

FAULT DETECTION AND CORRECTION FOR THE SOLAR DYNAMICS OBSERVATORY ATTITUDE CONTROL SYSTEM

Scott R. Starin, Melissa F. Vess, Thomas M. Kenney, Manuel D. Maldonado, and Wendy M. Morgenstern
Goddard Space Flight Center, Greenbelt, MD 20771

The Solar Dynamics Observatory is an Explorer-class mission that will launch in December, 2008. The spacecraft will operate in a geosynchronous orbit, sending data 24 hours a day to a devoted ground station in White Sands, New Mexico. It will carry a suite of instruments designed to observe the Sun in multiple wavelengths at unprecedented resolution. The Atmospheric Imaging Assembly includes four telescopes with focal plane CCDs that can image the full solar disk in four different visible wavelengths. The Extreme-ultraviolet Variability Experiment will collect time-correlated data on the activity of the Sun's corona. The Helioseismic and Magnetic Imager will enable study of pressure waves moving through the body of the Sun.

The attitude control system on Solar Dynamics Observatory is responsible for four main phases of activity. The physical safety of the spacecraft after separation must be guaranteed. Fine attitude determination and control must be sufficient for instrument calibration maneuvers. The mission science mode requires 2-arcsecond control according to error signals provided by guide telescopes on the Atmospheric Imaging Assembly, one of the three instruments to be carried. Lastly, accurate execution of linear and angular momentum changes to the spacecraft must be provided for momentum management and orbit maintenance.

In this paper, single-fault tolerant fault detection and correction of the Solar Dynamics Observatory attitude control system is described. The attitude control hardware suite for the mission is catalogued, with special attention to redundancy at the hardware level. Four reaction wheels are used where any three are satisfactory. Four pairs of redundant thrusters are employed for orbit change maneuvers and momentum management. Three two-axis gyroscopes provide full redundancy for rate sensing. A digital Sun sensor and two autonomous star trackers provide two-out-of-three redundancy for fine attitude determination. The use of software to maximize chances of recovery from any hardware or software fault is detailed. A generic fault detection and correction software structure is used, allowing additions, deletions, and adjustments to fault detection and correction rules. This software structure is fed by in-line fault tests that are also able to take appropriate actions to avoid corruption of the data stream.

INTRODUCTION

The Solar Dynamics Observatory (SDO) mission is the first Space Weather Research Network mission, part of NASA's Living With a Star program.¹ This program seeks to understand the changing Sun and its effects on the Solar System, life, and society. To this end, the SDO spacecraft will carry three Sun-observing instruments to geosynchronous orbit: Helioseismic and Magnetic Imager (HMI), led by Stanford University; Atmospheric Imaging Assembly (AIA), led by Lockheed Martin Space and Astrophysics Laboratory; and Extreme-ultraviolet Variability Experiment (EVE), led by the University of Colorado. The HMI will enable study of pressure

waves moving through the body of the Sun. The AIA includes four telescopes with focal plane CCDs that can image the full solar disk in four different visible wavelengths. The EVE will collect time-correlated data on the activity of the Sun's corona. Links describing the instruments in more detail may be found through the SDO web site.²

SDO is scheduled to launch in December, 2008. Its basic mission goals are to observe the Sun for a very high percentage of the spacecraft lifetime with long stretches of uninterrupted observations and to transmit these data continuously at a high data rate to a dedicated ground station. The SDO mission lifetime is 10 years, with a minimum mission success lifetime of 5 years. These science collection and longevity goals guided the design of the spacecraft bus that will carry and service the three-instrument payload. At the time of this publication, the SDO spacecraft bus is well into the integration and testing phase at the NASA Goddard Space Flight Center (GSFC). A three-axis stabilized attitude control system (ACS) is needed both to point the instruments at the Sun accurately and to keep the roll about the Sun vector correctly positioned relative to the solar rotational axis. This paper will describe the SDO ACS and the methods by which it can guarantee science performance throughout the mission despite any one of a number of possible failure scenarios.

HARDWARE COMPLEMENT

The SDO ACS has been designed to tolerate any single hardware fault and yet still retain capability to meet all requirements for science data quality. To this end, ACS sensors, actuators and computational capabilities have been selected and arranged both for performance and maximal redundancy. SDO ACS failure detection and correction depends to a large extent on hardware redundancies, so the hardware complement and its redundancies will be described in detail. Figure 1 shows a mechanical drawing of the SDO. Please see Reference 3 for more information on hardware placement.

ACS Sensors

The SDO sensor suite comprises sixteen Adcole coarse Sun sensors (CSS), one Adcole digital Sun sensor (DSS), two Galileo Avionica quaternion-output star trackers (ST), and three Kearfott Two-Axis Rate Assemblies (TARA). The CSSs are the only attitude sensors required in the most basic Sun pointing mode. The sixteen CSSs are divided into two independent sets of eight sensors each, and each set of eight can provide an adequate Sun vector with any seven sensors being functional.

For fine attitude determination, an on-board Kalman filter can provide adequate attitude knowledge with input from any two of the three fine pointing units—DSS, ST1, and ST2. To avoid multiple blockages, therefore, the STs are mounted nearly perpendicular to the SDO Sun-pointing axis (X axis), and far enough apart from each other that the Earth and Moon do not block both at the same time throughout the science collection phase of the mission. The TARAs are arranged with their insensitive axes orthogonal to one another, such that sensitive axes from two units are aligned with each of the three body axes of the Observatory. Thus, the loss of any single TARA still allows full three-axis rate information to be obtained.

In addition to the ACS suite described here, the ACS also makes extensive use of the guide telescopes (GT) mounted as part of the AIA instrumentation. Because the accuracy of images taken by SDO will be unprecedented, the ACS is expected to guide the spacecraft trusting the processed GT data as the best available knowledge of the Sun center. There are four GTs, with one mounted to each of the four science telescopes; the ACS only needs accurate information from one of the four GTs, identified as the controlling guide telescope (CGT), to perform its science control duties. Each GT is capable of providing attitude information relative to the Sun vector accurate to about 2 arcseconds within about 90 arcseconds of its centerline. This portion of the field of view (FOV) is called the linear range, and is required for accurate science data collection. Outside of that range, polarity is maintained as long as the center of the Sun remains within the 0.5-deg FOV, which allows sunlight to fall on one or more photodiodes. The controller that uses the GT can acquire the linear range given proper initialization in the full FOV, so the full FOV is called the acquisition range.




