NASA/TM–2004–206892, Vol. 28

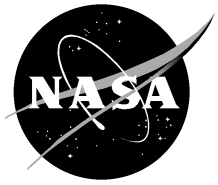

SeaWiFS Postlaunch Technical Report Series

Stanford B. Hooker and Elaine R. Firestone, Editors

Volume 28, Pointing Performance for the SeaWiFS Mission

Stephen Bilanow and Frederick S. Patt

National Aeronautics and
Space Administration

Goddard Space Flight Center
Greenbelt, Maryland 20771

March 2004

The NASA STI Program Office ... in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA's counterpart of peer-reviewed formal professional papers but has less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATION. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or cosponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and mission, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that complement the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results... even providing videos.

For more information about the NASA STI Program Office, see the following:

- Access the NASA SIT Program Home Page at http://www.sti.nasa.gov/STI-homepage.html

- E-mail your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA Access Help Desk at (301) 621-0134

- Write to:
  NASA Access Help Desk
  NASA Center for Aerospace Information
  7121 Standard Drive
  Hanover, MD 21076-1320

NASA/TM–2004–206892, Vol. 28

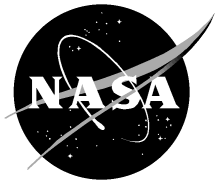

SeaWiFS Postlaunch Technical Report Series

Stanford B. Hooker, Editor
NASA Goddard Space Flight Center, Greenbelt, Maryland

Elaine R. Firestone, Senior Scientific Technical Editor
Science Applications International Corporation, Beltsville, Maryland

Volume 28, Pointing Performance for the SeaWiFS Mission

Stephen Bilanow and Frederick S. Patt
Science Applications International Corporation, Beltsville, Maryland

March 2004

ISSN 1522-8789

Available from:

NASA Center for AeroSpace Information
7121 Standard Drive
Hanover, MD 21076–1320
Price Code: A17

National Technical Information Service
5285 Port Royal Road
Springfield, VA 22161
Price Code: A10

Abstract

The onboard pointing performance of the OrbView-2 (OV-2) spacecraft for the first five years of the SeaWiFS mission is presented. Adjustments to the onboard attitude control system (ACS) since launch are described, and various issues and anomalies regarding the performance are discussed. Overall, this relatively low-cost spacecraft has performed quite effectively after various in-flight adjustments, however, a variety of sensor and computational anomalies have caused occasional minor pointing disturbances. Many of these disturbances have implications for the navigation processing performed for the science data by the Sea-viewing Wide Field-of-view Sensor (SeaWiFS) Project at the NASA Goddard Space Flight Center. Possible further adjustments to the OV-2 ACS have been investigated which could lead to improved pointing performance, and conclusions from these analyses are presented. Some of the various sensor and software anomalies are fairly well understood, but some others remain puzzling. Particularly vexing are various timing anomalies resulting from the way the clocks for three separate onboard processors are synchronized to each other and Global Positioning System (GPS) time. The pattern of occurrence of some of the sensor anomalies could merit further review and trending, which may be useful to monitor for any degradation in performance as the mission continues.

1. INTRODUCTION

The pointing control of the OrbView-2 (OV-2) spacecraft, which was built and launched by Orbital Sciences Corporation (OSC) and operated by OrbImage, has been adjusted in a number of major and minor ways since launch. The adjustments have been made to onboard tables, or in a few cases, via an onboard software patch or *poke*. The main goal has been to minimize control disturbances so that ground-computed attitudes can generally meet the accuracy requirements. A variety of sensor and onboard computational anomalies have been recognized, which occasionally contribute to onboard attitude errors and resulting pointing disturbances.

This report documents the onboard pointing and adjustments, the pointing stability of the spacecraft, and various anomalies. The following introductory sections present relevant background material about the spacecraft, the orbit, and data availability. Section 2 discusses the onboard attitude estimation and its errors, and adjustments made to improve the performance. Section 3 discusses the overall pointing stability in terms of the dynamics, disturbance sources, control responses, and control adjustments. Section 4 describes various anomalies, and Sect. 5 presents a discussion of key results and lessons learned.

1.1 Spacecraft Description

The OV-2 spacecraft (originally called SeaStar) is illustrated in Fig. 1. The single-imaging instrument—the Sea-viewing Wide Field-of-view Sensor (SeaWiFS)—is at the bottom of the structure, as shown in the nominally Earth-nadir pointing direction. Deployed solar array panels are at the top. The nominal flight direction is roughly into the page, so that the solar diffuser scope, on the left of the SeaWiFS instrument, and the back panel (shown with numerous boxes) face away from flight.

1.1.1 Coordinates and Attitude Angles

The attitude control system (ACS) coordinates used in this report are shown in Fig. 1 with respect to the spacecraft structure†. The ACS $\vec{X}$, $\vec{Y}$, and $\vec{Z}$ axes are also referred to as the roll, pitch, and yaw axes, respectively. When roll, pitch, and yaw are referred to as angles, they represent rotations about these axes from the nominal flight orientation: $\vec{X}$ is along the general direction of orbital motion, $\vec{Y}$ is opposite to the orbit normal direction, and $\vec{Z}$ is toward the nadir. The order of rotation, i.e., pitch–roll–yaw, is applied when computing the angles for the estimated attitude.

Sensor and actuator orientations are described in this document relative to these ACS axes, henceforth referred to as $\vec{X}$, $\vec{Y}$, and $\vec{Z}$.

1.1.2 ACS Overview

Being a relatively low cost mission, the OV-2 spacecraft does not carry instruments which are used for very accurate onboard attitude determination, such as, a gyroscope or star tracker. Instead, the spacecraft relies on alternative instrumentation: the two-axis digital sun sensor (DSS), the Earth horizon scanner (HS), and the three-axis magnetometer (TAM). Stability is maintained by a single-momentum wheel, which provides gyroscopic inertia along the pitch axis, $\vec{Y}$. The wheel speed is varied to

† The labeling for these axes is different than that used for the ground data processing (Patt 2002) for the SeaWiFS base (and also used elsewhere for OV-2 spacecraft and SeaWiFS instrument coordinates), where S_x is the axis normally along the geodetic nadir, S_y is the axis normally opposite the velocity direction, and S_z is the axis normally toward orbit normal. The definitions of roll, pitch, and yaw are nevertheless equivalent in the onboard and ground computations.

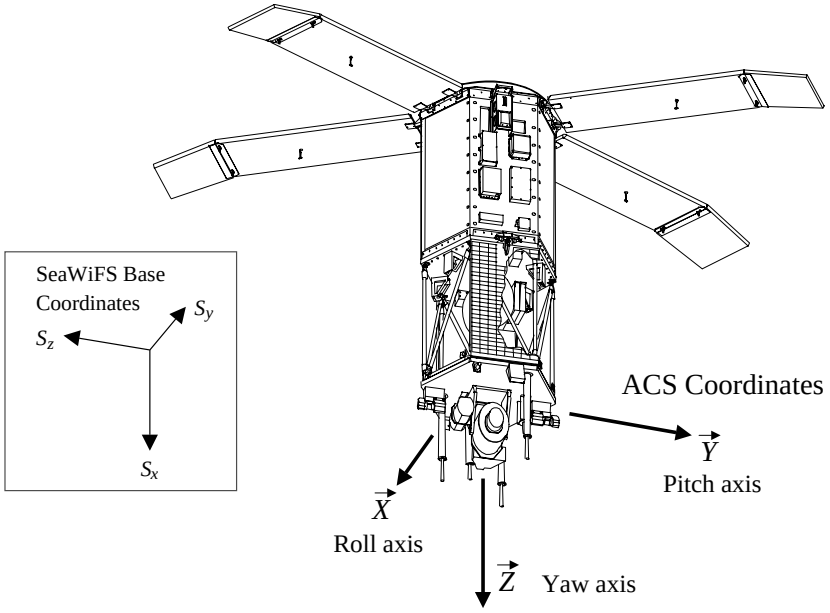

Fig. 1. OV-2 spacecraft with the SeaWiFS instrument and associated ACS coordinates and axes shown. The SeaWiFS base coordinates are nominally aligned with the ACS coordinates as follows: $S_x = Z$, $S_y = -X$, and $S_z = -Y$.

control pitch, while roll, yaw, and the total angular momentum are adjusted by magnetic coil (torque rod) interactions with the Earth's magnetic field. This control is described in more detail in Sect. 3.

1.1.3 Avionics Overview

The flight computer consists of three separate microprocessors:

a) Spacecraft control module (SCM) for attitude estimation and control;

b) Ground interface module (GIM) for telemetry data handling; and

c) Payload support module (PSM) for SeaWiFS image processing and data recording.

The SCM and GIM have reset several times during the mission, resulting in degraded ACS performance and temporary data loss, respectively. The effects of the SCM resets are discussed in Sect. 4.5.6.

There are also three electronics boxes of particular interest:

1) Battery charge regulator (BCR) controls battery charging;

2) Global Positioning System (GPS) produces output used by ACS for orbit propagation; and

3) Attitude control electronics (ACE) handles interfaces to ACS sensors and actuators.

The BCR switched autonomously from the primary (A side) to the redundant (B side) unit on 18 December 1998 (with an effect on the pointing as discussed in Sect. 3.1.6). Since that time, some processor resets have occurred which would have been expected to switch control back to the A side. Instead, the B side has remained in use, so it is believed that the A side failed.

The GPS unit provides new orbit estimates every 10 s, and these are filtered for use in an onboard orbit propagator; however, glitches occasionally cause onboard position estimation errors as discussed in Sect. 4.4. GPS error effects were severe enough early in the mission that a software patch was developed and loaded on board in December 1997 to better reject occasional invalid data, as discussed in Sect. 4.4. The backup onboard orbit source uses vectors loaded from the ground.

There are two ACE boxes, ACE-A and -B, and each one interfaces with a subset of the attitude sensors. The mission plan called for running only the primary (A) unit after initial in-orbit checkout, keeping the other box and associated sensors as cold backups. Various calibrations had to be completed before the switch to just the A side—referred to as *single-string operation*—could be implemented. This change was not finally implemented until just over a year of operations, in November 1998. After the redundant (B) unit was powered off, the sensor subset as noted in Sect. 2.1 was used. ACE-B and the backup attitude sensors are not expected to be activated again unless needed because of a component failure on the A side.

1.1.4 Control Modes and Safe Haven

The control modes used in the ACS flight code were defined by OSC, with names and descriptions as shown in Table 1 (not including some special modes for launch,

Table 1. Control modes encountered since early mission checkout.

Mode Name	*Description, Notes*
Fine Pointing	Used in normal mission operations. Takes data from all available attitude sensors.
Lunar Calibration	A 360° pitch rotation over about 30 min for sweeping the SeaWiFS instrument field of view (FOV) past the moon for calibration. Calibration maneuvers are done about once a month near full moon.
Despin	Entered if high attitude rates are detected.
Hunt for Sun Earth	Entered for attitude re-acquisition check prior to coarse point.
Coarse Pointing	Based only on local pitch and roll angle and rate measurements from HS data, thus not needing yaw estimates.

startup, low power contingency backup, and orbit raising). The normal mission mode is called *fine pointing*. The *lunar calibration* mode is only used during calibration maneuvers, nominally once each month. *Despin* and *hunt* modes have only been entered very briefly under a few anomalous conditions, with transitions to *coarse pointing* mode. The transition from coarse pointing mode back to fine pointing mode can be done by ground command. The spacecraft can switch autonomously to coarse pointing mode and back to fine pointing mode under certain conditions as described in Sect. 2.2.5.

The following discussion presumes the fine pointing mode, unless otherwise noted, as the operational mode in which science data are collected. Attitude estimation for this mode is discussed in detail in Sect. 2.2, and control for this mode is discussed in Sect. 3.1.

OSC also defined *safe haven* as a spacecraft mode, which is not an ACS control mode, but is an important general spacecraft mode for instrument safety, in which the SeaWiFS instrument is turned off. This is entered when errors are detected that might endanger the instrument especially by pointing it at the sun. Safe haven has been triggered in a number of cases when data checks in fine pointing mode erroneously estimated high attitude rates, sending the spacecraft briefly into despin and hunt modes before settling into coarse pointing mode. Safe haven mode is triggered by entry into despin or hunt modes and sets a flag to hold the ACS in coarse pointing mode.

There is also an ACE hardware-controlled *safe hold* mode for the case in which the flight computer is down, but that is not relevant for normal operations, and fortunately has not been needed.

1.1.5 Flight Software Adjustments

Normal adjustments to flight software parameters are accomplished by table loads—memory changes loaded by planned command to the spacecraft. There are a large number of parameters that can be adjusted this way, including various biases and alignment data for the attitude sensors. If a processor on board is reset (rebooted), default values are picked up and all changes must be reloaded. There is an option for storing changes permanently in the Electronically Erasable Programmable Read Only Memory (EEPROM), but that has not been implemented for most of the ACS adjustments discussed in this report.

Two flight code patches (modifications to the onboard software) have been implemented on board since launch:

1) GPS Patch: GPS errors initially were severe enough to make the onboard orbit unreliable. Additional checks on the GPS output were coded and implemented on 15 December 1997.
2) Lunar Calibration Patch: An error occurred during a lunar calibration maneuver on 10 July 1998, which caused the spacecraft to enter safe haven mode. Analysis showed the probable cause to be sensitivity of the algorithm to noise in the HS, which falsely detected Earth re-acquisition and loss. A check for a reasonable number, i.e., 5, of consecutive points on or off the Earth was added before mode changes took place.

In addition to these patches, one poke (in which one word in onboard memory was changed) was performed to change a value in memory that was not accessible by normal table loads. This was for limiting the DSS FOV and is discussed in Sect. 2.5.2.

One additional software patch was proposed to permit counteracting effects of an apparent magnetic dipole bias as discussed in Sect. 3.1.6, however, this patch was never tested and implemented.

Two flight code analysis tools were applied for much of the work discussed in this report. First, a PC-based dynamics simulation with the ACS flight code was developed by OSC prior to launch for analysis, and proved very useful after launch in analyzing and simulating some of the flight behavior. The C++ code and associated utilities were provided to the SeaWiFS Project and applied in various analysis efforts. All flight code adjustments were tested with this simulation, and the code was modified temporarily to simulate and analyze some of the anomalies.

Second, OSC built a spacecraft simulator, consisting of flight hardware components on laboratory benches, connected to microprocessors running the flight software. It

was known as *FlatSat* (for flat satellite). FlatSat was used by OSC and OrbImage for verification of flight code modifications, and simulation of the load process with interfaces to the ground system. FlatSat eventually encountered various maintenance problems that could not be addressed because of limited resources. By 2000, further use by OSC and OrbImage engineers was not possible; therefore, some additional flight code changes of lower priority have not been implemented.

1.1.6 State of Health Telemetry

The spacecraft State of Health (SOH) telemetry are included in the science data stream (Sect. 1.3) and are also available at low rates during non-science data collection (i.e., back orbit) periods. There are three subsets of interest for monitoring the ACS performance:

1) SCM ACS Control—Attitude system and various related parameters;
2) SCM ACE Agent—Attitude sensor and actuator data; and
3) SCM GPS Agent—GPS orbit-related data.

Daily plots of selected of SOH telemetry have been posted on the SeaWiFS Mission Operations Web site† since the start of the mission. These plots have been a key data source for analysis of the ACS performance. The plots can be reviewed to check ACS attitude and ancillary data for any GAC span in the mission. A few samples are used to illustrate the typical ACS performance and various anomalies throughout this report, and any data can be viewed at the Web pages as desired.

1.2 Mission Orbit Geometry

The mission orbit characteristics are summarized in Table 2.

Table 2. The OV-2 orbit characteristics are listed.

Parameters	*Description*
Altitude (nominal)	705 km (at start of data collection)
Inclination	98.2° (polar, sun synchronous)
Period	99 min
Equator Crossing	Local noon (descending)
Eccentricity	0.0° (near circular)

A key aspect of the OV-2 orbit is the local noon descending node. This means the spacecraft flies almost directly under the sun each orbit, a key reference position referred to hereafter as the "subsolar point." Because the spacecraft does not generally go exactly under the sun position, the subsolar point passage is defined more exactly as the point where the orbit position is closest to the sun direction. The inclination makes the orbit sun-synchronous, i.e., orbit precession nearly tracks the sun's motion in right ascension to maintain the orbit plane close to the sun line.

Details of the sun position relative to the orbit plane are illustrated in Fig. 2. This figure shows the position of the OV-2 orbit plane relative to the analemma—the track of the sun throughout the year relative to the zenith direction over the Earth's equator at noon. The horizontal axis represents relative longitude in minutes from local noon (with the Earth rotation rate of 4 min per degree), and the vertical axis shows degrees latitude. The middle diagonal line shows the orbit path after launch and initial orbit raising. The orbit plane then drifted slowly westward relative to noon until May 2000, and the diagonal line on the left shows the farthest westward drift. The reason for the slowing drift is that the orbit altitude was initially targeted slightly above the ideal altitude for an exactly sun-synchronous orbit, but the orbit has slowly decayed because of atmospheric drag. Since May 2000, the descending node has drifted eastward at an accelerating rate relative to local noon. The diagonal line on the right shows the orbit path on 31 March 2003. The sun elevation angle from the orbit plane (often called the "β angle") has stayed under 10° throughout the mission thus far. Twice per year, the sun elevation has crossed zero, but most of the year it has been about -3 to $-5°$. The orbit plane will drift away from this proximity to the sun direction unless an orbit adjustment is made, with the nodal drift relative to the sun gradually accelerating. No orbit adjustments are currently planned.

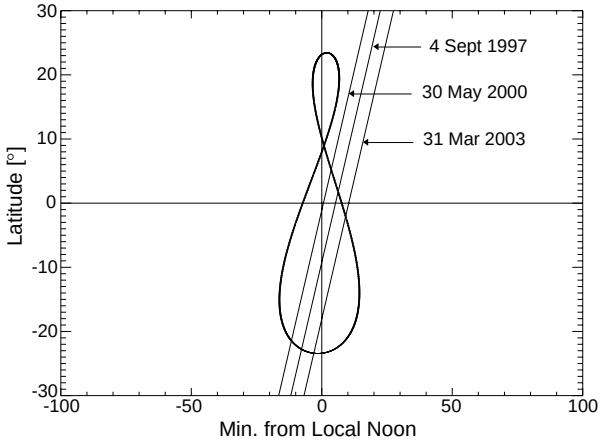

Fig. 2. Orbit plane and the path of the sun relative to local noon.

1.3 Data Collection

The SeaWiFS instrument collects science data on the sunlit side of each orbit, and SOH telemetry data are sent at various data rates according to the types of data collection. For example, while science data are not being collected around the back orbit, the ACS data are usually

† `http://plankton.gsfc.nasa.gov/tlm/`

only sampled once every 3 min for storage on board. This provides very limited visibility into events during that part of the orbit, and this has been a source of frustration for the analysis of various anomalies. Data are stored at the highest rate in the onboard recorder for very limited periods, and are available during limited contacts with ground stations. Understanding the data collection strategy and various sources of data is important to the discussion of the ACS performance throughout this document.

Figure 3 shows the geometry for the routine data collection period for SeaWiFS. Science data were originally collected on board OV-2 for 40 min on the sunlit side of the orbit between ±72.7° centered about the subsolar point. On 24 September 2000, the span was expanded slightly by 120 s earlier for northern data, and by 100 s later for southern data. The angular span, therefore, changed from +80 to −77.5° relative to the subsolar point, with a time span of 43.7 min.

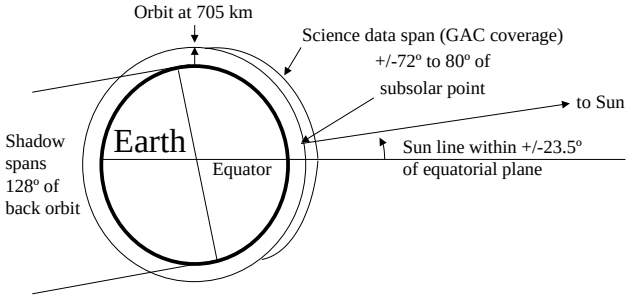

Fig. 3. Data collection span nearly centered about passage of subsolar point.

The full *science data span* shown in Fig. 3 is collected for global area coverage (GAC) data. Because of limited onboard data storage, this is stored on board at a reduced rate for a later downlink. Higher data rates are available for direct high resolution picture transmission (HRPT) broadcast to ground stations (which occurs simultaneously with GAC data collection), and for a limited amount of local area coverage (LAC) data, which is stored on board. For HRPT and LAC data, SeaWiFS image scans are collected at the rate of $6\,s^{-1}$. While ACS samples are taken at $2\,s^{-1}$, new ACS data are updated at the control cycle rate of once every 2 s. This is the rate at which the ACS onboard processor performs all of its basic calculations. For GAC data however, OV-2 stores every fourth SeaWiFS scan (three scans in 2 s) and ACS data only every 10 s, so transient events that affect the control can be missed. This limited sampling of the ACS has often been a source of frustration for visibility into anomalies and detailed ACS behavior. The visibility is much more limited on the night side of the orbit, where ACS data samples are usually only collected every 3 min. The types of data spans and their characteristics are summarized in Table 3. Further aspects of these various types of data collection and the tilt change are discussed in the following subsections.

1.3.1 GAC Scenes

The GAC scenes are the main baseline for the SeaWiFS global sampling. Because of limited space on the onboard recorders, every fourth scan of the SeaWiFS science data (three scans in 2 s) is sampled, and every fourth pixel along the scan from −45 to +45° cross track. The ACS data is taken less frequently than the science data, being sampled once every 10 s, thus, only every fifth ACS control cycle is sampled. Very short transient effects can occur in the missed samples that affect control and only leave secondary evidence of their occurrence. The GAC samples, however, generally provide a good picture of the overall pointing performance during each normal science data collection span.

1.3.2 LAC Scenes

There is very limited memory space on board for storage of the highest rates in the science data, therefore, these scenes are typically very short. They focus on areas of special scientific and calibration interest. They represent a relatively tiny fraction of the overall sampling time, and rarely catch events of interest for ACS review. In one case, however, LAC data collection was scheduled to show the ACS performance at a critical time. This was at a GPS week transition, Sunday 00:00 Coordinated Universal Time (UTC)†, to show the effect of a new limit loaded on board to reduce effects of a transient time tag error (Sect. 4.5.4.)

1.3.3 HRPT Scenes

HRPT data are very useful in providing the maximum data rates over various local areas, as the telemetry is received at various local ground stations. Coverage is not global, but there were many cases of attitude anomalies where HRPT data happened to be available to show the ACS behavior much more clearly. In HRPT data, every control cycle of the ACS is sampled—the 2 s interval during which new attitude estimates and commands are computed. Some tropical region stations have been particularly useful in showing detailed control responses around the subsolar region.

1.3.4 Back-Orbit Data

ACS data are sampled at very low rates in the back orbit—one point every 2 or 3 min—because of the limited recorder space and priority on the spans with science data collection. This has made analysis of numerous ACS performance issues and anomalies in the back orbit difficult, and with the luxury of hindsight, it seems clear that some higher rate sampling of selected data would have been beneficial. All of the housekeeping and engineering data need

† All times given in this report are in UTC unless otherwise indicated.

Table 3. Data span types and characteristics.

