The NEAR Rendezvous Burn Anomaly of December 1998

Final Report of the NEAR (Near Earth Asteroid Rendezvous) Anomaly Review Board

November 1999

The NEAR Rendezvous Burn Anomaly of December 1998

Final Report of the NEAR (Near Earth Asteroid Rendezvous) Anomaly Review Board

November 1999

©1999 The Johns Hopkins University Applied Physics Laboratory
11100 Johns Hopkins Road, Laurel, MD 20723-6099

NEAR Anomaly Review Board

/s/

Eric J. Hoffman, Chairman
JHU/Applied Physics Laboratory

/s/

Charles J. Gay
Litton Amecom

/s/

Ward L. Ebert
JHU/Applied Physics Laboratory

/s/

Chris P. Jones
NASA Jet Propulsion Laboratory

/s/

Michael D. Femiano
NASA Goddard Space Flight Center

/s/

Philip J. Luers
NASA Goddard Space Flight Center

/s/

H. Richard Freeman
NASA Goddard Space Flight Center

/s/

James G. Palmer
JHU/Applied Physics Laboratory

CONTENTS

EXECUTIVE SUMMARY 1

INTRODUCTION 2

- What Was Supposed to Have Happened 2
- Mission Operations Planning and Review 2
- What Was Observed to Happen 3

DISCUSSION 4

- Reconstructed Timeline 4
- Overview of Events 6
- Simulation Results 10
- A Walk Through the Fault Tree 14
  - *High Actual Momentum (Branch 1.1 of the Fault Tree)* 14
    - Momentum Stored in Fluids, or Slosh (Branch 1.1.1) 14
    - Slosh Modeling and Analysis 16
    - Other Torque on Spacecraft (Branch 1.1.2) 17
  - *Low Momentum Falsely Reported as High (Branch 1.2 of the Fault Tree)* 18
  - *Processor and Software Errors (Branch 1.3 of the Fault Tree)* 20
    - Processor Hardware Errors (Branch 1.3.1) 20
    - Processor Software Errors (Branch 1.3.2) 20
    - Code Inspection Team 20

FINDINGS AND RECOMMENDATIONS 23

APPENDICES

- Appendix A: NEAR Anomaly Review Board Membership and Charter **(Not online)**
- Appendix B: Overview of NEAR Propulsion System **(Not online)**
- Appendix C: Overview of NEAR Guidance and Control System **(Not online)**
- Appendix D: Overview of NEAR Autonomy System **(Not online)**
- Appendix E: Overview of NEAR Simulation Environment **(Not online)**
- Appendix F: Upgrades to Simulation to Support NEAR Anomaly Studies **(Not online)**
- Appendix G: NEAR Mechanical Response to LVA Transient **(Not online)**
- Appendix H: Reconstructed RND1 Event Timeline **(Not online)**
- Appendix I: The RND1 Timeline Reconstruction Process **(Not online)**
- Appendix J: Reconstructed Autonomy Command History **(Not online)**
- Appendix K: Reconstructed CTP1 and CTP2 Command Histories **(Not online)**
- Appendix L: Summary of Brassboard Simulations **(Not online)**
- Appendix M: References
- Appendix N: List of Acronyms

EXECUTIVE SUMMARY

On 20 December 1998 the Near Earth Asteroid Rendezvous (NEAR) spacecraft began the first and largest of a series of rendezvous burns required for capture into orbit around the asteroid Eros. Almost immediately after the main engine ignited, the burn aborted, demoting the spacecraft into safe mode. Less than a minute later the spacecraft began an anomalous series of attitude motions, and communications were lost for the next 27 hours. Onboard autonomy eventually recovered and stabilized the spacecraft in its lowest safe mode (Sun-safe mode). However, in the process NEAR had performed 15 autonomous momentum dumps, fired its thrusters thousands of times, and consumed 29 kg of fuel (equivalent to about 96 m/s in lost delta-v capability). The reduced solar array output during periods of uncontrolled attitude ultimately led to a low-voltage shutdown in which the solid-state recorder was powered off and its data lost. After reacquisition, NEAR was commanded to a contingency plan and took images of Eros as the spacecraft flew past the asteroid on 23 December. The NEAR team quickly designed a make-up maneuver that was successfully executed on 3 January 1999. The make-up burn placed NEAR on a trajectory to rendezvous with Eros on 14 February 2000, 13 months later than originally planned. The remaining fuel is sufficient to carry out the original NEAR mission, but with little or no margin.

A NEAR Anomaly Review Board (NARB) was formed to determine the reason for the rendezvous burn events and to make recommendations for NEAR and for similar programs. The cause of the abort itself was determined within 2 days of the event: the main engine's normal start-up transient exceeded a lateral acceleration safety threshold that was set too low. Compounding this error was a missing command in the onboard burn-abort contingency command script; this script error started the attitude anomaly. Fault protection software onboard NEAR correctly identified the problem and took the designed, preprogrammed actions. While the fault protection actions did prevent complete battery discharge before the spacecraft recovered its proper Sun-facing orientation, they did not prevent, and they possibly even exacerbated, the protracted recovery sequence.

The Board's investigation included a painstaking reconstruction of the post-abort timeline from the small amounts of data that remained following solid-state recorder powerdown. More than 128 simulations were run on a NEAR simulator containing ground hardware replicas of all six flight processors running the actual flight code. Additional simulations were run on a software-only simulator. These simulations show that the fault protection actions should have ended the attitude anomaly quickly. Although the simulation fidelity was substantially improved and extended during the course of this investigation, it is clearly deficient in some respect, since we are unable to duplicate the entire sequence of events that occurred in flight. An independent review of the flight code was also conducted, and suspect hardware and circuit elements were reviewed. The investigation established a good understanding of the events during approximately the first 47 min after the abort, but no explanation for the failure of onboard autonomy to quickly correct the problem. The Board found no evidence that any hardware fault or single-event upset contributed to the failure. Although software errors were found that could prolong and exacerbate the recovery, they by no means fully explain it.

The Board is unable to establish a complete explanation for the rendezvous burn events. Nevertheless, we include in this report observations and recommendations that could prevent a recurrence on NEAR or on other programs. These recommendations focus on improving quality control and configuration management within Mission Operations, making better use of NEAR's simulation capability, and taking certain defensive measures on the spacecraft.

INTRODUCTION

The Near Earth Asteroid Rendezvous (NEAR) spacecraft, after traveling for nearly 3 years, initiated a planned main engine burn on 20 December 1998 to begin the rendezvous with target asteroid 433 Eros. After a short settling burn, the main engine ignited. A fraction of a second into the planned 15-min burn, it aborted. The abort placed the spacecraft in safe mode to await instructions from the ground, as it is designed to do. But what happened next was definitely not expected. Within a few seconds, all communication from the spacecraft was lost; NEAR remained silent for 27 hours. When communications were reestablished, it was found that during the blackout period the spacecraft had encountered an attitude anomaly, experienced a low-voltage condition, and lost 29 kg of fuel.

The main engine abort meant that NEAR would soon fly past Eros rather than rendezvous with it. Mission Operations personnel and science team planners were prepared with a contingency picture-taking sequence to obtain some benefit from the flyby. On 23 December, more than 1000 images were taken as NEAR flew to within 4,000 km of Eros. Concurrently, the NEAR team worked to understand the cause of the abort well enough to safely command another rendezvous burn. The 24-min make-up burn of the main engine on 3 January was successful. It increased NEAR's speed by 940 m/s to catch up to the faster-moving Eros, which had overtaken NEAR during the flyby. A small clean-up burn followed on 20 January. These two burns placed NEAR on course to rendezvous with Eros 13 months later than originally planned. The make-up burn left NEAR with about one-third the fuel it would have had if the original burn had been successful. The remaining fuel is sufficient to carry out the original NEAR mission, but there is no longer any margin.

Once the spacecraft was safed and the NEAR mission rescued, it became important to fully understand the cause of the anomaly. A NEAR Anomaly Review Board (NARB) was formed to determine the cause of the abort and, particularly, the cause of the subsequent events. The NARB membership and charter are presented in Appendix A. This report sets forth the Board's findings and recommendations.

What Was *Supposed* to Have Happened

For NEAR to catch up with and rendezvous with the faster-moving Eros requires that the spacecraft be sped up with a series of four engine burns. The first rendezvous burn, RND1, would have increased the spacecraft's velocity by 650 m/s using the main, bipropellant large velocity adjust (LVA) engine (see Appendix B; Ref. 1 and Appendices B–E provide tutorial information on NEAR). This 15-min LVA burn was planned to begin at 5 p.m. Eastern Standard Time (EST) on 20 December 1998, when NEAR was 238,000 km from Eros. The LVA burn is preceded by a 3-min settling burn from NEAR's 22-N monopropellant thrusters to force the liquid oxidizer against its tank outlets. During the LVA burn, the 22-N thrusters provide vernier control.

If the LVA burn had terminated normally, a command sequence (the "clean-up macro") would have been executed to shut the propellant tanks and make a graceful transition back to attitude control using reaction wheels. The latter action also automatically selects the 4.5-N thrusters (the "B" thruster group) if an autonomous momentum dump is required, instead of using the larger 22-N thrusters (in the "A" thruster group) that had been commanded for use during the LVA burn.

Following the LVA burn on 20 December, a second burn 8 days later was planned. It would have increased NEAR's velocity by 294 m/s, at a distance of 21,000 km from Eros. This would have reduced NEAR's speed relative to Eros to less than 30 m/s. On 3 January, a small third burn was planned to reduce relative speed a further 22 m/s, at 5,000 km. At 10 a.m. EST on 10 January, NEAR would have performed one last small burn to begin orbiting Eros at an altitude of 1,000 km.

Mission Operations Planning and Review

This general rendezvous sequence was part of the mission design from inception. In preparation for the rendezvous, NEAR Mission Operations had begun building up its staff from an average of 7 full-time equivalents (FTEs) during the 3-year cruise phase to 21 FTEs just before the rendezvous. Detailed design of the command scripts for the four planned rendezvous burns and for the contingency burns began in September 1998. The designs were based on the successful Deep Space Maneuver (DSM) performed in July 1997, the only previous firing of the LVA in flight. It should be noted that, between the DSM and RND1 burns, Mission Operations introduced a new scripting tool (SEQGEN) that required certain prior scripts from the DSM to be brought manually into the new system.

An outside review board, chaired by the Jet Propulsion Laboratory (JPL), conducted an independent review on 27 October 1998 of NEAR's readiness to rendezvous with Eros. That board presented its findings to the NEAR team and issued its report on 14 November 1998, declaring NEAR "well on its way to a successful rendezvous and orbit" (Ref. 2). The board did raise 10 key issues, however, with recommendations for each, to "increase the robustness of the already excellent preparation." APL took positive steps to address each of the concerns raised and submitted written responses on 14 December 1998 to each of the issues raised (Ref. 3).

Mission Operations, working with guidance and control (G&C) and propulsion engineers, began generating the RND1 command script in November 1998. The scripts were internally reviewed for the last time on 7 December during a detailed command review. This review included members of G&C and Mission Operations, but did not include the spacecraft system engineer. The review's action items were documented by e-mail. The RND1 commands were uploaded on 16 December. This upload also included a contingency burn (RND1A) that could be executed the next day if RND1 was missed for any reason that could be corrected within 24 hours. Several other burn scenarios were planned should the primary and backup dates for RND1 not be feasible for some reason. Multiple 34-m and 70-m antennas from NASA's Deep Space Network (DSN) were scheduled to cover the real-time execution of RND1 to assure that adequate telemetry and tracking data were available for assessing maneuver performance.

What Was Observed to Happen

The sequence of small settling burns executed nominally beginning at $T - 200$ s on 20 December ($T = 0$ is the start time of the LVA burn). This sequence added about 6 m/s to spacecraft velocity, as planned. However, when the LVA burn began at $T = 0$, it shut down within a fraction of a second. RF communications from the spacecraft were lost at $T + 37$ s and not regained for 27 hours. The record of downlink Doppler measurements from the spacecraft (Fig. 1) shows these events up to loss of signal.

At 8 p.m. EST on 21 December, 27 hours after loss of signal, the DSN verified positive lock on a downlink signal from NEAR. The spacecraft was in the lowest of its safe modes, Sun-safe-rotate, in which the solar array normal is pointed at the Sun and the spacecraft rotates around the spacecraft–Sun line once every 3 hours (see Appendix D). On the second rotation, a command was sent to stop the rotation (Sun-safe-freeze mode). As data began to arrive at the Mission Operations Center, several things became clear:

- The LVA burn had aborted.
- A low-voltage trip had occurred, and the bus voltage (nominally 33 V) had gotten as low as 24.3 V.
- Detailed data stored in the solid-state recorders had been permanently lost when they were switched off during the low-voltage trip.

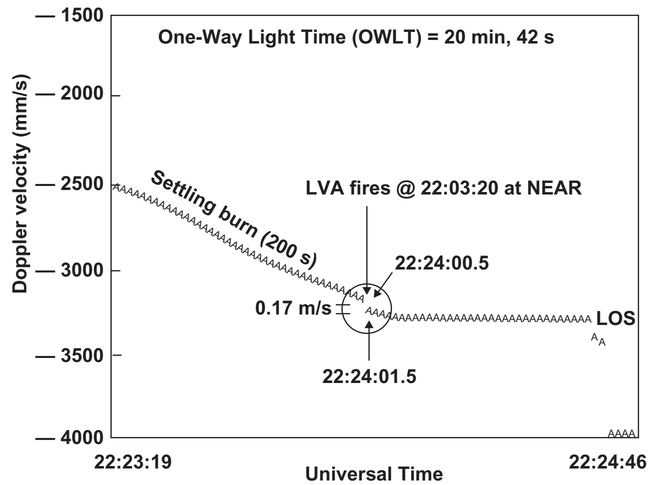

Figure 1. Spacecraft radial velocity, as measured by the downlink signal Doppler shift, shows the settling burn, the moment of LVA burn abort, and the subsequent loss of signal 37 s later. (Figure adapted from Ref. 4.)

- A significant amount of fuel (28.5 ± 1.5 kg) had been lost.
- The active Attitude Interface Unit (AIU) was now AIU2, rather than AIU1, which had been active at RND1 start.
- All three gyros had entered the Whole Angle Mode (WAM), suggesting that high body rates had been experienced in all three axes.

Within a day or two of recovery, the precipitating causes were understood well enough to permit a new burn, which took place successfully 3 January 1999. Data from thruster firings during this burn showed that no thruster had been damaged by cold firings that might have occurred during the RND1 event. Given the amount of propellant that was used, it can be theorized that there were few cold starts on any thruster because the time between firings was less than the thruster cool-down time of at least 2 hours. Aside from the fuel loss, the only permanent damage suffered by the spacecraft appears to be contamination of the Multispectral Imager's optics by residues from propellant lost during the anomaly. This contamination degrades the response at the shortest wavelengths. In-flight calibrations and ground processing can partially correct these effects, minimizing the impact on mission science.

The fuel lost during RND1 is equivalent to 96 m/s delta-v. DSN tracking results since RND1 indicate that a net delta-v of 12 m/s was imparted to the spacecraft after the settling burn and before the recovery. This small effect for a large expenditure of fuel is consistent with a tumbling spacecraft.

DISCUSSION

Formal NARB activities began at the first meeting on 21–22 January 1999. At this meeting, principals from the NEAR team provided tutorial briefings and answered the Board's questions. The NARB also conducted private interviews with many of the principals. To provide technical support to the NARB, an internal NEAR Anomaly Group (NAG) was established, headed by Dr. T. Strikwerda, supervisor of the Space Department's Mission Concept and Analysis Group. The NAG included a G&C analyst, the NEAR system engineer, two members of the Mission Operations team, and the NEAR performance assurance engineer. In late February a code inspection team was established to support NAG and NARB activities. This team, headed by M. White, supervisor of the Space Department's Flight Software section, was chartered to inspect the AIU and flight computer (FC) codes for possible software errors that could explain the observed events. The code inspection team consisted of three software experts and two additional analysts. This team also provided software support for the brassboard upgrades needed to support NAG simulations. The upgrades are described in Appendix F; the simulations and results are summarized in Appendix L.

Reconstructed Timeline

The initial NAG activities focused on trying to establish the exact timeline of events from the limited data available. Because all of the data stored in NEAR's solid-state recorders were lost during the low bus voltage event, the timeline had to be reconstructed by deduction and inference from the limited data stored in processor memories (Appendix I describes the process). These data consisted mostly of partial records of commands and autonomy rules executed and of min/max values (and their times of occurrence) for selected telemetry points. Timeline reconstruction took approximately 6 weeks. The detailed reconstructed master timeline is given in Appendix H. An abridged version is shown graphically in Figs. 2 (early events) and 3. The description of events given below is based on these data and on brassboard simulation results.