Figure 1: The Solar Dynamics Observatory

ACS Actuators

SDO guidance functions will be actuated by four Goodrich 70-Nms reaction wheel assemblies (RWA) and eight Ampac 5-lbf attitude control engines, or thrusters. The RWAs are arranged in a pyramidal structure, so that any set of three provides full three-axis control capability. The ACS thrusters are grouped into four pairs of thrusters, with one thruster of each pair linked to fuel and oxidizer by independent manifolds. In this way, the catastrophic failure of any one thruster can only require the closing of one manifold, leaving the other set of four capable of performing all necessary ACS tasks.

In addition to attitude control activities, the ACS is also responsible for the guidance and control of the two high-gain antennas (HGA) because it already has access to the necessary navigational information. Each antenna consists of a dish mounted on an elevation gimbal, with that mounted on an azimuth gimbal. Either antenna may fail, either from losing a gimbal or from other electrical failure, and the mission can still be completed. The greatest danger posed by the HGAs is the irradiation of the spacecraft itself; HGA FDC provides additional protection from that event.

Electronics and Microprocessors

A fully-capable copy of the command and data handling system of the SDO spacecraft bus resides on each of two independent microprocessors, here referred to as the main processors (MP). Only one of these MPs is in control of the spacecraft at any time. Each of these MPs also operates an independent copy of the attitude control task (ACT) and the on-board ephemeris. Most flight software tasks, including the ACT, operate on a 200-msec cycle, except that the ephemeris operates on a 1-sec cycle. The ACT samples sensor data and issues actuator commands at this 5-Hz rate, and it is this task that, if all goes perfectly nominally throughout the mission, will maintain attitude control of the spacecraft. There are also separate microprocessors that reside in each of two independent but cross-strapped attitude control electronics (ACE) boxes. These ACEs are always powered, but only one can be in control at any time. The ACE that is in control routes data from the CSSs, RWAs, TARAs, and various pressure and temperature sensors to the MPs. It also accepts and validates actuator commands to the RWAs and the propulsion system's valves and thrusters from the MP in control. The ACEs also run their software tasks on a 200-msec cycle. Similarly, there are two independent gimbal control electronics (GCE) boxes and two independent power service electronics (PSE) boxes.

If there is a disruption of ACT control over the ACE in control, that ACE will cease passing through ACT actuator commands and will instead begin issuing commands of its own to the reaction wheels. The control system running on the ACE is simplified, being dependent only on the CSSs and the RWA tachometer readings (during eclipse, if TARA signals are available, they are used to null rates). If the primary ACE itself is disrupted, such as from a power down due to a single-event upset (SEU), the other ACE will detect this state and assume control. The details of the algorithms for arbitration between the two ACE will be discussed in a later section. From the state of either ACE issuing safety commands, the Observatory may be recovered to nominal operations once one or the other MP is available for commanding.

ATTITUDE DETERMINATION AND CONTROL MODES

The ACS has four RWA-actuated modes and two thruster-actuated modes. More details about the ACS in general and the control modes in particular can be found in Refs. [3-6]. As discussed in the previous section, one RWA-actuated mode resides on the ACE microprocessors; this mode is called Safehold. The other five modes reside in the ACT. Sun Acquisition (SunAcq) performs an attitude function similar to Safehold, in that it simply maintains a power-positive, safe attitude with respect to the Sun using CSS signals. It differs from Safehold in that TARA signals are used for angular rate information at all times.

For all other modes, attitude determination (AD) is performed with some combination of the fine attitude sensors and propagation of TARA-derived rate information. An attitude solution may be initialized either by accepting a valid ST quaternion (nominal) or by uploading an estimate by ground command (available for testing and contingency). Once a solution is available, it may be simply propagated using rate sensors, as is always done in the thruster based modes, or it may be updated either using one preferred ST or by ground override command. The most accurate solution is obtained by combining all available fine attitude data from the two STs, the DSS, and the TARAs using a Kalman filter.

Whatever AD method is selected in the software, Inertial mode uses the solution for attitude error calculation against the target attitude in all three axes. Inertial mode has two sub-modes that differ only in the target calculation—one tracks a Sun-referenced target quaternion, using the on-board ephemeris to predict the appropriate inertially referenced quaternion for the Sun-referenced state, and the other maintains a commanded absolute inertially referenced quaternion. Science mode, during which most science data are collected, uses one of the specialized GTs to point a commanded science reference boresight (SRB) accurately at the Sun. The roll error about that SRB is calculated using the same methods as Inertial mode, except that the target is always Sun-referenced.

The thruster modes are called DeltaH and DeltaV. DeltaH is used to manage system angular momentum. With no magnetic torquers to gradually dump momentum, the thrusters must be used occasionally to remove momentum. To maximize time between uses of DeltaH, the mode allows a non-zero angular momentum to be placed into the body, which can be set opposite any predicted angular momentum change over a period of approximately four weeks. The attitude target for DeltaH is simply the attitude estimate at mode entry. DeltaV is used for changing or maintaining orbit parameters. It uses an absolute, inertially referenced target similar to Inertial's absolute targeting, and that target may be updated by command during a DeltaV maneuver.

Some transitions between modes are not allowed. By placing the in-control ACE into Safehold mode, the ACS mode running on the MP is ignored, so Safehold may be reached from any MP mode. Any MP mode may transition to SunAcq or to Inertial, including self-transitions. Science mode is the only other mode that may self-transition, and it may also be entered autonomously from Inertial mode when the Sun is in the field-of-view of the controlling guide telescope. DeltaH may be entered from SunAcq or Inertial mode. However, Science and DeltaV may only be entered from Inertial mode, with Science accessible only when Sun-referenced targeting is active and DeltaV accessible only when absolute targeting is active. These restrictions avoid large attitude changes occurring due only to misunderstandings of the two targeting sub-modes in Inertial. It is worth noting for this topic that thrusters are always disabled upon exiting DeltaH or DeltaV modes. Figure 2 is a diagram of the SDO ACS control modes and allowed transitions.




