

100 Questions for Technical Review

30 September 2005

Prepared by

P. G. CHENG
Risk Assessment and Management Subdivision
Systems Engineering Division

Prepared for

SPACE AND MISSILE SYSTEMS CENTER
AIR FORCE SPACE COMMAND
2430 E. El Segundo Blvd.
El Segundo, CA 90245

Contract No. FA8802-04-C-0001

Engineering and Technology Group

DISTRIBUTION STATEMENT: Distribution is limited to US Government agencies and their contractors only; Administrative or Operational Use, 30 September 2005. Other request for this document shall be referred to SMC/AX.

DESTRUCTION NOTICE: For classified documents, follow the procedures in DOD 5220.22-M, National Industrial Security Program Operating Manual (NISPOM), Paragraph 5, Section 7. For unclassified, limited documents, destroy by any method that will prevent disclosure of contents or reconstruction of the document.

100 QUESTIONS FOR TECHNICAL REVIEW

Prepared by

P. G. CHENG
Risk Assessment and Management Subdivision
Systems Engineering Division

30 September 2005

Engineering and Technology Group
THE AEROSPACE CORPORATION
El Segundo, CA 90245-2691

Prepared for

SPACE AND MISSILE SYSTEMS CENTER
AIR FORCE SPACE COMMAND
2430 E. El Segundo Blvd.
El Segundo, CA 90245

Contract No. FA8802-04-C-0001

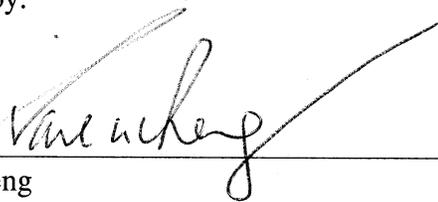
DISTRIBUTION STATEMENT: Distribution is limited to US Government agencies and their contractors only; Administrative or Operational Use, 30 September 2005. Other request for this document shall be referred to SMC/AX.

DESTRUCTION NOTICE: For classified documents, follow the procedures in DOD 5220.22-M, National Industrial Security Program Operating Manual (NISPOM), Paragraph 5, Section 7. For unclassified, limited documents, destroy by any method that will prevent disclosure of contents or reconstruction of the document.

THIS PAGE INTENTIONALLY LEFT BLANK

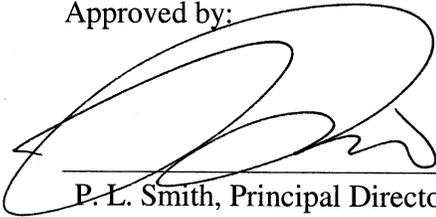
100 QUESTIONS FOR TECHNICAL REVIEW

Prepared by:



P. G. Cheng
Risk Assessment and Management Subdivision
Systems Engineering Division

Approved by:



P. L. Smith, Principal Director
Risk Assessment and Management Subdivision
Systems Engineering Division

THIS PAGE INTENTIONALLY LEFT BLANK

Abstract

Failure reports routinely trace the underlying cause to “engineering mistakes” and lament “inadequate reviewing.” Aerospace personnel participate in a variety of program reviews such as PDRs, CDRs, and MRRs. How can reviewers, in a few hours, find a mistake that has escaped years of design and quality checks by the contractor and program office?

Over the last several years we have published 100 “Space Systems Engineering Lessons Learned,” each describing some past incidents and the errors that contributed to them. The following 100 questions—each hyperlinked to the relevant lessons—will help reviewers check if proper engineering practices have been followed to prevent, catch, or mitigate similar errors. For example, Question 8-1 asks “Do the tests *independently* confirm development results?” If a reviewer had asked this question about Hubble, where a flawed optical instrument was used both to guide the mirror polishing and to verify the finished product, the infamous spherical aberration might have been avoided.

These questions are open-ended and not a comprehensive checklist (which would be impossible to create), and reviewers must use their expertise to tailor the questions for a particular situation. Still, if the response is “You know, we never thought about that, we better check it,” the reviewers have earned their pay!

Acknowledgement

Sincere thanks goes to the more than 40 Aerospace engineers who provided content for these lessons. Particular appreciation goes to Jon Binkley, Peter Carian, Dana Speece, and Ron Williamson for their prolific contributions as well as their assistance in creating the “100 Questions.”

THIS PAGE INTENTIONALLY LEFT BLANK

Contents

| | |
|---|----|
| Section 1: Requirements | 1 |
| Section 2: Heritage and “Qualification by Similarity” | 3 |
| Section 3: Analysis | 5 |
| Section 4: Failure Modes and Fault Management..... | 9 |
| Section 5: Embedded Software and Database..... | 11 |
| Section 6: Interfaces..... | 13 |
| Section 7: Parts, Materials, and Manufacturing Process..... | 15 |
| Section 8: Testing and Evaluation | 17 |
| Appendix A: Space Systems Engineering Lessons 1-100 | 21 |

THIS PAGE INTENTIONALLY LEFT BLANK

Section 1: Requirements

- 1-1 Are units and tolerances specified?
- Quantifying requirements reduces mistakes and surfaces manufacturing and test issues.
 - Use TBDs to highlight the need for further clarification, but clear them off in a timely manner.
 - Watch out for mistakes when two interfacing organizations use different units (English versus metric, for example), CAD/CAM protocols, or engineering practices.
 - Lessons: [73](#) and [76](#).
- 1-2 Is the specification's wording unambiguous?
- Avoid incomplete lists (typically ending with "etc."), vague words such as "to the best possible," passive voice, such as "the counter is set" (by whom?), and negative statements.
 - See <http://www.ntsc.navy.mil/Resources/Library/Acgguide/spec.htm>, a comprehensive "Guide to Specification Writing for U.S. Government Engineers."
 - Lessons: [4](#) and [12](#).
- 1-3 Should any statement be split up?
- Lumped requirements are difficult to trace. Some may be overlooked.
 - Lesson: [12](#).
- 1-4 How will it be demonstrated that each requirement is met?
- Each requirement should be traceable to a compliance matrix.
 - If a requirement is implemented by software, it must be linked to test cases.
 - Lesson: [19](#).
- 1-5 Does each requirement trace upward?
- Unnecessary or overtight requirements drive up costs.
 - Rationale for each derived requirement should be documented.
 - If a lower-level implementation affects a higher level, make sure other sub-systems will not be surprised.
 - Lesson: [85](#).
- 1-6 How are configuration changes tracked?
- Make sure requirement or design changes are coordinated, and reincorporate all redlinings and ad-hoc changes in the specifications.
 - Lessons: [97](#), [70](#), [53](#), and [64](#).

THIS PAGE INTENTIONALLY LEFT BLANK

Section 2: Heritage and “Qualification by Similarity”

- 2-1 Have all “heritage equipment” test and flight anomalies been resolved?
- The implication of each anomaly must be carefully addressed.
 - Lessons: [41](#) and [65](#).
- 2-2 Have catastrophic failures that involved similar technologies been reviewed?
- Lesson: [87](#).
- 2-3 Did the original analyst review the model’s application?
- Reusing a model without fully understanding underlying assumptions can be risky.
 - Lesson: [99](#).
- 2-4 Do previous analyses still apply?
- Changes in configuration or flight environment may invalidate the original analysis.
 - Parameters worth checking include temperature, power, electrical and mechanical stress, and flight duration.
 - Lessons: [95](#), [83](#), and [47](#).
- 2-5 Is the heritage design well understood?
- Lesson: [50](#).
- 2-6 Should an old unit recommissioned for flight be retrofitted?
- Design upgrades made while an old unit sat on the shelf should be considered.
 - Lesson: [57](#).
- 2-7 Have replacement materials and parts been fully qualified?
- It is not sufficient for the replacements to merely meet lot acceptance specifications.
 - Lesson: [14](#).
- 2-8 Should fault management circuits be redesigned?
- When a heritage unit is scaled up, key parameters such as start-up current and rise time may change.
 - Lesson: [84](#).

THIS PAGE INTENTIONALLY LEFT BLANK

Section 3: Analysis

- 3-1 Have all critical analyses been placed under configuration control?
- Design changes may invalidate the original analysis.
 - Lessons: [26](#) and [83](#).
- 3-2 Have designs been compared to similar, proven, equipment?
- Novel design approaches may entail risks.
 - Make sure subcontractors concur with the way their product is used.
 - Lessons: [82](#) and [99](#).
- 3-3 Has the analyst inspected the actual hardware?
- Sometimes the hardware is not what the analyst imagined.
 - Lessons: [81](#) and [26](#).
- 3-4 Can the manufacturing process meet design requirements?
- Make sure the manufacturing engineer reviewed drawings early on.
 - Use prototype and engineering models to discover problems early—issues found in late tests can be very expensive.
 - Lessons: [37](#) and [55](#).
- 3-5 Is the design tolerant of dimensional changes?
- Example: thermal mismatch and creep can cause dimension change, interference, and shorting.
 - Lessons: [52](#) and [47](#).
- 3-6 Is there any analysis that cannot be verified on account of “contractor proprietary data” or classified information?
- All proprietary processes should be thoroughly reviewed.
 - There are always ways to work with the classified issue.
 - Lesson: [23](#).
- 3-7 Was component qualification based on sufficient engineering data?
- That a few items worked is not sufficient—statistical data may be required to show margin of safety.
 - Instrumentation data may provide information to substantiate or disprove the analysis, which is more essential.
 - Lesson: [82](#).
- 3-8 Was the analysis unbiased?
- Do not throw out data points that do not fit a theory or could not be readily understood.
 - Lesson: [59](#).

- 3-9 Was the space environment fully accounted for?
- Examples: damping, radiation, charging, arcing, heat dissipation, refractive index, and microgravity.
 - Ground thermal insulation blanket to prevent space charge buildups.
 - Lessons: [41](#), [42](#), [10](#), and [75](#).
- 3-10 Has the electrical schematic been independently checked, from end to end?
- Mistakes sometimes occur between drawings.
 - Lesson: [68](#).
- 3-11 Can the harness be misconnected?
- Wiring and connectors should be designed to preclude mismating.
 - Lesson: [63](#).
- 3-12 Are mechanical load margins adequate?
- Immature state-of-the-art in the analysis of vibration, separation shocks, thruster imbalance, and dynamic load caused many failures.
 - Lessons: [11](#), [81](#), [33](#), [27](#), and [69](#).
- 3-13 Will excessive thermal or electrical loads damage hardware?
- Example: Relays can be welded shut by in-rush current and cause premature deployment.
 - Electronic output circuits should be self-limiting for worst-case failure currents.
 - Lessons: [19](#), [71](#), [87](#), [99](#), and [44](#).
- 3-14 Can unexpected time-dependent circuit behavior be accommodated?
- Start-up and turn-off transients can introduce problems such as EMI.
 - Lengths of transients, such as pyro firing pulses, should be bounded.
 - Lessons: [82](#) and [77](#).
- 3-15 Has a thorough safety analysis been conducted on each pyro event?
- Pyros impart a large and irreversible shock to the system and are involved in many mission failures.
 - Pyro design should be checked against available guidelines.
 - The effect of pyro shock on adjacent structures and circuits must be thoroughly validated.
 - If explosive bolt cutters are used, all ejected debris should be contained.
 - Lessons: [98](#), [89](#), [82](#), [77](#), [68](#), and [7](#).
- 3-16 Are deployables readily tested both in 0 g and in 1 g?
- Designs that work in 0 g but not in 1 g are difficult to verify.
 - Lessons: [20](#) and [42](#).

- 3-17 Will a malfunctioning valve cause a failure?
- Contamination in valves has led to numerous failures.
 - Make sure mistakes, such as software errors, in the valve controller will not disable the vehicle.
 - Lessons: [83](#), [65](#), [57](#), and [54](#).
- 3-18 Do moving units possess sufficient torque margins and clearance?
- Soft items such as cable and multi-layer insulation can move unexpectedly in the launch or space environment and cause interference.
 - Consider stiction in torque analysis.
 - Avoid structures that can snag soft items, and route wires to avoid pinching or snagging by a deployed structure.
 - Lessons: [78](#), [42](#), [70](#), and [9](#).
- 3-19 Will the solar array flutter?
- Conduct modal frequency analysis to avoid excessive vibration of the solar arrays upon entering or exiting the Earth's shadow.
 - Lesson: [13](#).
- 3-20 Has a worst-case analysis of EMI or crosstalk been conducted?
- EMI analysis should take into consideration the possibility of multiple boxes working in unison causing superimposition (such as in TDMA payloads).
 - Lesson: [86](#).
- 3-21 Are the power distribution and grounding schemes, including over-voltage and under-voltage limits, safe?
- Sneak paths should be eliminated.
 - All units should be protected from over- or under-voltage conditions from the power bus.
 - Current sensitive circuits should have over-current protection.
 - Power conversion and distribution units should be protected against over-current.
 - Components such as step motors and pyro circuits that experience sudden current changes should be isolated from all other current-carrying circuits.
 - Lesson: [98](#).
- 3-22 Are all known quirks of field programmable gate arrays (FPGAs) accounted for?
- FPGAs, used as anti-fuses, have demanding electrical design rules and software interface.
 - A NASA website <http://www.klabs.org/> describes common design mistakes.
 - Lessons: [77](#) and [100](#).

THIS PAGE INTENTIONALLY LEFT BLANK

Section 4: Failure Modes and Fault Management

- 4-1 Has the fault protection logic been independently verified?
- The fault management system (particularly the software) can be a source of single-point failures.
 - Example: Faulty sensor data may create a phantom problem and spoof the fault management system into taking precipitous actions such as resets.
 - Fault detection setting and responses should pass sanity checks. Endless resets, for example, are dangerous.
 - Lessons: [18](#), [36](#), and [43](#).
- 4-2 Will the satellite autonomous management system and the ground controller be provided with correct information?
- Inaccurate situation awareness can lead to wrong disposition.
 - Ensure subsystems report true status to the autonomy functions.
 - Lessons: [29](#) and [44](#).
- 4-3 Does the fault management design consider all operational possibilities?
- Example: solar array mispointing, engine abort, or eclipse transient.
 - Lessons: [36](#) and [38](#).
- 4-4 Is telemetry sufficient for all critical events?
- Knowledge for events such as separation can enable recovery.
 - Capture indelible records of system parameters in past events with, for example, strip chart records.
 - Lessons: [67](#) and [36](#).
- 4-5 Are multiple safeguards available during early operation?
- Problems frequently occur during early orbit operation.
 - Ground coverage must be ample.
 - The satellite should autonomously operate in case ground commands do not arrive promptly (due to erroneous position estimation, for example).
 - Lessons: [39](#) and [53](#).
- 4-6 Can a glitch trigger a crash?
- Systems should be designed to revert to “last known good state.”
 - Example: A momentary wiring short in the bus may reset all relays, with fatal consequences.
 - Lesson: [91](#).

- 4-7 How will the satellite handle battery undercharging?
- The satellite should be able to automatically shed non-essential loads under low voltage.
 - Even a partially deployed solar array should provide enough current to sustain the system.
 - The power regulator should be energized from the solar array, instead of being solely dependent on the battery for housekeeping.
 - Lessons: [53](#), [47](#), [67](#), and [30](#).
- 4-8 Can the fault management system itself survive major anomalies?
- Example: If a computer freezes, will fault correction software execute?
 - Lesson: [35](#).
- 4-9 Are contingency plans for on-orbit anomalies adequate?
- Contingency recovery plans, such as to correct the spacecraft's attitude, should be based on realistic timeline constraints and rehearsed.
 - Lesson: [60](#).
- 4-10 Can a problem in a primary unit cause the same failure in its backup?
- If the primary and redundant units share the same current feed, software, or processor, one flaw in the primary component can cause the backup to fail in the same way.
 - Lessons: [18](#) and [19](#).
- 4-11 Can serial safety devices (inhibits) fail simultaneously?
- Deployment mechanisms such as squibs or wax heater actuators should have separately driven safety devices lest one single error defeat both. Failure analysis of safety devices is particularly tricky.
 - Lessons: [77](#) and [100](#).
- 4-12 Can a device damage its neighbors?
- Example: EMI or shock from squibs and step motors.
 - Lessons: [89](#) and [100](#).
- 4-13 Does the design allow in-flight upgrades?
- On-orbit reprogrammability provides flexibility.
 - Lessons: [50](#), [30](#), and [23](#).
- 4-14 Can the on-board computer be safely reset?
- Executable software should be easily loadable even if the computer locks up.
 - Consider providing a backdoor receiver with default mode to overcome a computer lockup.
 - Lessons: [79](#).

Section 5: Embedded Software and Database

- 5-1 Will unexpected inputs cause the software to freeze or loop endlessly?
- Lessons: skipped sensor input data, data outside the expected range, or data that does not compute.
 - Software should ignore spurious inputs through filtering or limit checking.
 - Consider deliberately ignoring faults if there is no possible recovery.
 - Avoid permitting software to reset in response to errors. Consider error messages in telemetry instead.
 - All “IF” branches should provide an “ELSE” for the unexpected input.
 - Lesson: [18](#).
- 5-2 What happens if the software hangs up?
- Fault management logic must provide a way out.
 - Fault analysis must not assume perfect software.
 - Consider independent fault protection, such as hardware watchdog timers.
 - Lessons: [35](#), [36](#), and [18](#).
- 5-3 Can the computer get stuck during boot up?
- Do not let an error or malfunction prevent the computer from booting up.
 - Make sure watchdog functions will not cycle between start and reset.
 - Lesson: [79](#).
- 5-4 Will it be possible to remotely diagnose computer problems?
- Consider keeping debug utilities.
 - Avoid using dynamic memory allocation, which may complicate troubleshooting.
 - Lesson: [94](#).
- 5-5 Is every critical software under configuration control?
- All software, not just the software that is uploaded, that affects satellite behavior is critical and requires careful verification.
 - Changes should be tracked back to requirements and specifications.
 - Lesson: [73](#).
- 5-6 How are database parameters verified?
- Treat database loading as carefully as coding.
 - Ensure data entry procedures are free from human errors, and conduct independent verification of database integrity.
 - Lessons: [3](#) and [43](#).

- 5-7 Are command scripts formally controlled?
- A bad command sequence can be fatal.
 - Lesson: [29](#).
- 5-8 Will testing exercise all logic branches?
- Software should be tested over several days of equivalent mission time to find problem such as timing errors, overrunning counters, or unintended re-entries to “one-time” events.
 - Use automated tools to verify code paths.
 - All branches should be exercised and all parameters should be verified.
 - Lesson: [94](#).
- 5-9 How are reused or modified codes verified?
- Software changes should be controlled and retested as rigorously as hardware modification.
 - Issues understood by the original designer may be overlooked during modification.
 - Reused software should be compatible with the new application environment.
 - If a function in the reused code is not used, make sure it is completely disabled and has no output to downstream code.
 - Consider stripping off “dead” code.
 - Lessons: [18](#), [25](#), [48](#), and [79](#).
- 5-10 Has the flight software been tested with high-fidelity hardware in the loop, in the flight configuration?
- The ground test bed should be configured the same as the flight computer. At a minimum, the test bed should have the flight processor, flight memories, flight software, flight cables, flight power management equipment, and high-fidelity engineering model hardware.
 - Test beds should include test points for measuring all signal and control voltages and currents.
 - Lessons: [19](#), [36](#), and [53](#).
- 5-11 Are memory and throughput margins adequate?
- Check against unexpected data rates or excitation.
 - Ensure the system can handle a runaway sensor.
 - Lessons: [35](#) and [94](#).
- 5-12 Have all major events been scrubbed for out-of-sequence inputs?
- A signal arriving earlier or later than expected can trigger unintended timing conflicts.
 - Missing data may leave the system in an unknown state.
 - Lessons: [12](#) and [25](#).

Section 6: Interfaces

- 6-1 Have interface authority, end-to-end responsibility, and conflict resolution authorities been assigned?
- Examples: payload-to-bus, satellite-to-GSE, satellite-to-launcher, and satellite-to-ground interfaces.
 - Create a constructive mechanism to proactively distribute requirements, flow down error budgets, and assign footprints or connectors.
 - Lessons: [37](#) and [81](#).
- 6-2 Have potential incompatibilities between interfaces been analyzed early on?
- Independent analysis is often needed to overcome organizational barriers.
 - Lessons: coupled loads, nutational instability, and EMI.
 - Lessons: [2](#), [11](#), and [33](#).
- 6-3 Are handover procedures between two sources of control well defined?
- Two pieces of equipment vying for control (or each assuming the other is doing the job) can be dangerous.
 - Conduct thorough switching analysis to ensure fail-safe transfers.
 - Lesson: [86](#).
- 6-4 Are there items that could resonate with one another?
- Example: Spacecraft can mechanically resonate with the launch vehicle, causing fatigue damage.
 - Lesson: [11](#).
- 6-5 Do interfacing organizations use different engineering conventions?
- Use the engineering model to verify interface early.
 - Lessons: English/metric units, positive/negative polarity grounding.
 - Lessons: [73](#) and [93](#).
- 6-6 Are launch integration operations thoroughly planned?
- Ground support, typically involving several organizations, often causes confusion, or even equipment damage.
 - Consult the AIAA/NRO Space Launch Integration Recommended Practices.
 - Lessons: [71](#) and [76](#).

THIS PAGE INTENTIONALLY LEFT BLANK

Section 7: Parts, Materials, and Manufacturing Process

- 7-1 Are drawing tolerances compatible with manufacturing processes?
- Unnecessarily tight tolerances cause manufacturing difficulties.
 - Tolerance stack-up can result in improper fits, or even failures.
 - Provide sufficient stress relief points to prevent fatigue.
 - Lessons: [47](#), [8](#), [4](#), and [52](#).
- 7-2 Are honeycomb structures vented?
- Unvented honeycomb panels can entrap moisture and violently delaminate during ascent depressurization. Several failures have occurred as a result.
 - Lessons: [1](#) and [34](#).
- 7-3 Does any part, including those subcontracted, contain pure tin-plating or cadmium?
- Tin whiskers can cause shorts and arcing and have disabled several satellites.
 - Cadmium, commonly used to plate airborne equipment, outgases in space.
 - Audit vendor or subcontractor materials lists to ensure completeness.
 - Lessons: [5](#) and [49](#).
- 7-4 Are there separable flared fittings (B-nuts) or check valves in fluid lines?
- B-nuts and check valves can leak.
 - Lessons: [83](#) and [15](#).
- 7-5 Are cables, connectors, and circuit cards labeled and/or keyed to prevent mismatching?
- Mismatching can cause inadvertent shorting during testing, even flight failure.
 - Lesson: [63](#).
- 7-6 Can installations at launch site be readily verified?
- The pressure of prelaunch preparation often causes mistakes.
 - Lessons: [63](#), [61](#), and [43](#).
- 7-7 Will rework be difficult?
- Rework is a fact of life in our business.
 - Lesson: [57](#).
- 7-8 Are there procedures to prevent parts from being mixed up?
- Different parts may look alike.
 - Lesson: [51](#).

- 7-9 Did a significant accident occur during manufacturing?
- Make sure the MRB thoroughly investigated the anomaly before accepting the part as-is.
Lesson: [6](#).
- 7-10 Have the root causes of manufacturing problems been corrected?
- The corrective actions should be back annotated on drawings or shop orders to prevent recurrence.
 - Lesson: [64](#).
- 7-11 Can each work instruction be verified?
- Verification should be done independently.
 - Make sure rework instructions include verification.
 - Lessons: [32](#) and [88](#).
- 7-12 Could handling procedures damage delicate hardware?
- Lessons: composite pressure vessels, primary battery, optics, and cryogenic equipment.
 - Procedures used to handle satellites during test and integration should be reviewed by safety personnel.
 - Lessons: [28](#), [22](#), [88](#), and [54](#).
- 7-13 Will handling or testing procedures reduce hardware life?
- Example: running high-speed tests in air can destroy lubricants.
Lessons: [9](#) and [21](#).
- 7-14 Were excessive acceleration factors used to qualify design life?
- Acceleration factors larger than 5-10 should be independently ascertained.
 - Different degradation mechanisms, not susceptible to thermal acceleration, may operate in flight.
 - Lesson: [95](#).
- 7-15 Are procedures adequate to prevent foreign objects or debris from being left inside the hardware?
- Loose materials, such as wipe cloths, have led to numerous reworks or catastrophic failures.
 - Lesson: [90](#).
- 7-16 Are manufacturing facilities sufficiently clean?
- Examples of equipment requiring special care include valves, high-voltage electronics, heat pipes, and optics.
 - Watch out for contamination (especially chloride).
 - Lessons: [16](#), [75](#), [45](#), [65](#), and [41](#).

Section 8: Testing and Evaluation

- 8-1 Do the tests *independently* confirm development results?
- If testing reuses equipment, analysis, or algorithms from design or manufacturing, a source of single-point failure exists.
 - Manual adjustment, such as shimming and alignment, should be independently verified, too.
 - Lesson: [96](#).
- 8-2 Have results been analytically established before testing?
- Tests should be used to verify analysis, not for discovery.
 - All testing must be preceded by prototyping and analysis, followed by model correlation and updating.
 - Problems found during late testing can be very costly.
 - Lessons: [60](#) and [37](#).
- 8-3 Is the polarity (phasing) of equipment (hardware coupled with software) correct?
- Phasing mistake (particularly in the ACS subsystem) is one of the most common sources of failure.
 - Lessons: [60](#), [80](#), [43](#), [53](#), and [93](#).
- 8-4 Can a simple test be used to crosscheck an elaborate test?
- Although a simple test will not provide the necessary precision, it can prevent gross errors.
 - If the equipment can pass the more elaborate test, it should pass the simpler test easily. Therefore, failure to pass the simple test must be treated as a red flag.
 - Lessons: [96](#) and [80](#).
- 8-5 Has all test data been reviewed for trends, oddities, “out-of-family” values, and other indicators of anomalies?
- Test sets should collect data and enable automatic trending.
 - Excessive current draw during electrical test (suggestive of an impending short) and high G spikes (indicating intermittent rubbing) during acoustic testing should receive particular attention.
 - Many problems occur during the first temperature cycle. Therefore, the results after the first cycle should be scrutinized.
 - Lessons: [71](#), [39](#), and [19](#).

- 8-6 Are all test anomalies fully understood?
- Many flight failures first occur during tests but are mistakenly attributed to “random failures” or “test set malfunctions.”
 - Test equipment should be sufficiently powerful to enable unambiguous assignment of anomaly causes.
 - Lessons: [92](#), [38](#), [46](#), [55](#), and [56](#).
- 8-7 Have the test articles been fully inspected after testing?
- It is particularly important to inspect the hardware after vibration or acoustic tests, thermal cycling, or live pyro firing.
 - Lessons: [100](#), [66](#), and [7](#).
- 8-8 Do the tests cover all operating modes?
- Conditions worth checking include eclipse transits, cold start, safe-holding, load shedding, and recovery.
 - Simulate each operational mode through several cycles.
 - Lessons: [38](#) and [84](#).
- 8-9 Is the test equipment compatible with the test conditions?
- Example: Test equipment used inside the thermal vacuum chamber must be space qualified to prevent damage to the hardware or test facility.
 - Lesson: [49](#).
- 8-10 Are procedures adequate to prevent hardware from being damaged during testing?
- Test equipment should be properly maintained and calibrated.
 - Trial runs using limited force, current, or temperature should be made first and the responses characterized.
 - Hardware should be protected from sudden test equipment malfunctioning.
 - Vibration of large satellites should be avoided.
 - Lessons: [24](#), [66](#), and [74](#).
- 8-11 Do tests accurately simulate time-dependent (especially start-up) behavior?
- Before an analog circuit stabilizes, it can behave unpredictably.
 - Equipment that change status abruptly (step motors or pyro firing circuits, for example) often exhibit, or require, an unexpected time profile to overcome initial resistance or prevent premature decay.
 - The test set should be able to record transients!
 - Lessons: [82](#) and [77](#).

- 8-12 Does the test equipment allow sneak paths?
- Sneak paths via the test set can mask hardware deficiencies (by providing gratuitous grounding or power, for example).
 - If test equipment temporarily provides certain functions, independently verify that the hardware can operate on its own.
 - Test set sneak paths can also damage hardware.
 - Lessons: [58](#) and [72](#).
- 8-13 Have the units demonstrated an ability to start without the need of ground equipment (plug-out) or manual intervention?
- It is particularly important to check payload, GN&C, and C&DH processors to prevent endless looping.
 - Lessons: [79](#) and [84](#).
- 8-14 Does the design allow adequate inspection?
- Unambiguous inspection criteria should be developed before verification.
 - Lesson: [47](#).
- 8-15 Does the system being tested represent the flight configuration?
- Insert enough test points to compensate for items that could not be live-tested (thrusters and deployment mechanisms, for example).
 - Lessons: [85](#), [19](#), and [53](#).
- 8-16 Does the test inject sufficient off-nominal conditions to ensure the equipment is robust?
- Examples of off-nominal conditions include current spikes, sluggish separation wire breakage, and excessive data rate.
 - Lessons: [94](#), [44](#), and [56](#).

THIS PAGE INTENTIONALLY LEFT BLANK

Appendix A: Space Systems Engineering Lessons 1-100

THIS PAGE INTENTIONALLY LEFT BLANK

1

Honeycomb Structures Should be Vented to Reduce Delamination Risk

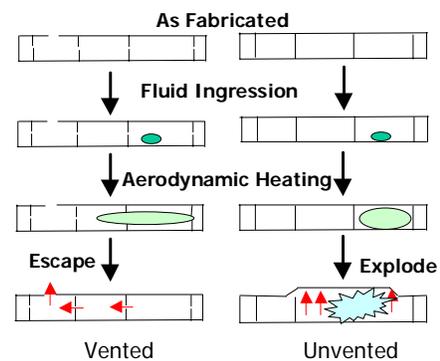
The Problem:

Several satellites have been destroyed when their honeycomb structures failed. Examples include:

- A NASA satellite was destroyed at T+103 sec when the payload fairing reached 600°F. During subsequent ground tests, the witness panels disintegrated (1964).
- A DOD rocket blew up shortly after launch. Later, the fairing's witness panel came apart when tested on ground (1966).
- Another DOD satellite was severely damaged upon launch. The fairing for the next flight was subsequently proof tested, whereupon it also burst (1981).
- Two solar array panels on a DOD program failed during qualification (1985).
- The massive hydrogen tank on an experimental reusable launch vehicle delaminated, eventually causing the program to be cancelled (1999).

