication that synthesizing waveforms differed *technically* from running applications. Nonetheless, the attractiveness of the SDR concept remained.

Field Programmable Gate Arrays

The Defense Advanced Research Projects Agency (DARPA) came up with a compromise solution: the use of hardware devices called field programmable gate arrays (FPGAs) to synthesize the waveforms. FPGAs, a new technology at the time, are a third kind of processing architecture (although like both GPPs and DSPs, they are commonly composed of Complementary Metal-Oxide Semiconductor (CMOS) transistor technology). An FPGA is composed of a large number of small identical logic elements, each typically consisting of a small lookup table and some memory. These logic blocks do not connect directly to each other as in a DSP, but rather to one or more of a series of parallel wires called *routing channels* that run around all sides of each block (see Figure 2). Connections between the internal logic of each block and the wires of each routing channel are controlled by transistors that can be opened or closed on command; there are more of these transistors in each region where horizontal and vertical routing channels cross each other. By opening and closing selected transistors, the small logic blocks can be connected in such a way that they can emulate virtually any other logic element, from a simple digital logic gate to a block of hardware memory to a multiplier or even a small microprocessor. Best of all, these connecting transistors can be commanded by a software program.

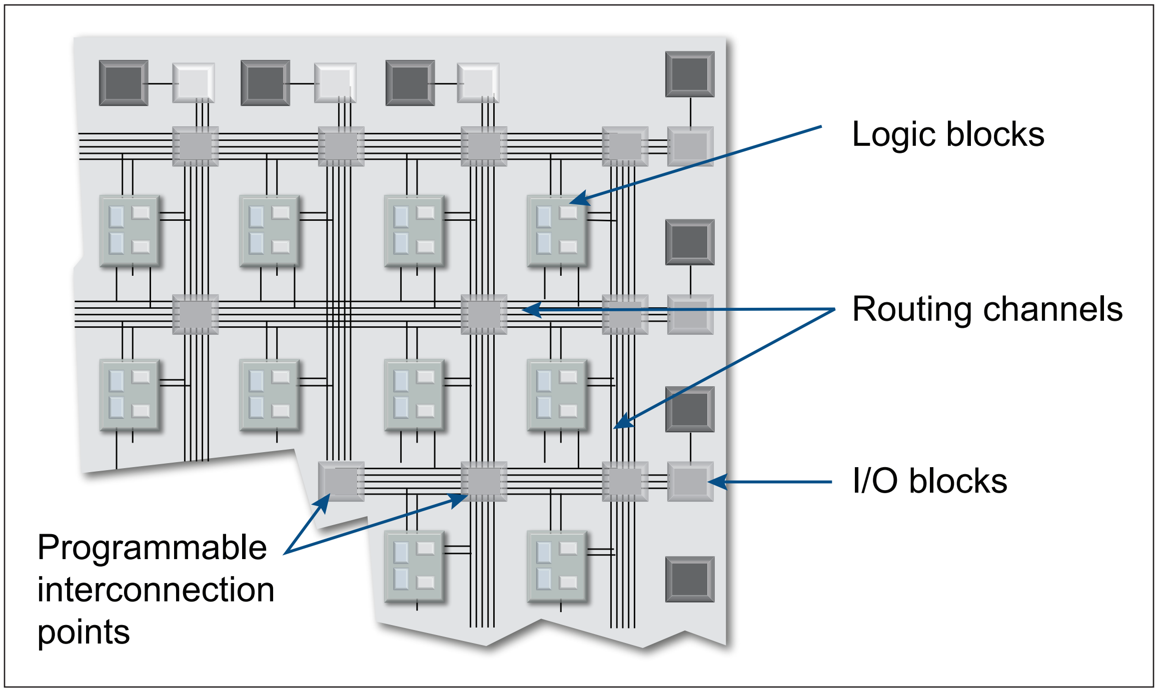

Figure 2. Notional FPGA Logic Architecture

Software packages called *hardware description languages* (HDLs) that look a lot like C or Ada could be used to “write circuits” on a given FPGA; this code could then, in principle, be recompiled to operate on different FPGAs, thus effectively offering the ability to port waveforms from one device to another. That is, the FPGA offers the SDR community a technology with comparable performance to a DSP, and yet is reprogrammable—perhaps more accurately, reconfigurable—like a GPP. After DARPA’s SpeakEasy program was considered to be a great success, the FPGA-based JTRS program was initiated.

Unfortunately, it quickly became apparent—to the engineers, if not the larger acquisition community—that the “sea of configurable logic blocks” approach was also technically inadequate. FPGAs require many times the number of transistors to perform the same function as a dedicated DSP, and each of these transistors draws electrical power and produces heat that must be dissipated. FPGA-based SDRs were found to draw a large amount of power—a problem for battery-powered tactical radios—and produced a large amount of heat. At some point, it also became apparent that, despite the superficial similarity of HDL languages to portable, high-level languages like C, they were in fact quite different. The portability of GPP-based languages depends on a certain similarity of the architecture in all cases and on the ability of the compilers to find some way, if not the optimal way, of implementing each program statement. HDL compilers, by their nature, are much more dependent on the underlying hardware. The JTRS community responded by trying to develop an intermediate layer between the waveform software and the FPGA-based hardware called Software Communications Architecture (SCA)—essentially an operating system and middleware. This was uncharted territory. When coupled with the real-time constraints of signal processing, it became clear that the amount of effort to port a waveform from one device to another was significantly greater than hoped.

The technology continued to evolve, however. New generations of devices that were still called FPGAs became available, but these new devices not only were larger in terms of the number of configurable logic blocks they contained, they also contained many pieces of intellectual property (IP). These IP elements were composed of anything from dedicated blocks of fast memory to dedicated hardware multipliers to microprocessors (see Figure 3). The "catch" is that, while these IP elements run faster and draw less power than strings of configured logic blocks, they are themselves essentially DSP elements, with the corresponding lack of flexibility. In fact, the trend in FPGA design has been for each vendor to create a wide diversity of products, each aiming at a different specialized market. And therein lies the problem: the key idea of the SDR is to rely on a single underlying card to generate any one of dozens of different waveforms, including some not yet invented. What we are finding is that FPGAs are not, apparently, a one-size-fits-all technology; in fact, they are almost the opposite.

But what about the flexibility-vs.-power consumption trade space of the newer generations of FPGAs, sometimes referred to as "System on a Chip" or SOC? Alas, the JTRS program tried these as well; the AMF JTRS design, when it ended, was based on a Xilinx Virtex 5 FPGA with hundreds of embedded IP elements. Although a number of different waveforms were

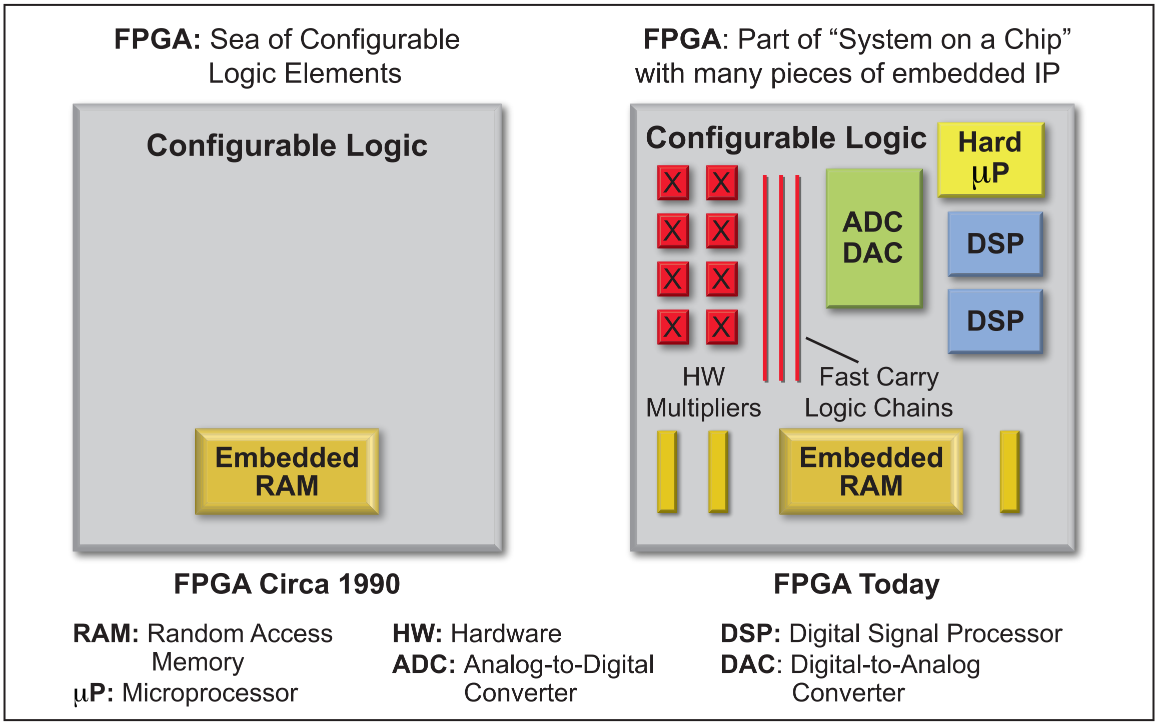

Figure 3. FPGA Architectures, Then and Now

demonstrated on the development hardware, the power consumption remained large, and porting waveforms remained challenging.