Data Span Type	*Typical Duration*	*ACS Sampling Intervals*	*Frequency of Data Collection*	*Storage Notes*
GAC	40–44 min	Every 10 s	Each orbit	Stored on board
LAC	1–4 min	Every 2 s	About 12 short intervals each day	Stored on board
HRPT	5–15 min	Every 2 s	About 5–20 independent ground locations each day	From available ground stations, can be noisy
Back orbit, normal data sampling	57–59 min	Every 3 min	Each orbit	Stored on board, saved only at OrbImage
Back orbit, high rate sampling	20–55 min	Every 2 s	Once a month or less since mid-2000	Stored on board by special scheduling

not have been sampled at the same rate, and if certain selected ACS parameters could have been stored at a higher rate, numerous anomalous events could have been better analyzed. The extra ACS data needed would still have represented only a small fraction of the science data volume, probably costing just a few seconds of LAC data. This could have been planned in the design phase, but after launch, the options for higher data rate collection in the back orbit were limited and required special scheduling.

After many cases where an anamolous event occurred in the back orbit that could not be discerned (because significant attitude motion can occur in between the 3 min between samples), an effort was made to schedule some selected data at higher rates to get additional insight in back-orbit ACS performance. Recorder allocations could be switched with some manual scheduling effort to collect ACS telemetry. A small amount (13 s) of LAC data allocation was removed from selected recorder dump periods, and with this allocation, about 100 min of high-rate ACS data could be collected. The first of these high-rate back-orbit spans was collected in August 2000, and thereafter, samples were collected about once a month for about two years, often targeting times of special interest. The frequency of back-orbit data collection has since been reduced.

1.3.5 Tilt Changes Schedule

The SeaWiFS instrument is tilted 20° to reduce sun glint. Near the subsolar point, the tilt is shifted from aft to forward. The position of the tilt change is staggered north or south of the subsolar point, switching every 2 d. In this way, global geographic coverage is completed every 4 d, with 2 d to fill in between swaths around the equatorial regions, and two more days where the gap from the tilt is shifted north or south. An effect of the tilt change is a slight temporary pitch disturbance of the spacecraft, about 0.2°, as discussed further in Sect. 3.1.7.

1.4 Ground-Computed Attitudes

The ground-determined attitudes produced by the SeaWiFS Project during the navigation processing are referenced several times in this report, as they are generally more accurate and give a better indication of the true spacecraft attitude for certain examples. The algorithms for ground processing of the attitude estimate are discussed in Patt 2002. This processing uses a Kalman smoother with DSS and HS data input. Some filtering of the sensor data is applied to reject spurious data points. The relatively simple dynamics assumption is that the inertial orientation of the spacecraft angular momentum axis (determined by the roll and yaw angles) is expected to change relatively slowly. This tracks well most of the attitude motions for OV-2, but smooths out some of the more rapid motion, such as that due to the pitch tilt change (Sect. 3.1.7), some pitch motion anomalies (Sect. 4.5.5), and cases of unusually high nutation (Sect. 3.1.2). Sensor alignments and biases for the ground processing have been adjusted based on SeaWiFS image data analysis (Patt and Bilanow 1999).

1.5 Mission Timeline Overview

Important milestones for the SeaWiFS pointing performance are summarized in Table 4. These milestones are discussed in other sections as indicated in the "'Notes" column of the table.

2. ONBOARD ATTITUDE

This section discusses the onboard attitude computation for the OV-2 spacecraft. Errors in the computation of the attitude onboard affect the stability of the instrument pointing, because the onboard control responds to transitory errors. This section focuses on the onboard attitude calculation and some of its error sources, while the effects of the errors in terms of the actual pointing motions is discussed in Sect. 3.

2.1 Attitude Sensors

For attitude sensing, the ACS has three two-axis DSSs (DSS-A, -B, and -C); two Earth HSs (HS-A and -B); and two TAMs (TAM-A and -B). The mounting geometry for the DSSs and HSs is illustrated in Fig. 4, and the orientation and coverage range of each sensor is summarized in

Table 4. Major milestones for SeaWiFS ACS adjustments.

Date	*Event Name*	*Notes*
10Oct1997	First HS bias adjust	The first and largest adjustment in the horizon scanners eliminated large (greater than 10°) yaw errors around the subsolar point in each orbit, which triggered significant nutation (Sect. 2.4.1).
7May1998	Torque rod gain change	The reduction of the coil command magnitudes led to less nutation being triggered at the subsolar point, but also larger roll–yaw variation over each orbit (Sect. 3.3.4).
11Nov1998	Single string	OV-2 switch to operation using only ACE-A with the sensor subset noted in Sect. 2.1. Earth oblateness effects became more important for the scanner data.
16Dec1998	BCR switch	The BCR switched autonomously to side B. Later it was noticed that this switch affected the spin axis attitude history in a way that indicated a probable change in the residual dipole (Sect. 3.1.6).
05Dec1999 and 07Dec1999	TAM weighting changes	The weighting of the magnetometer data was reduced to minimize magnetometer error effects around the subsolar point (Sect. 2.6.5).
17Feb2000 and 2Mar2000	Yaw hold region widening	The yaw hold region was widened to span the sun coverage gap at all seasons and also reduce noisy control and resulting attitude disturbances near the subsolar point (Sect. 2.2.4).

Table 5. More detailed descriptions of these sensors are presented in Patt 2002.

Each DSS has a 128×128° FOV. As shown in the figure, DSS-A is mounted on the front of the spacecraft, DSS-B is on the back, and DSS-C is on the top. The DSS-C FOV overlaps 64° with those of DSS-A and -B, and thus, provides some redundancy. With all three sensors operating, sun measurements are always obtained while the spacecraft is in sunlight at nominal attitude.

The boresight of the DSS is defined as its z-axis and the boresight directions in the ACS coordinates are noted in Table 5. The correspondence of each DSS coordinate with the $\vec{Y}$ coordinate is as follows for the nominal mounting: $\vec{Y}$ = x-axis for DSS-A, $\vec{Y}$ = $-x$-axis of DSS-B, and $\vec{Y}$ = y-axis of DSS-C. The remaining sensor axes, y for DSS-A and B, and x for DSS-C, measure the sun's direction around ACS $\vec{X}$-$\vec{Z}$ (nominal orbital) plane and complete the orthogonal triad.

The raw measurements of each DSS are an indication of the sun projection on the sensor's x- and y-axes, however, in this document the converted angular measurements are defined for each DSS as follows:

D_x is the angle of the sun line toward the DSS x-axis, from the plane defined by the DSS y- and z-axes.

D_y is the angle of the sun line toward the DSS y-axis, from the plane defined by the DSS x- and z-axes.

Further details of the DSS data handling are noted in Sects. 2.2 and 2.5.

The HSs have 45° scan cones, and are mounted on opposite sides of the spacecraft. HS-B views toward the orbit normal, and HS-A views opposite this direction. The scanner rotation axes are canted 5° downward from the nominal horizontal plane. The output of each HS is a measure of the Earth chord (angle between horizon in-crossing angle and out-crossing angle) and phase (average crossing angle) relative to an index. Further details of the horizon sensor data handling are noted in Sects. 2.2 and 2.4.

Each TAM measures three components (nominally orthogonal) of the local magnetic field vector. The TAMs are mounted with their z-axes parallel to $\vec{Z}$. The TAM-A coordinates are rotated 90° about the z-axis from the ACS frame, and the TAM-B coordinates are rotated −90°. This results in the TAM-A x-axis being parallel to $\vec{Y}$, and opposite the TAM-B x-axis, while the TAM-B y-axis is parallel to $\vec{X}$, and opposite the TAM-A y-axis. This correspondence is important for understanding the discussion of the TAM bias adjustments (Sect. 2.6). The output of each TAM is a measure of the Earth's magnetic field magnitude along each axis. For measurement references in this document, subscripts have been assigned that correspond to the ACS axes, so T_x, T_y, and T_z refer, for example, to the TAM measurements along the $\vec{X}$, $\vec{Y}$, and $\vec{Z}$ axes, respectively.

It is important to note that only a subset of the sensors (included in Table 5) has been in use since 11 November 1998, following the change to single-string operation.

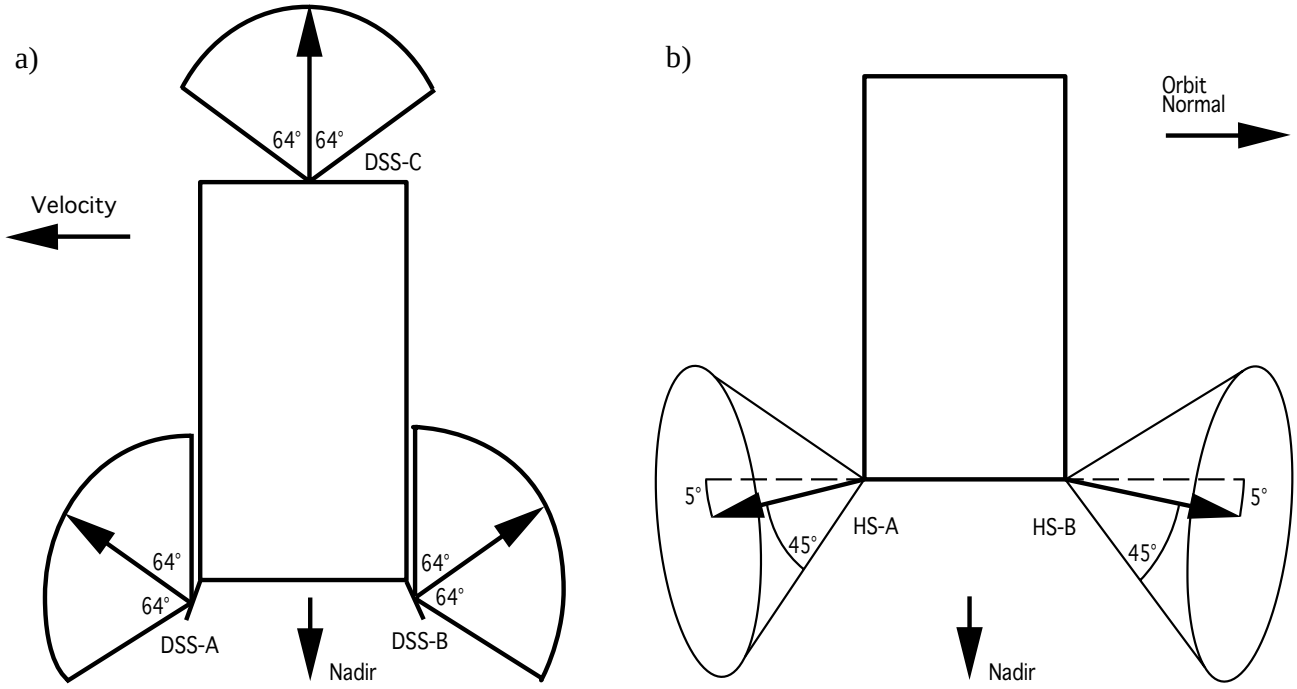

Fig. 4. Mounting geometries for the **a)** sun sensors, and **b)** horizon scanners.

Table 5. The attitude sensors, with a description of each, are summarized.

Attitude Sensor	*Description*		
	FOV	*Orientation*	*Coverage at Nominal Attitudes*
DSS-A†	128×128°	Centered in the $\vec{X}$-$\vec{Z}$ plane‡ 64° from $-\vec{Z}$ toward $\vec{X}$.	Gets coverage from the Earth shadow exit up to the subsolar point.
DSS-B†	128×128°	Centered in the $\vec{X}$-$\vec{Z}$ plane‡ 64° from $-\vec{Z}$ toward $-\vec{X}$.	Gets coverage from the subsolar point up to shadow entry.
DSS-C	128×128°	Centered in the $\vec{X}$-$\vec{Z}$ plane‡ along $-\vec{Z}$.	Coverage centered on the subsolar point.
HS-A†	45° scan cone	Canted 5° toward $\vec{Z}$ along $+\vec{Y}$.	Continuous coverage.§
HS-B	45° scan cone	Canted 5° toward $\vec{Z}$ along $-\vec{Y}$.	Continuous coverage.§
TAM-A†		Components on each axis.	Continuous coverage.
TAM-B		Components on each axis.	Continuous coverage.

† Limited set of sensors used in routine single string operation since 11 November 1998.

‡ Nominal orbit plane.

§ Coverage is continuous except during lunar calibration pitch maneuvers, because a segment of the scan that is normally away from the Earth is blanked.

As of this date, none of the inactive sensors has been reactivated, and barring a failure in one of the primary sensors, the mission will continue in this configuration using only DSS-A and DSS-B, HS-A, and TAM-A. This has had significant implications for the ACS performance.

2.2 Onboard Attitude Estimation

The ACS uses an implementation of the Quest algorithm (Shuster and Oh 1981) for attitude determination in the nominal fine pointing control mode. This is an efficient least-squares optimal estimate where the input is in the form of vector measurements. The sensor measurements are converted to vectors by onboard sensor models. For the DSS, angle measurements are obtained from output counts along the two axes, and it is straightforward to convert those to a unit vector representation in sensor body coordinates.

The HS measures the Earth chord and phase angle. To convert these measurements to a nadir vector for a single scanner, a simple spherical Earth model is used. There is an Earth oblateness correction model on board, but it was disabled early in the mission because of an unrelated onboard software problem. There is also a *dual scanner* solution available on board for roll estimates, which is less sensitive to certain Earth model parameters, and was used before the change to single-string operation in November 1998.

The TAM direct measurements are in vector form. The magnetic field strength measurements are obtained from the raw counts using different scale factors and calibration curves on each axis, and with coarse and fine calibration ranges (Sect. 2.5). Various sensor alignments, scale factors, and bias terms can be uplinked to the spacecraft by table loads.

Table 6. Sensor data weighting and relative accuracy in degrees.

Sensor	*Weighting*	*Accuracy*
DSS-A†	10000	0.01
DSS-B†	10000	0.01
DSS-C	10000	0.01
HS-A†	50	0.14
HS-B	50	0.14
TAM-A†	0.1 (until 5 Dec 1999)	3.16
	0.01 (5–7 Dec 1999)	10.0
	0.0001 (after 7 Dec 1999)	100.0
TAM-B	0.1	3.16

† Subset of sensors used for single string operations after 11 November 1998.

2.2.1 Measurement Weighting

Each vector measurement is weighted in the Quest algorithm according to a nominally expected sensor accuracy. The DSS measurements have the highest weighting, the HSs have the next highest, and the TAMs have the lowest weighting as summarized in Table 6. The TAM data weighting was reduced further after December 1999, as discussed in Sect. 2.6.6.

2.2.2 Dominant Vectors in Sun and Shadow

The weighting scheme means that in general, the most accurate sensors available dominate in fixing a reference axis in the attitude computation. When two of the same type of sensor are available, the average of their measurements is effectively used; thus, the orbit has two distinct regions:

1. In sunlight, DSS measurements fix one axis of the spacecraft along the reference sun direction. The rotation of the spacecraft about the sun line is then essentially computed from the HS measurements. (A small contribution from the TAM measurements is discussed in Sect. 2.6.6.) The DSS measurements dominate in giving pitch information, because the sun vector stays close to the spacecraft pitch normal plane. The DSS measurements also give yaw information, although near the subsolar point, uncertainty in yaw is strongly coupled with uncertainty about the sun line, which depends on the HS measurements as well. Roll information is mainly provided by the HS measurements.
2. In Earth shadow, the HS measurements dominate in fixing the $\vec{Z}$ (nadir) axis of the spacecraft, and the rotation about $\vec{Z}$ is given by the TAM measurements.

2.2.3 Poor Geometry Effects

When two vectors dominate in determining attitude, there is a well-known problem when these vectors have a very small angular separation. A simple formula for the error, η, in rotation about the primary reference vector, $\vec{V}$, as a result of position errors in the secondary vector, $\vec{W}$, can be given as

$$\eta \;=\; \frac{\delta_W}{\sin(\Gamma)}, \tag{1}$$

where δ_W is the angular error in $\vec{W}$ in the direction normal to the $\vec{V}$-$\vec{W}$ plane, and Γ is the separation angle between the vectors $\vec{V}$ and $\vec{W}$.

The errors are magnified greatly when the two vectors are nearly parallel or antiparallel. OV-2 has particular problems with this type of poor geometry effect in the following situations:

- First, where the orbit takes the spacecraft near the subsolar point, the sun–nadir geometry is poor, with vectors nearly opposite in the middle of the science data collection span. The effect of this on the yaw accuracy has been a significant concern, because the

DSSs and HSs are the most accurate sensors and heavily relied upon. A strategy taken to ameliorate this effect is discussed in the next subsection, but imperfect adjustments for this problem were the main issue for mission in-flight ACS adjustments.

- Second, the TAM–HS geometry is poor in the polar regions, where the magnetic field aligns with the nadir direction near the Earth's magnetic poles. This is only of concern during the Earth shadow periods, where science data are not collected. It means there are often very erroneous yaw attitude solutions on certain orbits during the shadow span, but this rarely results in a significant disturbance being carried over to the daylight science spans. When the TAM–HS geometry is particularly bad, i.e., the vectors are within 2.5°, the onboard three-axis attitude solution is flagged as invalid (Sect. 2.2.5).

2.2.4 Yaw Hold

It was recognized prior to the launch of OV-2, that the sun and nadir vector separation would be a problem near the subsolar point. The solution implemented in the onboard software provides what is known as a *yaw hold* region. (This is also referred to as the *yaw smoothing* region in the onboard software.) Within a period (table-adjustable) around the subsolar point, the yaw value is held at a constant value. Moreover, a simple lag filter smooths the computed yaw values, so that a smoothed value is held during the subsolar passage. The prelaunch value selected for the proximity to the subsolar passage was 3.0°. Based on the specified worst case accuracy of the HS of 0.05°, the worst errors just before reaching the yaw hold region were expected to be 1.0°, which seemed tolerable. As postlaunch data proved, however, the HS errors could be larger (Sect. 2.4.1).

About two years into the mission, the yaw hold region was significantly widened to help minimize various disturbances close to the subsolar region. Disturbances resulted from a gap in DSS coverage near the subsolar point after the switch to single ACE, and from other errors and noise effects near this poor sun–nadir geometry region. The proximity to the subsolar region for the yaw hold was expanded from the initial ±3.0° to ±10.0°.

The dates of the yaw hold region extension are summarized in Table 7. Onboard table-uploaded values define the yaw hold region, which are noted here as angles in degrees around the orbit with respect to the subsolar point as estimated on board. (The units used in the onboard loads are actually shifted by 180° and referred to as *sun latitude*.) Note that if the onboard estimate of the expected subsolar position is in error, as has happened, for example, because of GPS anomalies resulting in onboard orbit errors, then the region of yaw hold will not be centered well about the actual subsolar point. Any offset, however, is normally quite small. The reasons for the expansion of the yaw hold region are discussed further in Sects. 2.6.5 and 3.4.2.

Table 7. Yaw hold region extent is listed, where start and end entries are the orbit angles with respect to the subsolar point, as estimated on board.

Date	*Start*	*End*	*Notes*
04Sep1997	−3°	+3°	Initial on-orbit values
17Feb2000	−7°	+7°	First extension
02Mar2000	−10°	+10°	Second extension

2.2.5 Minimum Vector Separation

The three-axis attitude solution is only considered valid and used on board when the sun, nadir, or magnetic field directions meet a minimum separation criteria, which is nominally set (by table-loaded value) at 2.5°. This condition is not an issue near the subsolar point because HS and TAM data are both available there with good separation angles. In practice, this condition only causes invalid attitudes when OV-2 passes very close to the magnetic poles while in the Earth's shadow. This is the region where the magnetic field lines are almost along the nadir direction, and without a sun vector measurement available, an accurate attitude cannot be obtained. The magnetic pole is tilted about 11° from the Earth's axis. As the Earth rotates under the orbit plane, the magnetic axis passes through the orbit plane twice each day. One of these magnetic pole overflights, however, occurs closer to the daylight side of the orbit. The sun is visible for about 26° past the terminator crossing. At the solstices, the sun is 23.4° from the Earth's equator, so the shadow transition occurs just about 3° from the extreme latitude position over the winter pole. The result is that in December–January in the north and June–July in the south, the magnetic pole overflight slips into the shadow period. The minimum vector separation condition only lasts about 3 min or less a few times a day, typically around 02:00–05:00 UTC.

When the three-axis attitude solution becomes invalid, the control mode changes temporarily to coarse pointing mode. Sometimes a small control disturbance is induced by the event, which can lead to a small level of nutation (Sect. 3.1.2) at the start of the GAC span when occurring at high latitudes. A sample of this type of event is described in Sect. 2.6.2.

2.3 Error Signature Analysis

An approach was devised for determining the sensor biases for OV-2, based on analysis of the characteristic effect of biases (also known as the error signature) in the regions of poor attitude solution geometry. The error signature effect, which shows up in yaw, is illustrated and explained with a simple plane geometry model. This approach eliminated the need to develop any special software for the bias determination, as tools were not in place for this purpose prior to launch. The error effects were also verified with the PC simulator.

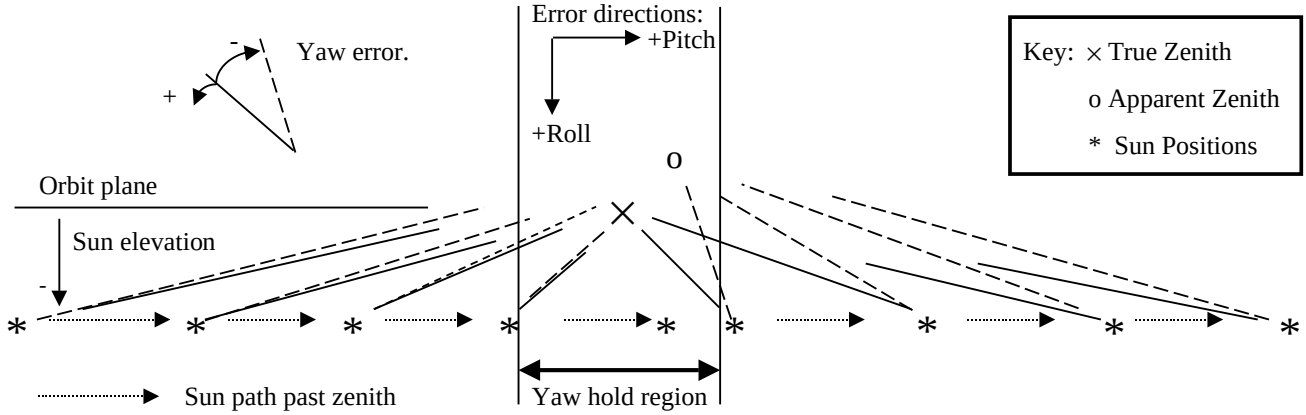

Fig. 5. Schematic diagram of yaw errors generated by bias.

2.3.1 Plane Geometry Model

A simple schematic of how a bias error can cause a yaw error around a subsolar passage is illustrated in Fig. 5, where the sun vector passes close to the zenith vector. (One can think of lying on top of the spacecraft looking at the zenith and seeing the sun go by in the path indicated.) The HS data determine the apparent zenith (opposite the computed nadir vector), as displaced from the true zenith. The yaw error is the rotation from the true sun–zenith line to the apparent sun–zenith line. This error changes rapidly, as shown, as the subsolar point is passed.

It is important to note that the relative biases between the sun and Earth measurements drive the errors. The errors illustrated here are in the apparent zenith direction, which is derived from the HS measurement, but errors in the apparent-sun versus true-sun direction have symmetric effects. This was demonstrated with simulations of biases for either or both. With both vectors biased the same (relatively small) amount, no significant yaw errors result at the subsolar passage.

