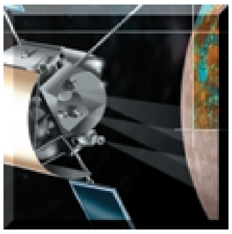

Quality and Reliability: APL's Key to Mission Success

Ward L. Ebert and Eric J. Hoffman

APL's Space Department has a long history of producing reliable spaceflight systems. Key to the success of these systems are the methods employed to ensure robust, failure-tolerant designs, to ensure a final product that faithfully embodies those designs, and to physically verify that launch and environmental stresses will indeed be well tolerated. These basic methods and practices have served well throughout the constantly changing scientific, engineering, and budgetary challenges of the first 40 years of the space age. The Space Department has demonstrated particular competence in addressing the country's need for small exploratory missions that extend the envelope of technology. (Keywords: Quality assurance, Reliability, Spacecraft, Spaceflight.)

INTRODUCTION

The APL Space Department attempts to be at the high end of the quality and reliability scale for unmanned space missions. A track record of nearly 60 spacecraft and well over 100 instruments bears this out as a success. Essentially, all of APL's space missions are one-of-a-kind, so the challenge is to get it right the first time, every time, by developing the necessary technology and addressing the risks inherent in innovation. Today, the nation is engaged in a debate about the validity of claims of success in meeting the challenge referred to as "better, faster, cheaper," coined by NASA Administrator Daniel Goldin. It is our claim that we have met that challenge, not by inventing a response to it, but by applying our traditional approach to space mission development.

Conventional wisdom says that in research and development it is not possible to predict or control all three critical parameters: performance, schedule, and cost. This claim is based on the premise that research and development is fundamentally an innovative endeavor, and any plan to carry out such work necessarily has steps with unknown outcomes. On the other hand, conventional management wisdom says that some degree of risk and unpredictability exists in any major undertaking, and the only element that differs in research and development is the larger degree of risks. More contingency in cost and schedule may be required, but fundamentally the trade-off decisions, planning, and control of a project are the same.

"Better, faster, cheaper" does not mean a trade-off among the three, but rather innovation and good technical judgment so that resources are applied to those factors that eliminate the most risk in each. The principles that serve well include the basics of getting a good, solid design early, as discussed in the previous article by Hoffman and Fountain (this issue), and controlling the quality and testing at each level of integration to verify design margins. Experience is a guide to efficiency and effectiveness in all these steps. Although the Near Earth Asteroid Rendezvous (NEAR) stands as an example of the new class of interplanetary scientific mission—developed in 3 years under cost caps of a few hundred million of today's dollars—it was developed the way we have always tried to develop space missions. It is obviously "faster" and "cheaper" than the billion-dollar interplanetary missions of past decades, but it is better technologically as well, employing a full suite of state-of-the-art instruments and fully redundant spacecraft subsystems.

UNDERSTANDING THE RISKS

In the 1960s, early spacecraft failures pointed to the need for more reliable launch services, better materials and assembly processes for electronics, and a better understanding of the space environment. Today's failures point to the need to improve the reliability of complex software and contend with longer and less controllable supplier chains for parts and materials. Experience (i.e., learning from mistakes and near misses) is essential both to success and to efficiency; however, the wisdom of the past is usually not enough, particularly for the types of challenges that define APL's mission. Keeping ahead of changes in technology means continuously undertaking new risks.

Table 1 highlights some major considerations of introducing changes and new designs in space systems, compared with replicating "heritage" designs. Heritage design refers to the reuse of qualified designs in a sufficiently similar environment that requalification is not necessary. The term heritage has been so greatly overused in an attempt to market products that prudence now demands careful scrutiny of any claims of heritage. One must always ask what differences of any sort there are in the interface specifications; the physical parameters of testing, storage, or operation; the availability of parts and materials (and in some cases people); and the length of time in testing, storage, or operation.

Some of the worst failures in the space business have been due to misapplied heritage arguments. For example, the failure of a $700M Ariane V rocket was attributed to the attempt to apply Ariane IV attitude control software to a greatly changed rocket. We at APL have had our share of less spectacular reminders when environments or suppliers change. Our worst experience came in the late 1970s with the launch of two TIP (Transit Improvement Program) spacecraft a year apart. Both failed to successfully complete the postlaunch deployment sequence because a flexible metal antenna had gotten very hot and melted a nylon keeper ring designed to prevent it from flapping around in the vehicle fairing. The nylon then cooled and bonded to the antenna before the attempted deployment. Because of the additional weight of these spacecraft, we had elected to eject the fairing at a lower altitude than our previous similar launches. What we had not accounted for was the aerodynamic heating of exposed surfaces immediately after deployment of the fairing. Once discovered, the problem was easily solved by changing the keeper material to Teflon.

Two cases come to mind in which procured components suddenly could not pass acceptance tests because of staffing changes at the suppliers. In both cases, the components had been built and tested originally to a specification that went unchanged for many years. In the case of a simple bellows actuator, the supplier's corporate knowledge had been lost through staff turnover. When units began failing in test, it became

Table 1. Risk/benefit trade-offs between heritage and new technology.

Considerations	Heritage design	New technology
Benefits	Usually costs less if application has not changed Failures are easier to understand Can usually meet shorter schedule	Compliant with interface constraints Adaptable to commercial off-the-shelf components Incorporates improvements in performance
Risks	Failure to bring along 100% of the heritage design Design insufficiently well documented to reproduce Loss of compatibility with new systems, environment Loss of quality with incremental adjustment Loss of necessary resources Inaccessibility of parts and materials	Cost/schedule initially less predictable Failure to pass qualification Failure modes escape qualification process

apparent that the specification did not include all that was necessary to make the design work. For whatever reason, the walls of the metal bellows had been made thicker, still within specification, but now too stiff to expand when fired. In the second case, an electric motor built for RCA/Astro-Electronics Division on the Transit Program[1] by one supplier could not be reproduced by another when the former closed its business. The original motors all met specification, and the original drawings accurately represented the original motor, but apparently to an insufficient level of detail. Their use in the reproduction effort, then, yielded a motor with significantly lower torque. RCA/Astro solved the problem by giving the flight unit to one of their own wizards who was able to "tweak" it into the specified range, which is probably what the original supplier had been doing all along.

