

STEREO Fault Protection Challenges and Lessons Learned

George J. Cancro and Michael D. Trela

Performing the fault protection task for the Solar TERrestrial RELations Observatory (STEREO) spacecraft forced the design team to deal with multiple issues that are either not emphasized or not encountered on single-spacecraft missions. The decisions made during STEREO development regarding redundancy, test philosophy, test execution, staffing, and data handling were directly impacted by the number of spacecraft. Looking back over the STEREO development, several lessons learned emerged in the area of fault protection for the design, testing, and operations of multispacecraft missions. This paper captures and documents seven lessons along with rationale and examples from the STEREO development in the hope that future multispacecraft missions can benefit from this information.

INTRODUCTION

Fault protection is defined as “the use of cooperative design of flight and ground elements to detect and respond to perceived spacecraft faults.”¹ The goal of fault protection is to achieve mission reliability objectives without violating program resources. Fault protection must achieve this goal by balancing project risk and the cost of developing, testing, and operating proposed fault detectors and remedies.

With multispacecraft missions, such as the twin Solar TERrestrial RELations Observatory (STEREO) spacecraft, the goal of fault protection is further complicated. Design,

testing, and operation decisions regarding redundancy, test philosophy, and operations staffing, for example, are directly impacted by the number of spacecraft. Performing the fault protection task for STEREO forced the design team to deal with multiple issues that are either not emphasized or not encountered on single-spacecraft missions.

During performance of the fault protection task for the STEREO spacecraft, several lessons emerged that will be beneficial to future multispacecraft missions. This paper discusses these lessons, which are divided into three sections based on three phases of develop-

ment: design, testing, and operation. For each phase, a brief summary of fault protection duties is provided. Following the summary, two to three detailed lessons from that phase are discussed. For each lesson, a high-level statement is made, and examples from STEREO are provided as the rationale for advocating the lesson.

DESIGNING TWO OF A KIND

Fault protection during the design phase of a program consists of analyzing potential faults and developing the architecture that can deal best with these faults throughout the entire mission. Fault protection architecture includes the organization, design, and interrelationships of hardware, software, and operations procedures to detect and respond to perceived spacecraft faults. At this phase, decisions about hardware redundancy and critical sequence design must be made to maximize reliability within project resources.

For the STEREO program, developing and launching two spacecraft with low project resource levels (both mass and dollars) forced early designs toward single-string redundancy schemes. To increase overall reliability under these conditions, functional and selective redundancies were applied.

The following sections discuss the lessons learned in applying functional and selective redundancy on the

STEREO program, lessons that can apply to any future multispacecraft mission.

Effects of Combined Science

Each STEREO spacecraft has 11 separate instruments in three different instrument suites. No hardware redundancy was allocated to either spacecraft because of project resource constraints; therefore, all instruments were single string. In single-spacecraft missions, science instruments can exhibit functional redundancy to the science objectives (i.e., the observations from one instrument or another can achieve the same level 1 science objective).

Because of the large number of instruments and the desire for combined science (i.e., to combine observations of similar instruments on both spacecraft simultaneously) on STEREO, an instrument-redundancy matrix comparing all science instruments on both spacecraft to the science objectives was generated. A section of this matrix is shown in Fig. 1.

For each set of science objectives, multiple instrument redundancy options were generated to achieve the science objective. The table exposes a twofold increase in functional redundancy with two spacecraft. For example, with STEREO, science objective 1A could be met with the COR1, EUVI, or S/WAVES instruments on either spacecraft. In addition, single-string

Suite		Required Science Measurement																				Redundancy Options			
		SECCHI										S/WAVES		IMPACT											
		COR1		EUVI		COR2		HI1		HI2		S/WAVES		MAG		SWEA		STE-UID		SEP			PLASTIC		
Instrument	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B			
Spacecraft	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B	A	B			
Science Objectives	1A		O																				1		
				O																				2	
					O																				3
						O																			4
												O													5
													O												6
	1B		O	X	O																			1	
			X	O		O																		2	
			X		X	X																		3	
			X	X	X																			4	
			O	X			O																	5	
			X	O			O																	6	
2C1		X	O			O																	7		
		X				X	X																8		
		X				X	X																9		
				O	X	O																	10		
				X	O		O																11		
				X		X	X																12		
F																		O	X	X			1		
																		X	O		X		2		
																		X		O	X		3		
																			X	X	O		4		
																							5		
																			O	X			6		
G																			X	O			1		
																							2		
																							3		
																							4		
																							5		
																							6		