What about the future? As logic element sizes shrink, both the power that the transistors require to change states (the so-called *active power*) and the time it takes them to do it decrease. Newer designs come with clever ways to remove the clock signal from transistors that are not in use. There are even new manufacturing technologies in which the transistors are grown "vertically" on the substrate as opposed to the planar approach of CMOS today; this offers the opportunity to lower power consumption even further. Can these improvements lead to a viable approach to the SDR? Perhaps, but every advance in this trade space seems to come with additional constraints. For example, as transistors get smaller, the active power consumption goes down, but the so-called static power goes up; the transistors essentially "leak" as long as they are powered up. Further, it appears that there is a tradeoff between this static power loss in small transistors and switching speed. It must also be noted that none of these new technological approaches addresses the problem of waveform portability, which is a primary driver in the preference for an SDR over, for example, a radio containing a different DSP (or different parts of a larger DSP) for each supported waveform.

Conclusion

The vision of an SDR—a communications device that can synthesize any waveform, including new waveforms, solely by changes to the software—has been, and remains, highly attractive. Although there have been some limited successes in the decade and a half since the initiation of the JTRS program, it seems clear that the task of creating an SDR that is both power-efficient and flexible enough to support new waveforms remains challenging. This is in part due to technical reasons: specifically, the difficulty of producing hardware that is both powerful enough to perform signal processing tasks in real time *and* flexible enough to support growth to future, undetermined waveforms. The original vision of the SDR might be technically achievable, but, it seems evident, we are not there yet.

Dr. Goeller is a Research Staff Member in IDA's Cost Analysis and Research Division. He holds a doctorate in physics from Rice University.

Dr. Bronson is a Research Staff Member in IDA's Cost Analysis and Research Division. She holds a doctorate in applied physics from Old Dominion University.

IMPLICATIONS OF CONTRACTOR WORKING CAPITAL ON CONTRACT PRICING AND FINANCING

Scot Arnold

The Problem

Designing and negotiating fixed-price, sole-source procurement contracts that motivate desired contract outcomes require an understanding of how fee structures and contract financing influence a contractor's return on equity.

The average customer trying to buy a common commercial product has the benefit of market-based intermediation for discovering prices or accessing purchase financing, such as leasing. In contrast, most large government acquisitions are done through negotiated sole-source firm fixed-price contracts where prices are based on the estimated cost of the item to be procured. Although the contractor often has an informational advantage on the item's cost, the government has an advantage by its ability to finance its purchases at a lower cost than all private financing.

Government procurement price negotiation is backed up—ideally—with a mutual understanding of what the item should cost to produce and government profit policy to guide fee structure and contract financing decisions.

Fee Structure and Contract Financing: Complementary Profit Policy Levers

Government procurement price negotiation is backed up—ideally—with a mutual understanding of what the item should cost to produce and government profit policy to guide fee structure and contract financing decisions. Profit policy has two main levers: contract fee (stated as a percentage of cost) and contract financing. The weighted guidelines method outlined in section 215.404-71 of the Defense Federal Acquisition Regulation Supplement (DFARS) provides a structured approach for determining the fee that should be paid to a contractor based on the expected effort and level of financial risk. Financial risk is based on the type of contract (fixed-price or cost-reimbursable), amount and type of contractor capital required (working vs. facilities), and the source of contract financing. The guidelines distinguish between different levels of government contract financing: private financing from the contractor, progress payments, and performance-based payments. Progress payments cover up to 80 percent of the incurred costs for partially completed work that is invoiced on a recurring short-term basis. Even with contract financing, a contract has a growing working capital balance that the government pays upon completion.

Profit policy has a dual role: to motivate contractor performance and to encourage and compensate contractors for putting capital at risk. Designing contracts to motivate desired outcomes is fraught with agency problems, such as adverse selection and moral hazard. The profit policy and other contracting rules aim to provide officials with tools to augment negotiations on cost and requirements. It is also important to maintain the long-term health of the defense industrial base because it is paramount to implementing defense policy. A "Goldilocks"—or "just right"—policy would provide sufficient, but not excessive, compensation for defense industry investment.

Why are fee structures and contract financing so important, and how are they related? The contract fee less non-reimbursable expenses is the contractor's profit. The contractor's shareholder value of that profit depends on how much of its equity was required to fund the contract execution. The important metric is the contractor's return on equity (ROE). Contractors can increase ROE by increasing profits, which can be very difficult to do, or by reducing the amount of equity needed. Debt is a common equity substitute, but it comes at a cost and with risks. The act of substituting debt for equity is to create leverage that effectively boosts ROE.

What Is the Relationship Between Leverage and Margin for a Firm?

Debt provides leverage that effectively boosts the profitability of a company when compared to its peers without debt financing. Technically, government-provided contract financing is a non-debt liability, but it provides the same leverage effect on equity returns. This is because contract financing allows a company to generate cash flows with much less capital than a firm without contract financing.

The Dupont formula can be used to better understand the relationship between the firm's ROE (and ultimately the firm's value to shareholders), its profit margin, and its capital structure (i.e., how much debt it holds). The formula is:

ROE = Return on Sales x Asset Turnover x Asset-to-Equity Ratio.[1]

Return on sales (ROS) is the profit margin ratio, while the latter two terms in the equation measure asset efficiency. Fewer assets with the same revenue improve efficiency, as does more debt with the same asset level. Figure 1 shows how a firm, with the same profit margin and asset turnover but two different capital structures, can have vastly different ROEs. On the right hand side of the chart, the firm has a high debt ratio and enjoys much higher ROE but at a greater risk of bankruptcy than it does on the left hand side with a low debt ratio.[2] This is because debt acts as a fixed cost and can lead to bankruptcy if sales drop too much.

But government contract financing is not debt and does not pose the

---

[1] In terms of definitions, ROE (profit/equity) = ROS (profit/sales) x Asset Turnover (sales/assets) x Asset-to-Equity Ratio (assets/equity).

[2] Readers who own a house are aware of the concept of how debt provides leverage to boost equity returns. A house that is sold for $110,000 a year after being purchased for $100,000 has a before tax ROS of 10%. If the owner borrowed 80% of the initial transaction, he now has a return on equity of 50%. On the other hand, anyone who remembers the recent great recession should also realize that debt exposes the property owner to bankruptcy or foreclosure risk.

