The Million-Mile Rescue

BY THE NASA SAFETY CENTER

The Solar Heliospheric Observatory (SOHO) is a major element of the joint International Solar Terrestrial Program between NASA and the European Space Agency (ESA). Originally a two-year mission to study the sun, from its deep core to the outer corona, and solar winds, the mission was later extended because of its spectacular success. This extension led to software code modifications meant to increase SOHO's operational lifetime. Instead, multiple errors in the new command sequences repeatedly sent the spacecraft into an emergency safe mode. SOHO's attitude progressively destabilized until all communication was lost in the early hours of June 25, 1998.

The mission was designed to maintain an orbit around the First Lagrangian point, the area where the combined and balancing gravity of Earth and the sun would keep SOHO's orbit anchored in the Earth–sun line. Once in this orbit, SOHO's attitude was generally stable and would use spinning reaction wheels controlled by an attitude-control unit (ACU) computer to autonomously adjust for internal or external disturbances. If the wheels reached a spin near their design limit, the ACU would automatically despin the wheels, use thrusters to stabilize attitude, and then reactivate the wheels to resume attitude control. During these maneuvers, the ACU would use one of three gyroscopes (Gyro C) to sense roll.

SOHO's second gyro (Gyro B) was used solely for fault detection—for example, to sense roll rates beyond some predetermined tolerance. If an excessive roll rate was detected, SOHO would enter a safe mode, where it ensured that its solar panels were facing the sun, temporarily suspended the ACU computer, and then awaited the ground commands it needed to restore normal operations. During one such recovery, ground controllers used the third gyro (Gyro A), instead of Gyro C, for roll-rate sensing.

Gyroscope Misconfigurations

Each gyro onboard SOHO was designed to be used only for its specific independent function. All three require periodic calibrations to account for drift bias, which results from mechanical wear, angular changes, or exposure to extreme temperatures. When drift bias occurs, ground engineers uplink the correct coordinates for each gyro to the spacecraft's onboard computer, allowing the spacecraft's attitude-control functions to operate accurately. The same mechanical and thermal wear that causes drift bias eventually makes the gyros non-operational, which became a concern when the SOHO mission was extended.

In February 1997, the flight operations team modified gyro command sequences in an attempt to address this wear issue. A command was written to deactivate, or spin down, Gyro A when not in use, which is any time other than the safe mode. The code was supposed to include a function to respin Gyro A upon entering safe mode, but this function was erroneously omitted in the new sequence.

The modification had been introduced with a mission operations change request in March 1997 but was not used in gyro calibrations until June 24, 1998. Therefore, even though the SOHO spacecraft had entered safe mode four times prior to June 24, the code modifications were not in use and did not affect successful recoveries by ground crews.

A later review revealed that these modifications were never properly documented, communicated, reviewed, or approved by either ESA or NASA. The change request itself was an internal flight-operations document only distributed within the team. The only testing performed was by a NASA computer-based simulator that verified each change separately, but not all together.

The software modifications contained a second critical error. The fault-detection setting on Gyro B was twenty times