Figure 1. Portion of the STEREO instrument-redundancy matrix showing how combinations of instruments from both spacecraft can meet science objectives. O, required for 240 days; X, required for 2 years. Sun–Earth Connection Coronal and Heliospheric Investigation (SECCHI) instruments: COR1 and -2, coronagraphs 1 and 2; EUVI, extreme ultraviolet imaging telescope; HI1 and -2, heliospheric imagers 1 and 2. S/WAVES, STEREO/WAVES. *In situ* Measurements of PArticles and CME Transients (IMPACT) instruments: MAG, magnetometer; SWEA, solar wind electron analyzer; STE-U/D, suprathermal electron telescope (upstream and downstream sensor heads); SEP, solar electron proton telescope; PLASTIC, PLAsma and SupraThermal Ion Composition Investigation.

instruments of both spacecraft form a physical redundancy to the science objectives. This is shown in the first two rows of objectives F and G, where the SEP instrument from Spacecraft A or Spacecraft B can be used to meet the objectives.

Using functional and physical redundancy across all spacecraft in a multispacecraft mission can increase the reliability of the mission and may enable reduction in cross-strapping and redundancy on individual spacecraft.

Lesson 1

Exploit the presence of multiple instruments on multiple spacecraft in combined science missions to increase overall mission reliability

Multiple Systems Tied Together

Developing a redundancy plan for the STEREO spacecraft involved developing a balance between reliability and cost. The cost of the STEREO mission precluded full redundancy; therefore, the fault protection plan used selective redundancy to achieve the greatest reliability for available resources.

Deciding what to make redundant in a selective-redundancy scheme is a difficult issue. For multispacecraft missions, the first essential place to allocate limited redundancy is to any single fault that can eliminate both spacecraft. Because STEREO launched with both observatories stacked upon each other, failure to separate the stack from the launch vehicle and failure to separate the spacecraft from each other are two prime examples of failures that eliminate both spacecraft.

To mitigate these failures, a reliability block diagram of the separation system was developed to ensure that the detection of stack separation and the initiation of spacecraft separation were free of single-point failures, as shown in Fig. 2. Redundancy was placed within the STEREO power subsystem to ensure that the detection and initiation contained no single-point failures. To mitigate against failure to separate the stack, requirements

were added to allow the bottom spacecraft (Spacecraft B) to release Spacecraft A upon command from the ground.

Lesson 2

Focus fault architecture and redundancy on single faults that can eliminate more than one spacecraft.

TESTING TWO OF A KIND

Fault protection during the testing phase is ensured by challenging the spacecraft at a system level to verify that the spacecraft can survive and continue the mission in light of one or more failures. This activity includes test planning, test execution, and test review.

For STEREO, testing two spacecraft provided some benefits for and some issues with the fault protection effort. The following sections discuss the lessons learned during the testing phase of the STEREO program.

Multispacecraft Test Planning

Testing two satellites that were nearly identical posed interesting questions regarding cost and schedule savings versus risk reduction with robust spacecraft testing. The type and amount of testing that each satellite would endure had to be determined.

Several space missions have performed testing programs with more manufacturing-like methods, in which a full set of system and environmental tests was performed on the first few spacecraft and then a small subset of these tests was run on later production models to demonstrate basic functionality or workmanship. A good example of this was the Globalstar test program.²

The STEREO program selected a test philosophy called “core and non-core.” Core tests were developed to test the base level of fault protection that could keep the spacecraft safe until the next available ground contact. These tests covered the “bootstrap” functionality of the STEREO fault protection system designed to protect the spacecraft if all other responses failed to