Figure 2: Allowed mode transitions for SDO ACS software

FAULT DETECTION AND CORRECTION

As has been pointed out in the previous sections, SDO is completely hardware redundant to allow the spacecraft to be single fault tolerant. SDO has a long mission life, and one of the primary goals of the SDO mission includes collecting long periods of uninterrupted science data. So, along with the redundant hardware, the spacecraft has also implemented on-board monitoring of the system such that if a fault is detected, corrective action can autonomously be taken to utilize the redundant hardware and place the spacecraft into a safe configuration until the ground can intervene to correct the fault. In some instances, software reconfiguration or operational changes may be needed to regain full capability. The failure correction philosophy is to configure the system for the possibility of any single fault, with flexibility to protect against subsequent faults after the first fault is understood. The on-board FDC is handled in two separate ways that work together to protect the system in an overall manner. The first way consists of checks and corrective actions taken within the flight software code itself (i.e. within the ACS sensor data processing or within the on-board ephemeris task). The second way is to monitor telemetry and initiate corrective action using separate, devoted software tasks.

In-Line Checking

From the ACS perspective, the first line of defense against faults resides within the code of the ACT, on-board ephemeris, and ACE. These in-line responses operate every time the code executes (5 Hz cycle for ACT and ACE, and 1 Hz cycle for ephemeris). On a cycle-by-cycle basis, they protect the system by preventing faults in one part of the ACS from propagating throughout the rest of the system. At their most basic, these checks protect the system from problems such as divide-by-zero errors or loss of accuracy in quaternions due to processor truncation. In addition, they also protect against some faults that could potentially be hardware failures. A simple example of in-line protection against hardware problems is the checks on the analog-to-digital converters (ADC) of the CSS signals. If the CSS ADC fails, flags are set that indicate that the CSS data for the current cycle are invalid and, therefore, the last valid CSS data are used in the SunAcq controller until new valid data are received.

As a more complex example, consider the ST data processing. As was mentioned earlier, the SDO STs are quaternion-out trackers. A portion of the data received from the STs includes a logical bit that indicates whether or not the ST identified and verified a valid quaternion. Within the ST data processing in the ACT, the input data are first checked that the ST-provided validity bit indicates a valid, verified quaternion. Then, the quaternions are checked to ensure that their norm equals one within a tolerance and that the difference between the current and previous quaternions does not exceed the expected movement of the spacecraft. If any of these checks fails, the ST data are considered invalid and will not be used in the AD solution for that cycle. In addition, the ST processing also checks ephemeris information against attitude to see if the Moon, Earth, or Sun lies within exclusion cones of the ST boresights. The STs function such that even if one of the bright bodies is within the exclusion zone, the ST could still provide a solution, but that solution may not be accurate enough to meet the SDO attitude determination requirements. So, if the ST processing predicts that the ST boresight is too close to a bright body, the in-line response will flag the ST quaternion as low quality, and so that quaternion will not be included in the AD solution for that cycle. Both the validity checks and quality checks protect Science and Inertial modes from using an AD solution corrupted by flawed ST data.

FDC Software Tasks

The two modules of software that implement the Failure, Detection and Correction functions on SDO are the Telemetry and Statistics Monitor (TSM) and Stored Command Processor (SC). Both of these modules reside in the MP and have considerable flight heritage on in-house GSFC missions. The TSM module performs monitoring of limits and upon detection of threshold crossings will send a command to the SC module to perform the corrective action.

The TSM module receives all spacecraft telemetry packets, and the receipt of telemetry packets drives the TSM module. Telemetry packets can originate from other spacecraft software modules or other spacecraft subsystems. Each TSM monitor point consists of a telemetry packet application identifier (APID), packet offset and telemetry monitor size and mask, 4 thresholds, any number of which may have their responses enabled, and an enable/disable flag for the TSM itself. Furthermore, each threshold consists of an enable/disable flag, a comparison operator (<, >, =, etc.), a threshold limit, a persistence, and the relative time sequence (RTS) that contains the corrective action should the TSM activate.

The SC module consists of an absolute time processor, used mostly by the mission operations team, and a relative time processor. The relative time processor stores RTSs and can initiate them via one command that calls the number of the desired RTS. An RTS consists of a series of spacecraft commands with time delays between them. The time delays may be any integer number of seconds. SDO has allocated sufficient memory to have 512 RTSs, each storing up to 300 bytes of commands. Command sizes are variable and follow CCSDS format; thus, the minimum command length is 8 bytes. Each command is also preceded by a 2-byte time delay in the RTS. Thus SDO RTSs can contain a maximum of 30 commands each. RTSs can be chained by including multiple RTS calls in another RTS if a command sequence does not fit in 300 bytes.

FDC Guidelines

Autonomous fault detection and correction for the SDO is designed to allow the spacecraft to operate safely without ground contacts or intervention for a minimum of 12 hours. The spacecraft places itself into a safe configuration in the event of an anomaly to allow the flight operations team (FOT) the opportunity to take full advantage of system redundancy and robustness. Areas where potential mission loss would be mitigated by autonomy include: loss of power (e.g. power subsystem, battery, solar array electronics, load shed); loss of attitude control (e.g. sensor failure, actuator failure); loss of communication with internal components or with the ground (1553 bus, commanding); component over or under temperature (e.g. propulsion, HGA system); and loss of control of specific hardware (switched power services, HGA gimbals, isolation valves).