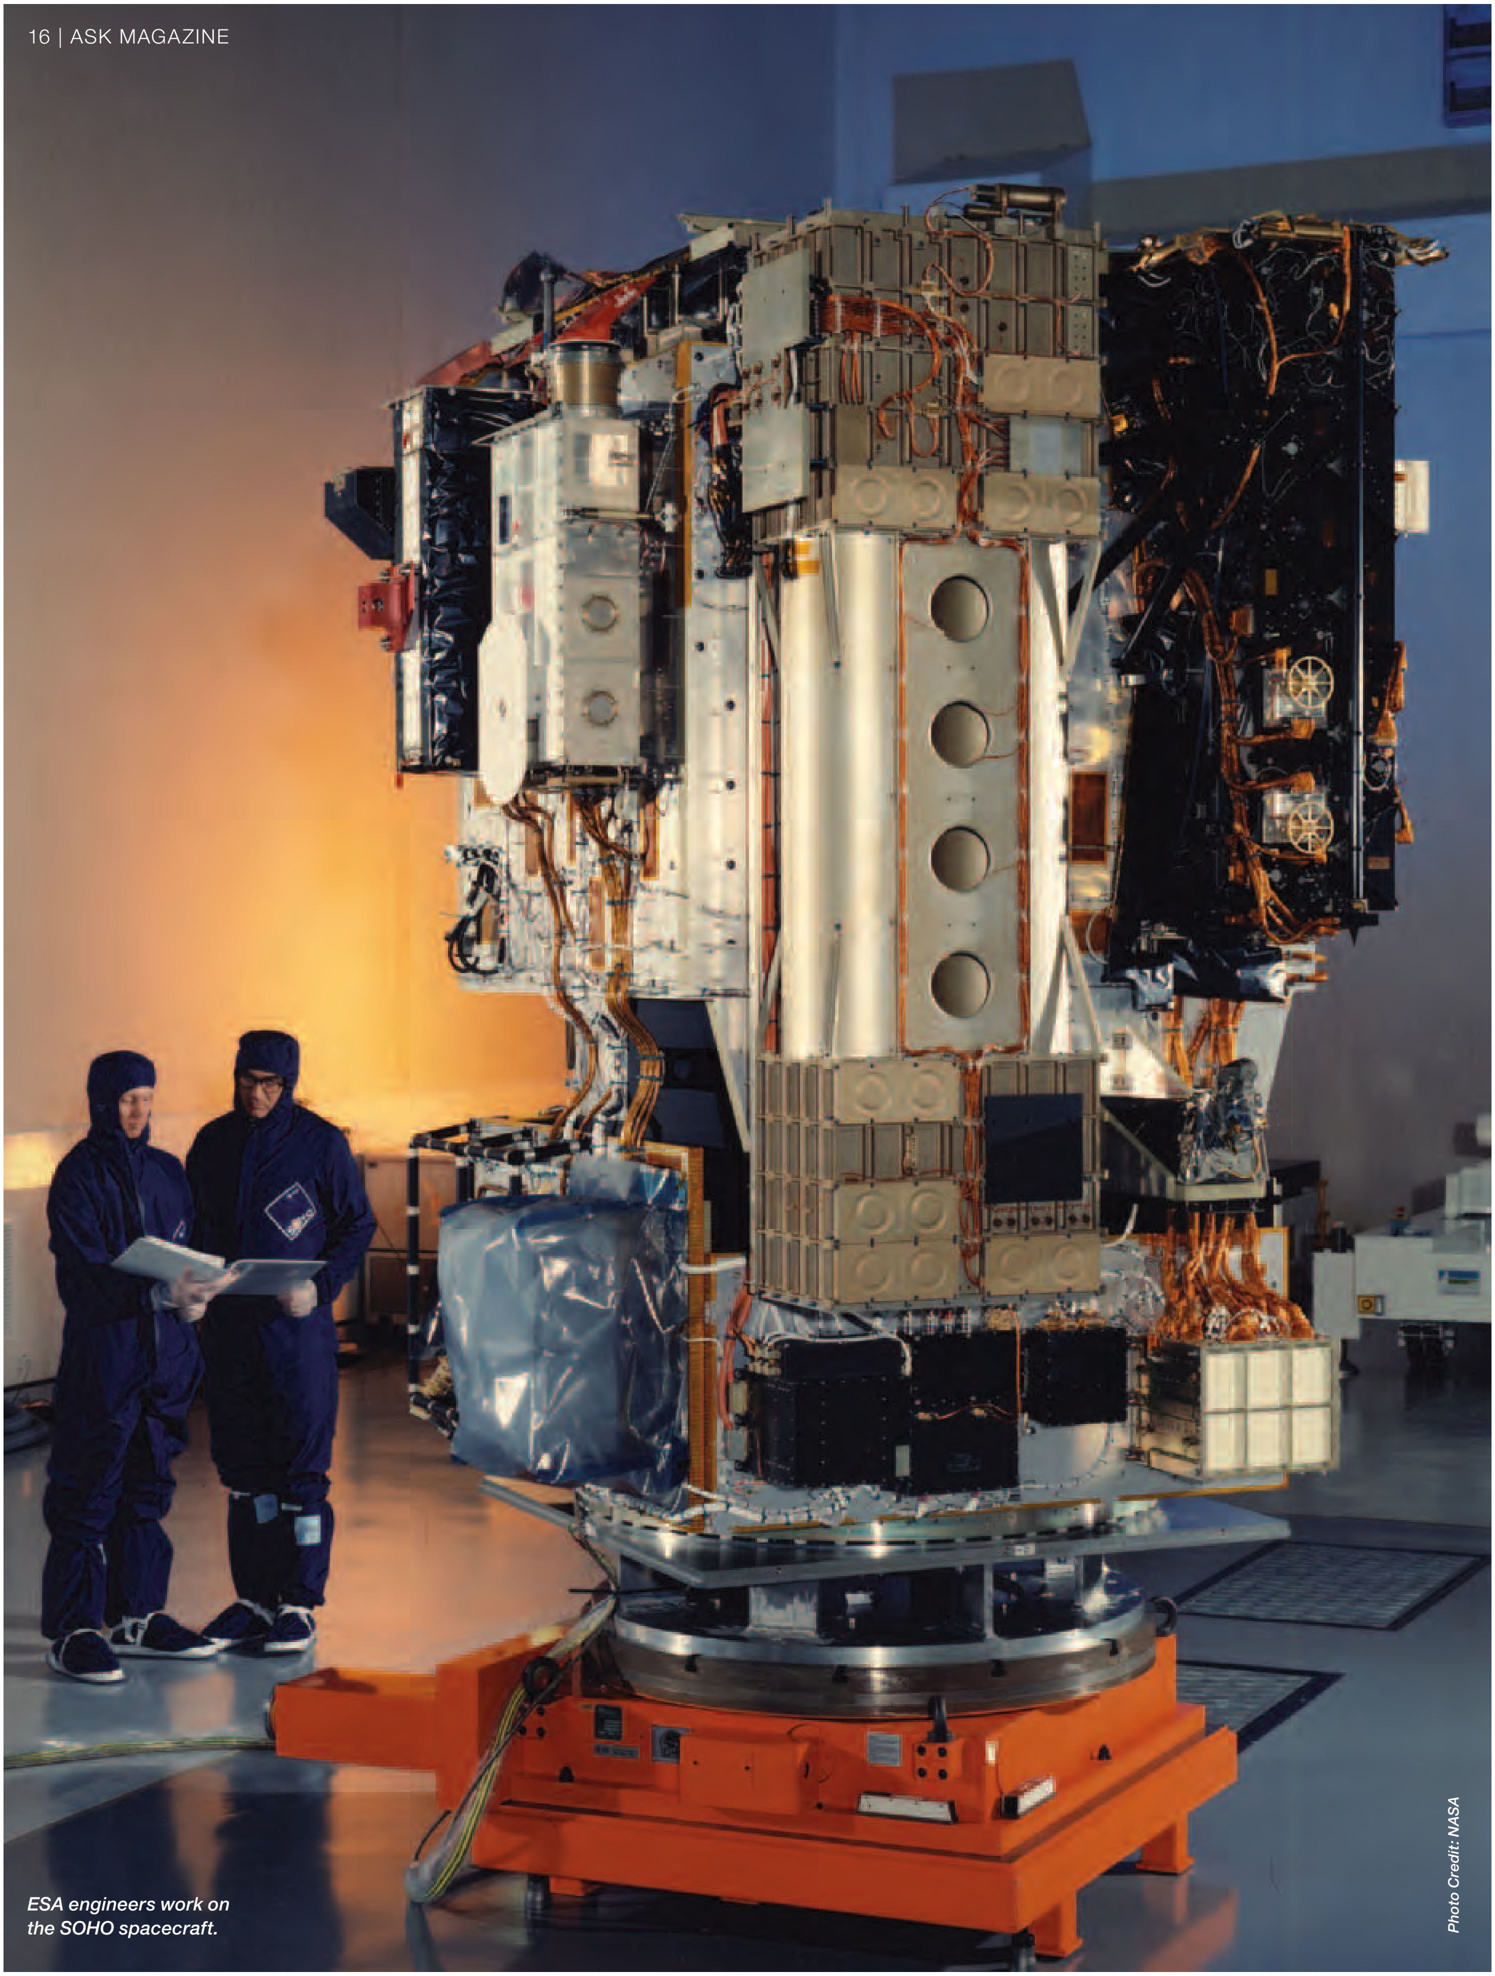

ESA engineers work on the SOHO spacecraft.