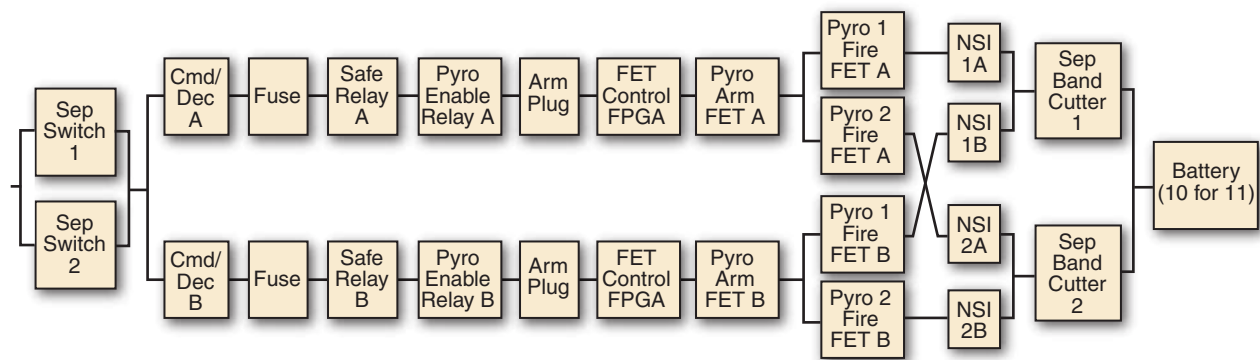


Figure 2. Reliability block diagram for separation of the two STEREO observatories. Cmd/Dec, command decoder; FET, field effect transistor; FPGA, field-programmable gate array; NSI, NASA Standard Initiator; Sep, separation.

remedy a problem. The non-core tests represented all other of the required fault protection functionality. The core test set had to be run on both spacecraft as a prerequisite for launch and also formed the regression suite. The non-core test set could be run on either spacecraft and did not have to be repeated if run successfully on one spacecraft. A total of 31 fault protection tests were generated (10 core and 21 non-core).

The core and non-core philosophy provided efficiency and flexibility, as described in the next section, by being able to take full advantage of the parallel STEREO integration and testing (I&T) program; however, the philosophy was not clear on the next step to take after a non-core test failed. Early in I&T, a non-core test involving a fault of the spacecraft onboard oscillator (time source) failed to correct the problem resulting in the spacecraft becoming nonresponsive. The root cause, determined to be a software bug, was fixed, and the test was rerun successfully. Because the test was non-core and succeeded, the test was never run again. However, later on in the program, the hardware simulator team used the fault protection core and non-core tests to validate the simulator. During the simulator validation, the oscillator test failed again. The root cause was determined to be later software releases that exposed the original flaw. Fault protection regression had been executed with the new software release, but because the oscillator test was non-core, it was not run as part of the regression.

The lesson here is a difficult one. The test was classified as non-core because an operational work-around existed; however, one can argue that the nonresponsive failure should have resulted in the addition of the oscillator test to the core regression suite. Either way, the test philosophy needs to be carefully developed for the effects and risks of test failure, and regression should be developed and communicated early.

Lesson 3

Develop a test philosophy early that not only discusses which test will be executed on which spacecraft but also plans for when tests fail and determines the regression test suite of the most critical fault scenarios.

Multispacecraft Test Execution/Review

The core and non-core philosophy was chosen because both spacecraft were being assembled simultaneously in I&T. A serial I&T would have been more conducive to a manufacturing-like approach. The parallel STEREO I&T, however, enabled the core and non-core test plan to use both spacecraft as test platforms, resulting in a twofold increase in test time available to the fault protection team. The parallel I&T also enabled the fault protection test team to avoid schedule hits associated with down time or periods when boxes had to be removed from the spacecraft. Non-core tests could be moved from one spacecraft to the next without slowing the schedule.

By the end of I&T, all 31 planned fault protection tests were executed as described in the original test plan. Considering that the 10 core regression tests were repeated three times on both spacecraft, all 21 non-core tests were executed on either spacecraft, and several tests had to be repeated because they did not achieve all test objectives on the first execution on the spacecraft, a total of 113 tests were executed on the STEREO program.

Per our test-review plan, all of the 113 test runs had to be reviewed to ensure that all objectives were met, all problems and failures were understood, and any behavior outside of the test plan was documented and understood. Because each test took from 4 to 8 h to execute and each STEREO spacecraft produces approximately 9 million telemetry values per hour, >4 billion telemetry values had to be reviewed.

