

Role of Software in Spacecraft Accidents

Nancy G. Leveson*

Massachusetts Institute of Technology, Cambridge, Massachusetts 02139

The first and most important step in solving any problem is understanding the problem well enough to create effective solutions. To this end, several software-related spacecraft accidents were studied to determine common systemic factors. Although the details in each accident were different, very similar factors related to flaws in the safety culture, the management and organization, and technical deficiencies were identified. These factors include complacency and discounting of software risk, diffusion of responsibility and authority, limited communication channels and poor information flow, inadequate system and software engineering (poor or missing specifications, unnecessary complexity and software functionality, software reuse without appropriate safety analysis, violation of basic safety engineering practices in the digital components), inadequate review activities, ineffective system safety engineering, flawed test and simulation environments, and inadequate human factors engineering. Each of these factors is discussed along with some recommendations on how to eliminate them in future projects.

Introduction

SOFTWARE is playing an increasingly important role in aerospace systems. Is it also playing an increasing role in accidents and, if so, what type of role? In the process of a research project to evaluate accident models, a variety of aerospace accidents were examined in detail that in some way involved software.^{1,2} Many of the factors were in common across several of the accidents. To prevent accidents in the future, we need to attack these problems.

The spacecraft accidents investigated were the explosion of the Ariane 5 launcher on its maiden flight in 1996; the loss of the Mars Climate Orbiter in 1999; the destruction of the Mars Polar Lander sometime during the entry, deployment, and landing phase in the following year; the placing of a Milstar satellite in an incorrect and unusable orbit by the Titan IV B-32/Centaur launch in 1999; and the loss of contact with the Solar Heliospheric Observatory (SOHO) spacecraft in 1998.

On the surface, the events and conditions involved in the accidents appear to be very different. A more careful, detailed analysis of the systemic factors, however, reveals striking similarities. Systemic factors are those that go beyond the specific technical causes, such as a flawed O-ring design in the Space Shuttle *Challenger* accident, and include the reasons why those failures or design errors were made. For *Challenger*, the latter include flawed decision making, poor problem reporting, lack of trend analysis, a “silent” or ineffective safety program, communication problems, etc. Systemic factors are those related to the overall system within which the technical device is developed and operated.

A difficulty was encountered in that several of the accident reports implicated the software, but then, for some unknown reason, never investigated the software development process in any depth to determine why the error was made. In some cases, it was possible to find information about the software development problems from sources outside the official accident investigation report. One conclusion from this observation might be that accident investigation boards must include more software experts and must more thoroughly investigate the reasons for the introduction of the errors and their lack of detection once introduced if we are to learn from our mistakes and improve our processes.

The accidents are first briefly described for those unfamiliar with them, and then the common factors are identified and discussed. These factors are divided into three groups: 1) flaws in the safety culture, 2) management and organizational problems, and 3) technical deficiencies.

The Accidents

Ariane 501

On 4 June 1996, the maiden flight of the Ariane 5 launcher ended in failure. About 40 s after initiation of the flight sequence, at an altitude of 2700 m, the launcher veered off its flight path, broke up, and exploded. The accident report describes what they called the “primary cause” as the complete loss of guidance and attitude information 37 s after start of the main engine ignition sequence (30 s after liftoff).³ The loss of information was due to specification and design errors in the software of the inertial reference system. The software was reused from the Ariane 4 and included functions that were not needed for Ariane 5 but were left in for “commonality.” In fact, these functions were useful but not required for the Ariane 4 either.

Mars Climate Orbiter

The Mars Climate Orbiter (MCO) was launched 11 December 1998 atop a Delta II launch vehicle. Nine and a half months after launch, in September 1999, the spacecraft was to fire its main engine to achieve an elliptical orbit around Mars and to skim through the Mars upper atmosphere for several weeks, in a technique called aerobraking, to move into a low circular orbit. On 23 September 1999, the MCO was lost when it entered the Martian atmosphere in a lower than expected trajectory. The investigation board identified what it called the “root” cause of the accident as the failure to use metric units in the coding of a ground software file used in the trajectory models.⁴ Thruster performance data were instead in English units.

Mars Polar Lander

Like MCO, Mars Polar Lander (MPL) was part of the Mars Surveyor program. It was launched 3 January 1999, using the same type of Delta II launch vehicle as MCO. Although the cause of the MPL loss is unknown, the most likely scenario is that the problem occurred during the entry, deployment, and landing sequence when the three landing legs were to be deployed from their stowed condition to the landed position.^{5,6} Each leg was fitted with a Hall effect magnetic sensor that generates a voltage when its leg contacts the surface of Mars. The descent engines were to be shut down by a command initiated by the flight software when touchdown was detected. The engine thrust must be terminated within 50 ms after touchdown to avoid overturning the lander. The flight software was also required to protect against a premature touchdown signal or a failed sensor in any of the landing legs.

Received 17 November 2002; revision received 2 June 2003; accepted for publication 9 June 2003. Copyright © 2003 by the American Institute of Aeronautics and Astronautics, Inc. All rights reserved. Copies of this paper may be made for personal or internal use, on condition that the copier pay the \$10.00 per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923; include the code 0022-4650/04 \$10.00 in correspondence with the CCC.

*Professor, Department of Aeronautics and Astronautics, Room 33-313, 77 Massachusetts Avenue; leveson@mit.edu.

The touchdown sensors characteristically generate a false momentary signal at leg deployment. This behavior was understood, and the flight software should have ignored it. The software requirements did not specifically describe these events, however, and consequently, the software designers did not account for them. It is believed that the software interpreted the spurious signals generated at leg deployment as valid touchdown events. When the sensor data were enabled at an altitude of 40 m, the software shut down the engines and the lander free fell to the surface, impacting at a velocity of 22 m/s and was destroyed.

Titan/Centaur/Milstar

On 30 April 1999, a Titan IV B-32/Centaur TC-14/Milstar-3 was launched from Cape Canaveral. The mission was to place the Milstar satellite in geosynchronous orbit. An incorrect roll rate filter constant zeroed the roll rate data, resulting in the loss of roll axis control and then yaw and pitch control. The loss of attitude control caused excessive firings of the reaction control system and subsequent hydrazine depletion. This erratic vehicle flight during the Centaur main engine burns in turn led to an orbit apogee and perigee much lower than desired, placing the Milstar satellite in an incorrect and unusable low elliptical final orbit instead of the intended geosynchronous orbit.

The accident investigation board concluded that failure of the Titan IV B-32 mission was due to an inadequate software development, testing, and quality assurance process for the Centaur upper stage.⁷ That process did not detect the incorrect entry by a flight software engineer of a roll rate filter constant into the inertial navigation unit software file.

The roll rate filter itself was included early in the design phase of the first Milstar spacecraft, but the spacecraft manufacturer later determined that filtering was not required at that frequency. A decision was made to leave the filter in place for the first and later Milstar flights for “consistency.”⁷

SOHO

SOHO was a joint effort between NASA and ESA to perform helioseismology and to monitor the solar atmosphere, corona, and wind. The spacecraft completed a successful two-year primary mission in May 1998 and then entered into its extended mission phase. After roughly 2 months of nominal activity, contact with SOHO was lost 25 June 1998. The loss was preceded by a routine calibration of the spacecraft’s three roll gyroscopes and by a momentum management maneuver.

The flight operations team had modified the ground operations procedures as part of a ground systems reengineering effort to reduce operations costs and streamline operations, to minimize science downtime, and to conserve gyro life. Though some of the modifications were made at the request of the SOHO science team, they were not necessarily driven by any specific requirements changes. A series of errors in making the software changes along with errors in performing the calibration and momentum management maneuver and in recovering from an emergency safing mode led to the loss of telemetry.⁸ Communication with the spacecraft was eventually restored after a gap of 4 months.

Flaws in the Safety Culture

The safety culture is the general attitude and approach to safety reflected by those working in an industry. The accident reports all described various aspects of complacency and a discounting or misunderstanding of the risks associated with software.

Success is ironically one of the progenitors of accidents when it leads to overconfidence and cutting corners or making tradeoffs that increase risk. This phenomenon is not new, and it is extremely difficult to counter when it enters the engineering culture in an organization. Complacency is the root cause of most of the other accident factors described in this paper and was exhibited in all of the accidents studied.

The MCO report noted that because the Jet Propulsion Laboratory’s (JPL’s) navigation of interplanetary spacecraft had worked

well for 30 years, there was widespread perception that “orbiting Mars is routine”⁴ and inadequate attention was devoted to navigation risk management and contingency planning. The MCO investigators found that project management teams appeared primarily focused on meeting mission cost and schedule objectives and did not adequately focus on mission risk. A recommendation common to several of the accident reports was to pay greater attention to risk identification and management.