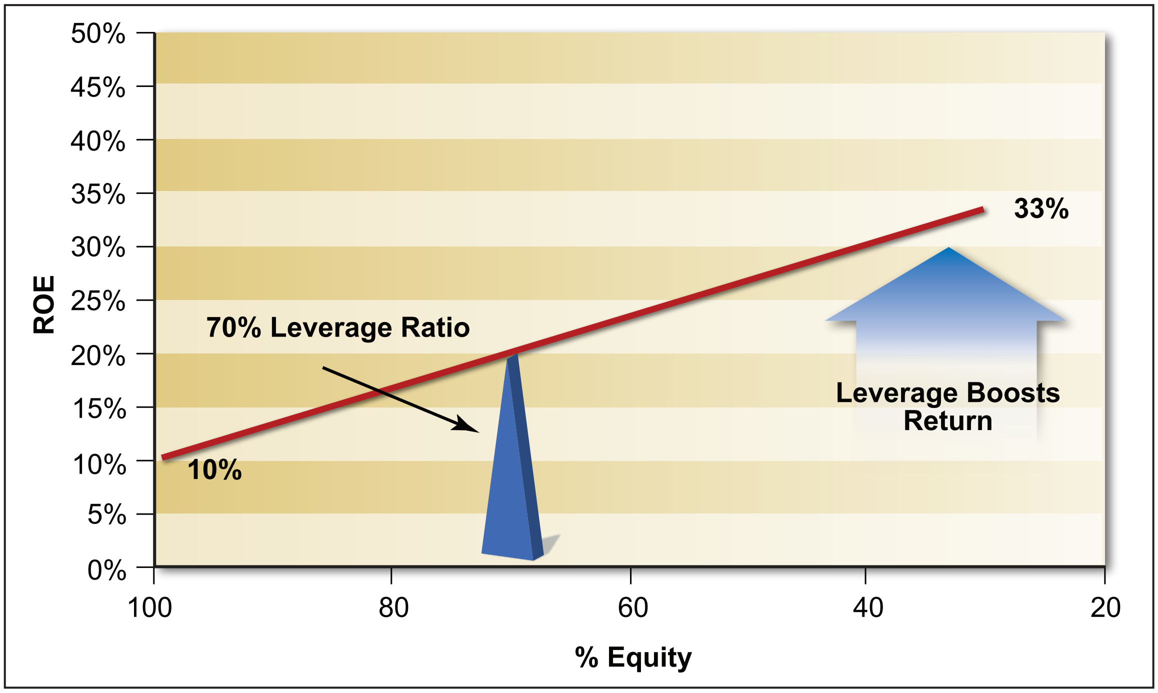

Figure 1. Comparison of ROEs with Two Different Capital Structures

same sort of risk to a contractor that uses it. Consequently, contract financing has the return-boosting leverage for equity but without the bankruptcy risk of debt—the implication being that this is due to the firm's ability to execute the contract without investing as much working capital as would otherwise be required. For example, assume a contract that costs $1 million to execute over the course of one year has a fee of 10 percent of cost, or about 9 percent ROS. If the firm had to borrow or use equity to pay for all of the materials, labor, and overhead, the ROE would be close to the ROS. But contract financing, such as bi-weekly progress payments, boosts the ROE for the same contract to about 68 percent.

Unlike most defense contractors, commercial industrials invest equity and debt capital into new plants, tooling, product design, and even dealer and customer financing in order to sell products. They might invest billions of dollars in a product before selling the first unit, but they generally have high profit margins that ultimately cover their cost of capital. A retailer by contrast, might not even own its inventory; rather, it uses customer and vendor cash to finance the cost of sales. Consequently, successful retailers have much less equity invested than industrials. Thus both types of firms can yield high returns with vastly different capital requirements.

Defense contractors share characteristics of both sectors: their products require large investments, but they can use customer funds to minimize equity and debt requirements. Most defense contractors have margins that are lower than commercial industrial firms but higher than pure retailers, and they have access to considerable government financing. In fact, contract finance can be so favorable that, as Christopher Kubasik, former Chief Financial Officer of Lockheed, said, "working capital will continue to be a great contributor to our cash."[3] The implication is that Lockheed's contract finance is providing a high level of cash for the company, whereas for most manufacturing companies, working capital is not a source—but a sink—for a firm's cash level.

The top defense contractors have exploited this financing strategy successfully, particularly over the past 10 years, as shown in Figure 2. This chart shows the ROE and defense industry average cost of equity since 2004.[4] During the past decade, prime contractors have easily made returns that exceeded their cost of equity. Yet at times, the defense industry will try to argue that single- or low- to-mid- double-digit percentage of cost contract fees provide insufficient returns when compared to other industries.[5] Focusing on margins and ignoring the power of contract financing provided by the government is misleading. Defense contractors have provided excellent shareholder returns, even though many of their contracts have single- or low double-digit margin rates.

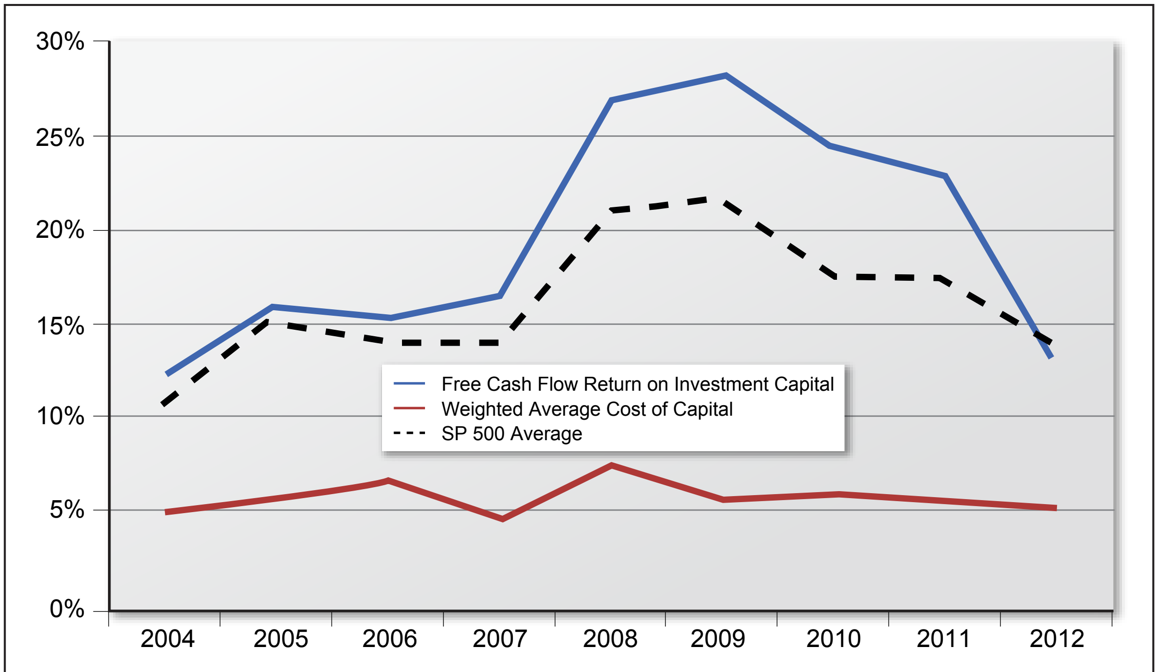

Data Source: S&P Capital IQ

Figure 2. Defense Contractor ROE vs. Cost of Equity Since 2004

---

[3] Christopher Kubasik made this statement during the third quarter 2006 earnings conference call. Working capital is invoiced sales that have not yet been paid, plus the cost of finished goods and work-in-process inventory, less what is owed suppliers.

[4] The industry is defined here as the top prime contractors: Lockheed Martin, General Dynamics, Raytheon, and Northrop Grumman. Boeing is excluded because it derives half of its revenue and earnings from commercial aircraft.

[5] Aerospace Industries Association, "Assessing the Health of the Defense Industry," 2005.

What Is the Relative Benefit to the Government of Using Financing or Fee?

Organizations must understand their relative advantages when it comes to financing. Just as large volume retailers use vendor, and in some cases customer, financing as leverage to boost earnings on low margins, the U.S. government has crafted financing policy to allow contractors to use contract finance at the cost of receiving lower profit margins.

At the top level, the Office of Management and Budget dictates that the discount rates used in financial decisions reflect the government's opportunity cost of capital. This policy ensures that decisions involving financing, such as a long-term lease, are biased away from using private financing. The DFARS also appears to provide a bias toward using government contract financing in lieu of contractor financing.

Consider the problem of financing working capital, which is the funding required to cover the contractor's operating cost until the sale is invoiced and paid. Without progress payments, contractors require enough working capital to cover the cost of executing the contract until it is finished and payment is received. Progress payments allow the contractor to receive partial payment every two weeks, and drastically cut the amount of capital the contractor must put at risk. Because government contract financing is a partial payment system, contractors must still fund at least 20 percent of the contract cost. This means that the longer the contract, even with progress payments, the longer the contractor must tie up its working capital and the higher the financing cost.

The relationship between the amount of fee required in order to cover financing cost and the progress payment rate is shown in Figure 3. Each line shows the minimum fee given the progress payment rate and contract term lasting from one to five years. In this case, the contractor is assumed to have a cost of equity of 10 percent. The slopes of each line represent the marginal amount of fee reduction possible for a unit increase in the payment rate.

The DFARS profit policy provides two guidelines to cover contractor financing cost. One is the working capital adjustment, which is tied to the progress payment rate, the length of the contract, and the prevailing interest rate; the other is a fee to cover "contract risk." Figure 3 shows the projected working capital adjustment, which assumes a 6 percent interest rate, as a blue trapezoid. The wedge thickness is bounded by contracts lasting 12 months on the bottom and 60 months on the top. Only at high progress payment rates does the working capital adjustment cover contractor costs. Contractors choosing to use their own capital are not fully compensated by the policy guideline fees.