Photo Credit: NASA

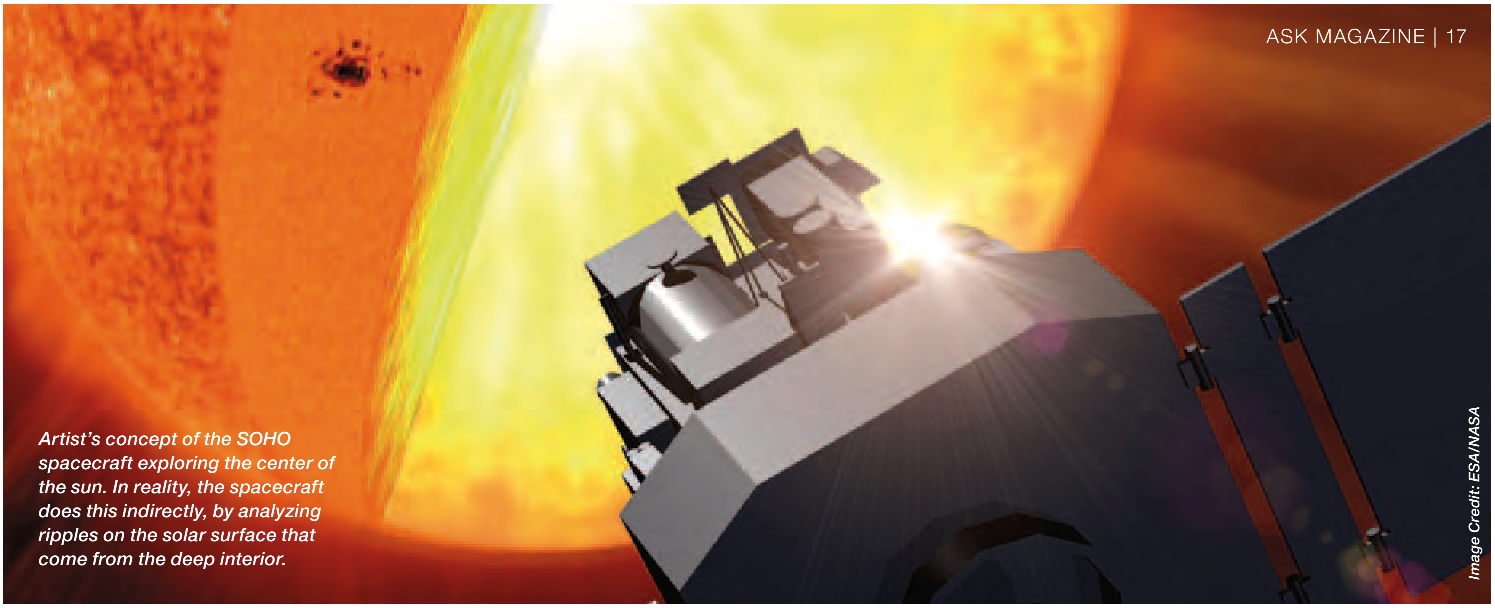

Artist's concept of the SOHO spacecraft exploring the center of the sun. In reality, the spacecraft does this indirectly, by analyzing ripples on the solar surface that come from the deep interior.

Image Credit: ESA/NASA

more sensitive than it should have been. This error triggered a mishap and sent SOHO into its fifth safe mode at 7:16 p.m. on June 24, 1998.

The recovery effort began immediately but was complicated by the aggressive scientific task schedule planned for June 24–29. The core SOHO team was already working on a compressed timeline without the luxury of additional support or contingency time. Ground controllers quickly discovered and corrected the error in Gyro B but did not notice that Gyro A had not reactivated. Shortly thereafter, as a normal part of the recovery sequence, all three gyros were recalibrated, and the ACU computer was activated to make any necessary adjustments using its thrusters. But when the computer attempted to correct for the drift bias on the spun-down Gyro A, it continuously attempted to correct for a perceived (but non-existent) roll-attitude error until the actual roll rate increased so significantly that Gyro B's fault detection accurately triggered another safe mode at 10:35 p.m. Again, recovery efforts began immediately.

Critical Decision Mistake

The safe mode recovery period was designed to give flight operations and engineering teams sufficient time to understand problematic anomalies before taking action. SOHO was programmed to store telemetry prior to any safe mode so it would be available for examination by ground crews. The operations procedures specifically stated that before attempting a recovery, Gyro A should be confirmed to be spinning and the telemetry should be analyzed. The SOHO operations team did not take advantage of this design feature; instead they chose to initiate recovery sequences almost immediately after each safe mode was triggered without checking either Gyro A's spin status or the telemetry data.