This self-protection capability meets several needs. In the event of difficulties in the ground telemetry collection and commanding system, the spacecraft can remain safe despite being out of contact for an extended period. Another important goal of on-board autonomy is to simplify ground operations by reducing the ground requirement to quickly react to unexpected situations. On-board FDC monitors are used to guarantee a power positive and stable attitude configuration. The FOT still needs to react but with less urgency, and this translates into economic efficiency. Finally, monitoring aspects of on-board FDC reduce FOT workload during mission critical events by tracking and summarizing key health and safety telemetry parameters. This in turn reduces the need to quickly analyze data recorder dumps, as the spacecraft FDC software identifies where key subsystem parameters may be out of expected ranges.

Fault detection and correction is bound to be a complex process in a closed-loop, multi-input, multi-output system like the SDO ACS. In designing the SDO ACS FDC, the goal was not to find any kind of true optimization between competing desires for flexibility, comprehensiveness, and simplicity, but instead to establish guiding principles that would minimize dangers while maximizing the likelihood of mission success. The guiding principles can be summarized as follows:

1. Autonomous monitoring may be done for any factor of interest, but responses are limited to health and safety concerns. Each autonomous response is taken to address some potential for hardware damage or observatory loss in the absence of corrective action, and conversely, each identified single-failure potential for loss has some response which should prevent loss from occurring.

2. The list of possible autonomous responses to contingencies should be as small as possible, and the responses should interfere with one another as little as possible. The ultimate goal of SDO FDC is to facilitate rapid recovery by a competent FOT, not to recover autonomously to a fully operational condition.
3. Failures on complex spacecraft often result in a cascade of hazardous conditions. A concomitant cascade of FDC responses would be unpredictable and may do more harm than good.

The application of these three principles will be pointed out throughout the remainder of this paper. The third principle bears some deeper discussion here. Establishing beforehand responses to an unknown cascade of failures is inherently difficult. On one hand, immediate response to a localized failure could potentially prevent the cascade. However, on the other hand, if the response is in a race with a better response for that failure, or is not adequate to the severity of the failure, the cascade may be worsened. Or, other responses that might have been useful may be less effective after other reconfigurations have occurred. The worst situation from a design perspective would be for autonomous responses to push the failure along, making a system unrecoverable when the initial failure need not have received an immediate response.

To avoid race conditions and detrimental cascades of failures, FDC monitors and associated responses are divided into four categories of widening influence: specific hardware faults, ACE faults, mode-specific performance faults, and system-level faults. Hardware-specific responses focus mainly on the hardware that has been detected to have failed, and attempt to establish a safe configuration with a minimum of change. ACE fault responses are mostly independent from the other FDC as they focus on selecting which of the two ACEs is more trustworthy. Mode-specific responses always issue commands to demote to a safer control mode, which are responses that do not interfere with other FDC. The most serious or dangerous system-wide conditions require the shortest persistence periods before activation of the appropriate RTS. They also encompass the possible responses of more specialized responses to mode-specific control performance or hardware-specific faults. The following sections discuss specific checks in these four categories and how they interact.

HARDWARE-SPECIFIC FDC

Most hardware-specific failure detection and correction is done in the ACT software, on a cycle-by-cycle basis. This in-line FDC consists of checking incoming sensor, actuator, and processor data against expectations, and accepting or rejecting that data as useful for further purposes, such as control calculations or attitude determination. Table 1 shows many of the hardware checks that are performed and their responses. If a condition merits an in-line response, no persistence is given for that response, and the response is numbered 0. If a response is effected by RTS, the persistence required by the devoted TSM is given in seconds or minutes, and each TSM-initiated response is numbered starting with 1.

Some responses are common to several types of data processing. When new data are not received in a cycle, or when validity limits are exceeded, a frequent response is to output the previously calculated valid data to the controllers, and to flag the data as invalid, so that the AD software will not use them in updating the Kalman filter. Another frequent response is to notify the ground of the fault. These are faults for which the FOT will collect statistics from the TSM software, but which do not represent an immediate threat to observatory health and safety. In these cases, persistence is unnecessary, because each outage is recorded by the on-board software and telemetered to the ground when requested.

Yet another common technique is to set a special flag when certain expected circumstances arise which invalidate a sensor's data. For example, the STs will sometimes be occulted by the Earth, and Condition ST-2 helps to manage the system response to those events (as discussed in an earlier section). However, since these occultations are expected daily in our mission orbit, FDC should not perform any autonomous reconfiguration. So, the flag set in response to Condition ST-2 is used to avoid tripping ST-3, which is meant to detect an unexpected loss of track. The same technique is used to prevent eclipses from being misinterpreted by FDC as Sun sensor failures.

Responses to GT data are all in-line, and fall into one of two categories. Either the condition is consistent with eclipse (GT-2), or the condition is clearly a fault. The GT is only used in Science mode, and operationally, the observatory should be switched out of Science mode before eclipse occurs. Since this condition is only a threat in Science mode, it is classified as a mode-specific test. The hardware responses only set flags that indicate the eclipse or the fault condition, and to which the mode-specific checks refer.

Actuator checks deal mainly with detection and prevention of faulty commands, though the propulsion system includes some basic safety checking. The check on the RWA speed is at first glance an exception to the principal of eliminating responses to non-safety items, as jitter is only a problem for the science data quality (though it is of large concern)⁷. However, since the operational plan should never allow system momentum to be high enough for the RWA speed to exceed jitter limits, RWA-1 does provide a notification of a failure in the complex angular momentum management required of the SDO ACS and the FOT. Note that all thruster commands are cancelled in the event of one command failing to meet standards. HGA commands are stopped whenever any of several quality checks is failed, and possible irradiation of the instruments draws the additional response of turning off the Ka-band transmitter.⁸