APL experience with 58 spacecraft, 120 instruments, and hundreds of delivered subsystems has revealed three key elements to achieving high reliability. First, and most important, is to get the design right. Retrospective NASA and Air Force studies of the causes of failure in spaceborne equipment have consistently shown *design error* as the single biggest contributor, accounting for 42% of the identified failures in one NASA study.[2] The Space Department's approach to reliability, therefore, begins with a strong emphasis on "design integrity," including a tradition of rigorous design reviews. The second important element is careful control and screening of all parts, processes, and workmanship used to produce the end item—the traditional domain of the reliability engineer. The last link in the chain is a thorough inspection and test program, and APL's mandatory test program[3] predates even the venerable MIL-STD-1540 Air Force test standard.

These lessons have been learned from our own, sometimes painful, experiences, as well as by comparing experiences with others in the field (benchmarking). The rules and standards are continually being reexamined in light of changes in technology, the "better, faster, cheaper" mandate, and the ascendency of mass-produced constellations of spacecraft. The reliability, configuration management, and test requirements for small, low-cost satellites are still a matter of vigorous debate. "Better, faster, cheaper" spacecraft, for example, can rarely afford total avoidance of single point failures. Redundancy must be applied with particularly good judgment, and sometimes functional redundancy or even single-string designs will be tolerated. But throughout this debate, it is important to remember that the *true* cost of a "cheapsat" is the stated cost *divided by the probability of success*. A satellite delivered at half price is no bargain if the probability of achieving its mission has also been halved. And the stated cost must include the cost of the launch vehicle and mission operations, not just the satellite itself. There is one last, intangible cost of failure: the public's loss of confidence, interest, and support for the space program. With all of these thoughts in mind, APL has always focused on mission performance assurance.

GETTING THE DESIGN RIGHT

Design is the most important contributor to high reliability. The principles that help ensure a reliable design (see the article by Hoffman and Fountain, this issue, for a fuller discussion) are summarized here:

- Keep the mission objective uppermost in mind at all times.
- Adopt a system engineering mentality.
- Stay organized amidst the complexity by writing things down.
- Adopt a "worst-case" mentality toward design parameters.
- Keep the design as simple as possible.
- Subject the design to a rigorous independent review.

Reliability apportionment and redundancy are key decisions that must be made early in the design of any system or subsystem. Although "reliability analysis" is often dismissed as an exercise in fiction, in fact the *relative* numerical answers from MIL-STD-217–type reliability analysis can help in comparing different system-level choices. It is particularly useful for comparing redundant versus nonredundant implementations and in conducting sensitivity analyses to show where additional reliability is needed.

Misunderstanding the environment is a common cause of design error, and, as we have seen, of misapplied heritage designs. The Laboratory is particularly well-positioned to assure a good understanding of the environment because Space Department engineers have always worked side by side with the scientists who study the space environment. Space Department scientists and engineers have written book chapters and taught short courses on the space environment. Our Reliability and Quality Assurance (R&QA) Group is particularly strong in radiation expertise, one of the more arcane and difficult aspects of operating in space. We run one of the nation's leading radiation test and evaluation programs, with the radiation data network-accessible to our engineers.

Again, conservative and worst-case design approaches pay off. Although we sometimes use statistical approaches to tolerancing, it is often simpler and not much more costly to simply design for the "worst case," so that the design will function even if every parameter simultaneously reaches its worst limit. A conservative approach to estimating environments is coupled with strict derating rules in determining how a particular part will perform in the environment. This approach is especially needed for the radiation environment, where a factor of 2 uncertainty is not unusual. Space Department derating rules are set forth in our Engineering Notebook and are one of the items examined at circuit-level design reviews.

CONTROLLING THE PARTS, MATERIALS, AND PROCESSES

Parts

The control of parts, materials, and processes is the traditional domain of the components engineer and what is thought of first when R&QA is mentioned. But the evolution of technology has brought important changes to this area. For example, in the early years of the space program, EEE (electrical, electronic, and electromechanical) piece part reliability was the principal lifetime issue and the focus of the most attention in reliability groups throughout the industry, including ours. Semiconductor parts (transistors, integrated circuits) were particularly troublesome in those days. More recently, the reliability of *all* EEE parts has improved to the point that other issues now present higher risk and consume more of our attention. There are several reasons for this general improvement in EEE part reliability, foremost being the demands of the automotive and consumer electronics industries, coupled with a better understanding of the physics of failure. The space industry is fortunate in being able to take advantage of this general improvement, as dedicated high-reliability parts manufacturers continue to exit the market (Fig. 1). The downside of this "riding the coattails" of commercial part development is the greater difficulty today in finding radiation-hard parts suitable for the space environment.

Our R&QA Group has always provided strong component engineering advice early in the design process and has pioneered innovative systems for purchasing, stocking, and managing flight parts. A key idea for maintaining parts discipline is the use of a preferred parts list (PPL); we use our own as well as NASA's PPL. We are also active members of the Space Parts Working Group and GIDEP (Government/Industry Data Exchange Program), nationwide forums for sharing parts information.