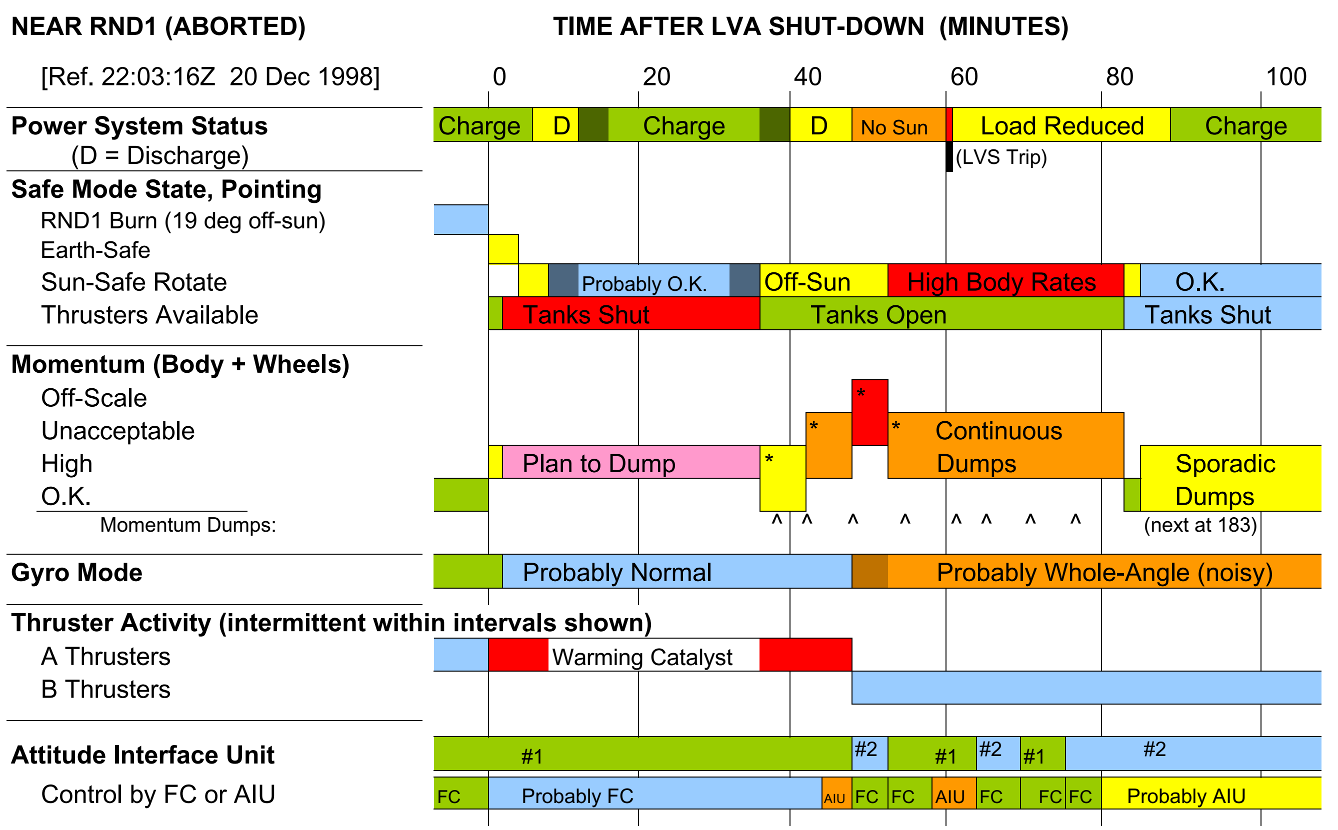

Figure 2. Early events in the RND1 timeline, measured from time of abort. Uncertainties in the timeline result from having to deduce and infer events from the small amounts of stored data remaining after the solid-state recorder was powered off during LVS.

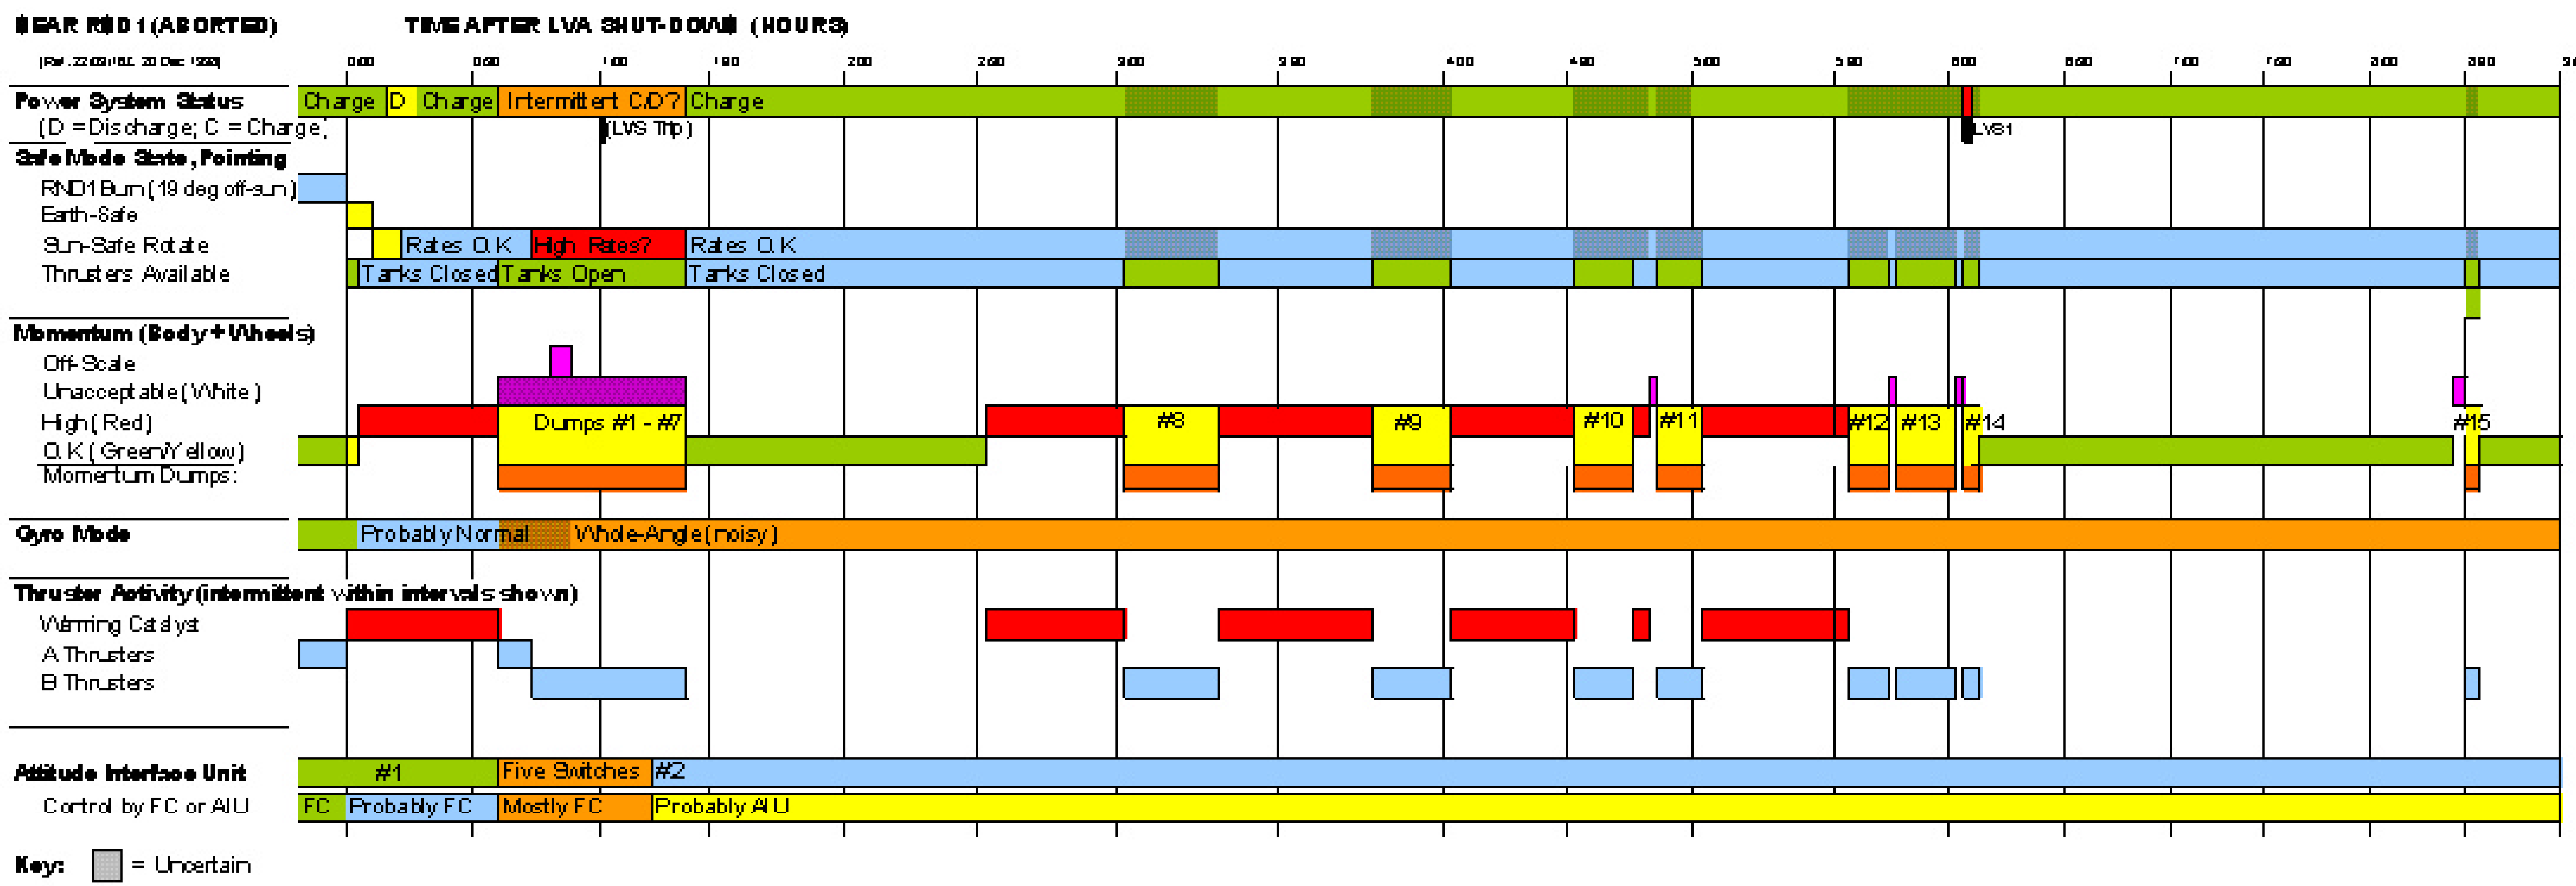

Figure 3. The complete RND1 timeline, derived from Table H-1 in Appendix H. Uncertainties in the timeline result from having to deduce and infer events from small amounts of stored data remaining after the solid-state recorder was powered off during LVS.

Overview of Events

The series of anomalous events began at LVA ignition, when the accelerometers sensed a lateral acceleration probably in excess of 0.15 m/s^2, which exceeded the threshold of 0.10 m/s^2 that had been set to detect uncontrolled lateral thrust during any burn. (Groups of four consecutive 100 sample/s accelerometer readings are averaged, then compared to a threshold.) This acceleration represented a reasonable mechanical response to the normally abrupt LVA start-up, and it can be concluded that the threshold was set too tightly. An LVA start-up transient had been observed on the only previous use of the LVA, the DSM burn of July 1997 (Fig. 4); however, the significance of this transient was not appreciated. In addition, NEAR's structural design, with a separate propulsion module cantilevered from the base of the spacecraft, heightens the mechanical coupling of this transient to the accelerometers (Fig. 5 and Appendix G); this bending response was not appreciated. In retrospect, the correct thing for the G&C software to have done would have been to ignore (blank out) the accelerometer readings during the brief transient period, or at the very least to increase the filtering and/or raise the threshold during that time. It is significant that NEAR is the first of JHU/APL's 58 spacecraft to use bipropellant propulsion, and RND1 was only the second time the LVA was fired in orbit.

The abort triggered execution of preprogrammed commands appropriate for a burn abort anomaly, including commands to place the spacecraft in Earth-safe mode (see Appendix D). The command system executed this 34-s script as expected. Meanwhile, the G&C software began a 19° slew from the attitude necessary for LVA burn to the Earth-safe attitude. In Earth-safe mode, the solar panels are aimed toward the Sun and the spacecraft is rolled around the Sun line to place the medium-gain (fan-beam) antenna toward Earth to await further instructions. The slew was initiated by thrusters rather than by reaction wheels, because the G&C system had

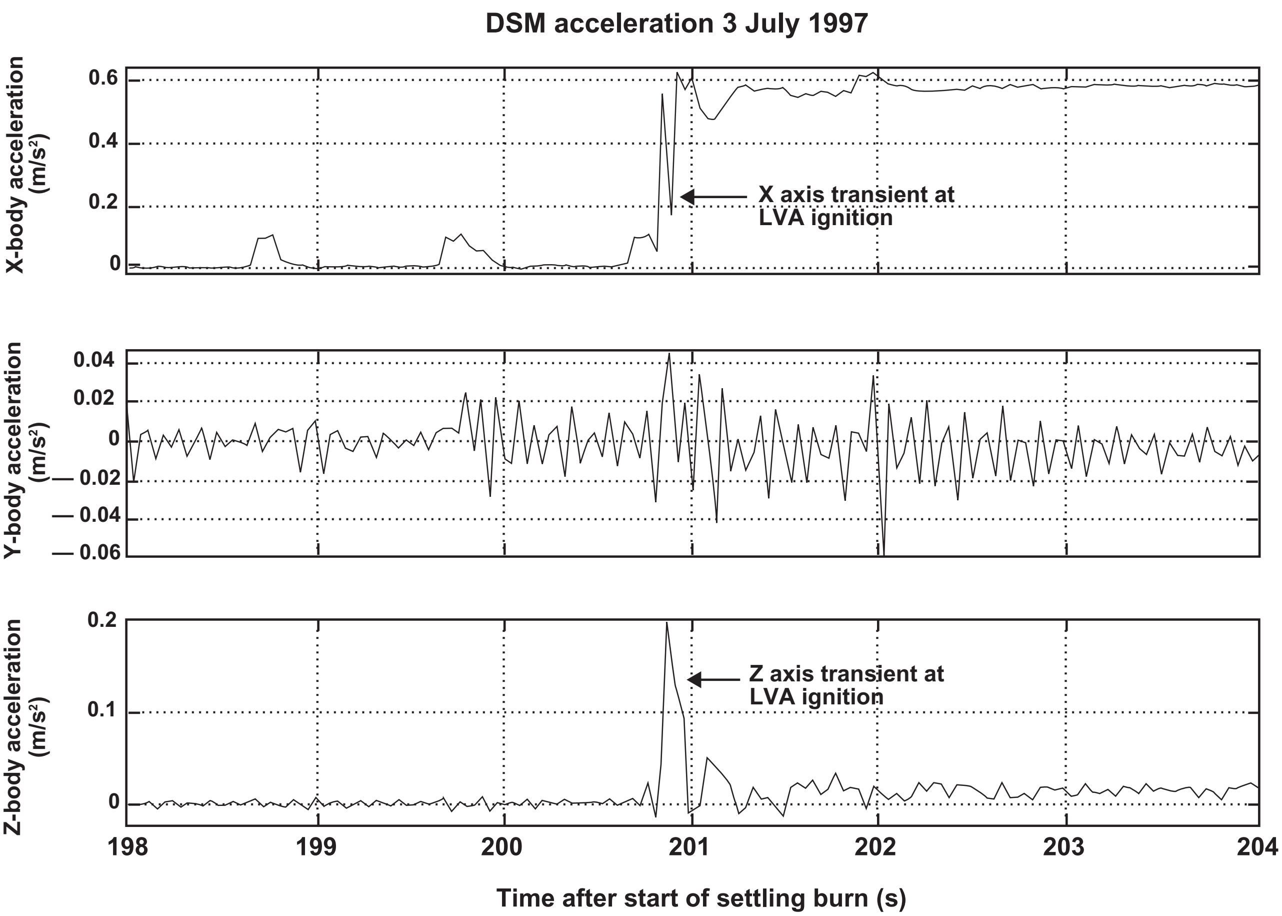

Figure 4. An acceleration transient similar to that which initiated RND1 occurred during the only previous use of the LVA engine, the Deep Space Maneuver burn in July 1997. The significance of the transient was not appreciated at the time.