The DFARS guidelines for "contract risk" also indicate that the government is biased against using private capital. For contracts without contract financing, the normal contract risk fee should be only two percentage points higher

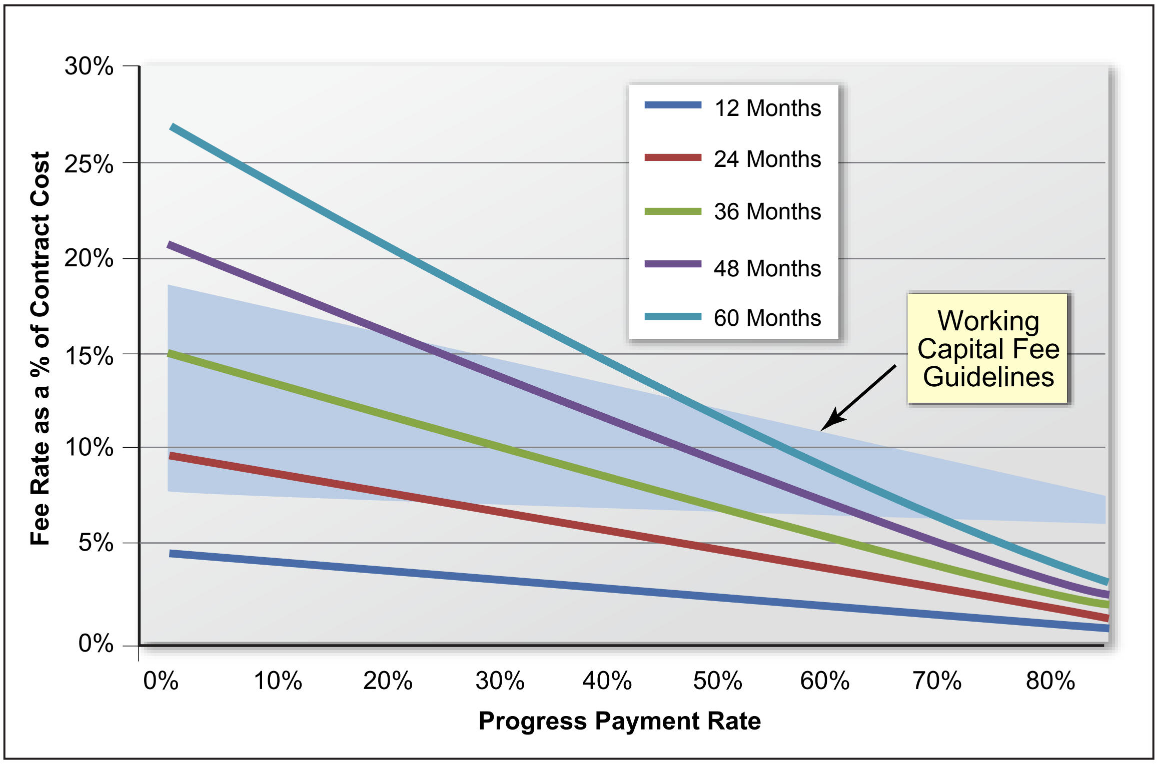

Figure 3. Minimum Working Capital Fee vs. Progress Payment Rate (Lines) and the Working Capital Adjustment Guidelines in DFARS 215.404-71-3 (blue trapezoid)

than for contracts with progress payments. One can infer the contractor's financing cost from Figure 3. For a given contract term, it is the difference between the fee at 80 percent progress payments and at zero percent progress payments; in all cases, it appears to be well in excess of two percentage points. Thus, contractors, particularly with long-term contracts, should prefer to use government contract financing.

Conclusion

Generally the government has a difficult problem negotiating cost-efficient sole-source procurement with fixed-price contracts. The contractor has an informational advantage on the cost of the contract that it may be able to exploit to gain higher profits than it might expect in a competitive market. It appears, however, that the government has developed a sound fee policy when it comes to contract finance. Clearly the government should use its long-established low relative financing cost to its advantage and lower its contracting costs.

This does not mean that the government should completely finance fixed-price contracts the way it does cost-reimbursable ones. By requiring the contractor to put some capital at risk, the contract has an embedded incentive to be completed as soon as possible. Furthermore, holding back some of the payment helps provide some surety protection that the contractor will complete the contract under its ceiling. But in spite of the low government financing cost, certain contractor financing structures have allure for government agencies that are underfunded. These structures often have high implied interest rates but allow agencies to make acquisitions without direct congressional appropriations. The government should maintain its discipline on this front and consider ways to better coordinate fee and contract financing policy across all agencies.

Dr. Arnold is a Research Staff Member in IDA's Cost Analysis and Research Division. He holds a doctorate in polymers from the Massachusetts Institute of Technology and an MBA from the University of Michigan.

THE MECHANISMS AND VALUE OF COMPETITION

James Dominy, Scot Arnold, Colin Doyle, Brandon Gould,
Bruce Harmon, Susan Rose, Robert Thomas, and Karen Tyson

The Problem

The incentives created by a competitive environment for acquisition of large, complex DoD systems present challenges for acquisition professionals.

Competition is a cornerstone of the federal government's acquisition processes. In an open marketplace, one would expect that effective competition would drive down the price of goods and services. However, Major Defense Acquisition Programs (MDAPs) are large, complex, and normally built to specific Department of Defense (DoD) requirements. Do competitions for MDAPs yield significant benefits? If so, what are those benefits? Drawing on previous conceptual and empirical analyses of competition, IDA examined how competition operates for MDAPs in an effort to aid acquisition professionals in understanding the incentives created by a competitive environment and the benefits the government can be expected to obtain.

Drawing on previous conceptual and empirical analyses of competition, IDA examined how competition operates for MDAPs in an effort to aid acquisition professionals in understanding the incentives created by a competitive environment and the benefits the government can be expected to obtain.

Competition for a Weapon System Franchise

Most of DoD's MDAPs are awarded under competitive conditions. Many of these are of the general nature of a franchise: DoD awards a contract to a firm for engineering and manufacturing development (EMD) of a system, and, upon successful completion of EMD, awards the firm a series of fixed-price contracts for serial production. In a pure form, then, a firm needs to win only one competition in order to lock in work that can extend years or even decades into the future.

But is this method of competition effective? After all, normally, most of the cost of acquiring a weapon system is in the production phase; those production costs are imperfectly known at the start of the EMD phase, and the competition occurs sometimes years before the first production lots are priced. There is no direct evidence that such competitions drive the price of the system down to cost (where cost includes a fair return on the contractor's capital). However, competitions for MDAPs do appear to provide a significant value to the government.

The mechanism by which this value is obtained is described in a seminal paper by William Rogerson.[1] Fixed-price production contracts incentivize the contractor to invest in cost savings methods in order to drive its cost below the price negotiated

---

[1] William P. Rogerson, "Profit Regulation of Defense Contractors and Prizes for Innovation," *Journal of Political Economy* 97, no. 6 (December 1989): 1284–1305.

in the contract. Although the government will eventually discover these cost savings and reduce its price offer on later contracts, the serial nature of the contracting process introduces a "regulatory lag"—subsequent contracts are negotiated before actual costs are known on predecessor contracts. Therefore, the contractor is able to retain the savings generated by its cost reduction efforts. This allows the firm to obtain prices high enough during the production phase to allow it to earn a return on invested capital greater than it could obtain in a similarly risky alternative employment; in economic terms, the supplier earns "rents."

Rogerson views these production-phase rents as "prizes" for firms to provide innovative solutions to DoD requirements. These "prizes" of potential returns above the competitive level in the production phase provide a strong incentive for competitors to propose innovative solutions in the EMD competition and then deliver the solution that DoD wants. Thus, the incentives help reduce (but not eliminate) the enormous governance problems associated with managing the development and production of complex weapon systems: the contractor has a strong profit incentive to see that the government receives a product that meets its definition of success.

In order to maintain the effectiveness of this incentive, however, the government must retain some ability to put the potential stream of rents at risk. The prize is not awarded when the EMD contract is signed, but rather earned lot by lot during the production phase; the last of the prize is not awarded until the last unit has been delivered and fully paid for. If the contractor is not able to design a system that meets the government's requirements in terms of cost, capability, and performance, the program is subject to termination prior to production. If cost, quality, and/or schedule cannot be maintained in production, the government has the option to terminate the program early or reduce quantities, thus limiting the opportunity to earn rents. The government also has other tools to place potential rents at risk, such as reintroducing competition through mechanisms such as dual sourcing production.

Competition for a Single Design-Build System

DoD may also develop and procure a system under a single contract, a method we term a "single development-build" program. Under this method, firms compete for a single contract to both develop and build a weapon system. This structure is normally utilized when the production phase involves only a limited number of systems, such as satellites.