A spreadsheet implementation of this plane geometry model was used to generate the simulation of the post-launch biases shown in Fig. 6. This model included the yaw smoothing filter and yaw hold (Sect. 2.2.4).

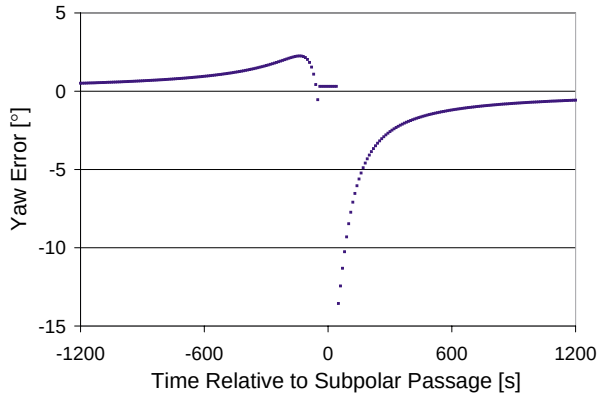

Fig. 6. Computed yaw from simple plane geometry model of biases.

The bias parameters in the above representation were chosen to roughly simulate the effects of the initial on-orbit effects (Sect. 2.4.1).

2.3.2 Subsolar Yaw Signatures

It proved useful to recognize the effects on yaw error of various bias combinations. Pitch bias effects have the same sign before and after the subsolar point, while the effects of roll biases have opposite signs. The other important parameters in the plane geometry model are the size of the yaw hold region and the elevation of the sun in the orbit plane. The effects of sun elevation changes can also be seen from the schematic (Fig. 5); pitch bias effects tend to diminish and disappear as the sun elevation approaches zero, while roll bias effects get slightly larger. In addition, pitch bias effects change sign as the sun elevation crosses zero.

The effects for various combinations of pitch and roll biases, and a $-3.0°$ sun elevation, are shown in Fig. 7. The characteristic shapes provide the error signatures.

The yaw error amplitude scales linearly with bias amplitudes over the range of tuning, so if the error signature shape is recognized, the bias adjustment amplitude can be readily estimated. In addition, using the plane geometry model, the pitch and roll biases can be determined from the error amplitude before and after the yaw hold. This was done to select bias values to try in the simulations; it worked very well initially, but the effectiveness of this approach broke down at finer levels of tuning as discussed below.

2.3.3 Error Signature Caveats

A limitation in the above approach is that it assumes the biases are the same before and after the subsolar point. This was not always the case for two reasons:

1. The DSS measurement biases were somewhat different just before and after the subsolar point as different sensors come into use. This had notable

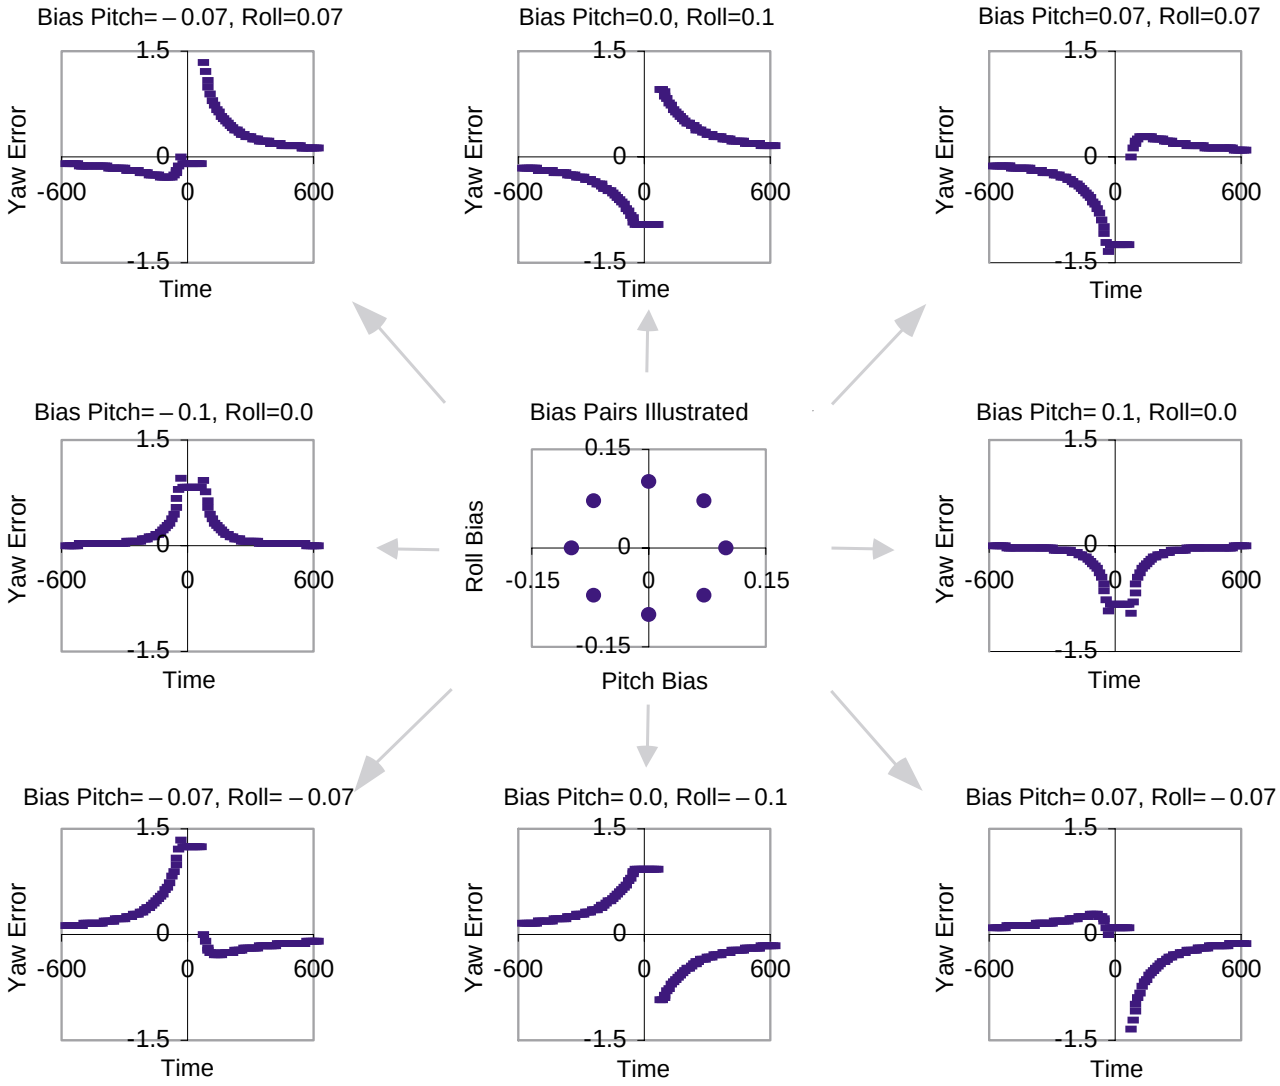

Fig. 7. Yaw error signatures at subsolar passage for various pitch and roll biases. The center plot is a map of the bias pairs whose signatures are illustrated around the periphery. Each solid circle (in the center plot) represents the bias pair associated with the signature, as directed by the gray arrows.

effects, as discussed in Sect. 2.5.3, and was relatively significant before some DSS alignment issues were addressed.

2. The HS errors evolve somewhat through the subsolar passage because of Earth oblateness effects that are not modeled on board (Sect. 2.4.3). This was not a significant effect for the yaw signature, but was noticeable after the yaw hold region was widened especially in the pitch bias differences at the beginning and end of the yaw hold (Sect. 2.4.5).

Partial- or half-signatures, refer to the indicated effect just before or after the subsolar passage. Similar effects, however, can be obtained from pitch or roll in the half-signature, so the bias combination is not clearly indicated without some constancy of effect assumed.

2.3.4 Magnetic Pole Signatures

A part of this same type of error signature can be seen in the Earth shadow period over the poor geometry region over the magnetic poles. The signature would vary significantly from orbit to orbit over each day as the magnetic pole position rotated with the Earth relative to the orbit. Only a part of the subsolar-type signatures are seen, because the spacecraft would come back into sunlight before,

or very soon after, the actual magnetic pole position was reached. As soon as the sun vector is added to the onboard attitude solution, the poor geometry conditions are eliminated. The magnetic pole error signatures are discussed in Sect. 2.6.2, with examples from flight data.

2.4 HS Bias Effects and Adjustments

This section discusses the HS onboard bias adjustments. Biases in the chord and phase angle measurements were chosen from a variety of possible HS model parameters that could be adjusted (Sect. 2.4.1). An initial set of adjustments was made shortly after launch, and these were revised for the start of single-string operation. Since that time, regular adjustments have been made to compensate for seasonal effects. Two factors that contribute to the need for bias adjustments are the lack of correction for Earth oblateness and the onboard calibration table. These topics are discussed below.

2.4.1 Initial Bias Adjustments

Following launch, very large errors in yaw appeared around the subsolar point. Routine science data collection was started on 18 September 1997, and typical onboard yaw estimates on the next day are shown in Fig. 8. The signature of the simulated error effects from Fig. 6 are clearly visible (along with an oscillation following the subsolar point from nutation, which is discussed in Sect. 3.1.2).

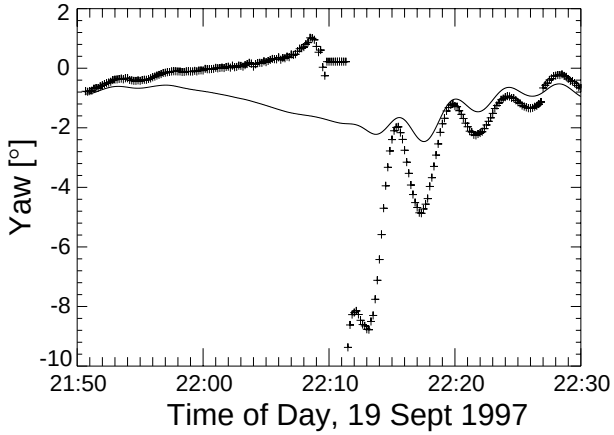

Fig. 8. Onboard-estimated and ground-computed yaw during one of the GAC spans on 19 September 1997. The onboard-computed yaw is denoted with the + symbol, and the ground-computed yaw is denoted with a solid line.

Biases were estimated based on these subsolar error signatures as discussed in Sect. 2.3.1. The biases that modeled the initial on-orbit error corresponded to 0.75° in pitch and 0.68° in roll; however, different parameters were adjusted on board. Model parameters that can cause a pitch or roll bias include sensor alignment parameters, HS scan cone angles, and horizon triggering heights. In addition, because it was the average of two HS angles that determined the effective nadir vector used on board, biases could be applied to any one sensor or both to achieve the same effect. The desired bias must be doubled when adjusting only one sensor in order to compensate for sensor averaging. For the first load, only HS-A was adjusted (somewhat arbitrarily), and the parameters adjusted were the phase and chord measurements. The phase has nearly a 1-to-1 correspondence with pitch, while the chord changes by about 2.2° per degree of roll for the nominal geometry on OV-2; thus, the actual loads were for 3° in chord and 1.5° in phase for HS-A.

After the first adjustment was made, the yaw error magnitude near the subsolar point was reduced dramatically, as illustrated in Fig. 9 (note the scale differences from Fig. 8). There were smaller but still significant yaw errors present, with a different error signature, indicating the effects of the remaining errors.

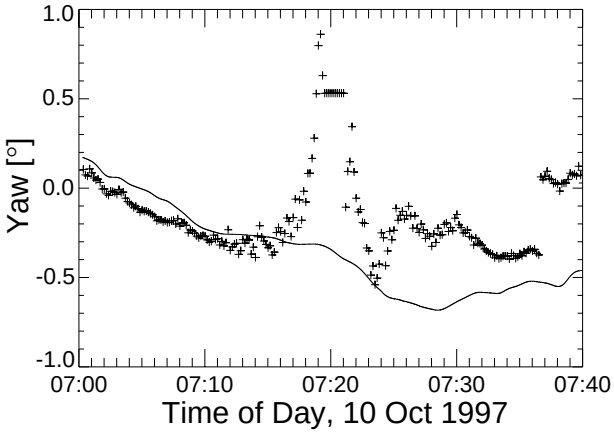

Fig. 9. Onboard-estimated and ground-computed yaw during one of the GAC spans on 10 October 1997. The symbols used are the same as in Fig. 8.

The second adjustment was applied to HS-B (again, arbitrarily) and amounted to approximately 0.05° in pitch and 0.16° in roll. A summary of initial bias loads is given in Table 8.

The next adjustment to the HS was needed in conjunction with a DSS adjustment. The error effects so far have been presented assuming the DSS is correct and the HS is biased, whereas actually it is the relative bias between them that produces an error signature at the subsolar point. As the simulator indicated, when adjusting the DSS-C alignment about the $\vec{Y}$ by 0.2°, the HSs need to be adjusted for a corresponding pitch bias, or else subsolar error signatures would be generated.

2.4.2 Single-String Adjustments

Additional bias adjustments were needed before OV-2 switched to using the single ACE and HS-A only in November 1998. Whereas the biases had been tuned so that the

Table 8. HS bias adjustments from launch through 1998.

Time *Date*	*SDY*†	*HS-A Chord*	*HS-A Phase*	*HS-B Chord*	*HS-B Phase*	*Adjustment Notes*
10Oct1997	283	−3.0°	−1.5°			Initial adjust from zero values.
20Nov1997	324			0.7°	0.1°	First fine tuning.
05Dec1997	339				−0.3°	Accompanied DSS-C pitch alignment adjustment.
06May1998	126				−0.5°	Related to seasonal changes.
28Oct1998	301	−0.6°		3.1°		Adjustment for valid single scanner chords.
30Oct1998	303		−1.8°		−0.6°	Adjustment for valid single scanner phases.
01Nov1998	305		−1.2°		0.0°	Readjust for single scanner phases.
13Nov1998	317	−0.7°				Postsingle-string tuning.
16Dec1998	350		−1.3°			Tuning with DSS-B adjust.

† Sequestial Day of the Year.

average nadir vector was accurate, additional adjustments were needed to ensure that each scanner would give an accurate vector on its own. Several issues had to be dealt with in this process, referred to as *balancing* the HS biases for single-string operations.

One problem was with visibility into what the attitude results would be with each sensor used separately. With both sensors on, only their total result was visible, and simulations showed that the same result could be obtained with a range of values. The onboard system had several complex steps that were quite different from the ground data processing stream, so they could not be compared directly. There were no tools to feed the raw measurements through the onboard simulation and simulate single-sensor results, however, two derived parameters included in telemetry were useful: the *local nadir roll* and *local nadir pitch*. These are computed on board to provide HS-specific pitch and roll measures to be used in the coarse pointing and despin modes. The documentation for these parameters stated that they would be based on the dual scanner solution when both scanners were available, but in fact, the code was implemented so that if HS-A was available, it was used first alone. This provided a straightforward way to tune the proper biases for HS-A, and by maintaining the overall bias balance, also get the correct tuning for HS-B.

2.4.3 Earth Oblateness Effect

As noted earlier, an onboard model for Earth oblateness was implemented but not used. This was because of a coding error in the background task where the oblateness correction was computed, that could cause computations to freeze up in the presence of very large onboard orbit errors. While the error-producing situation was considered unlikely, it did create some risk that was considered unacceptable for loss of attitude in anomalous conditions.

This oblateness effect causes some errors in the onboard computed pitch and roll, on the order of a few tenths of a degree. The effects on roll cancel out when two HSs are used, however, the effects become relevant for single-string operation when only one HS is used. With only HS-A in use, the modeled effect is to indicate a smaller roll by about 0.4° at the poles relative to the equator.

In pitch, the effect is about the same whether one or two HSs are in use. The maximum pitch error occurs at mid latitudes, with zero error at the equator and extreme latitudes. For the typical GAC span, the modeled effect has higher onboard pitch estimates by about 0.2° in the northern latitudes and lower pitch estimates by the same amount in the southern latitudes. Evidence from ground calibration of an Earth radiance model showed the actual oblateness effect for OV-2 seems slightly larger than the model values (Patt and Bilanow 2001).

Oblateness is, most probably, the dominant effect that shows up in the differences between onboard and ground-computed roll and pitch values in the single-string operating mode, as illustrated in Fig. 10. Other sensor alignment, calibration, and modeling differences have other subtle effects. The onboard errors and motions in pitch around the middle of the span are due to temporary pitch errors and the SeaWiFS tilt change (Sect. 2.4.5).

An obvious choice for OV-2 was to have the error be as small as possible, particularly at the subsolar point, and biases were chosen accordingly. As the subsolar point moves north and south with the season, however, the errors produced by oblateness at the subsolar point do change. This effect contributed to a need for seasonal adjustments of the biases (Sect. 2.4.5).

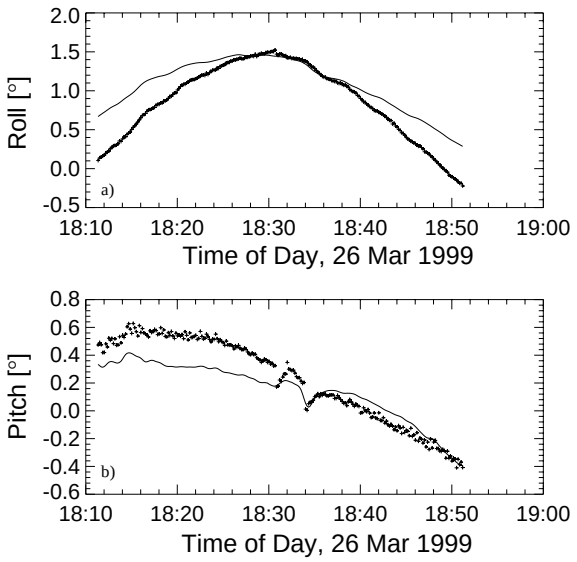

Fig. 10. Onboard- (+) and ground- (line) computed values for **a)** roll, and **b)** pitch for the GAC span on 26 March 1999. Differences are influenced by Earth oblateness effects not modeled onboard.

2.4.4 Onboard Calibration Curve

There are measurement-dependent adjustments to HS data that was implemented in tables onboard the spacecraft. Figure 11 shows one of these tables of corrections. There are four such tables onboard, one for the in-crossing angle edge and one for the out-crossing angle of each of the two HSs. Initially for ground processing at the SeaWiFS Project, it was decided to fit these with a polynomial, because it was suspected that the noise-like residuals from a low order polynomial fit were just the result of measurement errors in the calibration. It was difficult to obtain information about the original calibration. No reason was seen based on the hardware design to have these types of corrections on the raw HS measurements. One source, however, indicated that the onboard table was implemented to take out other table-type adjustments which were originally hard-coded in the sensor electronics. In this sense, the noise-like adjustments are offsetting the effects of other adjustments. For the normal operating range of sensor measurements, the variation in the table adjustments are small.

Because the HS phase is generally very steady, and the HS chord varies over a limited range, the sensor generally operates within one cycle of the curve shown in Fig. 11. In addition, it is usually at the same place on the curve each orbit, making the curve's effect hard to separate from other orbit frequency effects. Efforts were undertaken to try to determine if these measurement-dependent corrections could be estimated from the flight data, e.g., using lunar calibration data where the HS measurements change over a broad range. Eventually, this onboard calibration was also implemented in ground processing based on analysis of seasonal variations in accuracy.

A possible effect from the errors in this calibration curve is to cause some of the variability in the subsolar error signature throughout the day. The spin axis attitude currently seems to repeat an average trajectory over the course of the day, which varies within 1–2° from orbit normal. This is a relatively small change in the operating point on the curve, but it is estimated that this could affect roll by a few hundredths of a degree, and thus, have a few tenths of a degree effect on the subsolar yaw signature. More of the subsolar point error variability, however, is now thought to come from TAM errors and other sensor noise sources.

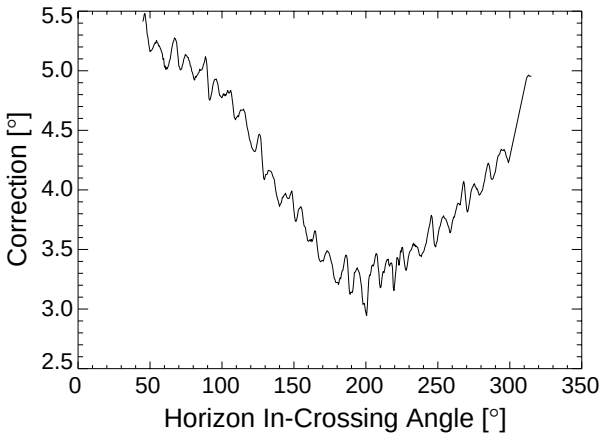

Fig. 11. Onboard correction from the calibration table for the HS-A in-crossing angle.

2.4.5 Seasonal Adjustments

It was found that the HS biases needed adjustment with season. One part of this adjustment could be attributed to Earth oblateness effects (Sect. 2.4.3) coupled with movement of the subsolar point 23.4° north and south of the equator. Other effects resulted from the interaction of the yaw hold with general seasonal variation in typical attitude history at the subsolar point. With the expansion of the yaw hold region on 17 February 2000, and again on 2 March 2000, the potential for a larger discontinuity in onboard roll at the end of the yaw hold region was apparent. At times of the year when the yaw angle changes most rapidly at the subsolar point, a proper bias avoids large jumps in yaw values at the end of the yaw hold, which can excite control disturbances (Sect. 3.5.2).

The seasonal adjustments are summarized in Table 9 for all adjustments since the switch to single-string operations using only HS-A. In the first year of single-string operation, the chord bias was adjusted between −0.6 in March to −1.0 in June, and back to −0.7 in October. The phase adjustment varied between −1.3 in December and −1.0 in June. After the switch to a widened yaw hold (17

Table 9. Seasonal HS bias adjustments in a single-string operation.

Date	*SDY*	*HS-A Chord*	*HS-A Phase*	*Adjustment Notes*
11Nov1998	315	−0.6	−1.2	Values at switch to single string.
13Nov1998	317	−0.7		Post-single string tuning.
16Dec1998	350		−1.3	Tuning with DSS-B adjust.
26Mar1999	85	−0.6	−1.2	
20Apr1999	110		−1.1	
04Jun1999	155	−0.8	−1.0	
17Jun1999	168	−1.0	−1.0	
03Sep1999	246	−0.9		
08Sep1999	251		−1.1	
26Oct1999	299	−0.8	−1.3	
28Oct1999	301	−0.7	−1.2	
08Dec1999	342		−1.25	
10Dec1999	344		−1.3	
25Feb2000	56		−1.2	
24Mar2000	84	−0.6		First adjust following widened yaw hold.
25Mar2000	85	−0.8		Correction for previous adjust.
17Apr2000	108		−1.0	
18Apr2000	109	−1.0		
12Oct2000	286	−0.8		
20Oct2000	294		−1.1	
14Nov2000	319		−1.2	
20Mar2001	79		−1.1	
08May2001	128	−1.0	−1.0	
09Oct2001	282	−0.8		
11Oct2001	284		−1.1	
04Dec2001	338		−1.2	
05Apr2002	95		−1.0	
24Apr2002	114	−1.0		
10Oct2002	283	−0.8	−1.1	
19Nov2002	323		−1.2	
21Apr2003	111	−1.0	−1.0	

February and 2 March 2000), the first chord bias adjustment was by 0.1. Because the widened yaw hold helped to reduce the level of subsolar disturbances, fine-tuning bias adjustments were done somewhat less often after that change. Since 2000, the chord values have typically been set to −1.0 in April and −0.8 in October. The phase adjustment varied between −1.0 in April–May and −1.2 in November–December.