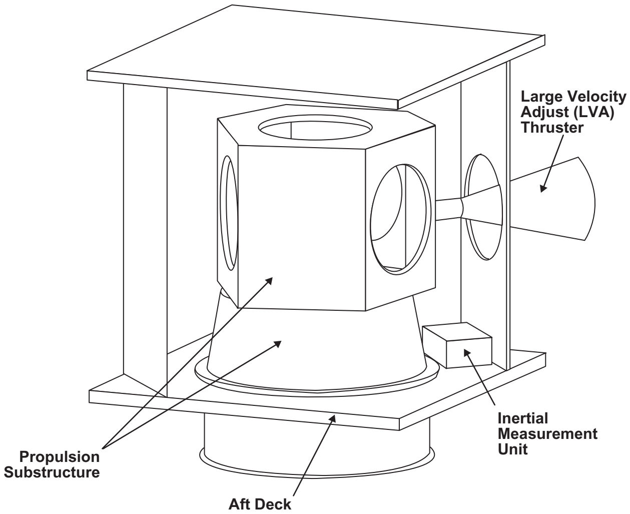

Figure 5. NEAR's structural design uses a separate propulsion module cantilevered from the aft deck of the main spacecraft structure. The resulting flexure during LVA ignition couples into the accelerometers located in the Intertial Measurement Unit. (Propulsion tanks and two side panels are removed for clarity.)

been configured to use the 22-N thrusters for attitude control during the burn. This instantaneous response by the G&C system to use *thrusters* for attitude control is perhaps a design weakness, since using thrusters for attitude control is generally to be avoided except during an actual delta-v maneuver. Simulations show that the thrusters imparted significant momentum to the system and established a body rate of about 1°/s toward the Sun.

Due to insufficient review and testing of the clean-up script, the commands needed to make a graceful transition to attitude control using reaction wheels were missing. The G&C system remained commanded to use the 22-N thruster group for attitude control. The command script did, however, close the fuel tanks and disable power to the thrusters (at T + 24 s). Without thruster power, the G&C system uses the reaction wheels by default. With the large momentum in the system, the reaction wheels could not stop the attempted slew quickly enough, and the spacecraft overshot the Sun. When the solar panels were not pointed at the Sun within 300 s of the initiation of Earth-safe mode, an autonomy rule demoted the spacecraft to Sun-safe mode.

During these Sun acquisition activities, the G&C momentum management software detected the large system momentum and initiated a 30-min warm-up of the thruster catalyst beds in preparation for an autonomous momentum dump (a so-called *Red* dump; see Appendix C). About 7 min into the warm-up period, a data structure error in the use of wheel speed data for the momentum calculation caused the warm-up timer to be reset. The error involved misreporting a wheel at its maximum speed as being at *zero* speed, causing the system to falsely compute a momentum below the safe (*Green*) level. The data structure error explains the "mystery" of the first dump occurring 37 min, rather than 30 min, after the first report of Red momentum. On the basis of brassboard simulation results and the general lack of any contrary indications, we believe that during the final 20–30 min of this 37-min warm-up period, the spacecraft attitude and rate probably stabilized into the desired Sun-safe mode.

When the catalyst bed warm-up period expired, the thrusters were enabled and the first momentum dump began (T + 00:37:52). The dump terminated after 285 s, just 15 s short of the time allowed. At the conclusion of the dump, the G&C system signaled the command processor to remove power from the thrusters. A brief period occurred during this handshake when G&C was not dumping momentum, but the thrusters were still powered because the command system had not yet responded to the G&C request for dump termination. This period was lengthened by a design error (bug) in the FC code, discovered during brassboard simulations 22–24. This bug keeps the thruster request high for several seconds too long and can also cause unnecessary AIU switchovers due to time-out. Had G&C been correctly configured to use reaction wheels to control attitude, leaving thrusters powered for this brief period would have had no ill effect. But the RND1 command script error left the G&C system commanded to use the 22-N thrusters to control attitude. During the period of the momentum dump, the spacecraft had drifted off Sun-pointing. With thrusters powered for a few seconds, the G&C system immediately attempted to regain Sun-pointing by firing thrusters. That imparted a "kick" of momentum back into the system, this time at a level high enough to trigger a momentum dump without catalyst bed warm-up (a so-called *White* dump). At T + 00:42:53, the G&C system began this second momentum dump, but failed to complete it in the allowed 300 s.* An autonomy rule booted the backup AIU and gave it control. The AIU switch-over finally reestablished the default operating conditions within G&C: attitude control using reaction wheels and autonomous momentum dumps using the smaller 4.5-N thruster set normally used for attitude control when the LVA is not actually firing.

*The G&C system has an internal timer limiting thruster use requests to 270 s. It was thought that this was long enough to dump momentum under nearly all conditions but still short of the Command and telemetry processor's (CTP's) 300-s autonomy limit. However, repeated simulations have shown that there is significant "overhead" before G&C starts dumping, making the 270 s marginally too long; that value has since been decreased on the spacecraft.

Given the range of unknown initial conditions shortly after $T = 0$, simulations have reproduced the observed spacecraft behavior reasonably well for the period up through the first two dumps (about T + 00:47:00), but not the behavior after that. From that point on, all simulations of RND1 that do not introduce some additional fault mode (e.g., a stuck-on thruster) result in the excess momentum being dumped and nominal attitude control being regained relatively quickly. In flight, however, the G&C system continued attempting to dump momentum. Autonomy interrupted this activity at 300-s intervals to switch back and forth between alternate AIUs in an attempt to end the excessive thruster use. After five AIU switches, the self-limiting autonomy rules left backup AIU2 in control and stopped attempting to limit the momentum dump duration. Analysis of the autonomy data indicates there were eight more momentum dumps over the next 7 hours. With autonomy rules no longer limiting their duration, many of these were very long (the longest over 1000 s). It was during this period that most of the fuel was probably consumed.

At T + 00:59:24, a Low Voltage Sense (LVS) trip occurred, indicating that attitude was not controlled to Sun-pointing for a considerable period. The exact period of time required for a fully charged battery to reach LVS* depends on many variables—such as heater duty cycles, wheel usage, and the exact battery temperature—that are unknown for RND1. To obtain a lower bound on the time, a worst-case analysis was performed. Using a low battery energy capacity and Sun-safe spacecraft loads iterated to match telemetry records (245 W), analysis suggests that at least 9 min of discharging (Sun angle greater than about 50°–60°) is required. This is well within the time (22 min) from the start of the first dump to the LVS, and coincidentally, is identical to the time between when the solar panels turned edge-on to the Sun and the LVS. It should also be noted that the "misreported wheel speed" error already described, which was responsible for extending the 30-min warm-up to 37 min, could also have caused more time to be spent at high Sun angles during the early part of the warm-up period. An example of this effect is seen in simulation 63 (Appendix L). In this simulation the misreported wheel speed error causes the Sun angle to exceed 50° for almost 7 min during the early part of the warm-up period. Since the battery recharges at only about 1/20th the rate of Sun-safe mode discharge, this simulation illustrates how the misreported wheel speed error can shorten the time to LVS.

* LVS occurs at 26-V bus voltage, corresponding to about 27.1 V at the battery.

The data also indicate that body rates increased dramatically, with system momentum going off the scale (>25.5 Nms, corresponding to about 4°/s) at T + 00:50:20. Another reason to believe the spacecraft experienced high body rates is that the gyros—all three axes independently—were found in Whole Angle Mode (WAM) when the spacecraft was recovered. This mode is entered when the gyro's normal strap-down mode (a locked force-to-rebalance loop) cannot be maintained; this occurs when body rates exceed about 11°/s (see Appendix C). This mode was added to NEAR's gyros specifically to deal with tip-off following launch vehicle separation. Analysis and tests have not established any other credible means of entry to WAM. No data allow us to pinpoint the time of entry into WAM. One data point does show that the gyros were not in WAM at T + 00:00:11, so we know that the abort did not directly cause the entry into WAM. The time of entry into WAM is of interest because our simulations show that once the gyros enter this coarse and noisy mode, momentum dump performance becomes extremely inefficient.

About T + 01:21:00 the spacecraft entered a quiet period (Fig. 6). The next momentum dump would not occur for another 1.5 hours. The only data point recorded during this period shows that the bus voltage recovered to 28 V at about T + 01:25:00, indicating that the arrays were pointing generally toward the Sun and charging the battery. With Sun-safe loads, charging would begin almost instantaneously when the panels were within 40° of the Sun, but we have no indication of how long this attitude was sustained. The quiet period ended at about T + 03:02:00, when the G&C system began another series of seven momentum dumps.

Over the next 3 hours there was a higher level of computed momentum, resulting in four Red dumps (those preceded by catalyst bed warm-up), three White (immediate) dumps, and one dump that started as Red but switched to White after about 7 min. As described above, with autonomy no longer intervening, some of these dumps were very long. The long periods of high calculated momentum probably resulted from the high WAM noise that was feeding the momentum estimate.

At T + 06:09:46, the bus voltage recorded its minimum value (24.3 V, less than the LVS trigger voltage), indicating continued drain on the battery during this period. There were no autonomy rules that would be triggered by this condition, so no other associated data were captured. However, we can conclude that the attitude was not well-controlled to Sun-pointing.

At about this same time, T + 06:10:00, the spacecraft entered another quiet period lasting more than 2 hours. We have no data points from this quiet period, but from the fact that the bus voltage stopped dropping we know

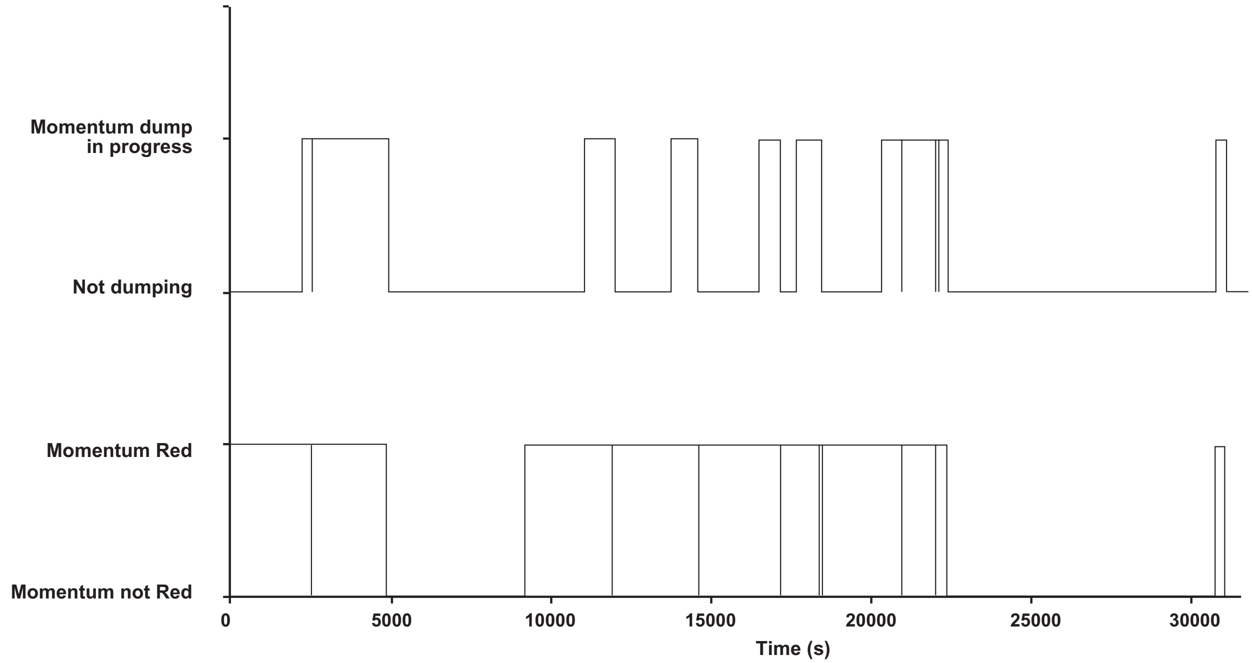

Figure 6. Extracted from Fig. 3 and Table H-1 in Appendix H, this figure illustrates that momentum dumps occurred in three distinct groupings, separated by quiescent periods up to 2 hours long. When momentum goes Red, a 30-min warm-up of the thruster catalyst bed heaters begins in anticipation of an autonomous momentum dump.

that the solar panels were mainly Sun-pointing. At about *T* + 08:31:00, G&C carried out one final White momentum dump of about 4.5-min duration. No further unexpected activity occurred.

In all, 15 autonomous dumps had occurred, an estimated 7910 s of dump activity. Estimating fuel use by counting "thruster seconds" shows that this thruster activity is about right to explain the 29 kg fuel loss. The details of these later dumps and the intervening quiet periods remain unexplained. At recovery, the spacecraft was in a stable Sun-safe mode attitude and rotation with a fully charged battery.

The LVS trip at *T* + 00:59:24 turned off the solid-state recorders, clearing their volatile memory of the detailed telemetry records. From the limited data available to reconstruct events, we have certainty about most of the spacecraft mode transitions and their timing, but no detailed information about thruster firings, G&C sensor data, spacecraft attitude and rates, or which of the two G&C computers (AIU or FC) was in control. Nominally, the G&C system operates with the FC generating the actuator commands and the AIU acting only as an interface unit and a checker. When an AIU is booted it immediately tries to go to this nominal "FC override" state. The G&C system has successfully entered the FC override state after AIU boot in all of our brassboard simulations. Furthermore, those simulations that most resemble the reconstructed timeline (30+ min Red warm-up followed by two dumps) show that the FC usually maintains control during the first and second dumps. Simulations with only one dump before the AIU switch usually have been in AIU control earlier, due to Sun keep-in (SKI) violation; the exact reason for this difference is unclear. At recovery, the AIU reported that it had taken control from the FC because of an SKI violation (an SKI violation occurs if solar panels are not pointed within 8° of the Sun for 300 s). Since the time-out for SKI is the same as the time-out for excessive thruster use that forced the AIU switches, we can be reasonably certain that at least the five momentum dumps following AIU switches were performed under FC control. The low voltage condition recorded in the data (indicating considerable periods spent significantly off Sun-pointing) suggest that the SKI logic probably triggered the AIU to take over control soon after the last switch. This is not certain, however, since the solar panels only had to pass within 8° of the Sun briefly to reset the 300-s SKI timer. All brassboard simulations show periods of control by both the FC and AIU, with the exact timing of the periods dependent on initial conditions.

Based on extensive simulations (a typical one is described below), high WAM noise in the gyro comes closest to replicating conditions during the periods of continuous momentum dumping. The noise levels required are 7–10 times higher than levels measured on the actual gyros with the spacecraft quiescent. These simulations show poor performance by the G&C system with high WAM noise, so much so that the spacecraft "wallows" significantly, often pointing many tens of degrees from the Sun (highest reported Sun angles by each of the AIUs during RND1 were 129° and 169°). However, in our simulations the system *never* settles down unless the WAM noise is subsequently reduced; yet we know that NEAR did settle down and that the noise at recovery was quite low.

Simulation Results

An extensive series of simulations was run to provide insight into the spacecraft behavior after the abort. Most of the simulations were conducted on the NEAR brassboard simulator. The simulation process required a substantial effort to repair accumulated brassboard hardware failures, improve fidelity, and add enhancements needed to test candidate causes for the anomaly. A total of 128 simulations were run over a 7-month period. A number of these were designed to explore various postulated branches on the fault tree (see fault tree discussion in the following section and Appendix L). A pattern soon emerged that fit the RND1 behavior up to the first AIU switch. Events after the first switch varied considerably depending on the failure mode tested. We have chosen one of the most recent simulations (number 126) to illustrate the "classic" behavior that seems to most nearly replicate the behavior inferred from the RND1 timeline reconstruction.