The benefits of competition for single development-build programs appear to be more limited than for competitions in franchise acquisitions because the incentives are different. Since there is only one contract award, it is more difficult to hold potential contractor rents at risk. This, in turn, reduces the incentives for the contractor to be responsive to the government's requirements, and to achieve program cost and schedule objectives. Thus, the government has to rely more heavily on other governance tools in order to attempt to achieve program success.

Conclusion

Competition for complex systems that require development and serial production offers an important tool to obtain the military capabilities that DoD requires. The potential rents available during the production phase provide a strong incentive for firms to remain in the defense market, to offer innovative solutions, and to remain responsive to the government's requirements during the (often extended) development and production period. However, the strength of this incentive effect is proportional to the ability to hold the production-phase rents at risk.

Mr. Dominy is a Research Staff Member in IDA's Cost Analysis and Research Division. He holds a master of arts degree in public administration from the University of Iowa and a master of science degree in management from the Massachusetts Institute of Technology.

Dr. Doyle is a Research Staff Member in IDA's Strategy, Forces, and Resources Division. He holds a doctorate in economics from the University of Maryland-College Park.

Mr. Gould is a Research Associate in IDA's Cost Analysis and Research Division. He holds a master's degree in public policy from the University of Virginia's Frank Batten School of Leadership and Public Policy.

Dr. Harmon is an Adjunct Research Staff Member in IDA's Cost Analysis and Research Division. He holds a master's degree in economics from the University of Cincinnati.

Dr. Rose is a Research Staff Member in IDA's Cost Analysis and Research Division. She holds a doctorate in economics from The Ohio State University.

Dr. Thomas is an Adjunct Research Staff Member in IDA's Cost Analysis and Research Division. He holds a doctorate in economics from the University of Pennsylvania.

Dr. Tyson was a former Research Staff Member in IDA's Cost Analysis and Research Division.

INITIATION AND EARLY MANAGEMENT OF ACQUISITION PROGRAMS

Royce Kneece and Gene Porter

The Problem

Unresolved problems in the early stages of Major Defense Acquisition Programs (MDAPs) frequently lead to negative program outcomes.

History tells us that problems in the initiation and early management of Major Defense Acquisition Programs (MDAPs) lead frequently to even greater problems in the later stages of these programs. Addressing issues in the early stages of acquisition can forestall such problems and significant cost growth.

Few topics in the DoD have generated as much controversy as the process for defining requirements for acquisition programs. Force size and mix and the manning and equipping of forces to perform the missions assigned to them within an acceptable level of risk ... are what we loosely call "requirements."

Defining Requirements for Acquisition Programs

Few topics in the Department of Defense (DoD) have generated as much controversy over the years as the process for defining requirements for acquisition programs. Our military forces should be designed to perform a set of military missions that are broadly defined by civilian authorities. Given those missions, one job of military commanders is to state what capabilities are needed—force size and mix and the manning and equipping of forces—to perform the missions assigned to them within an acceptable level of risk. Those capabilities are what we loosely call "requirements."

However, determining such requirements is fraught with difficulties. The projected future national security environment is at best uncertain and, indeed (as history shows), is likely to change. Yet another layer of uncertainty is the future military capabilities of both friends and enemies in a future scenario involving threats to U.S. national security. Thus, Pentagon decision makers must both appropriately discern requirements and determine the best ways to meet them within resource constraints and with an acceptable level of risk. Military commanders have a natural tendency to err on the high side when it comes to specifying requirements for mission success, since no commander wants to fail, and the region of uncertainty in defining capabilities needed in order to be successful in performing a complex military mission is generally large. These are some of the reasons that independent civilian oversight is necessary in establishing requirements for acquisition programs in order to strike an appropriate balance between a requirement, its technical feasibility, and the resources required to obtain it. (By "civilian oversight" we mean the staff in the Office of the Secretary of Defense (OSD)—the members of that staff comprise both civilians and military personnel, but the leadership is, with few exceptions, civilian.)

For an acquisition program, these factors translate into a proposed weapon system's required performance, the technology available to attain that performance, and the cost to acquire it (and thus its "affordability," i.e., is it reasonable to assume that funds can be made available to pay the projected costs?). Another important factor in the undertaking of a new acquisition program is *risk*. Risk in this sense can be defined as the probability that an acquisition program will not succeed—that is, required performance will not be achieved or will be achievable only at unacceptable cost. The fielding of systems that do not perform as needed (thus endangering mission success) is one potentially bad outcome. Another is the cancellation of a program after the expenditure of substantial resources, because of performance shortfalls and/or cost growth. The consequences are both waste and a failure to provide needed capabilities to the forces. While perhaps not as bad as those two outcomes, paying too much for the performance obtained is also undesirable. All these poor outcomes have "opportunity costs"—the other benefits that could have been obtained with the resources wasted or extra resources needed by a problem program.

These factors drive complexity into the requirements process for acquisition programs. There must be an effective interface between the process by which military commanders determine their requirements and the civilian oversight function to ensure that acquisition programs are affordable and that the value received will be worth the cost (i.e., cost-effectiveness).

A key component of the current process is the Joint Staff's Joint Capabilities Integration and Development System (JCIDS), which supports the Joint Requirements Oversight Council (JROC) in its statutory responsibility to support the Chairman of the Joint Staff in his role of advising the president and the Secretary of Defense regarding military requirements. Broadly speaking, JCIDS seeks to determine future capability needs through analytical processes that identify "capability gaps." The analytical results are presented in a Capabilities-Based Assessment (CBA). Currently, CBAs are normally performed by the Military Service that sponsors a proposed new acquisition program. Once gaps have been identified within a capability area, they are prioritized and assessed for potential solutions. If the best solution is deemed to be a new MDAP (i.e., a "new start"), the Service brings the proposal and the evidence supporting it to the JROC for approval. Upon obtaining JROC concurrence to start a new program, the sponsoring DoD Component presents the proposal to OSD for approval via what is known as a "Materiel Development Decision."

The JCIDS process has been the subject of much criticism, especially by the Government Accountability Office (GAO) and the Congress. As a result, the Congress has added provisions to the U.S. Code several times to strengthen JCIDS and the JROC.

A 2011 IDA paper recommended that an analytically based process, not overly dependent on Component analytical support, be used in conjunction with JCIDS. The recommended process would be:

- focused on the Secretary's priorities
- independent of sponsoring DoD Components
- focused on a clearly defined span of programs within capability/mission area (i.e., portfolio-based), and
- adequately resourced.

Acquisition Program Risks

Risks in acquisition programs are usually characterized as encompassing three aspects—cost risk, schedule risk, and performance risk. All three are closely related. Another important risk less frequently cited is affordability risk—closely related but distinct from cost risk. We will discuss these areas of risk in greater detail, keeping in mind that each area being discussed applies (in varying degrees) to the other areas of risk. Because of the close linkage, we discuss affordability risk immediately after cost risk.

Cost Risk

Cost risk is perhaps the most widely appreciated of these concerns. Cost overruns in DoD acquisition programs have not been uncommon.[1]

Continuing concerns about cost overruns helped motivate the Weapon Systems Acquisition Reform Act (WSARA) in 2009, which, among a number of provisions, required the establishment in DoD of an office dedicated solely to determining the "root causes" of acquisition program cost overruns—the office of the Director, Performance Assessments and Root Cause Analyses (PARCA). PARCA is required by law to perform root cause analyses for any program experiencing a "Nunn-McCurdy breach" (roughly speaking, a cost overrun exceeding a 15 percent increase in projected unit cost).

A recent assessment[2] of root causes of cost overruns by the director of PARCA, covering twelve programs, found the two most prominent causes were (1) unrealistic initial cost or schedule estimates (five programs), and (2) poor program execution performance (six programs). A less prevalent cause was changes in procurement quantities (three programs).

These causes are consistent with the findings of a 2009 IDA analysis—*The Major Causes of Cost Growth in Defense Acquisition*—of the causes of cost growth in eleven selected DoD acquisition programs, from the late 1990s through 2008. Figure 1 displays the cost growth documented by IDA researchers, and Table 1 summarizes the causes identified for the cost growth. The causes are highly interrelated—in fact, having the characteristics of a "Russian doll," because one must drill down even deeper to determine why DoD did not get the requirements right; why the programs proceeded into Engineering and Manufacturing Development before the technologies were proven, even though

---

[1] It may be small comfort, but cost overruns in DoD acquisition programs are, on average, no worse than other large-scale development and acquisition programs in non-defense areas.

[2] Available at http://www.acq.osd.mil/asda/docs/briefings/PARCA_General_Briefing.pdf.