The Cause:

Honeycomb panels for terrestrial applications are usually unvented—neither the panels nor the cores have holes. However, unvented honeycomb structures should not be used in space because aerodynamic heating during launch can cause temperature to rise dramatically. In an unvented design, entrapped fluid (e.g., moisture) can expand, turning each cell into a tiny pressure vessel that stresses the skin-to-core bonds. Debonding is apt to occur if the panel has a weak bond due to manufacturing defects.



Perforating honeycomb cells relieves pressure during ascent

Lessons Learned:

- Honeycomb structures for space systems should be vented whenever possible. The vast majority of spacecraft or launch vehicles use vented honeycomb structures, and these have not failed in space.
- If an unvented design cannot be avoided (e.g., to avoid contamination), it is necessary to adopt extensive development, verification, and quality assurance, including proof tests under applicable temperature and vacuum conditions. Aerospace experts are available to explain detailed design and quality-assurance requirements.

For more technical information, call S. R. Lin at (310) 336-7697.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

2

Perform Independent Mass Property, Stability Control, and Structural Load Analyses on Spacecraft and Launch Vehicles

The Problem:

Mistakes in determination of mass-property and control-stability analyses have caused a large number of launch failures. Examples include:

- Inappropriate reuse of aerodynamic coefficients (1994).
- Unanticipated structural vibration mode not filtered out (1995).
- Incorrectly simulated weight (1995).
- Underprediction of the load as well as an unexpected resonance due to wind shear (1992 and 1995).
- Unexpected increase in horizontal velocity (1996).
- Unaccounted roll mode caused by air-lit solid rocket motors (1998).

Flawed analysis has also led to numerous on-orbit anomalies.

The Cause:

Launching a satellite calls for extremely complex simulation of the mass, thermo-structural, fluid-mechanical, propulsion, and control properties (a single subsystem can easily involve over 100,000 equations). The state of the art in this area is far from robust: subtle assumptions, insufficiently sophisticated techniques, or human errors can all throw the results seriously off.

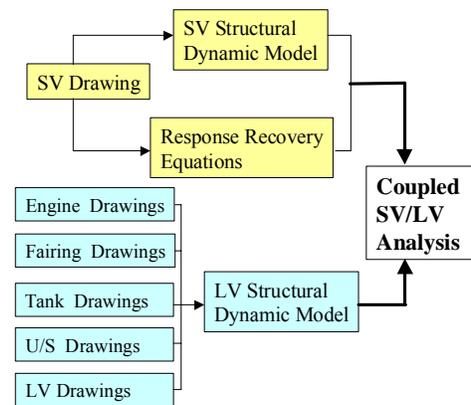
Moreover, when the satellite is integrated with the launcher, each organization must generate parochial models but each has little insight into each other's analytical process. Costly problems can easily arise without a clear settling of responsibility, especially with today's emphasis on proprietary data protection.

Lessons Learned:

- Inaccuracies on mass property, stability control, and structural loads continue to threaten mission performance.
- To ensure correct analysis, many programs require an independent analysis. These activities also help validate operational procedures, support flight anomaly resolution, and overcome the organizational issues. There have been no catastrophic failures in programs that abide by this policy, and several failures were averted thanks to independent analysis.

For more technical information, call Ray Skrinska at (310) 336-4001.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Integrating space vehicle (SV) to launch vehicle (LV) involves complex modeling; independent analysis is often necessary to overcome organizational barriers.

3

Rigorously Manage and Test Software, Including the Database

The Problem:

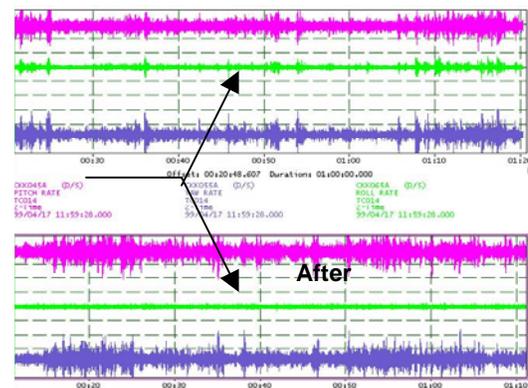
An expensive military satellite failed to reach the right orbit because a misplaced decimal point in the avionics database of the upper stage caused the reaction controller to fire excessively, depleting its fuel.

The Cause:

Multiple deficiencies in the software development, testing, and quality assurance (QA) processes allowed a single-point failure escape. Specifically:

- The process to create and test the constants database was poorly documented, fragmented, and not well understood. The control dynamics engineers created a new roll rate filter constant instead of using one that had been previously validated. This critical number was manually entered in error, slipped through visual inspection, and was not formally checked.
- The as-flown constant was neither independently verified nor validated due to a lack of overall software ownership. Many players were involved in the process, but none completely understood it.

Before wrong database was loaded



The wrongly placed decimal point caused the middle line to become flat. This anomalous reading was flagged at the launch site but fell through the crack.

As the program downsized, mission assurance functions were supposed to change from “oversight to insight.” This transition did not successfully take place, and the problem sneaked through all QA gates.

After the wrong constant was loaded, launch site personnel saw anomalous reading and tried to contact the designers. However, the issue was ignored. Even during the day of launch, the rocket showed a wrong response to the wind and to the rotation of the earth. A simple plot could have identified the problem and averted the failure.

Lessons Learned:

- One must test actual flight hardware and software.
- The integrity of software databases is no less critical than the source codes.
- The space business is extremely complex and human error cannot be completely eliminated. The system must be robust enough to catch the inevitable faults.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

4

Document Engineering Requirements As Clearly As Possible

The Problem:

Two very expensive mishaps occurred recently, in part, due to inadequate communications between the designers and the manufacturing operation:

- The combustion chamber of a rocket engine breached because an unclear requirement made it possible for a weak joint to pass quality assurance, leading to the loss of a \$230M commercial satellite.
- A DOD satellite was stranded in the wrong orbit because confusing drawing instructions led technicians to apply thermal protection tape in a way that prevented stage separation.

The Cause:

In the first incident, the seams of the engine are reinforced with many metal strips. The design requires the strips be brazed "80% per linear inch" (i.e., no big holes, see diagram), but the drawing only specified "80%".

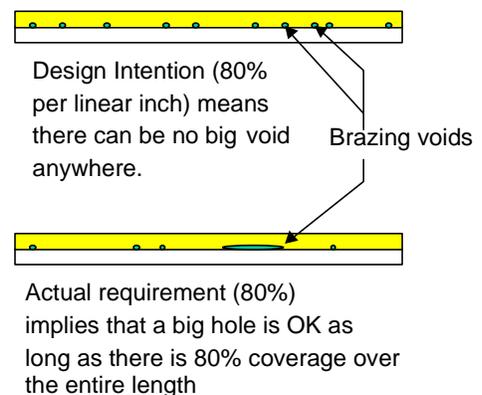
X-ray photos revealed that some strips were poorly brazed, but they were allowed to pass since the requirement was thought as "80% coverage averaged over the entire length of the reinforcement strip." The strips failed in flight.

In the second failure, the work instruction stated that the wrapping should be applied "within 0.5 inches of the mounting bracket flange" (instead of saying, e.g., no closer than 0.5 inches). The technicians, not knowing that the parts were to unfasten, applied the tapes as closely to the flange as possible, making separation impossible.

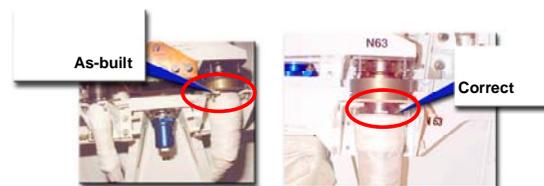
Lesson Learned:

- Engineers must clearly articulate their intentions and determine how the requirements should be interpreted or could be misconstrued. This is particularly true when making seemingly minor (Category II) changes.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Deleting the "per linear inch" phrase led QA to pass joints with low brazing coverage. In flight, the defective part caused combustion chamber to breach.



Thermal tapes were too tightly wrapped over the as-built connector and inhibited stage separation.

5

Avoid Pure Tin Plating

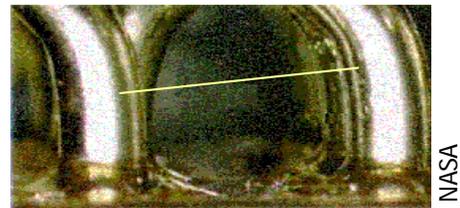
The Problem:

Pure tin plating can grow conductive filaments (whiskers) which have caused many problems. Examples include:

- In the late 1990s, at least four commercial satellites had problems with their spacecraft control processors (SCP), reportedly because whiskers grew on the relays and caused the power-supply fuses to blow. In three cases, both the primary and the redundant SCPs failed, and the satellites were lost.
- Again in the late 1990s, three DOD programs incurred costly delays: one discovered tin whiskers in an atomic clock, the second found tin whiskers on ground lugs, and the third saw tin whiskers forming inside thin-film capacitors.

The Cause:

MIL-STD-1547B bars several materials from space hardware. Among these "prohibited materials," tin is most noteworthy. Pure tin plating is often used commercially because it forms an excellent protective layer that accepts solder readily. Plating shops prefer pure tin over tin-lead to avoid lead disposal costs.



Tin whisker shorts

NASA

However, pure tin is liable to spontaneously form conductive whiskers, which can provide an unwanted conductive path and degrade hardware by causing shorts and even catastrophic arcing.

The whiskers appear unpredictably, without the need of an applied voltage or moisture (unlike silver dendrites), even in vacuum. It is impossible to ensure hardware integrity by inspection or by stress testing—the only way to prevent this problem is to eschew pure tin plating, fused tin, and alloys with very high (greater than 97%) tin contents.

Lessons Learned:

- Prohibit pure tin plating in both flight hardware and ground equipment but assume tin will be found.
- Ensure prime contractors flow down unambiguous plating requirements, and perform appropriate receiving inspections.
- Purge prohibited materials from project stores and standard catalog items, paying particular attention to the "commercial parts."
- Review subcontractor designs and part specifications to confirm that parts are safe.
- Apply conformal coatings on all exposed conducting surfaces wherever possible to inhibit shorts and vacuum arcing.

For more technical information, call Katherine Westphal at (310) 336-8794 or Steve Frost at (310) 336-7131.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

6

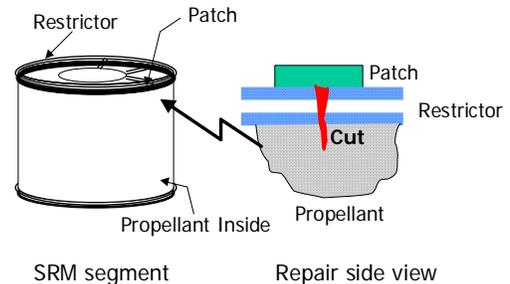
Following a Major Repair, Watch Out for Secondary Damage

The Problem:

In two launch failures, the Material Review Board (MRB) allowed repaired hardware to be used without taking secondary damage into full account. The first incident led to the destruction of three DOD satellites; the second mishap stranded another DOD satellite in a wrong orbit.

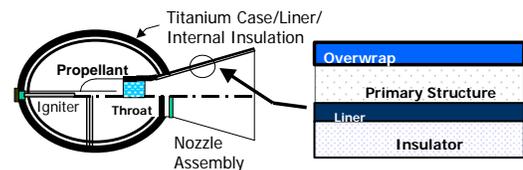
The Cause:

In the first incident, a large cut was made on a rocket segment during repair, and the slit was subsequently patched up. The engineers expected the cut to close by internal pressure, but it opened instead, allowing the flame to burn through the case. Afterwards, the manufacturer implemented several corrective actions to address the MRB repair process. The need to repair was eliminated by process changes, and other repaired segments were scrapped.



SRM segment Repair side view
Patching of deep cut allowed flame to burn through the case

In the second case, the fabrication of the apogee kick motor (AKM) nozzle involved wrapping a reinforcement layer over the primary structure in a bag, and heating the assembly under hydraulic pressure to cure. The bag broke, and the part came into contact with water. The contractor then machined off the semi-cured overwrap layer, laid up a new overwrap, and resumed production.



Apogee kick motor

Unfortunately, the part was not oven-dried—moisture was trapped in the primary structure and diffused back to the interface when the part was cured again. Not only was the mechanical strength lower as a result, but the interfacial adhesion between the primary structure and the overlap also became seriously degraded. During flight, the nozzle was unable to withstand the motor pressure and was ejected.

Lesson Learned:

- Ad-hoc repair processes tend to be much less defined and qualified than regular manufacturing operations. MRB reviews need to be more vigilant, and significant MRBs should be added to the readiness review process. In particular, the possibility of secondary damage must be taken into account.

For more technical information, call S. R. Lin at (310) 336-7697.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Perform High-Fidelity System Validation Tests for Pyrotechnics

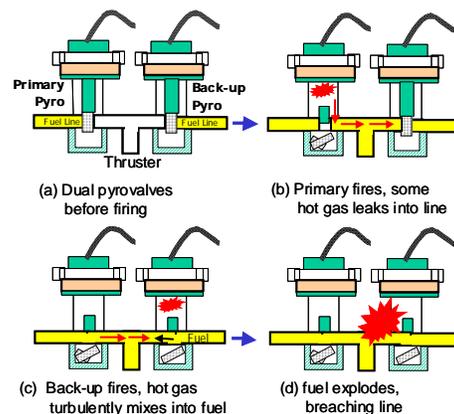
The Problem:

Explosive devices (pyros) are highly efficient, easily controlled, and can be readily stored. However, several anomalies occurred when pyros were turned on:

1. A science mission ended during the first orbit when its infrared telescope cover was unintentionally ejected, causing the loss of all cryogen (1999).
2. Three satellites, one for Earth observation, one for communication, and one for science, failed due to propulsion-system ruptures induced by pyros. A propulsive valve on a fourth similarly failed on ground (early 1990s).
3. An interplanetary probe almost fatally failed when the firing of a pyro initiator caused a voltage surge and induced a latch-up in the redundant memory board. The mission would have ended if the primary memory board had been affected (1989).

The Cause:

The telescope cover was ejected because a controller chip took a few milliseconds to warm up, during which a transient was generated. The designer did not take this known problem into account, and the design was not reviewed. Ground test failed to catch the flaw because a lab power supply was used, and its slower power rise time masked the transient. In flight, a relay applied power in two milliseconds, allowing the spurious firing to occur.



The four 1990 incidents involved dual "pyrovalves": the fuel-feed system incorporated two valves, the primary opening one second before the redundant. The second firing could lead to a blow-by of hot gas, igniting the propellant and breaching the fuel line. This problem escaped earlier tests that used an inert working fluid.

In the 1989 incident, the problem was not easy to spot, but could have been found if the engineering model had been tested with a simulated (non-explosive) pyro.

Lessons Learned:

- Pyros by themselves are very reliable, but the adjacent systems must be designed to withstand the mechanical or electrical shocks generated by the pyros.
- Tests should simulate flight configuration and functional performance.
- Post-test examinations of qualification or acceptance specimens should look for signs of inferred margin or incipient failure modes.

For more technical information, call Selma Goldstein at (310) 336-1013.

For comments on the Aerospace Lessons Learned Program including background specifics, call Paul Cheng at (310) 336-8222.

8

Solar Arrays Must Withstand Extreme Environments

The Problem:

Solar array mishaps have disabled numerous satellites. Examples include:

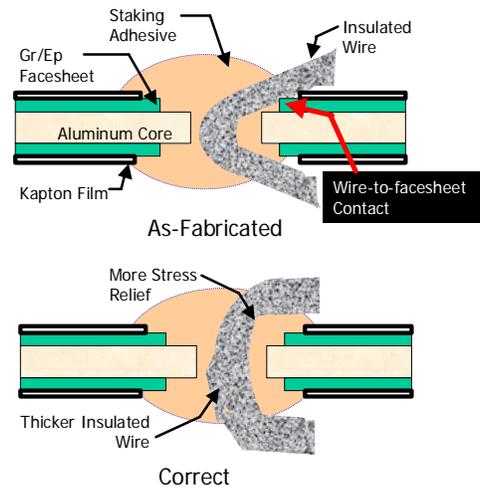
- Two Earth observation satellites failed due to shorts in the solar-array system, one in 1978 and another in 1993.
- In 1999, a technology demonstration spacecraft experienced excessive solar panel degradation that ended its mission prematurely.
- In the late 1990s, two commercial satellites suffered serious power losses, reportedly in solar storms.

The Cause:

Solar arrays contain many fragile elements, and are exposed to wide temperature fluctuations and other space hazards. They are thus particularly vulnerable to a host of problems that the designers must guard against. The mishaps above were caused by faulty materials, processes, and insufficient testing.

In the case of the commercial satellites, the wiring harnesses were squeezed into tight feed-through holes with sharp kinks and without sufficient strain-relieving loops. Temperature cycling, coupled with the movement of the adhesive, shifted the wires by several mils relative to the facesheets during each cycle.

With repeated heating and cooling, the insulation was abraded, as if by a saw. A short was inevitable, and was triggered by electrostatic discharges (ESDs) during weather storms. The problem could easily have been averted if the harness incorporated ample stress relief and thicker insulation.



Insufficient stress relief and insulation caused abrasion of wiring harness

Lessons Learned:

- Solar arrays should be carefully designed to prevent their fragile parts from being damaged by the hostile space environment.
- Satellites must be robustly designed to withstand the extremes of space weather as well as other space hazards.

For more technical information, call Robert W. Francis at (310) 336-6272.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Excessive Handling Can Destroy Solid Lubricant

The Problem:

Lubricants based on molybdenum disulfide (MoS_2) are used in gyros, drives, gimbals, or other moving mechanical assemblies. Several problems involving this lubricant have been noted, including:

- A microwave imager on a weather satellite catastrophically failed.
- A degraded sun sensor on another weather satellite caused excessive oscillation.
- The high-gain antenna on an interplanetary probe could be not fully opened.

The Cause:

MoS_2 has excellent properties in space, but it oxidizes in the presence of moisture. Hence, MoS_2 is degraded either by improper handling or by prolonged storage. Unfortunately, ground tests can fail to detect degraded lubrication because materials can behave differently on the ground than they do in space.

The imager problem occurred because manufacturing and storage exposed the labile lubricants in the slip-ring assembly to excessive oxidation. Furthermore, the part was stored for more than 11 years, causing more lubricant loss. The sun sensor problem was also traced to oxidation and contamination of the slip-ring materials during storage.

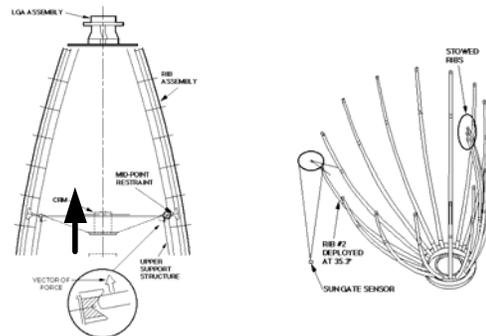
The high-gain antenna problem was caused by excessive handling (including vibration testing, rib pre-loading, and four cross-country trips) that dispersed the lube. Ground testing did not catch the problem because the vacuum test was not realistic and because the titanium pins got some lubrication (from the contaminants in the test chamber) not available in space.

Lessons Learned:

- Operation, testing, or storage of mechanisms under nonvacuum conditions must be performed with caution when MoS_2 dry lubricant is involved.
- Follow Aerospace's handling and storage guidelines to safeguard lubricants.

For more technical information, call Jeff Lince at (310) 336-4464.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



(a) High gain antenna unfurls like an umbrella. Excessive friction developed between the pin and the socket (inset) due to loss of lubricant.

(b, inverse view) The motor could not overcome the friction and stalled, and the antenna could not open.

10

Design Satellites to Withstand Space Weather, Regardless of Solar Cycles

The Problem:

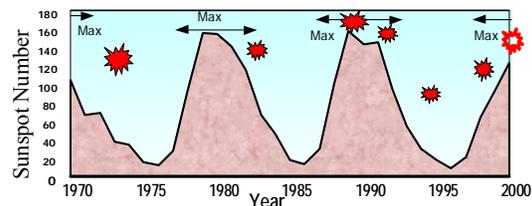
Space environment has caused hundreds of on-orbit anomalies, including:

- A military satellite lost power to its communications subsystem suddenly (1973).
- A weather satellite lost its primary instrument (1982).
- A foreign weather satellite lost attitude control (1988).
- A foreign communication satellite found its solar cells severely damaged (1991).
- A foreign commercial satellite was disabled for seven months after both reaction wheels failed (1994).
- A foreign communication satellite lost power (1997).
- A foreign science satellite was abandoned when increased atmospheric drag overpowered the attitude control system (2000).

The Cause:

The principal space weather hazards involve geomagnetic storms, which are stirred up when large numbers of solar particles hit the Earth's magnetic field. Storms can trigger an electrostatic discharge (ESD) in the spacecraft: all failures cited above except the last one involved ESDs.

Space weather hazards are often thought as mainly driven by the 11-year solar cycles. For example, there was extensive "satellite-killer" hype in the media in 2000 because one cycle peaked late that year. Conversely, some people associate periods of low solar activities with minimal weather hazards.



Space Weather Hazards Can Occur Outside of Solar Max

This belief is unfounded since space weather hazards and solar activity only marginally correlate. Geomagnetic storms can occur anytime, not just during the height of the solar cycles. Satellites can thus fail during valleys of solar cycle as easily as during peaks. Moreover, all storm prediction efforts, including new spacecraft designed to monitor solar activities, have been unsuccessful so far, and satellite operators cannot count on being forewarned of weather threats.

Lessons Learned:

- Spacecraft must be designed to withstand worst-case space environments as a matter of course.
- Satellites should be hardened against ESD, using well-established design guide-lines on structure, materials, shielding, cable interfaces, and circuits.

For more technical information, call Harry Koons at (310) 336-6519.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

11

Carefully Evaluate Satellite-Launcher Interface

The Problem:

An experimental spacecraft fell silent after having been successfully released from the launch vehicle. This failure was deemed to have occurred because unexpectedly high vibration developed in the launch vehicle before it was air-dropped, imparting stress in the satellite beyond its design limit.

The Cause:

This failure was caused primarily by a satellite-launcher interface problem:

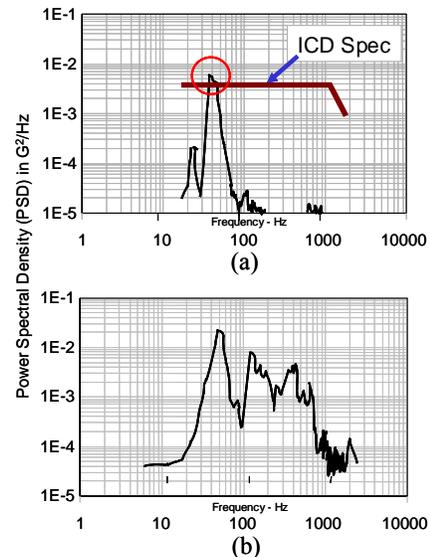
- The booster, while being carried by the launching airplane, vibrated at 40-50 Hz. In several previous flights, shaking went beyond the level spelled out in the Interface Control Document (ICD). As a result, the rocket contractor reduced the airplane's speed to minimize this problem. Still, vibration in this flight was double the specification.
- The satellite exhibited a structural resonance at 40 Hz. During factory test, this resonance amplified an acceleration input six-fold.
- The satellite contractor conducted the vibration acceptance test at a lower level than the ICD specification. A defect in the electronics or harness probably went undetected in the test, but propagated under a combination of excessive in-flight vibration and resonance to cause the failure.
- Both the launcher and the satellite prime contractors recognized the vibration issue and proposed to conduct a coupled-loads analysis. It was not performed because the program office, which served as the overall systems integrator, lacked funds.

Lessons Learned:

- Cables and connectors must be designed to withstand vibration-induced stresses.
- Margins must be reserved both in dynamic input estimation and in design.
- The interfaces among different organizations, particularly between the spacecraft side and the launcher side, frequently lead to problems. Independent analysis is advised to overcome organizational barriers (see Lesson No. 2).

For more technical information, call Robert Morse at (310) 336-2364.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Vibrational forces, expressed as power spectral density (PSD) in log scale (a) imparted on the spacecraft by the carrier airplane, and (b) as satellite's response toward an even level of excitation. Spacecraft resonated at the frequency where above-spec shaking took place.

12

One Requirement, One Statement

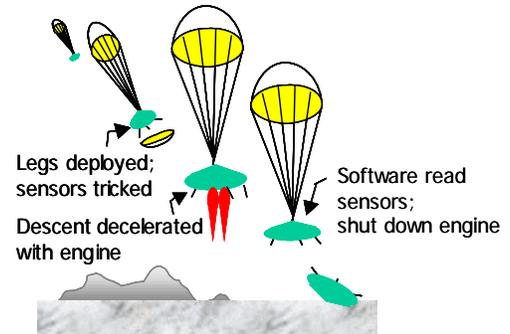
The Problem:

Contact to an interplanetary probe was lost.

The Cause:

As the lander parachuted down, it deployed three legs, each with a sensor designed to command the engine off upon touchdown lest the lander overturn.

Leg deployment shock could spoof the sensors into thinking the probe had landed. To prevent the confusion, the systems spec required: *“The sensors shall...(commence operation shortly before touchdown). However, the use of the sensor data shall not begin until...(after the leg deployment completes)...”*



Landing Sequence

This “*However...*” phrase was unfortunately not picked up by the software team or by other subsystems, and was not specifically tested at the system level. During descent, the deployment shock set off a status flag. When the touchdown sensing logic subsequently ran, it was misled into thinking landing already occurred. The descent engine shut itself off prematurely; the probe crashed.

The software walkthrough and integration/test did not detect this problem (logic flow diagrams could have helped). What’s more, a leg-deployment test failed to detect the fault because the sensors were improperly wired at first. A rerun of the deployment test, which might have caught the error, was not performed after rewiring.

Lessons Learned:

- Do not lump several requirements together—write them out separately so that each can be tracked individually. Negative statements (e.g., “Sampling shall *not* begin until...”) may cause misunderstanding and should be avoided.
- Systems engineers must take ownership of requirements and partition them to the appropriate subsystem. Whether or not a requirement is the software’s responsibility, for example, should not be left to the discretion of the software team.
- Systems engineering must ensure thorough end-to-end failure mode testing.
- The software review process should emphasize logic flow. Tests should exercise every requirement to see if there are conditions that could cause the software to fail.
- Test planning needs to consider transients or spurious signals.
- When important tests are aborted or are known to be flawed, they must be rerun after the errors are fixed. Repeat the test if any software or hardware involved are changed.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

13

Flexible Solar Arrays Are Susceptible to Thermally Induced Vibrations

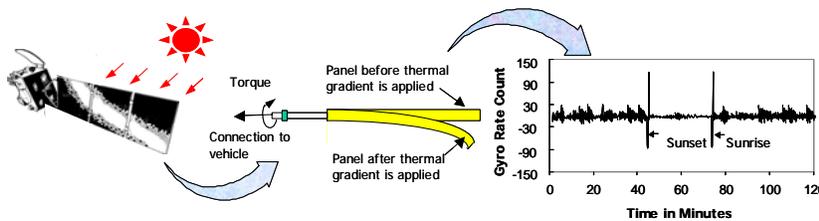
The Problem:

Thermally induced vibrations of spacecraft appendages have recurred numerous times. Resultant problems include:

- Two science satellites stopped spinning (early 1960s).
- Two Earth observation satellites showed large disturbances about the roll and yaw axes whenever the spacecraft entered or exited sunlight (early 1980s).
- A space observatory had to have its solar arrays replaced on-orbit because “jitters” interfered with star pointing (1993).
- A scientific satellite failed due to heating and expansion of the solar panels that damaged the structure (1997).

The Cause:

Spacecraft equipped with long appendages or solar arrays are susceptible to attitude perturbation upon entering or leaving the Earth's shadow, because large temperature gradients can develop around the boom. The sun-facing side of the boom or array can bend and create a torque on the satellite very rapidly, causing a flutter. Satellites with a single solar array are most susceptible.



Long appendages can deform and cause the spacecraft to shiver during eclipse transitions. Effective attitude control algorithms should be developed to address this concern.

The space observatory mentioned above, for example, employed flexible solar arrays with telescoping booms. A thermal gradient as much as 25-deg C developed around the boom circumference within one minute, causing the tip of the spar to defect by 20 cm.

Lessons Learned:

- Flexible solar arrays and supporting equipment are sensitive to thermal environment.
- Thorough thermomechanical analyses of the solar arrays, particularly on their modal frequencies, should be conducted.
- Control algorithms used to mitigate the effects of solar-array excitations should be refined.

For more technical information, call John Welch at (310) 336-6556.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

14

Look Beyond Specifications in Qualifying Materials by Similarity

The Problem:

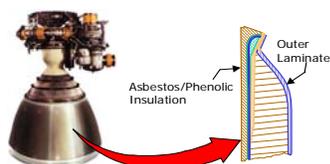
Numerous failures have occurred due to deficiencies in substitution materials that were thought to be similar to those originally specified. Some recent examples include:

- A rocket nozzle failed during test firing because a replacement insulator delaminated.
- The propulsion valves in a rocket broke down just before launch because the oxidizer reacted with a new cleaning solvent.
- A solar array would not open in space because radiation caused a rubber spacer to become sticky.