The official report on the MPL loss concludes that the pressure of meeting the cost and schedule goals resulted in an environment of increasing risk in which too many corners were cut in applying proven engineering practices and in the checks and balances necessary for mission success. “Lack of adequate risk identification, communication, management, and mitigation compromised mission success.”¹⁶

In the SOHO loss, overconfidence and complacency, according to the accident report, led to inadequate testing and review of changes to ground-issued software commands to the spacecraft, a false sense of confidence in the team’s ability to recover from a safe-hold mode (emergency sun reacquisition) from which a recovery sequence must be commanded and executed under ground operator control, the use of tight schedules and compressed timelines that eliminated any time to handle potential emergencies, inadequate contingency planning, responses to emergencies without taking the designed-in time to consider the options, etc. Protections built into the process, such as formal reviews of critical decisions by senior management and engineering staff, were bypassed. The functional content of an operational procedure was changed without appropriate documentation and review of the changes.

After two previous SOHO spacecraft retreats to safe mode, the software and procedures were not reviewed because (according to the accident report) higher priority had been assigned to other tasks. The report concludes that the success in recovering from the previous safe mode entries led to overconfidence by the operations team in their ability to recover and a lack of appreciation of the risks involved in entering and recovering from the safing mode.

The Ariane 5 accident report notes that software was assumed to be correct until it was shown to be faulty. As noted by the Ariane accident investigation board, the opposite assumption is more realistic. A similar attitude prevailed in Titan/Centaur operations. For example, on the day of the launch, the attitude rates for the vehicle on the launch pad were not properly sensing the Earth’s rotation rate (the software consistently reporting a zero roll rate), but no one had the responsibility to monitor that rate data specifically or to perform a check to see whether the software attitude filters were operating correctly. In fact, there were no formal processes to check the validity of the filter constants or to monitor attitude rates once the flight tape was actually loaded into the inertial navigation unit at the launch site. Potential hardware failures are usually checked up to launch time, but it may have been assumed that testing removed all software errors and no further checks were needed.

Although management may express concern for safety and mission risks, true priorities are shown during resource allocation. Although budget decisions are always difficult when resources are constrained, and budgets are almost always less than optimal, the first things to be cut are often system safety, system engineering, mission assurance, and operations, which are assigned a low priority and assumed to be the least critical parts of the project.

In the Milstar satellite loss, the Titan Program Office had no permanently assigned civil service or military personnel nor full-time support to work the Titan/Centaur software. They had decided that because the software was “mature, stable, and had not experienced problems in the past,”⁷ they could best use their limited resources available after the initial development effort to address hardware issues. The Titan program office had cut support for monitoring the software development and test process by 50% since 1994 and had greatly cut the number of engineers working launch operations.

The SOHO Mission Management Plan required that the NASA Project Operations Director be responsible for programmatic matters, provide overall technical direction to the flight operations team, and interface with the ESA technical support director. The position had been descope over time by NASA from a dedicated individual

during launch and commissioning to one NASA individual spending less than 10% of his time tracking SOHO operations. ESA was to retain ownership of the spacecraft and to be responsible for its technical integrity and safety, but they were understaffed to perform this function in other than routine situations. In both SOHO and MCO, the operations group did not have a mission assurance manager.

Complacency can also manifest itself in a general tendency of management and decision makers to discount unwanted evidence of risk. A “culture of denial”⁹ arises in which any evidence of significant risk is dismissed. In the MCO, SOHO, and Titan losses, warning signs existed that the software was flawed, but they went unheeded.

The problems experienced with the MCO software during the early stages of the flight did not seem to raise any red flags. During the first four months of the MCO mission, the ground software angular momentum desaturation files were not used in the orbit determination process because of multiple file format errors and incorrect spacecraft attitude data specifications. Four months were required to fix the files. Almost immediately (within a week), it became apparent that the files contained anomalous data that were indicating underestimation of the trajectory perturbations due to desaturation events. Despite all of these hints that there were serious problems in the software and perhaps the development process, reliance was still placed on the supposedly fixed software without extra manual checks or alternative calculations to check the results.

Three months before the loss of the SOHO telemetry, ground software problems had triggered an emergency sun reacquisition (a safe hold mode entered when there are attitude control anomalies) and a shortcut in the recovery from this emergency sun reacquisition led to a second one. A resulting recommended comprehensive review of the software and procedures had not been implemented before the accident because higher priority had been assigned to other tasks.

Engineers noticed the problems with the Titan/Centaur software after it was delivered to the launch site and they were reported back to Lockheed Martin Astronautics (LMA) in Denver, but nobody seemed to take them seriously.

Some of the complacency can arise from a misunderstanding of the risks associated with software. Throughout the accident reports, there is an emphasis on failures as the cause of accidents and redundancy as the solution. Accidents involving software, however, are usually “system accidents”¹⁰ that result from dysfunctional interactions among components, not from individual component failure. All of these accidents (as well as almost all of the software-related accidents known to the author) resulted from the software doing something wrong rather than the computer hardware or software failing to operate at all. In fact, in most cases the software or hardware components operated according to their specifications, that is, they did not fail, but the combined behavior of the components led to disastrous system behavior.

All of the accidents investigated for this paper displayed some aspects of system accidents. System accidents are caused by interactive complexity and tight coupling.¹⁰ Software allows us to build systems with a level of complexity and coupling that is beyond our ability to control; in fact, we are building systems where the interactions among the components (often controlled by software) cannot all be planned, understood, anticipated, or guarded against. This change is not solely the result of using digital components, but it is made possible because of the flexibility of software. Note that the use of redundancy only makes the problem worse: The added complexity introduced by redundancy has resulted in accidents that otherwise might not have occurred.

The Ariane 5 accident report notes that according to the culture of the Ariane program, only random failures were addressed and they were primarily handled with redundancy. The engineers designing the Ariane 5 inertial guidance system opted to shut down the computer when an exception was raised in an unnecessary function (the alignment function after takeoff): “The reason behind this drastic action lies in the culture within the Ariane programme of only addressing random hardware failures. From this point of view, exception-or error-handling mechanisms are designed for a random hardware failure which can quite rationally be handled by a backup system.”³ This approach obviously failed in the Ariane 5’s first flight

when both the primary and backup (redundant) inertial reference system computers shut themselves down, exactly as they were designed to do, while processing the same unexpected input value.

Software and digital systems require changes to important aspects of engineering practice. Not only are failures not random (if the term failure makes any sense when applied to something like software that is pure design separated from the physical realization of that design), but the complexity of most software precludes examining all of the ways it could misbehave. Also, the failure modes (the way it misbehaves) can be very different than for physical devices. The JPL MPL accident report, like others, recommends using failure modes and effects analysis (FMEA) and fault tree analysis (FTA) along with appropriate redundancy to eliminate failures. However, these techniques were developed to cope with random wearout failures in hardware and are not very effective against design errors, the only type of error found in software. Although computer hardware can fail, software itself is pure design, and thus, all errors are design errors and appropriate techniques for handling design errors must be used.

Management and Organizational Factors

The five accidents studied during this exercise, as well as most other major accidents, exhibited common organizational and managerial flaws, notably a diffusion of responsibility and authority, limited communication channels, and poor information flow.

Diffusion of Responsibility and Authority

In all of the accident reports, serious organizational and communication problems among the geographically dispersed partners are mentioned or implied by the recommendations. Responsibility was diffused without complete coverage and without complete understanding by anyone about what all of the groups were doing. Roles were not clearly allocated.

Both the Titan and Mars 1998 programs were transitioning to process insight from process oversight, reflecting different levels of feedback control over lower levels and a change from prescriptive management control to management by objectives, where the objectives are interpreted and satisfied according to the local context.

Just as the MPL reports noted that faster–better–cheaper was not defined adequately to ensure that it meant more than simply cutting budgets, this change in management role from oversight to insight seems to have been implemented on the Mars 1998 projects as well as the Titan/Centaur program simply as a reduction in personnel and budgets without assuring that anyone was responsible for specific critical tasks. For example, the MCO report says: “NASA management of out-of-house missions was changed from ‘oversight’ to ‘insight’—with far fewer resources devoted to contract monitoring.”⁴ One of the results of faster–better–cheaper was a reduction in workforce while maintaining an expectation for the same amount of work to be accomplished. In many of these accidents, the people were simply overworked, sometimes driven by their own dedication.

The process used in the Titan/Centaur program to develop the constants used in the flight software was neither well defined nor completely understood by any of the multiple players involved in that process. Procedures for creating and updating the database were not formally documented and were left to the flight software engineer’s discretion. The root problem is probably not the lack of documentation itself but the lack of anyone being in charge of the entire process. There were several people who performed part of the process, but they only completely understood their own specific part. The Accident Investigation Board could not identify a single process owner responsible for understanding, designing, documenting, controlling configuration, and ensuring proper execution of the overall software development process. Instead, responsibility was diffused among the various partners, without complete coverage. For example, the Centaur inertial navigation unit (INU) consists of two major software components developed by different companies. LMA developed the flight control system (FCS) software and was responsible for overall INU testing. Honeywell, Inc., developed the inertial measurement system (IMU) and was partially responsible for its software development and testing. The erroneous constants

were processed by the Honeywell-built IMU but were designed and tested by LMA. LMA, in turn, focused its flight software process on the FCS and not the IMS software and had little knowledge of IMS operations.