In simulation run 126 (shown in Figs. 7, 8, and 9), the event begins with the spacecraft pointed 19° from the Sun during the 200-s settling burn (point *a* of Fig. 7). The system momentum is highly variable due to the thruster firings (point *b*). At abort the spacecraft demotes from operational to Earth-safe mode (*c*). The G&C actuators selection was set to thrusters only for the LVA burn, and the thrusters now slew the spacecraft toward the Sun (*d*). Meanwhile, the command and telemetry processor (CTP) begins executing the burn-abort macro, which removes power from the thrusters (*e*). Without powered thrusters, the G&C system ignores

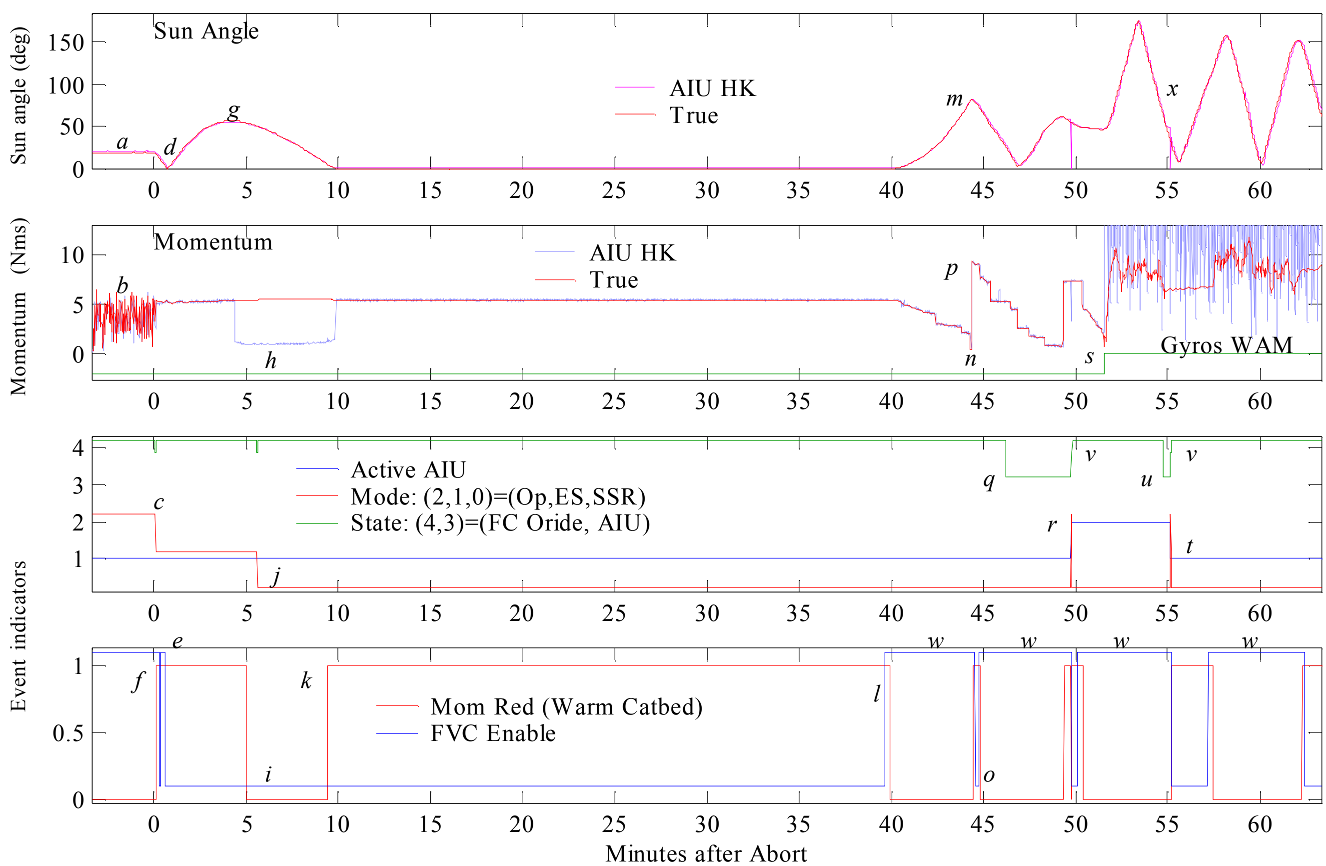

Figure 7. Simulation 126 illustrates the "classic" early timeline behavior. Shown are true and reported Sun angle and total momentum, as well as indicator flags showing which processor is in control and its state, as well as spacecraft mode. The bottom chart shows the momentum Red flag and the state of the thruster enable line.
Click here for a high-speed animation of the simulated spacecraft motions (first 90 min). Runtime 35 s, 2.2 MB.

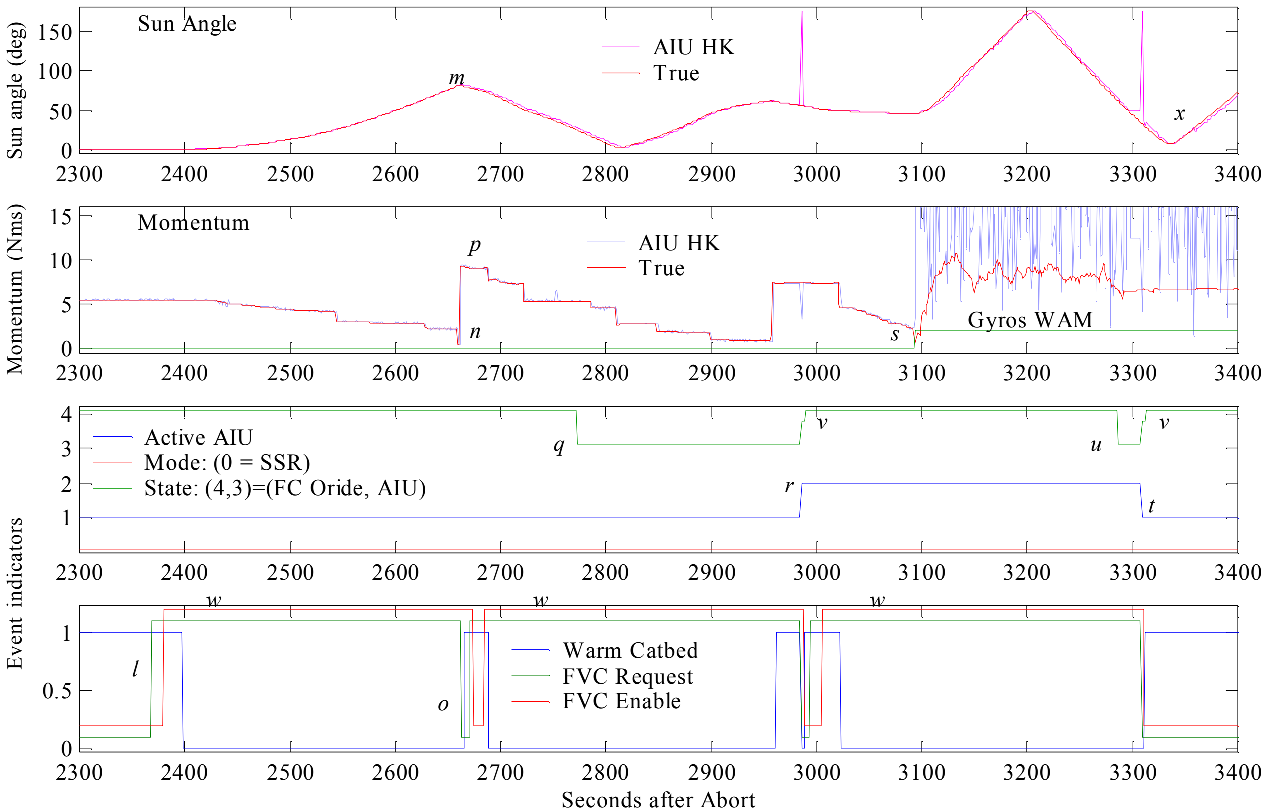

Figure 8. Expanded view of Fig. 7 at the time of the first momentum dumps. In this simulation, the gyro is manually forced into WAM at about 3100 s (with artifically high noise level).

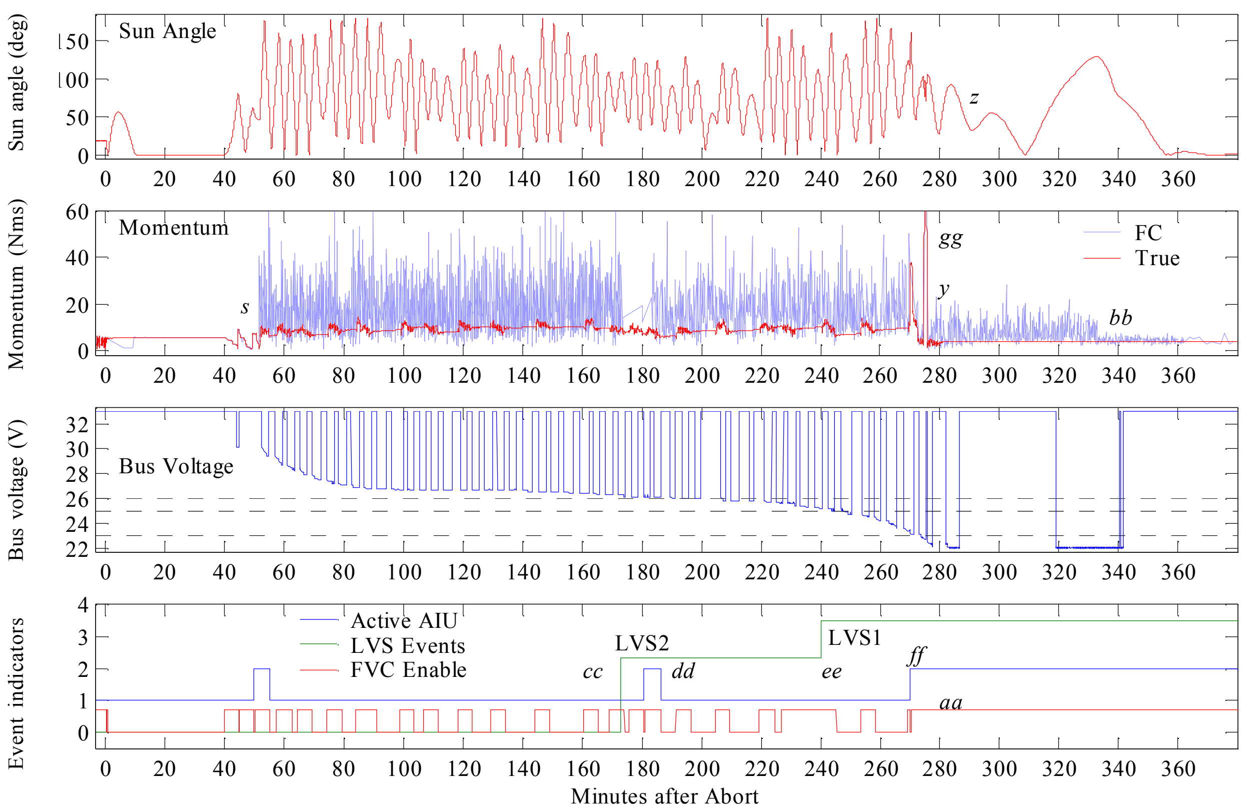

Figure 9. Illustration of the poorly controlled attitude and momentum when the gyros are in WAM. The third panel shows the power system status as simulated by a model within the brassboard (LVS trip point is 26 V). The bottom panel shows G&C processor and thruster status.

the commanded thrusters-only control mode and begins to control attitude with the reaction wheels. Since the G&C system is no longer executing a delta-v maneuver, it declares that momentum is too high (approximately 5 Nms in this case) and signals the command processor to warm the catalyst bed heaters (30-min warm-up) in preparation for a momentum dump (*f*). The spacecraft continues slewing toward the Sun, but with high system momentum the reaction wheels cannot slow it down quickly enough, and the spacecraft overshoots the Sun, reaching a Sun angle of about 56° (*g*). As the wheels spin up and one or more exceeds a software limit that had been erroneously set too low, they are declared "invalid" and removed from the momentum calculation (*h*) until they return within range. The momentum calculated without the "invalid" wheel speeds briefly dips below the Green limit, which clears the momentum dump request (*i*). At about the same time the AIU demotes the spacecraft to Sun-safe-rotate (*j*) because of a safe mode SKI violation, i.e., the spacecraft has pointed more than 8° away from the Sun for more than 300 s. As the spacecraft continues to rotate, the wheel speeds change and all four are again used in the momentum calculation, restarting the 30-min catalyst bed warm-up timer (*k*). This behavior explains the mystery of why the first RND1 momentum dump occurred at 37 min rather than the expected 30 min. This misreported wheel speed syndrome has been observed in several of the later simulations, and a special series of stand-alone, software-only simulations was run to better understand it.

The spacecraft gradually recovers a stable attitude and rotation rate during the 30-min catalyst bed warm-up period. The momentum dump begins at about 40 min with the *FVC_request* (fine velocity control thrusters requested) and *FVC_enable* (FVC thrusters powered) flags raised (*l*) and the warm-up flag lowered. During the dump, the attitude is not controlled well by the wheels as momentum is transferred from the wheels to the body; body rate is reduced by thrusters. Since the momentum is high and the dumping is inefficient, it takes some time to transfer momentum from the wheels to the body. The spacecraft drifts off the Sun (*m*) during this time. (Fig. 8 provides an expanded timescale view of events *m*–*v*.) The dump terminates after a 270-s time-out expires (*n*) with the momentum in the *Yellow* zone—a so-called *Yellow-Good-Enough* (YGE) dump.

When a dump terminates, the G&C system lowers the FVC_request flag to the command processor, which responds by unpowering the thrusters and lowering the FVC_enable flag (*o*). This process requires several seconds. Since the burn abort command script used for RND1 lacked the command to restore the reaction wheels as the attitude control actuators, the G&C control laws switch from dump mode to attitude-control mode using thrusters before the command processor can respond. Since the spacecraft is pointing far from the Sun, G&C fires thrusters again to restore Sun-pointing, imparting high momentum once again (*p*), and a second momentum dump begins, this time an immediate (White) dump because the momentum exceeds 6 Nms. During this second dump, the AIU takes control (*q*) because the FC has not cleared the SKI violation within the allowed 300 s. This dump also completes as YGE, but under AIU control, which apparently takes slightly longer. This time a command system autonomy rule detects that the dump request has exceeded 300 s. The autonomy rule powers AIU2 and gives it control (*r*). The boot-up of AIU2 causes a reset of the G&C parameters to nominal; most important, the control actuator setting is defaulted to reaction wheels. With the thrusters no longer used by G&C for attitude control (only for momentum dumping), no more momentum-imparting "kicks" occur.

In the actual RND1 abort on the spacecraft there were also two dumps prior to the first AIU switch. In other simulations with this "classic" behavior, there is only one dump under AIU1 control before the switch to AIU2. The difference cannot be fully understood from the brassboard simulations but may be related to a "race condition" between the AIU and the command processor autonomy rule, since the YGE dump plus overhead appears to be about 300 s, the same as the 300-s time-out of the autonomy rule for AIU switch. Since the purpose of a YGE dump is to terminate a momentum dump with an acceptably low momentum before the autonomy rules switch AIUs, this indicates that the YGE time-out (270 s) is too long.

The timing of the first two momentum dumps and the first AIU switch in the simulation correspond quite well with what is known about the behavior of the spacecraft following RND1. Some aspects of RND1, however, never occur in the simulation. First, the spacecraft rates never approach the value needed to cause the gyros to enter WAM. This implies either that the simulation is not replicating the actual RND1 behavior, or that some rate-independent means of inducing WAM exists. None has been found in this investigation (see "Low Momentum Falsely Reported as High [Branch 1.2]" in the section on the fault tree later in this report). Second, unless we induce some additional failure mode, the spacecraft always recovers in much less time than that observed in RND1. Without such intervention, the simulations all result in the momentum being reduced to acceptable levels and nominal, Sun-safe attitude and rates being achieved within a relatively short time after the switch to AIU2. Finally, this particular simulation does not encounter the extended time at poor Sun angle early in the warm-up period due to the misreported wheel speed error (simulation 63 in Appendix L shows an example of this effect).

When a failure mode is introduced after the first AIU switch, the simulation results vary widely (Appendix L summarizes all 128 simulations). After exploring a number of possible causes, we have found that only a highly variable and erroneous momentum calculation results in the repeated momentum dumps that characterized RND1. One possible cause of this type of error is noise in the rate measurements obtained from the gyros. Because the gyros were in WAM on recovery, and because the noise level in WAM is considerably higher than in normal mode, we extensively explored the effect of WAM on spacecraft behavior. We found that at noise levels measured in flight on the day NEAR was recovered (0.0014° to 0.0030° rms), the spacecraft was well controlled, and no repeated dumps occurred. At the Litton test values (0.0047° to 0.0135° rms [Ref. 5]), there are no repeated momentum dumps but the attitude does not fully stabilize. By trial and error, we found that a WAM noise level of 0.02° rms is about the minimum level that causes repeated momentum dumping.

In the simulation discussed here and shown in Figs. 7–9, the gyros are forced to WAM soon after the switch to AIU2 (point *s* on the figures), replicating the approximate time (relative to AIU2 turn-on) when the gyros may have entered this mode (this is approximately the time at which the momentum was reported to be at 25.5 Nms, a saturated value in telemetry). With the gyros in WAM, all subsequent dumps are very inefficient due to the noisy gyro outputs; that is, they waste time and fuel reducing momentum. At start-up an AIU attempts to promote the system to FC-override mode if it can (successfully in every case simulated to date) (*v*). In this particular simulation, the third dump, under FC control, also exceeds 300-s duration, causing the command system autonomy rule to switch back to AIU1 (*t*). Since this dump also fails to clear the SKI violation within the allotted 300 s, the AIU again takes control from the FC (*u*), just before the AIU switch. At this high setting for the WAM noise amplitude, the momentum estimate is very high most of the time. There are frequent reports of high momentum which trigger immediate (White) momentum dumps (*w*). A total of four dumps are visible in Fig. 7, but there are others beyond the range of the figure. Typically there are almost continuous dumps until the gyro noise is reduced (Fig. 9, point *y*). Note that the Sun angle is quasi-periodic during the high WAM noise period (*x*). This is typical of the simulations with high WAM noise and persists until the noise is substantially reduced. This behavior results from a subtle interaction of the unbalanced placement of the thrusters and the control system.