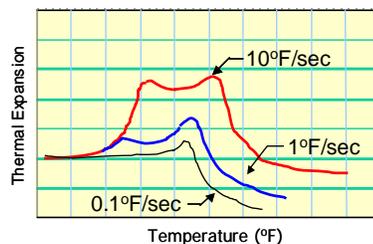
The Cause:

Programs sometimes must replace materials that are no longer available. It is often thought that if the substitute meets all the specifications, it can be accepted "by similarity." This approach can be risky; specifications usually only call out rudimentary requirements to facilitate incoming inspection—key tests used to qualify a material may be cumbersome to repeat, and are routinely left out of the spec as new materials lots are received.

In the first incident, a supplier problem prompted the contractor to select a replacement resin for the nozzle skirt. This new material met the applicable specification, had been used on other programs, and had passed an array of tests in the laboratory. However, test results of the new material were statistically different from the original material, and test conditions were not sufficiently flight-like: many properties were measured at room temperature, whereas the flight temperature approached 3000-deg F. Additionally, certain critical properties were not measured, and the vital thermal expansion test was performed at too low a heating rate.



Nozzle Skirt Insulation



The replacement material outgassed and delaminated during firing. This problem escaped qualification since slow heating rates (0.1-deg F/sec) used in the lab provided time for the gas to escape. Faster rates would have revealed the issue.

In a test firing, the flame burned through the new resin. At the time, two rockets having nozzles made from the new materials were already being prepared for launch. Potential losses of the satellites were narrowly averted.

Lesson Learned:

- Substitute materials should be tested under conditions that realistically simulate flight conditions and give results comparable to those exhibited by the original material.

For more technical information, call Wayne Goodman at (310) 336-5356.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

15

Avoid Separable Flared Fittings

The Problem:

Tubular fittings with flared ends, commonly referred to as B-nuts and designated as AN, MS, and MC types, are sometimes used as separable plumbing joints in rocket engines and spacecraft propulsion subsystems. These connectors are often found to leak during tests, and may be difficult to fix. Leaky fittings have also been implicated in several in-flight malfunctions, including the failure of a transfer vehicle.

The Cause:

Standard separable connectors are commonly used in ground systems to facilitate part replacement. B-nuts work by converting the applied torque into a stress that physically clamps and deforms the flared end of fitting until it fits tightly over the threaded element.

However, just as bolts in furniture can unscrew over time, the flared end of these fittings can undergo "stress relaxation" and become loose, resulting in a leak. Launch vibration can also pull the nuts back and cause leaks.

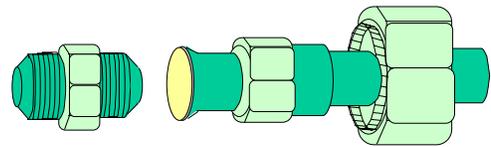
How fast the seals loosen depends on the manufacturing process, storage conditions, and other factors, but tests have shown that the applied torque can drop by one-third over a matter of weeks. Unless retightened (which can be difficult to do because the connectors may not be accessible), loose fittings can cause failures.

Lessons Learned:

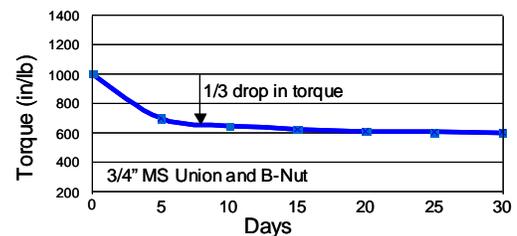
- Separable fittings in fluid lines should be avoided wherever practical in favor of permanent connections such as welded or brazed joints.
- Where separable connectors must be used, the fittings should have machined sleeves or redundant sealing surfaces. All separable connectors should be readily accessible at all stages of assembly and at the launch site to allow torque checks and repairs.
- All separable fittings should be torque-checked as close to launch as possible. If torque checks are not possible within 10 days prior to launch, locking devices that do not cause contamination should be used.

For more technical information, call Leon Gurevich at (310) 336-1268.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The flared-fitting seal relies on maintaining the clamping force high enough to deform the flare into a fit on the threaded elements.



The applied torque can drop substantially in a week and cause leaks to develop.

16

Systematically Monitor and Control Contamination

The Problem:

Contamination has degraded numerous radiators, thermal coatings, solar arrays, sensors, moving mechanical assemblies, and other components in space. Examples include:

- The sun-viewing bays of an interplanetary probe were 20-deg C hotter than anticipated.
- The radiator of a data-relay satellite became too hot.
- An instrument failed on-orbit when internal outgassing caused arcing.
- The focal plane on an early-warning satellite degraded.
- A satellite lost its orientation accuracy because three star trackers were fouled.
- The solar array output from five navigation satellites decreased more than expected.
- The wide-field planetary camera on a space telescope lost its ultraviolet capability. A similar camera degraded during thermal vacuum test.

The Cause:

Contamination is a serious risk during all phases of a spacecraft's life. Particulate can accumulate during manufacturing, testing, storage, and launch. Volatile materials can be released during vacuum tests or in space, and condense on critical surfaces. Some molecules can react with sunlight to deposit tenacious films that darken over time.

Contamination control has historically been performed on a "best effort" basis: all "low outgassing" materials were deemed acceptable in any application in any quantity, and manufacturing requirements were rather arbitrary.

Today's new sensors, which must be kept extraordinarily clean, require a quantitative contamination budget flowdown throughout the entire spacecraft lifecycle. Sophisticated monitors and models should be used to verify that derived cleanliness requirements are met.

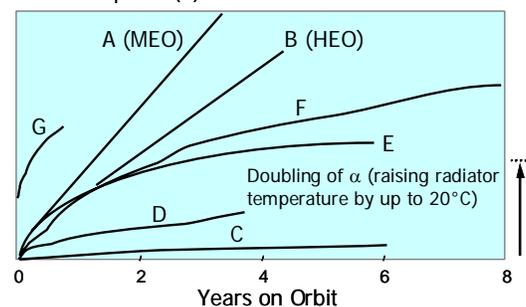
Lessons Learned:

- Recognize the importance of contamination-control engineering during every phase of development and hardware design.
- Perform contamination budget analysis, using tools derived from experimental data.
- Establish quantitative cleanliness requirements and apply cutting-edge processes to control particulate and molecular contamination.

For more technical information, call David Hall at (310) 336-5896.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Solar Absorptance (α) of Silverized Radiator Mirrors



Contamination of radiators makes electronics run hotter. Except for curves A and B, data was obtained from GEO satellites. Satellite C used a special design to reduce contamination.

Watch Out for the Return of Leonid Micrometeoroid Storms

The Problem:

When the Earth crosses a comet's orbit, tiny debris trailing the comet can trigger micrometeoroid outbursts and damage satellites. For example:

- A scientific spacecraft suffered a hit and lost substantial telescope capability (1991).
- A communication satellite lost its Earth sensor and had to be abandoned, probably due to a particle strike that triggered a power surge (1993).

The Cause:

Micrometeoroid showers occur several times a year, with dozens, sometimes hundreds, of particles per hour burning up in the Earth's atmosphere during a shower's peak.

Showers with 1000 or more particles per hour are called storms. The Leonid storms in 1966 exhibited a peak rate approaching 100,000 per hour. Leonid particles travel at speeds of about 70 km/sec and pose a significant threat to satellites.

Satellite operators can mitigate risks by:

- Turning telescopes away from incoming particles, adjusting solar panels, and orienting the satellite to minimize damage to internal hardware.
- Reviewing procedures for rebooting subsystems.
- Making sure experienced personnel are on duty during the storm.
- Turning off equipment sensitive to electrostatic discharge (ESD), and avoiding commanding the satellite or firing thrusters during storms.

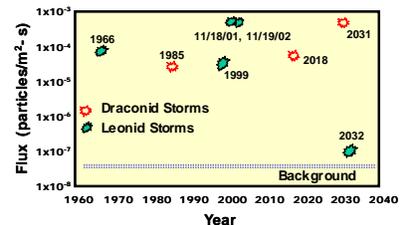
These techniques have proved successful. In the widely publicized 1998-2000 Leonid season, only a few minor anomalies were attributed to possible meteor strikes.

Lessons Learned

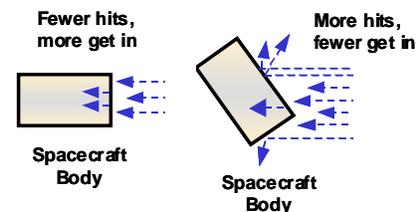
- Awareness of the space environment situation is vital.
- Advanced planning in anticipation of the coming storms is essential.

For more technical information, call Dave Desrocher at (719) 638-2280. A monograph from The Aerospace Press, *Dynamics of Meteor Outbursts and Satellite Mitigation Strategies*, discusses this issue in great length.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The next Leonid storms will occur in November 2001 and 2002. Each may have multiple bursts over approximately 16 hours. Long-term projections remain imprecise.



One way to reduce storm damage involves orienting the satellite to face the micrometeoroids at an oblique angle. Although more surface is exposed, particles will tend to glance off instead of penetrating into the spacecraft.

18

Make Sure Critical Software Performs in its Intended Environment

The Problem:

The 1996 maiden flight of a launch vehicle ended in a crash.

The Cause:

The launcher's flight control system, which had derived considerable heritage from the previous generation, used two identical inertial reference controllers, including a "hot" stand-by.

One function inherited from the legacy software computed the platform alignment before launch. This function was no longer needed in the new generation.

The new rocket flew a different trajectory, creating an alignment bias that was too large for the legacy code to compute. An "operand error exception" occurred.

Such errors are common, and are typically handled by software (for example, by inserting "likely" values). Unfortunately, although the programmers did identify the alignment bias input as one of the several variables capable of causing operand errors, they chose to leave it unprotected, probably supposing that there would be large safety margins.

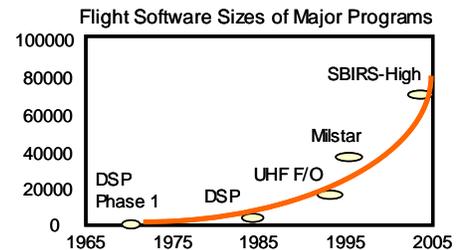
More tragically, the system was designed in the belief that any fault would be due to random hardware problems, and should be handled by an equipment swap. Thus, when the software detected the errant and irrelevant exception, it halted the active controller and switched to the backup. Of course, the backup immediately encountered the same error exception, and also shut down. The launch vehicle in essence destroyed itself even though both controllers worked perfectly.

Lessons Learned:

- Hardware redundancy does not necessarily protect against software faults.
- Mission-critical software failures should be included in system reliability and fault analysis.
- Software specifications should always include specific operational scenarios.
- Software reuse should be thoroughly analyzed to ensure suitability in a new environment, and all associated documentation, especially assumptions, should be reexamined.
- Extensive testing should be performed at every level, from unit through system test, using realistic operational and exception scenarios.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



As software takes over many functions that used to be controlled by hardware, code sizes increase almost exponentially. Software reliability thus poses a growing challenge and warrants more quality assurance efforts.

19

Be Sure that the Architecture Isolates Faults

The Problem:

A pair of scientific satellites was launched in late 2000, and in less than three weeks both stopped receiving commands. Both spacecraft failed due to improperly implemented software, compounded by a fault-intolerant power-distribution architecture

The Cause:

The root failure cause involved overheated relays, which should receive pulsed commands according to the system requirements. Software documents did not pick up this specification, and a constant voltage was supplied instead.

A status indicator relay coil shorted under continuous heating in vacuum and caused the circuit breaker of Receiver B to trip. Receiver A should have been isolated from this fault, but was not because it was joined to Receiver B via an “OR” diode. It thus also suffered a current surge and blew the fuse, preventing the ground station from controlling the satellite.

The architectural oversight escaped design review probably because the status indicator relays were not thought to be crucial. However, because these relays drew current from both receivers, a short in either of them would cause a catastrophic failure of the system.

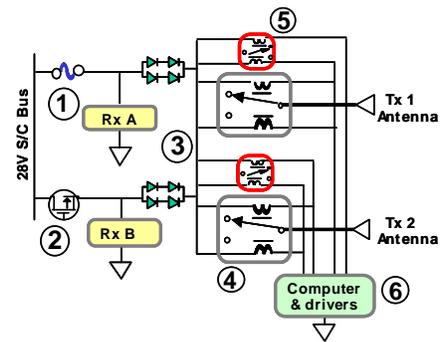
The continuous command fault was not detected during unit test because the test set software correctly drove the relays with pulsed signals. System test should have caught the error because the continuously powered coils drew five extra watts, a considerable amount in a low-power system. Unfortunately, the extra power draw was not noticed.

Lessons Learned:

- Create and use a verification matrix for all levels of test requirements.
- Inspect all test data for trends, oddities, and “out-of-family” values, even when all values are within expectation. Evaluate all indicators for potential impacts, should trends continue. Seek to explain all instances of anomalous data.
- Incorporate flight software into test at the earliest opportunity.
- Avoid sneak failure paths by keeping circuit designs straightforward.
- Use isolation resistors or downstream fuses to prevent a grounded component from bringing down the entire system.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



System Schematics (Simplified)

The bus power fed into one command receiver (Rx A) via a fuse (1) and into another via a circuit breaker (2).

The receiver power was, via “OR” diodes (i.e., the downstream circuits can draw current from either receiver, 3), tapped off by two transmitter (Tx)/antenna switches (4) and two commercial-off-the-shelf status indicator relays (5). The relays were commanded by the flight computer (6).

Thoroughly Analyze and Test Deployables

The Problem:

Troubles associated with deployables have affected numerous satellites. For example:

- A foreign satellite could not open its solar sail, causing attitude-control errors to build up and the mission to fail (1982).
- A comsat was abandoned after a solar array failed to deploy (1987).
- An interplanetary probe could not unfurl its high-gain antenna (1989).
- Two solar arrays of a comsat jammed, leading to an insurance claim of over \$200 million (1998).

In addition, several potential on-orbit catastrophes have been narrowly averted. Stuck deployables have been shaken loose by space-walking astronauts or by rocket burns. In 1991, the antenna on a comsat stuck and disabled the satellite for three months, until repeated on-orbit maneuvers finally freed it.

The Cause:

Deployables are complex mechanical equipment customized for each mission, and thus lack the heritage of testing and usage common to electronic devices. With deployables, robust design, thorough testing, and careful handling are vital.

The design must provide adequate force margins, including thermal and tolerance analyses, to overcome all resistances. The 1991 anomaly cited above was caused by interference from thermal blankets. A thermal blanket Velcro pad likewise snagged the magnetometer boom of another satellite in 1990.

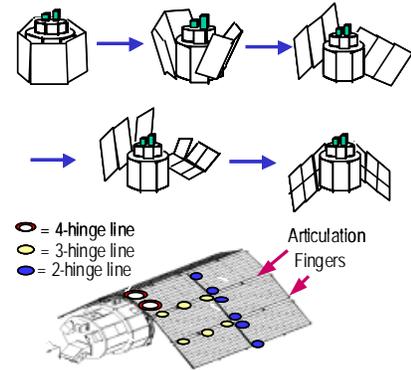
Testing is a major part of the deployment development effort. Special tests and off-loading fixtures (such as balloons or air bearings) are frequently required to demonstrate deployability in a zero-gravity environment. Some deployables cannot support their own weight on Earth, and require special testing accommodations.

Lessons Learned:

- Make sure the design can be effectively tested.
- Avoid unconventional designs, especially those involving complex motions.

For more technical information, call Brian Gore at (310) 336-7253.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Deployable design should not be so complex that it cannot be verified on the ground. The deployment scheme in the satellite depicted above was too complex to be tested, and The Aerospace Corporation had to run an in-depth analysis to verify it. Although the deployment proved successful in space, the contractor learned a lesson and decided to revert to simpler schemes in the future.

21

Prevent Loss of Lubricating Oil and Grease During Storage and Test

The Problem:

Many failures have been caused by mishandling of liquid lubricants (oils and greases), particularly during prelaunch storage. For example:

1. The reaction wheels on several navigation satellites malfunctioned.
2. Many instruments stopped functioning when their ball-bearing cages ran dry.
3. The focusing system in a space telescope developed high torque and had to be replaced in space.
4. A gyroscope stopped working during testing.
5. A sensor problem affected eight satellites, and caused an on-orbit failure.
6. A gimbal drive unit developed excessive noise.

The Cause:

Liquid lubricants are susceptible to physical loss and chemical degradation. Physical loss can occur by evaporation and migration. In the first mishap above, the satellites were stored longer than originally anticipated, and some oil was lost. Later builds switched to a less volatile oil, and stored the wheels separately from the satellites, with their spin axes oriented horizontally to limit migration.

Physical loss can also involve absorption. The second mishap occurred because the hardware surfaces are porous. Oil was absorbed into them and was no longer available for lubrication.

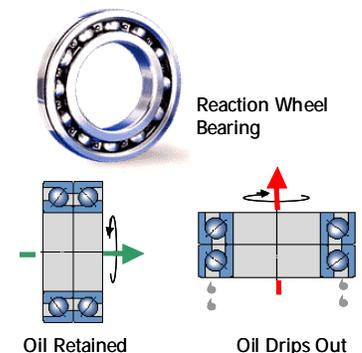
Oil and grease can also chemically degrade and lose their ability to lubricate. Unprotected lubricants have been known to polymerize (which caused mishap No. 3), oxidize (No. 4), react with titanium surfaces (No. 5), or dissolve plastics (No. 6).

Lessons Learned:

- Minimize oil evaporation and migration during hardware storage.
- Use enough oil to sustain storage and operation needs. If porous hardware requires lubrication, they should be thoroughly cleaned, protected from moisture, and stored in oil.
- Test high-speed moving parts in an inert environment to prevent oxidation.
- Perform materials compatibility analysis to avert chemical reactions.
- Check *NASA Mechanisms Handbook* (NASA/TP-1999-206988) for guidelines on mechanical assemblies.

For more technical information, call Steve Didziulis at (310) 336-0460.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The spin axes of gyros and wheels should be oriented during storage in such a way as to ensure oil retention.

Be Aware of Challenges in Silver/Zinc Battery Manufacturing and Deployment

The Problem:

Silver-zinc batteries have supplied power to many launch vehicles and upper stages over the years. These batteries are susceptible to a variety of problems during development and manufacturing. In the field, batteries have splashed operators with caustic chemicals, delayed launches, and caused a serious malfunction in an upper stage.

The Cause:

Launch vehicles rely on primary (non-rechargeable) batteries to power avionics, pyrotechnics, range safety, and other equipment. Silver/zinc batteries, the most common type, can be stored “dry” for several years until activated by the addition of the electrolyte. The activated batteries must be used within weeks or, at most, a few months.

Customized for the launch and space environment and for each particular program, batteries are hand-built in small lots. They are sensitive to operator changes, material alteration, contamination, and a host of factors during development.

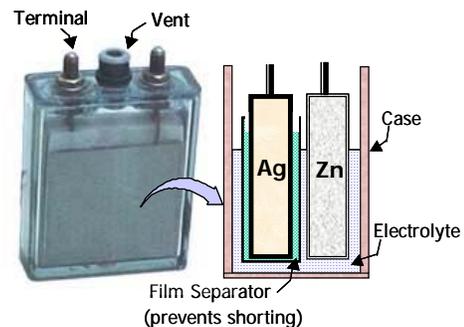
At launch sites, mishandling of batteries can allow caustic chemicals to escape. If too much electrolyte is added, batteries can spew or even start fires. The upper stage problem cited above, for example, occurred because electrolyte escaped from inadequately vented cells, causing a short to ground.

Lessons Learned:

- Design, documentation, manufacturing, storage, and field application of batteries require constant vigilance.
- Materials must be thoroughly screened before being incorporated in batteries.

For more technical information, call Margot Wasz at (310) 336-2141.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Batteries consist of numerous cells, each containing a silver electrode and a zinc electrode. One of the most common battery problems pertains to the plastic separators that wrap around the silver electrodes. Minor changes in the constituents of these items have led to incompatibility problems with the electrolytes, causing excessive shrinkage or chemical reactions.

Make Sure Requirements Are Developed Correctly

The Problem:

As a planetary probe neared its objective, a potentially crippling flaw was discovered—the designers had neglected to take the Doppler Effect into account.

The Cause:

After a seven-year journey toward one of the Saturn's moons, the probe will enter the moon's atmosphere, collecting data during descent for relay to the Earth via an accompanying orbiter.

As the probe speeds away from the orbiter, the data signal frequency will drop slightly, due to the Doppler shift. According to the Inquiry Board Report, this unavoidable frequency drop was overlooked from initial project requirement determination all the way through design specification of the orbiter's receiver. Extensive internal and external reviews failed to discover this oversight, in part due to a proprietary issue. Later, the design flaw escaped the system-level test because an incorrect frequency was used.

Two and half years after launch, a check-out of the probe indicated that the signal frequency was outside the receiver's bandwidth. Had the problem been unveiled on the ground, it could have been fixed with a simple software patch. Unfortunately, the software is not accessible in flight.

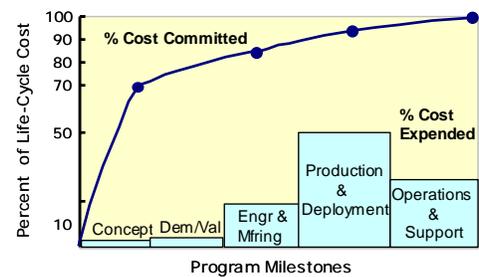
To minimize the Doppler shift, the flight trajectory had to be changed, at considerable expense in fuel, so that the orbiter will be farther away from the probe as it descends.

Lessons Learned:

- Formalize requirement development process and capture lessons.
- Provide adequate design margins and operational flexibility, such as the ability to use software patches.
- Make sure that the hardware or software a contractor wants to reuse from another program is indeed applicable and has a satisfactory flight history. Do not be deterred by the excuse that details are not available because the previous program was proprietary or classified—there are always ways to get around that hurdle.

For more technical information, call Mark Simpson at (310) 336-0159.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Most of the project's cost and performance are established by front-end decisions, but mistakes made there are difficult to catch. More resources, including the most experienced personnel, should be made available to ensure the early decisions are made properly.

Designers should thoroughly review the history of similar projects. If the probe designers had analyzed the requirements of other deep space projects, both the importance of the Doppler shift and the correct way to perform end-to-end test would have become obvious.

24

Safeguard Hardware Against Inadvertent Overtesting

The Problem:

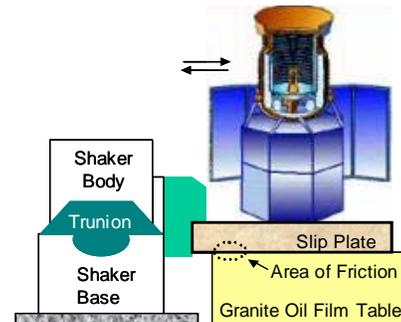
A satellite suffered considerable damage during vibration test because worn-out equipment misled the test operator into applying an excessive force.

The Cause:

Prior to vibrating the spacecraft, the operators first subjected it to a low-level calibration test to compute how much force should be applied to achieve the specified acceleration.

Unfortunately, the shaker was over 40 years old, and its trunion bearings had broken. The slip plate came into contact with the shaker table, resulting in an interference that attenuated the satellite's motion.

Unaware of the malfunction, the test engineer thought a much larger force needed to be applied to achieve the required acceleration. This force overcame the start-up friction, but overshot the acceleration by tenfold, damaging the spacecraft.



Friction during start-up can greatly exceed that during operation. This problem, known as stiction, frequently causes trouble. For example, when a tape drive is adjusted, the tape may not move until enough voltage to overcome the stiction is applied; but then the force is too large, and the tape suddenly runs wild.

Lessons Learned:

- Make sure that test facilities are maintained and checked.
- Implement overtest protection (such as over-temperature trip circuits in thermal chambers).
- Take risks of overtesting during vibration tests into account. In particular, large satellites should typically be acoustically tested instead of vibration-tested to prevent damage.
- Step up vibration tests from one-third to one-half of the full level so that the required force can be more accurately computed.
- Test procedures, set up, and data should be thoroughly checked to account for operator mistakes and avoid damage.

For more technical information, call Alan Peterson at (310) 336-0101.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Thoroughly Verify All Software Changes

The Problem:

A launch vehicle failed because part of a command line was left out of a software change.

The Cause:

The launch vehicle had flown successfully several times. This mission, however, had to be launched at a particular time. Accordingly, the time variable in the software was changed from *Reference Time* to *Fixed Time*.

Multiple updates to the ground software were made, including one that controlled a valve regulating the ground-supplied nitrogen and, indirectly, an attitude-control engine. This valve should have been closed shortly before liftoff.

Since the *Reference Time* no longer applied, an existing command, “If the state is Abort (or the state is Nominal and *Reference Time* is T-105 sec), close Valve X.” should have been updated to: “If the state is Abort (or the state is Nominal and *Fixed Time* is T-105 sec), close Valve X.”

Unfortunately, the conditional statement in the parenthesis was omitted, and the command became “If the state is Abort, close valve X.” Hence, the valve stayed open, and the engine malfunctioned.

The error went undetected because the change notice included several unrelated items, failed to explain why the control code was changed, and did not compare the was/is algorithms. In addition, not all logic paths, displays, and output commands were verified.

Lessons Learned:

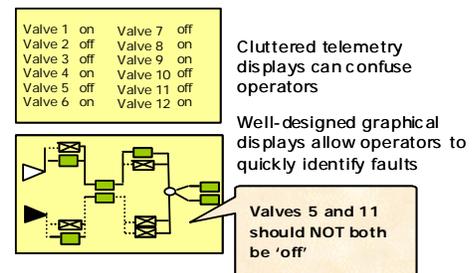
- A small software error can have catastrophic mission impacts.
- Software change processes require the same degree of rigor as the original development. Each change and associated rationale must be individually approved.
- Retest and regression testing should be formal and thorough. All logic paths affected by changes must be verified, and all results must be checked.
- Operational status, particularly off-nominal indicators, must be displayed effectively.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

The failed launch was rehearsed three times, during which the console operators could have spotted the open valve but missed it.

Graphical displays, summarized telemetry data, and error checking should be provided to allow operators to identify and diagnose faults.



To learn more about human factor engineering, see SMC Publication HM-RB-2001-1, “Human Computer Interface Display Conventions” on the “Documents” section of the SMC/AX Web site (http://ax.losangeles.af.mil/chief_engineer/).

Make Sure Hardware Analyzed Is Hardware Actually Built

The Problem:

A technology-demonstrator mission was terminated after only eight months because an oversight in thermal analysis was unrecognized by two projects.

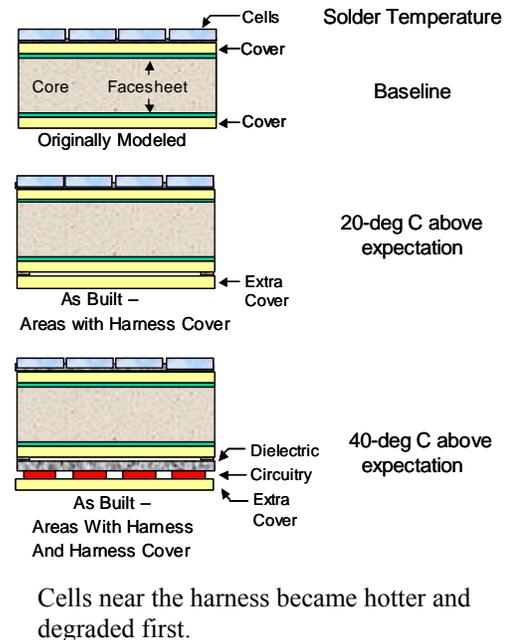
The Cause:

The solar array was originally designed for a “faster, better, cheaper” mission. Unfortunately, the thermal model did not account for the presence of harnesses and harness covers which, by preventing heat from radiating away, raised the temperature in the cells near the harness by as much as 40-deg C, causing stress in the solder joints of the cell interconnections. The joints cracked open, and the circuits failed.

The original mission never flew. However, the panel design was carried into this program without being revalidated, most likely because of resource constraints.

In retrospect, if a thermal analyst had actually looked at the hardware and seen the conspicuous harnesses at panel fold locations, the problem would have been caught right away.

Solar Panel Configuration: Modeled versus actual



Lessons Learned:

- Designers should be called back to inspect the products, to see if there are major differences between analysis and implementation.
- Modeling mistakes are not easily caught. Analysis does not negate testing.
- Do not cut corners on modeling or testing.
- Programs should insist that the analysts document their methodology and assumptions, and compare them against the actual hardware so that errors may be found.
- Do not rely on heritage designs until their flight experiences are thoroughly understood.

For more technical information, call David Gilmore at (310) 336-1897.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Control Propellant Balance

The Problem:

Dynamic instability caused by fluid imbalance has afflicted several satellites during orbit transfer maneuvers. Example include:

- A commercial communication satellite was stranded in a low orbit, and had to expend significant fuel in hundreds of thruster firings to reach a geosynchronous orbit.
- A foreign satellite failed to reach geostationary orbit.
- A military communication satellite wobbled unexpectedly (but was able to recover).

The Cause:

Propulsion control is a delicate task because many parameters, such as the flow rate of propellant in space, cannot be precisely modeled or controlled.

Several factors can trigger fluid imbalance:

- Improper fuel-load procedures. (This problem caused the first incident cited above).
- Differences in flow rates or valve responses can cause propellant to be drawn preferentially from one tank over another. (This problem probably caused the second mishap).

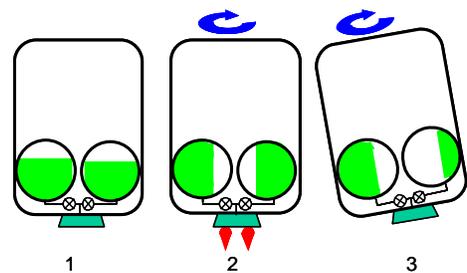
If one tank is cooler than the other, propellant will flow into the cooler tank from the warmer tank, causing imbalance.