Titan launch operations exhibited the same problems. The Space and Missile Systems Center Launch Directorate and the 3rd Space Launch Squadron had undergone personnel reductions and were also transitioning from a task oversight to a process insight role. That transition had not been managed by a detailed plan. According to the accident report, U.S. Air Force responsibilities under the insight concept were not well defined, and how to perform those responsibilities had not been communicated to the work force. There was no master surveillance plan in place to define the tasks for the engineers remaining after the reductions, so that the launch personnel used their best engineering judgment to determine which tasks they should perform, which tasks to monitor, and how closely to analyze the data from each task. This approach, however, did not ensure that anyone was responsible for specific tasks. In particular, on the day of the launch, the attitude rates for the vehicle on the launch pad were not properly sensing the Earth's rotation rate, but nobody had the responsibility to monitor that rate data specifically or to check the validity of the roll rate, and no reference was provided with which to compare. Thus, when the anomalies occurred during launch preparations that clearly showed a problem existed with the software, nobody had the responsibility or ability to follow up on them.

In MPL, there was essentially no JPL line management involvement or visibility into the software development and minimal involvement by JPL technical experts. Similarly, the MCO report suggests that authority and accountability were a significant issue in the accident and that roles and responsibilities were not clearly allocated. There was virtually no JPL oversight of LMA subsystem development. The MCO report says,

Line managers at the field centers need to be held accountable for the success of all missions at their centers. . . . The line management should be held accountable for asking the right questions at meetings and reviews, and getting the right people to those reviews to uncover mission-critical issues and concerns early in the program.⁴

For SOHO, a transfer of management authority to the SOHO Project Scientist resident at Goddard Space Flight Center left no manager, either from NASA or ESA, as the clear champion of spacecraft health and safety. Instead, the accident report concludes that the transfer encouraged management decisions that maximized science return over spacecraft risk. In addition, the decision structure for real-time divergence from agreed-upon ground and spacecraft procedures was far from clear: The flight operations staff was apparently able to change procedures without proper review.

The Ariane 501 accident report is almost totally silent about organizational structure problems: It does not describe the allocation of responsibility and authority for safety, nor does it mention any organizational or management factors that may have influenced the accident. There is one hint that there may have been problems, however, in a recommendation at the end of the report that says, "A more transparent organization of the cooperation among partners in the Ariane 5 programme must be considered. Close engineering co-operation, with clear cut authority and responsibility, is needed to achieve system coherence, with simple and clear interfaces between partners."³

Inadequate transition from development to operations played a role in several of the accidents. Engineering management sometimes has a tendency to focus on development and to put less effort into planning the operational phase. The MCO report states, "The overall project plan did not provide for a careful handover from the development project to the very busy operations project. Transition from development to operations—as two separate teams—disrupted continuity and unity of shared purpose."⁴

The operations teams (in those accidents that involved operations) also seemed isolated from the developers. The MCO report notes this isolation and provides as an example that the operators did not know until long after launch that the spacecraft sent down tracking data

that could have been compared with the ground data, which might have identified the software error while it could have been fixed. The operations crew for the Titan/Centaur also did not detect the obvious software problems partly because of a lack of the knowledge required to detect them.

Most important, responsibility for safety does not seem to have been clearly defined outside of the quality assurance function on any of these programs. All of the accident reports (except the Titan/Centaur) are surprisingly silent about their safety programs. One would think that the safety activities and why they had been ineffective would figure prominently in the reports.

Safety was originally identified as a separate responsibility by the U.S. Air Force during the ballistic missile programs of the 1950s and 1960s to solve exactly the problems seen in these accidents: to make sure that safety is given due consideration in decisions involving conflicting pressures and that safety issues are visible at all levels of decision making. An extensive system safety program was developed by NASA after the Apollo launch pad fire in 1967. However, the *Challenger* accident report noted that the system safety program had become "silent"¹¹ over time and through budget cuts. Has this perhaps happened again? Or are the system safety efforts just not handling software effectively?

One common mistake is to locate the safety efforts within the quality assurance function. Placing safety only under the assurance umbrella instead of treating it as a central engineering concern is not going to be effective, as has been continually demonstrated by these and other accidents. Whereas safety is certainly one property (among many) that needs to be assured, safety cannot be engineered into a design through after-the-fact assurance activities alone.

Having an effective safety program cannot prevent errors of judgment in balancing conflicting safety, schedule, and budget constraints, but a safety program can at least make sure that decisions are informed and that safety is given due consideration. It also ensures that someone is focusing attention on what the system is not supposed to do, that is, the hazards, and not just on what it is supposed to do. Both perspectives are necessary if safety and mission assurance are to be optimized.

Limited Communication Channels and Poor Information Flow

All of the accident reports mention poor information flow and communication problems except the Ariane 5, which includes very little information beyond the technical details. The Titan/Centaur accident report, for example, notes that "fragmentation/stovepiping in the flight software development process, coupled with the lack of an overall defined process, resulted in poor and inadequate communication and interfacing among the many partners and subprocesses."⁷ The report suggests that many of the various partners were confused about what the other groups were doing. For example, the LMA software group personnel, who created the database from which the erroneous load tape constants were generated, were not aware that the independent verification and validation testing did not use the as-flown constants but instead used default values. The company responsible for the independent verification and validation (Analex-Denver) did not know that the division actually doing the independent verification and validation (Analex-Cleveland) was only verifying the functionality of the design constant and not what was actually loaded into the Centaur for flight. The Defense Contract Management Command software surveillance personnel were not aware that the filter constants contained in the flight software were generated by a manual input and were never tested by LMA in their preflight simulation nor subjected to independent verification and validation by Analex-Cleveland.

All of the accidents involved one engineering group not getting the information they needed from another engineering group. The MCO report cited deficiencies in communication between the project development team and the operations team. For example, the report notes that "Critical information on the control and desaturation of the MCO momentum was not passed to the operations navigation team."⁴ As another example, a decision was made that the "barbecue mode"⁴ (a daily 180-deg flip to cancel angular momentum buildup) was not needed and it was deleted from the spacecraft operations

plan, but the operations navigation team was never notified. Communication was poor in the other direction, too. Throughout the first nine months of the MCO mission, concerns regarding discrepancies observed between navigation solutions were reported by the navigation operations team only informally and were not communicated effectively to the spacecraft operations team or project management.

A significant factor in the MPL loss was that test results and new information about the Hall effect sensors derived during testing was not communicated to all of the component designers that needed it. In general, system engineering on several of the projects did not keep abreast of test results from all areas and communicate the findings to other areas of the development project. The MPL report concludes that the effect of inadequate peer interaction was, in retrospect, a major problem that led to a breakdown in intergroup communications. Communication is one of the most important functions in any large, geographically distributed engineering project and must be carefully planned and fostered.

The Titan/Centaur accident also involved critical information not getting to the right people. For example, tests right before launch detected the zero roll rate, but there was no communication channel established for getting that information to those who could understand it. A guidance engineer at the launch site noticed the anomalous roll rates and called LMA in Denver, leaving a voice mail message to call her or her supervisor. She also sent an email to her supervisor at Cape Canaveral explaining the situation. Her supervisor was on vacation and was due back at the office the next Monday, but the engineer herself was scheduled to work the second shift that day. Two LMA engineers in Denver, the control dynamics engineer who had originally specified the filter values and his supervisor, listened to the voice mail from the launch site guidance engineer and called her supervisor, who had just returned from vacation. He was initially unable to find the email she had sent him during their conversation and said he would call back. By the time he called back, the control dynamics engineer who had created the filter values had left his supervisor's office. At no time did the LMA Denver engineers speak directly with the launch site guidance engineer who had originally noticed the anomaly.

SOHO had similar communication problems between the operations team and technical experts. For example, when a significant change to procedures was implemented, an internal process was used and nobody outside the flight operations team was notified.

In the Titan/Centaur and MCO accidents, there was evidence that a problem existed before the loss occurred, but there was no communication channel established for getting the information to those who could understand it and to those making decisions, or, alternatively, the problem-reporting channel was ineffective in some way or was simply unused. The MCO report concludes that project leadership did not instill the necessary sense of authority and accountability in workers that would have spurred them to broadcast problems they detected so that those problems might be "articulated, interpreted, and elevated to the highest appropriate level, until resolved."⁴ The report concludes that "Institutional management must be accountable for ensuring that concerns raised in their own area of responsibility are pursued, adequately addressed, and closed out."⁴

Researchers have found that the second most important factor in the success of any safety program (after top management concern) is the quality of the hazard information system. Both collection of critical information as well as dissemination to the appropriate people for action are required, but these activities were haphazard at best for most of the projects involved in these accidents. The MCO report concludes that lack of discipline in reporting problems and insufficient followup was at the heart of the mission's navigation mishap. Email was used to solve problems rather than the problem tracking system: "A critical deficiency in Mars Climate Orbiter project management was the lack of discipline in reporting problems and insufficient follow-up. The primary, structured problem-reporting procedure used by the Jet Propulsion Laboratory—the Incident, Surprise, Anomaly process—was not embraced by the whole team."⁴

For SOHO, critical information about the required operation of gyros used for changing the software was also provided informally to the flight operations team via email.