Faced with this daunting task, the STEREO fault protection team developed a new telemetry review tool called the STEREO Autonomy Visualizer (SAV)³ that enabled the team to examine up to 2500 telemetry values simultaneously and select and plot 100 telemetry points 30 min before and after the current review time. The user controls the current time in a TiVo-like manner, providing real-time random access to all the telemetry generated during the test. The effect of the addition of the tool into the test-review effort was significant, as demonstrated in Fig. 3.

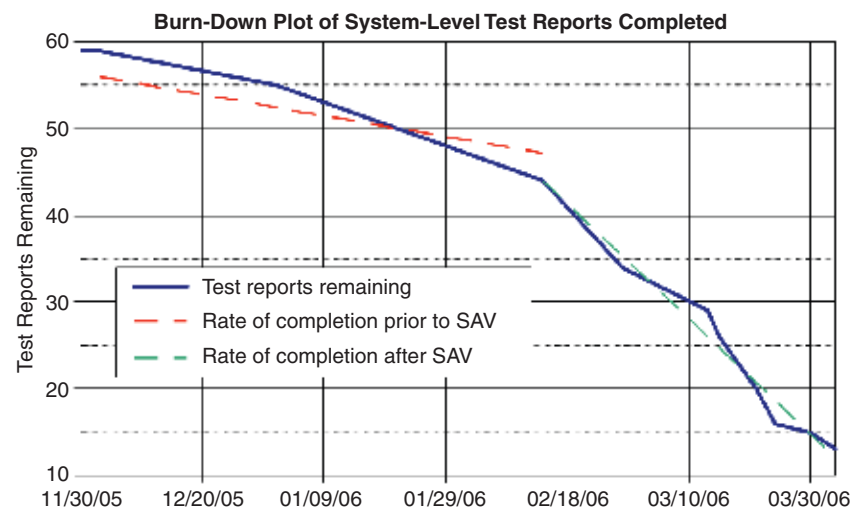


Figure 3. Test-review improvement with SAV.

Lesson 4

Take advantage of multiple spacecraft to achieve an n -fold increase in test time but also plan methods or automation to analyze n -fold more test data.

OPERATING TWO OF A KIND

Fault protection in the operations phase means engineering support is provided during critical events, both for diagnosis of on-orbit anomalies to determine failure and root cause and for maintaining and updating onboard fault protection routines.

For STEREO, performing fault protection during the operations phases revealed our poor assumptions that single teams and common implementations are sufficient for performing fault protection for multispacecraft missions. At the same time, however, flying multiples enhanced our ability to determine failure. The following sections discuss the lessons learned during the operations phase of the STEREO mission.

Issues with Staffing

Up until launch, the STEREO fault protection team was sized as if one spacecraft was under development, not two. With good planning and automation, the smaller team was able to succeed in accomplishing all objectives from design to testing. Nominally, it was thought that one individual would be responsible for monitoring the current state of the autonomy system on both spacecraft during launch. However, as the team prepared for launch, the idea that different faults could occur simultaneously

on both spacecraft drove an increase in staff to support critical operations.

Before, during, and after the launch, there were numerous activities that demanded the attention of the fault protection team. However, it was found that the long time span in which engineers were required to monitor and support the launch preparation and postlaunch activities led to the creation of three shifts: prelaunch (Launch – 13 h until Launch – 5 h), launch (Launch – 5 h until Launch + 7 h), and postlaunch (Launch + 7 h until Launch + 19 h). More importantly, the team decided that the work required for each spacecraft necessitated staffing one engineer per spacecraft on each shift. Both prelaunch and postlaunch activities are similar for both spacecraft but may or may not be synchronous in time; therefore, a single engineer responsible for both spacecraft might become distracted or unaware of the current state of one of the spacecraft. For this reason, it seemed best that one engineer remain focused on each spacecraft.

Lesson 5

Staff launch and simultaneous, critical events with a dedicated engineer for each spacecraft on each shift.