Lessons Learned:

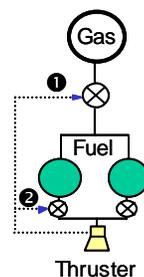
- Make sure tank loads are balanced.
- Use a single tank, if feasible, to avoid propellant migration.
- Ensure that attitude-control algorithms and mechanisms can correct dynamic instability caused by propellant imbalance.
- If possible, place a gas pressure regulator above the tanks, or latching isolation valves below each tank, to control propellant flow.

For more technical information, call Mark Mueller at (310) 336-5081.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



As satellites spin during transfer maneuvers, mass imbalances coupled with centrifugal forces can cause tilting. Severe tilt can divert the transfer thrust and prevent satellites from reaching their proper orbit.



Feedback loops can be designed to control gas pressure (1) or fuel flow (2) between the tanks to restore balance. The latter method is more precise.

28

Graphite/Epoxy Structures Are Easily Damaged by Processing Changes and Handling Mishaps

The Problem:

Two failures involving graphite/epoxy pressure vessels occurred recently:

- A launch vehicle crashed when one of its solid boosters ruptured.
- Two solid-rocket segments failed during hydroproof testing.

The Cause:

Graphite/epoxy composites are used for trusses, pressure vessels (such as nickel-hydrogen batteries and motor cases), and many other applications. Composite technology is relatively new. Minor variations in fiber, resin, and processing can dramatically affect product performance. Quality assurance is vital, yet difficult to achieve.

Graphite/epoxy pressure vessels, especially those incorporating high-strength fibers, are easily damaged. The launch failure was attributed to a handling mishap such as uneven lifting or an inadvertent impact. Unfortunately, damages are not readily detected—existing nondestructive testing procedures, based on ultrasonic scanning, is cumbersome and not 100-percent effective.

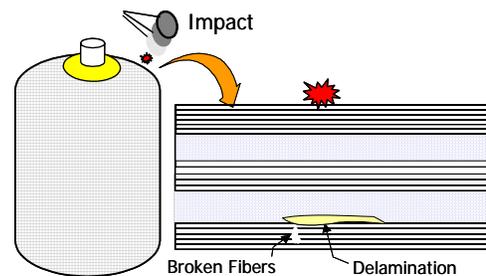
The rocket segment failures took place after the contractor altered materials to meet environmental regulation requirements and made several innocuous changes in the manufacturing processes. Although limited laboratory tests were satisfactory, the fibers wrinkled during winding, greatly reducing the composite's burst strength.

Lessons Learned:

- Protect graphite/epoxy pressure vessels from handling damages.
- Insist on safety margins and quality inspections for composite structures.
- Perform extensive requalification and acceptance tests to guard against subtle processing changes.

For more technical information, call S. R. Lin at (310) 336-7697.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



In addition to graphite/epoxy, Kevlar/epoxy structures are also easily damaged. In both cases, external impact usually leads to damage on the *inside* and can be difficult to detect.

Validate Changes in Command Script Configuration

The Problem:

Contact with a deep space observatory was lost (control was regained three months later following a dramatic rescue; see Lesson 30).

The Cause:

The spacecraft used three gyros:

- Gyro A, to control the safe mode;
- Gyro B, to detect faults; and
- Gyro C, for normal attitude control.

The flight software should turn on the normally off Gyro A when the satellite entered safe mode. Unfortunately, the engineer making a command procedure change did not know to implement the enable command. A loose change-control process failed to catch the error.

During a routine operation, Gyro B was accidentally set incorrectly, causing a false reading. The on-board computer detected B's error and put the satellite in safe mode. The fault on B was fixed, but control shifted from C to A.

Sensed rates from Gyro A (despun, reading zero) and B (active with variable readings) soon diverged, prompting the thruster to fire to try to null the nonexistent roll error. The effort was futile, and the satellite entered safe mode again two hours later.

The spacecraft was designed to survive in safe mode for at least 48 hours. Nonetheless, the operators did not pause to analyze why one anomaly followed on the heels of another. Side-stepping the required telemetry data check that would have indicated that Gyro A was in fact off, the operators mistook Gyro B's variable readings as a sign of a fault, and turned it off. With no functional gyro, control was soon lost.

Lessons Learned:

- Treat command-procedure changes with the same rigor as flight-critical software. This includes formal configuration management, peer review with knowledgeable technical personnel, and full command verification with an up-to-date simulator.
- Ensure change implementation timelines are consistent with staff workloads.
- Display spacecraft health and safety information clearly.
- Follow validated operations procedures, including review of all pertinent data.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The Lagrange Points

There are five Lagrange Points where gravitational attractions from the Sun and Earth balance each other. The loss of control occurred at the first Lagrange Point (L1, about 1.5 million kilometers from Earth), from which location the space observatory monitors solar activities. The L2 point, on the night side, is suitable for infrared astronomy.

Maximize On-board Reprogrammability To Enable Fault Recovery

The Event:

An observatory lost in deep space (Lesson 29) was brought back to life following three months of clever troubleshooting.

The Cause:

The salvage team faced daunting challenges. Following the loss of attitude control, the satellite's heaters had shut down, its batteries were drained, and its fuel had frozen. Insufficient bus power made it impossible to sustain a downlink long enough for the ground station to lock on, and rescuers were not even sure exactly which communication frequency would work.

The team hit upon the idea of borrowing the world's largest radar to transmit to the spacecraft, and using another big dish to receive return signals. They set up a special wideband analyzer over the Internet so that the downlink signal could be analyzed instantly.

The shot in the dark paid off—a faint heart-beat was received from the lost satellite. Only the carrier signal came, however, because the on-board receivers could not lock onto the uplink signal.

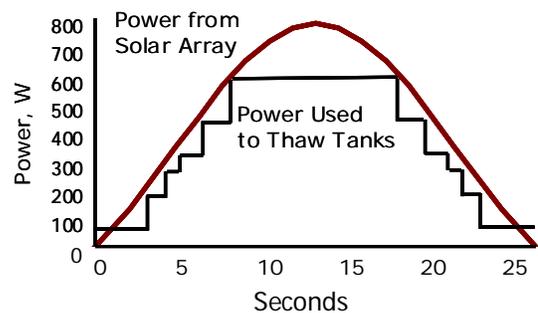
Ingenious commands, together with efficient power management, eventually brought the bus voltage up to 28 V, permitting controllers to monitor spacecraft status and thaw the propulsion system. An intricate attitude recovery maneuver was devised to allow the satellite to reacquire the Sun, and normal operations resumed. Remarkably, despite having been alternatively exposed to extremes of -120° and 100°C , all instruments survived!

Lessons Learned:

- Design into the satellite the flexibility to handle unforeseen emergencies, and provide emergency reset capability for major components.
- Add emergency protection of a satellite battery system, such as low-battery-voltage cut-out of nonessential loads.

For more technical information, call Julie White at (310) 416-7229.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Power-Efficient Thawing of the Hydrazine Tank

The fuel tank had to be warmed up before pipes and thrusters were, lest overpressure burst the lines.

Software changes allowed the battery to discharge current like a thermistor and turn on selective heaters whenever power became available. Because the flight computer was off during battery charging, the software patch had to be reloaded each time.

After fine-tuning, controllers managed to thaw the tanks with 48 heaters, using a peak power of over 500 watts!

31

Oxidation Can Cause Erratic Open Circuits In Solid State Devices

The Problem:

Several photodetector chips developed intermittent open (high resistance) circuits during integration.

The Cause:

This anomaly baffled experts because the chips, when returned to the foundry, often passed diagnostic tests. Also, investigators could find no mechanical defects (such as fractures) that might account for the open circuits.

An in-depth study revealed that the anomaly resulted from oxidation of the titanium diffusion barrier under the gold signal line. Titanium oxide can “switch” (jumping between conducting and insulating states) causing the circuits to open erratically.

The subtle flaw was caused by manufacturing imperfections that exposed the titanium layer to oxidation. The defect was not caught by the chip maker because the oxidation developed very slowly.

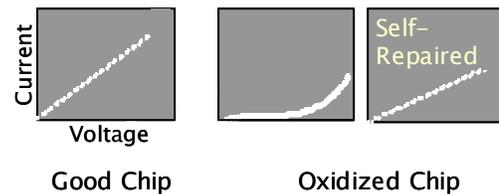
Other metals are susceptible to this problem. Oxidation of lead created excessive noise in a lead sulfide detector. Oxidation of nickel made some devices oversensitive to applied voltage or even shock. Nonlinear voltage-current behaviors were the cause in each case.

Lessons Learned:

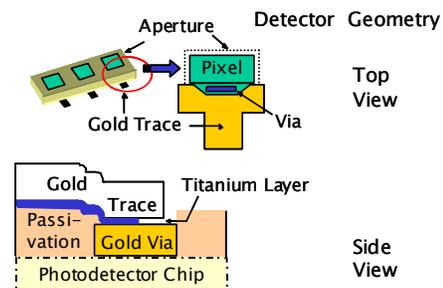
- Protect sensitive metal layers from oxidation (caused by over-etching, for example) during semiconductor fabrication.
- Use current-voltage profiles as a diagnostic tool—nonlinear high resistance usually indicates oxidation.

For more technical information, call Alfred Fote at (310) 336-6926.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



An applied voltage can sometimes heal the chips temporarily by pushing the oxide layer aside.



Incomplete coverage of the gold via by the trace exposed the titanium layer to inadvertent oxidation.

32

One Operation, One Verification

The Problem:

A prototype reusable rocket crashed because a technician forgot to reconnect a helium line.

The Cause:

The goal of the project was to demonstrate rapid turnarounds between vertical takeoffs and landings. A streamlined management approach kept paperwork to a minimum. A working vehicle was built in 18 months; a modified version had already flown three times before the incident.

The flyer was supported with four legs that were actuated by an on-board helium supply. During preflight preparation, each leg was deployed once so the control center could verify its deployment monitors. The helium line was then disconnected to vent the actuator, the legs stowed, and the helium line reconnected. Four technicians repeated this procedure on each leg.

Unfortunately, a technician forgot to reattach one helium line. The error was not detected because there was no procedure to check the integrity of the system after disconnection and reconnection. At landing, the leg failed to deploy, whereupon the vehicle toppled and exploded.

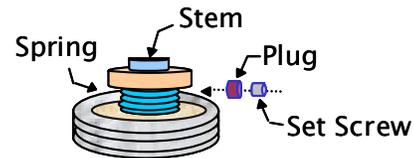
The investigators found that procedures were neither well developed nor rigorously applied. Operators and technicians used the procedures as guidelines instead of checklists. In fact, failure to reconnect happened once before. Although caught, the incident was not documented.

Lessons Learned:

- Implement a discrete verification step for each critical task.
- Avoid multiple tasks within a procedure (see Lesson 12).
- Ensure a fail-safe process by applying software technology, self-checking indicators, or positive feedback mechanisms to complex operations vulnerable to human errors.
- Document each near miss and correct its root cause.

For more technical information, call Ron Williamson at (310) 336-2149.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



A Similar Incident: Failure Caused by a Loose Screw

The precision regulator in a booster engine control system used a stem screw to modulate gas inlet. A set screw forced a nylon plug against the stem screw threads and prevented the stem from rotating.

The regulator was reworked to repair leakage during build. The rework instruction did not explicitly require set screw retorquing and verification. The loose set screw caused the stem screw to unseat. The launch failed.

Check Satellite-Launcher Compatibility As Early As Possible

The Problem:

A technology demonstrator satellite had to be substantially redesigned because the vehicle's stability during the orbit-transfer maneuver was not considered early on.

The Cause:

When a satellite spins, its components vibrate at a “nutation frequency” determined by the moments of inertia and by the spin rate. Flexible parts, such as whip antennas and fluids, will dissipate the rotational energy, particularly if these parts resonate near the nutation frequency. Energy dissipation may lead to increased coning angles, even a flat spin.

Nutational growth caused several early satellites to malfunction. Although well understood in general today, it remains a challenge whenever spinning upper stages are used—because fuel motion and burning complicate the analysis, the satellite should be designed with extra margins to prevent the stack from entering a flat spin during orbit transfer.

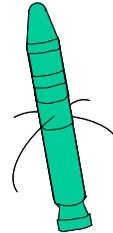
The upper stage selected by this program spins. Unfortunately, the contractor failed to pay attention to the issue during preliminary design, despite advice from experts. The instability could have been mitigated by simply modifying the satellite propellant tanks. However, because the problem was recognized late, numerous costly modifications became necessary. The project was almost cancelled.

Lesson Learned:

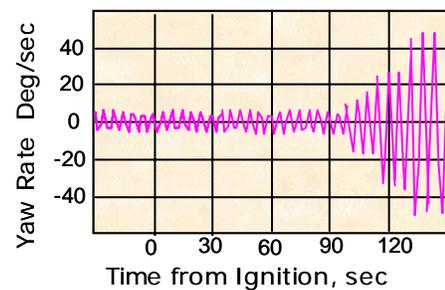
- Ensure interface problems between the satellite and launcher, such as dynamic instability, are analyzed early on in the design process (see Lessons 2, 11).

For more technical information, call David Stampleman at (310) 336-2243.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The first American satellite, Explorer 1, went into a flat spin because its flexible antennas triggered nutational growth.



As shown here, solid upper stages, which this mission used, are more prone to instability. The satellite contractor did not recognize this risk in part because the launch vehicle contractor failed to formally communicate this requirement. The design changes kept the instability in check during flight, and the satellite reached the correct orbit.

34

Safeguard Hardware Against Inadvertent Overtesting (II)

The Problem:

A satellite launch had to be postponed by several months because an antenna panel delaminated.

The Cause:

The antenna assembly, based on a honeycomb sandwich structure, was undergoing a thermal vacuum test. An operator set the heater voltage too high, causing the panel to be subjected to 100-deg C instead of the planned 61-deg C.

Similar overheating problems had occurred before at this facility, and an automated temperature limiter or alarm on the test equipment would have averted the mishap. However, motivation to invest in facilities or training was low because the program was coming to an end.

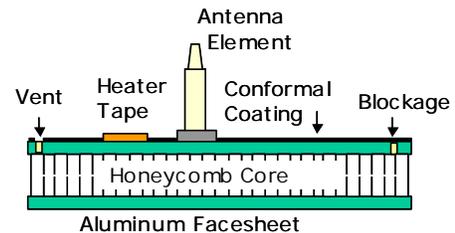
Overheating prompted pressure to build up within the sandwich cells. Unfortunately, four of five venting holes in the facesheet were inadvertently blocked by conformal coating because the operators were not provided with clear assembly instructions. The trapped pressure caused the panel to rupture.

Lessons Learned:

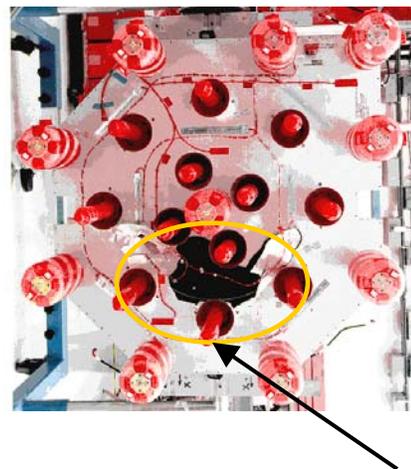
- Implement overtest protection (see Lesson 24).
- Correct the root cause of operational mistakes.
- Incorporate visual guides or overlays as part of process control procedures.
- Honeycomb sandwich structures for space structures should be vented. Otherwise, when heated, trapped air and moisture can expand, creating pressure and causing delamination (Lesson 1).

For more technical information, call Susan Ruth at (310) 336-6765.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Antenna Construction



Delamination Area (1' x 2')

35

Implement Independent Fault Protection

The Problem:

A deep-space mission ended prematurely after excessive thruster firing depleted its fuel.

The Cause:

This spacecraft was developed by a highly motivated group operating under a rigid cost cap and tight schedule. Flying just 22 months after being funded, it successfully circled the moon and demonstrated many technologies.

Soon afterward, however, a maneuver triggered a numeric overflow in the processor, causing it to erroneously fire its thrusters and freeze. A “watchdog timer” algorithm should have stopped the thrusters from continuously firing, but did not execute because the computer had already crashed. By the time ground operators regained control, all the fuel was gone.

A hard-wired timer, which would have stopped thruster firing, was not implemented due to the tight schedule. Time pressure also prevented the software from being fully tested, and many changes had to be uploaded as faults were discovered.

The overflow error had occurred thousand of times (without causing malfunctions) because the project had to settle for an inadequate but available processor. Software changes had been written to correct the problem, but the overstretched staff could not handle operations, anomaly analysis, and software repair at the same time, and the change was not loaded.

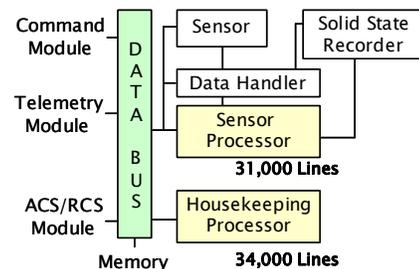
Four years later, another interplanetary probe encountered a similar anomaly. Fortunately, engineers learned the lessons from the previous incident; the precautions they took allowed them to successfully complete the mission (see Lesson 36).

Lessons Learned:

- Apply independent fault protection for critical software functions.
- Implement exception handling to protect the flight processor from aborts due to data handling errors (see Lesson 18).
- Do not cut corners in testing critical flight software.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



A Rushed Job

Over 65,000 lines of flight code (only 20% inherited) were developed in 17 increments within one year, leaving little time for thorough testing.

36

Implement Independent Fault Protection (II)

The Event:

An interplanetary probe recovered from a major anomaly.

The Cause:

The spacecraft, designed to rendezvous with an asteroid, employed extensive autonomy because ground intervention during an emergency would take too long. The designers studied the history of an earlier project, which terminated prematurely after a data error depleted on-board fuel (see Lesson 35).

Three years into the flight, an engine burn aborted. A missing command in the burn-abort contingency command script prevented a graceful transition into the safe mode, and a series of anomalies ensued. Communication was lost for 27 hours before the flight computer regained control.

The initial script error was not caught during software tests. Hardware-in-the-loop simulation could not test abort scenarios because the brassboards were difficult to use. Exactly how the anomalies propagated is unclear because a bus undervoltage wiped out data from the recorder, nor could the anomalous behaviors be reproduced on ground.

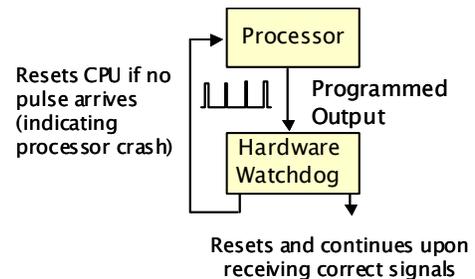
During the emergency, the spacecraft fired its thrusters thousands of times. Fortunately, the fuel loss was tolerable because the thrusters were hardwired to fire only for fractions of a second. The mission was saved because the designers took precaution against fuel depletion during a software crash, a lesson learned from the previous failure.

Lessons Learned:

- Create extensive, realistic nominal and anomalous operational scenarios for testing at every level, from unit through system test.
- Implement robust simulators, including hardware-in-the-loop, for testing critical flight software functions.
- Apply independent fault protection, such as hardware watchdogs, to mitigate risk in real-time systems, where errors can be so deeply buried as to be practically undetectable.

For more technical information, call Richard Adams at (310) 336-2907.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Watchdog Scheme (Simplified)

The processor feeds a series of programmed pulses into the hardware timer, which will reset itself and await the next input. If the expected “heartbeat” does not arrive, the watchdog knows that the processor has probably crashed and intervenes (such as by initiating a fault protection routine).

Aim for Realistic Schedules in Development Projects

The Problem:

A sophisticated instrument was delivered five years behind schedule.

The Cause:

Combining three previously separate sensors and aiming for greater sensitivity, this instrument densely packed together diverse technologies. The developer contracted for delivery in three years, even though two heritage systems each took eight years to build.

Soon after program start, the spacecraft prime contractor issued unexpectedly stringent interface requirements. The preliminary instrument design had to be substantially altered to meet new weight, volume, and vibration constraints.

More features (such as stiffer structures) had to be added, but design flexibility was limited due to volume constraints. In compensation, cutting-edge electronics had to be deployed, but the vendors could not deliver them on schedule due to manufacturing difficulties.

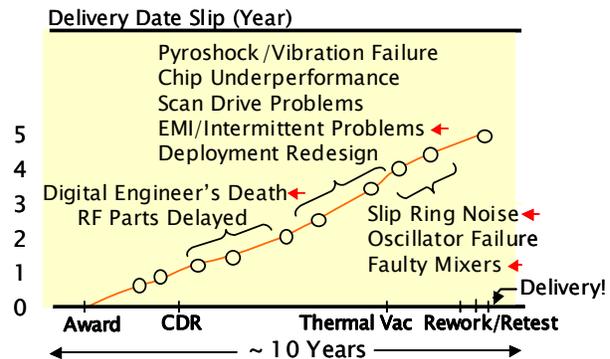
The contractor adopted first-pass-success schedules—the design went into manufacturing directly, skipping prototyping. Problems surfaced late (such as during thermal vacuum testing), and were discovered sequentially. Despite the contractor’s heroic effort, it took eight years before the product was delivered.

Lessons Learned:

- Provide a detailed interface specification as early as possible.
- Foster a cooperative working arrangement among contractors and proactively maintain realistic power, weight, and volume reserves.
- Create engineering models so that problems can be discovered early.

For more technical information, call Alfred Fote at (310) 336-6926.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Slim margins, unproven technology, tight schedules, and fixed cost conspired to incrementally push the delivery date.

Items marked with arrows each impacted the schedule by between 9 and 18 months.

38

Do Not Ignore Unexplained Test Anomalies

The Problem:

A power regulator had to be pulled from a spacecraft.

The Cause:

A new block of satellites, extensively upgraded in its power system, exhibited several unusual anomalies during system testing. Although the contractor managed to work around the anomalies, the program office was uncomfortable with the design robustness and requested an independent analysis.

A preliminary independent simulation did not find anything odd. The analyst continued to refine the model without spotting a “smoking gun” that would account for the problems, and most people became skeptical as to whether a problem in fact existed.

The analyst’s persistence paid off, however, when he found a circuit instability that induced the glitches. Moreover, the design flaw would have caused the solar array voltage to oscillate when the satellite exited from eclipse, overstressing power components and triggering immediate mission loss. The contractor verified the subtle flaw and pulled the already integrated unit from the satellite, averting a disaster.

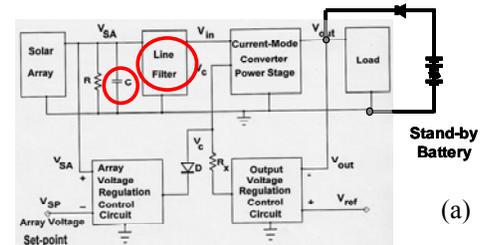
The problem was not found earlier because the test configuration was not sufficiently flight-like—resonance was quickly damped out by a one-ohm dummy load resistor. On orbit, the solar array’s high impedance would have made it impossible to keep the resonance in check.

Lessons Learned:

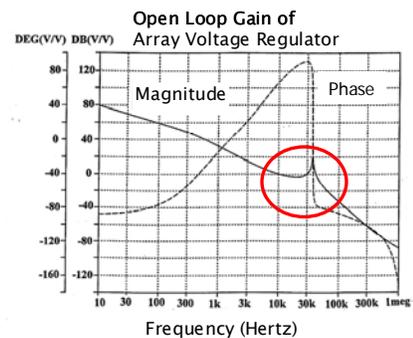
- Test under all operating conditions—not only sunlight and eclipse operation, but transitions, safe-hold mode, loadshed mode, and recovery mode.
- Strive to understand implications of test anomalies.
- Ensure perceptive instrumentation, lest test-set glitches cast doubt on results.
- Minor design changes in power supplies can result in disastrous consequences. Double-check design changes, and perform independent analysis where practical.

For more technical information, call Kasemsan Siri at (310) 336-2931.

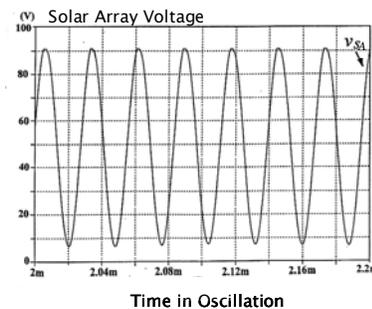
For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



(a)



(b)



(c)

The line filter and feed-through capacitor (a) combined to resonate at a “crossover frequency” (b). The array would suffer sustained oscillation (c) and fail.

39

Thoroughly Review Test Data for Early Indicators of Anomalies

The Problem:

A satellite lost attitude control when defective heater circuitry caused a fuel line rupture.

The Cause:

Shortly after launch, the propellant began to freeze. After a few days on orbit, repeated freeze/thaw cycles fractured a line, and all propellants were lost.

A review of the test record revealed that a heater had ceased functioning on the ground, but the defect was not noticed.

Cutbacks in ground support prevented continuous satellite monitoring during early operations. The anomalous thruster temperatures were recognized several days too late by controllers. If the problem had been spotted earlier, the satellite could have been saved by firing the thruster before its line froze.

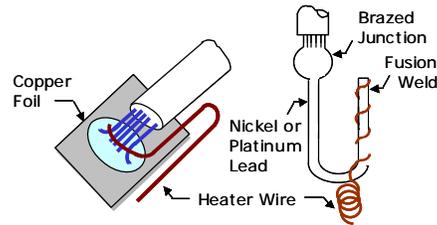
In the wake of this failure, numerous design and operation changes were implemented, and the propulsion thermal control system on all subsequent flights performed successfully.

Lessons Learned:

- Carefully inspect all test and operational data for trends, oddities, and “out-of-family” values, even when all values are within preset limits. Evaluate all indicators for potential impacts, should trends continue. Seek to explain all instances of anomalous data (see Lesson 19).
- Make sure that experienced operators closely monitor the satellite’s health during early operations.
- Provide ground-commandable back-up heaters.
- Install heaters to fill/drain lines, and provide temperature monitors for all propellant lines and valves.

For more technical information, call Lori Crosse at (310) 336-5821.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Damage of the wiring at the heater lead (left) probably caused this failure. A more robust configuration (right) was used in all subsequent flights.

Failure Indicator Available But Missed

| | Thruster Temperature (°F) |
|------------------------------|---------------------------|
| Normal | 100+ |
| During First System Test | 104 |
| During Four Subsequent Tests | 77-83 |

Although the heater failed during early ground tests, the problem was not recognized because temperature limit checks were set to accommodate test environment changes, not to verify heater performance.

Later tests and operations used computer-controlled stepwise limit checks to highlight anomalous behaviors early.

40

Avoid Radio Frequency Interference

The Problem:

Signals from one program inadvertently interfered with another program.

The Cause:

This project, driven by a unique requirement, provides radio frequency (RF) intersatellite links among its fleet.

The RF crosslinks were originally designed to null toward Earth to prevent appreciable amounts of emission from reaching the ground.

Over the years, however, the original requirement was forgotten, and the next generation of satellites no longer nulled toward Earth.

At a conference, an analyst fortuitously noticed that another program's downlinks used the same frequency band. A quick calculation showed that this program would suffer interference from the crosslinks. The impact is minimal at present, but will increase as crosslinks on the first program multiply.

The remaining satellites on the ground had to be reengineered to reduce leakage toward the ground.

Lessons Learned:

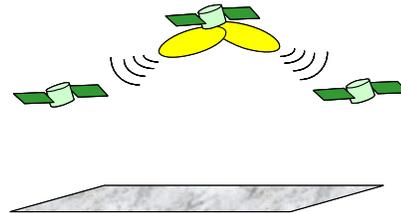
- Understand why requirements exist in legacy designs before discarding them.
- Coordinate spectrum planning with authorities (for example, Manager of Spectrum Allocation at the Space Command), because not all frequency usages are public information.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Emission from crosslinks can reach Earth and interfere with other users.



The emission problem can be cured by phasing the signals in the array to place a null toward Earth.



41

Carefully Consider the Implication of Test Failures Beyond the Narrow Issues at Hand

The Problem:

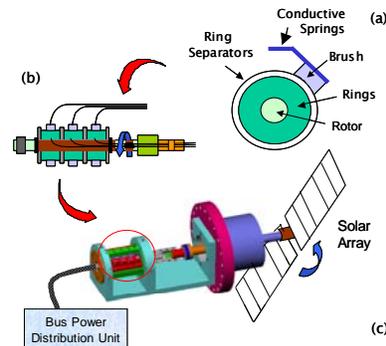
The electric power system on a satellite suddenly failed.

The Cause:

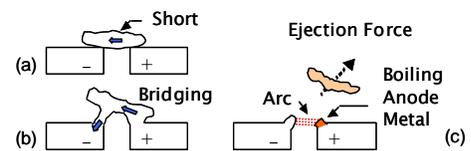
The slip rings, wired with opposite polarities on adjacent brushes and therefore prone to arcing, were destroyed after debris induced a short.

Several mistakes led to the faulty design:

- Having chosen a bus with an excellent flight history, the program focused virtually exclusively on the payload. The bus in fact had to be extensively modified—rotating arrays, for example, were put on the aft end of the satellite for the first time, requiring new array drive electronics. Yet, the program was too firmly set in the idea of a standard bus to grasp the risks.
- The slip ring design provided practically no internal clearance between adjacent brushes, making it apt for debris to cause a short. The design was accepted because another project had flown it.
- The other project, however, had rewired the rings to keep the same polarities next to each other after encountering a short during launch-simulating vibration tests. Notified of the change, the first program felt that the change did not apply because its slip rings were unpowered during launch.
- Slip ring arcing was also observed during ground test of a control moment gyro by the same contractor working on yet another project. Unaware of this incident, the designers did not consider shorting in the reliability analysis or in part selection. The program also deleted thermal vacuum test of the slip rings to save money.