In the Titan/Centaur loss, the use of voice mail and email implies that either there was no formal anomaly reporting and tracking system or the formal reporting procedure was not known or used by the process participants for some reason. The report states that there was confusion and uncertainty as to how the roll rate anomalies should be reported, analyzed, documented, and tracked because it was a "concern" and not a "deviation."⁷ There is no explanation of these terms.

In all of the accidents (except for Ariane, where anomaly reporting is not mentioned), the existing formal anomaly reporting system was bypassed, and informal email and voice mail was substituted. The problem is clear but not the cause, which was not included in the reports and perhaps not investigated. When a structured process exists and is not used, there is usually a reason. Some possible explanations may be that the system is difficult or unwieldy to use or it involves too much overhead. Such systems may not be changing as new technology changes the way engineers work.

There is no reason why reporting something within the problem-reporting system should be much more cumbersome than adding an additional recipient to email. Large projects have successfully implemented informal email processes for reporting anomalies and safety concerns or issues to system safety personnel. New hazards and concerns will be identified throughout the development process and into operations, and there must be a simple and nononerous way for software engineers and operational personnel to raise concerns and safety issues and get questions answered at any time.

Technical Deficiencies

These cultural and managerial flaws manifested themselves in the form of technical deficiencies: 1) inadequate system and software engineering, 2) inadequate review activities, 3) ineffective system safety engineering, 4) inadequate human factors engineering, and 5) flaws in the test and simulation environments.

Inadequate System and Software Engineering

For any project as complex as those involved in these accidents, good system engineering is essential for success. In some of the accidents, system engineering resources were insufficient to meet the needs of the project. For example, the MPL report notes that insufficient system engineering during the formulation stage led to important decisions that ultimately required more development effort than originally foreseen, as well as inadequate baseline decisions and hazard identification. In others, the process followed was flawed, such as in the flowdown of system requirements to software requirements or in the coordination and communication among project partners and teams. As just one example, the MCO report notes that navigation requirements were set at too high a management level, and that there was insufficient flowdown to the subsystem level and inadequate validation of the requirements.

The Centaur software process was developed early in the Titan program, and many of the individuals who designed the original process were no longer involved in it due to corporate mergers and restructuring and the maturation and completion of the Titan/Centaur design and development. The accident report notes that much of the system and process history was lost with their departure, and therefore, nobody knew enough about the overall process to detect that it omitted any testing with the actual load tape or knew that the test facilities had the capability of running the type of test that could have caught the error.

Preventing system accidents falls into the province of system engineering; those building individual components have little control over events arising from dysfunctional interactions among components. As the systems we build become more complex (much of that complexity being made possible by the use of computers), system engineering will play an increasingly important role in the engineering effort. In turn, system engineering will need new modeling and analysis tools that can handle the complexity inherent in the systems we are building. Appropriate modeling methodologies will have to include software, hardware, and human components of systems.

Given that software played a role in all of the accidents, it is surprising the reports reflected so little investigation of the practices

that led to the introduction of the software flaws and a dearth of recommendations to fix them. In some cases, software processes were declared in the accident reports to have been adequate when the evidence shows they were not.

The accidents all involved very common system and software engineering problems, including poor specification practices, unnecessary complexity and software functions, software reuse without appropriate safety analysis, and violation of basic safety engineering design practices in the digital components.

Poor or Missing Specifications

Almost all software-related aerospace accidents (and accidents in other industries) have been related to flawed requirements and misunderstanding about what the software should do: The software performed exactly as the designers intended (it did not fail), but the designed behavior was not safe from a system viewpoint.¹² There is not only anecdotal but some hard data to support this hypothesis. Lutz examined 387 software errors uncovered during integration and system testing of the Voyager and Galileo spacecraft.¹³ She concluded that the software errors identified as potentially hazardous to the system tended to be produced by different error mechanisms than nonsafety-related software errors. She showed that for these two spacecraft, the safety-related software errors arose most commonly from 1) discrepancies between the documented requirements specifications and the requirements needed for correct functioning of the system and 2) misunderstandings about the software's interface with the rest of the system. This experiential evidence points to a need for better specification review and analysis.

All of the reports refer to inadequate specification practices. The Ariane accident report mentions poor specification practices in several places and notes that the structure of the documentation obscured the ability to review the critical design decisions and their underlying rationale. Inadequate documentation of design rationale to allow effective review of design decisions is a very common problem in system and software specifications. The Ariane report recommends that justification documents be given the same attention as code and that techniques for keeping code and its justifications consistent be improved.

The MCO report contains little information about the software engineering practices but hints at specification deficiencies in statements about "JPL's process of cowboy programming" and "the use of 20-year-old trajectory code that can neither be run, seen, or verified by anyone or anything external to JPL."¹⁴

The MPL report notes that the system-level requirements document did not specifically state the failure modes the requirement was protecting against (in this case possible transients) and speculates that the software designers or one of the reviewers might have discovered the missing requirement if they had been aware of the rationale underlying the requirements. The small part of the requirements specification shown in the accident report (which may very well be misleading) seems to avoid all mention of what the software should not do. In fact, standards and industry practices often forbid such negative requirements statements. The result is that software specifications often describe nominal behavior well but are very incomplete with respect to required software behavior under off-nominal conditions and rarely describe what the software is not supposed to do. Most safety-related requirements and design constraints are best described using such negative requirements or design constraints. In addition, the requirements flowdown process for MPL was clearly flawed, and the rationale for requirements did not appear to be included in the specification.

Not surprising, the interfaces were a source of problems. It seems likely from the evidence in several of the accidents that the interface documentation practices were flawed. The MPL report includes a recommendation that in the future "all hardware inputs to the software must be identified. . . . The character of the inputs must be documented in a set of system-level requirements."¹⁵ This information is usually included in the standard interface specifications, and it is surprising that it was not.

There are differing accounts of what happened with respect to the MCO incorrect units problem. The official accident report seems to

place blame on the programmers and recommends that the software development team be provided additional training in "the use and importance of following the Mission Operations Software Interface Specification (SIS)."¹⁴ Although not included in the official NASA MCO accident report, Oberg in a paper on the accident¹⁴ claims that JPL never specified the units to be used. It is common for specifications to be incomplete or not to be available until late in the development process.

A different explanation for the MCO units error was provided by the developers.¹⁵ According to them, the files were required to conform to a Mars Global Surveyor (MGS) heritage software interface specification. The equations used in the erroneous calculation were supplied by the vendor in English units:

Although starting from MGS-heritage software, the coded MGS thruster equation had to be changed because of the different size RCS [Reaction Control System] thruster that MCO employed (same vendor). As luck would have it, the 4.45 conversion factor, although correctly included in the MGS equation by the previous development team, was not immediately identifiable by inspection (being buried in the equation) or commented in the code in an obvious way that the MCO team recognized it. Thus, although the SIS required SI units, the new thruster equation was inserted in the place of the MGS equation—without the conversion factor.¹⁵

This explanation raises questions about the other software specifications, including the requirements specification, which seemingly should include descriptions of the computations to be used. Either these did not exist, or the software engineers did not refer to them when making the change. Formal acceptance testing apparently did not use the (MGS) software interface specification because the test oracle (computed manually) used for comparison contained the same error as the output file.¹⁵

Complete and understandable specifications are not only necessary for development, but they are critical for operations and the handoff between developers, maintainers, and operators. In the Titan/Centaur accident, nobody other than the control dynamics engineers who designed the roll rate constants understood their use or the impact of filtering the roll rate to zero. When discrepancies were discovered right before the Titan/Centaur/Milstar launch, as noted earlier, nobody understood them. The MCO operations staff also clearly had inadequate understanding of the automation and, therefore, were unable to monitor its operation effectively.

The SOHO accident report mentions that no hard copy of the software command procedure set existed, and the latest versions were stored electronically without adequate notification when the procedures were modified. The report also states that the missing software *enable* command (which led to the loss) had not been included in the software module due to a lack of system knowledge of the person who modified the procedure: He did not know that an automatic software function must be reenabled each time gyro A was despun. The information had been provided, but via email. Such information, particularly about safety-critical features, obviously needs to be clearly and prominently described in the system specifications.

Good specifications that include requirements tracing and design rationale are critical for complex systems, particularly those that are software controlled. They also must be reviewable and reviewed in depth by domain experts.

Unnecessary Complexity and Software Functionality

One of the most basic concepts in engineering critical systems is to keep it simple. "The price of reliability is the pursuit of the utmost simplicity. It is a price which the very rich find most hard to pay."¹⁶

The seemingly unlimited ability of software to implement desirable features often, as in the case of most of the accidents examined in this paper, pushes this basic principle into the background: Creeping Featurism is a common problem in software-intensive systems:

"And they looked upon the software, and saw that it was good. But they just had to add this one other feature." . . . A project's specification rapidly becomes a wish list. Additions to the list encounter

little or no resistance. We can always justify one more feature, one more mode, one more gee-whiz capability. And don't worry, it'll be easy—after all, it's just software. We can do anything.

In one stroke, we are free of nature's constraints. This freedom is software's main attraction, but unbounded freedom lies at the heart of all software difficulty (Frank McCormick, unpublished essay).

All of the accidents, except MCO, involved either unnecessary software functions or software operating when it was not necessary. The MCO report does not mention or discuss the software features.