Initially, all bias adjustments were checked with PC simulations, but later, after the seasonal pattern was understood, adjustments were made based on the error patterns in the flight data. The simulations did not exactly reproduce the flight data values, but the error signatures were very closely represented. Values could be chosen to get the typical subsolar error signature seen in flight, and the change relative to a minimal error simulation evaluated. The reverse of the change needed to simulate the error was loaded on board to remove the error.

After the yaw hold widening, the previously minimal error condition showed a larger jump in yaw, because the actual yaw value would change significantly during the hold period. In this case, an appropriate scanner chord bias was added so there would be a smaller discontinuity after the hold period. This revised approach to the bias adjust was adopted because the discontinuity could induce a control disturbance and nutation (Sect. 3.5.2). It was also observed, however, at certain times of the year the discontinuity often provided an impulse that actually helped reduce residual nutation (Sect. 3.5.2).

Figure 12 shows the yaw and pitch around the subsolar region before and after the adjustments with the widened yaw hold. As noted previously, a chord bias was added to remove the discontinuity after the yaw hold. This bias also had the effect of increasing the yaw error after the yaw hold region, however, this was necessary to reduce the discontinuity.

Pitch discontinuities in the subsolar region result from a gap in DSS coverage (Sect. 2.5.1). Pitch is essentially

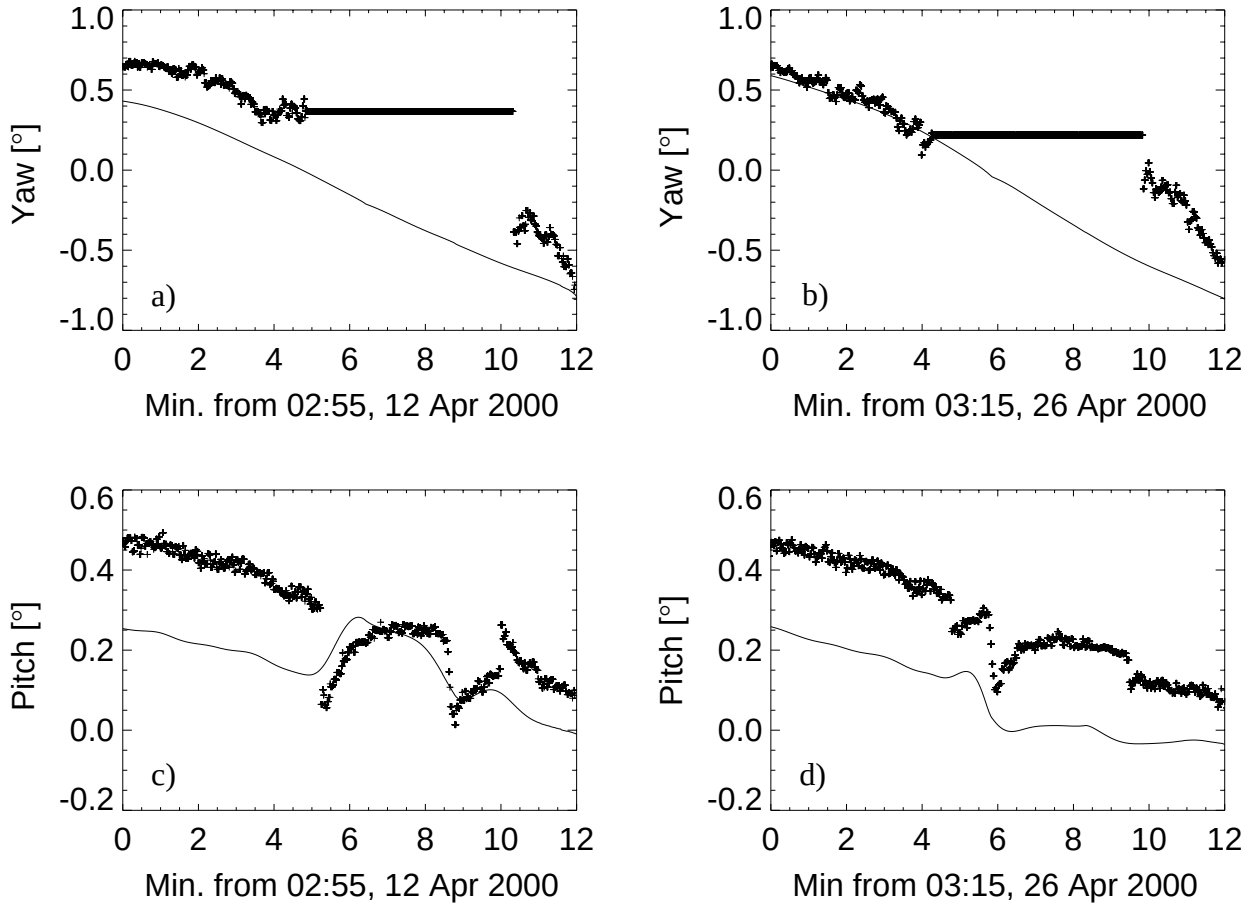

Fig. 12. Onboard attitude (+ symbols) sampled near the subsolar point at high data rates (2 s per sample) before and after HS bias adjustments with the widened yaw hold region: **a)** yaw before adjustments, **b)** yaw after adjustments, **c)** pitch before adjustments, and **d)** pitch after adjustments. The adjustments were 0.2° in chord (which mainly affects yaw) and 0.2° in phase (which mainly affects pitch in the sun coverage gap). These biases were loaded on board in the intervening period between HRPT data spans on 12 and 26 April 2000. Ground-computed attitudes (lines) are also shown.

determined from the HS data inside the gap, and from DSS data outside the gap; thus, two discontinuities show up in the span shown in Fig. 12, along with the tilt change event. The highest available data rate, 2 s, helps show the different character of the tilt change event, which takes about 13 s.

Because of Earth oblateness effects on the HS, the error in the onboard pitch changes by about 0.1° through the DSS coverage gap; thus, pitch discontinuities cannot be eliminated with a fixed phase bias adjustment. The jumps are roughly minimized for the 27 April 2000 sample, which splits the difference between the ideal bias for gap entry and gap exit. On 13 April, before the phase bias adjustment, the pitch computed in the coverage gap was consistently lower. Adjustment in the HS phase bias to follow the seasonal changes in Earth oblateness effects are done not more than four times a year, because the pitch motion caused by discontinuities less than about 0.2° are not a great cause for concern.

2.5 DSS Bias Effects and Adjustments

A summary of the onboard DSS alignment adjustments is given in Table 10; the issues analyzed with each adjustment are discussed in the following subsections. The adjusted parameters were the onboard alignment angles, which may be understood as a z–x–y rotation sequence from the sensor coordinates previously described in Sect. 2.1. The first rotation is a twist about the boresight. The final two rotations can be interpreted as adjustments about the $\vec{Y}$axis, and elevation toward the $\vec{Y}$axis, with the order as indicated depending on the sensor mounting. Section 2.5.1 describes how the FOVs of the DSSs overlap, and the sequence of coverage transitions; the transition effects in the presence of relative alignment errors is also discussed. The FOVs were limited by an onboard adjustment initially to better see the transition errors, and were limited again later in the mission (Sect. 2.5.2) to avoid shadowing effects. Eventually, alignment angles computed

Table 10. Sun sensor alignment adjustments. The order of rotation is z–x–y from the sensor coordinates (Sect. 2.1) in each case, however, the order of the table headers vary because of the different DSS mountings: *twist* is a rotation about the sensor boresight; *pitch* moves the boresight (which is near the $\vec{X}$-$\vec{Z}$ plane for each DSS) about the $\vec{Y}$ axis, which affects pitch measurements directly; and *elevation* moves the boresight toward the $\vec{Y}$ axis, which mostly affects the apparent sun elevation from the orbit plane. Blank entries indicate there was no change made in the given parameter.

Date	*DSS-A*			*DSS-C*			*DSS-B*			*Notes*
	z *Twist*	x *Pitch*	y *Elev.*	z *Twist*	x *Elev.*†	y *Pitch*†	z *Twist*	x *Pitch*	y *Elev.*	
Prelaunch	0.255	.025	.101	0.0	0.0	0.0	−0.22	−0.134	−0.053	DSS-C not measured
09Nov1997							−0.35		0.7	Yaw adjust of B
05Dec1997						−0.2				Pitch adjust of C
05May1998	−0.045							0.066		Twist to A; pitch to B
16Dec1998								0.22		Single string; pitch to B
30Mar1999							−0.16	0.30	0.57	Based on ground align

† Note the reversed headers as explained in the caption.

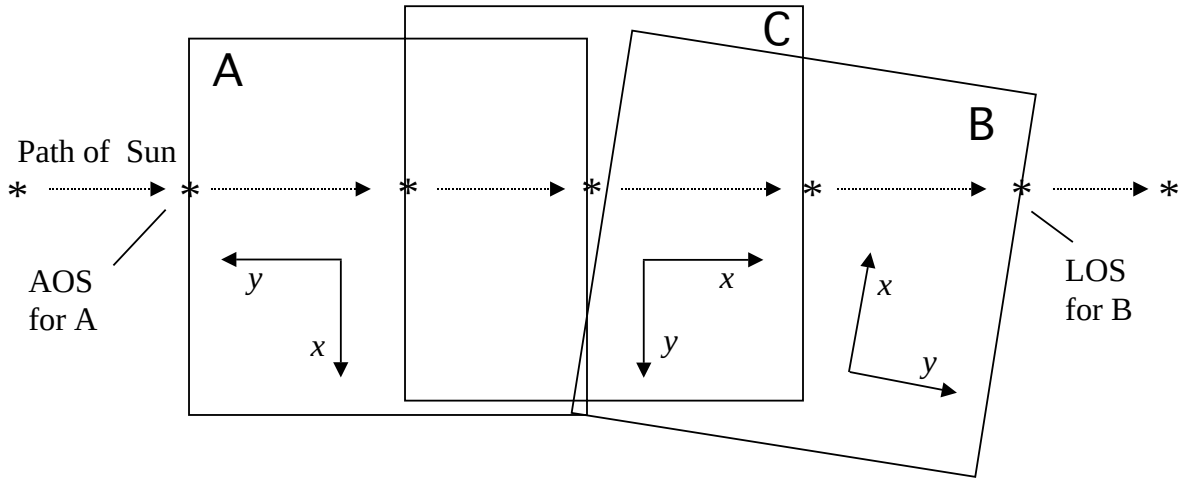

Fig. 13. DSS FOVs schematic illustrating the overlaps, typical sun path, and sensor coordinates for DSS-A, -B, and -C.

for use with ground processing were used as a basis for onboard adjustments. The DSS calibration curve adds small corrections whose accuracy has not been verified.

2.5.1 Overlaps and Gaps

A helpful way to view the DSS FOV overlaps is to lay them out flat, unwrapping about the $\vec{Y}$ in a projection on the sky as shown in Fig. 13. For this body-centered view, looking out toward the sky as in Sect. 2.3.1, the apparent sun direction moves from left to right across the middle of each sensor FOV. The discussion below refers to the acquisition of signal (AOS) and loss of signal (LOS) for the sun in each sensor. The sun encounters the sensors in the order A, C, and then B. DSS-C overlaps both DSS-A and DSS-B for approximately 64° of sun vector arc. Although the DSS-A and DSS-B FOVs would nominally just meet, in practice, a gap was observed between them that varied with the sun elevation (Sect. 4.1.2). That gap, and the sensor misalignments, are exaggerated for the illustration. The sensor D_x and D_y coordinates are labeled as shown.

2.5.2 Software FOV Limit

A DSS FOV limitation was first introduced on 17 December 1997. That change was allowed to lapse after an SCM reset on 30 June 1998, but was reintroduced on 15 March 2000. The original motivation was to help distinguish the sensor AOS–LOS transition effects, especially around the subsolar point, where several FOV changes take

place in close succession. This also served to remove concerns about the ends of the DSS calibration curves and any FOV edge effects (Sect. 4.1.1). The FOV limitation was reintroduced after the realization that shadowing effects were affecting data at the FOV edges in January–February, as discussed further in Sect. 4.1.2.

It was first thought that the FOV limits could be implemented with an onboard table load, however, it was performed with a software poke. There are parameters in the tables for limiting the FOV size used on each axis, but it was discovered that because of a software error, those parameters are not actually used. OSC flight software programmers (who at that time were also working on a GPS data filtering patch) identified an alternative means for accomplishing the FOV limit by using a validity check in the software that tests the computed sun vector in the sensor frame. In the onboard code, this limit is applied as follows:

$$D \;=\; \sqrt{1 + \tan^2 D_x + \;\tan^2 D_y}, \qquad (2)$$

where D_x and D_y are the DSS angles.

The original check flagged the DSS data as invalid if the value of D was less than 0.99 or greater than 20. The upper limit was changed to 1.75, and this had the effect of limiting the effective sensor FOV to a cone centered on the boresight with a 55° half-angle. The FOV limitation did help separate the errors during the first few months of the mission, although it was still difficult to distinguish AOS–LOS transition jumps from the other noise sources near the subsolar point.

2.5.3 Transition Effects

A discontinuity in the average sun direction computed onboard results at each acquisition and loss of data from each of the sensors. These discontinuities provided the first obvious clues of the DSS alignment errors. For example, the jumps in yaw angles on board near the end of each of the GAC spans early in the mission (e.g., Figs. 8 and 9) were due to the LOS of DSS-C data.

The approach taken for adjusting the onboard alignments was initially a stepwise process, in which the most obvious errors indicated by discontinuities were progressively removed. This process was partially guided by the SeaWiFS ground processing results, but it mainly relied on effects at the sensor AOS and LOS. The goal for the initial adjustments was to have the sensors agree at the AOS and LOS points. Fixing errors at these so-called *tie points* would, theoretically, have the sensors be consistent everywhere else.

The first adjustment was applied to DSS-B, guided by analysis from the ground processing stream, which indicated this sensor had the largest magnitude of errors. The second adjustment, to DSS-C in pitch, was guided by evidence in the HSs of real pitch changes at the sensor AOS and LOS.

2.5.4 Single-String Adjustments

The need for the next adjustment to DSS biases became clear after DSS-C was turned off. Any bias effects would have been cut in half by the presence of DSS-C, averaging the results for the observed sun direction through the subsolar point (during DSS-A LOS and DSS-B AOS). With DSS-C off, pitch motion at the DSS-B AOS was observable, and a pitch alignment correction to adjust it was estimated.

After this last bias adjustment was done, it was realized that there was still a residual signature in yaw just after the subsolar point. It looked like the right half of the signature from Fig. 7 for either a $-0.1°$ pitch bias or a $+0.1°$ roll bias. It was observed that this error could be accounted for by a small adjustment in either the twist or the elevation of DSS-B. (The yaw signature could also be simulated by a pitch bias, but that would also show up as real motion via the HSs.) A different approach, therefore, was taken for selecting further refinements.

2.5.5 Ground-Computed Alignments

The final set of alignment adjustments reported in Table 10 was based on the ground alignments translated to give an adjustment to DSS-B relative to DSS-A. The effects were simulated and seemed to indicate the correct error. The relative error was translated to provide just a correction to DSS-B. It was gratifying to note that the alignments accumulated on board by various small corrections had converged close to those determined on the ground using island targets (Patt and Bilanow 1999). These loads were uplinked, and worked as expected. It was later estimated that there was probably a small residual pitch error in DSS-B of about 0.04°, but another load for this small adjustment did not seem warranted.

In addition, further adjustments to both of the onboard DSSs currently in use were considered to make them agree with ground alignments. This last step was not pursued, however, because only the relative alignment errors have a significant effect on board.

2.5.6 Onboard Calibration Curve

There is an onboard lookup table for each sensor axis based on the manufacturer's calibration. On the ground, a nominal relation for DSS measurements was used whereby the raw measurements are proportional to the tangent of the angle, and the raw counts were biased or scaled for fine tuning bias adjustments. A sample plot of the difference between the onboard and ground calibration for one of the axes is shown in Fig. 14. The validity of these table adjustments is not clear, but the differences are small enough that they are not critical to the onboard accuracy. To the extent that the high frequency noise-like features of the correction are not accurate, this may lead to some additional noise in the onboard pitch and yaw estimates.

Table 11. The magnetometer bias adjustments. Blank entries indicate there was no change made in the given parameter. Note the correspondence between the TAM and ACS coordinates described in Sect. 2.1. The units are in nanoteslas times 10^{-4} (nT$\times 10^{-4}$).

Date	*TAM-A x-axis (Along $\vec{Y}$)*	*TAM-A y-axis (Along $-\vec{X}$)*	*TAM-B x-axis (Along $-\vec{Y}$)*	*TAM-B y-axis (Along $\vec{X}$)*	*Notes*
Prelaunch	−0.0484	−0.0320	−0.0320	−0.0036	
19Aug1997	−0.4684		−0.4520		Made x-axes agree
20Oct1998	−0.57	−0.17	−0.35	−0.24	Adjust for yaw spikes
15Dec1998		−0.07			Tuning in single string

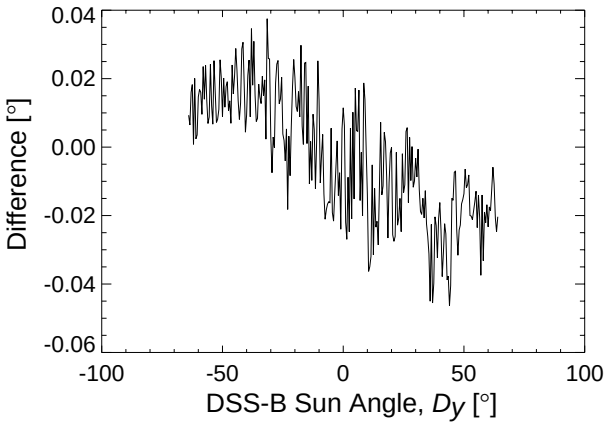

Fig. 14. The difference between the onboard table and ground formula for the DSS-B D_y angle measurement.

2.6 TAM Bias Effects and Adjustments

A summary of the onboard TAM bias adjustments is given in Table 11, and the issues associated with each adjustment are discussed in the following subsections. These biases were applied via the zeroth order parameters in the TAM onboard calibration polynomials, because more direct onboard bias adjustments were not implemented correctly in the onboard code. The sensor axes for these adjustments were different than the ACS-defined axes, as noted in Table 11. The TAMs were only roughly calibrated, with constant offsets in two of the three axes. Because TAM errors mainly affect onboard attitudes in the back orbit where science data are not collected, these calibrations were of lower priority.

2.6.1 Initial Bias Adjustments

These biases were applied to the constant term in the polynomial calibration used on board. That is why the initial prelaunch values were not zero. Another bias term available in table loads only allowed the same bias to be used in all three axes because of a coding oversight. It is also noteworthy that there are actually different calibration formulas applied at large values beyond the nominal ranges, but the formulas for the extended conversions from sensor counts to field measurements were not updated. This led to some discontinuities in the computed measurements at maximum field values, but these were not of significant concern (Sect. 4.3.1).

The first load was done soon after launch by OSC engineers to fix a large inconsistency between the $\vec{Y}$ components of TAM-A and TAM-B. Both components were shifted by 4,200 nT relative to their initial values. The cause of the discrepancy was not known.

2.6.2 Back-Orbit Error Signature

The effects of TAM errors on the attitude are mainly important when DSS data are not available, which is mostly in the shadow period of the back orbit. In addition, the nadir and magnetic field geometry is most sensitive to biases near the magnetic poles, similar to the way the sun–nadir geometry is sensitive to biases at the subsolar point. Over the magnetic pole, the field lines are along the nadir direction. This problem is of secondary concern, because the SeaWiFS science data are collected well inside the sunlit portion of the orbit. Often, the magnetic pole overflight is also inside the sunlit portion of the orbit, but not always.

The Earth magnetic poles are displaced from the geographic pole (by about 10° in the Northern Hemisphere and 15° in the Southern Hemisphere), and therefore, they are overflown by the 98.2° inclination orbit of OV-2. An overflight of the magnetic pole is not of concern unless OV-2 is in the Earth's shadow at the time. OV-2 reaches the Earth's shadow at about 26° past the terminator crossing. With the sun up to 23.4° from the equator at the solstices, the shadow entry can be as close as 2.6° to the northernmost or southernmost point in the orbit. OV-2 can, therefore, pass over the magnetic poles in the Earth's shadow during winter in the Northern or Southern hemispheres.

A sample of the yaw error signature for a back-orbit data span is provided in Fig. 15. The start and end of the shadow period (from 01:19:04–01:54:12) are clearly marked by discontinuities in the onboard calculated yaw. In this case, the spacecraft was within 2.5° of the expected magnetic pole position just before the end of the shadow period. At this pole proximity, as noted in Sect. 2.2.5, the yaw is set to zero (from 01:53:16–01:54:12) because the

three-axis attitude estimate is considered unreliable. The region of poorest attitude estimation geometry (Sect. 2.2.3) occurs nearest the north magnetic pole. The north magnetic pole overflight can be in shadow because the date is near the Northern Hemisphere winter solstice; meanwhile, the southernmost latitude overflight at the beginning of the span is in sunlight. The yaw error gets progressively worse as the north magnetic pole is approached. The causes of the large outlier values (−43.6 and 11.5° at 01:52:58 and 01:53:14, respectively) are not known. (The feature at about 01:48 is caused by low frequency TAM noise, which is discussed in Sect. 2.6.4, along with general TAM modeling errors.) This signature, as the magnetic pole is approached, is similar to effects from an HS pitch or roll bias near the subsolar point, as discussed in Sect. 2.3.4.

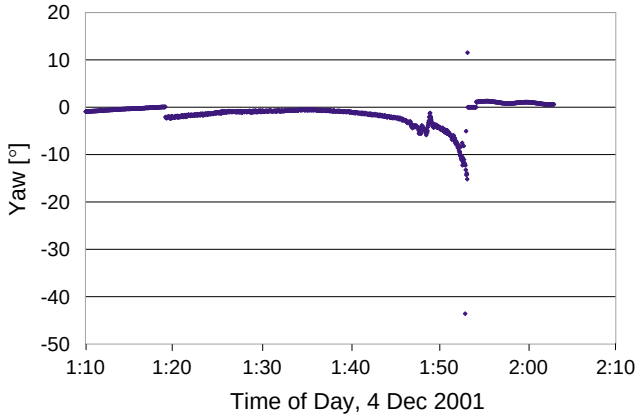

Fig. 15. Yaw angles computed on board during the back-orbit span sampled at high data rates on 4 December 2001.

This yaw error signature changes significantly throughout the day as the Earth rotation causes the magnetic pole to change position relative to the orbit plane and shadow entry point. There is also an important seasonal difference in the yaw error signatures: in December, the large errors occur in near the north magnetic pole just before the end of shadow, while in June, the large errors occur near the south magnetic pole just as the shadow period starts. Figure 16 shows this different characteristic, along with the variation throughout the day.

The effects of TAM biases first became quite apparent as the June 1998 solstice was approached (Fig. 16). This includes data from the back orbit, outside the science data collection period. Data are sampled on board every 3 min during the back-orbit period (Sect. 1.2.4), and so data samples are well separated in the plot where the yaw error magnitude changes rapidly during the shadow period. At this low data rate, the worst case of the yaw errors may not have been sampled (as in the high-rate data plot shown in Fig. 15), so the errors may well have exceeded 15° during parts of this day. Errors exceeding 50° were seen on other days.