Slip rings connect rotating solar arrays to the bus.



Shorting of slip rings is fairly common—improperly lubricated brushes can easily abrade conductive slivers out of the rings. The voltage gap across adjacent brushes exacerbated shorting by triggering an arc, which wrecked every anode in its path.

Lessons Learned:

- Thoroughly evaluate the heritage and applicability of using “existing” or “flight-proven” equipment, especially if modifications have been made.
- Include shorting in analyzing potential failure modes of power systems.
- Apply manufacturing and handling practices that minimize slip ring damage.

For more technical information, call Jeff Lince at (310) 336-4464.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

42

Account for Electrostatic Interaction in Structural Analysis

The Problem:

The performance of a communication satellite significantly degraded.

The Cause:

The satellite deployed a new phased-array antenna, consisting of multiple microstrip elements made of copper circuits over dielectrics. A large thermal blanket, used for the first time on this type of antenna, shielded the elements from the Sun.

The sunshield was not adequately supported—too few tensioners were provided to keep the blanket taut under Earth's gravity (1 G). The sunshield was installed loosely, often touching the antenna elements. Nevertheless, no attempt was made to compare antenna performance before and after blanket installation on ground, because the cover was expected to recover from drooping once in orbit.

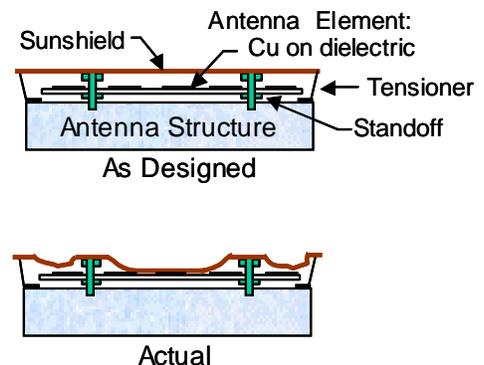
Unfortunately, an electrostatic charge built up in the ungrounded dielectrics of the antenna. The resulting electrostatic attraction overpowered the insufficiently applied tension, keeping part of the blanket in contact with the elements. The phased-array's gain degraded due to dielectric coupling and shorting to the conductive layer of the sunshield.

Lessons Learned:

- Be aware of the propensity of dielectrics to pick up an electrostatic charge in space.
- Thoroughly review the potential impacts of the space environment on flight hardware.
- Whenever possible, a design's operation in space (0 G) should be designed to be verifiable under 1 G test conditions.
- Test the entire system in the final flight configuration.

For more technical information, call Harry Koons at (310) 336-6519.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The sunshield curled toward the antenna due to charges that accumulated in the insulators. Notice that electrostatic attraction can take place even though one surface (the sunshield in this case) is grounded.

43

Do Not Circumvent Processes Designed to Catch Human Errors

The Problem:

A satellite was placed into a moderately degraded orbit.

The Cause:

During launch preparations, operators made final measurements of the spacecraft's inertial measurement unit (IMU). The readings, together with factory calibration data, were used to control the satellite's orientation during ascent.

Unlike all the other inputs loaded to the satellites, the IMU measurement and calibration data could not be verified in a testbed because the readings had to be made just before launch. Therefore, a procedure was set forth to avert mistakes: one operator was required to transcribe the calibrations numbers from the factory printout, another would verify the entries.

An engineer supervising the keyboard operators copied the calibration data from the computer printout onto a scratch paper, leaving the original printout in his office. He gave the scratch paper to the operators, telling them that it was suitable. The data were typed in and verified.

Unfortunately, the engineer left out a symbol, and the orbit insertion went awry!

Lessons Learned:

- Ascertain software databases as thoroughly as the source codes (see Lesson 3).
- Verify software algorithm and database on a simulator whenever possible.
- Double-check manually entered data against original sources.
- Automate data transfer and checking whenever possible to minimize human error.

For more technical information, call Julio Rivera at (310) 336-3287.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

\dot{R}_n versus \bar{R}_n

The First Software-Related Crash

An incorrect formula in the ground software led to the failure of Mariner I in 1962.

Ascent control required velocity smoothing, or "R dot bar n" where R stood for radius from a tracking antenna, the dot for the first derivative (i.e., the velocity), the bar for averaging, and n for the increment.

The bar was left out of the handwritten equations provided to the programmer, causing the guidance computer to be coded to process raw velocity instead. Confronted by fluctuating telemetry, the computer sent erratic correction signals, forcing a smoothly ascending booster to veer off course.

44

Beware of Sneak Paths Through Test Equipment

The Problem:

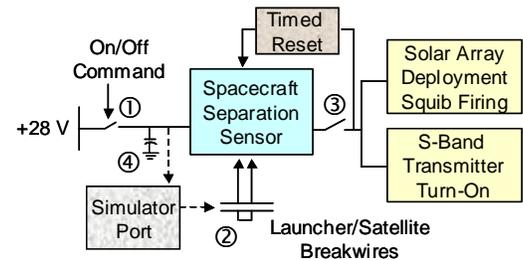
Two days before launch, a satellite spontaneously tried to deploy.

The Cause:

Baffled engineers found that the separation sensor unexpectedly powered up. Even then, it should not have turned on. Unexplained internal flaws inside the unit, which had operated nominally up to that day, threatened to scrub the mission.

Not wanting to spend millions of dollars to return the satellite to the factory, the program sought help from an outside expert, who found:

- The functional test was unable to detect whether the power relay was open or not.
- The test set inadvertently enabled the sensor, as if the breakwire had opened.
- The sensor could turn on only if powered quickly.
- The anomaly first occurred when the bus was powered up too fast by mistake, but appeared again after the power was properly reapplied.



Simplified Separation Electronics Schematics

A latch in the separation sensor (powered via relay ①) opens after the satellite breaks away from the launcher (②), deploying the solar array via relay ③.

Failure of relay ①, due to the addition of a filter ④, formed a sneak path (dashed line) via the simulator port, triggering the prelaunch anomaly. Premature separation in fact could not occur in flight because the port is not used.

The analyst traced the anomaly to a noise filter added to the input line. The filter caused an overcurrent, welding the relay shut and powering the sensor up. Welding in fact occurred on a relay installed in this same spot once before, but no corrective action was taken.

Energizing the bus too fast during ground test created a current strong enough to turn on the sensor and start the deployment sequence. After an abort, the problem recurred upon a nominal restart because the sensor timer had not yet reset.

Once understood, the concern vanished—the relay would be closed in flight and the sneak path would be blocked by the flight plug. The satellite flew successfully.

Lessons Learned:

- Determine and correct the root cause of all failures.
- Trace the flow of power and signals from source to load during troubleshooting.
- Provide a mechanism to independently validate the status of critical components.
- Inject unexpected conditions (such as a closed relay, current surge, and sluggish separation wire breakage) during reliability analysis to discover lurking failure paths.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Guard Against Chloride Contamination Due to Manufacturing Process Changes

The Problem:

Two heat pipes suffered significant performance degradation in system-level test.

The Cause:

Analysis of the failed units revealed particulate materials, hydrogen gas, and internal etching. Obviously, the ammonia working fluid had reacted with the aluminum tubing—a problem that had not occurred in recent memory.

The problem was eventually traced to a minor manufacturing procedure change. After machining, the vendor previously wrapped the end of the tubing with aluminum foil to keep dust out. It replaced this untidy-looking procedure with dust caps. Apparently the tubing scratched the common polyvinyl chloride (PVC) caps, lodging some debris in the assembly.

Unfortunately, chloride in the PVC catalyzed ammonia's decomposition. The entire batch of heat pipes had to be removed.

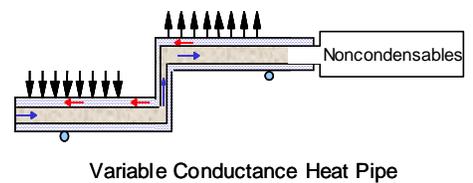
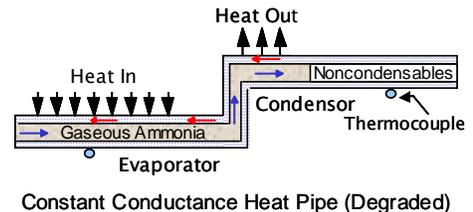
The impact of this procedure change was not evident to the manufacturer. The caps were first used on a batch of variable conductance heat pipes instead of the more common constant conductance type, and the variable conductance mode masked the noncondensable problem. Moreover, this batch passed vendor acceptance test because the test was made within two days of ammonia charge, before noncondensables had a chance to build up.

Lessons Learned:

- Heat pipes are highly sensitive to minor materials and process changes.
- Seemingly minor process alterations can have catastrophic side effects.
- Allow sufficient time before conducting tests of chemical degradation.

For more technical information, call Robert Prager at (310) 336-5582.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Noncondensables diminish heat rejection efficiencies of constant conductance heat pipes.

A Similar Incident

An engine suffered severe leak during recent ignition testing because the chamber was cleaned with over-the-counter detergent. Chloride in the cleaner induced stress corrosion, cracking the tubes.

46

Make Sure Test Equipment Is Sufficiently Capable

The Problem:

A power regulation unit underwent five months of acceptance tests due to an inefficient setup.

The Cause:

The unit under test, consisting of eight DC-DC power stages, exhibited major glitches during vibration. Based on sketchy data, the manufacturer assumed a short had occurred in the output stage, and replaced all suspected parts.

The same anomalies recurred during a second vibration test. Now the vendor believed that the first power stage was at fault.

An independent simulation showed that neither scenario was credible, and it was recommended that full instrumentation as well as computerized data collection be implemented. The manufacturer did not do this.

Four more rounds of vibration failures ensued. Not only did the root cause remain elusive, the equipment's vibration life was almost depleted. The manufacturer proposed to replace the entire first power stage, which would have seriously impacted the program schedule.

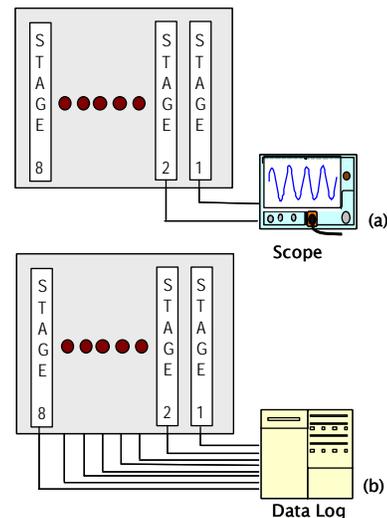
Exasperated, the program office made the manufacturer run one more test with full instrumentation. Right away, an insidious short in the current sensor was found. Within a few weeks, the repaired unit passed.

Lesson Learned:

- Budget for high fidelity, reproducible, functional tests to facilitate troubleshooting.

For more technical information, call Dave Caldwell at (310) 336-6344.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Troubleshooting was hampered because the test set (a) could not monitor all channels. Also, the reliance on oscilloscopes made data collection inefficient. Digital data collection from all ports (b) solved the problem in a few days.

Housekeeping (as opposed to hardware-related) glitches in facility, software, equipment, or connectors routinely account for the majority of discrepancy reports, unnecessarily impacting program schedule.

47

Review Hardware Reusability When Configuration Changes Affect Margins

The Problem:

A satellite failed two weeks after launch when a battery charger shorted.

The Cause:

The short took place between the grounded radiator and the electronics-mounting heatsink that was at the solar array potential.

The radiator was isolated from the heatsink by thin adhesive and anodization layers only. A tolerance buildup, after repeated temperature excursions, drove the mounting screws through the anodization, causing the short. Conductive debris could also have bridged the heatsink to the box's walls.

Several factors contributed to the mistake:

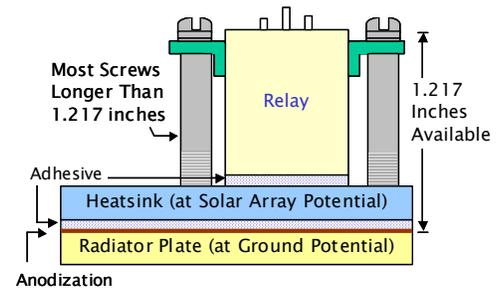
1. The charger had flown on many spacecraft over 20 years. Far from robust, the units were handled meticulously in the past. The failed box, on the other hand, was treated routinely, and not thoroughly inspected.
2. Two scientific instruments added to this mission caused the system to run 10 deg C hotter, exacerbating the tolerance problem by, for example, flexing the box walls. Unfortunately, the units were not requalified.
3. The survival mode software, which could have shed the load and provided time to diagnose the problem before the spacecraft batteries were depleted, was not enabled.

Lessons Learned:

- Recognize that workmanship plays a large role in the space hardware, and reliability may be compromised when undertrained personnel assemble heritage equipment.
- Computerize manufacturability analysis, including interface tolerance buildup, dynamic interference, and ease of inspection on all packaging designs.
- Provide automatic fault management mechanisms so that a single defect will not bring down the entire system.

For more technical information, call Robert Tsutsui at (310) 336-3273.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



A Vulnerable Packaging Design

An inspection of the hardware destined for the next flight revealed that many screws were too long to fit into the space between the relay mount and the radiator plate, making a short virtually inevitable.

Moreover, the heatsink barely cleared the unit walls. Because the heatsink was not conformally coated, debris such as a loose solder ball could also have caused a short.

48

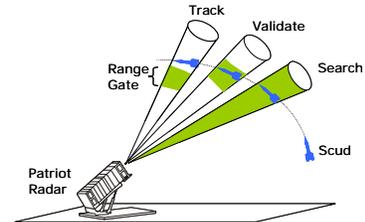
Thoroughly Reverify Software When Requirements Change

The Problem:

The Patriot defense system failed to intercept a Scud missile.

The Cause:

As the Patriot detects a threat, its radar beam narrows for better tracking. The fire controller extrapolates the trajectory (the position where the object should appear next) to commence locking on the target. Trajectory calculations require knowledge of time. Time is updated in the system clock every tenth of a second. A pair of 24-bit integers (31 x 0.1 sec, 32 x 0.1 sec, and so on) are converted to a floating point number before computation. Because 0.1 cannot be fully expressed in binary digits, it is truncated, with a loss in precision by one part per million.



Cumulative precision loss let the radar look in the wrong place (range gate) for the Scud.

When Patriots were brought to the Gulf, the software was modified to track faster Scuds. A change was made to convert clock-time more accurately, but was not inserted everywhere it was needed in the software. The elapsed time between two radar pulses, which used to be based on two clock readings containing canceling arithmetic errors, now contained a systematic error because the truncated time of one pulse was subtracted from a more accurate time of another pulse.

The Patriots, designed for mobile defense, were expected to shut down for redeployment or maintenance after no more than 14 hours. In the Gulf, they were operated continuously from fixed positions. As the radar clock ticked, the error accumulated. The Army became aware of the drift, modified the software, and alerted the field units to periodically reboot so the clock could start anew.

Unfortunately, the instructions arrived the day after a Scud hit an Army barracks and killed 28 soldiers. The battery protecting the base had been in operation for over 100 consecutive hours, during which the timing inaccuracy had grown to the point where the Patriots could no longer lock on the Scud.

Lessons Learned:

- Reverify software performance when its intended environment changes (Lesson 18).
- Thoroughly analyze the impact of loss of precision.
- Ensure change analysis is complete and changes are comprehensively verified.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

49

Equipment Intended for Use in Simulated Space Environments Should Be Space-Rated

The Problem:

A flight payload was damaged during thermal vacuum testing.

The Cause:

An inspection of the flight hardware showed that the test cable, wrapped in microwave-reflective tapes, suffered corona discharging and overheated.

Flight-qualified cables have safety features, such as built-in vents in the connectors. Unfortunately, neither the test cable nor its connectors were vacuum qualified. Corona started in the connector, and the cable out-gassed. A destructive resonance, known as multipaction breakdown, set in, and ignited parts on the payload.

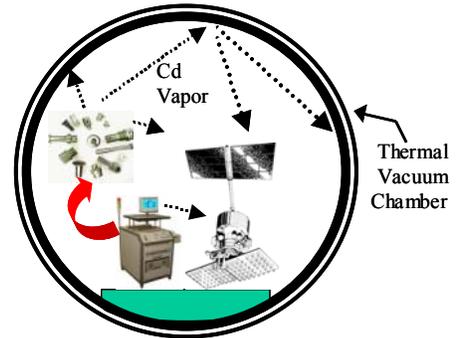
The accident was not caught during thermal vacuum testing operations because no one from the payload supplier monitored the test and because the satellite was not fully instrumented.

Lessons Learned:

- Perform formal design reviews on ground-test equipment intended for use in space-like environments.
- Test radio frequency equipment in vacuum to 6 decibels over the expected input level (to account for unfavorable signal return) to ensure operational safety.
- Monitor flight hardware during test lest overstressing cause damage.
- Improve interfaces between payload engineers and bus engineers, particularly during system level tests.

For more technical information, call Tom Darone at (703) 633-5134.

For comments on the Aerospace Lessons Learned Program including background specifics, call Paul Cheng at (310) 336-8222.



A Similar Incident

A test set scheduled for use in the thermal vacuum chamber contained cadmium-plated parts. Cadmium, commonly used to plate military components, sublimates in vacuum and is not allowed in space. If the test had gone ahead, the cadmium could have contaminated not only the spacecraft being tested, but also the chamber and future satellites!

Virtual Cross-strapping Extends Satellite Life

The Event:

A government satellite, almost deorbited after losing both primary and redundant gimbal control, was brought back to operational status.

The Cause:

A power supply failure in the A-side caused the payload gimbal control to be switched to the B-side. Later, the B-side was disabled when its position sensor malfunctioned. An initial analysis indicated an in-orbit fix was impossible.

An engineer who worked on the original gimbal development was brought in to assist safing the satellite for de-orbit. Drawing on his experience with similar programs and on this gimbal's design, the engineer realized that there was a secondary command path for the gimbal motor, which would make it possible to cross-strap the functioning components of both sides. Calculations showed that the spacecraft's design margins would support this fix without significantly compromising mission status.

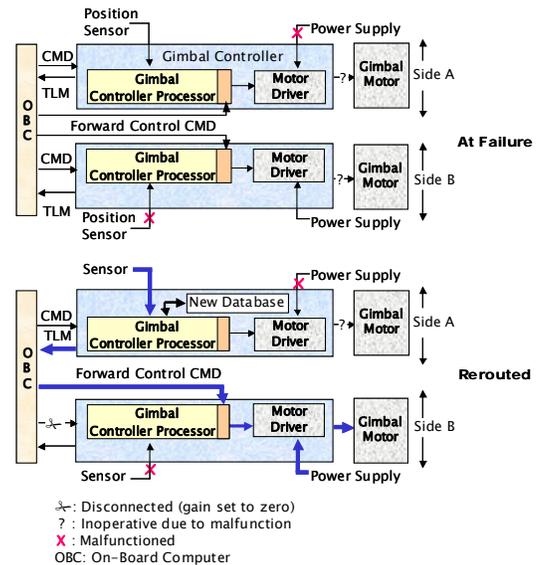
The new control laws were programmed into the controller processor, and the mission was restored.

Lessons Learned:

- On-board reprogrammability provides enormous flexibility (see Lesson 30).
- In a tight spot, seek cross-program wisdom from diverse organizations.
- Capture knowledge of heritage designs and look for novel ways to take advantage of design features.

For more technical information, call Hiroshi Shibata at (310) 336-5036.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Recovery Strategy

The gimbal controller design included a path to forward-control nonlinear motor driver behavior.

The rescue scheme fed commands, derived from sensor A data and calculated by the processor using new control laws, into the motor controller B via this route, bypassing the processor B.

51

Review Troubleshooting Process When Encountering Surprising Test Results

The Problem:

An attitude control unit exhibited unrepeatability performance degradation.

The Cause:

In the middle of the acceptance test, a production unit failed. Engineers could not identify the cause.

Eleven days later, the problem abruptly vanished. An all-out effort, lasting over four months, failed to recreate the anomaly, driving the contractor to consider tearing the unit apart.

It turned out that the unit, slightly modified from a product designed for another project, looked identical to the other except for the part number on the nameplate. Both operated on the same test set and were equipped with identical connectors.

Units for both programs, by chance having the same serial number, were stored in identical carrying cases and stowed side by side in the same storage cabinet. Apparently, a technician had removed the wrong unit from the cabinet to test. During the intensive troubleshooting effort, nobody checked the label of the unit under test!

Lessons Learned:

- Consider using bar codes in production control.
- Incorporate design features, such as colored cables, to preclude human errors.
- Don't overlook simple human errors when confronting unexplained problems.

For more technical information, call Tom Fuhrman at (310) 336-6596.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



A Similar Incident

A thermal vacuum test was delayed because two rolls of Kapton tapes were mixed up.

Both rolls of tape came from the same supplier and looked exactly the same. However, the roll inadvertently used to attach insulation blankets contained an adhesive that was based on silicone instead of on low-outgassing acrylics. The satellite had to be baked and pumped for a long time before silicone outgassing subsided.

Protect Cryogenic Systems Against Thermal Expansion Mismatch

The Problem:

An expensive instrument performed poorly and failed early.

The Cause:

The instrument used a dewar filled with solid nitrogen to cool the detectors. Between filling the dewar and launch, cold helium was pumped through coils to keep the nitrogen from thawing.

Soon after the dewar was attached to the optical system, the cameras were found to be out of focus, albeit within the adjustable range. An investigation panel concluded that the dewar had deformed due to thermal expansion mismatch but approved the launch.

Unfortunately, defocusing worsened on orbit, and a camera became disabled. Moreover, the cryogen depleted rapidly, ending the mission.

Apparently, part of the camera light baffle, attached to the inner wall, expanded forward and touched the other part of the baffle, which was attached to the outer shield. The thermal short accelerated cryogen loss and increased deformation.

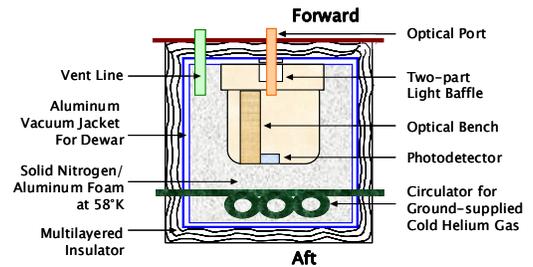
The unanticipated impact of repeated cooling cycles was not recognized because there was no prototype testing. During optics installation, an “alarmingly small clearance” was reported, but neither the designers nor the first investigation team conducted an interference analysis.

Lessons Learned:

- Perform in-depth modeling and thermal cycling tests on cryogenic systems, which are delicate equipment involving complex physics and material behavior.
- Provide adequate tolerances for thermal expansion mismatch (using flexible links, for example).
- Be extra vigilant when stretching the state-of-the-art.

For more technical information, call Martin Donabedian at (310) 336-6315.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



How Deformation Occurred

Because the aft part, though which super-cold helium was pumped, was colder than the forward part, forward nitrogen could sublime and refreeze aft, eliminating ullage space.

After helium flow stopped, the tank warmed up. The large CTE differential (700 ppm/°K for solid nitrogen, 17 ppm/°K for aluminum) probably forced the dewar to yield. Progressive deformation gradually closed the gaps between the baffles.

53

Test Hardware and Software Together

The Problem:

A satellite lost power shortly after launch.

The Cause:

The satellite used magnetic torquers for attitude control, a common approach.

Installation constraints made it necessary to mount one of the torque coils with a phase opposite of that of the other two coils. Unfortunately, this configuration was not reflected in the software reused from another mission, resulting in a sign error.

The mistake was not caught because the software was reviewed only at a top level. Moreover, the attitude control test to verify coil wiring was hardware-only. An end-to-end test, which would have detected the fault, was deemed too costly.

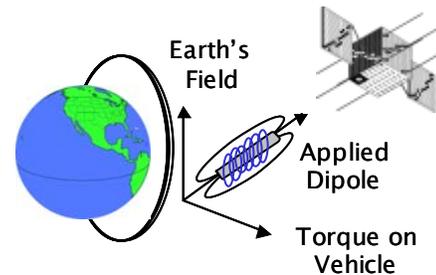
In orbit, the phase reversal caused the solar array to be steered away from the Sun. Limited ground station coverage made it impossible to diagnose the problem soon enough to prevent the battery from being drained.

Lessons Learned:

- Rigorously control configuration, especially at hardware/software interface.
- Always ascertain torquer polarity.
- Provide sufficient ground station coverage in early operation.
- Design battery protection to keep the satellite alive long enough for troubleshooting by implementing automatic load shedding and by configuring solar panels so that even a partially deployed array could keep battery charged.

For more technical information, call Tom Fuhrman at (310) 336-6596.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Magnetic torquers are coils wound around an iron core. Passing a current through the coils creates a magnetic dipole which interacts with the Earth's magnetic field and generates a feeble torque. Reversing the current flow (phase) produces the opposite effect.

Torquer polarity mistakes occur often. The orientation of large coils are easily verified with a magnetometer (essentially a compass). Background noise can make checking small torquers difficult.

54

Design and Handle Cryogenic Equipment with Great Care

The Problem:

A cryogenic dewar containing liquid helium exploded on ground.

The Cause:

Liquid helium freezes air. If air leaks into helium containers and blocks vent lines, internal pressurization can set off violent failures.

Blockage may occur when containers are brought to a lower altitude (for example, after being carried down from a mountaintop observatory). Also, since helium boils extremely readily, any heat ingress can cause the pressure to rise rapidly. Accidents involving cryogenic equipment are therefore fairly common.

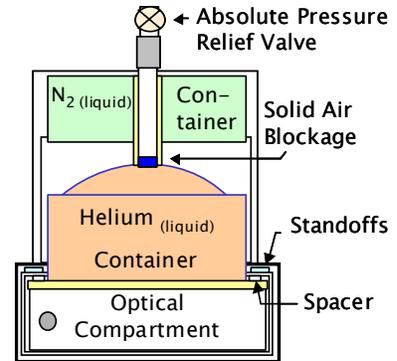
In this incident, a leak allowed air to freeze, plugging the vent line following a plane trip. Subtle structural flaws caused a thermal short. The pressure rose quickly, and the tank burst.

Lessons Learned:

- Review and follow operating and transportation procedures associated with cryogenic equipment to ensure safety to personnel, flight hardware, or facilities.
- Provide a graceful failure mechanism, if possible, to prevent catastrophic failure.
- Design for containment—make sure the cryogenics that unexpectedly boil off can be constrained within the vessel.
- Provide redundant vent paths.
- Design for convenient disassembly to aid inspection and maintenance.
- Service absolute pressure valves often—never exceed vendor specifications. Test valves before every field operation.

For more technical information, call John Hackwell at (310) 336-6041.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The exact cause of this accident could not be ascertained. The leak sprang due to contaminants accumulated in the valve, or fatigue of internal parts. The container was damaged before, which probably sheared off a spacer and tilted the container slightly. When the blockage formed, internal pressure pushed the tank into contact with the outer shroud, causing an unexpected thermal short. A small helium leak could have taken place too.

Do Not Dismiss Test Anomalies as Random Events—Find Out Why (I)

The Problem:

Two commercial satellites failed to deploy during the same Space Shuttle mission.

The Cause:

Both satellites suffered identical mishaps—the carbon/carbon nozzles on their kick motors came off a few seconds into firing.

Three other nozzles failed in a similar manner during qualification tests. Unfortunately, these failures were attributed to deficiencies in materials and workmanship. The flight incident investigation report also blamed the two failures on undetected flaws in the material used to fabricate the exit cones. The fundamental problem was not diagnosed.

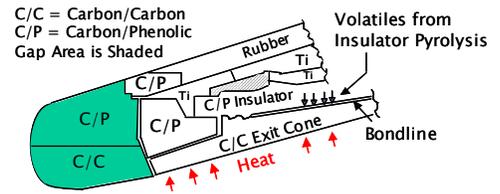
Because the motors were slated for government applications, Congress asked for an independent investigation. Finally, the root cause was discovered: charring of the unvented carbon/phenolic insulator created gaseous pressure within the exit cone. Since permeabilities inside the insulating materials are highly variable, the gas sometimes became trapped, forcing the exit cone to buckle. The problem could have been avoided simply by placing vent grooves in the bondlines.

Lessons Learned:

- Exhaustively search for the root cause of failures.
- Conduct fully instrumented tests.
- Provide sufficient thermal and structural margins to allow for material, manufacturing, and processing fluctuations.

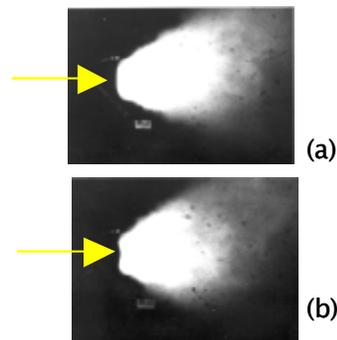
For more technical information, call S. R. Lin at (310) 336-7697.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Exit Cone Collapse

The independent investigation prompted NASA to conduct its own instrumented firing, which proved the buckling scenario. Prior to firing, the cone curled toward the left. It became vertical (Photograph A) and started to curl toward the right (Photograph B). The cone failed shortly afterwards.



56

Do Not Dismiss Test Anomalies as Random Events—Find Out Why (II)

The Problem:

A solar array drive failed soon after deployment.

The Cause:

The problem occurred because of a seemingly minor design tweak. The addition of an electromagnetic interference (EMI) filter allowed transient noises from the bus to propagate into the drive electronics. A spike blew the fuse for the H-bridge that controlled the motor.

During thermal vacuum testing in the months preceding this on-orbit failure, two other satellites in the same block also blew their controller fuses. Unfortunately, even though the previous block of satellites never encountered this problem, the project did not investigate the root cause. The damaged parts were simply replaced, allowing the satellites to be bought off.