The team observed that Gyro B's readings of an excessive roll rate did not agree with Gyro A's nominal roll-rate reading, but the flight operations crew still failed to notice that Gyro A was not spinning. In a quick decision, the flight operations manager incorrectly concluded that it was Gyro B (and not Gyro A) that was faulty.

Gyro B was shut down, which rendered the fault-detection capability inactive. When control was returned to the onboard computer for the recalibration sequence of recovery, roll thruster firing resumed and sun-pointing errors eventually resulted in pitch and yaw thruster firings. This produced unstable spinning of the spacecraft that exceeded allowed limits and triggered another safe mode at 12:38 a.m. on June 25.

SEVERAL FACTORS CONTRIBUTED TO SOHO'S MISHAP: CHANGE CONTROLS WERE LACKING, PROCEDURES WERE NOT FOLLOWED AS WRITTEN, AGGRESSIVE SCHEDULING OVERTASKED THE TEAM, AND NOT ENOUGH STAFF WAS AVAILABLE TO HANDLE THE PLANNED SCIENCE TASKS AND SUBSEQUENT RECOVERY MODES.

The critical software errors in the modified gyro command sequence meant that SOHO's gyros were configured incorrectly and caused the onboard computer to erroneously fire its thrusters until the spacecraft destabilized. This was exacerbated by the decision to shut down a gyro believed to be malfunctioning in favor of a gyro that was actually inactive.

Within minutes, SOHO's attitude diverged beyond control. Power, communications, and telemetry were all lost. By 12:43 a.m., SOHO was officially lost in space.

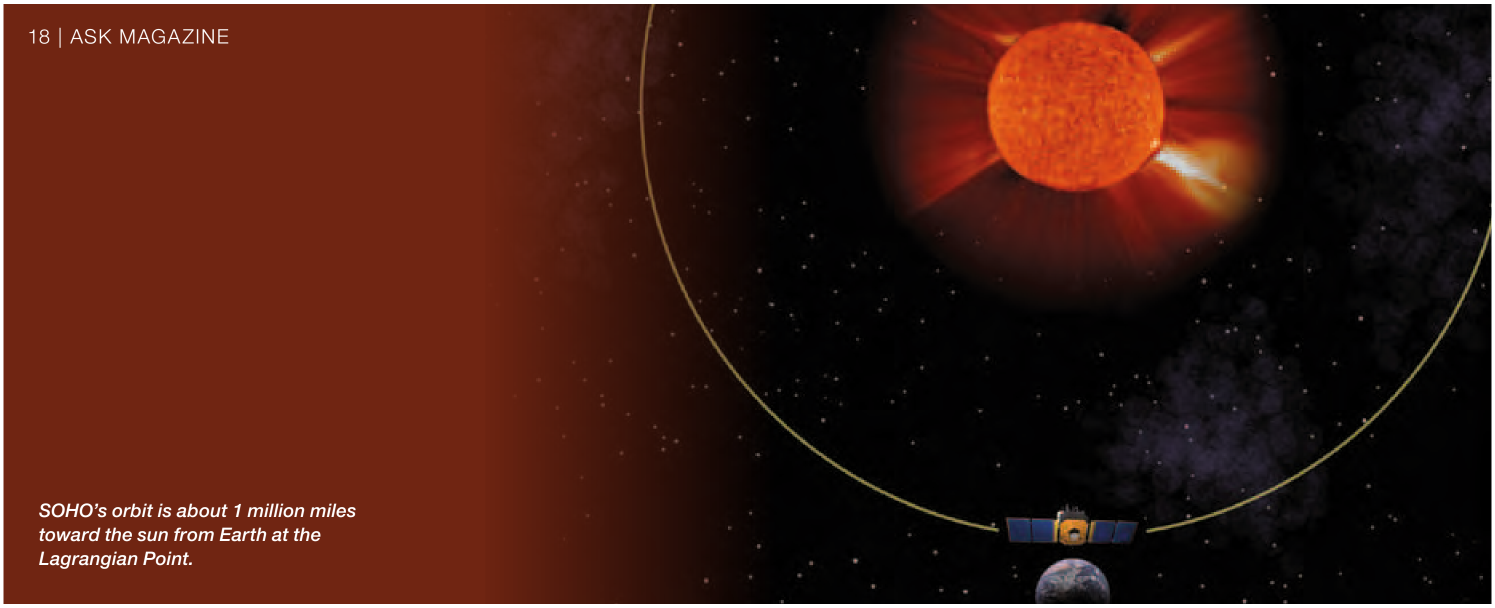

SOHO's orbit is about 1 million miles toward the sun from Earth at the Lagrangian Point.