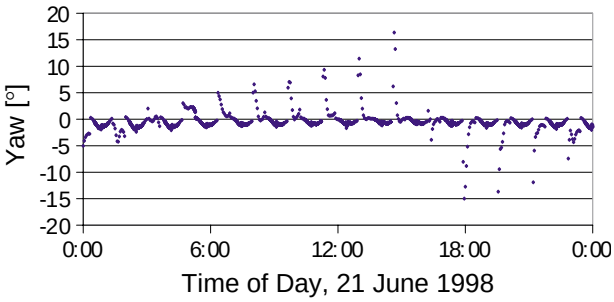

Fig. 16. Yaw angles computed on board, on 21 June 1998, showing large errors during the shadow periods due to TAM biases. Sampling is sparse during the large errors because of the usual 3 min intervals between samples taken in the back-orbit periods.

The peak yaw errors seen in Fig. 16 occur at the first data sample inside the Earth's shadow. This is where the spacecraft must rely on nadir and magnetic field vectors for attitude, and vector separation geometry is poor while the spacecraft is closest to the south magnetic pole. The pattern of maximum yaw errors, from orbit to orbit throughout the day, can be understood from the effects of the Earth's rotation on the relative nadir–magnetic field geometry. As the magnetic pole rotates with the Earth relative to the spacecraft's shadow entry point, the magnetic field vector separation angle from the nadir varies accordingly. The change in the sign of the error close to 18:00 is similar to the effect of a sun elevation change on a pitch bias effect when passing the subsolar point (Sect. 2.3.2). In this case, a TAM bias along the component parallel to the spacecraft velocity causes a yaw error of opposite sign, when the magnetic field direction at the shadow entry point changes as the magnetic pole rotates under the orbit plane between subsequent orbits.

2.6.3 Bias Adjustment Results

Adjustments to the TAM biases were decided upon based on simulations done in consultation with the OSC ACS engineers. It was shown that yaw error signatures very similar to those in Figs. 15 and 16 could be simulated by adjusting the T_y component of the bias. Simulations showed an apparent T_x component adjustment was indicated as well. The T_x components were adjusted differently for the two TAMs to give more consistency between the two.

After the bias adjustments were loaded on 20 October 1998, yaw errors at night were reduced as expected, but another small effect was noted in sunlight: the subsolar yaw error signatures changed noticeably. This indicated how in regions where the sun–nadir geometry was poor, the TAMs could have an influence that was not always beneficial, as will be discussed further in Sect. 2.6.6. It was also later noticed that this bias adjustment caused a

small shift in the average spacecraft $\vec{Y}$ (spin axis) direction (Sect. 3.1.1).

After the switch to the single string (using TAM-A only) an additional adjustment was needed to reduce yaw errors during the Northern Hemisphere's winter season. It was clear that further TAM calibration could be desirable, and that the constant biases might be refined with seasonal adjustments because seasonal effects in the back-orbit error patterns were apparent. Further refinements of this type were not pursued, however, because of two things:

1) The errors in the shadow period were less important than concerns for the daylight span where science data is collected, and
2) Other systematic model errors that were hard to accommodate became apparent, as discussed in the next section.

2.6.4 Model Errors and Low Frequency Noise

It was clear from analysis of the TAM data that there are scale factor and higher order errors present in the calibration tables, as well as misalignment errors. Many missions have solved for a combined scale factor and misalignment matrix for TAMs as well as biases, but no such effort was considered for OV-2. (Designing tools and testing data interfaces for such a calibration is usually done in prelaunch preparations, while postlaunch analysis resources are limited.) In any case, the calibration curves show nonlinear effects that could be hard to model, and consistent results could not be easily obtained. A further complication came with the understanding that the onboard model for the Earth's magnetic field uses 1980 International Geomagnetic Reference Field (IGRF) coefficients, which are significantly out of date. Hashmall and Sedlak (1997) note that errors from the field model can be significant. TAM bias tuning for OV-2 was limited to minimizing the errors at the most sensitive place in the orbit by a simple bias adjustment, with simulations showing the signatures of errors in those places.

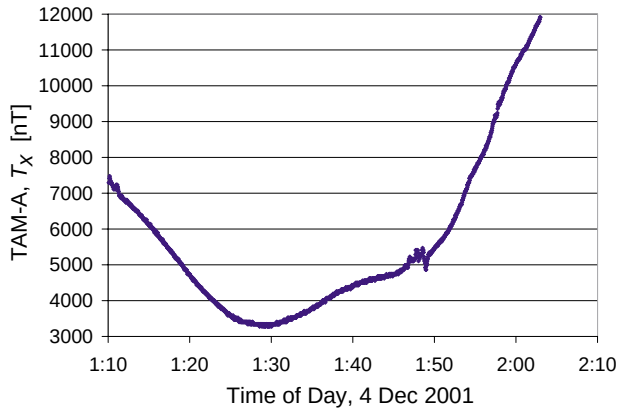

Fig. 17. TAM-A T_x measurements during the back-orbit span sampled at high data rates on 4 December 2001.

Low or intermediate frequency noise in the TAM data show up as irregularities in what is normally expected to be a gradually changing field direction. An example is shown in the T_x measurement in Fig. 17. This is the same time span as Fig. 15, and it can be seen how the variations in T_x at about 01:48 correlate with changes in the onboard yaw estimate at the same time. Many irregularities were later associated with Birkeland currents (Sect. 4.3.2).

2.6.5 DSS Subsolar Gap Errors

The effects of TAM biases can also be seen on the sunlit side of the orbit. The effect was indicated with the load of the TAM biases, but became especially clear after the switch to the single string and the appearance of a gap in daylight DSS coverage. In January 1999, as the sun elevation reached a maximum, the gap in the coverage between DSS-A and DSS-B increased (Sect. 4.1). During this time, the gap extended outside the yaw hold region. This provided visibility on the HS-TAM yaw solution for a brief period in the middle of the day. At a certain time every day, the yaw error would be a as much as 3° negative, and then an orbit or two later, it would exceed 1° positive. This is illustrated for three consecutive orbits in Fig. 18.

The time of day with the largest errors corresponds to the South Atlantic Anomaly (SAA) passage, where the magnetic field is weakest. Constant biases can be expected to have their largest angular effect when the field magnitude is smallest, so that may contribute to the error here. That the errors change so much over a few hours might make the 1980 IGRF model a suspect error source. These daytime errors are much smaller than the yaw errors seen in the nighttime, but were of more concern because they occur during science data collection. Note that the discontinuities introduced by the DSS coverage gap, as illustrated in Fig. 18, were later eliminated by the widening of the yaw hold region, as discussed in Sect. 2.2.4.

2.6.6 TAM Weighting Reduction

The variability of the TAM errors discussed in the previous section helped motivate a decision to further lower the relative weighting of the TAMs in the Quest algorithm (Sect. 2.2.1). The *de-weighting,* as it was called, was done in two stages, as noted in Table 6. The first step made a smaller, more conservative adjustment just to make sure no unexpected results occurred. The final change reduced the TAM-A presumed accuracy relative to the HS by a factor of over 700. This ensures that the DSS and HS data dominate the attitude solutions when available. This subsection describes the rationale for that change in more detail.

The influence of the TAM errors on the subsolar errors can be described as follows. When the sun–nadir geometry degrades, the TAM data have more influence via the Quest algorithm. The influence is determined by the relative weights used. The assumed sensor accuracy is the

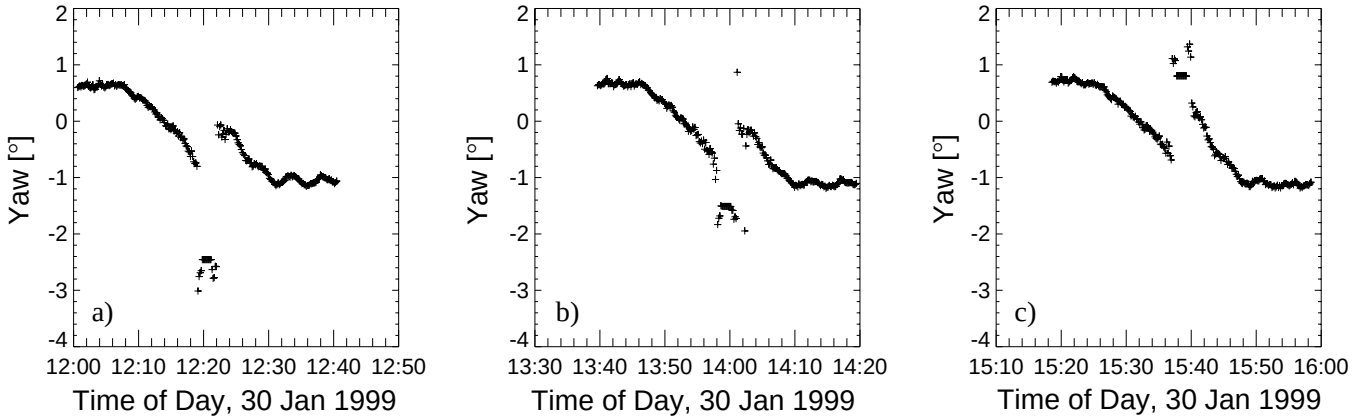

Fig. 18. HS-TAM yaw solutions are shown during the sun coverage gap on 30 January 1999 for three consecutive GAC scenes, for the following times of the day: **a)** 12:00–12:50, **b)** 13:30–14:20, and **c)** 15:10–16:00.

inverse of the square root of the weight used, and so was originally about .14 and 3.2 for the HS and TAM, respectively. From geometrical considerations already discussed, the accuracy is reduced by the sine of the separation angle in computing the rotation direction to another vector. With the nadir within 5° of the sun direction, the accuracy provided in yaw by the HS data is 1.6. This implies that the HS only provides about twice the accuracy of the TAM in yaw information very close to the subsolar point. TAM errors can, therefore, noticeably influence the yaw solution at the subsolar point.

The results shown in the previous subsection suggest that a significant part of the subsolar yaw signature variability is produced by the effect of TAM errors. The worst TAM yaw errors only occur at certain times of day, and their effect is reduced to roughly a third of what it would be without HS data. The amount that the subsolar signature changes at these times of day, even with DSS data present (e.g., Sect. 3.4.2) is consistent with the data weighting and the errors that come from the TAM data. A further reduction in the TAM weighting was, therefore, desirable to reduce yaw errors at the subsolar point. The reductions were implemented using in table loads in December 1999, as noted in Table 6.

3. POINTING STABILITY

This section discusses the natural and disturbance-induced attitude motion of OV-2 and the control system responses. Because of momentum wheel gyroscopic stability along the $\vec{Y}$ (pitch) axis, the motion has different behavior in roll and yaw compared to the behavior in pitch. Moreover, roll and yaw control is coupled with momentum management, which in turn affects pitch indirectly.

The actuators used to affect the pointing are the magnetic coil electric current pulses, and the momentum wheel torques. The coils used onboard OV-2 operate with a pulse width modulation. The coil current is turned off for at least 200 ms of every 2 s control cycle so that the generated dipole does not corrupt the TAM measurements of the Earth's magnetic field. The wheel speed is adjusted to control rotations about the $\vec{Y}$. Momentum management and roll–yaw adjustments for the spacecraft angular momentum vector are accomplished by the coil generated dipole interaction with the Earth's magnetic field. There are individual magnetic coils oriented along $\vec{X}$, $\vec{Y}$, and $\vec{Z}$, with commanded moments C_x, C_y, and C_z, respectively.

The following sections discuss various aspects of the spacecraft dynamics, some details of the control, typical attitude motion, and control adjustments.

3.1 Dynamics Overview

This section examines the dynamic behavior of the spacecraft attitude. The characteristics of gyroscopic stability are described, along with the associated problem of nutation. This is followed by summaries of disturbance sources and correlated alignment effects, and descriptions of the specific effects of magnetic dipoles, the BCR switch, and tilt changes.

3.1.1 Gyroscopic Stability

The momentum wheel rotating about $\vec{Y}$ introduces gyroscopic stability because of conservation of angular momentum. The momentum wheel rotates in a positive right hand sense about the $-\vec{Y}$(negative pitch) axis, referred to hereafter as the spin axis of OV-2. The spacecraft body also rotates at 1 RPO† about this axis, which nominally points in the positive orbit normal direction. The total angular momentum of the spacecraft results from both the wheel and body rotations, but most of the angular momentum is stored in the wheel.

† Revolutions per orbit.

Roll and yaw are coupled because the spin axis tends to stay fixed in inertial space. This axis does not stay strictly fixed because it is acted on by various torques; however, the maximum rate at which the axis changes is quite slow relative to the data sampling rate. The largest torques affecting the spin axis come from the coils (Sect. 3.1.7) and with the largest commanded moments sustained during some anomalies, the spin axis changes at a rate of roughly $1^\circ\,\mathrm{min}^{-1}$. Typically, this axis changes much more slowly, moving less than a degree during any orbit.

Roll–yaw coupling causes a typical orbit period sinusoidal time history in both roll and yaw, where yaw leads roll by 90° of orbit phase. If the spin axis stays along the orbit normal direction, roll and yaw will have zero values; however, if the axis is fixed in a direction 1° away from orbit normal, the roll and yaw will have sine wave time histories with amplitudes of 1°. The quarter orbit, roll–yaw coupling is an important property for a spin stabilized spacecraft, which is often used in attitude modeling and control (Headrick 1978).

A sample of how the roll and yaw time histories approximately follow this sinusoidal behavior is shown in Fig. 19 for a 24 h period, on 25 June 2000. The figure is based on the GAC data collected for that day, and therefore, only shows parts of the overall orbit period behavior. Points along sinusoidal waves of constant amplitude are overlaid, with yaw leading roll by 90° of phase, modeling a spin axis in a fixed position relative to the orbit normal direction. The spacecraft roll and yaw curves vary a small amount in amplitude and phase. The roll amplitudes vary from about 2.0–2.5°, which can be taken as an indication that the spin axis varies about that distance from the orbit normal direction. The spin axis is not exactly fixed relative to the orbit normal, but Fig. 19 illustrates how this simple model is an approximation, at least for this time frame of the mission.

The orbit normal direction is not fixed in inertial space, but precesses about 1° per day because of the sun synchronous orbit. The declination of the orbit normal is −8.2°, based on the 98.2° orbit inclination, and the right ascension of orbit normal increases 0.986° per day. Movement of the spin axis is required at a minimum to track the orbit normal. Following early mission control tuning (Sect. 3.3) the spin axis has tended to stay higher in declination than the orbit normal, and lower in right ascension, but there is drift in this axis due to various effects that are discussed in the following sections.

Typically, a regular daily pattern in the roll and yaw histories is seen, and thus, in the estimated spin axis relative to orbit normal. The pattern changes gradually with the seasons of the year, and more abruptly with ACS changes.

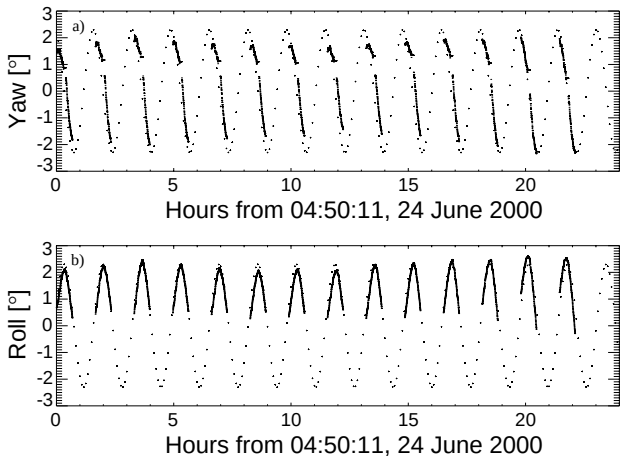

Fig. 19. Onboard **a)** yaw, and **b)** roll, from GAC spans for a 24 h period with a sinusoid pattern overlaid (dot symbols) based on a simple constant-spin axis model. The yaw leads roll by 90° phase as the onboard data fall approximately along different parts of the sinusoids during the 40% of the orbit in each GAC span.

One notable change in the general direction of the spin axis relative to orbit normal occurred with the changes in the onboard TAM biases that were discussed in Sect. 2.6.3. It is expected that the attitude estimates in the back-orbit span help steer the spin axis to the position where it stays, on average, during the GAC span. Because there were fairly large yaw errors in the shadow period (Sect. 2.6.3), which changed pattern throughout the day, they could be expected to cause variations in spin-axis pointing via control system responses to the errors.

The change and direction of the spin axis is illustrated with data from the ground estimates of the attitude for selected orbits on 19 and 21 October 1998 (Fig. 20). GAC scenes were selected to represent the typical range of spin axis variation over the course of each day, and the apparent motion shown by the connected line during each scene is due to a combination of real motion and attitude estimation error. Panel a) shows the spin axis relative to orbit normal during three GAC scenes on 19 October. The three orbits shown capture the typical range of motion of the spin axis relative to orbit normal for an extended period leading up to the TAM bias load.

Random variations in the spin axis path result from noise sources in the attitude estimate. The ground estimate of the attitude effectively assumes some stability in the spin axis relative to orbit normal in the dynamics model, but via the Kalman smoothing algorithm (Patt 2002), drift in the spin axis is estimated.

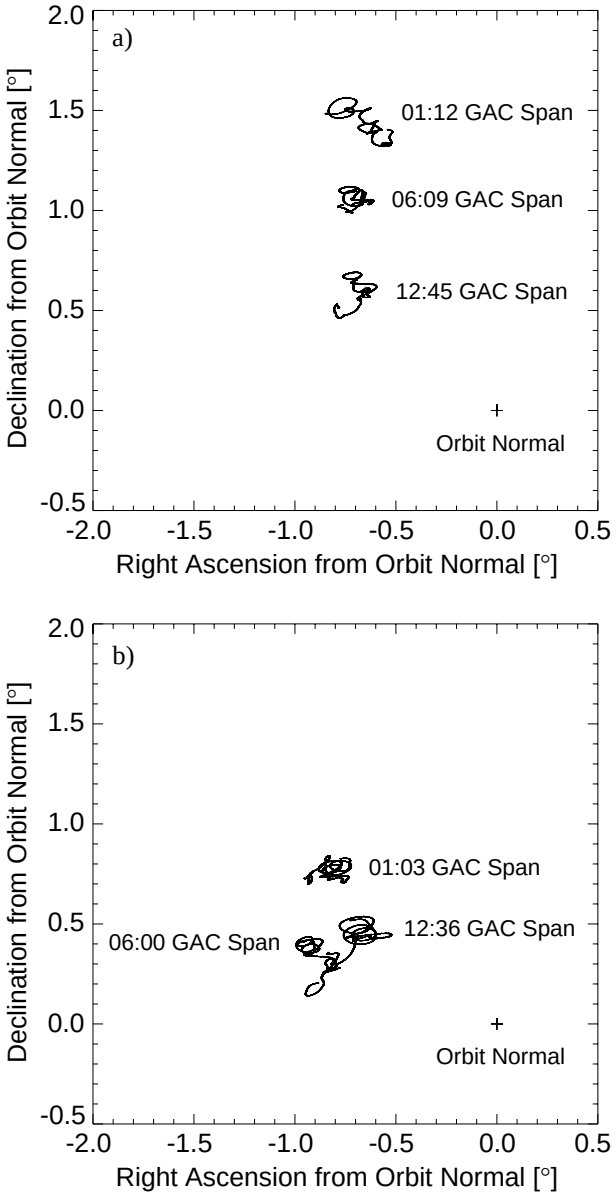

Fig. 20. Ground-estimated spin axis positions relative to orbit normal before and after TAM bias changes, where **a)** is during three GAC scenes on 19 October 1998, and **b)** is during three GAC scenes on 21 October 1998. Note the shift in the mean position and variation over the day, which was generated after the TAM bias changes.

Panel b) of Fig. 20 shows similarly timed orbits for 21 October. The average spin axis has moved down in declination and slightly in right ascension (relative to orbit normal). The variation over the day was smaller, and followed a different pattern—the same pattern repeating for many days following the bias change. The reduced motion over the day is probably a result of reducing the yaw errors in the back-orbit span.

3.1.2 Nutation with Dual Spin

A regular feature in the OV-2 roll and yaw behavior is a rocking motion with a period of about 5 min, which is referred to as nutation. This has a different characteristic than nutation as defined for rigid body dynamics. (Rigid body nutation occurs when the angular velocity vector is not along a principal axis, causing an offset between the angular velocity and angular momentum, which results in wobbling of the body.)

The OV-2 is a momentum-biased system (referred to as either a gyrostat or a dual-spin spacecraft—a body attached to an independently spinning mass). In this configuration, the majority of the system angular momentum is produced by the spinning mass, i.e., the momentum wheel. Nutation results whenever the spacecraft angular velocity is not entirely along $-\vec{Y}$ (assuming that the $-\vec{Y}$ and momentum wheel axes are perfectly aligned). Pure pitch rotation is an idealization, not achievable in practice, so nutation occurs in a momentum-biased spacecraft at some level nearly all of the time. The extent of nutation depends on two factors: the size and frequency of disturbances, which produce angular velocity along the $\vec{X}$ and $\vec{Z}$ axes; and the effectiveness of the control system in *damping* (or attenuating) nutation (along with other factors, e.g., energy-dissipating spacecraft components).

Nutation can be described in terms of torque-free spacecraft dynamics (Chobotov 1991, and Beard and Plett 1978). Although OV-2 is not torque-free, the typical torques applied by the coils are small, so for short periods, the angular momentum is nearly constant. The effect of a $\vec{X}$ or $\vec{Z}$ component of angular velocity (hereafter referred to as *transverse angular velocity*) is to precess the momentum wheel axis. To conserve angular momentum, the transverse angular velocity (and its associated momentum) precesses in the opposite direction. A detailed treatment of this phenomenon is beyond the scope of this document; however, in broad terms, the precession rate of the transverse angular velocity (which determines the nutation frequency) depends on the total $\vec{Y}$ axis angular momentum and the ratio of the $\vec{Y}$ and transverse axis moments of inertia.

The total angular momentum (combined spacecraft and momentum wheel) is about 20 times the spacecraft pitch angular momentum alone (as estimated from the OV-2 lunar calibration pitch maneuvers). The $\vec{Y}$ and $\vec{X}$ moments of inertia are nearly equal, although the $\vec{Z}$ moment is significantly smaller. The spacecraft pitch rate is 1 RPO, which for the orbit period of about 99 min is about $0.00106\,\mathrm{s}^{-1}$. In simple terms, this produces a nutation frequency of about 20 times the orbit rate, or about $0.02\,\mathrm{s}^{-1}$, corresponding to a nutation period of about 300 s, or 5 min. Although this description is greatly simplified, it explains the basic character of the nutation.

The transverse angular velocity that produces nutation has been triggered by a variety of sources during the mission. Most of these have resulted from large discontinuities or errors in the onboard attitude angles, that result

in significant, persistent transverse torques being applied. The triggering events tend to occur either in the back orbit (with nutation observed at the start of the GAC data collection) or at the subsolar point. Several such phenomena are described later in this section and in Sect. 4. One unique type of event (the key change anomaly, Sect. 4.6.1) resulted from a physical disturbance of the spacecraft.

The OV-2 nutation typically has amplitudes around 0.1°. Various anomalies (Sect. 4) have occasionally produced nutation up to 1°. Early-mission attitude control (described below) included frequent large nutation. The control system works to actively damp nutation via rate feedback to the coils (Sect. 3.2). Other passive energy dissipation (e.g., fuel slosh, magnetic eddy currents) may also provide damping, but these are very difficult to quantify. Small nutation will usually disappear during a GAC scene, while larger nutation takes longer. Typical nutation can be seen in Fig. 21, which shows the onboard roll and yaw angles for a GAC scene on 24 June (SDY 175) 2003. The nutation can be seen in both angles at the start of the scene, with an amplitude of 0.1° or less, and gradually diminishes to the point of imperceptibility at the end of the scene.