Figure 9 shows the entire time span of this simulation. WAM was initiated under the control of the brassboard operator, since the simulations do not produce rates high enough to trigger WAM. Because momentum dumping never stops when the noise level is high (and after RND1, the repeated momentum dumps *did* stop), WAM noise levels are manually stepped down to three distinct levels. The difference between true and FC-estimated angular momentum in Fig. 9 is a measure of the level of WAM noise. At initial WAM entry (*s*) the WAM noise was set artificially high (0.02° rms), a value chosen by trial and error as about the minimum level that causes repeated momentum dumping. This value is about 7–10 times higher than levels actually observed on NEAR's gyro. At 275 min (point *y*) the noise level is reduced to the values measured during gyro acceptance testing (Ref. 5). The tumbling rate decreases significantly (*z*) and repeated momentum dumping stops (*aa*), but the attitude does not fully stabilize. At 330 min (*bb*), the WAM noise is further reduced to the values estimated from actual flight data when NEAR was recovered. With this final reduction in noise, the spacecraft at last achieves stable Sun-safe attitude and rate. One postulated mechanism for this variation in noise amplitude is that WAM noise depends on rate or voltage. If true, the dependence must be a strong one, since the simulated rates (about 1°/s) are small. Further, in flight, NEAR had two long quiet periods (1.5 and 2.3 hours) in between repeated momentum dumps. Although this could be simulated by stepping the WAM noise *up*, as well as down, we can think of no credible mechanism that would cause this to happen in flight.

Also shown in Fig. 9 is the spacecraft bus voltage as modeled in the Bench Test Equipment (BTE) portion of the brassboard (see Appendix E). The bus is at 33 V whenever there is sufficient array current, generally when the Sun angle is less than 90°. But as the spacecraft tumbles, the arrays point away from the Sun and the bus drops to the battery voltage, which continues to decline since the net array current is insufficient to charge the battery. At about 170 min the bus drops below 26 V, triggering the LVS event (LVS2, point *cc* of Fig. 9). LVS sheds a number of power system loads and also reenables the autonomy rules that cause AIU switching for incomplete momentum dumps. Two more AIU switches are seen shortly thereafter (*dd*). Many momentum dumps occur after these, but they do not trigger AIU switches because of the self-limiting autonomy (see Appendix D). At about 240 min, the bus drops to 25 V (*ee*), triggering the LVS1 event (*ee*). This event does not reenable AIU switching.

The fifth AIU switch (*ff*) in this simulation is due to an autonomy rule triggered by the bus voltage going below 23 V for more than 40 s. We know for certain that this rule, which initiates numerous drastic actions including power cycling of the gyros and turning off of the FC, did *not* fire in flight. The large momentum spike (*gg*) after the autonomy rule fires is because a thruster firing command is erroneously sent to the brassboard truth model repeatedly during the autonomy rule activation, causing the (simulated) thruster to be full ON for over 30 s. As soon as the autonomy rule completes and AIU2 is running normally, the excess momentum is dumped correctly. This thruster event is interesting—even alarming!—but it did *not* happen in flight because the hardware thruster drivers turn OFF in the absence of a valid AIU. This does not happen on the brassboard, which erroneously responds to stale firing commands. In the actual RND1 event, the fifth and final switch was due to the two command processors being imperfectly synchronized, so that the last AIU switch occurred twice (once from each command processor).

In flight after RND1, the LVS event occurred sooner and the five AIU switches occurred much closer together than in this simulation. From this infomation, we suspect that the brassboard power model has lower power drain than the real spacecraft, or that the real Sun angle history included less time at low Sun angle, or both.

A Walk Through the Fault Tree

In the course of the anomaly investigation numerous "mysteries" presented themselves. At one time the "Mysteries List" had eight distinct questions. For example, why did the catalyst bed warm-up for the first dump take 37 min instead of the design value of 30 min? This mystery was finally explained by the misreported wheel speed error described earlier. Why was 23 min of FC Sun vector usage reported (suggesting that the G&C system lost Sun sensor output for that long)? One hypothesis, which there was no safe way to test, was that thruster plumes from a tumbling spacecraft might confuse the digital solar aspect detectors (DSADs) when the Sun is far from array normal. Another speculation was that a flap of insulation may have covered a DSAD. Ultimately, this mystery was explained by the effects of DSAD "keyholing" (the DSAD fields of view are ±64° square cones with "keyholed" regions of invisibility up to almost 7° above the array plane). One by one the questions were answered, leaving the principal unexplained mystery: why did we have such an inefficient recovery, involving repeated momentum dumps over such a long time? To attack this question, the fault tree in Fig. 10 was constructed. This section systematically explores the branches of the fault tree.

The repeated momentum dumps (Fig. 10) could result from any combination of the following:

- Spacecraft momentum actually continued to exceed the Red value (branch 1.1 of the fault tree).
- Spacecraft momentum was actually low, but a sensor falsely reported it as exceeding Red (branch 1.2).
- A processor was corrupted (bad algorithms, code, data structures, hardware, or single event upsets) (branch 1.3).

High Actual Momentum (Branch 1.1 of the Fault Tree)

Momentum Stored in Fluids, or Slosh (Branch 1.1.1)

Could high system momentum result from momentum stored in fluids (i.e., slosh)? NEAR's three fuel tanks have internal diaphragms which, among other things, control slosh. But NEAR's two oxidizer tanks (OT), located one above the other on the spacecraft *Z*-axis, have no such internal propellant management device. Furthermore, NEAR's tanks were approximately half full at RND1 (Table 1), roughly the worst case for slosh. One postulated scenario is: (a) the settling burn "organizes" the 82 kg of oxidizer into a contiguous mass; (b) after RND1 abort, the spacecraft slew sets this mass in motion; (c) something approximating a spherical shell of oxidizer mass continues to rotate, coupling momentum to the body through viscous stresses at the OT interior walls; (d) the peak motions thus induced throw the gyro into WAM.

During the design of NEAR, slosh was analyzed for propulsion subsystem supplier Aerojet by Southwest Research Institute (SWRI) and reported in Ref. 6. This analysis established the parameters used for the slosh model and helped design a fuel-efficient settling burn strategy. The analysis also showed, among other things, that the time for the wall–liquid momentum coupling postulated in (c) above could take on the order of 15 min, which might help explain some of the delayed effects that occurred later in the RND1 timeline. Thus we felt it important to include slosh dynamics in the RND1 simulations and analysis.

Table 1. NEAR's propellant tanks.

	Fuel	Oxidizer
Number of tanks	3	2
Tank nominal radius	0.279 m (11 in.)	0.243 m (9.56 in.)
Propellant management device	diaphragm	none internal
Ultimate capacity of tanks	276 kg	173 kg (22% ullage)
Load at launch	215 kg	109 kg (37% ullage)
Load at RND1	165 kg	82 kg (47% full)

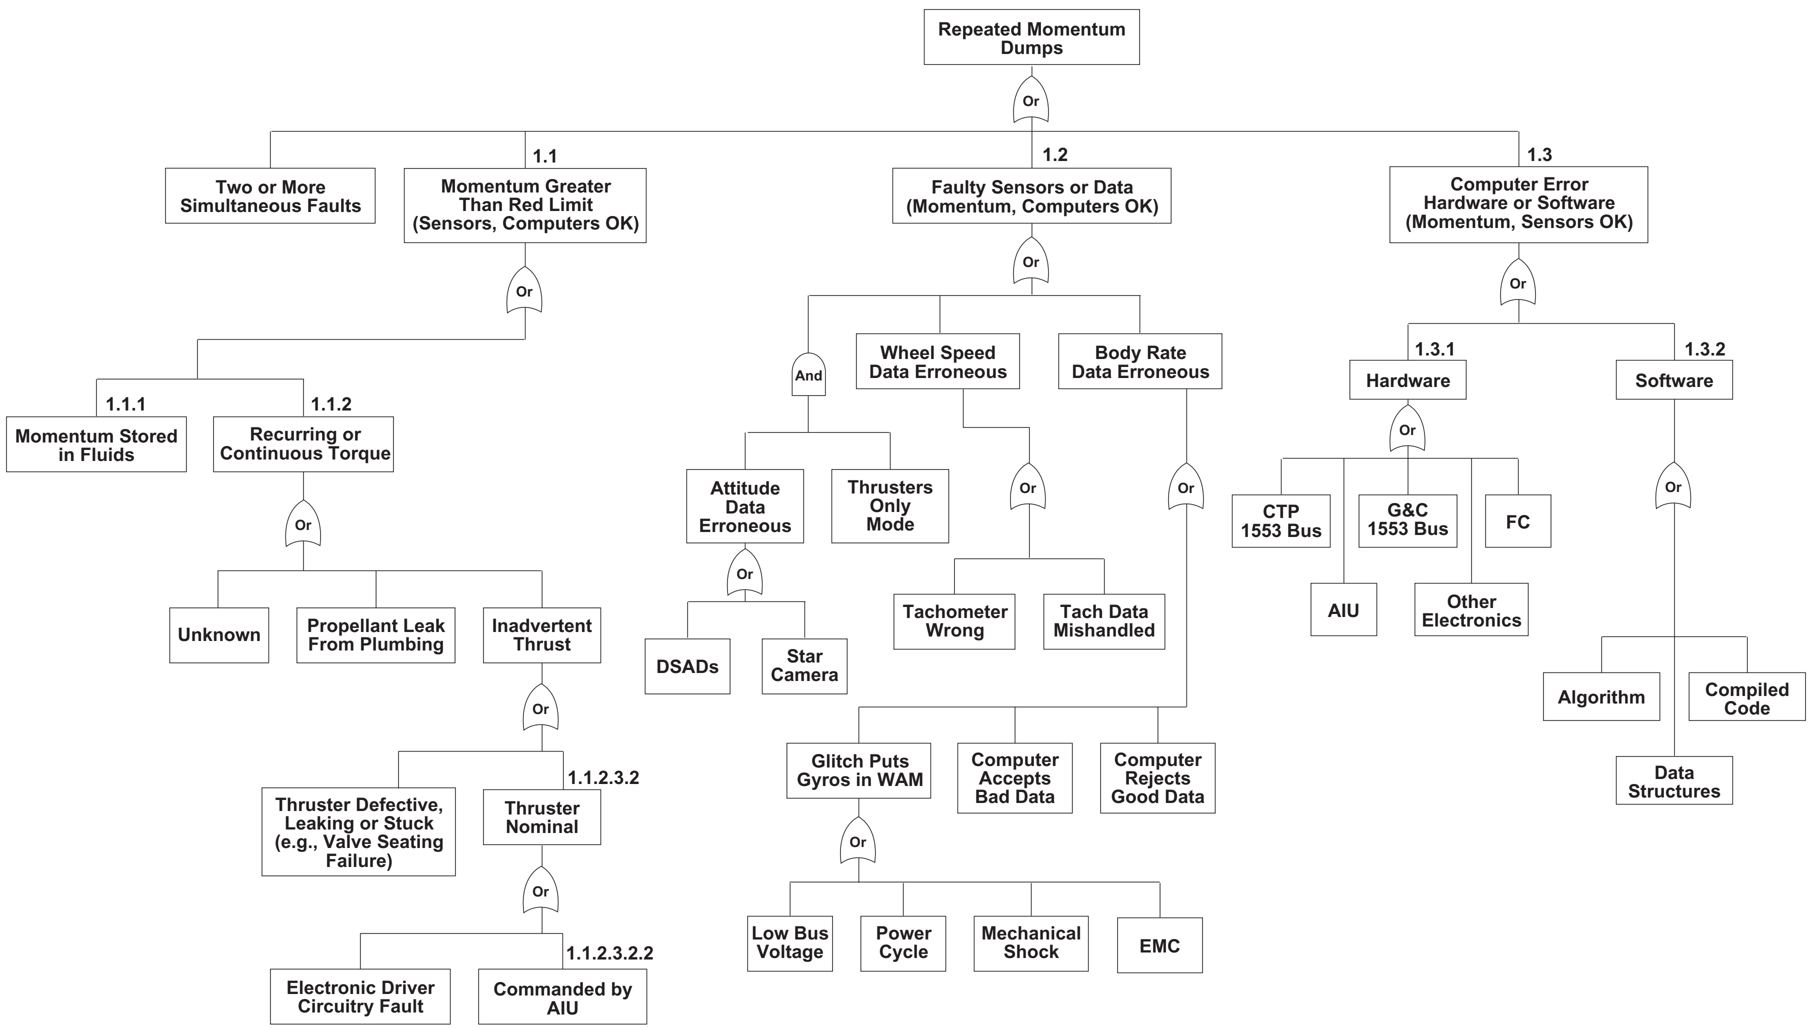

Figure 10. Fault tree of possible causes of the RND1 repeated momentum dumps. Certain events result from AND-ing or OR-ing other events. Numbered events are discussed more completely in the text.

Slosh Modeling and Analysis

Propellant slosh is modeled in the simulations by spherical pendulums in each of the three fuel and two oxidizer tanks (OT). This commonly used pendulum model is two-dimensional and does not model rotational slosh. Because the slosh model is computationally intensive, it was usually disabled on the brassboard to ensure that the simulation would run in real time within the allowed time slice. Investigations on the stand-alone, software-only simulator found that at high body rates and low fuel (as after the RND1 abort), the OT pendulum states were diverging numerically. Thus, for most brassboard RND1 abort simulations, revised slosh parameters were used that made these stable for very high body rates, high enough to induce WAM. The fuel tank pendulums were not changed (the upgrades to the brassboard slosh model are further discussed in Appendix F). Various large-momentum test cases were run with these parameters, in which the gyros enter WAM and the system dumps momentum, but slosh dynamics do not numerically diverge.

For reference, Table 2 shows the OT parameters that were changed, listing the old and new values used on the brassboard.

Table 2. Oxygen tank slosh parameters used in original and revised brassboard models.

Parameter	Description	Old	New
K_S	Spring constant	0.004 Nm/rad	0.01 Nm/rad
AlfaMax	Max angular acceleration	10.0 rad/s^2	0.4 rad/s^2
ZetaOx	Damping ratio (unitless)	0.004	0.01
DminOx	Min damping coefficient	0.015 Nm/rad/s	0.020 Nm/rad/s

It is important to recognize that these parameters, old or new, are somewhat arbitrary. For example, there is no strong physical basis for a restoring spring in the OT pendulums. However, in nearly zero-*G*, it seems intuitively reasonable that the liquid in the OT will migrate to a preferred location in the tank. A weak spring was added at the pendulum pivot to satisfy this intuitive desire, and its spring constant (K_S) is entirely arbitrary (as long as it is small). When parameters were adjusted by trial and error to improve performance, K_S was one of the parameters adjusted, and the "new" value listed in Table 2 is simply the final value used. Recent experiments have shown that K_S is not critical and can be set to zero without qualitatively changing OT slosh behavior. AlfaMax and DminOx are clipping limits introduced in the code to prevent runaway behavior. It seems intuitively reasonable that the angular acceleration of the liquid cannot be arbitrarily large, but the value of this limit for the pendulum model is strictly ad hoc. The damping ratio, ZetaOx, is traceable to the original slosh report (Ref. 6) and was chosen empirically to produce slosh oscillation time constants in settling burn simulations that were about as predicted by the report. ZetaOx is unitless, however, and the actual damping coefficient (torque per unit rate) is calculated in the code as a function of propellant mass, applied acceleration, and geometry as well as ZetaOx. The clip limit DminOx was introduced to prevent numerical exceptions in these calculations.

More recently, after most RND1 brassboard simulation activity had been concluded, the small trajectory corection maneuver scheduled for NEAR on 20 October 1999 (TCM20) was simulated on the stand-alone simulator. It diverged numerically with slosh enabled but worked fine with slosh disabled. On investigation, the latest problem was traced to the fuel tanks, which have diaphragms ("bladders") to hold the liquid against the outlet even if the tank is inverted. The slosh pendulum characteristics vary with the amount of propellant remaining and the current acceleration vector. The inverted-tank pendulum model (i.e., when acceleration causes fuel to "fall" away from the outlet) was diverging at low fuel load. The fix applied (after much iteration) is to always use the upright tank model whenever the remaining propellant is below a threshold. This is a reasonable compromise, since slosh has diminishing effect on vehicle dynamics as propellant decreases; i.e., larger approximations are reasonable for smaller effects. This worked well for the TCM20 simulation.