The Laboratory instituted and operates a flight parts storage facility to maintain residual parts for future use and to stock long-lead and widely used flight-qualified electronic parts. These parts can be provided on short notice, assuring that "faster, cheaper" programs need not compromise parts quality. Many of the parts are specialized, radiation-hard flight components no longer available from outside. APL's flight parts staging is managed with MFG/PRO,[4] a state-of-the-art, on-line Manufacturing Resources Planning II (MRP II) system that is used by many major international corporations (e.g., Hewlett-Packard, Johnson Controls). It provides parts status information to our engineers and project managers at their desktops. If parts substitutions are needed, engineers can verify their designs quickly with parts on hand, saving money and time and further reducing risk. The NEAR Program, among others, benefited from this innovative process; 63% of NEAR's parts were on hand by the preliminary design review. By avoiding parts delivery delays, we preserve schedule time for careful fabrication, inspection, and test, thereby reducing risk throughout the program.

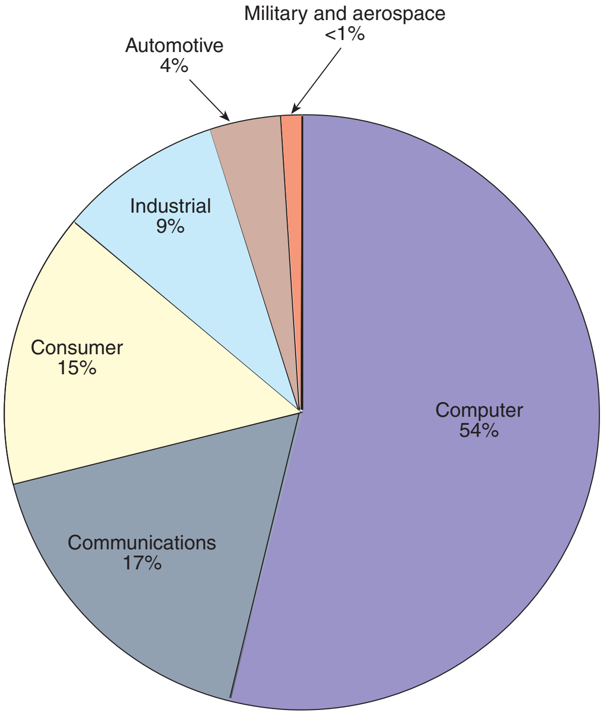

Figure 1. In 1965, two-thirds of the U.S. microelectronics industry was driven by military and aerospace markets. Today that share has declined to <1%.

Figure 2 is a view of the staging and kitting area for parts procured to support the TIMED (Thermosphere-Ionosphere-Mesosphere Energetics and Dynamics) Program. All parts have passed all inspection and screening requirements. This storage space is secured and environmentally controlled. Advanced purchases of long-lead parts to support future programs are made under the guidance of an internal Space Parts Advisory Panel; the parts are then transferred to programs on a chargeback basis. Programs save additional money by buying from the stockroom only the quantity needed (not minimum lots), with instant backup if more parts are required.

Shown in Fig. 3 is our R&QA Group's state-of-the-art, $1.5M Sentry SX-100 digital integrated circuit (IC) tester (the first on the East Coast), which permits in-house testing of the most advanced ICs. It provides fast turnaround, reduces schedule risk by ensuring proper function prior to board assembly, and facilitates failure analysis. The proximity (10 m) of the tester to our cobalt-60 irradiator ensures accurate parametric characterization following radiation dose exposure.

Plastic-encapsulated microcircuits (PEMs) were once widely forbidden in flight hardware in favor of hermetically sealed parts. This was the rule in the APL Space Department, too, until the NEAR Program forced us to find reliable ways to fly PEMs. PEMs, in fact, were a "mission-enabling technology" for NEAR—the mission could not have been performed had we not been able to find a reliable way to fly PEMs in NEAR's solid-state recorder. Through this activity, the APL R&QA Group established itself as a leader in the reliable application of PEMs in space, and remains so to this day.[5–7]

Materials

An unfortunate incident on the Active Magnetospheric Particle Tracer Explorer (AMPTE) Program caused us to tighten up our materials control process. During vibration acceptance testing of the AMPTE spacecraft, a support bracket failed. Analysis showed that it had been fabricated inhouse from the wrong (mislabeled) material. Fortunately, no other hardware was damaged. Following that incident, our materials stockroom was purged of all materials that did not bear a Certificate of Compliance guaranteeing correct identification; all materials since have been purchased with a certificate. This incident also illustrates the value of vibration testing (and the importance of keeping your eyes and ears open during the test).

Similar rules apply to organic materials, materials with limited shelf life, and critical fasteners. Critical fasteners are defined and identified early in a design phase, and a special process controls their purchase, acceptance, and use. Of course, our designers strive to eliminate from the design as many critical fasteners as possible.

Figure 2. The staging area for flight-ready electronic parts currently holds over 1 million piece parts with a purchase cost of over $12M. The inventory is managed using a commercial Manufacturing Resources Planning (MRP II) system, MFG/PRO.