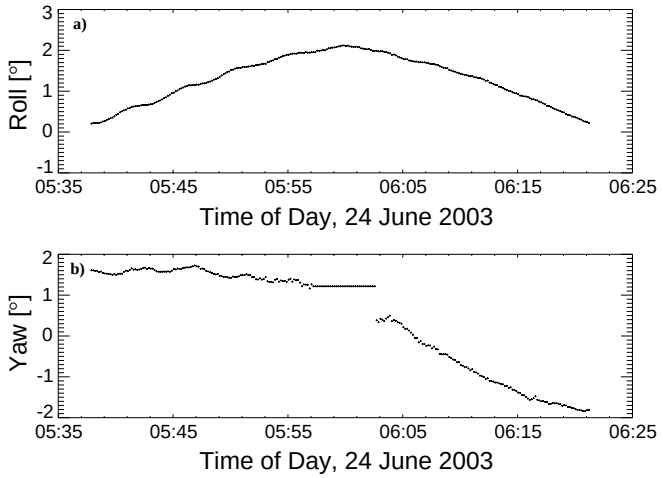

Fig. 21. Attitude angles with typical nutation: **a)** onboard roll angle, and **b)** onboard yaw angle.

An example of large nutation occurred on the first GAC scene of 15 July 1998, following an encryption key change. In this instance, the onboard roll and yaw angles (Fig. 22) showed nutation of approximately 1° amplitude coming out of the back orbit. The ground-computed angles for this orbit (shown in Fig. 22, solid lines) show nutation with a smaller amplitude; this is because the ground navigation algorithms do not include a nutation model, and the Kalman filter tends to underestimate attitude motion on short time scales. The amplitude decreases during the scene, but has not completely damped out at the end of the scene.

In the yaw angle, the nutation amplitude appears to be larger at the start of the scene; this is expected from the smaller moment of inertia about $\vec{Z}$. Nutation in yaw is not observable near the subsolar point because of the poor geometry (Sect. 2.2.3) and because of the yaw hold region (onboard) and the effects of the Kalman filter (ground). The magnitude of the nutation in this scene allows it to be seen in the attitude rates as well (Fig. 23), both on board (dot symbols) and ground computed (lines). The characteristics are similar to those of the angles; the ground values have a smaller amplitude, and yaw has a larger amplitude than roll.

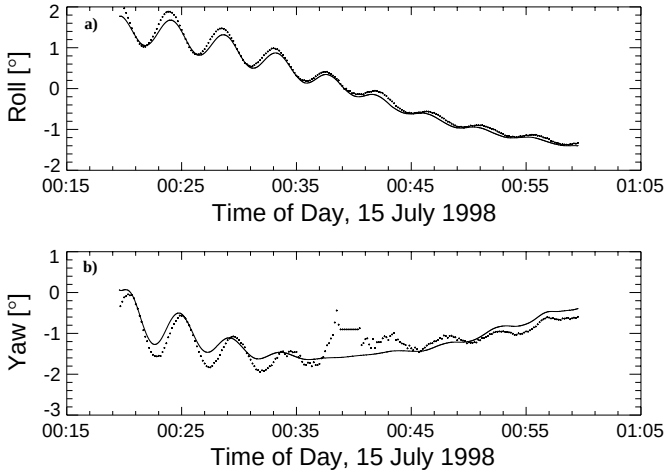

Fig. 22. Attitude angles with large nutation are shown: **a)** onboard (dots) and ground-computed roll angles (solid line), and **b)** onboard (dots) and ground-computed (solid line) yaw angles.

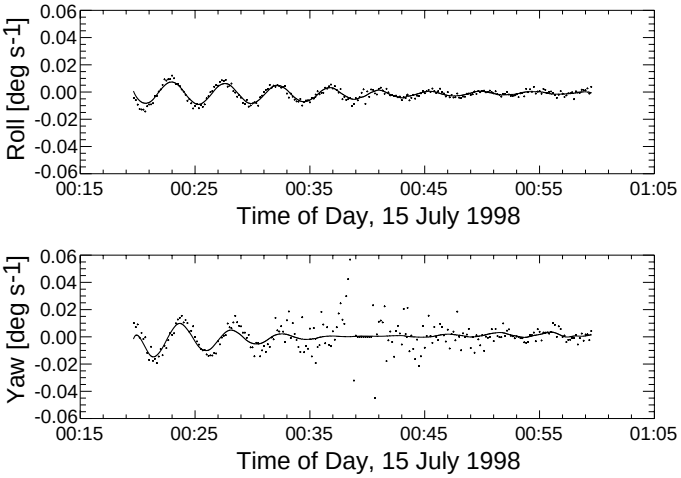

Fig. 23. Attitude rates with large nutation are shown: **a)** onboard and ground-computed roll rates, and **b)** onboard and ground-computed yaw rates. (The symbols are the same as in Fig. 22.)

The coil commands, C_y and C_z (Fig. 24), show the control system response to nutation; the C_y command shows an obvious oscillation, reflecting the control system attempting to reduce the nutation, while the C_z coil command shows a small variation with the same period. The

relationship of the coil commands to the angle and rate errors is described in Sect. 3.2.

The ground-estimated spin axis path resulting from the large nutation amplitude is illustrated in Fig. 25. This looping motion of the axis is reduced by the damping effect of the coils. The roll oscillations lead the yaw by about 90° of nutation phase angle.

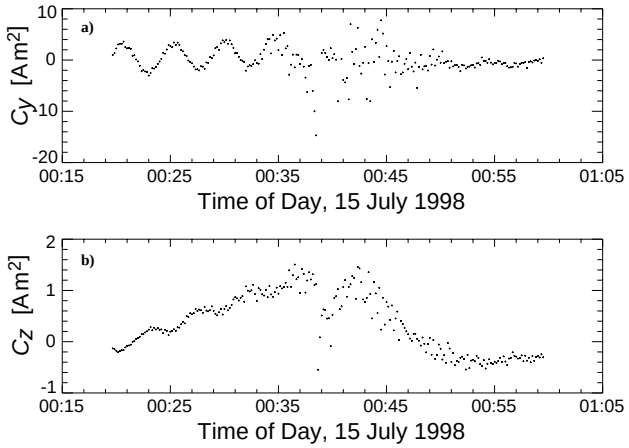

Fig. 24. Coil commands in response to large nutation are shown: **a)** C_y, and **b)** C_z.

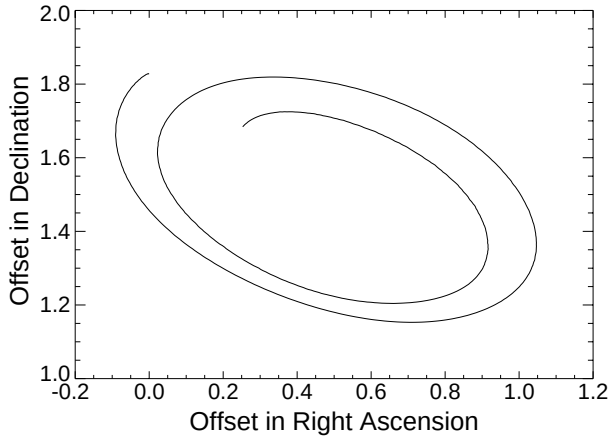

Fig. 25. Path of ground-estimated pitch ($\vec{Y}$) axis relative to negative orbit normal with large nutation in the first 10 min of the GAC scene starting at 00:19 on 15 July 1998.

3.1.3 Disturbance Sources

Sources of torques generally affecting the spin axis are summarized here and important effects are discussed further in subsequent sections. The largest effects on the OV-2 motion are from the commanded magnetic dipole, and the results are not always beneficial, as discussed with various disturbance-generating events in Sect. 3.3 and certain anomalies as discussed in Sect. 4. Another dominant effect is understood to be a residual magnetic dipole, as discussed further in Sect. 3.1.5.

Magnetic Dipole, Commanded: The coils are commanded to create a dipole on board that applies a torque to the spacecraft, as the magnetic forces tend to align the dipole with the Earth's magnetic field. It is used for control and momentum management, but also as a source of disturbance when commanded in response to erroneous ACS input.

Magnetic Dipole, Residual: A residual dipole can result from a number of occurrences:

a) Free current loops in the onboard power system and instrument wiring,

b) Structure-charging interactions with the space environment, or

c) A bias in the coil commands.

A residual dipole is thought to be present on OV-2, which affects the spin axis offset from the target orbit normal axis.

Aerodynamic Torque: Torques result from offsets between the center of pressure and center of mass as the thin atmosphere at the orbit altitude interacts with the spacecraft. This is not expected to be significant for OV-2, however, there are no data on this and effects might be correlated with alignment errors (Sect. 3.1.4). Spacecraft are usually designed to minimize the center-of-pressure to center-of-mass offset. Because there is more axial symmetry about $\vec{Z}$, and $\vec{X}$ is the flight direction, an offset would be most likely in the $\vec{Z}$ body direction because of effects of the solar arrays. The arrays are nominally *edge on* to the wind, but the surface cross section does change with pitch variations.

Gravity Gradient Torque: This torque is nominally very small as OV-2 flies with its long axis vertical, and this is a stable position for gravity gradient effects. During the lunar calibration maneuvers, the gravity gradient torque has noticeable effects as the spacecraft rotates about the $\vec{Y}$ axis. While the maneuver rate would nominally be constant from the wheel speed control, the gravity gradient torques slightly affect the motion, giving noticeable variations in the pitch rate during the maneuver.

Solar Radiation Pressure: Torques result from solar radiation pressure, as offset from the center of mass. This is not expected to be significant for OV-2.

3.1.4 Correlated Alignment Effects

Errors in alignments can have similar effects to dynamic motions, especially as it concerns certain orbit frequency effects. Key associations that are relevant for a momentum-biased Earth-pointing spacecraft are shown in Table 12. The misalignments here are between the reference frame for the spacecraft, and a coordinate frame with the wheel momentum axis exactly along $\vec{Y}$.

Note two particular associations: a constant roll torque has the same effect as a negative roll misalignment error,

Table 12. The effects of alignments and torques on attitude histories. Blank entries indicate there is no change to that parameter directly from the given cause.

Cause	*Effect on Pitch*	*Effect on Roll*	*Effect on Yaw*	*Effect on Spin Axis Time History*
Pitch misalignment	Pitch offset	nc	nc	nc
Roll misalignment		Roll offset		1 RPO cycle, with $\vec{Y}$ farthest south at maximum latitudes.
Yaw misalignment			Yaw offset	1 RPO cycle, with $\vec{Y}$ farthest south at ascending node.
Constant pitch torque	†	‡	‡	‡
Constant roll torque		Roll offset, leads to negative yaw offset after 1/4 orbit, so roll does not accumulate.	Yaw motion as momentum axis direction is precessed from roll torque.	1 RPO cycle precession toward $\vec{X}$ so that $\vec{Y}$ is farthest north at maximum latitude.
Constant yaw torque		Roll motion as momentum axis direction is precessed from yaw torque.	Yaw offset, leads to roll offset after 1/4 orbit, so yaw error does not accumulate.	1 RPO cycle precession toward $\vec{Z}$ so that $\vec{Y}$ is farthest north at ascending node.

† Wheel torques counteract body pitch torques, but the momentum increase requires additional dumping.
‡ Coil commands for momentum dumping can indirectly affect the roll, yaw, and spin axis time history.

and a constant yaw torque has the same effect as a negative yaw misalignment error. In the torque case, the momentum axis of the spacecraft is actually changing in inertial space, while in the misalignment case it is just the spacecraft reference $\vec{Y}$ axis offset from the momentum axis that prescribes a circle because of the 1 RPO rotation of the spacecraft body. Nevertheless, these two effects can be indistinguishable, i.e., an alignment adjustment could compensate for real motion of the wheel spin axis at the orbit frequency.

3.1.5 Dipole Effects

Residual magnetic dipole effects have a twice-per-orbit period effect on the motion of the spin axis on a momentum biased Earth-pointing spacecraft, and thus, are generally distinguishable from the orbit period effects discussed above. It is noteworthy, however, that these effects are highly correlated over a relatively short data span and even notably correlated for the 40 min GAC scenes.

The general nature of a dipole effect on the spin axis can be understood by reviewing how the magnetic field direction changes in the spacecraft body frame around a typical orbit. The basic near-Earth magnetic field is illustrated in Fig. 26. The Earth's magnetic field direction is roughly along $\vec{Z}$ near the poles and along $\vec{X}$ near the equator; and cycles once during each orbit in ACS coordinates and twice during the orbit in inertial coordinates. A residual dipole along $\vec{Y}$ will tend to apply a torque to align the spacecraft with this field, and thus the torque axis will cycle twice each orbit in inertial coordinates while staying perpendicular to $\vec{Y}$ and the field direction. Because of the gyroscopic stability about $\vec{Y}$, the effect of this torque is to precess the spin axis at twice the orbit period in a looping path. The torque will act to precess the momentum vector toward the torque direction, incrementally steering the spin axis, which will loop twice each orbit. If the field magnitude and dipole were constant, and the field stayed exactly in the $\vec{X}$–$\vec{Z}$ plane, this would be a simple circular looping of the spin axis. In reality, there is more complexity to the motion induced. For example, because the field strength is greater over the poles than over the equator, the torque tends to be greater over the poles and the motion there is faster, tending to move the spin axis more in one direction.

About 18 months into the mission, it was inferred that a residual dipole on the spacecraft was probably having an important effect on the roll and yaw history. This was demonstrated by simulations showing how a commanded dipole bias could generate key aspects of the off-orbit normal pointing trends and certain more detailed variations

in the spin axis time history over each orbit. A bias of very reasonable amplitude (based on experience with other spacecraft missions) of $2.0\,\mathrm{A\,m^2}$ along $\vec{Y}$, made the axis move off of orbit normal in a way very similar to that seen in flight. Furthermore, the simulated spin axis drifted relative to orbit normal over a day in a way similar to the flight data, because of effects of the magnetic field rotating with the Earth under the orbit plane. The key clue that this was a dominant effect came with the BCR switch, which is discussed next.

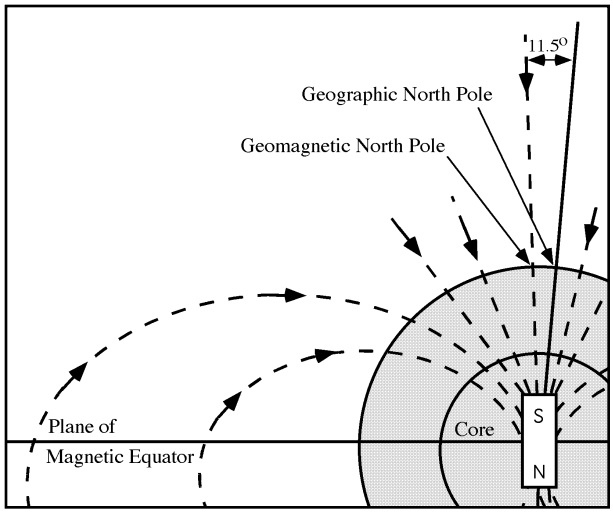

Fig. 26. Simple dipole model of the Earth's magnetic field.

3.1.6 BCR Switch Effects

A change in the attitude history occurred after a switch to the alternate BCR on 16 December 1998. Sometime after this occurred, it was noticed that the path of the spin axis in inertial space during each data span had changed characteristics precisely when the BCR switch occurred. Previously, ground computations had shown the spin axis to change relatively little during the orbit, and in fact, some of the sensor alignment parameters had been tuned based on the assumption that the spin axis was not expected to move very much. After the BCR change, a looping path was followed. The estimated spin axis history before and after the switch is illustrated in Fig. 27.

Given that the GAC spans represent about 40% of the orbit, the path showed a twice-per-orbit frequency, which is a characteristic of a residual magnetic dipole. After discussions with OSC engineers about the functions of the BCR, it was realized that one possible effect of different wiring paths to the solar arrays for each BCR might be a different residual magnetic dipole.

Based on this observation, a $\vec{Y}$ dipole bias was tried in simulation, and it was found that it had two important effects:

1. The twice-per-orbit looping motion of the spin axis was generated.
2. The overall average position of the spin axis was driven away from orbit normal in a very similar manner to its position in actual flight data.

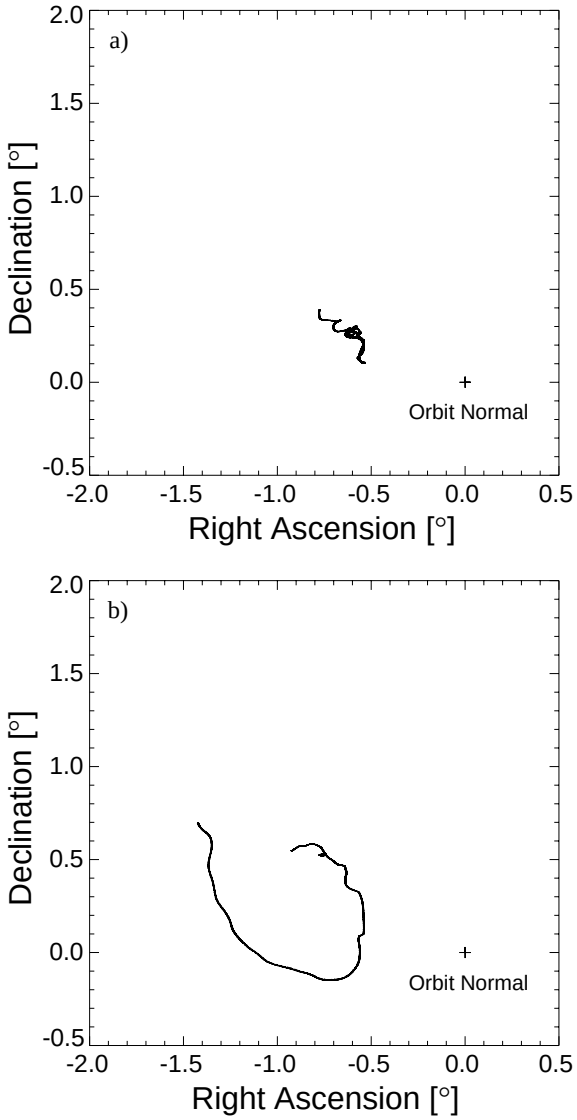

Fig. 27. Ground-estimated spin axis relative to orbit normal before and after BCR change on 16 December 1998, where **a)** is before, and **b)** is after. After the BCR switch, the looping path at roughly twice per orbit frequency was present.

It is believed, therefore, that a residual magnetic dipole is present, which influences the attitude history; however while the first effect above was noted only after the BCR switch, the second effect was seen seven months earlier after the coil gains were adjusted, as discussed in Sect. 3.3.4. This discrepancy is not fully understood, however, it is suspected that some biases in sensor alignments may have compensated for some of the dipole-induced attitude motion. This would be similar to the way more exactly-correlated motions may be compensated (Sect. 3.1.4). Because the alignments are calibrated against ground-truth

image data (Patt and Bilanow 1999), it does not matter that they are compensating for motion versus real alignment errors. The various effects cannot always be decoupled, however, as long as the model parameters lead to accurate calculation of the SeaWiFS pointing, the model is effective.

It seems likely, therefore, that a dipole bias has been present since launch, but that the dipole effects became more pronounced in causing off-orbit normal pointing of the spin axis after the coil gains were changed (Sect. 3.5.2), and the dipole changed magnitude at the BCR switch. If the dipole changes magnitude, then the twice-per-orbit motion induced should change magnitude proportionally. If other model parameters were effectively compensating for some amount of dipole-induced orbit motion, then the additional motion still produces apparent twice-per-orbit motion.

The overall effect of the BCR switch is probably more complex than a constant dipole effect. There could be some time dependence of a dipole generated by solar array currents as the charging levels change around each orbit. In any case, the overall motion of the OV-2 spin axis is also driven by time-varying responses of the control coil commands.

3.1.7 Tilt Change Effects

When the tilt of the SeaWiFS instrument is changed, this induces pitch motion on the spacecraft. This is illustrated in Fig. 28 with high sample rate data from an HRPT pass. The span shown was taken from an orbit where the tilt change was scheduled several minutes after the subsolar passage, so this sample does not also have pitch changes due to other subsolar region transition effects (e.g., Figs. 12c and d). The tilt change is a rotation of the instrument about $\vec{Y}$, which takes about 13 s to complete. The dynamic effect of the instrument motion is to impart a pitch motion to the spacecraft, resulting in about a 0.2° change in pitch during the tilt change. After this change, the spacecraft control system reacts to the pitch error, and drives the pitch back to the normal value over the course of about a minute.

In the ground processing, the pitch motion is smoothed by the Kalman filter in the attitude estimation and from effects of the sensor data smoothing. The current ground estimation, therefore, does not show the pitch motion during the tilt change being as rapid as it is in reality.

3.2 Control Torques Overview

This section discusses the control torques used for onboard attitude control.

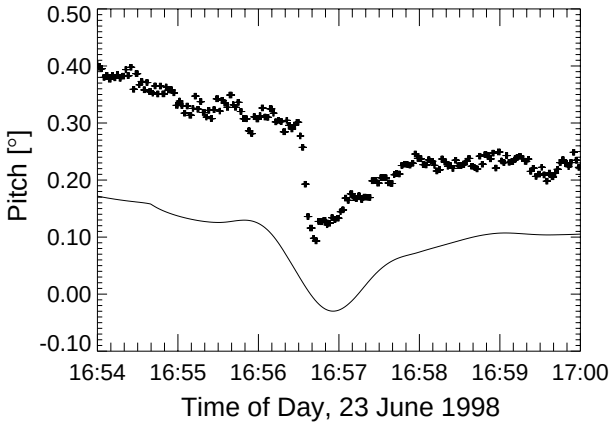

Fig. 28. Onboard (+) and ground-computed pitch (line) during a tilt change sampled at high data rates. The HRPT samples are every 0.5 s, but the onboard value only changes with the control cycle every 2.0 s. The ground estimate smooths the rapid dynamic motion with most of the real pitch motion occurring within 4 s during the most rapid part of the tilt change in the SeaWiFS instrument pointing.

3.2.1 Roll and Yaw Torque Computation

Basic roll–yaw stability is established by the momentum wheel, and therefore, the coil commands only need to apply sufficient torque to slowly steer the momentum axis to zero roll and yaw, and damp out effects from nutation. Momentum is stored along the $\vec{Y}$ axis, so $\vec{X}$ and $\vec{Z}$ torques are required for this steering. In the fine pointing mode, the desired torques are first computed about each axis, and for roll and yaw, the torques are chosen according to the roll and yaw errors and rates:

$$N_x = -0.0165\delta_x - 0.0248\delta_z - 2.4300\omega_x + 0.0096\omega_z, \quad (3)$$

and

$$N_z = 0.0248\delta_x - 0.0164\delta_z + 0.0161\omega_x - 1.8800\omega_z, \quad (4)$$

where N_x is the desired roll torque, N_z is the desired yaw torque, δ_x is the roll error, δ_z is the yaw error, ω_x is the roll rate, and ω_z is the yaw rate.