Both the Ariane and Titan/Centaur accidents involved software functions that were not needed, but surprisingly the decision to put in these unneeded features was not questioned in the accident reports. The software alignment function in the reused Ariane 4 software had no use in the different Ariane 5 design. The alignment function was designed to cope with the unlikely event of a hold in the Ariane 4 countdown: The countdown could be restarted and a short launch window could still be used. The feature had been used once (in 1989 in flight 33 of the Ariane 4). The Ariane 5 has a different preparation sequence and cannot use the feature at all. In addition, the alignment function computes meaningful results only before liftoff: During flight, it serves no purpose, but the problem occurred while the function was operating after liftoff.

The Ariane accident report does question the advisability of retaining the unused Ariane 4 alignment function in the Ariane 5 software, but it does not question whether the Ariane 4 software should have included such a nonrequired but convenient software function in the first place. Outside of its effect on reuse (which may reasonably not have been contemplated during Ariane 4 development), a tradeoff was made between possibly delaying a launch and simplifying the software.

The MPL accident also involved software that was executing when it was not necessary to execute, although in that case the function was required at a later time in the descent sequence. The report states that the decision to start the software early was based on trying to avoid transients in CPU loading. The tradeoff in this case may have seemed justified, but executing software when it is not needed or including unnecessary code raises risk significantly. The MPL independent accident investigation report also noted that requirements creep was a factor in the accident.

The Titan/Centaur accident report explains that the software roll rate filter involved in the loss of the Milstar satellite was not needed but was kept in for consistency. The same justification is used to explain why the unnecessary software function leading to the loss of the Ariane 5 was retained from the Ariane 4 software. Neither report explains why consistency was assigned such high priority. Although changing software that works can increase risk, executing unnecessary software functions is also risky.

For SOHO, there was no reason to introduce a new function into the module that eventually led to the loss. A software function already existed to perform the required maneuver and could have been used. There was also no need to despin gyro A between gyro calibration and the momentum maneuvers.

In all of these projects, tradeoffs were obviously not considered adequately (considering the consequences), perhaps partially due to complacency about software risk. The more features included in software and the greater the resulting complexity (both software complexity and system complexity), the harder and more expensive it is to test, to provide assurance through reviews and analysis, to maintain, and to reuse in the future. Well-reasoned tradeoff decisions must include a realistic assessment of the impact of these extra difficulties. Software functions are not free.

Software Reuse or Changes without Appropriate Safety Analysis

Reuse and the use of commercial off-the-shelf software (COTS) is common practice today in embedded software development. The Ariane 5 software involved in the loss was reused from the Ariane 4. According to the MCO developers, the small forces software was reused from the MGS project, with the substitution of a new thruster equation.¹⁵ Technical management accepted the "just like MGS"⁴ argument and did not focus on the details of the software. The

SOHO software was changed without appropriate analysis of the changes.

It is widely believed that because software has executed safely in other applications, it will be safe in the new one. This misconception arises from confusion between software reliability and safety: Most accidents involve software that is doing exactly what it was designed to do, but the designers misunderstood what behavior was required and would be safe, that is, it reliably performs the wrong function.¹² The blackbox (externally visible) behavior of a component can only be determined to be safe by analyzing its effects on the system in which it will be operating, that is, by considering the specific operational context. The fact that software has been used safely in another environment provides no information about its safety in the current one. In fact, reused software is probably less safe because the original decisions about the required software behavior were made for a different system design and were based on different environmental assumptions. Changing the environment in which the software operates makes all previous usage experience with the software irrelevant for determining safety.

The problems in the Ariane 5 software arose from an overflow of a 16-bit integer variable. The horizontal velocity of the Ariane 4 cannot reach a value beyond the limit of the software, that is, the value will always fit in 16 bits. The initial acceleration and trajectory of the Ariane 5, however, leads to a buildup of horizontal velocity five times more rapid than that of the Ariane 4, leading to an overflow. For some unexplained reason, the Ariane 5 requirements and specifications did not include the Ariane 5 trajectory data. This omission was not simply an oversight: The accident report states that the partners jointly agreed not to include it. The report provides no reason for this very strange decision, which negatively impacted the ability to detect the problem during reviews and testing.

The philosophy of the Ariane 5 program, as stated in the accident report, that it was not wise to change software that had worked well on the Ariane 4 unless it was proven necessary to do so is well founded: Errors are often introduced when software is changed, particularly by those who did not originally write it, as occurred with SOHO. However, in this case, there is a tradeoff with safety that needs to be carefully evaluated. The best solution may not have been to leave the software as it was but to rewrite it from scratch. Such a decision depends on the type of change required, the specific design of the software, and the potential effect of the feature on system safety. The cost of the decision to reuse the Ariane 4 software without rewriting it (in terms of both money and program delays) was much greater than the cost of doing a proper analysis of its safety. The same is true for the Centaur software with respect to the losses resulting from leaving in the unneeded filter function. If the cost of the analysis of reused (or COTS) software or of changes to software is prohibitively expensive or beyond the state of the art, then redeveloping the software or completely rewriting it may be the appropriate decision.

A reasonable conclusion to be drawn is not that software cannot be reused, but that a safety analysis of its operation in the new system context is mandatory: Testing alone is not adequate to accomplish this goal. For complex designs, the safety analysis required stretches the limits of current technology. For such analysis to be technically and financially feasible, reused software must contain only the features necessary to perform critical functions, another reason to avoid unnecessary functions.

COTS software is often constructed with as many features as possible to make it commercially useful in a variety of systems. Thus, there is tension between using COTS vs being able to perform a safety analysis and have confidence in the safety of the system. This tension must be resolved in management decisions about specific project risk; ignoring the potential safety issues associated with COTS software can lead to accidents and potential losses that are greater than the additional cost would have been to design and build new components instead of buying them.

If software reuse and the use of COTS components are to result in acceptable risk, then system and software modeling and analysis techniques must be used to perform the necessary safety analyses. This process is not easy or cheap. Introducing computers does not

preclude the need for good engineering practices nor the need for difficult tradeoff decisions, and it almost always involves higher costs despite the common myth that introducing automation, particularly digital automation, will save money.

A similar argument applies to changing existing software. Changes to software appear to be easy. Whereas it is indeed easy to change software, it is very difficult to change it correctly, and the difficulty increases over time. The MCO software was changed to include a new thruster equation, but the 4.45 correction factor (the difference between the metric and imperial units), buried in the original code, was not noticed when the new vendor-supplied equation was used to update the software. Modifications to the SOHO command procedures were subjected to very little testing and review, perhaps, because they were considered to be minor.

The more changes that are made to software, the more the original design erodes and the more difficult it becomes to make changes without introducing errors. In addition, the assumptions and rationale behind the design decisions are commonly not documented and are easily violated when the software is changed. To prevent accidents, all changes to software must be thoroughly tested and, in addition, analyzed for their impact on safety. Such change analysis will not be feasible unless special steps are taken during development to document the information needed.

The environment in which the system and software are operating will also change over time, partially as a result of the introduction of the automation or system itself. Basic assumptions made in the original hazard analysis process must be recorded and periodically evaluated to assure they are not being violated in practice. Incident and accident analysis are also important, as well as performance monitoring and periodic operational process audits.

Violation of Basic Safety Engineering Practices in the Digital Components

Although system safety engineering textbooks and standards include principles for safe design, software engineers are almost never taught them. As a result, software often does not incorporate basic safe design principles, for example, separating and isolating critical functions, eliminating unnecessary functionality, designing error-reporting messages such that they cannot be confused with critical data (which occurred in the Ariane 5 loss), and reasonableness checking of inputs and internal states.

Consider the MPL loss as an example. The JPL report on the accident states that the software designers did not include any mechanisms to protect against transient sensor signals nor did they think they had to test for transient conditions. Runtime reasonableness and other types of checks should be part of the design criteria used for any real-time software. The Ariane 5 software appears not to have included simple and fairly standard range checks and overflow checks (without using exception handling) that would have been simple to implement in ADA (the programming language used), which was designed to provide the facilities required to perform real-time checking of software.

The SOHO, Titan/Centaur, and MCO accident reports are silent about whether the code included features to prevent the problem but were implemented incorrectly, or, alternatively, such features were omitted. From the outside, it appears that safety-related design features could easily have been used with minimal impact on hardware resources. System and software hazard analysis procedures can provide the information needed to implement the critical features with little cost in terms of programming or execution time. We need to start applying safe design principles to software just as we do for hardware.

The practices involved, however, may need to be very different than those used for hardware, as witnessed by the misunderstanding of the effectiveness of redundancy in the Ariane 5 accident. Traditional system engineering techniques to protect against accidents caused by component failure usually involve increasing the reliability (integrity) of the individual components, including redundancy to protect against component failure, and building error margins into components to protect against analysis or manufacturing flaws. None of these techniques are useful in protecting against system accidents or software-related accidents, where the problems result

from system design flaws: Increasing component integrity will have no effect on accidents where the software operated exactly as it was required and designed to do, but the requirements were incorrect. Every one of the accidents considered in this paper fit this description, except perhaps the Titan/Centaur loss, which resulted from a typographical error in a software load tape.