Advantage of Operating Multiples

One great advantage to flying multiple identical spacecraft is the ability to compare data from the similar operations of different spacecraft to gain insight on all spacecraft. A perfect example of this was the high-gain antenna (HGA) deployment as shown in Fig. 4.

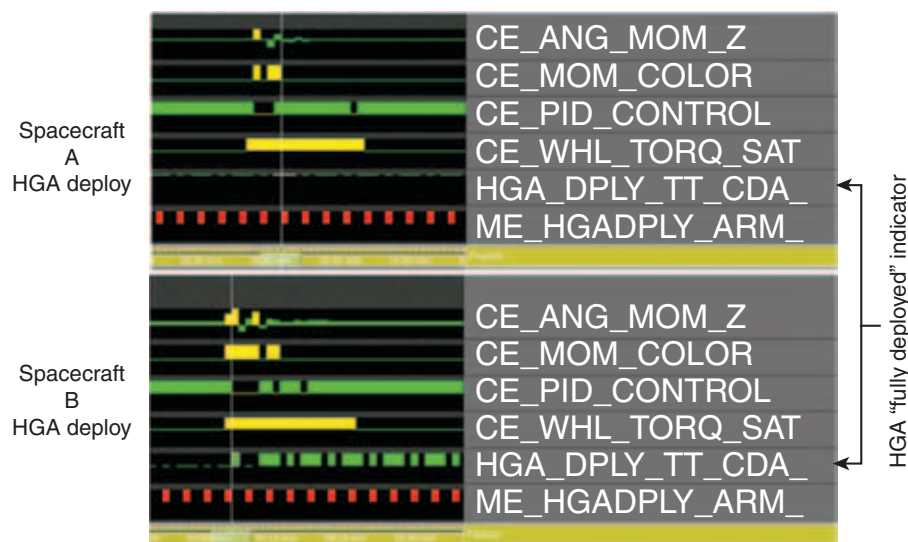


Figure 4. HGA deployment on both spacecraft shows how STEREO A has an indicator failure and not a failed deployment.

STEREO B deployed its HGA during the postlaunch period. The deployment was detected by hardware switches indicating the release of the dish and achievement of the fully deployed state. The deployment can also be detected in the guidance and control states of the spacecraft, where the deployment induces a rotation in the spacecraft body, an increase in overall momentum, and a measured response in the onboard reaction wheels to control the momentum increase, as shown in the plots at the top of Fig. 4.

When STEREO A attempted the deployment of its HGA, the “fully deployed” hardware switches indicated that the HGA had not fully deployed. The reading of this switch caused alarm because it was uncertain whether the switch was providing a false-positive or a false-negative reading. This issue was solved by comparing the body rate data of each spacecraft at the time of the HGA deployment. The signature of the body rates and other G&C states were similar on both spacecraft, however, indicating that the HGA had fully deployed. Using the data from both spacecraft simultaneously enabled engineers to deduce that one of the switches on Spacecraft A had failed. Had the data from the second spacecraft not been readily available, it is likely that the operations team would have delayed using the HGA for a significant period of time until a series of checkout tests could be performed to verify whether the deployment had been successful.

The ability to compare data plots across spacecraft provides tremendous additional information. To use this information, mechanisms must be in place to rapidly make the comparison. In addition, the mechanisms must have the ability to manipulate the telemetry to align events rather than examine the telemetry based on time, because even identical spacecraft cannot conduct operations synchronously. In the HGA example, the time between events was >30 min. The fault protection team had to align the G&C telemetry responses to properly interpret the event.

Lesson 6

Use data from one spacecraft in a set to diagnose a fault on another spacecraft in the set.

Deviation from Commonality

All efforts to maintain commonality between multiple spacecraft missions can be easily disrupted following the first on-orbit failure. For STEREO, commonality was broken 6 months after launch with the failure of the inertial measurement unit (IMU) on Spacecraft A.

The IMU failure on Spacecraft A forced the spacecraft to operate on the backup IMU while Spacecraft B

continued to operate on the prime IMU. Because the STEREO fault protection team did not fully understand the mechanism causing the IMU failure, the team was forced to devise a way for Spacecraft A to survive a second IMU fault. Before the solution was loaded to Spacecraft A, the team considered whether the same solution should be loaded to Spacecraft B strictly for the purposes of maintaining commonality.