Table 1: Hardware-Specific Fault Conditions and Associated Responses

ID	Condition	Response	Persistence
CSS-1	CSS counts outside valid range	0) Use previous value	In-line
CSS-2	7 out of 8 CSS are dark (ACE in control)	0) Set CSS Eclipse Flag	In-line
DSS-1	No DSS sun presence	0) Use previous value and disable use of DSS data in KF for this cycle	In-line
DSS-2	No DSS packet update	0) Use previous value and disable use of DSS data in KF for this cycle	In-line
ST-1	ST data packet not received	0) Use previous data & disable use of ST data in KF for this cycle	In-line
ST-2	ST Boresight is within occultation cone of Sun, Moon, or Earth	0) Flag data as low quality & disable use of ST in KF for this cycle	In-line
ST-3	Star tracker is not tracking and not occulted	1) Command ST to reset and return to track	5 sec
TARA-1	Change in integrated angle exceeds limit	0) Use previous value for control & AD	In-line
TARA-2	Angular rate comparison: Primary TARA rates vs backup TARA rates vs primary ST rates	1) Two out of three voting: a) Reconfigure to back-up TARA if back-up TARA agrees with ST b) Disable this check if back-up TARA or ST gets voted out	5 sec
TARA-3	Primary TARA rates vs backup TARA rates vs DSS-derived rates	1) Similar to TARA-2 response	5 sec
TARA-4	TARA motor current outside limits	1) Notify ground	N/A
GT-1	Photodiode error condition detected	0) Use previous value and flag GT as invalid	In-line
GT-2	No Sun on GT	0) Use previous value and set GT acquisition flag to FALSE	In-line
GT-3	No GT Packet Update	0) Use previous value and flag GT and Controlling GT data as invalid	In-line
RWA-1	RWA speed outside allowed range for jitter	0) Set a flag warning that jitter might affect the science data	In-line
RWA-2	More than one RWA is powered off	1) Power On All RWA	1 sec
RWA-3	Commanded torque and change in measured speed do not compare for one RWA	1) Disable all "disable wheel" RTSSs, and disable and power off the RWA	15 sec
PROP-1	Commanded thruster counts are invalid for main engine or any thruster	0) Cancel all thruster commands and set flag to indicate commands cancelled	In-line
PROP-2	Thruster command counts echo not equal to commanded counts	0) Cancel all thruster commands and set flag to indicate commands cancelled	In-line
PROP-3	Engine valve temperature exceeds limit	1) Go to Sun Acq mode	1 sec
HGA-1	Elevation gimbal axis angle outside range	0) Command 0 pulses to elevation actuator and continue tracking in azimuth	In-line
HGA-2	Expected antenna position not equal to measured position	1) Send HGA Stop commands	1 sec
HGA-3	Invalid or absent GCE data packet	0) Send 0 pulses to both gimbals 1) Send HGA Stop commands and power off Ka-Band transmitter	0) In-line 1) 30 sec
HGA-4	HGA is pointing in keep-out zone (i.e. danger to science instruments)	1) Send HGA Stop commands and power off Ka-Band transmitter and gimbal drivers	1 sec

ATTITUDE CONTROL ELECTRONICS (ACE) ARBITRATION

SDO incorporates two electrically identical ACE boxes; Figure 3 shows an electrical block diagram of the ACE. Each ACE interfaces to the observatory over the 1553 bus. The ACE provides the interface between a number of ACS sensors, all of the ACS actuators, and the observatory. The ACE also serves as the platform for the Safehold control mode that is executed by the Subsystem Data Node (SDN) contained in each ACE. Since a transition into Safehold is not considered a failure in itself, it was a goal that the ACEs be able to arbitrate control authority between the ACT and themselves without the intervention of the main processor.

Each ACE receives and conditions the signals from 8 of the 16 CSS. After conditioning, the CSS signals are converted to digital values for use by the Safehold controller and supplied to the MP for use in other control modes. A field programmable gate array (FPGA), located on the SDN, controls the conversion. Since a fully determined Sun vector can be achieved with any 7 of the 8 CSSs, each ACE is considered single-fault tolerant of a CSS failure.




Figure 3: Block diagram of the SDO Attitude Control Electronics

The ACEs also receive the positive and negative integrated rate pulses from all six TARA axes. The pulses, which represent accrued displacements of 0.5 arcsec, are accumulated in six registers in the FPGA on the SDN card. The value in the register represents the total angular displacement about that axis since the register was last reset. The registers are provided to the ACT and used by Safehold to derive rates that are only used when the observatory is in eclipse.

Each ACE contains all the necessary interfaces to drive the valves and thrusters of the propulsion subsystem. The ACE provides the command interface to the pyrotechnic valves. Some of these valves are used to pressurize the propellant delivery system after separation from the launch vehicle. Others are used to convert the propulsion system to “blow-down” mode after arrival in mission orbit. A final pyrotechnic valve will be used at the end of the mission to reopen the pressurization path and allow the system energy to be minimized for disposal. The ACE also supplies the command and telemetry interface to the propulsion system latching valves. These latching valves can be actuated repeatedly to open or isolate various parts of the propulsion system throughout the mission. Finally, the ACE incorporates the necessary drive electronics to command the ACS thrusters and main engine. Although all of these propulsion system interfaces are implemented in the ACEs, the Safehold controller does not interact with them. All normal propulsive operations are accomplished through the ACT only.

The ACE interface to the four RWAs is through the reaction wheel interface card (RWIC). The tachometer pulses from each RWA are collected in four registers on the RWIC. Each RWIC also has four latching relays that are used to arbitrate control of the four RWAs. The relays allow either analog torque commands or ground reference voltage to be connected to each ACE RWA output. Since each RWA accepts the inputs from both ACEs in a summing junction, it is necessary to ground the unused input. An ACE has the ability to assume control of the RWAs by switching the state of its own four relays to provide control signals and switching the four relays in the other ACE to ground. The ACEs do not have the capability to give control to the other ACE. Only the in-control ACE supplies the analog torque commands from either the ACT or Safehold, depending on arbitration results with the MP.

To establish a simple monitoring system between the ACEs, each ACE sends a two-bit status counter to the other, referred to as the “I’m OK” or IMOK signal. The transmit and receive registers are both located in the SDN FPGA. These counters normally cycle from 0 to 3, incrementing each ACE cycle, when the ACE is operating normally. These counters are an integral part of the ACE arbitration scheme.