Misunderstanding the differences between the failure modes in hardware and those in software has led to attempts to increase reliability and safety through software redundancy, often in the form of independent groups writing multiple versions of the software with majority voting on the outputs. This approach is based on the assumption that such versions will fail in a statistically independent manner, but this assumption has been shown to be false in practice and by scientific experiments, for example, see Knight and Leveson.¹⁷ Common-cause (but usually different) logic errors tend to lead to incorrect results when the various versions attempt to handle the same unusual or difficult-to-handle inputs, as in the Ariane 5. In addition, such designs usually involve adding to system complexity, which can result in failure itself. A NASA study of an experimental aircraft with two versions of the control system found that all of the software problems occurring during flight testing resulted from errors in the redundancy management system and not in the control software itself, which worked perfectly.¹⁸ We need to devise protection features for software that reflect the failure modes of software and not those of hardware.

Inadequate Review Activities

General problems with the way quality and mission assurance was practiced were mentioned in the MCO, MPL, and Titan reports. Quality assurance (QA) often becomes an ineffective activity that is limited simply to checking boxes signifying the appropriate documents have been produced without verifying the quality of the contents. The Titan/Centaur accident report makes this point particularly strongly. According to the report, the LMA Quality Assurance Plan is a top-level document that focuses on verification of process completeness, not on how the processes are executed or implemented. Its basis was the original General Dynamics Quality Assurance Plan with updates to ensure compliance with ISO 9001. Following this plan, the LMA Software Quality Assurance staff only verified that the signoff report containing the software load constants had all of the proper signatures; they left the constant generation and verification process to the flight software and control dynamics engineers. Software QA involvement was limited to verification of software checksums and placing quality assurance stamps on the software products that were produced.

Review processes (outside of QA) are also described as flawed in the reports, but few details are provided to understand the problems. The Ariane 5 report states that reviews including all major partners in the Ariane 5 program took place, but no information is provided about what types of reviews were held or why they were unsuccessful in detecting the design flaws. It does say that the limitations of the inertial reference system software were not fully analyzed in reviews, and it was not realized that the test coverage was inadequate to expose such limitations. For example, the report states "there is no evidence that any trajectory data was used to analyze the behavior of the three unprotected variables."¹³ Given that the specifications did not include the Ariane 5 trajectory data, it is unlikely such reviews would have uncovered the problem. An assumption by the Ariane 5 developers that it was not possible to perform a complete system integration test made reviews, including simulation and analysis, even more important. The Ariane accident report recommends including external (to the project) participants when reviewing specifications, code, and justification documents and to make sure that these reviews consider the substance of arguments, rather than check that the verifications have been made. The latter seems to imply that reviewers and assurance personnel may have simply ticked off on checklists that activities had been performed without reviewing their quality and substance.

The MCO report recommends that NASA "conduct more rigorous, in-depth reviews of the contractor's and team's work,"¹⁴ which

it states were lacking on the MCO. The report also concludes that the operations team could have benefited from independent peer reviews to validate their navigation analysis technique and to provide independent oversight of the trajectory analysis. There is no mention of software QA activities or the software review process in the MCO report.

In the MPL descent engine control software reviews, apparently nobody attending was familiar with the potential for spurious Hall effect sensor signals. The independent MPL accident report noted that JPL and LMA failed to ensure adequate independent reviews.⁶

The SOHO accident report states that the changes to the ground-generated commands were subjected to very limited review. The flight-operations team placed high reliance on ESA and Matra Marconi Space representatives who were quite knowledgeable about the spacecraft design, but there were only two of them, and neither was versed in the computer language used to define the commands. A simulation was performed on the new compressed SOHO timelines, but the analysis of a problem detected during simulation was still going on as the new procedures were being used.

The Titan/Centaur report was the only one to mention the existence of an independent verification and validation review process by a group other than the developers. In that process, default values were used for the filter rate constants, and the actual constants used in flight were never validated. An engineer at the launch site detected the anomalous roll rate data at the time of tower roll back, but he was not able to identify the problem, partly because (according to the accident report) he had no documented requirements or procedures to review the data and no reference with which to compare.

Concerns about proprietary software is increasingly preventing external reviewers familiar with the spacecraft from a systems viewpoint from reviewing the software, for example, see the Cassini/Huygens communications link enquiry board report.¹⁹ The problem is not limited to aerospace but is occurring in all industries. Adequate review is not possible when the system engineers do not have access to the complete design of the spacecraft.

In general, software is difficult to review and the success of such an effort is greatly dependent on the quality of the specifications. However, identifying unsafe behavior, that is, the things that the software should *not* do and concentrating on that behavior for at least part of the review process, helps to focus the review and to ensure that critical issues are adequately considered.

Such unsafe (or mission-critical) behavior should be identified in the system engineering process before software development begins. The design rationale and design features used to prevent the unsafe behavior should also have been documented and can be the focus of such a review. This presupposes, of course, a system safety process to provide the information, which does not appear to have existed for the projects that were involved in the accidents studied.

As mentioned earlier, almost all software-related accidents have involved incomplete requirements specification and unhandled or mishandled system states or conditions. The two identified MPL software errors, for example, involved incomplete handling of software states and are both examples of very common specification flaws and logic omissions often involved in accidents. Such errors are most likely to be found if spacecraft and subsystem experts participate actively in the reviews.

Software hazard analysis and requirements analysis techniques and tools exist to assist in finding these types of incompleteness. To make such a review feasible, the requirements should include only the externally visible (blackbox) behavior; all implementation-specific information should be put into a separate software design specification (which can be subjected to a later software design review by a different set of reviewers). The only information relevant for a software requirements review is the software behavior that is visible outside the computer. Specifying only blackbox behavior (in engineering terminology, the transfer function across the digital component) allows the reviewers to concentrate on the information of importance to them without being overwhelmed by internal design information that has no impact on externally observable behavior.

The language used to specify the software requirements is also critical to the success of such a review. The best way to find errors in the software requirements is to include a wide range of disciplines and expertise in the review process. The reviewers must be able to read and understand the specifications without extensive training and, ideally, the notation should not differ significantly from standard engineering notations. While formal and executable specification languages have tremendous potential for enhancing our ability to understand the implications of complex software behavior and to provide correct and complete requirements, most of the languages created by computer scientists require too much reviewer training to be practical. A high priority on readability and learnability has not been placed on the development of such languages.

Ineffective System Safety Engineering

All of the accident reports studied are surprisingly silent about the safety programs and system safety activities involved. The MCO report cryptically says only that the “safety and mission assurance group was not properly engaged during the operations stage.”⁴ All of the reports, however, do cite inadequate hazard analysis or component criticality analysis. The MCO report notes that there was an absence of a process, such as FTA, for determining what could go wrong during a mission to assist in design, reviews, and contingency planning. The MPL report also notes inadequate hazard identification: During MPL development, FTA and other hazard analysis activities were used “inconsistently.”⁵

There did appear to be a criticality analysis performed on the Ariane and MCO projects, albeit flawed ones. The Ariane 5 report recommends reconsidering the definition of critical components, taking failures of software components into account (particularly single-point failures). Clearly the handling of an overflow in the alignment function by turning off the guidance and control computer was flawed. Unfortunately, not enough information is given about how the criticality analysis was performed to determine why it was unsuccessful.

Like other systemic factors, the Ariane 5 report is silent about the tradeoff analysis process. The report describes the decisions that resulted but not the process used to reach those decisions. For example, the fact that several software variables were unprotected is not the root of the problem: There is always a limit as to how much checking and protection is feasible. Spacecraft computers are limited by size and weight, which in turn limits processor power and memory, and timing constraints may preclude extensive checking software. The problem was that there was a lack of effective analysis to determine which variables should be protected. Important information that is missing includes how the analysis and trade studies were performed and what additional information or additional analysis techniques could have allowed better decisions to be made.

The MCO accident report also concludes there was “inadequate identification of mission-critical elements throughout the mission.”⁴ The criticality of specific elements of the ground software that impacted the navigation trajectory was not identified, for example.

The source of these problems may stem from the common practice of applying the same engineering analysis techniques developed for electromechanical systems, for example, FMEA and failure modes and effects criticality analysis (FMECA), to the new software-intensive systems. This approach will be limited because the contribution of software to accidents, as noted earlier, is different than that of purely mechanical or electronic components. In particular, software does not fail in the sense assumed by these techniques.

Often hazard analyses simply omit software, and when included, it is often treated superficially at best. The hazard analysis produced after the MPL loss is typical. The JPL report on the loss identifies the hazards for each phase of the entry, descent, and landing sequence, such as propellant line ruptures, excessive horizontal velocity causes lander to tip over at touchdown, and premature shutdown of the descent engines. For software, however, only one hazard, “flight software fails to execute properly,”⁵ is identified, and it is labeled as common to all phases. Such a software hazard is not very useful; it is equivalent to substituting the general statement—hardware fails to operate properly—for all of the hardware-related hazards

noted earlier. Identifying specific unsafe software behaviors as is done for hardware, such as software prematurely shuts down the descent engines, would be much more helpful during development in identifying ways to mitigate risk.