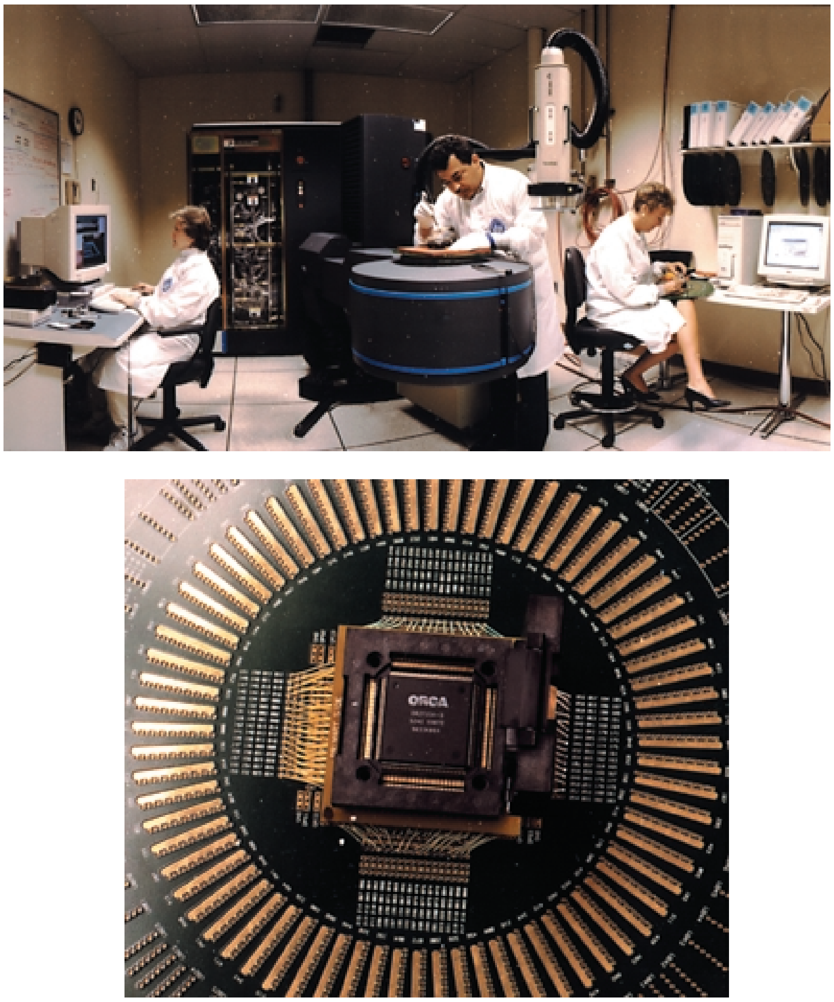

Figure 3. The state-of-the-art VLSI digital integrated circuit tester (top), a Sentry SX-100, can test parts with up to 448 input/output pins at clock speeds of 100 MHz and lower pin-count devices up to 200 MHz. Shown on the test head (bottom) is a 240-pin field-programmable gate array.

Processes

Most APL fabrication processes are controlled by our sister department, the Technical Services Department (TSD). Several volumes of controlled processes are in the midst of being transferred from notebooks to our Intranet. These processes are referenced by designers, provide direction for the fabricators, and are the basis for inspections. Repeatability is guaranteed, and with everyone "singing from the same sheet," little room for argument exists between the parties.

There are processes other than hardware, of course. Guidelines for our software development process are given in the Engineering Notebook; these are currently being revised to capture lessons learned from recent programs. Even the development and preparation of proposals for new programs is a "process," and following it is one way a small organization like the Space Department is able to compete in substantial Announcements of Opportunity.

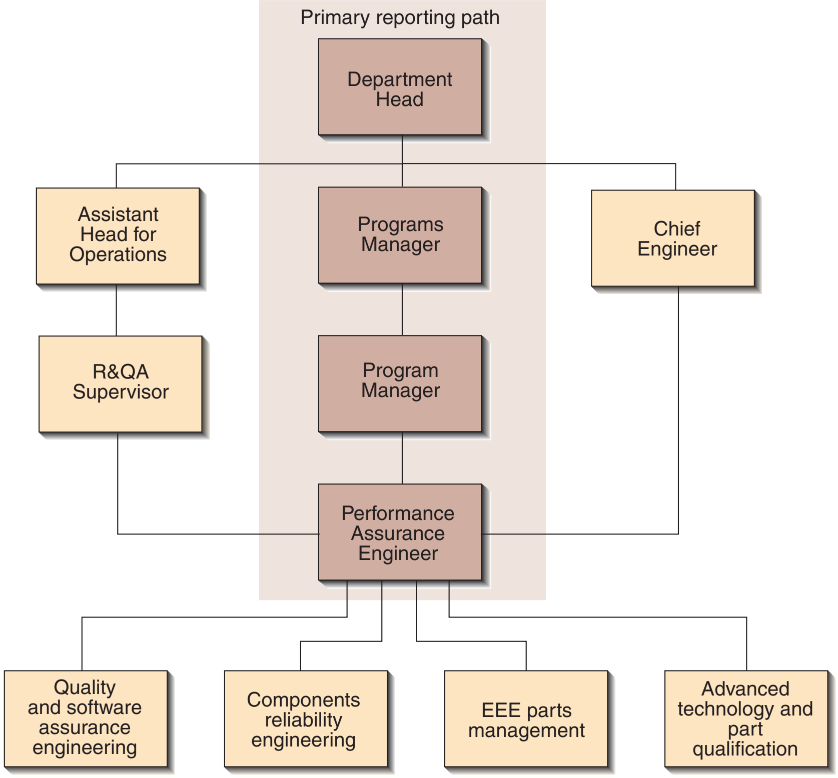
Primary reporting path
Department Head
Assistant Head for Operations
Programs Manager
Chief Engineer
R&QA Supervisor
Program Manager
Performance Assurance Engineer
Quality and software assurance engineering
Components reliability engineering
EEE parts management
Advanced technology and part qualification

Figure 4. The primary and secondary reporting paths through which quality assurance, safety, risk assessment, and other information flows from the R&QA staff to Department management. This configuration recognizes that the primary responsibility for quality and quality reporting lies with the Program Manager. The other paths are always open and are used for regular formal written documents and occasionally to alert Department management to urgent concerns.

ORGANIZING FOR PERFORMANCE ASSURANCE AND RELIABILITY

A distinguishing characteristic of APL's approach to program management is the high degree of responsibility held by the Program Manager for all aspects of performance. The article by Chiu and Dassoulas (this issue) details this philosophy. The performance assurance arena is no exception; the Performance Assurance Engineer (PAE) is assigned to a program just like other lead engineers, and is expected to bring a strong inclination toward ownership of the final product and value-added assurance efforts. Figure 4 depicts the primary and secondary reporting paths in this system. The PAE (and all R&QA support) is on direct charge to the program, with primary responsibility to support the Program Manager. This arrangement fosters teamwork, problem solving, and efficiency and has served us well through the years.