Image Credit: NASA Goddard Space Flight Center

The Million-Mile Rescue

Within hours, investigation teams at both ESA and NASA had been assembled. On June 28 they convened at Goddard Space Flight Center to begin recovery efforts. Based on the last few minutes of telemetry, simulations predicting possible trajectories for SOHO indicated that the spacecraft would diverge and escape into a solar orbit if it was not recovered by mid-November. By a stroke of good fortune, calculations also indicated that, in roughly ninety days, the spin of the spacecraft would naturally align the solar arrays with the sun for about half a spin period, giving the recovery team an opportunity to regain control over SOHO. On July 23, using the Arecibo radio telescope in Puerto Rico in combination with NASA's Deep Space Network in California, the team was able to locate the spacecraft's radar echoes and confirm both its location and spin rate.

The flight operations team uplinked commands to SOHO for twelve hours a day, searching for any signs of return communication. On August 3, contact was established. Over the next two months, SOHO was progressively restored to normal operating mode. On September 25, about ninety days after contact was initially lost, SOHO was fully operational. Remarkably, all twelve scientific instruments remained in complete working condition despite having been subjected to temperatures from -120°C to 100°C during the three-month ordeal.

Lessons Learned for NASA

Several factors contributed to SOHO's mishap: change controls were lacking, procedures were not followed as written, aggressive scheduling overtasked the team, and not enough staff was available to handle the planned science tasks and subsequent recovery modes. As a result, key engineers were preparing for upcoming science tasks rather than assisting with safe-mode recoveries. Recovery efforts were rushed in order to return the spacecraft to its science operations as quickly as possible. Ironically, the prioritization of science over spacecraft safety contributed to the loss of science operations for three months and risked the total loss of SOHO.

It is important that modifications or updates to procedural scripts on future NASA missions have formal approval before implementation, and the entire script (not just the modification) should be revalidated. Operational timelines should also be planned and validated *before* implementation—not in parallel with implementation—with the proper attention and reserve given to contingency planning and safety. There should be sufficient time for coordinating tests and simulations so they do not conflict with management and operations of real-time, on-orbit events.

The health and safety of a spacecraft are critical in achieving any scientific or operational goals. To keep the spacecraft healthy, the team needs to be healthy. Reassess staffing levels periodically, strengthen staff as needed, and provide the capability for the extra support required by contingency operations. This can be difficult in extended operations on missions that have limited budget flexibility, but it is important. In any case, operations teams must be well trained on the systems they will be required to use and should practice responses to emergency situations. ●

This article is adapted from a NASA safety-awareness training document based on information available in the public domain. The findings, proximate causes, and contributing factors identified in this case study do not necessarily represent those of the agency. Sections of this case study were derived from multiple sources listed under References. Any misrepresentation or improper use of source material is unintentional.

REFERENCES

1. F.C. Vandenbussche, "SOHO's Recovery: An Unprecedented Success Story," *ESA Bulletin 97*, March 1999, sohowww.nascom.nasa.gov/operations/Recovery/vandenbu.pdf
2. SOHO Mission Interruption Joint NASA/ESA Investigation Board, *Final Report*, August 31, 1998, umbra.nascom.nasa.gov/soho/SOHO_final_report
3. NASA Public Lesson Learned 0664, "SOHO Mission Interruption Joint NASA/ESA Investigation Board," December 1, 1999, www.nasa.gov/offices/oce/llis/0664.html
4. Kathryn A. Weiss, Nancy Leveson, Kristina Lundqvist, Nida Farid, and Margaret Stringfellow, "An Analysis of Causation in Aerospace Accidents," June 25, 2001, csel.eng.ohio-state.edu/woods/accident_reports/NASA/leveson_soho.pdf