Providing the information needed to make safety-related engineering decisions is the major contribution of system safety techniques to engineering. It has been estimated that 70–90% of the safety-related decisions in an engineering project are made during the early concept development stage.²⁰ When hazard analyses are not performed, are done only after the fact (for example, as a part of quality or mission assurance of a completed design), or are performed but the information is never integrated into the system design environment, they can have no effect on these decisions, and the safety effort reduces to a cosmetic and perfunctory role.

Take SOHO, for example. A hazard analysis surely would have shown that the roll rate and the status of gyros A and B were critical, and this information could have guided the design of feedback channels to the operators about the current gyro state. A rigorous system safety process also would have triggered special safety analyses when changes were made to the SOHO operational procedures involving safety-critical components. In addition, a strong system safety program would have ensured that high priority was given to the analysis of previous emergency sun reacquisitions, that greater controls were placed on safety-critical operational procedures, and that safety-related open operational reports, such as the one reporting the difficulty the SOHO operators were having in reviewing telemetry data, did not stay open for four years and instead were tracked and resolved in a timely manner.

The description of the MCO problem by the developers¹⁶ says that the best chance to find and eliminate the problem existed at the early stages of development, but the team failed to recognize the importance of the small forces ground software and it was not given the same attention as the flight software.

The Titan/Centaur accident provides another example of what happens when such analysis is not done. The risk analysis, in that case, was not based on determining the steps critical to mission success, but instead considered only the problems that had occurred in previous launches. Software constant generation was considered to be low risk because there had been no previous problems with it. There is, however, a potentially enormous (perhaps unlimited) number of errors related to software, and considering only those mistakes made previously, while certainly prudent, is not adequate.

Not only is such a fly–fix–fly approach inadequate for complex systems in general, particularly when a single loss is unacceptable, but considering only the specific events and conditions occurring in past accidents is not going to be effective when new technology is introduced into a system. Computers are, in fact, introduced to make radical changes in functionality and design possible. In addition, software is often used precisely because it is possible to make changes for each mission and throughout operations: The system being flown today is often not the same one that existed yesterday. Proper hazard analysis that examines all of the ways the system components (including software) or their interaction can contribute to accidents needs to be performed and used in the original development and when making changes during operations.

At the same time, system-safety techniques, like other engineering techniques, need to be expanded to include software and the complex cognitive decision making and new roles played by human operators.¹² Existing approaches need to be applied, and new and better ones developed. If system hazard analysis is performed before software implementation (not just before testing, as is recommended in the MPL report), requirements can be analyzed for hazardous states and protection against potentially hazardous behavior designed into the software logic from the beginning.

The MCO accident report recommended that the NASA Mars Program institute a classic system safety engineering program: 1) continually performing the system hazard analyses necessary to identify explicitly mission risks and communicating these risks to all segments of the project team and institutional management; 2) vigorously working to make tradeoff decisions that mitigate the risks to maximize the likelihood of mission success; and 3) regularly com-

municating the progress of the risk mitigation plans and tradeoffs to project, program, and institutional management.⁴ The other spacecraft accident reports, in contrast, recommended applying classic reliability engineering approaches that are unlikely to be effective for system accidents or software-related causal factors.

One of the benefits of using system-safety engineering processes is simply that someone becomes responsible for ensuring that particular hazardous behaviors are eliminated if possible or their likelihood reduced and their effects mitigated in the design. Almost all attention during development is focused on what the system and software are supposed to do. System safety and software safety engineers are responsible for ensuring that adequate attention is also paid to what the system and software are *not* supposed to do and verifying that hazardous behavior will not occur. It is this unique focus that has made the difference in systems where safety engineering successfully identified problems that were not found by the other engineering processes.

Flaws in the Test and Simulation Environments

It is always dangerous to conclude that poor testing was the cause of an accident. After the fact, it is always easy to find a test case that would have uncovered a known error. It is usually difficult, however, to prove that the particular test case would have been selected beforehand, even if testing procedures were changed. By definition, the cause of an accident can always be stated as a failure to test for the condition that was determined, after the accident, to have led to the loss. However, all of the accidents reports studied mention omissions that reflect poor decisions related to testing, particularly with respect to the accuracy of the simulated operational environment.

A general principle in testing aerospace systems is to fly what you test and test what you fly. This principle was violated in all of the spacecraft accidents, especially with respect to software. The software test and simulation processes must reflect the environment accurately. Although implementing this principle is often difficult or even impossible for spacecraft, no reasonable explanation was presented in the reports for some of the omissions and flaws in the testing for these systems. One example was not testing the Titan/Centaur software with the actual load tape prior to launch. Another example was the use of Ariane 4 trajectory data in the simulations of the Ariane 5 software even though the Ariane 5 trajectory was known to be different. The Ariane inertial reference system specification, in addition, did not include operational restrictions arising from the implementation and needing to be considered during test.

Testing of SOHO operational procedures was primarily performed using a simulator, but the simulator had not been maintained with all of the onboard software changes that had been implemented on the spacecraft, essentially making such testing useless. The MPL touchdown sensing software was not tested with the lander in the flight configuration. Various other types of test cases were also not included. The JPL MPL report concludes,

The laboratory needs to reinforce the system-level test principle of “test as you fly, and fly as you test.” Departures from this principle must be carefully assessed, and if they are determined to be necessary, alternative measures, such as independent validation, should be incorporated. Such items must be reflected in the project risk management plan, communicated to senior management for concurrence, and reported at reviews.⁵

For both the Ariane 5 and Mars '98 projects, a conclusion was reached during development that the components implicated in the accidents could not be tested and simulation was substituted. After the fact, it was determined that such testing was indeed possible and would have had the ability to detect the design flaws. For example, the Ariane 5 report says,

It would have been technically feasible to include almost the entire inertial reference system in the overall system simulations which were performed. For a number of reasons, it was decided to use the simulated output of the inertial reference system, not the system itself or its detailed simulation. Had the system been included, the [design error] could have been detected.³

The same occurred with the Titan/Centaur accident, where default and simulated values were used in system testing, although the real roll rate filter constants could have been used. Like Ariane, the Titan/Centaur engineers incorrectly thought the rigid-body simulation of the vehicle would not exercise the filters sufficiently. In addition, the corporate consolidation/evolution process had led to a loss of knowledge that the test bed could use the actual roll rate filter constants. The current engineers thought only default constants could be used in the test bed. Even the tests performed on the Titan/Centaur right before launch (because anomalies had been detected) used a default set of constants to filter the measured roll rate instead of the actual constants and, thus, were unsuccessful in detecting the error. After wiring errors were discovered in the MPL testing process, for undisclosed reasons the tests necessary to detect the software flaw were not rerun.

Not all problems in testing can be traced to the simulation environment, of course. There have been cases of spacecraft losses involving inadequate, inappropriate, and ill-suited testing. A basic problem is one of piecemeal testing and not testing at the system level for system-level effects and emergent behavior. The rush to get the ground software operational after a problem was discovered postlaunch in MCO resulted in the testing program being abbreviated. For example, the testing of the ground software changes did not test the correctness of the file formats but simply made sure that the file could be moved across on the file server.¹⁵ At a 10 November 1999 press conference, Alfred Stephenson, the chief accident investigator, admitted, "Had we done end-to-end testing, we believe this error would have been caught." However, the rushed and inadequate preparations left no time to do it right. The problem lies not only in testing, but in relying on software that had not been adequately tested without additional manual or other checks to gain confidence. The same lack of system test also contributed to the WIRE (Wide-Field Infrared Explorer Mission) spacecraft, where a contributing cause cited in the accident report was that no system level end-to-end test with live pyrotechnic devices in the as-flown configuration had been done.²¹

Another common problem in software testing is inadequate emphasis on off-nominal and stress testing. Most software-related accidents have involved situations that were not considered during development or were assumed to be impossible and not handled by the software. For example, the error in the Ariane exception handling software might have been found if off-nominal testing had been performed on the Ariane despite the fact that an incorrect assumption about the Ariane trajectory had been made in the specifications. The recent emphasis on use cases in software testing exacerbates the lack of off-nominal and stress testing. Inadequate and unsafe software logic is much more likely to be found when inputs violate the specifications than when they simply test the expected cases.

Better system testing practices are needed for components containing software (almost everything these days), more accurate simulated environments need to be used in software testing, and the assumptions used in testing and simulations need to be carefully checked.

Inadequate Human Factors Design for Software

Understanding the impact of software design on human error is still in the early stages. Accidents, surveys, and simulator studies, however, have emphasized the problems pilots are having in understanding digital automation and have shown that pilots are surprisingly uninformed about how the automation works.^{22,23}

Human factors experts have written extensively on the potential risks introduced by the automation capabilities of glass cockpit aircraft. Among the problems identified are causing mode confusion and situational awareness difficulties, providing inadequate feedback to support effective monitoring and decision making, creating overreliance on automation, shifting workload by increasing it during periods of already high workload and decreasing it during periods of already low workload, being clumsy or difficult to use, being opaque or difficult to understand, and requiring excessive experience to gain proficiency in its use. Not all of this information seems to have affected engineering practice.