ISO 9000 Compliance

APL has undertaken the task of modifying its formal quality systems to meet the challenge put forth by Dan Goldin to NASA contractors to become compliant with ISO 9000[7] standards. Announcements of Opportunity in the Discovery, Earth Science Systems Pathfinder, and Midex Programs have stipulated that offerors have quality systems "consistent with" ISO 9000, and to truly be consistent with this standard we must do some things a bit differently. The first change was made in 1998 when the Director of the Laboratory issued a new version of the APL Quality Assurance Plan. APL's quality system serves a diverse array of sponsors and is necessarily flexible so that it can be tailored in proportion to the cost, scope, and risks of each task. Every task at APL is assigned a Quality Assurance Requirement Level that specifies, based on the task's criticality, the framework of quality assurance activities and documentation necessary.

Within this framework, the Space Department imposes standard practices not shared by other departments within the Laboratory, although aspects of Space Department requirements drive the performance of design, fabrication, and procurement activities carried out in support of our programs by other departments. Each major task, such as spacecraft or instrument development, requires a Product Assurance Implementation Plan (PAIP). This is often a single document tailored to meet sponsor requirements, but it will occasionally call for separate documents such as a Safety Plan, Software Quality Assurance Plan, Contamination Control Plan, or Configuration Management Plan. The PAIP is actually a part of the Program Plan, and Departmental Practices and Procedures define the required scope and approvals for the PAIP. When a significant portion of the work is subcontracted or procured, an additional document, called the Procured Product Assurance Requirements (ProcPAR), is required and becomes part of the subcontract requirements. It provides a means for consistent flow-down of the program quality assurance requirements to our subcontractors.

The general quality system structure described here has been in place and has evolved over several decades. The step to ISO 9000 compliance is evolutionary, but challenging in several aspects. The Space Department has completed a "gap analysis," which is the first step toward assuring full ISO 9000 compliance (specifically ISO 9001 for our design, development, test, integration, and operations work), and we are in the midst of closing the identified gaps. For example, the requirement to independently and regularly audit top-level procedures and document conformance to the PAIP and ProcPAR is not presently formalized in our system, and is accomplished only rarely. Lower-level reviews of designs, inspection, testing, rework, nonconformance, failure analysis and resolution, and so forth are well handled.

Process discipline is, of course, the key idea within ISO 9000, and it begins with documenting the good practices we have already described in this article. The difficulties in designing very top-level procedures are rooted in the nature of our business, in particular, the wide diversity of sponsor quality assurance requirements and the nonrepetitive nature of nearly all the work we do. These concerns are not unique to APL, nor do they represent insurmountable obstacles. One would imagine that the community of research and development centers and laboratories will converge to similar approaches to blend common sense with the intent of ISO 9000.

Inspection and Testing

A common misconception is that ISO 9000—with its emphasis on repeatable processes—allows you to "fire all the inspectors." In fact, a well-designed inspection and test program is vital to delivering high-reliability products.

Hardware inspection begins at the piece part level. Particularly today, with some ICs costing $25K, we cannot take the chance of populating boards with bad parts. Going back to around 1980—and based on some parts incidents on the Global Positioning System Package (GPSPAC) Program—the Department instituted a requirement for 100% inspection of parts going into our flight hardware. More recently, we have relaxed that requirement to accept test data from certain well-trusted suppliers in lieu of APL retest. Our parts testing generally meets NASA requirements and sometimes includes a burn-in period of 168 hours. In those occasional cases when adequate screening cannot be implemented at the part level, we can subject the assembled hardware to environmental stress screening to eliminate potential workmanship and "infant mortality" failures. We also require our assembled flight hardware to accumulate a minimum number of continuous failure-free hours of operation, typically ≥120 hours at the board and box level and 200 to 500 hours at the spacecraft system level.

As hardware is fabricated (typically by TSD), the operators self-inspect and an independent TSD inspector inspects before delivery. Space Department lead engineers are also trained in basic at-delivery inspection techniques, and certified inspectors from the Space Department's R&QA Group provide further inspection at key points. Off-site and subcontractor inspections are performed by the R&QA Group or by trusted subcontract inspectors. If a sponsor requires it, APL's on-site Defense Contract Management Command Office can provide government source inspection at critical points. All parties use APL's workmanship standards document[8] to reduce subjective interpretations of acceptability. At the start of each hardware development effort, an inspection program is outlined to assure effective—but not overkill—inspection.

The APL Space Department maintains a rigorous, mandatory program for testing flight hardware and software, with written standards for qualification and acceptance testing. As noted earlier, the Laboratory has traditionally taken a conservative approach to testing in order to minimize overall program risk. This includes a thorough, hierarchical program of board- and box-level tests and subsystem tests, including full environmental exposure, so that units are delivered flight-qualified for overall spacecraft integration and testing (I&T). Having our own in-house facilities for vibration and thermal vacuum (TV) testing permits a rigorous test program with low schedule risk. If acoustic testing or large TV facilities are required, we can use the facilities at NASA/Goddard Space Flight Center (GSFC), and have done so successfully on several programs. Our test margins are set by Department standards and are similar to those in MIL-STD-1540 (for example, 3 dB on spacecraft vibration, ±10° on predicted temperatures). All test equipment used with flight articles is fully calibrated to National Institute of Standards and Technology traceability, and we maintain a rigorous system of problem/failure reporting and closure.

In the "good old days," programs often had the luxury of qualifying a prototype or qualification unit before moving on to the flight unit. In today's budget-constrained era, a "protoflight" test philosophy is more common: the first unit built and tested is the one that flies. A protoflight test program may, for example, use qualification test levels for flight durations. All spacecraft instruments, subsystems, components, and software are delivered to the Laboratory fully qualified before spacecraft integration. The only exceptions are the solar arrays, harnesses, and thermal subsystems. Their qualification is completed at the spacecraft level with a protoflight-level sine vibration test of the fully assembled spacecraft, conducted at APL. Our policy requires all flight hardware and software to be onboard for spacecraft environmental testing—we express this as "test what you fly, fly what you test." Preenvironmental and preshipment reviews are held at key points to independently review the test history as well as the test plans.