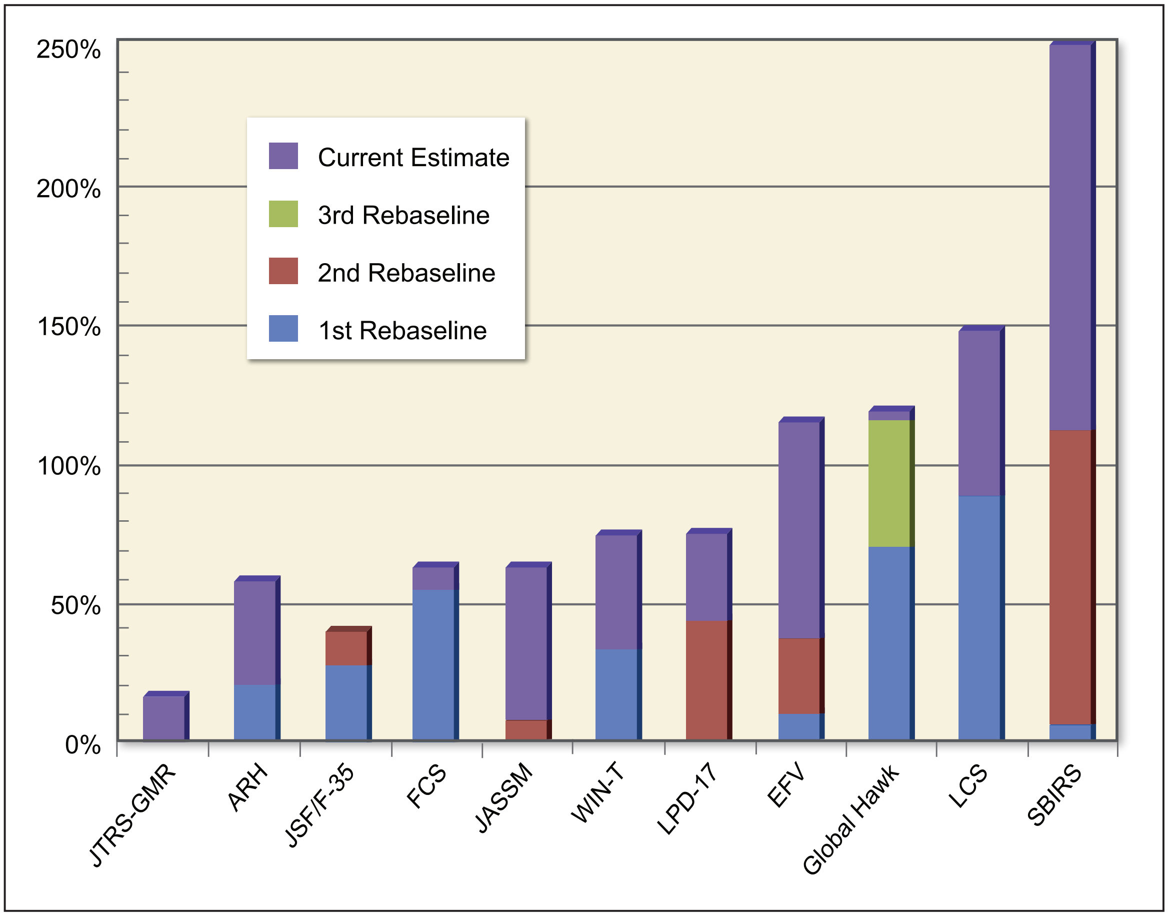

Figure 1. Growth in Program Acquisition Unit Cost (through 2009) in Selected Acquisition Programs

DoD policy explicitly prohibited it; why there was a lack of adequate early systems engineering; and why program schedules were accelerated or compressed unrealistically.

A "low ball" initial cost estimate is a sure way to cause a subsequent cost overrun. Thus a program can be a great success, but be subjected to much criticism simply because the initial cost estimate was faulty.

These problems occur for complicated reasons. First, as discussed earlier in this article, are the problems in the requirements definition process. Second, the problem of proceeding into full-scale development with immature technologies is, quite simply, management failure—a willingness to accept high risks that technologies would mature in time because of a perceived urgency in getting a capability fielded. (It is encouraging that the PARCA analysis cited above did not ascribe this cause to *any* of the twelve programs assessed.) Another article in this publication details technical issues with the Joint Tactical Radio System (JTRS). In June 2002, DoD approved JTRS for entry into full-scale development even though none of its critical technologies was at the required

Table 1. Areas of Weakness Causing Cost Growth in Programs Investigated by the 2009 IDA Analysis

	1. Top Management Activities			2. Initial Program Definition and Costing			
	a.	b.	c.	a.	b.	c.	d.
	Implementation of Policies	Implementation of Acquisition Reforms	Contractor Selection, Oversight, and Incentivization	Requirements Processes	Immature Technologies	Systems Engineering	Schedule Compression and Concurrency
ARH			●	●	●	●	●
EFV		●	●	●		●	●
FCS	●		●	●	●	●	
Global Hawk	●	●	●	●		●	●
JASSM		●	●	●		●	●
JSF				●		●	●
JTRS	●			●	●	●	
LCS		●	●	●		●	●
LPD-17		●	●		●	●	●
SBIRS	●	●	●	●	●	●	●
WIN-T	●				●	●	

readiness level. For the Future Combat System, only *three* of 44 critical technologies were mature; yet it also was allowed to move into full-scale development.

There is yet another, more subtle, reason for initially underestimating the cost of a new start. The program will more likely be approved if the cost estimate appears "reasonable and affordable." Thus, there is a natural tendency for proponents of new programs to err on the low side.

The 2009 IDA analysis noted that deficiencies in early systems engineering were due primarily to reductions in qualified systems engineers in government program management offices, requiring an excessive dependence on contractor-provided systems engineering. Shortcomings in early systems engineering have a direct impact on cost risks because the system being costed is not properly defined. The most salient case in point among the eleven programs considered in the 2009 IDA report is the Future Combat System, which was so poorly defined as to make cost estimation almost an exercise in speculation. However, the PARCA briefing, based on more recent analyses, cited only one of twelve systems as suffering from "unanticipated design, engineering, manufacturing or technology issues" (i.e., systems-engineering-related issues), indicating, perhaps, that greater attention is now being paid to more realistic front-end systems definition.

Affordability Risk

Affordability risk is the prospect that an acquisition program will become "unaffordable" at some time after its initiation. What does "unaffordable" mean, considering that, given a DoD budget in the vicinity of $500 billion, virtually any individual program should be affordable? It means that DoD resource allocators are not, or are no longer, willing to devote the resources needed to execute the rest of the program. That could happen because of program cost growth or because of changing priorities (driven by changes in threat or strategy), or—what is usually the case—a combination of both. In other words, a program becomes unaffordable if it no longer seems worth what it is projected to cost.

Affordability risk is also inherent at program initiation if no realistic affordability assessment is made, or if it is made and not acted on. For example, it was evident (based on briefing materials presented at the FCS Milestone B review to approve entry into full-scale development) that the Army would not be able to afford to complete the program as planned with the funds for investment projected to be available for the Army. Apparently, the decision makers either ignored that fact or believed that additional resources could be made available for Army investment. A problem in making affordability assessments for such programs is that usually most funding needs will occur beyond the five- to six-year fiscal planning horizon that DoD uses—i.e., the Future Years Defense Program (FYDP). Because no approved fiscal projections exist for the time period when the demand for funds will be the greatest, decision makers are not forced to confront the problem.

Affordability is usually a contributing, but seldom the only, factor leading to program cancellation. Affordability was cited as a concern for many of the acquisition programs that have been canceled over the past several years. These cancellations have resulted in significant inefficiencies—the sunk money on FCS alone at its cancellation has been estimated at $19 billion or more. (Some of that money was spent developing technologies that will find application elsewhere, so arguably it was not all wasted.)

Schedule Risk

As noted, cost and schedule risks go hand-in-glove. For the programs examined in the 2009 IDA report, schedule growth, measured by the estimates of the time required for full-scale development, averaged 80 percent for the eleven programs, while unit cost growth averaged 94 percent. Attempts to accelerate programs have frequently backfired, resulting in longer, rather than shorter, execution times. A prime example is Global Hawk, for which a perceived urgency to field the system rapidly for operations in Afghanistan led to concurrency in testing and production, and resulted in the need for expensive rework of systems post-production, which ultimately delayed fielding. Schedule risk also correlates strongly with the use of immature technologies. When the technologies fail to mature as anticipated, programs must either slow down development to await technology maturation, or seek alternative technologies—either way, the schedule will likely slip.