In the end, the team found that it was beneficial to keep the configuration as similar as possible on both spacecraft because the new changes did not add any significant risk to Spacecraft B. This policy of maintaining commonality is the safest for multiple spacecraft because “noncommon” spacecraft lead to “noncommon” operations, during which operator error could be more common because an operations caveat is forgotten or applied to the wrong spacecraft. However, it is evident that maintaining commonality may not always be possible; if a situation requiring significant modifications to the normal operating state of one spacecraft arises, the risk associated with deploying the new operating state on the other spacecraft could be greater than the risk of managing two different spacecraft configurations.

Preparing for deviations involves not only developing philosophies on the steps to take when on-orbit failures occur on one of the spacecraft but also evaluating tools, parameters, and documentation. For example, STEREO maintains documentation of the onboard autonomy system on a website to enable fault protection, systems, and operations staff to understand detections and reactions to faults. However, this single website was never designed to differentiate between Spacecraft A and Spacecraft B. This issue illustrates why preparing for deviations early in the design phase is most prudent.

Lesson 7

Departure from commonality during operations is inevitable.

CONCLUSIONS

The design, testing, and operation of the twin STEREO spacecraft have yielded several interesting lessons for the development and execution of future fault protection efforts. As NASA moves forward, we expect more and more missions to contain multiple spacecraft. Consideration of these lessons early in any future effort will be beneficial to the program and ensure that future fault protection teams are adequately prepared for the differences between single-spacecraft and multispacecraft missions.

REFERENCES

- ¹National Aeronautics and Space Administration, *Fault Protection*, NASA Preferred Reliability Practices, Practice Number PD-EC-1243 (Oct 1995).
- ²Costabile, V., Discepoli, A., Fiorentino, C., and Morelli, G., "New Spacecraft Assembly, Integration and Test Approach for Globalstar Satellite Constellation," in *Proc. 16th International Communications Satellite Systems Conf.*, American Institute of Aeronautics and Astronautics (AIAA), Washington, DC, pp. 1320–1330 (1996).
- ³Cancro, G., Turner, R., Nguyen, L., Li, A., Sibol, D., Gersh, J., Piatko, C., Montemayor, J., and McKerracher, P., "Innovative Interactive Visualization System for Analyzing Spacecraft Telemetry," in *Proc. IEEE Aerospace Conf.*, Big Sky, MT, 10.1109/AERO.2007.352952 (2007).

The Authors

George J. Cancro is the Assistant Group Supervisor of the Embedded Applications Group in the Space Department. He holds a B.S. in engineering science from The Pennsylvania State University and an M.S. in mechanical engineering–astronautics from the George Washington University. Before joining APL in 2002, he worked at the NASA Jet Propulsion Laboratory and NASA Langley Research Center on projects such as Mars Global Surveyor Aerobraking and the Dawn Mission to Vesta and Ceres. Since joining APL, he has worked as a System Engineer on the MESSENGER (Mercury Surface, Space Environment, Geochemistry, and Ranging), New Horizons, and STEREO missions, as a Project Manager for the NASA Small Satellite (SmallSat) project, and as a Principal Investigator of two research projects in the areas of autonomy and telemetry visualization. He currently is Principal Investigator of a research project investigating spacecraft tactical commanding and an advisor to the NASA Constellation program in the area of fault detection, isolation, and recovery. His areas of interest include modular software, hardware/software architectures, fault protection, and spacecraft autonomy.



George J. Cancro



Michael D. Trela

Michael D. Trela has been a member of the Space Systems Applications Group at APL since 2005. He received a B.S. in aerospace engineering from the University of Notre Dame in 2004 and an M.S. in aeronautics and astronautics from Stanford University in 2006. He has worked as a Fault Protection Engineer on numerous APL missions, including STEREO, MESSENGER, New Horizons, and Solar Probe. He also works as a Propulsion Analyst for the MESSENGER program. He currently is the Principal Investigator of a research project to improve systems engineering methods for verification, validation, and testing. For further information on the work reported here, contact George Cancro. His email address is george.cancro@jhuapl.edu.