As more sophisticated automation has been introduced into spacecraft control, and control of safety-critical functions is increasingly shared between humans and computers, the same problems found in high-technology aircraft are appearing. Ground controllers rely on indirect information about the object they are controlling. Their effectiveness depends on the information they are given and the format in which it is presented. Many of the problems found in human automation interaction lie in the human not getting appropriate feedback to monitor the automation and to make informed decisions.

For the three accidents studied that involved operators (MCO, Titan, and SOHO), all of the accident reports mentioned human factors issues. Neither the MCO nor the Titan/Centaur mission operations personnel understood the system or software well enough to interpret the data they saw as indicating there was a problem in time to prevent the loss.^{4,7} As noted earlier, the MCO navigation operations team did not know that the spacecraft sent down tracking data that could have been compared with the ground data to detect the error and to validate the changes in the software. Complexity in the automation combined with poor documentation and training procedures are contributing to these problems.

Problems abound in the design of the interfaces between humans and automation. As noted earlier, the SOHO operations personnel had filed a report about the difficulty they were having interpreting the telemetry data using the interface they were given. For diagnostic purposes, SOHO was designed to store within its onboard computer the last three telemetry frames that precede entry into the safe mode, but the ground operators were unable to display this information in a usable format. If the information had been available to them, the operators could have determined the status of the gyros and avoided the loss. No explanation for why the problem was not fixed is provided in the accident report. Several other places in the SOHO accident report also hint at the controllers not having the information they needed about the state of the gyros and the spacecraft in general to make appropriate decisions. The diagnosis of the wrong gyro as the bad one and its subsequent deactivation raises many important questions about the information provided to the operators that are not answered in the accident report.

The SOHO report also says that the software had been modified incorrectly by the operations team due to lack of system knowledge of the person making the changes. The flight operations team was not sufficiently versed regarding details of the spacecraft design and its idiosyncrasies, and this problem became worse as turnover occurred and there was no time for new personnel to take advantages of training sessions.

Complexity in the automation combined with inadequate documentation and training procedures are contributing to the problems we are seeing. Sometimes the incorrect assumption is made that introducing computers lessens the need for in-depth knowledge by operational personnel, but the opposite is true. Some of the spacecraft accidents where operators were implicated involved a failure to transfer skill and knowledge from those who operated spacecraft in prior missions, and they were, therefore, unable to detect the software deficiencies in time to save the mission. Either the design of the automation we are building needs to be improved from a human factors engineering viewpoint or new training methods are needed for those who must deal with the clumsy automation and confusing, error-prone interfaces we are designing.

Conclusions

This paper has described common factors noted in the accident reports for five recent software-related spacecraft losses: Ariane 501, MCO, MPL, Titan IV B-32, and SOHO communications. Some suggestions for ways to avoid these problems were provided throughout the paper, although a comprehensive discussion of how to solve these very complex problems is beyond its scope. The first step in preventing accidents, however, is learning from those that have already occurred.

Complacency and misunderstanding software and its risks were at the root of all these accidents. Software presents tremendous potential for increasing our engineering capabilities. At the same time,

it introduces new causal factors for accidents and requires changes in the techniques used to prevent the old ones. We need to apply the same good engineering practices to software development that we apply to other engineering technologies, while at the same time understanding the differences and making the appropriate changes to handle them. There is no magic in software: It requires hard work and is difficult to do well, but the result is worth the effort.

Acknowledgments

This research was partially supported by a grant from the NASA Ames Research Center Design for Safety program and by the NASA IV&V Center Software Initiative program.


References

- ¹Leveson, N. G., "Evaluating Accident Models Using Recent Aerospace Accidents: Part I. Event-Based Models," Massachusetts Inst. of Technology, Engineering Systems Div. Technical Rept., ESD-WP-2002-06, Cambridge, MA, June 2001.
- ²Weiss, K., Leveson, N. G., Lundqvist, K., Farid, N., and Stringfellow, M., "An Analysis of Causation in Aerospace Accidents," *Proceedings of the Digital Aviation Systems Conference*, 2001, pp. 137-147.
- ³Lions, J. L., (Chairman) "Ariane 501 Failure: Report by the Inquiry Board," ESA, July 1996.
- ⁴Stephenson, A., "Mars Climate Orbiter: Mishap Investigation Board Report," NASA Nov. 1999.
- ⁵JPL Special Review Board, "Report on the Loss of the Mars Polar Lander and Deep Space 2 Missions," Jet Propulsion Lab., California Inst. of Technology, Pasadena, CA, March 2000.
- ⁶Young, T. (Chairman), "Mars Program Independent Assessment Team Report," NASA, March 2000.
- ⁷Pavlovich, J. G., "Formal Report of Investigation of the 30 April 1999 Titan IV B/Centaur TC-14/Milstar-3 (B-32) Space Launch Mishap," U.S. Air Force, 1999.
- ⁸NASA/ESA Investigation Board, "SOHO Mission Interruption," NASA, Aug. 1998.
- ⁹Hopkins, A., Managing Major Hazards: *The Lessons of the Moira Mine Disaster*, Allen and Unwin, Sidney, NSW, Australia, 1999.
- ¹⁰Perrow, C., *Normal Accidents: Living with High-Risk Technology*, Basic Books, New York, 1984.
- ¹¹Rogers, W. P., "Report of the Presidential Commission of the Space Shuttle Challenger Accident," U.S. Government Accounting Office, Washington, DC, 1986.
- ¹²Leveson, N. G., *Safeware: System Safety and Computers*, Addison Wesley, Boston, 1985.
- ¹³Lutz, R. R., "Analyzing Software Requirements Errors in Safety-Critical, Embedded Systems," *Proceedings of the International Conference on Software Requirements*, IEEE Publications, Piscataway, NJ, 1992, pp. 53-65.
- ¹⁴Oberg, J., "Why the Mars Probe Went Off Course," *IEEE Spectrum Magazine*, Vol. 36, No. 12, 1999.
- ¹⁵Euler, E. E., Jolly, S. D., and Curtis, H. H., "The Failures of the Mars Climate Orbiter and Mars Polar Lander: A Perspective from the People Involved," American Astronautical Society, Paper AAS 01-074, 2001.
- ¹⁶Hoare, A. A., "The Emperor's Old Clothes," *Communications of the ACM*, Vol. 24, No. 2, 1981, pp. 75-83.
- ¹⁷Knight, J. C., and Leveson, N. G., "An Experimental Evaluation of the Assumption of Independence in Multi-Version Programming," *IEEE Transactions on Software Engineering*, Vol. SE-12, No. 1, 1986, pp. 96-109.
- ¹⁸Mackall, D. A., "Development and Flight Test Experiences with a Flight-Critical Digital Control System," NASA TP 2857, Nov. 1988.
- ¹⁹Link, D. C. R. (Chairman), "Report of the Huygens Communications System Inquiry Board," NASA, Dec. 2000.
- ²⁰Frola, F. R., and Miller, C. O., "System Safety in Aircraft Acquisition," Logistics Management Inst., Washington, DC, Jan. 1984.
- ²¹Branscome, D. R. (Chairman), "WIRE Mishap Investigation Board Report," NASA, June 1999.
- ²²Sarter, N. D., and Woods, D., "How in the World Did I Ever Get into That Mode?: Mode Error and Awareness in Supervisory Control," *Human Factors*, Vol. 37, No. 1, 1995, pp. 5-19.
- ²³Bureau of Air Safety Investigation, "Advanced Technology Aircraft Safety Survey Report," Dept. of Transport and Regional Development, Canberra, Australia, June 1996.

D. L. Edwards
Associate Editor

Economic Principles Applied to Space Industry Decisions

Joel S. Greenberg, Princeton Synergetics, Inc.



This is not an economics book. It is a book about the application of economic principles and concepts in decision making related to space activities. The book is primarily tutorial and elaborates upon concepts and methodology and their applications. Emphasis is placed upon applications with typical results of performed analyses presented to demonstrate concepts and methods.

The use of mathematical and simulation models serves as the underpinning for much of the presented materials. The specific models considered have been selected to demonstrate the role that a structured thought process can play in the decision process. Since most decisions relating to technology development, product design, capital expenditures, and investments involve uncertainty and risk, a number of the selected models, developed methodologies, and presented examples explicitly and quantitatively consider uncertainty and risk.

The objective of this book is to put economic analysis into perspective with respect to real-world decision making in the space industry. It will expand the perspective of the reader with respect to the type of tools and analyses that might be brought to bear on complex business and government problems.

Contents:

Introduction • Investment Decisions • RLV Economics • Space Operations • Licensing and Regulatory Issues • Beyond Space: Energy and Gaming • Appendix: Estimating the Likelihood of Investment

Progress in Astronautics and Aeronautics Series

2003, 480 pages, Hardback

ISBN: 1-56347-607-X

List Price: \$100.95

AIAA Member Price: \$69.95

Publications Customer Service, P.O. Box 960

Herndon, VA 20172-0960

Phone: 800/682-2422; 703/661-1595

Fax: 703/661-1501

E-mail: warehouse@aiaa.org • Web: www.aiaa.org



American Institute of
Aeronautics and Astronautics