The control arbitration can be broken into two distinct elements. The first is ACE arbitration, that is, which ACE is in control and sending commands to the RWA. The second is Safehold arbitration, that is, whether the ACE is passing MP commands to the actuators or is in Safehold mode and sending its own RWA commands, ignoring the MP actuator commands.

Normally, for ACE arbitration the in-control ACE is determined at observatory power-on and the startup of the processors. During startup each ACE confirms communication with its RWIC; if successful, the ACE will begin providing IMOK signals at completion of power-on. The in-control ACE is then determined based on the ACE with the majority of the eight relays. In the

unlikely event of a tie, a TSM/RTS combination on the MP selects the backup ACE. Upon establishing the majority, the in-control ACE will then take control by pulling all relays to it (and grounding the other ACEs relays). This ACE will also disable its response to a loss of IMOK signals from the backup ACE. The backup ACE will also transmit its IMOK; however, it will enable its response to a loss of IMOK from the in-control ACE. In subsequent cycles each ACE will transmit its IMOK so long as analog-to-digital conversion (ADC) is successful for all eight CSSs. If either a hardware or software error causes the IMOK from the in-control ACE to be interrupted, an IMOK failure counter on the backup ACE is incremented. If this counter reaches 60 ACE cycles, the backup ACE will interpret this as an indication of failure of the in-control ACE and pull control of the RWAs. If the IMOK from the in-control ACE is reestablished for 10 cycles prior to the backup taking control, the failure counter will be reset to 0 and the backup will continue monitoring the IMOK signal from the in-control ACE.

The same FPGA that sends and receives the IMOK also performs the CSS ADC and provides the interface to the TARA integrated rate pulses, and is therefore critical to proper Safehold functionality. Therefore, a test was established to reduce the risk that an ACE with a malfunctioning FPGA would take control. This is the FPGA read back test. The value written into the register that provides the IMOK to the other ACE is read back by the transmitting ACE. If this value is the same as was written, the FPGA is assumed to be functioning normally. A bad read back on the in-control ACE has no effect except an indication in telemetry. A failed read back on the backup ACE, however, will result in any value of IMOK received being treated as correct. The FPGA read back test can be overridden by command.

For Safehold arbitration, the in-control ACE monitors the health of the MP similarly to how the backup ACE monitors the in-control ACE. The MP is normally in control, and the in-control ACE normally passes MP commands on to the actuators with a small amount of conditioning and quality checking (Normal mode). Should the IMOK signal from the MP be interrupted for more than a table-defined number of ACE cycles, the in-control ACE will replace the MP RWA torque commands with those generated from its own Safehold task (Safehold mode). The ACE may also be commanded to Safehold, with the same result. Recovery from Safehold is accomplished by reestablishing observatory command through one of the MPs and commanding the in-control ACE back to Normal mode.

CONTROL MODE PERFORMANCE

The hardware-specific and ACE arbitration FDC is designed to address the known failure paths in the SDO ACS design as understood by designers and systems engineers. However, the possibility of an overlooked path from a single fault to mission loss leads to additional checks that the ACS is performing correctly. FDC based on control mode performance is qualitatively similar to performance requirements on the five ACT modes. Attitude errors and angular rates should normally remain low, and FDC checks these quantities. However, the actual performance values set as goals for the design and testing of those control modes are not necessarily the standards used in the detection of failures. For instance, though the control performance requirement on the Inertial mode controller is in the range of arcseconds, no real hazard exists inherently in attitude excursions even out to the level of degrees. Instead, the FDC performance failure levels are set to be enough outside of expected performance, as based on analysis and

simulation, that excursions beyond those levels are not expected except in the case of hardware failure.

Table 2 shows the FDC based on control mode performance. The format is similar to the hardware-specific FDC. The ACT control modes that run on the MP are described above, but are listed again here for easy reference: Sun Acquisition (SA), Inertial (IN), Science (SCI), DeltaH (DH), and DeltaV (DV).

FDC conditions based on control mode performance all depend on at least two logical criteria to activate: 1) each must meet its described condition in telemetry, and 2) the ACT must be operating in the specified mode. So, the RTS in response to IN-2, for example, will not be commanded unless both the listed condition is true for the entire persistence period and the ACT is in Inertial for the entire persistence period. Because Sun Acquisition mode is expected to take some time to acquire its target, as opposed to all other modes, which essentially begin on their targets, an additional mode restriction is placed on the SA-1 performance check. Before the TSM indicates failure at all, the MP must have been in Sun Acquisition mode for at least 30 minutes. This restriction prevents normal operation of Sun Acquisition, which can take up to 30 minutes to acquire the Sun, from triggering the SA-1 TSM and issuing 5 failure event messages per second for 30 minutes during normal operations.

The responses to control mode performance faults are nearly all mode demotions; the one exception—SCI-4—is useful more for monitoring of an unexpected condition than for automated response. Sun Acquisition mode is the usual mode of choice, though sometimes Inertial mode is selected. Two conflicting ideas are at work in this design. One idea is to demote in incremental steps, allowing each new state the possibility of recovering a stable configuration and lengthening the amount of time before all possible automated responses are exhausted. The other is to isolate the system as much as possible from the total condition at the time of the fault. Both are sensible, and for the SDO ACS, the second is usually the proper choice. The SDO Sun Acquisition mode is extremely robust to failures of any of the hardware complement. It is capable of maintaining adequate power on the solar arrays even in many double-fault conditions. Since it is primarily dependent on the CSSs, and since no other MP mode uses those sensors, Sun Acquisition provides the most complete safe harbor from hardware problems that may disturb other control modes. The exceptions to this reliance on Sun Acquisition are cases in which it is highly likely that a software glitch would cause the fault, such as if the MP warm resets, or in the special case of a Science mode GT problem. In the case of a GT problem, Inertial mode is a safe harbor because it does not use the GTs in any way. Of course, a failure of Sun Acquisition to acquire the Sun can only be addressed by hoping the problem lies in the ACE in control or in a severe TARA fault and switching to Safehold on the backup ACE, which is independent of the TARAs.