With these fixes to the fuel-tank pendulum models, additional simulations were run of RND1 post-abort behavior. We found that slosh dynamics were adequately stable with the original OT pendulum parameters (the "Old" column of Table 2), except for Max Angular Acceleration (AlfaMax). With these values it is still possible to produce a situation that has unstable OT slosh dynamics, resulting in overall system dynamic instability and eventual numerical blowup. The case most thoroughly investigated has an arbitrary high body rate introduced, which ordinarily results in an autonomous momentum dump and recovery. As AlfaMax is increased, slosh-induced oscillations increase, and the dump takes longer to complete, until eventually the simulation diverges. We know the divergence is not due to an unstable interaction between the control laws and the slosh dynamics, because for high enough AlfaMax it diverges *even with no control at all* from either thrusters or wheels. This result is clearly physically meaningless because it has increasing (in fact diverging) momentum and energy with no applied external torque. In these divergent cases the OT pendulum frequencies are above

2 Hz, much higher than is supported by any known data. As long as the OT slosh pendulums are constrained not to exceed about 1 rad/s^2 of angular acceleration (AlfaMax), the simulation remains stable and reasonable and can be used to study control performance in the presence of slosh-like disturbances. In fact with AlfaMax = 1, the slosh states slowly diverge with no control at all, but are damped with control enabled. This lends confidence that NEAR's control laws are not the source of the instability, and in fact can even deal with instability in physically unsupportable slosh dynamics.

Stand-alone simulations of an AIU-only B-thruster autonomous momentum dump case were used to explore the effect of the various slosh parameters. This case introduced an arbitrary tumble rate of about 9°/s about the *Y*-axis, which triggered an autonomous momentum dump. With slosh disabled in the simulation, the dump completed (momentum Green) about 50 s after it began. With slosh parameters set at values that are reasonable, but believed pessimistic, we found that the dump also returned to Green but took longer, about 125 s. There were noticeably greater oscillations in angular rate and momentum estimates, and hence in thruster activity, for the slosh case. So it is evident that slosh does affect the behavior of momentum dumping and that our simulation does model this. However, there is clearly significant uncertainty in the slosh model. Any simulation results must be considered as qualitatively representative of reality at best.

This incomplete situation for slosh is not entirely comfortable. Some interaction between thruster control and slosh is clearly possible, and we still cannot rule out that slosh was a factor in the RND1 post-abort misbehavior. It is probably not possible to answer definitively and quantitatively how much the effect is or was. Very possibly, slosh contributed to some of the inefficient or apparently ineffective dumping, particularly after the gyros entered WAM. However, with any physically defensible set of slosh model parameters tried, we have not been able so far to demonstrate that control–slosh interaction can produce anything like the body rates necessary to induce WAM. This work is still in progress because of its importance for future missions. The NARB's report was not held up, however, to await a satisfying conclusion to this particularly difficult modeling problem.

Given the uncertainty and difficulty of analyzing slosh dynamics, we questioned the NEAR methodology for modeling and analyzing slosh and compared it with what other spacecraft programs have done. The consensus from several peer organizations (Orbital Sciences, Litton Amecom, Lockheed Martin Astronautics, the Naval Research Center, Goddard Space Flight Center, Aerojet, and the Jet Propulsion Laboratory) is that the pendulous mass approach used for slosh modeling on NEAR is common in the industry. Furthermore, slosh is generally considered a more critical concern for spinning spacecraft. For 3-axis stabilized spacecraft such as NEAR, most analysts do not consider it a major concern for stability, although it is of interest for jitter.

Flexible body interactions must also be considered. Since the second trajectory correction maneuver (1996), NEAR simulations have modeled these effects (force and torque reactions on the main spacecraft body). Solar array flexure is the principal source of these effects. It is modeled as single-cantilever bending modes of each of the four solar panels, with frequency and damping empirically tuned to match flight data. The resulting effect is seen as sinusoidal oscillations at about 2 Hz, excited by thruster firings and decaying with a time constant of about 10–30 s. These oscillations are evident in both accelerometer and gyro data on all burns. Typical amplitudes of spacecraft reaction are less than about 10 μrad/s of angular rate and about 0.1 μg of acceleration. Flexure modeling was activated in our simulations whenever fuel slosh was enabled, that is, in all RND1 brassboard simulations after number 51. The small flexure effect, like slosh, cannot of course explain one of the principal mysteries of RND1: why did the spacecraft begin dumping momentum again after long quiescent periods (as Fig. 6 shows)?

Other Torque on Spacecraft (Branch 1.1.2)

Could there be a recurring or continual torque on the spacecraft? Absent repeated hits from outside the spacecraft, something that is utterly implausible, such torque could result from either a propellant or a pressurant leak or from propellant actually exiting a thruster nozzle. There has been no evidence of a propellant or pressurant leak before or after RND1, so a leak involving any part of the propulsion system other than a thruster seems most unlikely. There are two ways propellant could erroneously exit a thruster: a thruster valve could actually leak, or a thruster could be erroneously commanded to open.

NEAR's thruster valves (Wright Components P/N 18207 for the 4.4-N thruster and P/N 18208 for the 22-N thruster) contain two series-redundant, normally closed valves. Each valve has its own seat and return spring and is activated by an independent solenoid coil. The coils are tied together at the propulsion subsystem interface and driven by circuitry in the AIU. A fine mesh (20-μm) absolute filter at the valve inlet protects the upstream valve seat, and the same kind of filter screen on the injector protects the downstream seat. The redundant valves are leak-checked individually prior to integration into the propulsion subsystem, but once their coils are wired in parallel they can no longer be tested independently. Could a thruster valve be held fully or partially open by a particle, possibly trapped *between* the filters at the time of manufacture? The postulated scenario is that a particle would hold open the valve at some partial thrust level until it was either flushed clear by the propellant flow or eventually pounded into the soft elastomeric seat. Given that both of the series valves would have to be held open by particles at once, and that there has been no evidence of thruster leakage either before or after RND1, this possiblity seems most unlikely. Nevertheless, we explored ways such a thing could occur and also explored the possible consequences.

Extensive discussions were held with the supplier of the NEAR thrusters (Primex) regarding the possibility of thruster valve leakage. The same topic was discussed with other thruster valve makers, including Moog and Marotta. The anecdotal evidence indicates that particle-induced leakage has in fact occurred in the past with almost everyone's thruster valves, including Wright's. One mechanism is for the particles to shed from an internal component, such as a spring, and become trapped between the protective filters. Such particles seem to be caught by the extensive predelivery ground testing and have not been documented to cause problems in orbit. Another mechanism occurs when a spacecraft is vibration-tested and launched with a thruster facing up (as two NEAR thrusters were). Vibration can cause the iridium-coated alumina catalyst particles to grind against each other; the "fines" so generated can pass through the downstream valve outlet filter into the seat area, where they can agglomerate into a larger particle. Such particles would be flushed clear in the first few thruster firings. In both cases, of course, this would have to affect both series valves, an implausible event.

Despite the unlikelihood of thruster leakage from any cause, we did explore its likely consequences in a series of simulation runs (Appendix L). A series of simulations, beginning with run 44, was conducted in which various FVC thrusters were held open for varying amounts of time and at various levels of thrust, from 100% to 10% of rated thrust. While an A-thruster held 100% open for several seconds would impose sufficient body rate to drive one gyro axis into WAM, no leaky thruster (20% thrust) case caused momentum that was not eventually reduced by the G&C system.

In addition to a thruster valve seat leaking, software or drive electronics failures could erroneously command a thruster ON, which would have essentially the same effect as a leak. Branch 1.1.2.3.2 of the fault tree (Fig. 10) results in one or more FVC thrusters being falsely directed to open, imposing a momentum-raising torque on the spacecraft. The AIU circuitry that drives the thrusters was reviewed. This circuitry consists essentially of digital one-shots that hold a thruster open for only 20 or 40 ms. Holding a thruster open longer than this requires a repeated series of software commands. This approach, which was implemented in digital *hardware*, was specifically designed to prevent a "Clementine-type" loss-of-propellant failure.* Midway through our investigation, NASA's WIRE spacecraft was lost due to a pyrotechnic anomaly that was subsequently traced to a misapplication of Actel Field Programmable Gate Arrays (FPGAs). Because such FPGAs are also used in the AIU digital one-shot circuitry, we again had this circuitry carefully reviewed, this time by a senior circuit designer particularly experienced with Actel usage. Neither review revealed any design flaw in the AIU thruster-driving circuitry that might hold a thruster open.

The other way that thrusters could be erroneously commanded ON is by software. This branch (1.1.2.3.2.2) is discussed later.

Low Momentum Falsely Reported as High (Branch 1.2 of the Fault Tree)

NEAR's G&C sensors measure attitude (using DSADs and the star camera), body rate (using the gyros), and wheel speed (using tachometers). "Bad data marked as good" is a known deadly failure mode for any control system. For this reason, the NEAR G&C and C&DH (command and data handling) subsystems contain numerous cross-checks on data quality.

DSAD data were examined up to and through RND1. There was no evidence of wrong vectors being reported. There was some evidence of slight imbalance between the DSAD heads, but the imbalance was still within specifications. A series of simulations with failed DSAD heads did not replicate anything similar to the RND1 behavior. The star camera appears to be irrelevant to the RND1 discussion.

Wheel speed errors can occur from tachometer errors (for which there is no evidence) or through mishandling by the software of a properly reported wheel speed. We know that NEAR *was* vulnerable to this latter situation due to an improperly set software data structure value. Simulations on the stand-alone, software-only simulator have shown that this error can cause

*The Clementine spacecraft suffered a catastrophic loss of propellant on 7 May 1994 that prevented it from accomplishing its goal of flying by and photographing the asteroid Geographos. Clementine's hydrazine supply was depleted by thrusters being erroneously held open for 11 min by the spacecraft's control processor. Instead of hardware one-shots, Clementine used 100-ms thruster protection timers implemented in software (Ref. 7); this software contained an undetected bug.

some, but by no means all, of the features observed in RND1 (see the earlier discussion of the misreported wheel speed error, under Simulation Results). However, this error has not been simulated on the brassboard simulator over a full range of initial body rates and system momentum.

Body rate is measured by the 3-axis gyroscope package. The G&C system added software filtering to reject gyro noise, which at delivery was higher than specified noise levels. This filtering was turned off (intentionally) at RND1 to raise the bandwidth for the burn.

Gross errors in gyro output seem to be sufficiently well detected by G&C software tests on the 1553 bus messages and on the values themselves. However, if the gyro should enter WAM its output becomes extremely noisy, a type of "bad data reported as good." All three gyro axes were found in WAM when NEAR was recovered following RND1. We know that if one gyro axis enters WAM, the other two active axes are not automatically set to WAM—instead they can enter WAM only if the resulting spacecraft motions cause high rates around those axes. Simulation 89 and others showed that, with one gyro axis intentionally kicked into WAM, the other two axes are not automatically pulled into WAM. The most obvious explanation for a gyro entering WAM is that body rate did actually exceed the value at which the gyro's force-to-rebalance loop can no longer maintain lock (roughly 11°–14°/s) . Of course that leaves the question of where the required angular momentum came from. (High angular momentum was discussed earlier, under "High Actual Momentum, Branch 1.1 of the Fault Tree.")

The exact *time* of entry into WAM was of great interest, for it could give valuable clues as to the cause. We know for certain, from the limited telemetry records, that the gyros were *not* in WAM at $T + 11$ s. If WAM was induced by high body rates, it must have occurred after the first crossing of the momentum telemetry reading of 25.5 Nms (saturation value) at $T + 50$ min, corresponding to roughly 4°/s. It probably occurred well before the first rate violation, at $T + 01{:}21{:}00$.

The NARB was particularly interested in exploring causes other than body rate for a gyro to enter WAM. Postulated mechanisms included mechanical shock, data bus errors, electromagnetic interference (EMI), and the effects of the sagging supply voltage. We were also interested in exploring the hypothesis that the output noise amplitude in WAM might be rate- or voltage-dependent. This would provide a mechanism for higher-than-measured WAM noise values that cause repeated momentum dumps in our simulations and fit the observed RND1 behavior of eventual recovery at low noise levels. Since body rates during RND1 simulations with high WAM noise are actually quite low (<1°/s), any rate dependence would have to be very strong to produce the seven- to tenfold increases in noise values that the simulations show are required to reproduce the behavior.

NAG members and an APL control circuits specialist met with representatives of the gyro supplier (Delco, now Litton) on 22 April 1999 to review circuit schematics and to brainstorm and explore possible non-rate-induced entries into WAM. Low bus voltage alone can be ruled out, since the gyros were operated in prelaunch tests down to 23 V without entering WAM, and we were able to repeat those tests at Litton. EMI from thruster firings alone can probably be ruled out, since numerous burns on NEAR did not result in the gyros entering WAM. However, EMI from nearby thruster solenoids, *in the presence of reduced bus voltage*, was a plausible hypothesis. Mechanical shock can be ruled out as the mechanism because we do know that the gyro was not in WAM for the first 11 s after the LVA engine abort. The gyro cannot be commanded into WAM through 1553 bus actions alone. However, the gyro *can* be placed in WAM through its test port. This was actually done once inadvertently during spacecraft thermal vacuum testing, but in an unusual breach of procedure no Problem/Failure Report was written for this anomaly. In any event, this test port is capped off and unused in the flight hardware. We decided that a series of tests would be most useful to explore the effects of body rate and supply voltage on the noise and to explore the hypothesis about the combined effects of EMI and bus voltage.

Unfortunately, the NEAR program did not procure an engineering model nor any other gyro than the flight unit. As the NEAR unit was a unique configuration (superseded by newer designs), Litton had no units available that emulated the NEAR gyro and its WAM. It was not until quite late in the investigation that Litton was able to assemble a test article resembling the NEAR gyro, including WAM and the proper voltage DC/DC converters. This test gyro was available only for a few days, after which it had to be surrendered to another program. APL engineers prepared a test plan (including a day of EMI conducted susceptibility tests), and an APL representative traveled to Litton to direct and witness the tests. Unfortunately, Litton was not able to provide a rate table or the equipment and cable needed for EMI testing during this narrow window of opportunity. Therefore we were unable to determine the noise dependence on body rate or to attempt to "kick" the gyro into WAM through EMI. The only data taken were noise versus supply voltage data (Ref. 8). Noise showed only a weak dependence on supply voltage over the range experienced during the RND1 events. Most significantly, Litton data showed that the measured noise levels in WAM, at least at near-zero body rates, were much lower than the levels required in our simulations to produce repeated momentum dumping.

Processor and Software Errors (Branch 1.3 of the Fault Tree)

Processor Hardware Errors (Branch 1.3.1)

A *hardware* failure in either of the three processor types (FC, AIU, CTP; two units of each), if it occurred, would have to have been intermittent or transient, because there is no evidence of such failure before or after RND1. One class of hardware-related possibilities explored by the NARB is bus errors, particularly on the G&C and CTP 1553 buses. For example, could instructions to or data from the gyro become bit- or byte-shifted and misinterpreted? Since the FC maintains a quality check by keeping track of the order that data arrive, this might also require the FC's 1553 message counter to be reset.

The code inspection team examined this issue. The FC resets its 1553 buffer only at start-up. The 1553 bus protocol is word oriented, with 20 bits transmitted per message. Extra parity bits are checked to assure that each word is transferred correctly. The entire message is also checked, for example, to ensure that the correct number of words is transferred each time. The FC checks this status word for each transfer, and then clears it so as not to be fooled by stale data. The FC software also checks the status word in each message from the inertial measurement units (gyros and accelerometers). If words or bytes were shifted, the FC would reject the entire message. If the *entire* message is in error, both the AIU and the FC would receive bad data. Although hypothetically possible, there is no evidence that this ever happened; the probability is about of the same order as that of the AIU or FC software getting momentarily corrupted and then spontaneously fixing itself without resetting a processor.

A final type of hardware error is the single-event upset (SEU). Both the RTX2010 processor used for the CTP and AIU and the Honeywell 1750A FC processor are extremely hard to single-event effects. The FC protects its memory against SEUs with error detection and correction (EDAC); the CTP and AIU memories are inherently SEU-resistant. Prior to launch, APL had calculated an upset rate of one every 300 days for the FC, and at least two orders of magnitude better than that for the RTX2010. The vast majority of SEUs are either corrected on the fly by EDAC or are trapped and cause a reset. No processor resets had been observed in the 3 years prior to RND1, and we know for certain that no processors reset during RND1. It is possible (but not provable) that the FC1 spontaneous reset of February 1999 (see below) was caused by an SEU; this would be consistent with the calculated rate. In any event, it seems highly unlikely that SEUs played any role in the RND1 events.

Processor Software Errors (Branch 1.3.2)

"Software" error is taken loosely here to mean a number of things: an error in an algorithm; an error (bug) in compiled flight code (in either the AIU, FC, or CTP processors); or a "data structure" error in the tables of coefficients that are uploaded to the spacecraft. We have already seen that it was a data structure error (setting an acceleration parameter too tight) that precipitated the entire RND1 event; another data structure error can cause the misreported wheel speed problem. Together, the FC and the AIU comprise about 80,000 lines of source code.