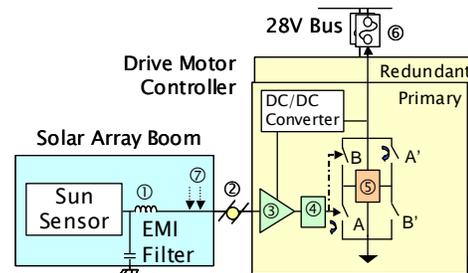
An earnest analysis would have identified the leakage from the EMI filter, and the on-orbit failure would have been avoided.

Lessons Learned:

- Define and implement a verification plan.
- Perform a worst-case circuit analysis to meet defined interface requirements.
- Always ascertain the root causes of ground test anomalies (Lesson 55).

For more technical information, call Walter Dennis at (310) 416-7207 or Steve VanWormer at (703) 633-5213.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Signals from the Sun Sensor passed through the EMI filter, ①, the slip rings, ②, and the amplifier, ③, to the controller ④. The controller oriented the boom by alternating the motor ⑤ between two states (A, A' transistors on the H-bridge open, B, B' closed; and B, B' open, A, A' closed).

The grounded EMI filter, coupled with a circuit not designed for fast switching, allowed transient noises from the chassis to momentarily turn all transistors on, blowing the fuse, ⑥.

Installation of a resistor (⑦) eliminates the noise problem.

Protect Propulsion System from Contamination

The Problem:

A launch was delayed for many months.

The Cause:

Following a guidance system malfunction, the satellite had to be removed from the launch vehicle. Off-loading of the toxic propellant caused a problem: the legacy satellite had no gravity drains, and the thruster valve was not robust. Neither the original valve vendor nor the system manufacturer was still in business, nor could the build paper be located to help find a good solution.

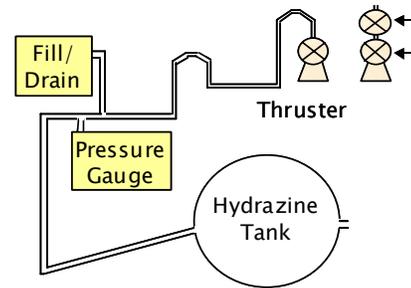
The decision was made to pump out most of the fuel, fix the guidance unit, and restack the satellite. Unfortunately, before refueling could start, a valve failed. Carbon dioxide in the air leaked in and reacted with hydrazine, forming corrosive carbazic acid and fouling the line. The entire propulsion system had to be replaced.

Lessons Learned:

- Consider retrofitting legacy hardware with proven design upgrades. Anticipate out-of-sequence operations, such as rework, during hardware design.
- Design propulsion systems to accommodate ground handling by including features such as low point drains to facilitate fuel removal.
- Archive manufacturing documents.

For more technical information, call Mark Mueller at (310) 336-5081.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Fuel System (Simplified)

The higher location of the fill/drain port in the legacy propulsion system prevents gravity draining, and the single seat valve is prone to leak. Dual seat valves (right), typically used in new designs, would have prevented air ingress unless both valves leaked.

A Similar Incident

An ICBM, refurbished to launch satellites, suffered a performance degradation recently after its turbine seal leaked, allowing ammonia in the exhaust gas to react with the lubricant, plugging the filter and blocking lubricant circulation.

The problem, chemically alike the thruster contamination, was addressed in the follow-on generation of the rockets, but the original units were not retrofitted.

58

Guard Against Sneak Paths Through Ground Test Equipment

The Problem:

The primary side of an instrument failed shortly after launch.

The Cause:

The instrument had parallel redundant power pins, but the power plug on the bus had only single pins for source and return. The flight cable had to be spliced so redundant conductors could be crimped into the same socket. The circuit opened because of broken solder joints at the current supply board, loose contacts, or defective crimps.

A subtle test issue hid this single point failure. The instrument needed a long time to stabilize, and was therefore kept on during ground testing by an external power supply with battery backup. On the test stand, the instrument operated normally, despite the faulty cable, by drawing power from the external power supply.

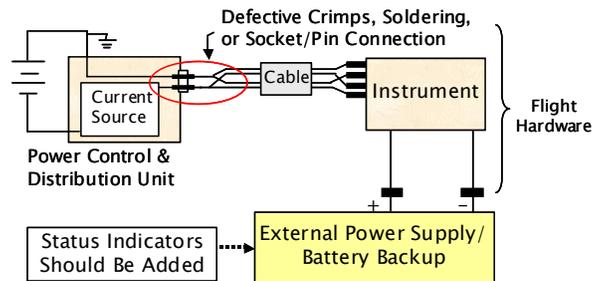
The flaw would likely have been caught if the test equipment provided metering to show the unit was unexpectedly drawing power from it.

Lessons Learned:

- Independently confirm hardware performance for functions temporarily provided by test equipment.
- Use a breakout box to check harness connector paths, and directions and magnitudes of currents flows.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Test Setup (Simplified)

Similar Examples

A flight box was not grounded by mistake. The problem was missed because the test equipment supplied grounding.

59

Lesson from Challenger: Understand Your Data!

The Problem:

Vital O-ring data was ignored before the Shuttle lifted off on a freezing morning.

The Cause:

During a pre-launch telecon, 34 engineers debated for hours over whether to delay the launch, out of the concern that cold weather might compromise the seals.

Citing O-ring anomalies at both 75 deg F and 53 deg F launches, some engineers argued against launch. But because damage occurred both hot and cold, managers perceived no temperature effect. The launch went forward.

The Post-Challenger Investigation Commission found that in presenting the flight history, the engineers omitted data from flights in which the O-rings remained intact, mistakenly thinking that successful flights did not provide any evidence about risk.

If presenters had plotted data from all flights, nobody would have missed the effect of temperature on the O-rings!

Lessons Learned:

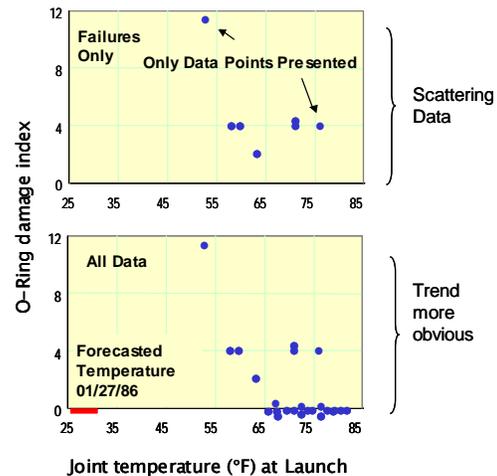
- Consider all relevant information.
- Develop a coherent explanation of engineering data to help audience analyze risks.
- Display data cogently (see *Visual Explanations* by E. Tufte, for example).

For more technical information, call Jon Binkley at (310) 336-7787.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

| <u>MOTOR</u> | <u>O-Ring</u> | |
|--------------|---------------|---|
| DM-4 | 47 | Tested on horizontal platforms in Utah |
| DM-2 | 52 | |
| QM-3 | 48 | The only 2 launches (of 24) shown |
| QM-4 | 51 | |
| SRM-14 | 53 | Forecasted temperature for the Challenger |
| SRM-22 | 75 | |
| SRM-25 | 29 | |
| | 27 | |

A table of temperature data presented during pre-launch telecon included irrelevant information but only selective flight experience. The audience was misled.



O-ring Damage History

Anomalies rarely occurred in warm days, but routinely took place during launches below 65°F.

60

Tests Are for Verification, Not for Discovery

The Problem:

A satellite started to tumble shortly after deployment.

The Cause:

The spacecraft used magnetic torque rods to stabilize body spins. During the Guidance and Control (G&C) subsystem test, an analyst misinterpreted the meaning of the Earth's magnetic poles and set the flight software incorrectly. The error went unnoticed because the coil test had no expected polarity values—the configuration was determined based on the measured responses.

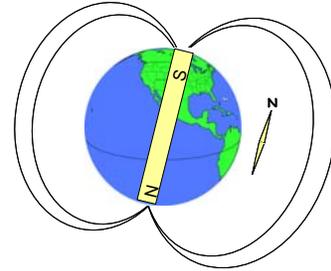
After separating from the launcher, the satellite began to wobble. Fortunately, the lead G&C engineer was prepared. Having heard many horror stories about torque rod phase mistakes, he had spent the previous day making contingency plans. Within half an hour, he reversed the controller gain, stabilizing the satellite.

Lessons Learned:

- Expected test results should be established in advance of the test. Deviation from expected results should raise a flag, and be thoroughly investigated before making any changes.
- Rigorously manage software development, especially on requirements, interfaces, and configuration control.
- Plan for contingencies, using a top-down fault tree (ask “what happens if the satellite failed to de-spin?” for example).
- Double-check torquer signs (Lesson 53).

For more technical information, call Tom Fuhrman at (310) 336-6596.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The Earth as a Magnet

Opposite magnetic poles attract. The north pole of magnet needles points to the Earth's magnetic South Pole, also called the geomagnetic North Pole!

61

Do Not Assume a Situation Is Acceptable Simply Because Nothing Is Said About It in Documents

The Problem:

A separation failure sent a launch vehicle tumbling out of control.

The Cause:

Following stage-1 separation, a small interstage ring surrounding the stage-2 nozzle also had to be jettisoned. Equipped with three guide tracks, this ring was supposed to slide along three foam blocks attached to the gimbaled nozzle without striking it.

One of the foam skids had to be installed just days before launch, through an access panel with little visibility. The technician reported to an on-site engineer that the foam felt too tight. Seeing no inspection criteria in the sparse launch-site processing instructions, the engineer assumed the tight fit was OK. He did not realize that the installation was off-center, nor query the designers as to possible consequences.

During ascent, the nozzle was commanded to a position that further pushed the foam against the guide track. Staging unleashed the strain in the foam, which jammed the interstage on the nozzle. The mission failed.

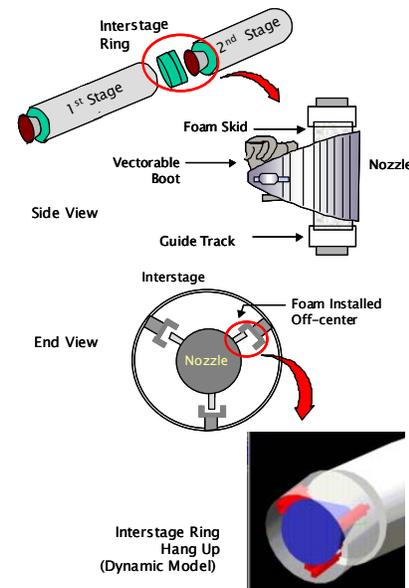
Several design changes, such as rounding the foam blocks, were later made to reduce the friction between the foam and the track—something not previously considered.

Lessons Learned:

- Double-check designs against possible misinstallation.
- Make sure field-assembled hardware can be inspected.

For more technical information, call Andy Shearon at (310) 336-1762 or Brian Gore at (310) 336-7253.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



An on-board video camera captured the interstage hang-up, enabling the investigation team to create a dynamic model and to replicate the problem on a mock-up.

62

Test as You Fly

The Problem:

A battery exploded on orbit.

The Cause:

The silver-zinc (Ag-Zn) battery, powering an arcjet experiment, fires pulses at a peak power of 30 kW and a current approaching 180 A. It takes about 24 hours to replenish the battery between discharges.

Silver-zinc batteries, used in all launch vehicles, are typically run for just a few minutes. Although some Ag-Zn batteries are rechargeable, they are not intended for arduous duty cycles. In particular, prolonged use of Ag-Zn batteries in space is apt to cause electrolytes to spill, forming a metallic zinc bridge through which a large current can flow. This problem led to a serious malfunction in an upper stage (Lesson 22).

The designers overlooked these issues. Qualification tests did not fully simulate the operation scenarios, and all ground firings were performed with a fresh battery at atmospheric pressure in an upright position. In actuality, the cells are partially discharged and laid on their sides during launch, making spills more likely.

In orbit, leakage triggered a violent short. The plastic case ignited, and the battery blew up.

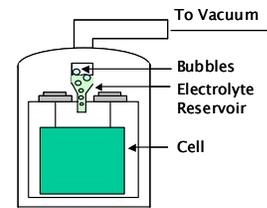
“Ultimately, this anomaly occurred because of a programmatic philosophy to minimize cost,” said the failure report. “All failure scenarios could have been ruled out if enough testing had been done.”

Lessons Learned:

- Analyze prior incidents of equipment malfunction.
- Review all aspects of battery application—do not regard batteries as simple plug-and-play items.

For more technical information, call Doug Chism at (310) 336- 6375.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The Hazards of Activated Ag-Zn Batteries

Dry silver/zinc batteries are activated by adding electrolytes in a vacuum environment. Once filled, internal reactions can lead to frothing and spattering. Launch depressurization and continuous discharging heat up the cells, causing more spills.

Serious mishaps had occurred, even on the ground. Several years ago, a launch delay caused a battery to exceed its wet life. Days later, it caught fire. Apparently, drops of escaped electrolyte made their way along the power wires via capillary action, shorting a connector.

Verify Field Installations of All Single-Point-Failure Items

The Problem:

A suborbital launch failed because the second stage would not start.

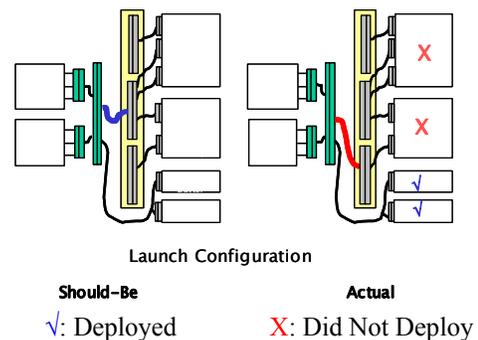
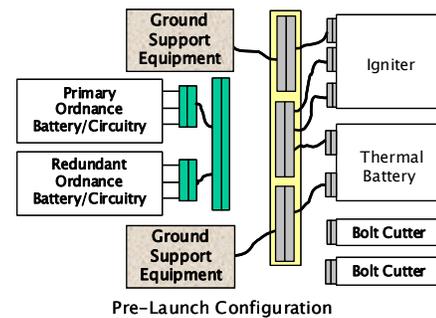
The Cause:

After the first-stage burn, two bolt cutters were fired, successfully jettisoning the spent stage. However, neither the second-stage motor nor its thermal battery ignited upon command.

The igniter and the thermal battery shared an ordnance connector which, due to range safety rules, had to remain detached until just before launch. Adjacent to the ordnance connector was a ground power receptacle.

Just prior to launch, the ground power umbilical was removed. Subsequently, the harness slated for the ordnance connector was instead mated by mistake into the neighboring power plug! Although both sides of the connection were male, their shell types and pin configuration allowed an unintentional fit.

The error was not caught because, unlike most Air Force programs, an end-to-end test with a load to verify circuit performance was not performed, nor was a quality assurance checklist used.



Lessons Learned:

- Simplify interfaces, commands, and procedures in prelaunch operations lest the hectic pace cause errors.
- Verify final assembly operations, particularly on single-point-failure risks. Pay particular attention to possible connector mismating.
- Do not allow primary and redundant sides of critical circuits to join in a single-point-failure area.

For more technical information, call Bruce Wendler at (310) 336-5475.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

64

Review Out-Of-Flow Processes to Ensure No Steps Are Bypassed

The Problem:

The temperature of an antenna dropped below expectation in certain conditions.

The Cause:

A legacy antenna had a radiator that was oversized for this mission. Thermal designers specified that the excess area should be covered with multi-layer insulation (MLI).

A veteran engineer, conducting a walkaround prior to the system-level thermal vacuum test, discovered that the MLI was missing. The blanket was installed.

After the test was completed, the temporary MLI was removed in preparation for installation of the flight MLI. Unfortunately, the final integration order still neglected to include the MLI reinstallation instruction.

Meanwhile, the old hand retired. His replacement did not spot the missing MLI, and the antenna was flown without the blanket.

Lessons Learned:

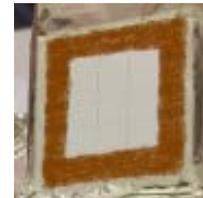
- Make sure corrections in engineering drawings or work instructions are back annotated in all applicable drawings and shop orders (including subsequent builds and units that have been distributed).
- Conduct final walkthroughs in the presence of the most experienced personnel.
- Keep good records of all “non-flight” installations.

For more technical information, call Todd Dickey at (310) 336-5352.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Antenna Without MLI



MLI Installed

A Similar Incident

A satellite used active louvers to control the baseplate temperature of an instrument.

The system, including the louvers, underwent thermal vacuum testing, after which the louvers were removed. They were temporarily reinstalled, without being connected, for fit check.

The louvers were left in place, without anyone realizing that the connector remained unattached. Pre-shipment checks did not verify the mate status because the connector was not accessible.

Running too hot in space, the instrument suffered significant degradation.

65

Perform Thorough Post-Flight Analysis

The Problem:

A launch vehicle lost control.

The Cause:

The investigation board traced the mishap to a solenoid valve in the thrust vector actuators. Apparently, microscopic metal shavings, created during the assembly and adjustment and dispersed during ascent, jammed the spool shut for eight seconds—time enough to ruin the mission.

In a previous launch, this valve stuck open. In another, it seized up twice, once open, once closed. Minor anomalies occurred two other times, but all previous flights succeeded.

Since a valve that is stuck open is manageable, these earlier troubles were disregarded. But a sticky valve can as easily fail closed as open. The blockage proved lethal.

“It is recommended that procedures for dealing with flight and ground test anomalies be reviewed. This recommendation is necessarily the least specific of those arising from this investigation, but may be the most significant,” concluded the board.

Lessons Learned:

- Track down the root causes of anomalies and consider implications beyond the narrow issues at hand (Lesson 41).
- Unexpected hardware behavior implies a failure to understand the application. Safety cannot be inferred just because the mission succeeded since the problem may be much more severe next time (rephrased from *Personal Observations on the Reliability of the Shuttle*, by Richard P. Feynman).

For more technical information, call Keith Coste at (310) 336-0032.

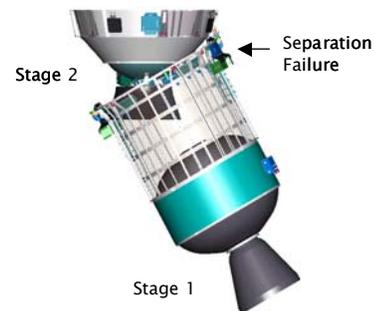
For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

A Similar Incident

Misleading instructions on drawings led assemblers to wrap thermal tapes too close to a separation connector (Lesson 4). The stage jammed (see diagram below), stranding the satellite.

Eleven previous flights were subsequently reviewed; all showed the same hang-up. Seven, in fact, were saved only because the floating connectors were jolted apart when they hit the allowable stops. The mission right before the failure had the narrowest escape.

The warning signs were not pursued.



66

Thoroughly Analyze All Environmental Load Paths and Develop a Detailed System Dynamic Model

The Problem:

A solar array broke on orbit.

The Cause:

Four solar array paddles were attached to the spacecraft with aluminum brackets. Three brackets were stiffened with gussets, but interference from surrounding components prevented a gusset from being added to the fourth bracket.

During vibration testing, the flexible hinge channeled most of the force into the release mechanism at the other end of the paddle, damaging a latching clevis. The problem would have been recognized had the paddle been instrumented or the component inspected after test. Unfortunately, the program did not adequately analyze dynamic loads during environmental testing and launch.

Loads during upper-stage burn exceeded nominal, and the clevis and bracket came loose. The paddle was left dangling by its cabling. The attitude-controlling magnetometer malfunctioned, whereupon the satellite turned away from the Sun, draining the battery.

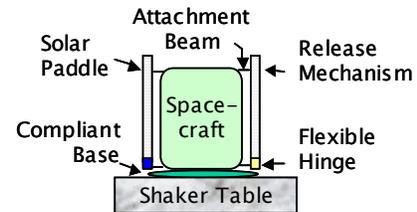
The satellite was rescued later (Lesson 67).

Lessons Learned:

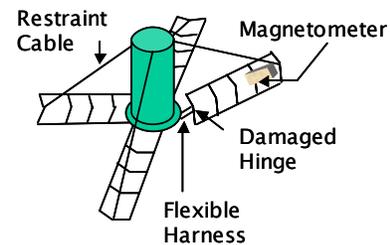
- Provide extra margins to accommodate excessive launch shocks that occasionally occur, especially with new launch vehicles (Lesson 11).
- Independently review dynamic loads analysis prior to test.
- Adequately instrument the unit, subsystem, and vehicle during environment tests.
- Check all data and inspect critical parts for damage after tests.

For more technical information, call Julia White at (310) 416-7229.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Test Configuration (Side View)



As-Deployed

67

Provide Design Flexibility to Enable Emergency Recovery

The Event:

Despite a damaged solar array (Lesson 66), a satellite was recovered.

The Cause:

When one of the solar paddles came loose, the magnetometer attached to it was disabled. Lacking autonomous attitude control, the satellite turned away from the Sun, and the battery drained. Ground controllers could not contact the satellite.

Fortunately, a video from the launcher showed that the failure might be survivable. Operators persevered.

Weeks later, a downlink arrived. As it happened, the satellite had rotated such that Earthshine could partially replenish the battery!

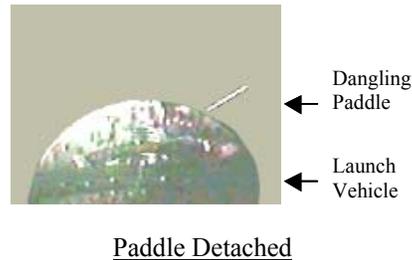
All non-emergency functions were commanded off to allow the batteries to fully charge. With its torquers manually controlled from the ground, the satellite was reoriented toward the Sun and spun up nominally. Full operation started three months after launch.

Lesson Learned:

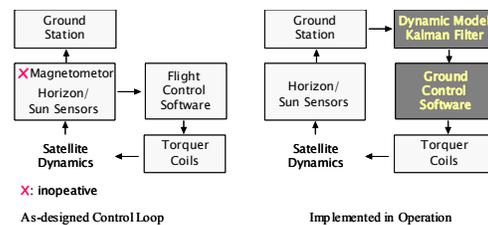
- Provide as much telemetry as possible on launch vehicles, especially on separation events. Without knowing how the satellite malfunctioned, controllers would likely have given up before the downlink was received!

For more technical information, call Tom Fuhrman at (310) 336-6596.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Paddle Detached



Pointing Information Recovery

Accurate attitude knowledge, especially during orbit night when most of the observations were made, posed the next challenge—the satellite no longer rotated as a rigid body; even the spin axis orientation was uncertain.

The program created a non-linear rigid-body model. Using Sun sensor and horizon crossing indicator data as input, an algorithm incorporating Kalman filters calculated the satellite attitude to 0.25° accuracy, even during most of the orbit nights when direct sensor readings were unavailable. Most mission requirements were met.

68

Insist On End-to-End Ownership to Verify Interfaces

The Problem:

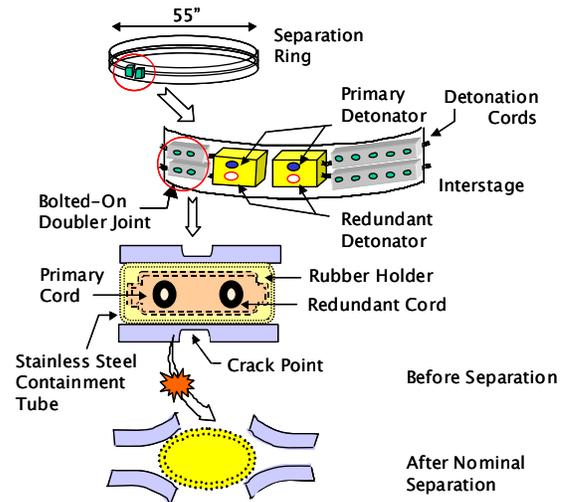
An uncontrolled explosion during the release of a satellite damaged the Space Shuttle.

The Cause:

The separator's harness, which plugged into four closely located detonators, was designed wrong.

The "Fire 1" command should go to the port and starboard primary detonators, followed by a "Fire 2" command to the backups a fraction of second later. A successful firing of the primary cord will cut off the backup signal, preventing excessive explosion.

Instead, the "Fire 1" signal was routed to the port detonators for both the primary and backup cords. A simultaneous shock broke the containment tube, hurling debris through the shuttle bulkhead. Fortunately, nothing critical was hit.



Separation Mechanism (Simplified)

The mistake was not caught despite hundreds of hours of reviews and tests because the separate drawings were never put together into a single, end-to-end, schematic. "Even after the occurrence of the separation system anomaly, detecting the design error through drawing reviews was difficult," reported the investigation panel.

Investigators also found that the documentation describing the mechanical and electrical subsystem interfaces was inadequate. Labeling of the components was "incomplete and confusing." Verification tests were flawed—designed to ascertain that the separator was built to the (flawed) design, instead of demonstrating the intended function. Discrepancies raised during the critical design review were not properly resolved.

Lessons Learned:

- Develop end-to-end diagrams for electrical and mechanical interfaces, including software driven interfaces.
- Clearly label each connector to avoid mismatching.

For more technical information, call Selma Goldstein at (310) 336-1013.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Protect Solid Rocket Grain Structure from Destabilizing Gas Flow

The Problem:

A prototype solid rocket motor exploded during prequalification firing.

The Cause:

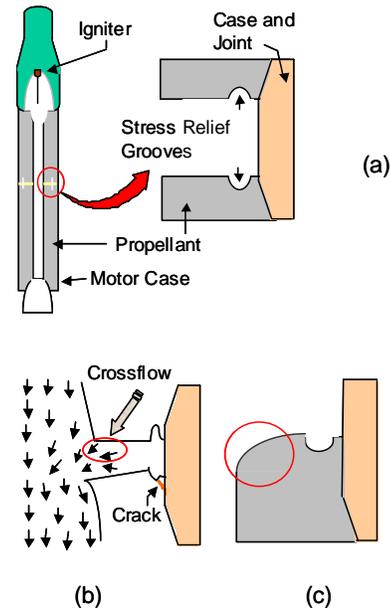
Mixing of combustion gas streams created a turbulence near the interstage joints, causing the soft propellant grains to crack. Blocked in the main bore by slumped propellant, the gas burst the case.

The contractor did not realize that the grain deformation should be taken into account even though a similar problem occurred in another solid motor (a lesson not shared). A subscale flow test would have revealed the dynamic instability problem. Unfortunately, the contractor bypassed this step.

In the wake of this failure, a sophisticated model was developed so that the impact of gas flow on the grains could be evaluated. A redesigned motor successfully passed the full-scale firing.

Lessons Learned:

- Conduct adequate subscale testing.
- Study post-test and post-flight anomaly reports from similar programs.



The original design (a) constricted flow at the segment joint. The grain cracked (b), further raising chamber pressure.

Chamfering of the forward grain face (c) eliminated the chokepoint.

For more technical information, call Nat Patel at (310) 336-6473.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

70

Late Modifications Require Careful Revalidation

The Problem:

A jammed tether prevented a satellite from being deployed from the Shuttle.

The Cause:

Post-flight inspection found that a bolt protruded into the path of a traveling ball nut.

The bolt was installed at the launch site, after an eleventh-hour analysis uncovered a design error: an overlooked thermal design change took away the cold plate's ability to carry loads, and altered the satellite's mass properties. By the time the problem was found, the tether deployment mechanism had been validated, and the satellite had been integrated.

Under severe pressure to improvise a fix, engineers overlooked the interference caused by the bolt because assembly drawings were not current (no updates were required until after three modifications), nor did drawings provide a direct view of the interference path.

The original design engineer, thousands of miles from the Cape, could not see firsthand how the modified hardware fit. The modification was not tested, and the change review considered only the loads.

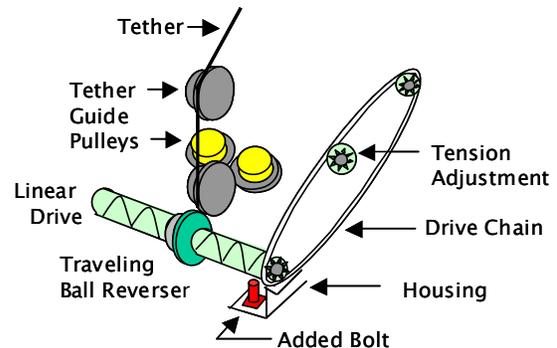
If the load inadequacy had been discovered sooner, it could have been corrected by simply making the fasteners larger.

Lessons Learned:

- Perform thorough analysis and testing of late hardware changes. Pay particular attention to system-level impacts.
- Update structural analysis following design changes to find problems earlier.
- Avoid assessing design changes from a narrow, discipline-oriented view.

For more technical information, call Brian Gore at (310) 336-7253.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Tether Mechanism (Simplified)

71

Make Sure Ground Support Equipment Cannot Damage Flight Hardware

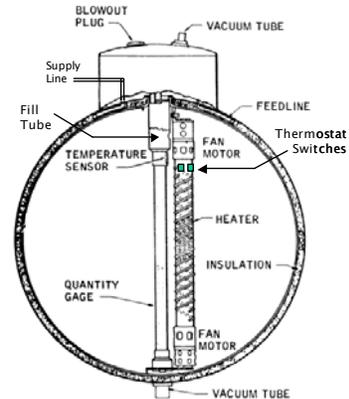
The Problem:

An oxygen tank on Apollo 13 blew up.

The Cause:

A month before launch, the spaceship was stacked on top of the Saturn V booster and moved to the pad. Countdown rehearsal began.

Operators filled two liquid oxygen tanks in the service module, then pumped in gaseous oxygen to empty them. One tank could not drain. Apparently, a handling accident had jarred loose an internal fill tube, preventing the gas from reaching the tank bottom and displacing the liquid. An internal heater, designed to maintain tank pressure in flight, was used to boil off the remaining liquid oxygen.



Oxygen Tank (Simplified)

Initially required to operate from 28 bus volts, the tanks had been reengineered to accept ground power at 65 volts. Unfortunately, the redesign overlooked two bimetallic thermostats protecting the heater circuits, and neither qualification nor acceptance testing exercised them.