Performance Risk

Performance risk is the chance that key performance characteristics required of the system will not be obtained. Again, this risk is tightly intertwined with both cost and schedule risk. When performance shortfalls become apparent, either remedial steps will be needed or the user will have to accept lesser capabilities (or both). Frequently the cause is ambitious initial requirements. An excellent example is the Marine Corps Expeditionary Fighting Vehicle (EFV), with a requirement to skim through modest waves at speeds up to 25 knots, and operate ashore as an armored fighting vehicle carrying a reinforced squad of 17 marines over land at speeds up to 45 miles per hour. These requirements ultimately proved to be unattainable in a vehicle that was sufficiently reliable and affordable. After two Nunn-McCurdy breaches, numerous test failures, and more than $7 billion spent, the program was canceled by Secretary Gates in 2011.

Similar stories can be told for the Future Combat System, Global Hawk, Joint Strike Fighter, Littoral Combat Ship, and Space-Based Infrared System, among those examined in the 2009 IDA analysis. However, of the twelve programs covered in the PARCA briefing cited above, only one was scored as having unrealistic performance expectations contributing to cost growth.

Conclusion

The 2009 IDA report concluded that, for the programs examined, cost growth could have been greatly reduced or eliminated if policies and procedures in place had been more rigorously followed.

The establishment of the PARCA office institutionalizes the type of analysis performed by IDA in 2009. In fact, since 2009, we have seen a significant decrease in acquisition program cost growth—annualized growth in the estimated Average Procurement Unit Costs for all MDAPs averaged 5.7 percent per year through December 2009, while such growth has averaged 3.3 percent per year since then. And this is before the full impact of cost control efforts such as PARCA and the Deparment's Better Buying Power initiatives can be felt. Because of the nature of the challenge (i.e., complex systems that must operate in stressful environments), there will always be risks of cost growth in many DoD acquisition programs. Nonetheless, there is cause for optimism that DoD is doing a better job today of addressing and managing cost growth and the associated risk of achieving needed defense capabilities within available resources.

Dr. Kneece is an Adjunct Staff Member in IDA's Strategy, Forces and Resources Division. He holds a doctorate in mathematics from the University of Maryland.

Mr. Porter, an Adjunct Staff Member in IDA's Strategy, Forces and Resources Division was Director, Acquisition Program Integration in the Office of the Under Secretary for Acquisition, Technology and Logistics and, earlier, Principal Deputy Director, Program Analysis and Evaluation.

Sources:

Porter, Gene H., Brian G, Gladstone, C. Vance Gordon, Nicholas S. J. Karvonides, R. Royce Kneece, Jr., Jay Mandelbaum, and William H. O'Neil. *"The Major Causes of Cost Growth in Defense Acquisition."* IDA Paper P-4531. Alexandria, VA: Institute for Defense Analyses, December 2009.

Porter, Gene H., C. Vance Gordon, and R. Royce Kneece, Jr. *Improving the 'Front-End' of the DoD Acquisition Process.* IDA Paper P-4710, Volume 1. Alexandria, VA: Institute for Defense Analyses, June 2011.

Selected Past Issues

Best Publications in the Open Literature

- The Saddam Tapes: The Inner Workings of a Tyrant's Regime 1978-2001
- A New Methodology for Estimating Nerve Agent Casualties
- Rotorcraft Safety and Survivability
- Secure Cloud Based Computing
- Orbital Maneuver Optimization Using Time-Explicit Power Series
- Comparison of Predicted and Measured Multipath Impulse Responses

Security in Africa

- Trends in Africa Provide Reasons for Optimism
- China's Soft Power Strategy in Africa
- Sudan on a Precipice
- A New Threat: Radicalized Somali-American Youth
- Chinese Arms Sales to Africa
- Outsourcing Imagination: The Potential of Informal Engagement Networks in Africa
- Defense Environmental Cooperation with South Africa

Challenges in Cyberspace

- Cyberspace – The Fifth and Dominant Operational Domain
- Transitioning to Secure Web-Based Standards and Protocols
- Information Assurance Assessments for Fielded Systems During Combatant Command Exercises
- Supplier-Supply Chain Risk Management
- Internet-Derived Targeting: Trends and Technology Forecasting
- Training and Educating the DoD Cybersecurity Workforce

Today's Security Challenges, Part II

- A Framework for Irregular Warfare Capabilities
- Bridging the Interagency Gap for Stability Operations
- Developing an Adaptability Training Strategy
- Force Sizing for Stability Operations
- Planning Forces for Steady State Foreign Internal Defense and Counterinsurgency
- Test and Evaluation for Rapid-Fielding Programs
- Understanding Security Threats in East Africa

Today's Security Challenges, Part I

- Supporting Warfighting Commands
- Detecting Improvised Explosive Devices
- Building Partner Capacity
- Combating the Trans-South Atlantic Drug Trade
- Countering Transnational Criminal Insurgents
- Using Economic and Financial Leverage
- Understanding the Conflict in Sudan

Resource Analyses

- Evaluating the Costs and Benefits of Competition for Joint Strike Fighter Engines
- Analysis and Forecasts of TRICARE Costs
- Cost Savings from the Post-Cold War Consolidation of the Defense Industrial Base: A Case Study of the Shipyards
- The Effects of Reserve Component Mobilization on Employers
- Does DoD Profit Policy Sufficiently Motivate Defense Contractors?
- Auctions in Military Compensation

Focusing on the Asia-Pacific Region

- Making American Security Partners Better Resource Managers
- Collaborating with Singapore to Better Understand the U.S. and Asian Defense Environments
- Intellectual Outreach to the Muslim World: The Council for Asian Terrorism Research
- Inside North Korea
- Red Teaming for Terminal Fury
- Promoting Interagency Cooperation in Shaping U.S.-China Relations
- Extending Trilateral Cooperation in Dealing with Disaster
- Developing Human Capital in China – Implications for the United States
- Nanotechnology in the Pacific