Table 2: Control Mode Performance Fault Conditions and Associated Responses

ID	Condition	Response	Persistence
SA-1	Sun angle exceeds 15 deg <u>and</u> ACS mode timer exceeds 30 minutes	1) Go to Safehold on other ACE	60 sec
IN-1	Attitude error exceeds 5 deg	1) Go to Sun Acquisition Mode	60 sec
IN-2	Angular rate exceeds 0.3 deg/sec	1) Go to Sun Acquisition Mode	30 sec
SCI-1	Angular rate exceeds 0.3 deg/sec	1) Go to Sun Acquisition Mode	30 sec
SCI-2	CGT outside linear range	1) Go to Inertial mode	5 min
SCI-3	CGT outside of acquisition range or invalid	1) Go to Inertial mode	10 sec
SCI-4	Roll attitude error exceeds 10 arcseconds	0) Set a flag warning that the attitude excursion might affect the science data	In-line
DH-1	Time in DeltaH exceeds 7.5 min	1) Go to Sun Acquisition Mode	0.6 sec
DH-2	Magnitude of the system angular momentum error exceeds limit	1) Go to Sun Acquisition Mode	0.6 sec
DH-3	Attitude error exceeds 5 deg	1) Go to Sun Acquisition Mode	0.6 sec
DH-4	MP performs a warm restart	1) Go to Previous Mode	Immediate
DH-5	ACE in control performs warm restart	1) Go to Sun Acquisition Mode	Immediate
DH-6	Thruster commands cancelled (see PROP-1,2)	1) Go to Sun Acquisition Mode	0.4 sec
DV-1	Time in Delta V exceeds commanded time	1) Go to Inertial Mode	0.6 sec
DV-2	Angular rate exceeds limit	1) Go to Sun Acquisition Mode	0.6 sec
DV-3	Attitude error exceeds 5 deg	1) Go to Sun Acquisition Mode	0.6 sec
DV-4	MP performs warm restart	1) Go to Inertial Mode	Immediate
DV-5	ACE in control performs warm restart	1) Go to Sun Acquisition Mode	Immediate
DV-6	Thruster commands cancelled (see PROP-1,2)	1) Go to Sun Acquisition Mode	0.4 sec

SYSTEM HEALTH AND SAFETY

The big net that catches any faults that fall or cascade through the other three levels is the FDC for system health and safety. These tests check the specific quantities that either represent direct threats to the spacecraft, such as large angular momentum or Sun angle values, or that involve multiple inputs such that the error source may be difficult to trace, such as convergence of the Kalman filter. Table 3 is divided into 4 sub-categories: hazardous physical conditions (HAZ); ephemeris checks (EPH), which mainly feed into other FDC; checks on the status of the extended Kalman filter (KF); and cross-checking of attitude sources using Sun vectors (ATT).

The first three HAZ checks look at the system angular momentum. HAZ-1 and HAZ-2 are essentially the same test, but with different limits. The lower limit has a slower response, and would respond to a gradual build-up of momentum that is not likely to be caused by a concentrated external torque. Conversely, the higher limit has a very rapid response, with the assumption being that, since the lower limit test was not triggered, there is an actual leakage from the propulsion system (SDO has no dewars, so the propulsion system is the only potential source of rapid gas expansion). The third test looks for the same possibility, but tries to catch it before it reaches a high momentum. HAZ-2 and HAZ-3 have the same response, so there is no danger in having these two tests racing each other—in fact, that is the intention.

Ephemeris faults all result in the ephemeris validity flag being set to Invalid until reloaded by the FOT. Then, KF-1 does not use any sensor inputs for any cycle in which the ephemeris is flagged as invalid. Other Kalman filter checks similarly exclude one or all of the three fine attitude sensors (ST #1, ST #2, and DSS) from use in the filter, either for one cycle or until reset by the FOT. KF-8 then sweeps up those results into a holistic check on the health of the fine attitude pointing. The Kalman filter can really only be trusted to provide accurate attitude solutions when it receives accurate data from any two of the three fine attitude sensors. One sensor is sufficient for safe operation, though not for meeting knowledge requirements, but this is not an expected state for the ACS and may indicate a more insidious problem. So, KF-8 monitors whether the filter is in fact receiving suitable data from two sensors. If a sufficient time passes without two accurate measurements, then KF-8 responds by changing to Sun Acquisition mode, which is the safest MP mode, since it does not use the Kalman filter at all. Some of the other KF conditions, such as the divergence of the filter, are an immediate threat; continued divergence of the Kalman filter can result in “not-a-number” operations, which can eventually cause reset of the MP software. These conditions will also send the MP to the safety of Sun Acquisition mode until the Kalman filter problem can be isolated. If any collection of human error, ground or flight software error, subtle hardware error, or unfortunate radiation or dynamic events sufficiently corrupts the KF in any of these several ways, then one or more KF conditions will be met, and the spacecraft will end up in Sun Acquisition.

The ATT conditions are all interconnected, and are only checked when the spacecraft is not in eclipse. There are four sources of calculated Sun vectors in the onboard software: the CSS Sun vector from ACE A (CSS-A), the CSS Sun vector from ACE B (CSS-B), the DSS Sun vector, and the Sun vector estimated in the SDO body frame by rotating the inertial-referenced Sun location from the ephemeris into the body reference frame using the current attitude quaternion. Each of the four conditions has one of these Sun vectors being voted out by two other sources that agree with each other to within a few degrees. Then, the responses isolate the ACS from the faulty measurement. If CSS-A is faulty, then that ACE cannot be trusted if Safehold were needed. Since ACE A is the nominally in-control ACE, FDC places ACE B in control of the spacecraft using its Sun vector measurements, which the test has effectively validated against the DSS measurement. If instead the ACE B Sun vector is faulty, then transitions to ACE B for other FDC reasons are disabled until the problem can be understood. If the DSS is faulty, it is excluded from the KF, with the potential consequence of meeting the condition for KF-8 if, for example, ST-1 is occulted. Finally, if the AD solution is faulty, Sun Acquisition is entered as the only MP mode that is independent of the AD processes.