It is important that failures and anomalies be explained and corrected and that trends from possibly widely separated failures be spotted. Our formal system of problem/failure reports (P/FRs) begins upon the first end-item acceptance test of any unit. Prior to that, lead engineers keep all test histories in logbooks. P/FRs are closed in a formal process that emphasizes corrective and preventive actions which are tracked for trends and reported on-line and at all test reviews. Before launch, the R&QA Group Supervisor and the Chief Engineer conduct a final, independent review of all P/FRs. After launch, a similar P/FR process is initiated to deal with any mission operations errors, failures, and anomalies.

Software Validation

Software is not only an ever-larger cost consideration on each program, it is increasingly where much of the reliability risk resides. Since flying the first reprogrammable computer in space, APL has long recognized the importance of correct flight and ground software. Documented software quality assurance guidelines form an important part of our design integrity program. In meeting these requirements for a particular program, a tailored Software Assurance Plan is prepared. All software is developed using a repeatable process that includes formal documentation for software requirements, software design, software test and validation, and configuration management. Our software development process is based on the recommendations developed by the NASA/GSFC-sponsored Software Engineering Laboratory.[9]

Software figures prominently in our conceptual, preliminary, and critical design reviews, in addition to the more detailed software design reviews and structured walkthroughs. Independent validation and verification (IV&V) is used for especially critical code; for example, we implemented IV&V on the Midcourse Space Experiment (MSX) (through a subcontract), NEAR (through another APL Department), and the Advanced Composition Explorer (ACE) (through NASA/Goddard). A Space Department software process engineering team provides education on current processes and solicits inputs for continuous improvement. Our R&QA Group also has several American Society for Quality (ASQ)–certified software quality engineers to provide independent quality assurance.

Radiation Effects Assessment

By 1962, APL had launched eight spacecraft; two failed to reach orbit, one reentered the atmosphere a month after launch, and three ceased functioning either because of or shortly after the Starfish Prime[10,11] atmospheric nuclear weapon test in July 1962. Of the other two, one failed prior to the Starfish test and the other, Transit 4A, continued to operate for years.

The high failure rate of launch vehicles in those early years (30%) was soon overcome, but there was much to be learned about the radiation environments, both natural and enhanced, in which future spacecraft would need to operate. The Starfish Prime blast, over a megaton at a 248-mile altitude, knocked out electrical systems in Oahu some 800 miles away, and provided much data about the survivability of both space and ground assets in the event of a nuclear exchange. Figure 5 shows the effect of that blast on the solar cells onboard two of APL's spacecraft. Great concerns were raised about spacecraft lifetimes in orbit that shaped the early technological research and experimentation. We were one of a large number of organizations involved in trying to understand the issues of spacecraft reliability and

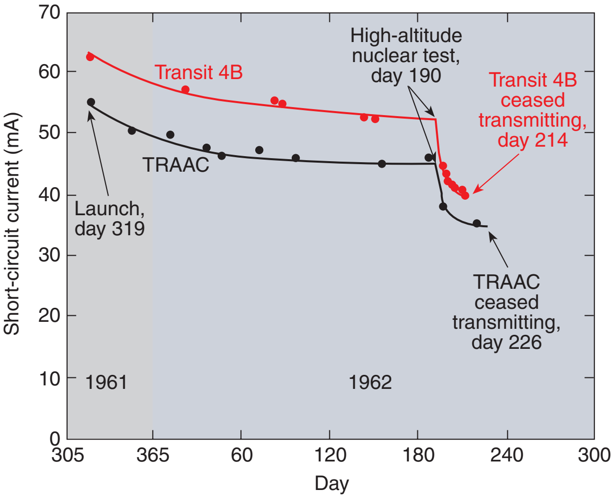

Figure 5. Short-circuit current versus time for two solar arrays in orbit at the time of the Starfish Prime high-altitude nuclear test. Degradation after the blast was due to exposure of the solar cells to greatly increased numbers of energetic charged particles trapped by the Earth's magnetic field.

develop systems that could operate in wartime; it was as a result of that one test that the IEEE started its annual Conference on Nuclear and Space Radiation Effects in the early 1960s.

APL conducted on-orbit reliability experiments beginning with the series of Transit spacecraft called 5E, launched in 1963 and 1964. Experimental circuitry was designed to measure and communicate the degradation of solid-state devices such as transistors. Eventually this work led to sufficient understanding of radiation effects to allow qualification by analysis and ground testing. Our cobalt-60 radiation source is in service today next to the Sentry 100 VLSI tester (Fig. 3), minimizing the delay between irradiation and parametric characterization of complex electronic parts.

The Laboratory continued to develop and improve the reliability of the Transit spacecraft. Lifetimes originally measured in months at inception in 1964 increased to more than a decade by the late 1960s when the first-generation operational spacecraft (Oscars) were produced by RCA's Astro-Electronics Division with technology transfer from APL. Meanwhile, we had been advising the Navy to develop a more robust series of second-generation spacecraft, applying newly developed reliability and survivability technology. The Navy agreed, leading to the design of an extremely sophisticated spacecraft that could survive autonomously through prescribed levels of high-energy radiation, electromagnetic pulse, and the secondary effects produced by such a burst of energy. Most challenging in this effort was the verification by ground test, for these effects are not easily reproduced in a controlled environment. Eventually, what began as a surprise failure mode for spacecraft was reduced to a manageable and predictable degradation, though still on the list as a potential killer.