Two main tools were used to search for all three types of software error:

- The series of simulations already described, which ran actual RND1 flight code on ground hardware replicas of the flight AIU, FC, and CTP computers
- Auditing by the code inspection team

Code Inspection Team

The brassboard used for simulations had exact copies of the RND1 flight code in its AIU and FC processors. In theory, if flight software problems contributed to the RND1 behavior, they should reveal themselves through the simulations. However, there is always the possibility that the simulations might not exercise a particular region of code in a reasonable number of tries. So a software inspection team was formed to analyze the G&C flight software for defects that might have contributed to the RND1 behavior deduced from the timeline reconstruction. The team looked for design flaws as well as coding errors such as mishandled exceptions or interfaces, faulty control logic or timing, and so forth. The code inspection team leader was a member of the NAG and attended the NAG meetings. Simulation results provided clues on areas of the code to examine first and what to look for. For example, a suspiciously large spike in Sun angle seen in simulation 92 was analyzed by G&C and the code inspection team and led to the discovery of an AIU code error in the way the solar array Sun vector is handled. In addition, the code inspection team provided valuable support for necessary software fixes and upgrades to the brassboard.

The team initially consisted of three software experts and two G&C-cognizant analysts; it was later reduced to three members. The team had a mix of NEAR developers and others who were completely independent of the

NEAR development. Original NEAR code developers provided consultation to the team as needed. "Corporate memory" of the NEAR code development was therefore available, but was kept from "contaminating" an independent evaluation of the architecture, design, and implementation of the G&C software.

The team's approach was twofold. First, while the rest of the NAG was reconstructing the timeline, conducting simulations, and generating hypotheses, the code team independently reviewed the entire software in detail, in some cases having to reverse engineer parts of the G&C code because of inadequate documentation. Once all team members were fully educated on the G&C software design and implementation, *what if* scenarios and *use case* analyses were performed against the design and code. Second, in response to specific NAG hypotheses, the directly relevant software that might have contributed to the anomaly was scrutinized for defects using code walk-throughs and design reviews. All G&C software was examined, but it should be pointed out that a complete code review of all 80,000 lines of G&C code was beyond the resources of this investigation.

The G&C software consists of two major components: AIU code version 1.06 and FC code version 1.11, the versions active during RND1. The AIU processor is an RTX2010, with software consisting of approximately 21,000 lines of C and 10,000 lines of assembly code. The FC is a 1750A running approximately 42,000 lines of Ada and 7,000 lines of assembly code. In addition to flight code, the inspection team also scrutinized the autonomy rules that safed and operated the spacecraft during the anomaly. Appendices C and D give overviews of NEAR's Guidance & Control and Autonomy systems, respectively. An area of particular interest to the NARB was the behavior of processors during AIU switches. An AIU requires 7–9 s to boot up; during this time, no processor is in control of the G&C actuators. However, the code inspection team confirmed that software-driven thruster activity *cannot* take place during this time; this was also confirmed by brassboard simulations.

The code inspection team concluded that the G&C software is sound in design and implementation. That is not to say that no defects were found: the team found a total of 17 errors, including 9 in the compiled code. In addition, certain poor software design or programming practices were identified, although they did not directly impact the health of the spacecraft. Defects ranged from expressions having identical constants to "save" variables being overwritten under certain conditions to poorly designed exception handlers. Each error or defect was analyzed, and many were studied in simulations. For example, two different series of stand-alone simulations were run to force gross errors in momentum estimates so that the effects of the misreported wheel speed error could be studied. Initially, the AIU-only cases erroneously added momentum to a near-zero momentum system, raising it above the White limit. But all cases recovered to a low-momentum, stable safe mode in less than 20 min. This error and others were discussed in the section "Simulation Results" earlier; some of these software errors could have moderately prolonged the post-RND1 recovery or made it less efficient (more wasteful of fuel). However, in none of the simulations did the software errors alone fully account for the 15 autonomous momentum dumps, the prolonged recovery, and the amount of fuel used in the actual RND1 event. Table 3 identifies the defects found and their disposition.

Before leaving the discussion of processor hardware and software, we should note one final point. On the evening of 24 February 1999, Mission Operations discovered that, following completion of a normal pass and return to passive momentum management the day before, NEAR's FC1 processor (running version 1.11 code) had spontaneously rebooted, demoting the spacecraft to Earth-safe mode. Because Mission Operations had failed to update the orbit stored in EEPROM, the Sun angle was way off. This violated SKI rules, so NEAR soon demoted to Sun-safe mode and AIU2 took over. Fortunately, there was no fuel use or LVS (therefore, solid-state recorder data were preserved). An intensive examination by the code inspection team failed to uncover the cause of this reboot, the first ever for any NEAR processor. Unfortunately, Mission Operations had not, during the 2 months since RND1, downloaded a memory image of suspect processor FC1. With the reboot, that opportunity was now lost, and the team had little data to work with. Included in the team's investigation was a margin analysis of the FC 1.11 code. Although some timing margins were tight, they all appeared to be sufficient, in the worst case. Following the reboot, NEAR was operated on FC2 (and its prior 1.10 code) until the team's investigation was completed. When NEAR was returned to FC1 and the 1.11 code on 4 August the spacecraft immediately demoted to Sun-safe mode, this time because Mission Operations had failed to *activate* the new orbit. While this event did not shed much light on possible processor hardware faults, it does illuminate some shortcomings in Mission Operations' error rate.

Table 3. G&C software and data structure defects identified.

	Unit	Description	Resolution	Status
1	AIU	Yellow momentum is safe, so continuing a post-abort dump too long is undesirable	Change AIU momentum "Yellow-Good-Enough" limit time-out to 120 s from 270 s to finish dump ASAP	Data structure changed and uploaded
2	AIU	When switching to safe mode there is a potential to still be using thrusters	Force AIU to enable wheels when going to safe mode	Software Change Request generated and being worked
3	FC	Green momentum limit of 0.5 Nms is unnecessary and takes too long to reach	Change FC Green limit to 0.9 Nms	Data structure changed and uploaded
4	AIU	AIU Green limit should be higher than FC's (changed to 0.9 Nms, item 3 above), so AIU will not prolong a dump that FC has finished	Change AIU Green limit to 1 Nms	Data structure changed and uploaded
5	FC	A 30-min wait (Red) dump after an abort can cause problems; immediate (White) is preferable after abort	Change FC White momentum limit to 4 Nms for quicker dump in case of abort	Data structure changed and uploaded (for burns only; changed back to 6 Nms after burn)
6	AIU	A 30-min wait (Red) dump after an abort can cause problems; immediate (White) is preferable if abort	Change AIU White momentum limit to 4 Nms for quicker dump in case of abort	Data structure changed and uploaded (for burns only; changed back to 6 Nms after burn)
7	AIU	FC could drive spacecraft through a large angle before the AIU takes over	Shorten AIU time for AIU-only Sun-safe-rotate due to SKI violation to 30 s from 300 s	Data structure changed and uploaded
8	AIU/ FC	High wheel speeds can be flagged as invalid and report as zero speed	Change upper limit on data structure value controlling wheel speed upper and lower limit tolerances	Data structure changed and uploaded
9	FC	An incomplete FC-to-AIU High Rate message could be sent	Check the return status of Build_Msg and flag output message as error	Problem/Failure Report and Software Change Request generated and being worked
10	FC	An invalid Sun vector could be sent to the AIU by the FC	FC_Sun_Stat should be initialized at the beginning of its processing section to indicate an invalid Sun Vector	Problem/Failure Report and Software Change Request generated and being worked
11	FC	Under some conditions a reset of the Sun-angle Sun filter can erroneously reset the DSAD/FC filter	Do not overwrite the DSAD/FC filter when the Sun-angle Sun filter is reset	Problem/Failure Report and Software Change Request generated and being worked
12	FC	A request from the FC to AIU for demotion to Sun-safe-rotate mode may not be granted	Expression had the same words (Sun-safe-freeze mode) twice instead of Sun-safe-freeze mode and Sun-safe-rotate mode	Problem/Failure Report and Software Change Request generated and being worked
13	AIU	FVC_Request discrete is being lowered at an inopportune time	Clear FVC_Request bit at a better time	Software Change Request generated and being worked
14	AIU	AIU can report a ±1-s error or the same second twice in a row due to AIU and CTP clocks not being synchronized	Implement AIU clock walk fix developed for TIMED mission	Problem/Failure Report and Software Change Request generated and being worked
15	FC	FC can report a ±1-s error or the same second twice in a row due to its estimation scheme using MET and UT	Implement the time tag calculation summarized by the following formula: Current_UT + .040 * minor_frame_number – message_age – TickFudge	Problem/Failure Report and Software Change Request generated and being worked
16	AIU	Autonomy rules require quite a few seconds to execute; 13 s time-out is too short	Change AIU FVC_Request time-out to 60 s from 13 to prevent its going inactive too soon	Data structure changed and uploaded
17	AIU	Violation of compiler vendor data passing method	Copy the global in a local data store and pass the local to the function	Problem/Failure Report and Software Change Request generated and being worked

FINDINGS AND RECOMMENDATIONS

Despite the loss of spacecraft data when the solid-state recorder was powered off during the LVS event, it was possible to reconstruct a general understanding of the post-RND1 timeline. More than 128 simulations were run on the NEAR brassboard simulator, an independent review of the flight code was conducted, and suspect hardware and circuit elements were reviewed to attempt to explain the protracted RND1 recovery. Although software errors were found that could contribute to the protracted recovery, they do not fully explain it. We were unable to find a complete explanation for the post-RND1 events. Nevertheless, it is possible to make observations and recommendations that might prevent a recurrence on NEAR or on other programs.

Finding 1. The precipitating event for the RND1 anomaly was the burn abort. The burn aborted because the main engine's normal start-up transient exceeded a lateral acceleration safety threshold that was set too low. A start-up transient very similar to that of the RND1 burn had been observed during the only previous use of the LVA engine, the DSM burn of July 1997. The NEAR team failed to recognize the significance of this transient and its potential increase in importance with decreasing propellant mass. NEAR's structural design, with a separate cantilevered propulsion module, combined with the placement of the accelerometer package, exacerbates the coupling of the transient into the accelerometers.

An error in the burn-abort command script initiated the spacecraft's anomalous attitude motions and the protracted recovery events. A formal review of the RND1 scripts had been held, as required, but failed to catch the missing command. RND1 simulations were run on the stand-alone, software-only simulator, but this simulator does not emulate the actions of the autonomy rules and so did not catch the script error. The RND1 command scripts were simulated on the NEAR brassboard simulator, but abort cases were not run. The brassboard is difficult to use and can give "strange" results; this discourages thorough testing of command scripts. When abort cases were simulated *after* RND1, every simulation exhibited the disastrous effect of the abort command script. It is likely that without this script error, NEAR could have been recovered in time to execute one of the backup rendezvous burns, allowing Eros orbit insertion in January 1999 as planned.

The safing modes on NEAR are overly complex, involving frequent hand-overs of control among multiple computers and too many logical paths and branches. This complexity adds risk to an anomaly such as that of RND1. If a safe mode must be computer-based (analog control might be better), the processor modes and supporting software must be *simple* and readily de-buggable so that the safe modes can be relied upon. Nevertheless, NEAR's autonomy system performed every action it was designed to perform, and there is no evidence that it did anything that it was not supposed to do. The autonomy system was tested prior to launch (in particular, burn aborts, momentum dumping while in Sun-safe mode, and excessive thruster use scenarios were tested). However, it should be remembered that the autonomy rules were tested with *earlier versions* of the AIU and FC codes, not those running during RND1. The failure of the autonomy system to prevent the kind of fuel-loss scenario it was designed to guard against demonstrates a clear design flaw.

Recommendation 1a: The NEAR mission cannot survive a repeat of the RND1 events. The NEAR system engineer must be given the responsibility and authority to form a team to (1) create and test a burn abort command macro that minimizes the risk of a future burn abort leaving the spacecraft with high system momentum, and (2) redesign the autonomy response to excessive thruster use.

Recommendation 1b: The RND1 event was precipitated by an incorrect parameter. Parameter (data structure) errors can be every bit as deadly as errors in compiled code, and they need the same careful design, review, testing, and configuration management. G&C data structures should be placed under configuration control that considers the risk/benefit of changed parameters and reviews the adequacy of analysis and testing of new values.

Recommendation 1c: Formal review of command scripts can catch many, but not all, errors (particularly if the reviewers are not independent). For critical events, complete testing of all command scripts on the NEAR brassboard simulator is required for safety. For example, burn abort cases, not just success cases, must be run on the simulator. Future programs must have simulators with sufficient fidelity and user-friendliness to encourage this practice.

Recommendation 1d: The NEAR system engineer should be required to review and sign off all critical scripts. This was once the practice within Mission Operations, but at some point the practice ceased.

Recommendation 1e: Regression testing of the entire autonomy system (including interactions between the autonomy rules and the AIU, FC, and CTP processors) should be carefully considered whenever G&C software or autonomy rules change. The original 21 scripts used for NEAR software acceptance testing should still be available to support regression testing.

Finding 2. There is no evidence that the prolonged and fuel-wasteful momentum dumping was caused by hardware failure.

- The propulsion system (and its drive electronics) did not cause the prolonged recovery. Although in theory a leaking thruster that spontaneously recovered could explain the initiation of the prolonged recovery sequence, there is no evidence that this occurred, and the Board considers it extremely unlikely. Simulated thruster failures do not reproduce the observed RND1 behavior. Furthermore, the Board feels there is no technical reason to avoid using NEAR's LVA engine again, should that be necessary.
- We found no evidence that DSAD or star camera failures contributed to the prolonged recovery. Simulated DSAD failures do not reproduce the observed RND1 behavior.
- We found no evidence of hardware failures in any of the six processors or in their interconnecting buses.
- We found no evidence that single-event upsets contributed to the prolonged recovery. Given the hardness of the NEAR processors, and the benign radiation environment, the probability of an undetected SEU is extremely small.

Finding 3. The NARB investigation uncovered 17 software errors, taken here to mean errors in algorithms or in uploaded data structures as well as in compiled code. Most of these were minor and do not necessarily need to be corrected. However, some are significant and contributed to the event, while not fully explaining it.

Recommendation 3a: NEAR's safing logic was not well designed to handle the burn abort, in that it allowed thrusters to fire immediately, before the command system's autonomy rules could change the attitude control actuators, reconfigure tanks, or change data structures. Although the design was demonstrated to be workable prior to launch, it has undesirable side effects that could have been avoided. Changing this now on NEAR entails unacceptable risk, but future programs with similar attitude control systems should consider this interaction more carefully.

***Recommendation 3b**:* The G&C design for momentum dumping temporarily abandons the requirement for Sun-pointing. The rationale behind this choice (perhaps driven by schedule pressures) was to have a single momentum dumping algorithm handle the immediate post-launch tip-off momentum dump as well as the operational dumps. Under stressed failure conditions this turns out to have been an unwise choice. Changing this design now on NEAR entails unacceptable risk, but this should be carefully evaluated for future missions.

The FC code version 1.11 includes an "enhanced burn" capability that maintains attitude for commanded (but not autonomous) momentum dumps. When redesigning the autonomous momentum dump safing logic, a careful evaluation should be made of which course is least risky: using this capability for autonomous dumps, or relying more heavily on ground commands.

***Recommendation 3c**:* The "misreported wheel speed" data structure error, in which a wheel at maximum speed is reported as zero speed, is serious and has been fixed.

Recommendation 3d: Having so many of the autonomy rule and G&C time-outs set to 300 s is confusing and can cause dangerous "race" conditions. Consider changing the time-outs to more easily distinguishable values. Verify by analysis and simulation that no negative effects result for other time-outs or autonomy functions.

***Recommendation 3e**:* The G&C subsystem should incorporate a coarse ephemeris to use in the event that Mission Operations fails to upload a valid orbit. (This recommendation has been partially implemented on NEAR by storing coarse ephemeris data in EEPROM in case of processor reset. The system will also use its last known position when the current orbit expires. This should no longer be a reason for NEAR to drop into Sun-safe-rotate mode.)

***Recommendation 3f**:* Momentum limits (Green, Red, etc.) were set too tight in some cases, causing or contributing to the loss of extra fuel. Adjust these momentum threshold levels to minimize fuel loss (already done on NEAR).

***Recommendation 3g**:* A more robust safing design would have been unaffected by the Mission Operations script error that allowed the thrusters to be used briefly when transitioning between momentum dump and attitude control modes. On NEAR, G&C does not produce the telemetry that would allow the command system autonomy rules to find and correct this error. Such "belt and suspenders" type safety checks must be implemented more intensively on future programs.

Recommendation 3h: The AIU is allowed to go into the FC control state ("FC Override") on boot, even if the system had been in AIU-only mode prior to boot. This violates the NEAR safing philosophy of no transitions from a lower to a higher operating mode without ground intervention. Again, changing this design now on

NEAR entails unacceptable risk, but this should be corrected on future missions having a processor architecture like that of NEAR.

Recommendation 3i: The use of a "Red Teamer" who continually tries to "break" the software design should be considered.

Finding 4. The NEAR team is to be commended on the speed with which they determined the proximate cause of the abort itself. The rapid diagnosis of the precipitating events, coupled with quick work from Mission Operations and the rest of the NEAR team, enabled the successful re-burn 2 weeks later that rescued the mission. The NEAR team won an American Institute of Aeronautics and Astronautics Space Operations and Support Award for the successful recovery effort.