The constants in the above formula are used whenever roll and yaw error magnitudes are below 5°, which is the normal situation. They were set using the results of control simulations and analysis prior to launch, and cannot be adjusted from the ground in table loads, although all of the input, δ_x, δ_z, ω_x, and ω_z can be scaled by ground-loadable factors. Because the target attitude is zero roll and yaw, δ_x and δ_z are the same as the onboard calculated roll and yaw angles, respectively. The roll and yaw rate feedback is clearly designed to damp the nutation, with roll rates leading to an opposing roll torque (when possible) and yaw rates leading to an opposing yaw torque. The

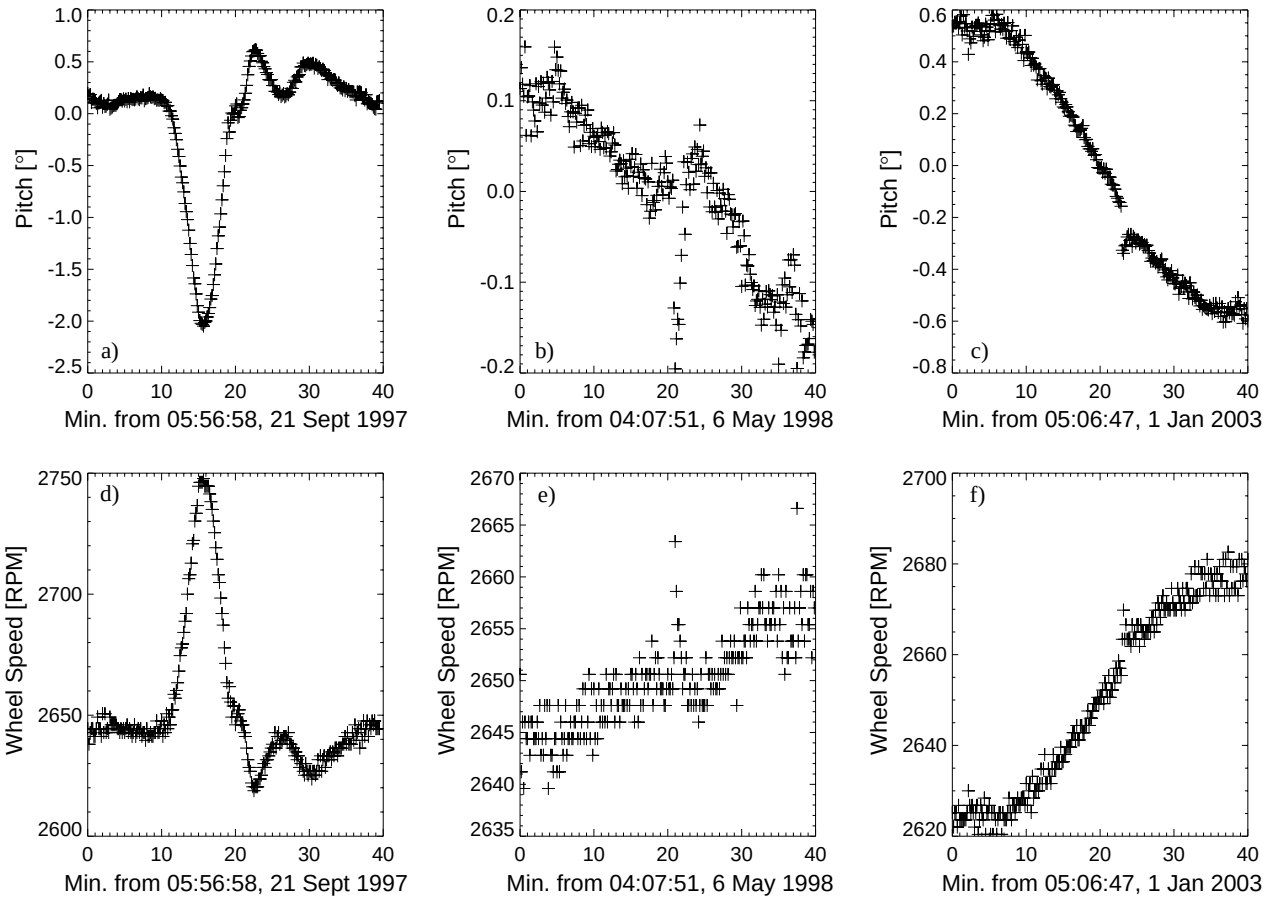

Fig. 29. Pitch and wheel speed correlation are illustrated for three sample GAC spans: **a)** and **d)** are the pitch and wheel speed for 21 September 1997, respectively; **b)** and **e)** are the pitch and wheel speed for 6 May 1998, respectively; and **c)** and **f)** are the pitch and wheel speed for 1 January 2003, respectively.

roll and yaw error response strategy is less obvious, as optimal steering of the spin axis overall must respond to both orbit precession and the coupling of roll and yaw. The numbers were developed and tested by OSC based on extensive (many orbits and days) simulations. Different constants are used in case of roll and yaw errors greater than 5°, but that case is not relevant for normal operation.

3.2.2 Pitch and Wheel Torques

The formula for the $\vec{Y}$ (pitch) torque depends on the momentum wheel rate as well as the pitch error and rate:

$$N_y \;=\; -0.0196\delta_y - \;2.19\omega_y + \;0.00017\omega_{\rm w}, \qquad (5)$$

where N_y is the desired pitch torque, δ_y is the pitch error, ω_y is the pitch rate, and $\omega_{\rm w}$ is the wheel rate difference from a nominal rate. Here, the wheel speed feedback provides the torque bias to maintain the same average wheel speed over the long term. Pitch and pitch-rate feedback provide torques to contribute to the pitch angle control, however, the overall pitch angle is also controlled with the separate pitch wheel speed control loop.

Torques are applied to the momentum wheel to help control the spacecraft pitch. As the wheel is sped up or slowed down, momentum is transferred to the spacecraft body about $\vec{Y}$. The formula takes input from the pitch errors and pitch rates, and also the current wheel speed:

$$N_{\rm w} \;=\; 0.542\delta_y \;+\; 24.6\omega_y \;+\; 0.0037\omega_{\rm w}, \qquad (6)$$

where $N_{\rm w}$ is the momentum wheel torque.

The above formula implies that for steady state conditions (where no wheel speed changes are needed, and the pitch rate is zero), the pitch error will be correlated with the wheel speed. This is indeed the case, as illustrated in Fig. 29. This shows pitch and wheel speed for three data spans representative of different phases in the mission. Note there are different vertical ranges in the plots for the different days. The first span was during the early mission period when subsolar disturbances caused large changes in the wheel speed around the subsolar region. The second span shows the typical pitch and wheel speed variation during a GAC span after the initial bias adjustments to reduce the subsolar disturbances, but before the control gain changes (Sect. 3.3.4). The third span shows

the typical behavior with wider pitch and wheel speed variation, which was typical throughout most of the mission after these control gain changes. In each case, a strong correlation is present with the pitch changes mirroring the wheel speed changes.

3.2.3 Coil Commands Torque Offset

An important aspect of momentum management control by coil generated dipoles, is that the ideal desired torque direction is generally not possible at a given position in the orbit. Magnetic torques can only be applied perpendicular to the Earth's magnetic field vector, $\vec{B}$, where the generated dipole tends to align with $\vec{B}$ like a compass needle. Because of this constraint, the actual torques generated can be offset from the ideal desired torques.

The ACS computes the coil commands as proportional to the vector cross product of $\vec{B}$ (more precisely, the measured value, $\vec{T}$) with the desired torque vector:

$$\vec{C} \;=\; \frac{\vec{T} \times \vec{N}}{\vec{T} \cdot \vec{T}}, \tag{7}$$

where $\vec{T} = [T_x, T_y, T_z]^T$ is the TAM-measured field (averaged over TAM-A and TAM-B when both are in use) and $\vec{N} = [N_x, N_y, N_z]^T$ is the desired torque vector. This gives a vector that is perpendicular to the Earth's magnetic field. Because the generated dipole can only provide a torque about an axis perpendicular to the Earth magnetic field, in this way the coil use is more efficient. Any component of the desired torque, however, which happens to lie along the $\vec{B}$ direction is lost. Control is maintained over the long term because $\vec{B}$ varies in direction over each orbit and throughout the day.

This momentum management approach is known as the *cross product law*, and some controllers using this approach set a minimum separation between the desired torque axis and the $\vec{B}$ direction, avoiding conditions where the control geometry is poor. The OV-2 control system does not do this, however. The consequences for regions of poor control geometry are discussed in Sect. 3.5.1.

It turns out that, while most of the time, desirable torques are obtained, the actual torque applied can be nearly perpendicular to the desired direction as computed by the formulas in Sect. 3.2.1. Even so, some of the torque effect may be desirable, but the torques will tend to get very noisy as the attitude and rate measurement noise propagate into the computed commands. Typically for OV-2, this means that when a roll torque is needed near the Earth's equator (where the roll axis tends to be along the $\vec{B}$ field direction), OV-2 may actually get a torque primarily in pitch or yaw that is also noisy.

The most relevant command for the roll–yaw control turns out to be C_y, because $\vec{B}$ tends to stay near the $\vec{X}$-$\vec{Z}$ plane for the nominal OV-2 orbit and orientation. If $\vec{B}$ happens to be along $\vec{Z}$ (as is typical near the poles) then a desired roll torque is also achieved by the C_y command, and if $\vec{B}$ happens to be along $\vec{X}$ (as typical near the equator) then a desired yaw torque is achieved using the C_y command. A dipole along the $\vec{Y}$ axis does not affect pitch or the momentum magnitude, but acts to precess the momentum axis in the desired direction. In addition, the C_x and C_z coils contribute to the momentum management, as either one (or both) provides efficient torques about the $\vec{Y}$ axis.

3.3 Control Adjustments

This section discusses adjustments performed for the onboard control. The main goal of the changes was reducing the overactive torques being applied to the spacecraft, thereby reducing the transient motions induced in the OV-2 attitude. Reducing disturbances was a key goal so that the ground-computed attitude could track the onboard attitude more reliably. Changes for pitch control that were supported early in the mission by OSC engineers are noted as well. Changes made to the control gains via table loads are summarized in Table 13.

3.3.1 Pitch Gains Change

There was a change to the pitch control loop gains very early in the mission before the spacecraft underwent orbit raising. It is noted here for completeness even though the change was prior to the operational period for science data collection. It might be taken as an indicator how sensor noise levels were underestimated in modeling prior to launch. Initially, the wheel torques were hitting maximum values in both directions, so control output gains were cut in half for the C_y coils and the wheel torque command. After further analysis for dealing with the wheel saturation–chatter problem, the output gains factors were set back to 1.0 and input gain factors were selected (on 28 August 1997) of 0.08 on the pitch rate feedback to the wheel torque, and 0.5 on the wheel rate feedback.

There was discussion during the first year of operations that it might be beneficial to reduce the wheel response further, but this was not considered critical and not implemented. A possible benefit of looser wheel control would be slower response to transient pitch measurement errors.

3.3.2 Rate Feedback Increase

A few days after operational data collection was initiated, an increase in the rate feedback for roll and yaw was implemented by OSC engineers in an effort to more quickly reduce nutation, which was being triggered by large yaw errors at the subsolar point (Sect. 2.4.1). While some more rapid nutation damping may have been achieved, there was also much larger attitude motion in response to the attitude error at the subsolar point. In addition, the coil output saturated even more frequently than before the change. This is illustrated in Fig. 30. Note that after the gain

Table 13. Summarized here are the changes made to the control gains via table loads.

Date of Change	*Control Parameters*	*Adjustment Factor*	*Description of Change*
28Aug1997†	Pitch rate input Wheel speed input	0.08 0.5	Wheel chatter fix (Sect. 3.3.1). This was made permanent in EEPROM around November 1997.
20Sep1997	Roll rate input Yaw rate input	5.0 5.0	Rate feedback increase (Sect. 3.2).
23Sep1997	Roll rate input Yaw rate input	1.0 1.0	The rate-feedback gains were returned to the at-launch values.
08Oct1997	Roll input Yaw input	0.1 0.1	Roll–yaw feedback reduction (Sect. 3.3.2).
07May1998	Coil output for $\vec{X}$, $\vec{Y}$, and $\vec{Z}$	0.25	Coil output reduction (Sect. 3.3.4).

† The exact date is not certain for this change.

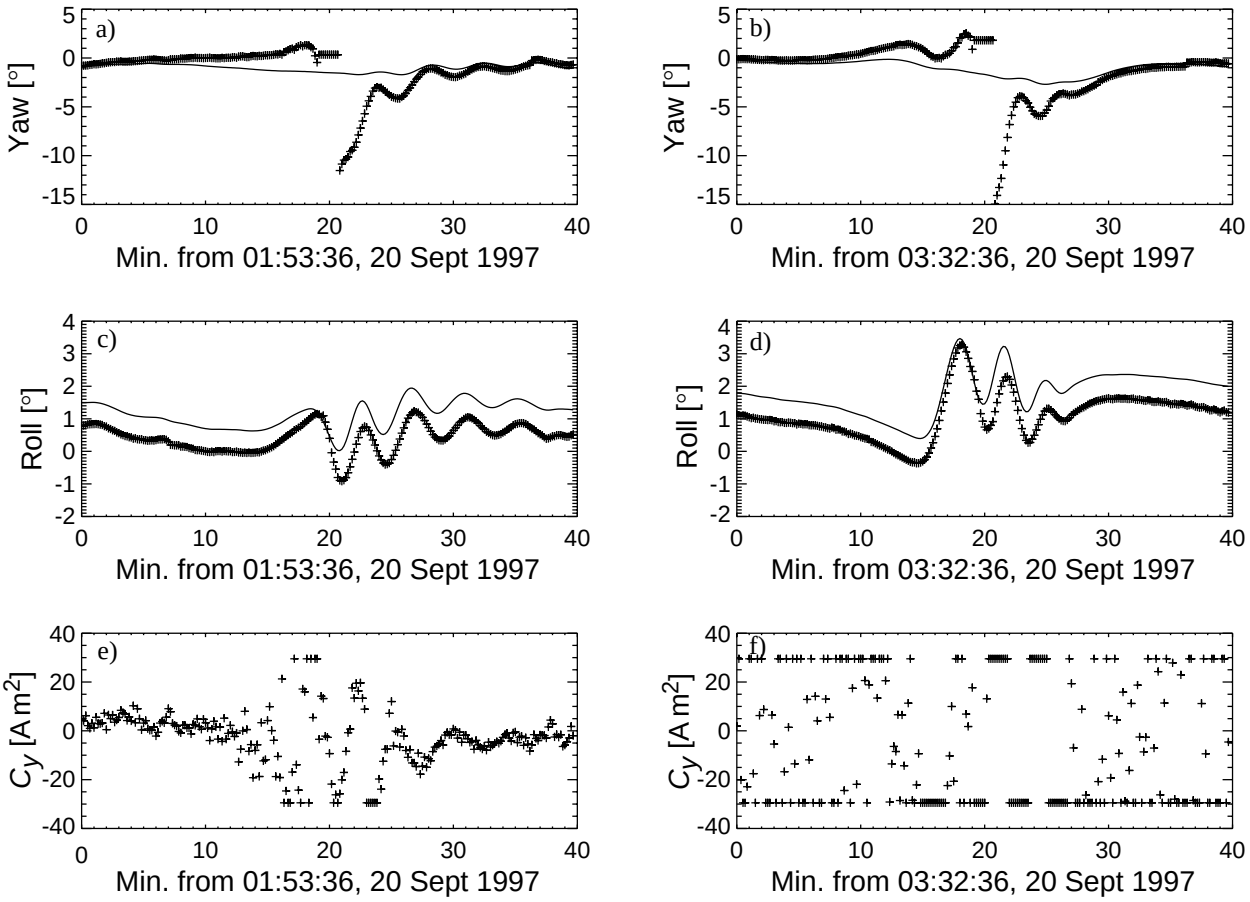

Fig. 30. Effects of the 20 September 1997 rate feedback increases. Data from consecutive GAC scenes just before and after the change are shown, respectively, for **a)** and **b)** yaw, for **c)** and **d)** roll, and for **e)** and **f)** C_y coil commands. Note that nutation damping was more rapid, but the coil commands often were maximum, and subsolar disturbances were larger with this change.

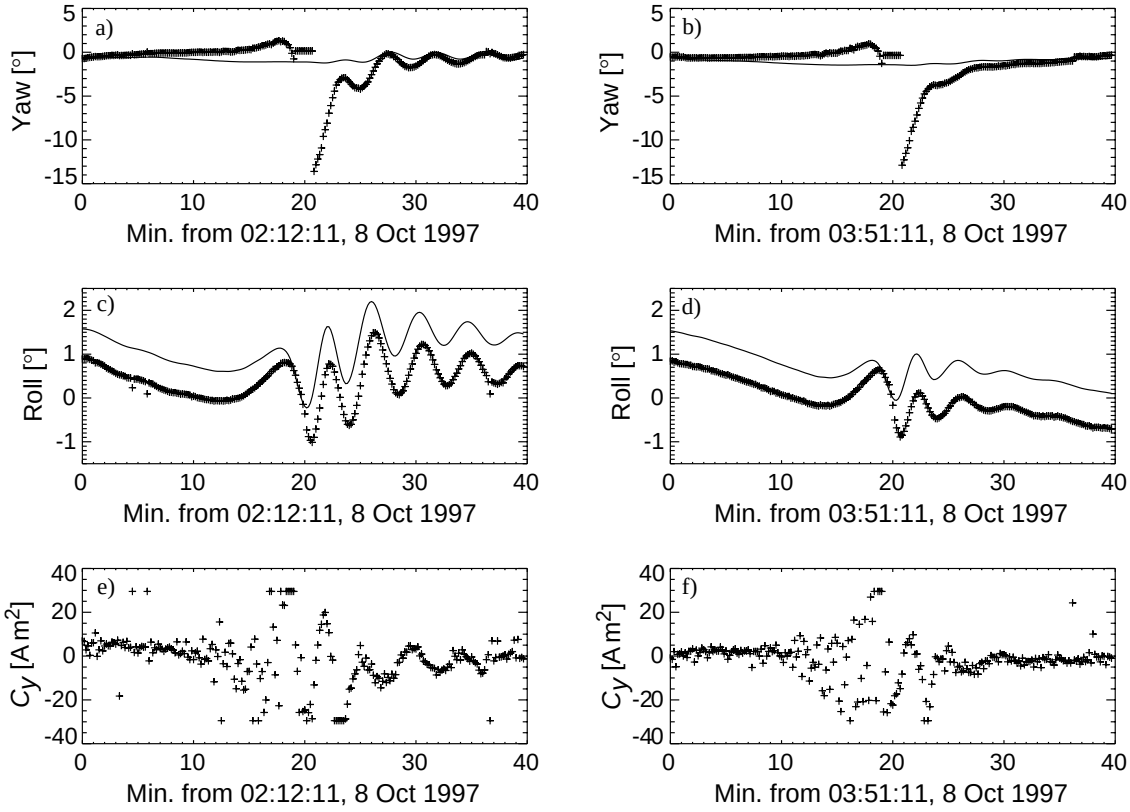

Fig. 31. Effects of the 8 October 1997 roll and yaw feedback reductions. Data from consecutive GAC scenes just before and after the change are shown, respectively, for **a)** and **b)** yaw, for **c)** and **d)** roll, and for **e)** and **f)** C_y coil commands. Note that less nutation was generated in the subsolar passage from the yaw error and the nutation was damped more quickly.

change, the roll motion was larger before being damped out, and very noisy maximum torque commands persisted throughout the orbit. This was not judged to be desirable behavior, and so the roll and yaw rate feedback was returned to its nominal value on 23 September 1997.

3.3.3 Roll–Yaw Feedback Reduction

An adjustment to the roll–yaw control was implemented on 8 October 1997, which proved effective—a reduction of the roll and yaw feedback by a factor of 0.1. The results are shown in Fig. 31.

Note that the nutation damped more quickly without a significant level of change in the torque commands. With this adjustment, roll–yaw feedback apparently interfered less with the rate feedback that acted to reduce the nutation.

This control change was loaded 2 d before the first HS bias adjustment, which reduced the large yaw error at the subsolar point and also helped to reduce the nutation generated at the subsolar point. Further tuning of biases, as discussed in Sect. 2, also reduced errors at the subsolar region, but disturbances to the attitude were still readily generated in this region, as is discussed in more detail in Sect. 3.4.1. It was observed that the coils would generate maximum commands around the subsolar region, and apparently, the torque commands were driving some of the control disturbances. Bias errors at the subsolar point contributed to the control disturbances initially, but even as the bias errors were reduced, control disturbances were often still present. This sensitivity remained a source of concern that helped motivate investigation of further gain reductions.

3.3.4 Coil Output Reduction

Nine months after launch, the roll–yaw control gains for the coil commands were adjusted to reduce the overactive torque commands near the subsolar point. This was decided in consultation with OrbImage, OSC, and SeaWiFS Project support engineers. It was argued that gyroscopic stability of the momentum wheel should keep the spacecraft fairly stable, and control did not need to be so active, especially in response to transitory errors. Following a series of bias adjustments to the DSSs and HSs on consecutive days, the roll and yaw feedback to the control torque algorithm was reduced by a factor of 0.25 on 7 May 1998. The resulting onboard attitude and coil activity for sample data spans before and after this change are shown in Figs. 32 and 33, respectively.

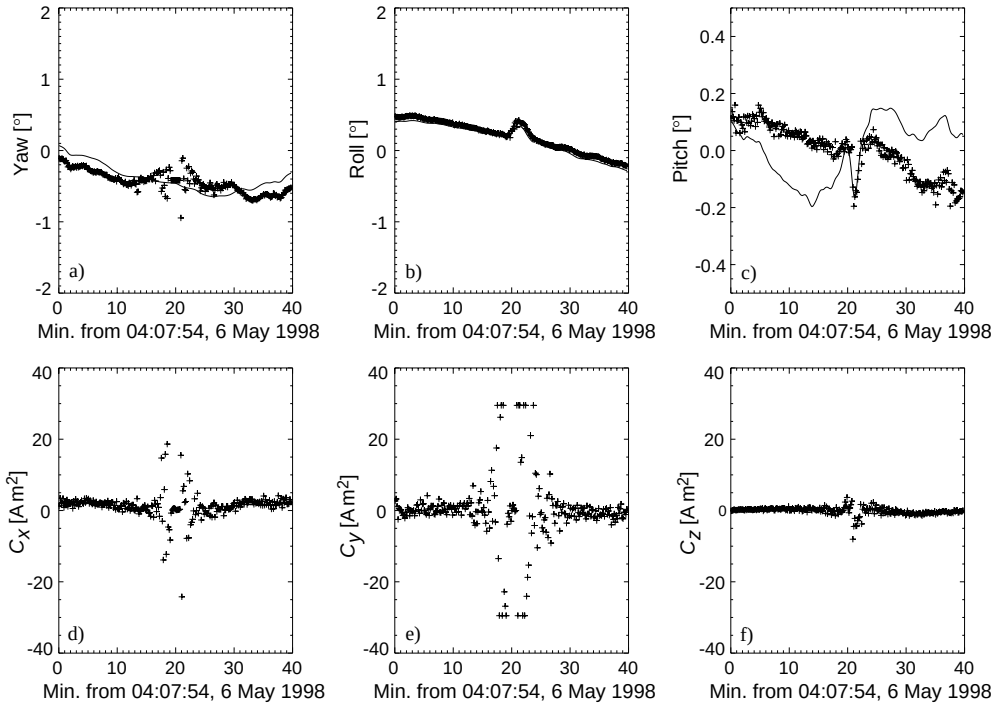

Fig. 32. Onboard attitude angles and coil commands (+) on 6 May 1998 from a representative GAC span before the coil command reductions: **a)** yaw, **b)** roll, **c)** pitch, **d)** C_x command, **e)** C_y command, and **f)** C_z command. Ground-computed yaw, pitch, and roll are also plotted (lines) in **a)**, **b)**, and **c)**, respectively. Note the C_y command is particularly noisy and often reaches maximum levels near the subsolar point in the middle of the GAC span where roll motion was generated.