Somewhat more benign and predictable, the natural space radiation environment has been a hazard for spacecraft electronics since the early days of the space program. Initially, total ionizing dose damage, in which the cumulative effect of long-term exposure leads to changes in the electrical characteristics of devices, was the primary cause of radiation-induced failure. Standard techniques were developed for producing less sensitive semiconductor devices and for shielding others.

However, in the mid-1980s, as electronics became smaller and more highly integrated, a new class of effects appeared that were different from total dose effects in that failures were induced by the interaction of a *single* energetic particle with a device. These "single event effects" are all produced by localized charge generation in a semiconductor device. The most common of these is the change in logic state of a memory cell or latch, often called a single event upset. At worst, this can cause a computer to perform an unintended operation or stop functioning entirely. Another major concern is the "latch-up" phenomenon in which a high current path is formed in an IC by the interacting particle. This usually permanently damages the affected device by self-heating, and in any case the device will not function until the conducting path is eliminated.

To enable the use of high-performance microelectronics in space, test techniques and mitigation strategies for single event effects have been developed at APL and elsewhere. In space, these effects are caused by energetic protons and heavy ions of both solar and galactic origin. APL has participated in the development of the Single Event Upset Test Facility at Brookhaven National Laboratory, which has become the premier facility for ground simulation of the space single event effects environment. Single event effects testing is expensive and time-consuming but necessary, particularly for devices not designed to survive in space. To reduce costs and speed testing, APL built a computer that is easily adapted to testing almost any kind of device. It allows great flexibility with minimum hardware changes, and can be easily integrated into the Brookhaven facility for beam control and monitoring.

Single event effects occur at the device level but have system-level impact. On a number of occasions, we have implemented system-level effects mitigation in spacecraft that require the performance of a sensitive device to accomplish a mission. These include watchdog timers for processors, error detection and correction devices for memories, current-limiting latch-up detection circuits, redundant processors, and triple-voting latches in gate-array circuits.

Nearly every APL spacecraft and instrument since the Topex altimeter has used single event effects mitigation. To date, these techniques have ensured that single event effects have not adversely affected spacecraft performance.

Overall System Reliability Assurance

Having discussed some aspects of spacecraft reliability engineering, let us digress to the broader subject of system reliability. Looking back at the life cycle of the Transit System, spanning the years from 1958 to 1996, we can see roughly 20 years of growth and improvement and 20 years of maintenance and phase-out. Here we will examine this latter phase. GPS was long planned to replace Transit in providing navigation fixes for the Poseidon/Trident fleet, and the schedule was sufficiently flexible that it became driven as much by funding as by technical requirements. After all, the Navy had not anticipated the longevity of the Oscar spacecraft series, and the dozen spares in storage in New Jersey seemed more than adequate, at least in quantity.

The obvious concern that the aging spacecraft in storage were somehow becoming less reliable was addressed primarily by periodic testing, which never yielded any evidence of deterioration. In fact, the handling probably posed a greater threat than age. The more serious issues did not lie within the "long gray line" of shipping containers that held the spacecraft and solar panels, but with the flight batteries and the launch vehicles. The batteries were worrisome because of design changes, particularly in separator materials, which were claimed to improve life span and reliability but were not proven under simulation of the particular parameter ranges of the Oscar spacecraft. By the mid-1980s it was understood that the Oscars were exhibiting a very long mean-time-to-failure, and that the usual failure mode was power system degradation. The Navy undertook life testing of the new batteries, but useful results would not appear early enough to support any change in plans.

So what came to drive major decisions and move money in the mid-1980s were two external constraints. GPS had slipped its original operational availability from 1982 to 1993, resulting in several extensions of the Transit Program, eventually through 1996. At the same time, NASA's plans to phase out the Scout launch vehicle were proceeding and the Navy could rely only on the vehicles it chose to acquire and store. The Scout solid rocket motors had to be used long before the end of the program. This prompted the Navy to (1) ask APL to carefully analyze the reliability of the on-orbit constellation of satellites, accounting for the known reliability of individual satellites, and (2) ask RCA/Astro, the spacecraft production contractor, to modify the launch configuration to accommodate *two* spacecraft per rocket. The latter effort included the necessary electronic modifications to operate the Oscar spacecraft as on-orbit spares, a novel idea at the time. RCA responded by equipping eight of these spacecraft with a commandable alternate downlink frequency that would not interfere with the signals from the operational spacecraft, and which thereby preserved the "always on" electrical and thermal balance that had proven so reliable on orbit.

Using the results of the APL constellation simulation, the Navy was able to develop a revised launch plan based on realistic failure rates for spacecraft that were both age- and health-dependent. The plan accounted for the shelf-life restrictions on vehicle motors and accommodated Transit phase-out schedules and funding profiles as they changed from year to year. What began as relatively simple statistical reliability models gradually evolved to a constellation simulation that accounted for orbit plane drift. It also accounted for the possibility of higher failure rates during wartime if nuclear devices were used that generated an enhanced radiation environment (as was seen in the 1962 high-altitude nuclear blast).

An interesting outcome of this investigation of the reliability of the operational Oscar spacecraft was the statistically significant dependence of failure rate on age. The demonstrated average time to failure of 14 years tells only part of the story. During the first 5 years (the original design life), the failure rate was 4 times lower than that. This, of course, was balanced by a higher failure rate for spacecraft which had passed the 15-year mark. Recognizing this fact was important to the outcome when the constellation was heavily populated with aging spacecraft. Figure 6 shows a two-parameter model of the failure rate using a Gamma function (other functional forms such as Weibull give similar results). This behavior is obvious logically, but contrary to the psychological feeling developed by watching nothing fail for long intervals. The main goal of these probabilistic representations of individual spacecraft was to combine all known information about the health of each asset plus best guesses at the unknowns into a probabilistic assessment of mission objectives involving the entire constellation. This assessment was essentially a level of confidence that (even with several years of extension of the Transit Program due to GPS delays) the exposure of Fleet submarines while surfacing to obtain position fixes would be sufficiently minimal. The Navy, as a result, did not need to seek new launch services to support the final years of the Transit Program.