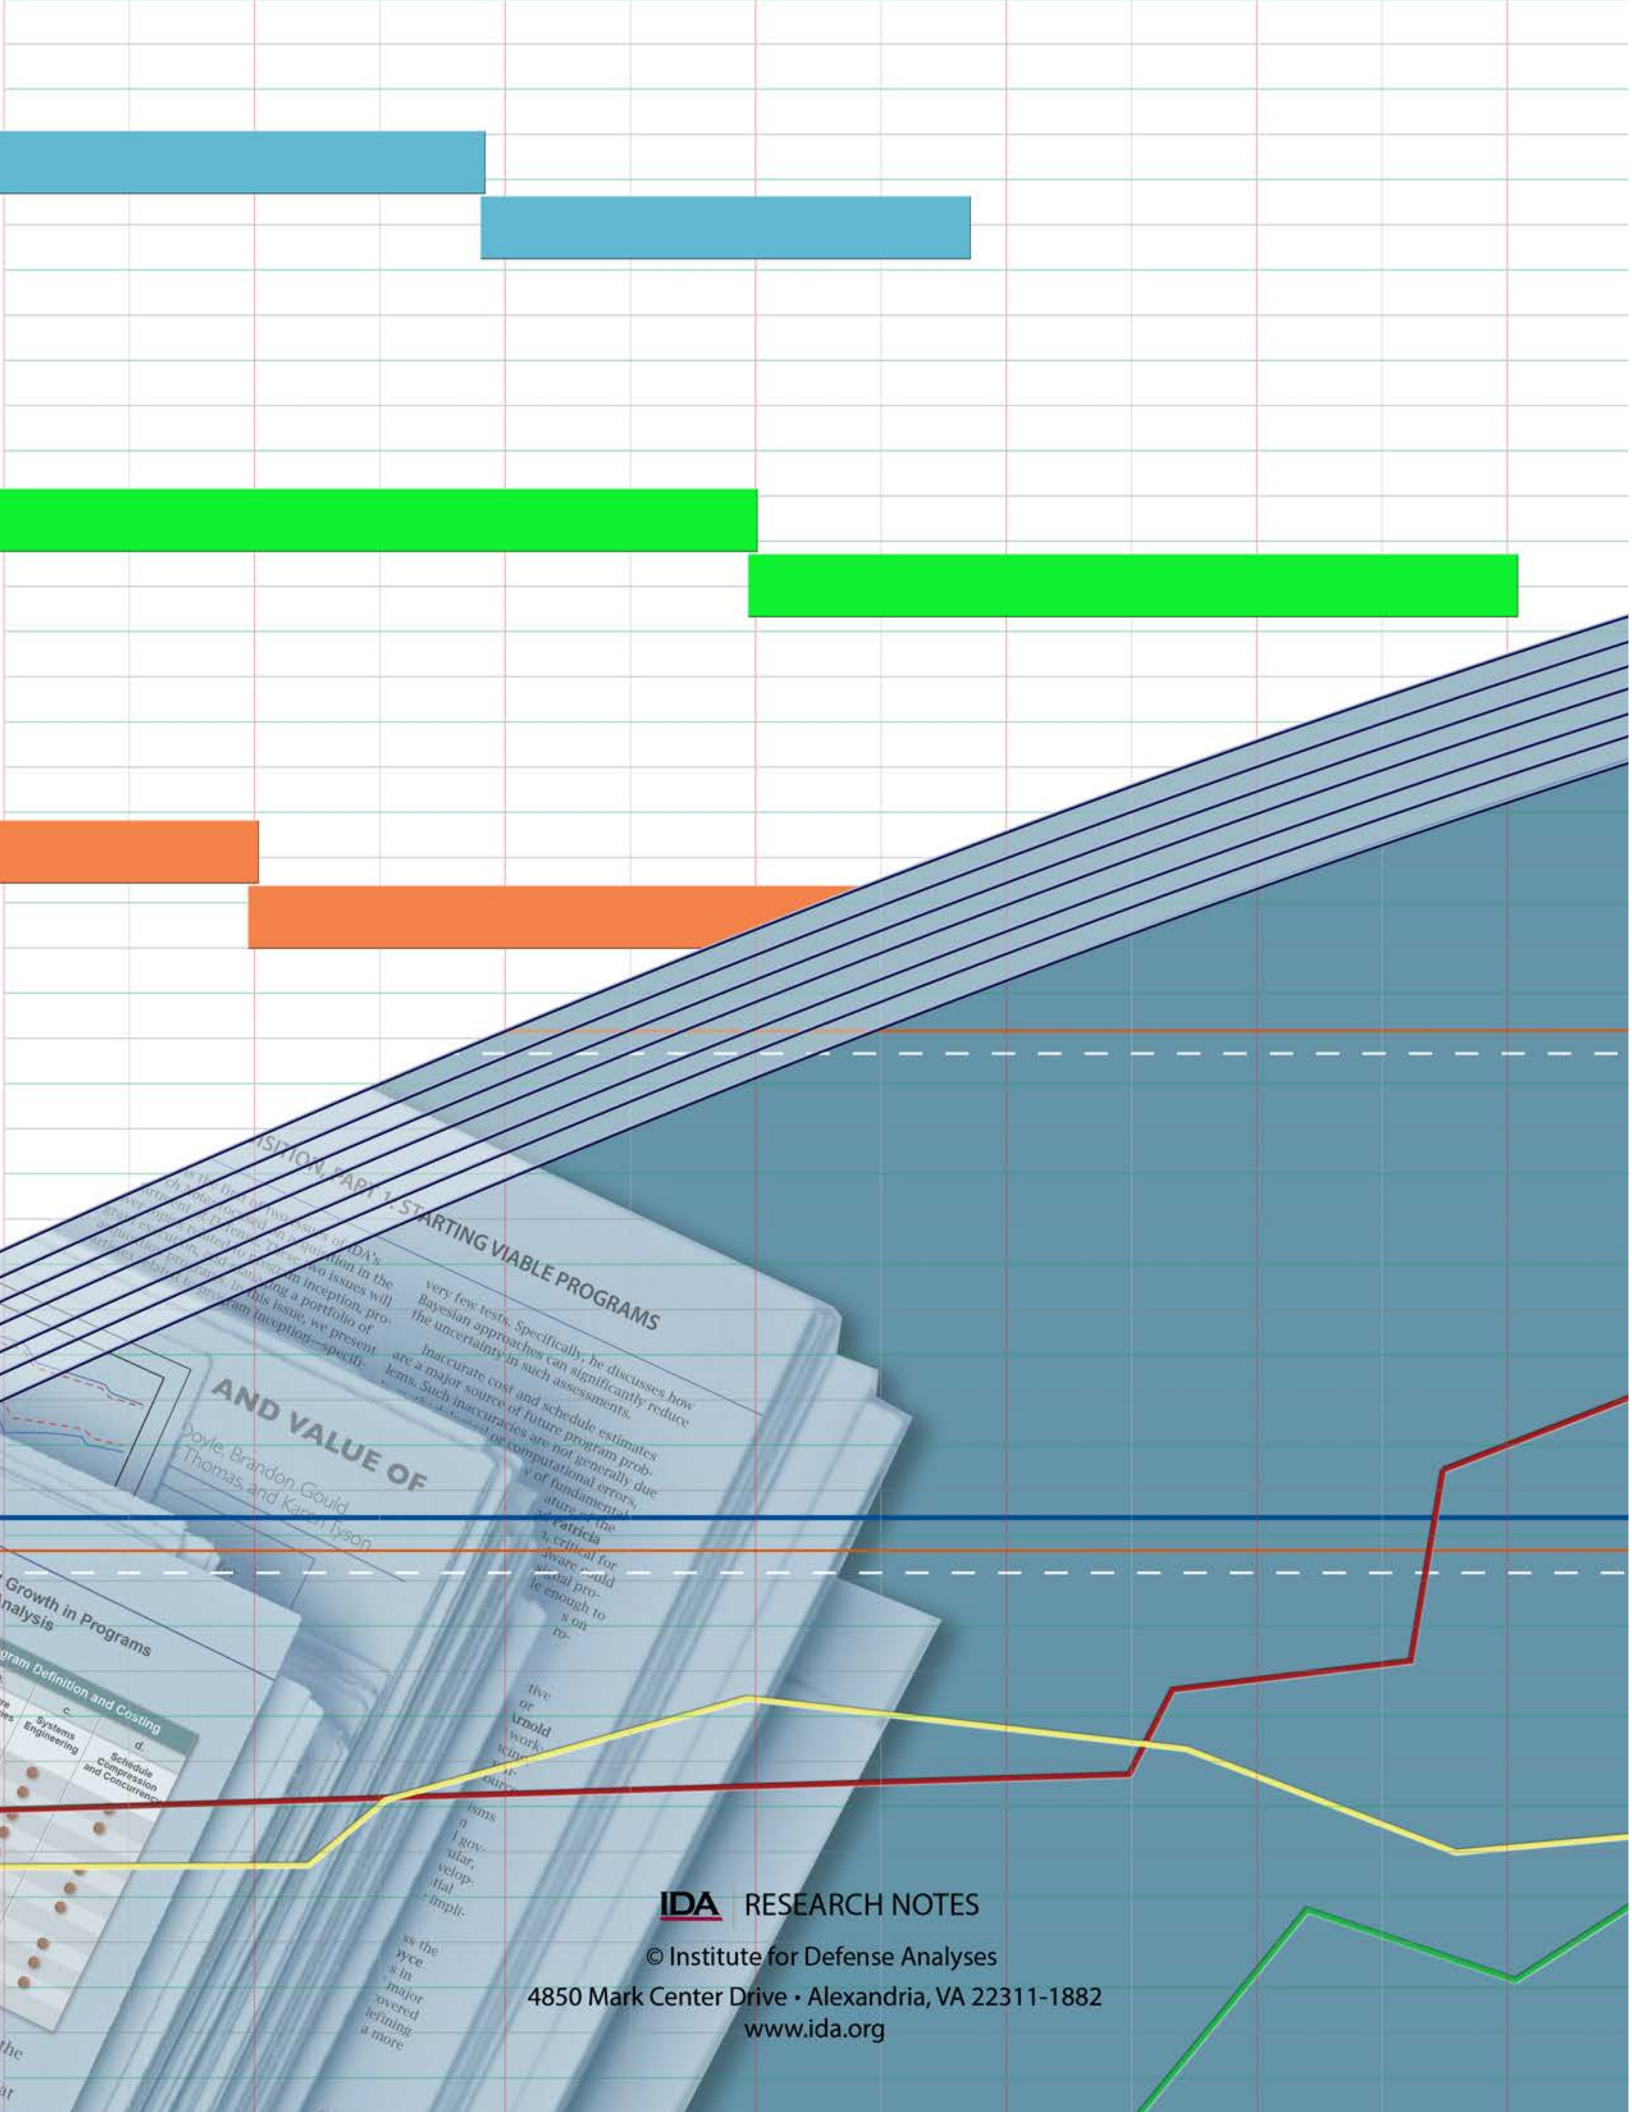

IDA | RESEARCH NOTES

4850 Mark Center Drive • Alexandria, VA 22311-1882
www.ida.org