As the detanking proceeded, the temperature rose. The bimetallic switches began to open, but the higher voltage immediately induced arcing across the contacts and welded them shut. With no one monitoring the current, the heater ran for eight hours. The temperature reached 1000°F, severely damaging the insulation on the power wires leading to a fan motor. When the astronauts activated the fan en route to the Moon, a short touched off the infamous explosion.

Lessons Learned:

- Ensure heritage thermostats and relays properly function when the system is redesigned for higher voltages.
- Provide ample test instrumentation to validate that all components of a system are functioning properly, and always check for unplanned current draw (Lesson 19).
- Individual heater circuits should not draw more than two amps to prevent thermostats from being damaged by self heating (each of the Apollo 13 switches drew six amps).
- Thoroughly test subsystems that are not exercised until they are integrated into the main spacecraft (such as propulsion lines) during system thermal vacuum test.

For more technical information, call Bill Fischer at (310) 336-5198.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Prevent Failures in Support Equipment from Propagating into Flight Boxes

The Problem:

A transmitter was damaged during test.

The Cause:

The test set incorporated 15 separate power supplies with various voltages. To automatically record data from each test point, a computer addressed the power supplies via a bank of relays. The commercial test unit did not isolate each monitor point.

A relay on the 5-volt line did not disengage after being scanned, remaining tied, via the monitor's internal bus, to all power supplies subsequently scanned. Exposed to as high as 31 volts during the following scan, the 5-volt flight circuits were damaged.

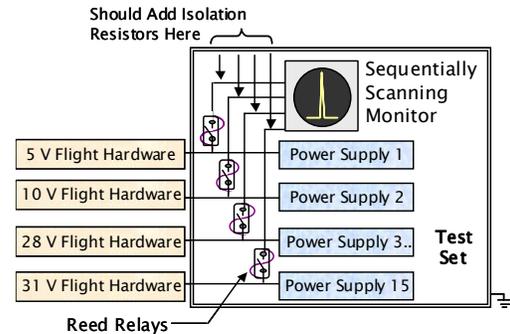
The original safety analysis of the ground equipment did not consider the impact of a failure on the flight box. Isolation resistors between the power supply lines and the scanner inputs would have averted the damage.

Lessons Learned:

- Buffer test point outputs so shorts in test will not damage flight hardware.
- Implement abort logic in automated test equipment to prevent damage if a failure occurs.
- Thoroughly understand the inner workings of any item that interacts with flight hardware.

For more technical information, call Ron Williamson at (310) 336-2149.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Test Arrangement (Simplified)

Reed Relays

Reed relays, commonly used in control circuits, consist of two overlapping iron strips enclosed in a glass tube. The contacts are readily closed with a magnetic field applied via the surrounding coils.

The strips should spring back to their normally open position after the field is turned off, but residual magnetism or magnetic contaminants sometimes keep them stuck closed.

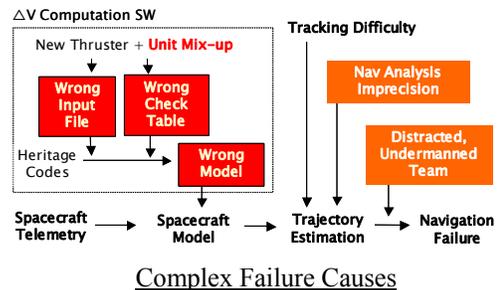
Trace All Software Changes Back to System Requirements and Specifications—Do Not Simply Modify the Code

The Problem:

A spacecraft broke up near Mars.

The Cause:

En route to Mars, the probe would fire its thrusters to unload the reaction wheels. Ground controllers planned the burns with a thruster model, reused from a successful mission.



A thruster change made it necessary to update this model, which specified thruster input in Newton-sec. The thruster vendor—the same for both missions—used lb-force-sec. In the original model, engineers correctly added the 4.45 conversion factor to the vendor’s equation. Overlooking the interface specification and seeing no warning in the code comments, the follow-on team simply made a substitution.

Labeled as non-mission critical, the ground software—without the conversion factor—was not rigorously reviewed; the “truth” table, computed manually for acceptance testing, contained the same mistake. Interface with the navigation function was informally tested only to ensure that it could move across servers.

Only one, occasionally two, engineers navigated the spacecraft. Two months before orbit insertion, radar returns projected a path too close to Mars. Unfortunately, as the probe neared Mars, poor observation geometry from Earth reduced tracking precision. The flight team, confident with their navigation ability, decided against raising the orbit.

Not until aerobraking, after Martian gravity had captured the probe, was it possible to calculate the spacecraft’s true position. Only then did the controllers realize the probe was 100 kilometers off course!

The successful reflight listed both English and metric units on all interface control documents, adopted a more robust navigation method, and used six full-time navigators.

Lessons Learned:

- Any software that commands a satellite is mission critical, even though it may not be embedded in the flight vehicle.
- Validate changes in mission-critical software with more vigor than the original development (Lesson 25, 29, 47). Rigorous formal testing is essential.
- Always specify the units in requirements and Interface specifications.
- Generate expected results used in verification tests independently, in accordance with system requirements.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Understand Why Warning Lights Come On Before Disabling Them

The Problem:

During a satellite test, the thermal vacuum chamber suffered a pressure burst.

The Cause:

An investigation revealed that the helium refrigeration system unexpectedly shut down. The unit had sprung a small leak during a previous test.

As the test progressed, the helium leak rate increased, causing the pressure in the turbine wheel inlet to oscillate.

Vibration caused the alarm setting in the pressure regulator to drift down. The alarm went off a few days into the test.

Knowing that the equipment was working well within its normal range, the testers returned the alarm level to near the factory-set level. The engineers unfortunately did not realize that the regulator's emergency shutdown sensor could drift down, too.

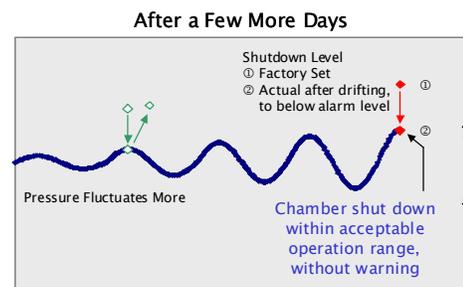
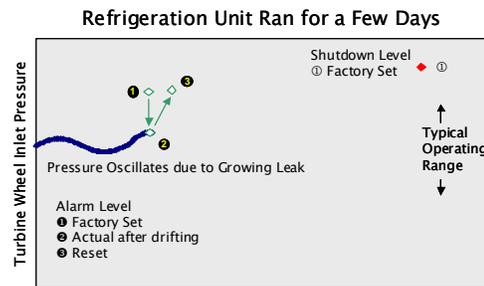
A few more days into the test, while the turbine was still operating within its normal range, the oversensitive emergency shutdown sensor tripped. Lacking an operator override or other means to gracefully degrade, the turbine switched itself off. Since the alarm setting had been adjusted up, the malfunction came without warning. The satellite could not be powered off first, and corona discharging set in. Luckily, robust hardware design practices prevented serious damage.

Lessons Learned:

- Operate environmental tests with the same degree of care as space operation (Lesson 49).
- Develop test contingency plans and failure-mode-and-effect-analyses for ground support equipment (for example, analyze the likelihood of contamination in case the thermal vacuum facility loses power).
- If turning off a piece of test equipment can endanger flight hardware, such equipment must not be allowed to shut down autonomously.

For more technical information, call David Homco at (310) 336-5800.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Turbine Pressure Trend

75

Protect High-Voltage Equipment from Contamination

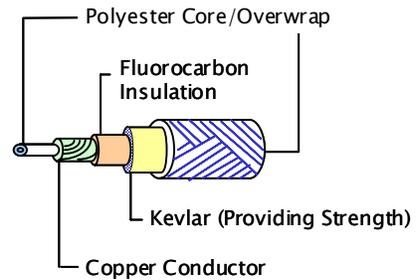
The Problem:

A satellite was lost when the tether deploying it was severed by arcing.

The Cause:

Inspection of the recovered tether fragment revealed contamination, pinholes, and other defects. Debris was also found on the deployment mechanism.

Apparently, the underlayers of tether experienced severe compression loads while wound on the reel. The insulation layer flattened, causing debris to puncture through.



Tether Construction

As the deployed tether flew through the Earth's magnetic field, a potential of several thousand volts was generated along its conductive core. An exposed spot attracted a spark from a nearby pulley. Because the mechanism housing was insufficiently vented (Lesson 49), the arcing continued, burning down the tether.

To avoid fatal arcing, the program fabricated the insulator layer with great care. Unfortunately, subsequent processing was performed in a regular shop, making contamination inevitable.

“Excellent designs can be defeated through quite common cleanliness and handling violations,” concluded the investigation board.

Lessons Learned:

- Design high-voltage equipment to withstand mishandling.
- Properly vent enclosed areas to eliminate corona and arcing caused by outgassing and pressure buildup.
- Thoroughly test the entire circuit if a high voltage is expected.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

76

Make Sure Someone Takes Responsibility for Each Interface

The Problem:

A space probe was damaged on the launch pad.

The Cause:

The probe, developed by one agency (A), required another agency (B) to provide launch-pad cooling.

Neither agency bothered to assign interface responsibilities. The requirements were not spelled out; the design and operational procedures were not placed under configuration control. Communications faltered.

Agency A faxed agency B a gas-flow value, which it intended as the not-to-exceed limit. The nominal value was buried in a thick review package.

Seeing only the faxed number, agency B made certain it could be met by making several procedural changes, such as narrowing the cooling duct, without considering the effect of too much air. On the pad, excessive air flow tore a hole in the probe's insulation.

The investigation board found that in five years the two organizations missed catching the problem 26 times. "The actions taken were logical, based on the knowledge available to the people taking action. The incident was entirely due to inadequate or imprecise information exchange," said the board.

Lessons Learned:

- Check ground operation procedures and support equipment to avoid damage to flight hardware.
- Ensure interfaces between two organizations are worked out in detail, agreed to by both sides, and documented.
- Bound each requirement within a range.

For more technical information, call Susan Ruth at (310) 336-6765.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

The Importance of Stating TBDs

Agency B's cooling plan stated that the equipment would be set to "agency A value" or "desired" flow rate." The two partners reviewed the plan step by step, never realizing that this number had not been agreed upon.

Stating "set to TBD \pm TBD units (agency A value to be supplied)" would have raised a flag and avoided the misunderstanding.

Make Sure Sequential Safety Devices Operate Independently

The Problem:

A science mission ended during the first orbit.

The Cause:

The aperture cover’s design called for its pyro circuits “safed” prior to being sequentially “armed” and “fired.”

A design feature in the controller chip invalidated all the programming circuits for a few milliseconds upon powering up. All outputs, including “ARM” and “FIRE”, were momentarily asserted. The cover blew open prematurely; the cryogen escaped.

The chip would manifest this start-up problem only after having been turned off for several hours. Although power cycled many times during component testing, it was never unpowered long enough to reveal the problem.

The use of a slow, non-flight-like, power supply during unit testing masked the spurious output: during the transient period there was not enough voltage to close the arming relays. Later, anomalies repeatedly occurred during system testing. Unfortunately, because the pyro simulator was very sensitive, a load delay was fitted to the test equipment to filter out spurious triggers, unintentionally preventing the actual start-up glitch from being recorded. The warning signs were ignored.

At launch, the chip had been powered down for weeks. Not only did it go awry but, because power to the pyro box was applied via a fast relay, sufficient voltage had also built up to complete the arming circuit. The FIRE switch, commanded by the same controller and therefore not truly independent, set off as well, ending the mission.

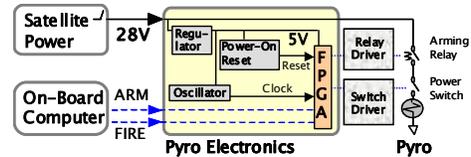
This controller chip had caused troubles before, prompting NASA to issue an application note. However, the contractor and the field engineer from the vendor did not know about it. “[We need] an information hotline, set up on an industry-wide lessons learned web page,” suggested the engineers later.

Lesson Learned:

- Beware that many programmable devices do not follow their truth tables at power-on— see <http://www.klabs.org/> for more information.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Timing Issue in the Safety Mechanism

After the bus power is switched to the pyro box via a relay, the controller (a field programmable gate array, FPGA) should be safed and initialized at the direction of an oscillator clock.

It took 30 milliseconds for the local voltage to rise and another 25 milliseconds for the safing clock to start, but only 15 milliseconds for the transient to occur.

78

Thermal Blankets And Tie-down Cables Can Jam Mechanisms

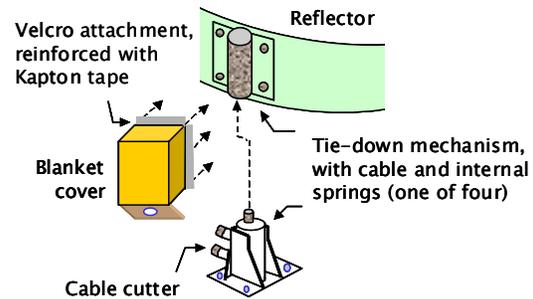
The Problem:

An antenna reflector on a communication satellite could deploy.

The Cause:

The reflector was tied down to the bus deck during launch with four cables. When the cables were pyrotechnically cut, the two hinged reflector booms failed to deploy.

Later, ground testing showed that the pocket-shaped thermal blankets covering the tie-down mechanisms expanded during the ascent, fouling the wrap cable. The spring-loaded hinges did not have enough force to overcome this interference.



Antenna Reflector (Simplified)

Fortunately, the satellite was designed to collect sufficient solar power even when the arrays were stowed, making it possible to spin and nutate the spacecraft in progressively more drastic maneuvers. Using ingenious ways to control the orientation, the operators were able to force the hinges open without damaging the satellite. The reflector opened a month later.

Lessons Learned:

- Anticipate the errant movement and expansion of flexible materials, such as wires and blankets.
- Allow thermal blankets to vent whenever possible.
- Avoid protrusions or sharp edges that can snag soft items.
- Indicate the presence of soft goods on top-level assembly drawings to draw attention to the risks of interference and obstruction problems.

For more technical information, call Robert Postma at (310) 336-7228.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Make Sure Software and Hardware Engineers Communicate with Each Other

The Problem:

An experimental spacecraft lost its computers.

The Cause:

The satellite, hitchhiking on the qualification flight of a launch vehicle, was designed and built in one year.

The bus software was checked out against the engineering model without incident, but was not tested against the payloads until the spacecraft was already loaded onto the host vehicle. It was then discovered that a payload performed very sluggishly.

Three launch-support engineers worked 14 hours a day for a week to adjust the bus memory-management functions. They created several software patches, one of them contained a wrong boot-up parameter. The mistake was not caught because the software developer did not consult with the processor engineers, nor verify the changes in the engineering model.

The software was loaded into the primary processor, which right away halted. Assuming a faulty primary memory was the cause, and again not enlisting the CPU expert's help, the engineers loaded the same code in the backup computer. It froze, too.

The computer could be physically reset. But by this time it would take several days to remove other experiments to reach the frozen computer, possibly delaying the flight. The host mission refused, and the hitchhiking project could only watch the launch, knowing its computers had already died.

The project manager traced the failure to poor communication between the software and hardware personnel, because the software team worked in isolation.

Lessons Learned:

- Make sure no single parameter error or single spacecraft malfunction can cause endless cycling (for example, by enabling the watchdog function to switch to a recovery mode after a few “try agains”).
- Double-check last-minute code changes (Lesson 43).
- Problems in embedded systems are not always due to random hardware defects. Pause and think before inflicting the same software flaw on the redundant side (Lesson 18).

For more technical information, call Lan Nguyen, at (310) 336-2146.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

A “Deadly Embrace” by the Watchdog

The computer uses an independently clocked watchdog function (Lesson 36) to enable switching to the redundant CPU if the primary side malfunctions (for example, due to radiation damage).

The final software mistakenly set the watchdog counter to 0.1 s, but it took the hardware about a third of a second to boot. The CPU could not finish booting before being reset, and was stuck in an endless loop.

80

Check, Double-check, and Triple-check Torquer Phases

The Problem:

A magnetic torquer sign error was caught just one day before launch.

The Cause:

The attitude control engineer who calculated the fields induced by the applied current made an error in an equation, which reversed the predicted torques.

The engineer left the project, and his successor, misunderstanding the vendor's drawing notes, installed all three coils upside down. The second error, which could have been easily discovered with a compass, was masked by the faulty truth table.

Fortunately, the prime contractor's president had concerns with a delay in generating solar power (Lesson 53). As a result, the attitude control components relating to sun acquisition were thoroughly scrutinized.

To alleviate prelaunch work load, the customer paid to bring back the original attitude control engineer. Rechecking his own calculations, he spotted the sign error one day before launch.

Lessons Learned:

- Don't overlook simple tests that can discover problems early.
- Whenever possible, conduct independent analyses.
- Document attitude control coordinate frames early in development to avoid mistakes.

For more technical information, call David Voelkel at (505) 846-8380 or Geoffrey Smit at (310) 336-1602.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Two Other Mistakes on This Mission

1. The calculated moments of inertia, which should have been referenced against the center of gravity, were instead referenced against the origin point on the drawing. The mistake was caught by an independent analysis (Lesson 2).
2. The star tracker misbehaved on-orbit because the vendor altered its coordinate convention but the change notice was not heeded.

81

Designate A Responsible Engineer for Complex Equipment

The Problem:

A satellite lost part of its primary structure one minute after liftoff.

The Cause:

The micrometeoroid shield that enclosed the spacecraft was peeled off by aerodynamic loads.

The 1200-pound shroud was supposed to fit tightly to the satellite body during ascent and then extend five inches after reaching orbit. The contractor delegated the development of this complex hardware to its structures department without putting a project engineer in charge.

Coordination suffered. Not having been told that the shield must fit tightly during launch, the structural and manufacturing engineers made it light but fragile. Without looking at the actual hardware, project engineers assumed that design criteria were met and saw no aerodynamic concerns. All dynamic tests were waived.

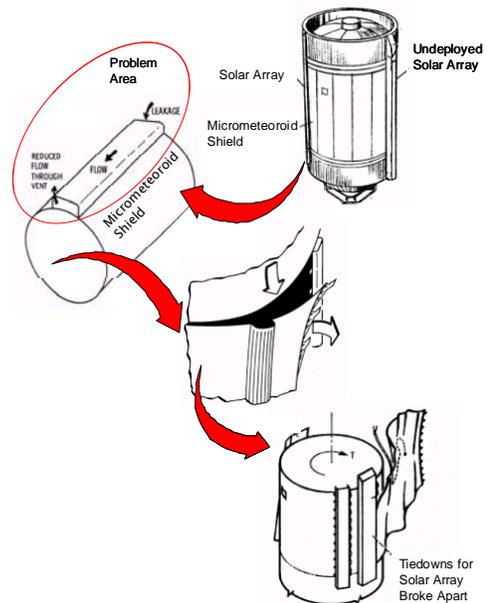
The investigation board blamed the failure on a lack of systems engineering leadership and chided the engineers for “believing that a drawing is the real world.” The board concluded that “positive steps must always be taken to assure that engineers become familiar with actual hardware.”

Lessons Learned:

- Designers should inspect actual hardware (Lesson 26).
- Analysis does not obviate the need to test.

For more technical information, call Susan Ruth at (310) 336-6765.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Damage Mechanism

Supersonic air rammed through a supposedly sealed tunnel on the shield, generating excessive lift that broke the shield as well as a nearby solar array.

82

Understand Transient Behavior of Analog Circuits

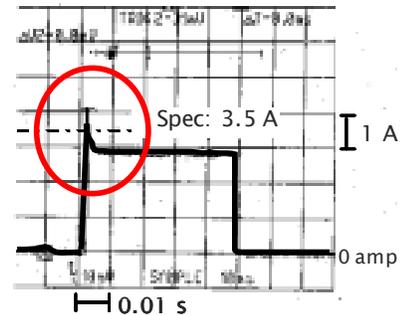
The Problem:

A pyro device failed to fire on orbit.

The Cause:

The incident stumped engineers because pyro units rarely malfunction, and two certification units fired successfully.

An outside expert pointed out that when current passed through the bridgewire, ohmic heating raised its resistance. Because the firing circuit was designed as a constant voltage output, current and power dropped off ($P = V^2/R$) just enough to thwart ignition.



Current vs. Time (Malfunctioning Unit)

Most pyro unit outputs are current-limited with series resistors, or energy-limited with capacitor discharges. Few engineers realize that the bridgewire resistance can change within the hundredths of a second it takes to heat the bridgewire enough to ignite the charges. In fact, the initiator specification only stipulated the firing current, not how long the pulse should hold. The designers, who did not know how pyro circuits typically work, used a constant, low-voltage approach that turned out to be vulnerable.

A lack of fidelity in design verification hid this mistake. During simulation tests, a resistor was used to emulate the initiator, and the current was steady because the resistance did not change. A fast-blow fuse, which more accurately simulates the load, would have revealed the resistance change.

The design was certified based on only two live firings, during which no current trace was recorded. In retrospect, the successes were purely a matter of luck—there was just enough current margin for success 60 percent of the time. If more units had been fired, or if instrumentation had been used, the inadequacy would have been found.

Lessons Learned:

- Check time-dependent circuit behavior, and bound transients in specifications.
- Do not qualify a design solely because a unit worked. Measure circuit parameters and verify that positive margins exist.
- Analyze instrumentation data, which can provide more engineering information such as postfire conduction (which may drain flight battery).
- Understand how circuits are typically designed and tested before inventing novel approaches.
- Qualify pyro devices by conducting lot acceptance testing.
- Review the Pyroinitiator User's Guide published by NASA (JSC-28596A).

For more technical information, call Ron Williamson at (310) 336-2149.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

83

Put Critical Analyses Under Configuration Control

The Problem:

Two upper stages failed for the same reason within 16 months.

The Cause:

Before launch, liquid helium was circulated through the cryoengines so they could start smoothly. During the boost phase, aerodynamic turbulence shoved air into the helium feed port. A malfunctioning check-valve allowed the air into the frigid engine, where it froze and jammed the turbo pump.

A check-valve, instead of a more secure shutoff-valve, was used in the duct because an air flow computation indicated that no pressure differential would exist within the line. But subsequent design changes created a pressure gradient. Because the aerodynamic analysis was not placed under configuration control, there was no requirement to recheck the calculations to confirm that the check-valve would still suffice.

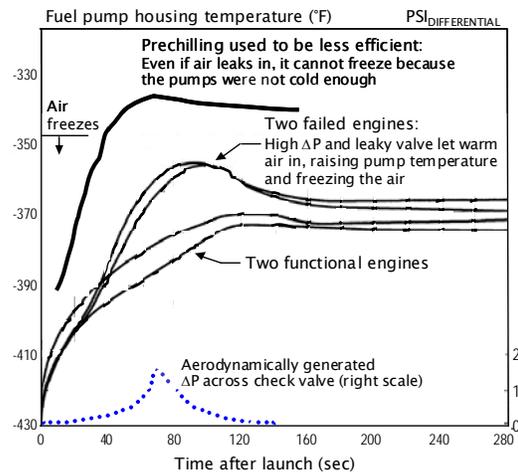
After the first failure, engineers tore apart several pieces of hardware in stock and found residual Scotch-Brite in numerous joints. Under considerable schedule pressure, they concluded that the failure was caused by contamination. The second investigation team examined more than 1200 potential causes before finding the actual cause.

Lessons Learned:

- Do not assume the first, easiest explanation is the correct one.
- Refrain from using check-valves as sole means for isolation, as they can chatter or leak (the check-valve design and assembly process on this launcher was particularly prone to seize in the open position). See *Check-Valve Reliability in Aerospace Applications*, NASA Preferred Reliability Practice No. PD-ED-1267, for additional information.

For more technical information, call Robert Foust at (865) 932-0366.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Telltale Thermal Telemetry Signatures

The failure cause was found in out-of-family data from successful flights between the two failures. Notice that a process change, chosen to reduce development costs, chilled the engine so much that ingressing air could freeze.

84

Check Start-up Circuit Behavior, Particularly at Low Temperatures

The Problem:

The primary side of an onboard computer would not turn on.

The Cause:

The computer received analog housekeeping inputs via numerous multiplexers inside the data interface unit (DIU).

During power-up, multiplexer chips can draw twenty times more current than during steady operation. The designers did not notice this start-up surge partly because it is only significant at low temperatures.

Following a safe-hold event, the onboard computer tried to reboot when it was unusually cold. The current draw exceeded the limit set on the fault-tolerance circuit, preventing the primary DIU, and consequently the primary computer, from starting.

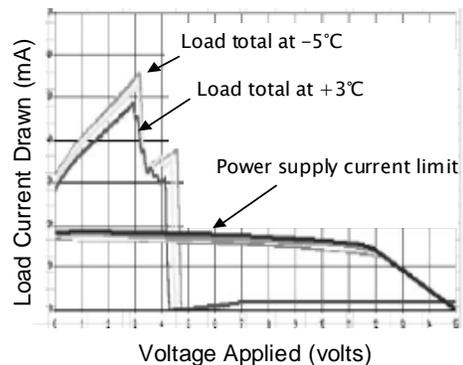
The current limiter did not have large enough margins because the DIU was inherited from an earlier design that supported fewer multiplexers. The low temperature DIU test was manually controlled, and the engineers did not realize that the unit took longer to boot than the time limit programmed into the computer.

Lessons Learned:

- Use fault-tolerance circuits to protect upstream assets, not load units. Better yet, use dual-level current limiters to protect load units during ground tests. But for flight, protect only the source circuits.
- Redesign fault-tolerance circuits when the load units have been substantially altered.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Effect of Temperature on Turn-on Loads

The multiplex chips draw 0.25 mA during operation, but as much as 5 mA during cold power up.

When the current draw exceeds the power source's capability, the unit would continue trying to reboot. The primary computer timed out; its back-up finally succeeded in booting after the chips warmed up.

85

Systems and Software Engineering Should Actively Coordinate

The Problem:

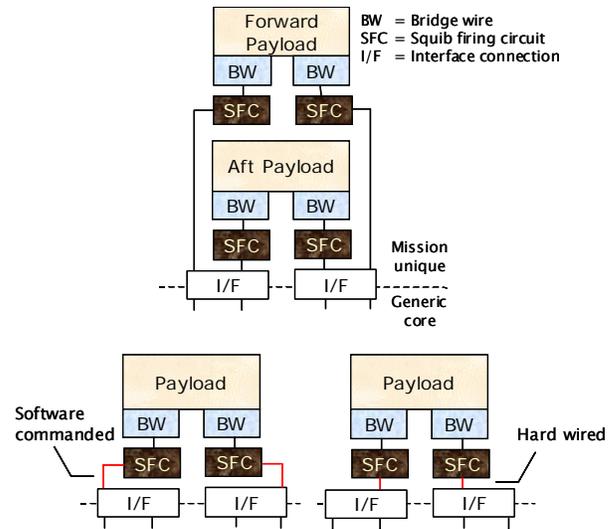
A satellite could not be deployed.

The Cause:

The payload separation system was designed to accommodate two satellites, but only one satellite flew on this mission.

The mission specification had the separation commands sent to the “forward” position. An engineer redlined the commands to “aft” to simplify wiring. Unfortunately, this change was not incorporated in the final mission specification.

Not realizing that the informal redline had fallen through the cracks, the hardware group designed an incompatible harness. The drawings were released as a new baseline, making it difficult to detect crucial changes. Several systems engineering departments could have checked the compatibility of the final design to overall requirements, but none did—the key mission specification was developed by software engineers and was not placed under systems engineering’s jurisdiction.



Separation Configuration

(Top) For two payloads
 (Bottom) For the failed mission

The mistake was not discovered on the ground because the generic systems test activated both positions, allowing the miswired ordnance verification unit to appear working.

Lessons Learned:

- Test the specific configuration that will be flown (Lesson 3).
- Conduct tests and reviews to validate that the requirements are met, rather than that the drawings are correctly implemented.
- Actively involve systems engineers in software development activities, and formally control all system (including software) interfaces.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Hand-Over Logic Tree Must Be Unambiguous

The Problem:

A suborbital launch was inadvertently terminated less than a minute after liftoff.

The Cause:

During the flight, control had to switch from the ground to a downrange airplane. Commands were sent via three analog channels: A, B, and C. Ground used tones A and B; and the airplane used tones B and C.

Tone B was the “ALIVE” signal, and a combination of tones B and C meant “ARM”. Once armed, if the onboard receiver loses the “ALIVE” signal, it would assume that something went awry and abort the flight. By having the airplane take over control using the ARM signal, the handover plan put the flight in danger.

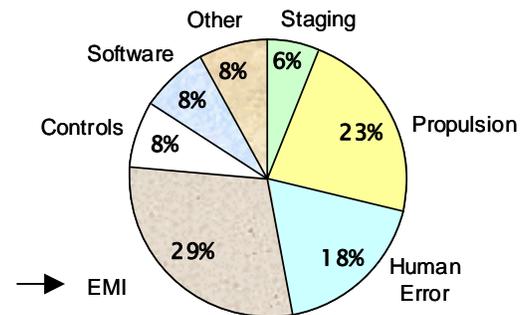
The onboard receiver detected tone C from the airplane and armed. However, it could not immediately lock into the airborne transmitter because plume attenuation caused incoming ground transmission to fluctuate. While the receiver dithered, land and airborne B tones became momentarily out of phase. The phase-looped oscillator in the receiver lost lock, spoofing the self-destruct mechanism into thinking it lost the “ALIVE” heartbeat. The launcher blew itself up.

Lessons Learned:

- Conduct redundancy switching analysis to ensure a fail-safe transfer between multiple, or redundant, controllers. Postulate all credible failure paths (such as part failure, start-up transients, latch-up, overvoltage, and EMI) and determine the effect on the switching process. Make sure glitches in one unit will not propagate across interfaces.
- Guard against radio frequency (RF) interference from multiple sources.