CONCLUSION

The APL Space Department has earned a 40-year reputation as a high-quality provider. Sometimes we are even accused of being overly careful and conservative. It is easy to see how this mindset developed from our

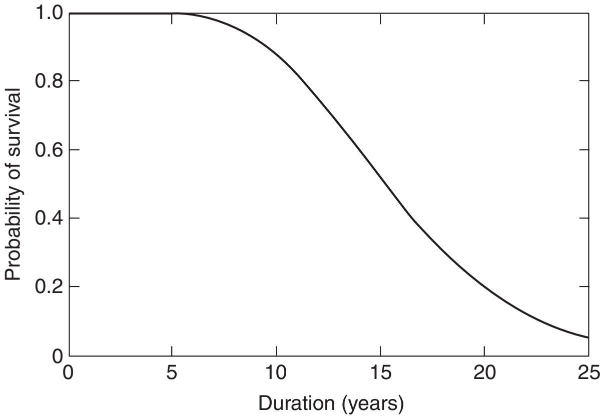

Figure 6. The Gamma function is a useful mathematical model for the probabilistic survival of systems that include "wear-out mechanisms" (batteries or gears) or have some tolerance to subsystem failures (such as complex systems with cross-strapped redundant components). It reflects a likelihood of failure that increases with time, in contrast to the simpler exponential function which applies when the likelihood of failure is constant in time. The curve shown was fitted to actual data for the Oscar spacecraft (Transit Program).

first big task: developing a strategically important operational system for the Navy. In the present environment of overconstrained government space budgets and the "better, faster, cheaper" mandate, it is as important as ever to ensure that innovative mission design is coupled with careful, reliable execution of that design. In this way we try to assure that missions are successfully accomplished, that the customer's objectives are met, and that the taxpayer's money is well spent.

REFERENCES

[1] Danchik, R. J., "The Navy Navigation Satellite System (Transit)," *Johns Hopkins APL Tech. Dig.* **5**(4), 323–329 (1984).

[2] Williams, W. C., "Lessons from NASA," *IEEE Spectrum*, 79–84 (Oct 1981).

[3] Nagrant, J., *Integrated Test Specification for Space Payload Equipment*, SDO 2387-1, JHU/APL, Laurel, MD (Feb 1982).

[4] QAD, Inc., product literature, available at http://www.qad.com/ (accessed 28 Jul 1999).

[5] Hoffman, E. J., and Ash, W. M. III, "Reliable Application of Plastic Encapsulated Microcircuits for Small Satellites," in *Proc. 8th Annual AIAA/USU Conf. on Small Satellites* (Aug 1994).

[6] Moor, A. F., Casasnovas, A., and Purwin, S. R., "The Case for Plastic-Encapsulated Microcircuits in Spaceflight Applications," *Johns Hopkins APL Tech. Dig.* **20**(1), 91–100 (1999).

[7] *Quality Systems—Model for Quality Assurance in Design, Development, Production, Installation, and Servicing*, American National Standard ANSI/ISO/ASQC Q9001-1994 (Aug 1994).

[8] *Electro-Mechanical Hardware Workmanship Standards Manual*, TSD-STD-800.1, JHU/APL, Laurel, MD (Jan 1991).

[9] NASA/GSFC Software Engineering Laboratory Web pages, available at http://sel.gsfc.nasa.gov/ (accessed 28 Jul 1999).

[10] *Operation Dominic*, available at http://www.fas.org/nuke/hew/Usa/Tests/Dominic.html (accessed 28 Jul 1999).

[11] *United States Nuclear Tests, July 1945 Through September 1992*, DE95006143, U.S. Department of Energy, Washington, DC (Dec 1994).

ACKNOWLEDGMENT: The authors wish to thank James D. Kinnison for his assistance in the description of single event effects and Harry K. Charles, Jr., for reviewing this article.

THE AUTHORS

WARD L. EBERT is the Assistant Department Head for Operations in the Space Department. He holds an A.B. in mathematics from Princeton University and a Ph.D. in mathematics from Case Western Reserve University. He joined APL's Space Department in 1969 and has served as Group Supervisor for the Guidance and Control Group and subsequently the Reliability and Quality Assurance Group. Dr. Ebert was involved in the development of the Navy Navigation Satellite System's orbital mechanics and ground software, and later the Transit spacecraft development and operations. During the 1980s he served as System Engineer in support of the Navy's launches of three Nova spacecraft and four stacks of two Oscars each, called SOOS (Stacked Oscars On Scout). His e-mail address is ward.ebert@jhuapl.edu.

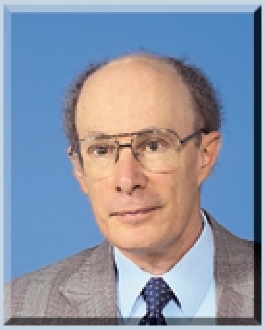

ERIC J. HOFFMAN received degrees in electrical engineering from M.I.T. and from Rice University, and joined APL in 1964. He has performed system engineering for many of APL's space communication and navigation systems, as well as for entire spacecraft. He has supervised communications and navigation design activities and led satellite conceptual designs. As Space Department Chief Engineer, he provides technical guidance for Department programs, promotes system engineering, and sets standards for design integrity, design review, configuration management, and test. Mr. Hoffman has taught space systems design courses at the Naval Academy, The Johns Hopkins University, National Taiwan University, NASA, and NSA. He has authored 60 papers on these topics and is a co-author of *Fundamentals of Space Systems*. Mr. Hoffman is a Fellow of the British Interplanetary Society and an Associate Fellow of the AIAA. His e-mail address is eric.hoffman@jhuapl.edu.