On the other hand, the role Mission Operations played in causing the RND1 failure cannot be ignored. Furthermore, Mission Operations was less than successful in protecting and acquiring spacecraft data needed to support a diagnosis: certain data that might have been valuable to the diagnosis were lost because Mission Operations overwrote spacecraft command memories, failed to downlink processor memory images, or made other operational errors. There was no plan or procedure in place for recovering from an anomalous event, despite several previous entries into safe modes during cruise.

- Although detailed lower-level procedures exist, there is a general lack of top-level procedures, and of adherence to those in place. For example, there is no established and enforced policy of having all critical command loads fully simulated. The requirement that the system engineer sign off all critical command scripts is not enforced.
- Configuration management is inadequate. For example, the version of the burn-abort command script used during the DSM *did* contain the critical command to return the G&C to reaction wheel attitude control; when or why this command script was changed is unknown. As the NAG simulations began it was discovered that there were *two* versions of FC code version 1.11. Flight code was stored on a network server in an uncontrolled environment. The tool Mission Operations used to unpack data structures had an error that gave wrong values for certain data structures (the tool had not been fixed, despite having caused a problem in the past). That these errors were not discovered until the NARB investigation suggests that data structure values were not routinely checked, although the NARB finds them to be critical for correct operation of the spacecraft.
- As NEAR entered its 3-year cruise phase, the role of the system engineer declined relative to that of the mission managers. The comments, suggestions, and requests of the system engineer in operational matters carried no particular weight and were often ignored. The former NEAR software lead engineer, now in another department, is not a reviewer of even critical operations. These two individuals are the principal repository of knowledge of NEAR's overall system design. Not having them fully involved in the decisions and reviews reduced the chances of eliminating errors during the prerendezvous activities.
- The conservative "belt and suspenders" mind-set exhibited by the engineers toward flight hardware prior to launch is missing from Mission Operations. Planning does not sufficiently consider how failures or errors would affect the operations. The spacecraft is not always placed in the best possible state to recover from a failed operation. The lowest-risk approach to each operation is not always used, even when there is little or no impact to using that approach. Some of the risk-reduction practices that were established for critical operations prior to launch and were used during early operations had simply been abandoned by the time of RND1.
- Problem/Failure reporting, although generally well-disciplined prior to launch, is not taken as seriously after launch.

Recommendation 4a: Better configuration management is needed in Mission Operations, including configuration management of the brassboard simulator, command scripts, data structures, autonomy rules, ground software tools, and other critical items currently outside the Configuration Control Board's (CCB's) purview. At a minimum, the CCB should have cognizance over changes to autonomy rules and certain critical data structures. An individual responsible for maintaining configuration-controlled copies of the flight software and important ground software should be identified (a software librarian function).

Recommendation 4b: NEAR operations require technical oversight by a knowledgeable system engineer. System knowledge and a system viewpoint beyond the level of expertise of the Mission Operations staff must be incorporated into future NEAR operations planning and review.

Recommendation 4c: Protecting diagnostic data must be made a priority. This requires written, tested post-recovery procedures, and staff trained in their use, to assure that valuable spacecraft data are not lost.

Recommendation 4d: Establish and enforce procedures to assure that the spacecraft will be prepared for recovery from failed operations. This includes setting command loss timers and resetting important data capture buffers in the command processor and elsewhere before critical maneuvers. Be fully prepared to capture important diagnostic data if something goes wrong.

Recommendation 4e: Implement a plan to routinely collect and possibly analyze all the G&C telemetry formats of the G&C system, not just the ones requested by G&C engineers. This will assure adequate data for analysis of anomalies.

Recommendation 4f: The post-launch Problem/Failure Reporting system must be adhered to with no exceptions, to assure that important lessons are not lost and that dangerous trends are spotted in time.

Recommendation 4g: A near-current orbit should always reside in EEPROM onboard the spacecraft. As the February 1999 FC1 reboot event demonstrated, failure to do so can cause the spacecraft to drop unnecessarily from Earth-safe to Sun-safe mode.

Recommendation 4h: A general reemphasis on Mission Operations process and quality control is needed.

Finding 5. The NEAR simulation environment is deficient in two respects. First, the most complete simulation of critical G&C actions takes place on the stand-alone, software-only simulator. This simulator is not tied directly to the scripts executed by Mission Operations. The parameters and commands that are in effect on the spacecraft during an operation cannot be guaranteed to be those used during the pre-operation simulation on this simulator. Therefore the simulation results may imply nothing about the outcome of a given operation.

Second, the NEAR brassboard simulator, where the actual effect of Mission Operations scripts can be determined, is inadequate. It was assembled late in the program from breadboards, engineering models, and ground support equipment that would otherwise have been dispersed. Because Mission Operations never accepted "ownership" of the brassboard simulator, its documentation was skimpy, and its configuration management was poor. Its hardware and dynamics do not always faithfully emulate the spacecraft, and its operation is frustrating and difficult. Mission Operations and G&C must share the brassboard equipment, requiring frequent reconfiguration. A high degree of skill and knowledge is required to configure and initialize the brassboard simulator correctly; without this careful setup, "strange" results may be obtained. Strange results were often ignored because more often than not they would turn out, after lengthy analysis, to be caused by the brassboard itself. The brassboard's user-unfriendliness discouraged Mission Operations from running many cases (such as the abort cases prior to RND1). Still, the brassboard simulator remains the best and only tool for completely debugging critical operations. Its use prior to RND1 would have uncovered the burn-abort command script error.

Recommendation 5a: Guidelines must be developed and adhered to for deciding which scripts must be simulated on the brassboard, and which test cases must be run (for example, which failure mode cases, if any). All strange results for these cases must be analyzed, even if most turn out to be brassboard artifacts.

Recommendation 5b: Procedures and scripts for configuring and initializing the brassboard simulator should be developed to minimize the number of brassboard runs that give strange results, make the most efficient use of the brassboard itself, and encourage its confident use by the Mission Operations team.

Recommendation 5c: The "simulation environment" is of critical importance for reviewing and checking command uploads and diagnosing anomalies. It needs to be considered early in a program, treated (and funded) as seriously as any other subsystem, and then maintained as a mission-level resource.

Recommendation 5d: Careful consideration needs to be given to which parts of the simulator are "hardware in the loop" and which are emulated by software. Software emulations can uncover problems only in those areas anticipated by the software designers, which can limit the simulator's utility in anomaly investigations.

Recommendation 5e: Subsystems that participate in autonomous safing must be emulated with the highest fidelity.

Recommendation 5f: Consideration should be given to creating a tool to tie the software-only simulation to the actual command scripts used by Mission Operations. This tool could be of value to both NEAR and TIMED.

Finding 6. The gyros are critical to the G&C system operation and performance. At recovery, all three gyro axes were found in Whole Angle Mode (WAM) with an acceptably low noise amplitude. There is no evidence that they entered WAM by any mechanism other than high body rate. However, we were unable to test the gyros for EMI susceptibility or for rate-dependence of the WAM output noise amplitude. An EMI test had been run during acceptance testing of the gyro, but the test data could not be located by either APL or Litton. We could not, therefore, rule out the *possibility* that the gyros are sensitive to EMI in the presence of low bus voltage.

Recommendation 6a: Future programs should maintain closer oversight of vendors during acceptance testing, carefully specify test deliverables at contract initiation, and maintain better test data reporting and archiving.

Recommendation 6b: Future programs should consider purchasing an engineering model gyro along with the flight unit and retaining it postlaunch. That would have been very helpful to this investigation (as well as possibly useful in the brassboard simulator). It is particularly important when the unit is the first of its type to fly, as the hemispherical resonant gyro was for NEAR.

Finding 7. The design of the G&C software for both attitude control and for momentum management ignores the dynamics of onboard fluids ("slosh" effects). With NEAR's oxidizer tanks loaded approximately half full at RND1, and with no internal propellant management device, they were nearly at a worst case for slosh. Realistic slosh modeling is particularly difficult, and numerical instabilities can often mask or mimic physical effects. The spherical pendulum model used for NEAR is more or less standard in the industry. With this model, simulations indicate that slosh may have contributed to some of the inefficient or apparently ineffective momentum dumping, particularly after the gyros entered WAM. However, to date we have been unable to demonstrate that control–slosh interaction can produce anything like the body rates necessary to induce WAM using any physically defensible set of slosh model parameters tried. If slosh played a major role in the post-RND1 difficulties it would be fortunate, for as NEAR approaches Eros only 6 kg of oxidizer remain in the tanks, precluding any repeat of oxidizer slosh effects.

Recommendation 7a: APL's MESSENGER spacecraft propulsion architecture will be similar to NEAR's: bipropellant, no oxidizer tank propellant management device, and the requirement for settling burns. It will have a somewhat higher initial liquid mass fraction, and the bipropellant thrust will be higher. APL must carefully consider slosh dynamics in the MESSENGER control system design.

Recommendation 7b: A better slosh model is needed for future missions—the NEAR model is too complex, too difficult to code, has numerical problems, and provides little or no physical insight.

Finding 8. Much of the difficulty in diagnosing the RND1 anomaly arose because all of the solid-state recorder data were lost when the recorder was powered off during the LVS event. These recorder data contained detailed thruster firing, G&C sensor, spacecraft attitude and rate, and G&C state information that would have made this investigation immensely easier. Although extremely unfortunate for the NARB, the decision to shed the solid-state recorder was necessary to achieve power balance under LVS conditions. Some features were included in NEAR's system design to capture data even without the recorder. For example, the onboard processors were programmed to capture in memory autonomy rules and the data that triggered them, commands sent by the command processors, a min–max telemetry summary, and various counters and flags. These features did prove useful in the investigation, although some were rendered less useful by Mission Operations' failure to reset important buffers prior to RND1.

Recommendation 8a: To preserve diagnostic data, solid-state recorders should have nonvolatile memory or be the last load shed in a low-power situation. Future designs could preserve critical diagnostic data in a truly nonvolatile portion of the recorder, in a region backed up by a small independent battery, or in a lower power mode.

Recommendation 8b: Valuable CTP memory space for diagnostic data could be saved by recording just the first autonomy command in a macro, not the entire macro.

Recommendation 8c: The min–max recording method was inadequate for storing some key G&C data because it operated only on bytes, and the G&C telemetry often packed several single or double-bit status flags into a single telemetry byte. Information was difficult to interpret at best and confusing or useless at worst. The G&C also zeroed all its telemetry upon boot, so that the minimum data points recorded for the G&C carried little information. The min–max data summary can be extremely useful and should be extended to address these inadequacies in future designs.

Recommendation 8d: Critical diagnostic data could be preserved in the processor memories to a greater extent than was done on NEAR. With the advances in memory density since NEAR's design freeze, this should be easier for future programs.

Recommendation 8e: The telemetry system should have a special shortened frame for use during emergency mode operations (any entry into safe modes). This frame would include only data needed for resolution of the anomaly (e.g., spacecraft housekeeping and G&C data), eliminating, for example, instrument suite data. This abbreviated telemetry frame should be stored in the solid-state recorder. Autonomy rule housekeeping snapshots should be transferred to the recorder as soon as the snapshot buffer is full.

Finding 9. Having the gyro in Whole Angle Mode with a high noise level while the G&C system is under AIU control is a dangerous combination because of the AIU's limited ability to filter the noisy WAM output. At the noise levels measured in flight, the AIU is completely capable of controlling the spacecraft, as demonstrated by its performance at recovery. However, we do not have records of the noise levels during the RND1 periods of high body rates, and the NARB was unable to measure the dependence of noise amplitude on rate or other conditions that might cause it to increase.

Recommendation 9a: Consideration should be given to reducing the risk from this combination by, for example, automatically re-setting the gyro from WAM when it is safe to do so, or through changes to control loop gain or other parameters while in WAM. It may be possible to implement this by autonomy rule changes only. Of course, any such changes are risky and the probability that high body rates actually exist must be weighed before such a course is followed. If a change is contemplated, it must be carefully reviewed by the CCB, other concerned parties, and most especially, by the system engineer. Designing, analyzing, and testing changes to these rules is one example of why a full-time postlaunch system engineer is needed.

Finding 10. NEAR's multiprocessor architecture (two each of three different processors) is cumbersome, and redundancy complicates it even further. Although the Board did not find that the design contributed to the RND1 abort, it adds to the risk of mis-synchronization and other issues that arise as processors shut down, boot up, and transfer control. Having separate AIU and FC processors was felt to be necessary because no single processor was capable of handling the entire job back when NEAR had to make its selection (the NEAR Preliminary Design Review was April 1994).

Recommendation 10a: For future programs, simplify the controller modes, with particular emphasis on all safe modes. Consider consolidating control authority in a single processor, rather than distributing it across multiple processors.

Finding 11. Mission Operations was staffed at 20 people at the time of RND1, one short of the number planned for operations at Eros. In addition, two people in another group supported mission design and G&C aspects of mission operations. During most of the cruise phase, Mission Operations staffing averaged 7 people, several less than the number originally planned. While staffing shortfalls did not *directly* contribute to the RND1 event, they may have indirectly contributed through the inadequacy of procedures, training, and simulation. A general reemphasis on quality control should not require an increase in staff.

Finding 12. Following launch, NEAR returned $3.6 million to NASA, a rare underrun for this first Discovery program that APL was justifiably proud of. In retrospect, however, it might have been wiser for NASA to have redirected this $3.6 million toward activities to further reduce mission risk. These could have included improving the state of the brassboard, replicating the brassboard so that it didn't have to be shared, regression testing flight software, properly documenting Mission Operations procedures, and so forth.

APPENDIX M: REFERENCES

1. “The NEAR Mission,” *Johns Hopkins APL Technical Digest*, special issue, Vol. 19, No. 2 (1998).

2. P. E. Doms, *NEAR Eros Encounter Mission Operations Review—Final Report,* Report to A. Carro, JPL 900-98.010 (16 Nov 1998).

3. T. B. Coughlin, *Responses to Peter Doms' Report,* Memorandum to A. Carro, JHU/APL TSSD-17899 (14 Dec 1998).

4. D. W. Dunham, R. W. Farquhar, J. V. McAdams, B. G. Williams, J. K. Miller, C. L. Helfrich, P. G. Antreasian, and W. Owen, Jr., “Recovery of NEAR's Mission to Eros,” *50th International Astronautical Congress,* IAF Paper 99-Q.5.05, Amsterdam, The Netherlands (Oct 1999).

5. D. Zaida, *NEAR Whole Angle Model*, Litton internal memorandum to D. Beisecker (25 Feb 1999).

6. F. T. Dodge, *Final Report: Propellent Dynamics and PMD Design for the Near Earth Asteroid Rendezvous (NEAR) Spacecraft*, Southwest Research Institute Project 04-6297 (Apr 1994).

7. S. C. Lee, *How Clementine Really Failed and What NEAR Can Learn,* JHU/APL memorandum S1I-94-092 (26 May 1994).

8. P. Tennant, “NEAR Anomaly Support: Inertial Measurement Unit EDU Test Report,” Litton internal memorandum to R. Tanabe (23 July 1999).

9. S. M. Krimigis, “Appointment of NEAR Anomaly Review Board,” JHU/APL memorandum SDO-10944 (6 Jan 1999).

10. D. F. Persons, “NEAR NARB Structural Analysis Summary,” JHU/APL memorandum SEM-99-1-1536 (18 Feb 1999).

APPENDIX N: LIST OF ACRONYMS

AIU	Attitude interface unit
BTE	Bench test equipment
CCB	Configuration Control Board
C&DH	Command and data handling
CTP	Command and telemetry processor
DSAD	Digital solar aspect detector
DSM	Deep space maneuver
DSN	Deep Space Network
EDAC	Error detection and correction
EEPROM	Electronically erasable read-only memory
EMI	Electromagnetic interference
EST	Eastern Standard Time
FBV	Fuel bleed valve
FC	Flight computer
FPGA	Field-programmable gate array
FT1, FT2	Fuel tanks 1 and 2
FTE	Full-time equivalent
FTR	Force-to-rebalance
FVC	Fine velocity control
G&C	Guidance and control
HRG	Hemispherical resonant gyroscope
IMU	Inertial measurement unit
JPL	Jet Propulsion Laboratory
LVA	Large velocity adjust
LVS	Low voltage sense
MET	Mission elapsed time
NAG	NEAR Anomaly Group
NARB	NEAR Anomaly Review Board
NASTIE	NEAR attitude system test and integration equipment
NEAR	Near Earth Asteroid Rendezvous
Nms	Newton-meter-second
Ns	Newton-second
NTO	Nitrogen tetroxide
OT	Oxidizer tank
RF	Radio frequency
RND1	NEAR rendezvous burn 1
SEU	Single-event upset
SKI	Sun keep-in
SWRI	Southwest Research Institute
TCM	Trajectory correction maneuver
TIMED	Thermosphere-Ionosphere-Mesosphere Energetics and Dynamics
UT, UTC	Universal Time, Coordinated Universal Time
WAM	Whole-angle mode
YGE	Yellow-good-enough