Table 3: System Health and Safety FDC

ID	Condition	Response	Persistence
HAZ-1	Total angular momentum exceeds 35 Nms	1) Go to Sun Acquisition Mode	10 sec
HAZ-2	Total angular momentum exceeds 50 Nms	1) Close propellant isolation valves and remove power from thruster coils	3 sec
HAZ-3	Total angular momentum is increasing faster than 3 Nms per 200-msec cycle	1) Close propellant isolation valves and remove power from thruster coils	0.4 sec
HAZ-4	CSS Sun Angle exceeds 45 deg	1) Go to Sun Acquisition Mode	5 min
EPH-1	Ephemeris Uplink Error	0) Reject Ephemeris Load	Immediate
EPH-2	Magnitude of the spacecraft position vector is outside the specified range	0) Set ephemeris validity flag to invalid	In-line
EPH-3	Ephemeris propagation has been running more than 20 days	0) Set ephemeris validity flag to invalid	In-line
KF-1	Ephemeris is flagged as invalid	0) Exit KF Routine for 1 cycle	In-line
KF-2	Kalman Filter is not converged	1) Reset KF 2) Disable KF and Go to Sun Acquisition	1) 5 min 2) 10 min
KF-3	KF covariance diagonal has a negative value	0) Exit KF for one cycle 1) Reset KF 2) Disable KF and Go to Sun Acquisition	0) In-line 1) 5 sec 2) 10 sec
KF-4	KF covariance has diverged	0) Exit KF for one cycle 1) Reset KF 2) Disable KF and Go to Sun Acquisition	0) In-line 1) 5 sec 2) 10 sec
KF-5	KF drift bias correction exceeds limit	1) Reset KF 2) Disable KF and go to Sun Acquisition	1) 30 sec 2) 60 sec
KF-6	ST coarse residual exceeds limit <u>but</u> ST is not occulted by Earth, Sun or Moon	0) Drop sensor measurement from KF 1) Disable KF use of ST and Reset KF	0) In-line 1) 10 min
KF-7	DSS coarse residual exceeds limit <u>but</u> SDO is not in eclipse	0) Drop sensor measurement 1) Disable KF use of DSS and reset KF	0) In-line 1) 10 min
KF-8	For two or more of the fine attitude sensors (ST-1, ST-2 or DSS), for any reason, the sensor measurement was not used by the KF	1) Go to Sun Acquisition Mode	10 min (1 min)
ATT	The following 4 tests compare Sun vectors from 4 sources: CSS-A, CSS-B, DSS, AD. A test automatically passes if one of its sources shows no Sun presence.		
ATT-1	CSS-A not equal CSS-B and CSS-A not equal DSS and CSS-B equal DSS	1) Go to Safehold on ACE-B, because CSS-As are suspect	1 sec
ATT-2	CSS-A not equal CSS-B and CSS-B not equal DSS and CSS-A equal DSS	1) Disable RTSs that go to Safehold on backup ACE	1 sec
ATT-3	CSS-A equal CSS-B and either CSS-A not equal DSS or CSS-B not equal DSS	1) Disable DSS from KF indefinitely	1 min
ATT-4	CSS-A equal CSS-B either CSS-A not equal AD or CSS-B not equal AD	1) Go to Sun Acquisition Mode	1 min

Conclusion

Fault detection and correction on the SDO is designed in recognition of the value both of simplicity and comprehensiveness. All foreseen potential single faults are addressed, but in adherence with principles that encourage simplicity. These principles result in an approach that limits FDC use for any but recognized threats to health and safety of the observatory, simplifies the universe of possible automated responses, and establishes an organization of tests that should allow any cascading failures to settle into a communicative, power-positive, thermally safe attitude. The FDC organization as developed divides the tests into four broad categories, each with its own focus, and each able to respond to problems caused by other FDC in a stabilizing way. Hardware-specific checks attempt to isolate hardware faults before they propagate through the ACS. ACE arbitration FDC provides a self-policing algorithm for determining which of the onboard processors is most fit for attitude control. Mode performance FDC monitors a very narrow field of factors against the minimally safe operation of each control mode. System health and safety check cross-protect the observatory by addressing direct threats to operation due to potentially complex interactions in what is necessarily a complex attitude control system.

References

- ¹<http://lws.gsfc.nasa.gov>, NASA Living With a Star web site. Responsible official: Mary DiJoseph
- ²<http://sdo.gsfc.nasa.gov>, NASA SDO web site. Responsible official: Dean Pesnell
- ³Scott R. Starin, Kristin L. Bourkland, Kuo-Chia Liu, Paul A. C. Mason, Melissa F. Vess, Stephen F. Andrews, and Wendy M. Morgenstern. Attitude Control System Design for the Solar Dynamics Observatory. Flight Mechanics Symposium, 2005.
- ⁴Kristin L. Bourkland, Scott R. Starin, and David J. Mangus. The Use of a Gyroless Wheel-Tach Controller in SDO Safehold Mode. Flight Mechanics Symposium, 2005.
- ⁵Melissa F. Vess, Scott R. Starin, and Wendy M. Morgenstern. Use of the SDO Pointing Controllers for Instrument Calibration Maneuvers. Flight Mechanics Symposium, 2005.
- ⁶Paul A. C. Mason and Scott R. Starin. SDO Delta H Mode Design and Analysis. International Symposium on Space Flight Mechanics, 2007.
- ⁷Kuo-Chia (Alice) Liu, Thomas Kenney, Peiman Maghami, Pete Mule, Carl Blaurock, and William B Haile. Jitter Test Program and On-Orbit Mitigation Strategies for Solar Dynamics Observatory. International Symposium on Space Flight Mechanics, 2007.
- ⁸Kristin L. Bourkland and Kuo-Chia (Alice) Liu. A Jitter-Mitigating High Gain Antenna Pointing Algorithm for the Solar Dynamics Observatory. International Symposium on Space Flight Mechanics, 2007.