For more technical information, call Ron Williamson at (310) 336-2149.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



AIAA 2000-3578

Watch Out for Radio Interference

A study of missiles converted for suborbital or space launches found that the largest cause of failure was electromagnetic interference (EMI).

87

Avoid Repeating Other People's Mistakes

The Problem:

A launcher's maiden flight failed.

The Cause:

The launch vehicle, unlike most other systems, did not recycle hydraulic fluid, but drained it at the nozzle exit plane instead.

The spent oil dripped into the exhaust plume and caught fire. Recirculated by external air flow into the aft area, the flame damaged an uninsulated guidance cable, sending false signals into the thrust vector controller. The vehicle veered off course.

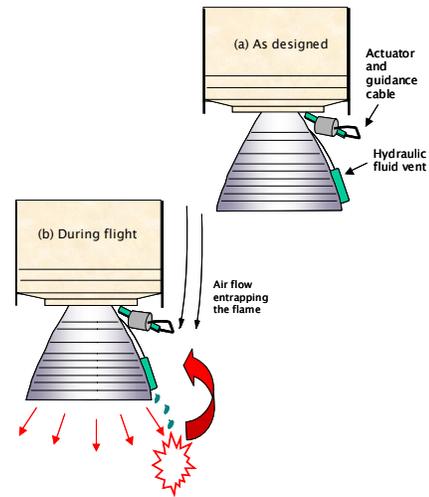
Four years earlier, another rocket crashed because excessive engine heat destroyed guidance cables. The investigation board concluded that jettisoned hydraulic oil could have dripped into the exhaust and contributed to the mishap. Several programs thereafter changed designs to keep fluid clear of the plume and to add insulation.

Unfortunately, even though the motor supplier of this failed vehicle also built the motor that went awry four years earlier, the lesson was not heeded.

Lesson Learned:

- Study past failures that involved similar technologies and implement appropriate corrective actions.
- Ensure subcontractors discuss relevant lessons with the primes.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Plume Safety

As a rocket ascends, decreasing atmospheric pressure causes its flame to spread out.

The designers of this failed launcher conducted static firings, but did not run sufficient computational fluid dynamics modeling. Thus, they did not anticipate the conflagration or the need to protect the cable.

88

Verify Each Operation Step

The Problem:

A piece of flight hardware was damaged during its integration to the launch vehicle.

The Cause:

During launch vehicle erection, the Stage III, spin table, and the satellite were contained in a canister and bolted to the Stage II. After the guidance systems were connected, a technician had to remove the bolts before the canister could be lifted.

To indicate that he was to start unbolting, the technician put both thumbs up and shouted “ready.” The crane operator heard “Randy,” his name, and mistakenly interpreted the gesture as a command to hoist. The shackled stack was raised up; the spin table suffered structural damage.

The error took place because:

1. Not realizing the lift operation could be hazardous, the foreman allowed an uncertified technician to direct the crane. A properly trained rigger would have avoided making an ambiguous “thumb-up” sign.
2. The operating procedure did not require anyone to verify that the bolts had indeed been removed. The crane driver should have been taught to ask for the restraining pin, for example, first.
3. The procedure did not specify communication protocol.

Lessons Learned:

- Implement a discrete verification step for each critical task.
- Require positive confirmation before hazardous commands can be acted upon.
- Do not deviate from written procedures.
- Handle space hardware carefully.

For more technical information, call Norman Lagerquist at (310) 336-2362.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

A Similar Incident

As a thunderstorm approached a launch pad, workers draped a rain shield over a satellite being processed in the White Room.

The shield consisted of overlapping strips of waterproof cloth, secured with adhesive tapes. The installation instructions stated, “ensure both top and bottom sides of seam are taped.” Nonetheless, the lower side was neglected, nor was there a verification.

Rainwater poured through the building’s leaks. The weak rain shield collapsed, drenching the satellite. Launch had to be delayed for years.

89

Prevent Hardware Fratricide

The Problem:

A payload fairing did not open in flight.

The Cause:

The shroud was deployed with two sets of explosive-driven springs. The primary circumferential squib should have fired first, followed by, at 22 millisecond intervals, its backup; the primary longitudinal ordnance; and its backup.

The circumferential cut thrust the nose-cone forward and pulled the longitudinal firing plugs apart. An unfavorable tolerance buildup, plus an unexpectedly large forward motion of the fairing, disconnected several pins. The longitudinal split did not take place.

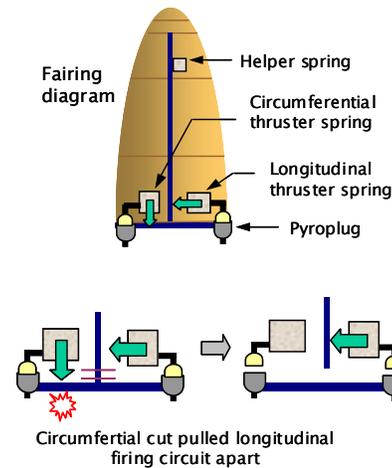
The fix involved adding a heritage locking mechanism to prevent the connector halves from moving apart during firing. When the shroud starts to unlatch forward and outward, lanyards attached to a bracket mounted above the plug pull the fastener open.

Lessons Learned:

- Ensure the neighboring units survive after the primary device operates.
- Qualify ordnance devices in their operational environment.

For more technical information, call Selma Goldstein at (310) 336-1013.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

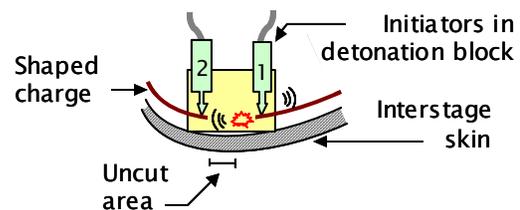


Prior Problems Missed

A review of a previous mission revealed that several non-critical pins had disengaged. Unfortunately, these warning signs were not heeded and the connectors were not redesigned (Lesson 65).

Similar Incident

A launcher used shaped charges to separate the stages. The initiator on one end fired first, disabling the other end of the charge and preventing the structure underneath the damaged initiator from tearing apart. The vehicle jackknifed.



Account for All Loose Materials

The Problem:

A large engine partially melted during a test firing.

The Cause:

Investigators found that a large piece of sealing tape, routinely used during engine assembly, blocked the fuel injector and caused the turbopump to overheat.

The investigation board reprimanded the manufacturer for not having a disciplined process to handle, or account for, loose materials. The processing paperwork was not traceable, making it difficult to know what work was done on which part.

In this case, the build log supposedly documented tape removal and independent verification. The Investigation Board discovered, however, that tape reportedly taken out was repeatedly found during postfire inspection or engine rebuild.

Lessons Learned:

- Make sure loose, nonserialized materials (such as wipe cloth) used during assembly are carefully accounted for.
- Correct the root cause of in-process anomalies (Lesson 32).
- Keep accurate records of all “nonflight” installations.
- Take photos frequently during assembly.
- Design hardware to minimize areas that cannot be easily inspected, and avoid the use of potential contaminants whenever possible.
- Keep hardware closed when access is not needed.
- Review out-of-flow processes to ensure no steps are bypassed (Lesson 64).

For more technical information, call Dana Speece at (310) 336-5021 or Gary Shultz at (310) 336-2342.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Other “Foreign Object Damage” Incidents

- Debris contamination spoiled five foreign launches between 1990 and 1999, including several caused by rags clogging propulsion lines.
- Debris such as paper clips left in RF cavities repeatedly caused test failures on a satellite program. The contractor finally developed an electromagnetic probe to sweep all cavities before they were sealed.
- A jet engine contractor suffered several failures caused by bolts or tools being left inside test units. The management subsequently required an inspector to go inside the inlet to check for debris using a flashlight.

Right after the new procedure was implemented, the engine blew up. The flashlight was left behind. (From “Augustine’s Laws.”)

91

Ensure Critical Systems Are Tolerant of Transient Power Loss

The Problem:

A first-stage engine shut down soon after liftoff.

The Cause:

Immediately before the mishap, the bus current spiked twice. Evidently, a power cable had a breach in its insulation layer, and momentarily grounded. The engine relay box lost power, and numerous relays controlling the propulsion valves dropped out, disabling the engine.

By design, the relays lock on their own contacts during flight, which depends on a continuous supply of electricity to retain their running configuration. If the power is lost, even for an instant, the relays unlatch with no means to recover.

The vulnerability to a transient short had been recognized by the contractor for years. Unfortunately, even though many design improvements had been made elsewhere, such as in the propulsion system, little attention was given to this single failure point.

Lessons Learned:

- Ensure the onboard computer retains “most recent state” information so that if a glitch causes the loss of “present state” data, the vehicle can revert to a survivable configuration.
- Anticipate wiring problems, and provide redundant power sources to critical systems, including lock-in power circuits to prevent hardware reset.
- Recognize the need to address weaknesses in nonpropulsive systems.

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

A Lesson Not Learned

After this incident, the contractor redesigned the 30-year old control electronics to provide redundant power and guidance. A sister launch vehicle program, however, did not make a similar change.

Years later, the second program suffered a failure. Apparently, a defective power cable shorted intermittently, causing the guidance computer to reset and the inertial measurement unit to lose reference.

The launcher had miles of wires—forty-four repairs had been made on this particular vehicle alone. In retrospect, it was clearly impossible to inspect out every wiring defect, and the decision not to provide redundant power proved costly.



Cabling defects led to the most costly unmanned launch failure

Rigorously Determine the Root Causes of Test Failures

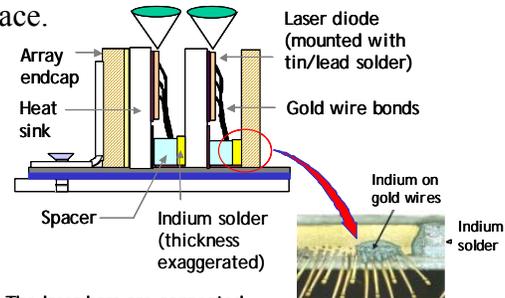
The Problem:

The primary laser in an instrument failed after a month in space.

The Cause:

The laser pump consisted of several diodes mounted on heatsinks, soldered together into stacks. Apparently, the indium solder contaminated the gold bondwires, forming an insulating layer of intermetallics. In orbit, the corroded bondwires suffered from thermo-mechanical fatigue and cracked.

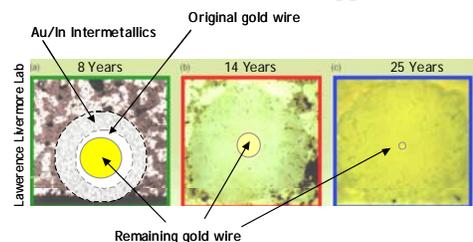
Lasers have not flown in space often. The design of this laser was derived from a previous program and was procured commercially. In retrospect, the vendor's internal processes and controls were not up to par for space applications. The new design was more vulnerable because current density in the contaminated bondwires increased by 40 percent, intensifying thermal loads in the wires. Several years of launch delay made the degradation worse.



The laser bars are connected in series through gold plating, wirebonds, and indium solder.

Laser Array Stacks (Simplified)

Au/In Reaction in Terrestrial Applications



During qualification, the bondwires broke several times. The vendor replaced the defective components and asserted that the failures would not recur. A laboratory analysis, which would have discovered the root problem, was requested but not carried out.

Lessons Learned:

- New technologies require rigorous qualification, analysis of design changes, and a thorough understanding of failure modes.
- Audit a vendor's manufacturing process, conduct destructive physical analysis of sample parts, and ascertain the root causes of all anomalies.
- Review the materials and processes for each new application drawing.
- Guard against known materials incompatibilities (gold/tin intermetallics can embrittle solder joints, for example).

For more technical information, call Renny Fields at (310) 336-6973.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

93

Always Ascertain the Direction of Current Flow

The Problem:

Contact with a satellite was lost soon after launch.

The Cause:

The satellite consisted of a domestic instrument module and a foreign service module. A design mistake in the foreign unit caused the solar panels to be connected backwards.

The domestic instrument supplier, in charge of system integration, checked the interface between the solar panels and battery, but only verified the magnitude of current, not its direction—engineers might have become confused as to how the current should flow because the foreign unit grounded positively but the American unit grounded negatively.

Once in orbit, the battery drained, ruining the mission.

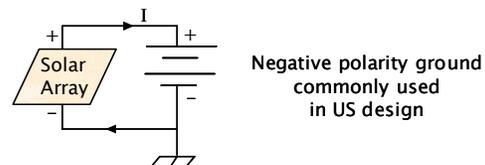
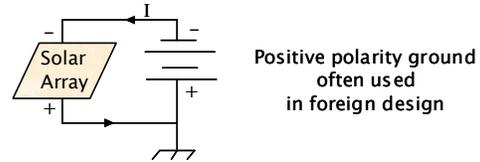
"It's always the simple stuff that kills you," lamented the lead engineer.

Lessons Learned:

- Make sure that engineers understand how the system or component should function during test.
- Thoroughly verify interfaces of subcontracted items, particularly when the suppliers use different engineering conventions.
- Use an engineering model to verify interfaces early.

For more technical information, call Ron Williamson at (310) 336-2149.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Polarity Confusion

A Similar Incident (from "Augustine's Laws")

A preflight check found two hardware modules wired in the opposite polarity. Both subcontractors reversed their cables. The launch failed.

Provide Debug Features in Flight Software to Assist Anomaly Resolution

The Problem:

An interplanetary probe lost some scientific data due to occasional system resets.

The Cause:

Driven by demanding mission requirements, the designers used a commercial, realtime, multiple-tasking operating system.

An esoteric “priority inversion” problem took place during science operation and caused some data loss. This glitch was not caught on the ground because the Earth-pointing antenna performed better than expected, allowing more frequent downlinks than originally planned.

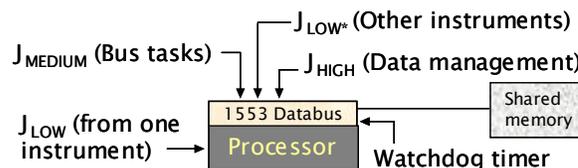
Fortunately, debugging tools, written during code development, were embedded in the software. With extensive support of the vendor, the project was able to reproduce the problem in the laboratory and identify the cause. A quick fix allowed the mission to successfully conclude.

Lessons Learned:

- Ensure that commercial software, especially the operating system, allows access to internal information and is compatible with development debug tools.
- Test for off-nominal conditions, both “better” and “worse” than expected (for example, at higher throughput rate), to see if the system misbehaves.
- Leave debug capabilities embedded in the operational system.
- Shared functions must be thoroughly tested, especially for timing.

For more technical information, call Suellen Eslinger at (310) 336-2906.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



The Priority Inversion Problem

Because the bus and instruments share the processor, job allocation is vital. The highest priority is given to data management, followed by bus tasks and by science activities. If data management tasks cannot complete within the watchdog’s 125 millisecond cycle, an anomaly is assumed and the computer is reset.

Data from the bus and payloads flow through a 1553 data bus, but one instrument is processed directly. That sensor shares a software function with the transaction manager—not a prudent design but normally not a problem. Access to this resource was controlled with a key. If a data manager job (J_{HIGH}) starts late in the cycle, it may find a job from this instrument (J_{LOW}) still in process. If J_{HIGH} also requires the shared software function, it must pause for the key.

When a communication job (J_{MEDIUM}) initiates during the short interval, however, it preempts J_{LOW} , preventing the key’s release. The system watchdog timer starts the next cycle, finds J_{HIGH} unfinished, and resets the system.

Turning on “priority inheritance” options for that particular thread (giving high priority to J_{LOW} in light of jobs blocked by it) solves this problem. This option is not normally used as default due to performance concerns.

Ensure Heritage Designs Can Operate in the New Application Environment

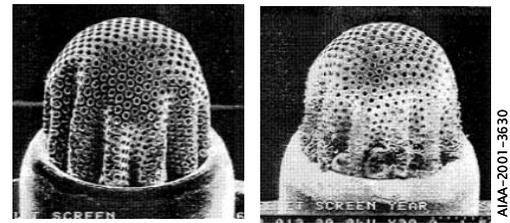
The Problem:

An interplanetary probe mysteriously failed.

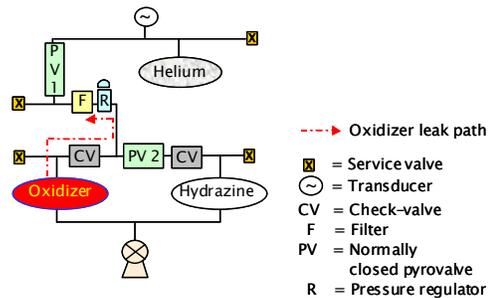
The Cause:

The incident occurred when the vehicle, having completed a year-long flight, pressurized its propulsion system in preparation for an orbit-insertion burn. The propulsion system had been used for apogee boosting in numerous GEO satellites without incident. Extensive testing could not reproduce the failure.

Years went by. Then, in a program review, a propulsion expert heard that a commercial restrictor contained a brazing alloy that is incompatible with oxidizer vapors. The same part had been used in the failed spacecraft; a failure mechanism finally dawned on him.



Restrictor exposed to oxidizer vapor
 (a) After 30 days (b) After one year showing extensive corrosion



Simplified Propulsion Schematic

Evidently, the oxidizer vapor can pass through the check-valves and cause a very slow corrosion. Normally a negligible amount, so much debris would accumulate on this long-duration mission that when the pyrovalves fired, the debris was shaken into the restrictor orifices and kept the regulators open. Helium rushed out, bursting the line.

The incompatibility was not recognized at the time because the restrictor’s materials list did not include the braze. In fact, if the expert had not made the connection, two more probes would have been launched with the same flaw.

Lessons Learned:

- Avoid relying on short-term tests (days to months) to confirm long-term reliability.
- Audit vendor material lists to ensure completeness.
- Account for vapor diffusion in propulsion subsystem design.

For more technical information, call Mark Mueller at (310) 336-5081.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

Tests Must Independently Verify Development Results

The Problem:

A space telescope was out of focus.

The Cause:

The telescope's primary mirror was polished with the aid of a "null corrector." Lights that are shone on a perfect mirror, when reflected through the corrector, should form straight interference patterns.

The corrector was set up with a positioning rod capped on one end. A light beam passed through a small aperture in the cap to focus on the rod's tip, and a lens was placed at the other end of the rod.

Unfortunately, a speck of antireflective coating chipped off the rod's cap, and the focusing beam was aimed at the cap instead of the rod. The lens was misplaced; the mirror was misshapen.

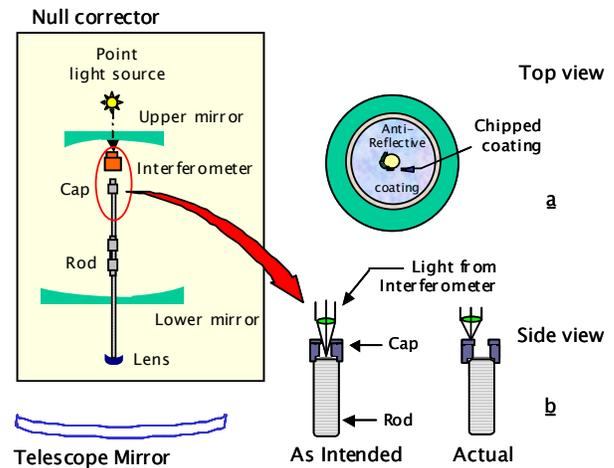
Because the contractor used the corrector not only as a manufacturing tool but also as the sole referee standard, it could not detect the mistake. In fact, each of two pieces of auxiliary optics suggested gross errors. However, confident that the new-technology corrector was better, the engineers ignored the red flags.

Lessons Learned:

- Use simple tools to crosscheck elaborate tests.
- Scrutinize test equipment, analysis, or algorithms reused from design or manufacturing for possible single-point failure.

For more technical information, call Julie White at (310) 416-7229.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Mirror Manufacturing Process (Simplified)

Missing coating (view a above) near the cap aperture caused the operator to aim the light at the cap instead of at the rod (view b above).

Operators Failing to Call Attention to the Problem

The misfocusing prevented the metering rod from reaching the lens, but the technicians simply extended the rod by inserting a few washers.

"That in itself should have alerted people...because clearly there should not be a need for any unexpected washers to be added," said the investigation board.

97

Control Hardware and Software Configurations Before, During, and After Tests

The Problem:

A satellite pointed toward the Sun with the wrong axis.

The Cause:

As the satellite exited eclipse for the first time, it should have pointed a vector 35 degrees off the z-axis toward the Sun. Instead, it wobbled, while pointing the x-axis to the Sun. Fortunately, one of the solar wings was illuminated, giving the engineers time to recover.

The next day, an examination of a photo taken at the launch site revealed that two Sun sensors were mounted ninety degrees off. A software change quickly fixed the problem.

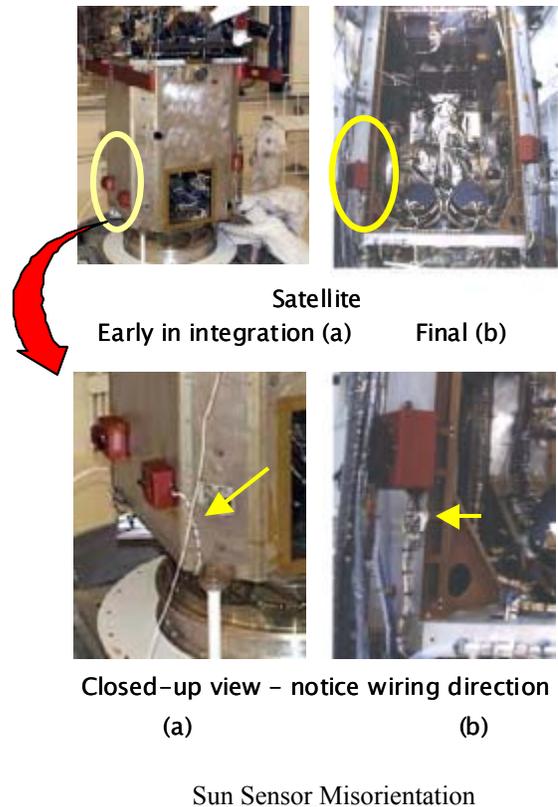
The Sun sensors were mounted on the main access panel in the intended direction during verification testing, before the panel was attached to the spacecraft. When the panel was being installed, however, the mechanical engineers found that the sensor cables were too short to mount the sensors “as hung.” Seeing no control document on the sensor configuration, they turned the sensors sideways, without informing the guidance and control (G&C) engineers of the change.

Lessons Learned:

- Always ascertain G&C actuator phasing (Lessons 53, 60, 80).
- Ensure domain engineers own all aspects of their subsystems.
- Conduct end-to-end testing in the flight configuration.
- Take plenty of photographs during assembly.
- Document G&C subsystem-level alignment. See Guideline GD-ED-2211 from NASA Technical Memorandum 4322A, for example.

For more technical information, call Geoffrey Smit at (310) 336-1602.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Guard Against Post-Firing Conduction of Pyro Initiators

The Problem:

The redundant memory board on a spacecraft filed.

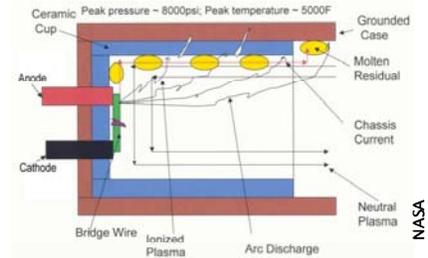
The Cause:

During an orbit insertion maneuver, the satellite fired several explosive bolts to jettison a solid rocket.

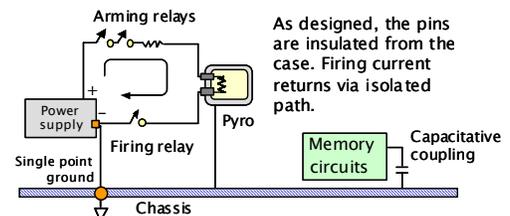
The burning pyro propellant formed a conductive plasma, shorting to the chassis-grounded case. A voltage surge rippled through the input protection diode in the backup memory circuit, causing upsets. If the primary memory had latched, the mission could have failed.

Lessons Learned:

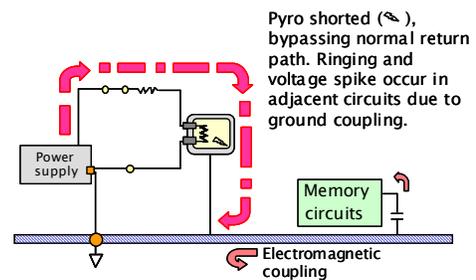
- Protect firing circuits against sneak currents and line-to-ground shorts. Components such as step motors and pyro circuits that experience sudden current changes should be isolated from all other current-carrying circuits including electrical power, electrical control, RF transmission lines, and monitoring circuitry. For additional information, see *Electromagnetic Interference Analysis of Circuit Transients*, NASA Preferred Reliability Practice No. PD-AP-1308, for example.
- Check circuit designs against *Electroexplosive Subsystem Safety Requirements and Test Methods for Space Systems* (MIL-STD-1576), *NASA Standard Initiator User's Guide* (JSC-28596A), and *Electrical Grounding Architecture for Unmanned Spacecraft* (NASA-HDBK-4001).



Post-firing Conductive Mechanism



As designed, the pins are insulated from the case. Firing current returns via isolated path.



Pyro shorted (∞), bypassing normal return path. Ringing and voltage spike occur in adjacent circuits due to ground coupling.

Simplified Bus Grounding Architecture

Other Post-Fire Conduction Conditions

Post-fire plasma shorts can drain batteries. See *Journal of Spacecraft and Rockets*, 36, 586-590 (1999).

Drive elements can be disabled by residual current, and should be inspected after ground live tests. In one case, an inspection found a damaged fusing resistor, which would have prevented in-flight firing.

Between 3% and 5% of firings result in conduction.

For more technical information, call Ron Williamson at (310) 336-2149.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

99

Have the Model's Originator Check the Analysis

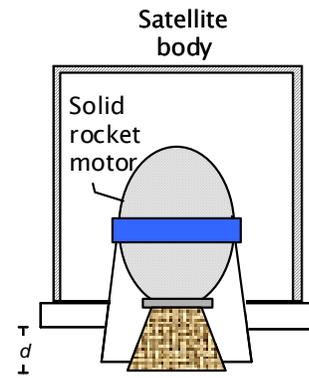
The Problem:

A spacecraft broke up after firing its embedded solid rocket motor.

The Cause:

The contractor bought the motor off-the-shelf and learned that another company had flown a similar design. It obtained that company's thermostructural analysis, but did not refine its own model, nor ask the original analyst for support.

The analyst had also presented the results of his analysis in a conference. A diagram published in the proceedings showed the nozzle was deeply buried inside the spacecraft (the distance from the structural base to the nozzle mouth, $d_{\text{heritage mission, reported}} = 6.03$ inches). The engineers used this information to justify the final design, which submerged the motor deeper ($d_{\text{new mission}} = 4.95$ inches) and did not thoroughly shield the spacecraft against plume heating.



Satellite Diagram (Simplified)

The accident investigation board subsequently found that the spacecraft would suffer massive heating from the motor exhaust plume and disintegrate. The motor vendor estimated that heating would be almost two orders of magnitude higher than expected by the contractor. Why was the design, qualified by similarity, so far off?

It turned out that the motor in the previous mission was actually more extended ($d_{\text{heritage mission, actual}} = 11.03$ inches). The distance shown in the conference paper was an error! The author knew about the mistake but unfortunately did not know the contractor relied on his publication instead of the model, which did not include this erroneous diagram.

Lessons Learned:

- Double check all analysis models, assumptions, methods, and predictions.
- Develop a rigorous process for using experience as a basis for accepting further designs and equipment.
- Have the original analyst review final product (Lesson 26).
- Make sure key subcontractors accept how their product is being used.

For more technical information, call Dan Perez at (310) 336-2734.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.

100

Make Sure Safety Mechanisms Are Truly Independent

The Problem:

A satellite suffered a near-catastrophic short.

The Cause:

Following launch, the spacecraft turned on a set of wax heaters for three minutes to activate the release actuators on the solar arrays.

Later, a design error in a field-programmable gate array (FPGA) inside the power controller caused the primary heaters to be reactivated. After ten minutes, the overheating primary elements shorted to the secondary elements, and subsequently to the bus structure. The short circuits drew hundreds of watts, at a current level several times the power board's design limit.

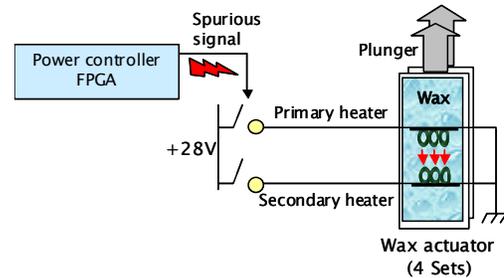
Fortunately, the heater traces burned open, saving the power distribution unit from permanent damage. Otherwise, the mission would have ended.

Lessons Learned:

- Ensure safing mechanisms will prevent one design error from causing a cascade of irreversible failures (Lesson 77). In this case, one error could have activated all the heaters, and the solar arrays might have been deployed prematurely.
- Check for failure mechanisms during extended operation even if that is not the intended application. If prolonged operation leads to catastrophic failure, provide circuit interrupts, time-out protection, or a graceful degradation mechanism (Lesson 19, 71).
- Review special design requirements for FPGAs (Lesson 77).

For more technical information, call Peter Carian at (310) 336-8215.

For comments on the Aerospace Lessons Learned Program, including background specifics, call Paul Cheng at (310) 336-8222.



Actuator Diagram (Simplified)

Ensure Independent Safety Mechanisms

The ARM and FIRE relays in Diagram (a) below can prematurely close on one FPGA error. Separate drivers (b) should be used.