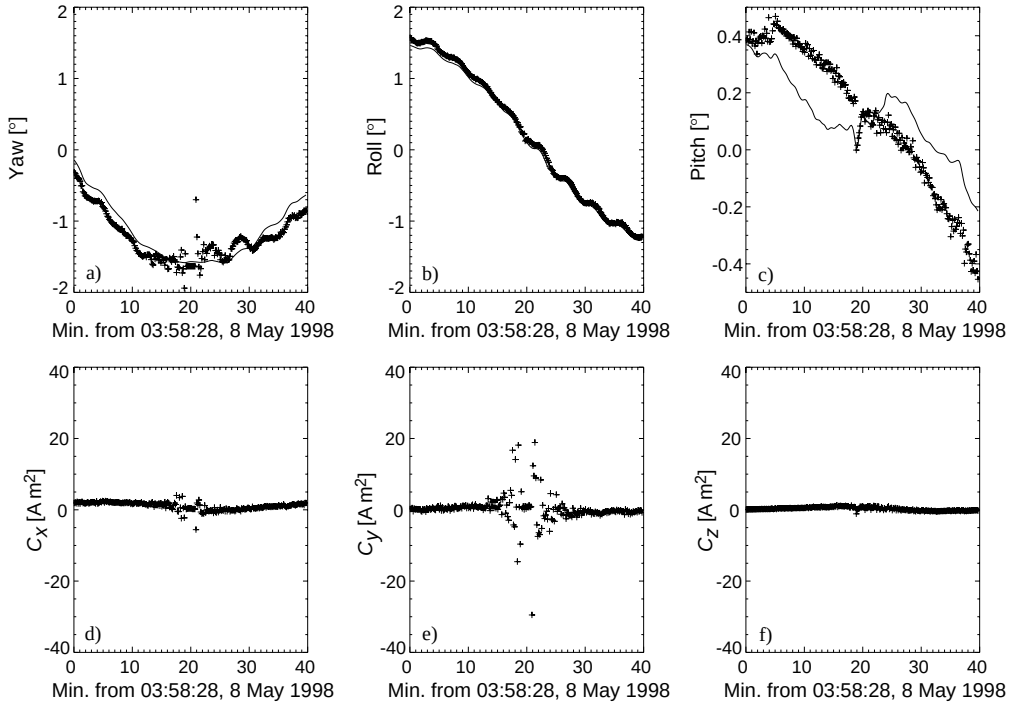

Fig. 33. Same as in Fig. 32, but for 8 May 1998, where **a)** yaw, **b)** roll, **c)** pitch, **d)** C_x command, **e)** C_y command, and **f)** C_z command. Note that larger yaw and roll angles result from the average spin axis moving farther from orbit normal, and nutation damping is slower, but maximum coil commands are less frequently encountered and less motion is initiated near the subsolar point in the middle of the GAC span.

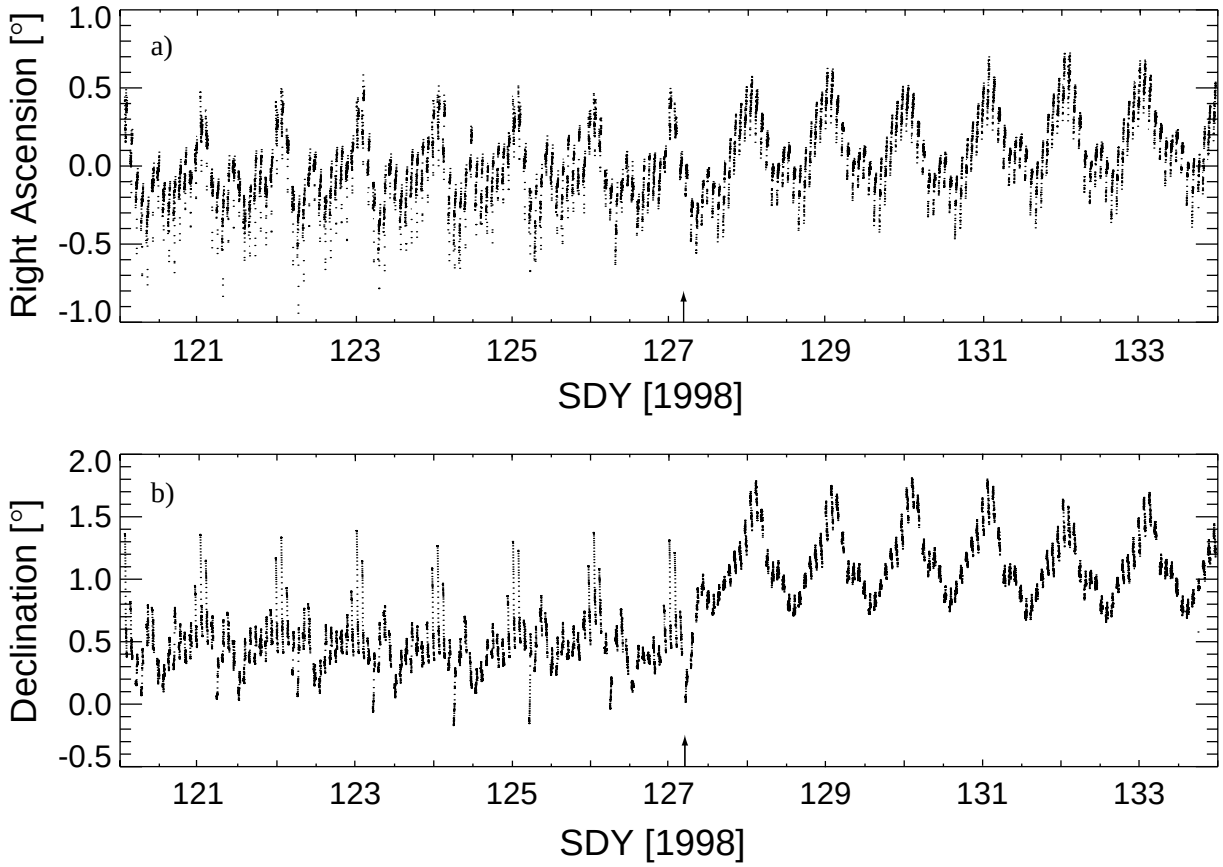

Fig. 34. Spin axis variation relative to orbit normal in **a)** right ascension, and **b)** declination, for the period around the time of the coil command output gain change, which occurred on 7 May 1998 (SDY 127) at approximately 04:30 (vertical arrow).

The coil activity following this gain change was still noisy around the subsolar point, but it reached the maximum less often. The mid-orbit roll bump disappeared; this occurred regularly at the subsolar point with much larger magnitude before certain residual sensor biases were reduced.

A significant side effect from this gain change was that the overall roll and yaw variations around the orbit grew significantly. The roll and yaw history started to roughly follow an orbit period sine wave, with yaw leading roll by a quarter orbit phase. The amplitude of this sine wave initially varied slowly from about 1.0–1.5° over the course of each day. This seemed to imply that the momentum axis was fairly steady, but had drifted away from the orbit normal direction.

Figure 34 shows the variations and dramatic shift in the spin axis position following this gain change on 7 May 1998 (SDY 127). The shift was mainly in the declination of the spin axis. The spin axis variation over particular GAC spans was much reduced, and the diurnal pattern of spin axis drift changed. Overall, the spin axis moved farther from orbit normal in a consistent direction, staying on average about 1.0–1.5° north of orbit normal. This shift created the sine–wave-like pattern in roll and yaw (Sect. 3.1 and Fig. 19). The diurnal changes in the spin axis position help modulate the amplitude and phase of the sine wave patterns in roll and yaw.

In addition, the pitch angle showed a wider range over the data collection span, correlated with wider wheel speed variations (see Figs. 29b, and 29e, as compared to 29c and 29f). The looser control of pitch, roll, yaw, and the system momentum was a natural result of turning down the overall control moments. It was a reasonable tradeoff to reduce the transient effects of errors around the subsolar region. The roll and yaw angles of this amplitude did not affect the science data collection, and the added stability of $\vec{Y}$ made the ground attitude estimation more reliable.

3.3.5 Proposed Coil Calibration Fix

A change to the ACS system that would allow the spin axis to be positioned closer to orbit normal was investigated. There was every reason to believe it would work and OSC, OrbImage and SeaWiFS Project Office engineers all agreed in principle, but it required a patch to the flight code that was never implemented.

The coil calibration was reviewed after the significance of the dipole bias was understood (Sect. 3.1.6), and an im-

portant realization was made. Because of an error in the onboard code, the prelaunch computed calibration tables for the coils were not being used on board. The calibrated value was computed, however, a simple linear scale factor was applied instead for the command. Fortunately, the calibration was still adequate for overall control, but probably contributed to the dipole bias on board by effectively producing a bias in the command.

A proposal was made to patch the onboard code, to allow a bias to be added via the calibration tables loaded from the ground. A ground-commanded bias could compensate for any dipole bias present. In fact, the bias could have been tuned to help provide some desirable spin axis precession in right ascension, so lower overall control torque activity would have been required. At this point in the mission, however, problems were encountered with keeping the FlatSat spacecraft simulator operating so patches could be tested; resources were not available to pursue the patch.

3.4 Adjustments Affecting Stability

A number of adjustments affecting onboard attitude determination (Sect. 2), were done with the primary purpose of improving the stability of OV-2, i.e., reducing the amount of unnecessary motion and nutation. Section 3.4.1 briefly reviews effects of these changes with regard to their effect on attitude estimation, and Sect. 3.4.2 describes the effects of certain changes on nutation variability.

3.4.1 Attitude Effects Overview

Several ACS attitude estimation changes that affected stability are summarized below.

Subsolar error signature reduction (Sect. 2.4.1): The first bias load was an important change done to eliminate the initiation of large levels of nutation around the subsolar region. Subsequent tunings gradually improved the situation further, but seasonal changes needed tracking and other efforts were needed to reduce the extreme sensitivity to errors in the subsolar region.

Sensor handover adjustments (Sects. 2.4.4 and 2.5.3): At transitions in the AOS and LOS of the DSSs, pitch motion would be initiated in response to new values of the pitch solution. This was not as much concern as the nutation-generating errors, especially after it was understood that ground attitudes could adequately track this motion. After the switch to single-string operation, these events were given attention in the DSS coverage gap.

TAM bias adjustments (Sects. 2.6.2 and 2.6.3): It was noticed that large yaw attitude errors resulting from TAM calibration or magnetic field modeling errors could initiate nutation in the back orbit that continued into the GAC scenes. The bias adjustment in October 1998 helped marginally, however, the errors which varied diurnally and seasonally could not be readily fixed. Fortunately, the levels of nutation coming out of the back-orbit span were usually not large enough to be of concern.

TAM weighting reduction (Sect. 2.6.6): This change significantly reduced the diurnal variability of the subsolar error signature. The bias tuning was applied more evenly through each day, and nutation was not triggered on a regular basis during parts of each day.

DSS FOV limitation (Sect. 2.5.2): This change eliminated occasional outliers at the very edge of the FOV, which caused nutation. (Outlier effects on torques are discussed further in Sect. 3.5.2.)

Yaw hold expansion (Sect. 2.2.4): This provided a very important final reduction in subsolar disturbances by reducing the sensitivity to the worst parts of the poor geometry regions.

3.4.2 Nutation Variability

Specific control- and attitude determination changes have played a key role in reducing the variability of the roll–yaw attitude disturbances, although some random element to the motion remains. With the coil commands being noisy occasionally, this is not surprising. What is perhaps surprising is the regularity of the effects in spite of the noise. With the subsampled GAC data easily missing brief transients in the torques, it is not easy to analyze what causes some of the variability.

A sample diurnal variability in the subsolar error signature is illustrated in Fig. 35. This preceded the TAM de-weighting change, and shows how the subsolar error signature could vary during the day and cause noticeable nutation for some GAC spans. A sample of some continued variability of the levels of nutation after the yaw hold region expansion is illustrated in Fig. 36.

It is believed that the remaining variability results from the noisy coil commands in the subsolar region. One of the biggest benefits of the yaw hold expansion was some reduction of the random effects from noisy torques. The nature of this noise is best illustrated at high data rates. Using merged HRPT data, two samples of all the torque commands through the subsolar region before and after the extended yaw hold region are given in Fig. 37.

Because there is sometimes residual nutation when entering the subsolar region, the random torques at the subsolar region do not always hurt the nutation levels. In some cases, these torques reduce existing nutation, however, the random changes are generally undesirable for their lack of predictability.

3.5 Coil Noise and Transients

This section further discusses the causes and effects of the noisy coil commands and their role in attitude disturbances.

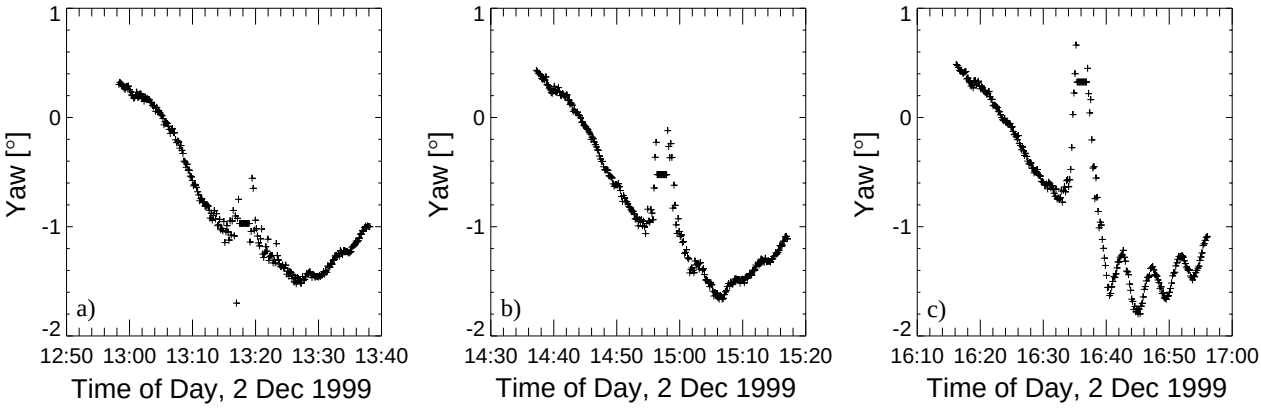

Fig. 35. Variation in subsolar error signature for three consecutive orbits on 2 December 1999.

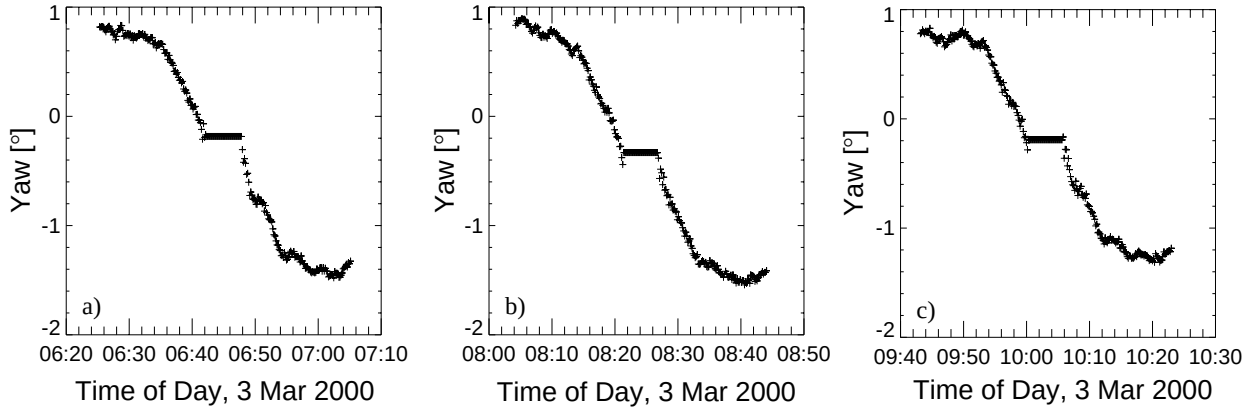

Fig. 36. Variation in subsolar error signature for three consecutive orbits on 3 March 2000.

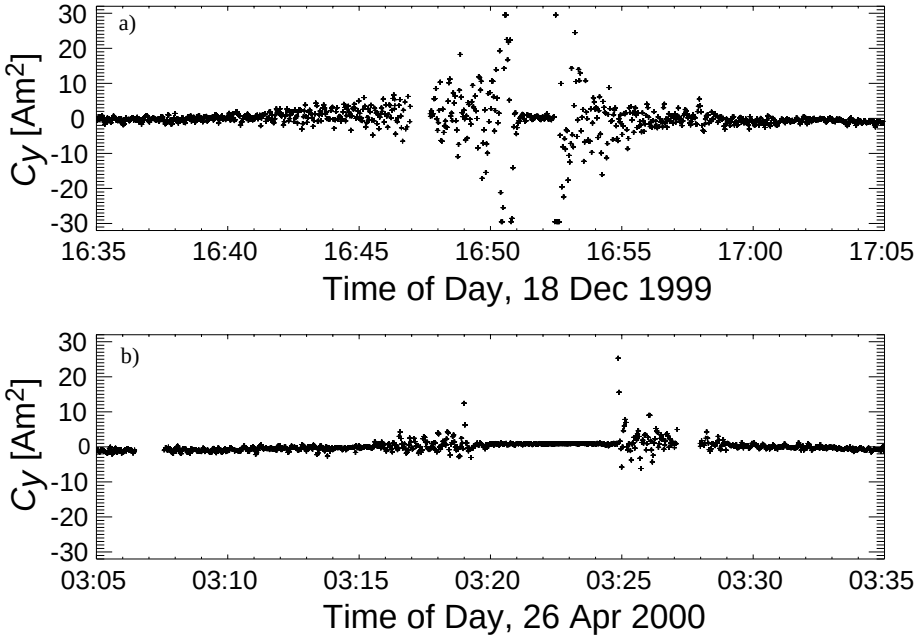

Fig. 37. Representative coil commands C_y shown at high data rates around subsolar point **a)** before, and **b)** after yaw hold extension.

3.5.1 Poor Control Geometry Effects

The coil commands are very noisy, especially C_y, as previously illustrated in sample plots. The subsolar region is especially noisy because of two effects:

1) Poor roll control geometry (Sect. 3.2.2); and
2) Noise in the yaw angles from the poor attitude estimation geometry (Sect. 2.2.3).

In the shadow period, the first of these is important, particularly for certain orbits every day. In some shadow periods, the coil command noise is highly amplified even where the attitude determination geometry is not particularly bad. This is thought to happen because in the Earth's shadow, yaw estimates are generally noisier, because of the reliance on TAM data. The quantization in the TAM data causes noticeable noise in yaw, however, the real issue is the resultant noise in the yaw-rate estimates.

Figure 38 shows the C_y commands during a sample back-orbit span, for which the yaw angle measurements were also illustrated in Fig. 15. The coil commands are very noisy even though the yaw angles are not. These noisy commands are found to be highly correlated with the noise in the yaw rate estimates (Fig. 39) during the noisiest section of the shadow period span (around 01:30). The data from a 4 min span were used for this plot. A one-sample (2 s) time lag between the rate data and the command data had to be accounted for to show the strong correlation, because the coil command computation (Sect. 3.2) relies on the rate computation from the previous control cycle.

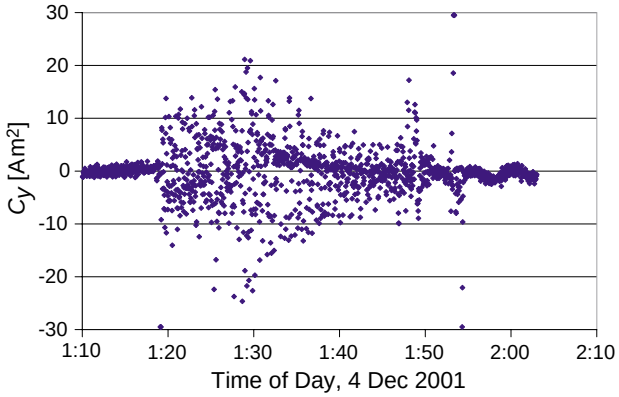

Fig. 38. Coil command C_y for high rate back-orbit span on 4 December 2001.

The two shorter periods of noisy commands toward the end of the shadow period in Fig. 38 can be associated with features in the yaw attitude (Fig. 15):

a. At 01:49, there were transient errors in the TAM data as were illustrated in Fig. 17.
b. At 01:53, there were large yaw errors in the region of poor yaw attitude geometry near the north magnetic pole, just before the shadow period ends.

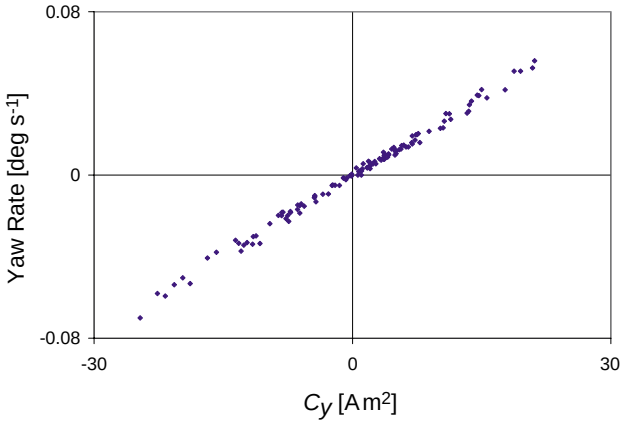

Fig. 39. Coil command C_y correlation with yaw rates for the 4 min period of 01:28–01:32 on 4 December 2001.

The same type of correlation as shown in Fig. 39 can be demonstrated for C_y commands near the subsolar region, except the slope is opposite because of the different field direction in ACS coordinates. These results indicate that the yaw rate feedback dominates in creating noisy conditions in the C_y commands.

Reduction of the yaw rate feedback, or further reduction of the C_y command output, is probably not desirable, as it is a key aspect of nutation damping, which is effective, especially at high latitudes. Adjustment of the rate filtering parameters was considered as a way to reduce the noise, but this would require careful analysis and may not provide much benefit. The fact is that the very noisy commands do not seem to cause much disturbance in the back orbit. The random fluctuating values seem to have little effect, as sustained torques in one direction are needed to initiate significant attitude motion or nutation. The noise adds some random level of motion, but the amount is not problematic.

3.5.2 Coil Transient Effects

Often the coils show momentary large commands in response to transient errors in the onboard attitude estimation. This is a source of disturbances in the actual attitude; it can initiate nutation because of sustained torque in one direction. These torques do not average out like the noisy data just discussed; thus, understanding the general effects of events that trigger these transients in the coils is helpful. Many of these types of events are anomalies (Sect. 4), but some routine examples where the coils affect the spacecraft motion are described below.

Attitude rate data are important input to the coil commands, and the onboard rate filtering has an effect on the persistence of the coil commands following a jump or a spike in the onboard attitude. This is illustrated in Fig. 40 at a high data rate for two cases: a) a jump in yaw (from 12 April 2000), and b) a spike in roll (from 3 March 2000). The respective attitude changes (Figs. 40a and b), are

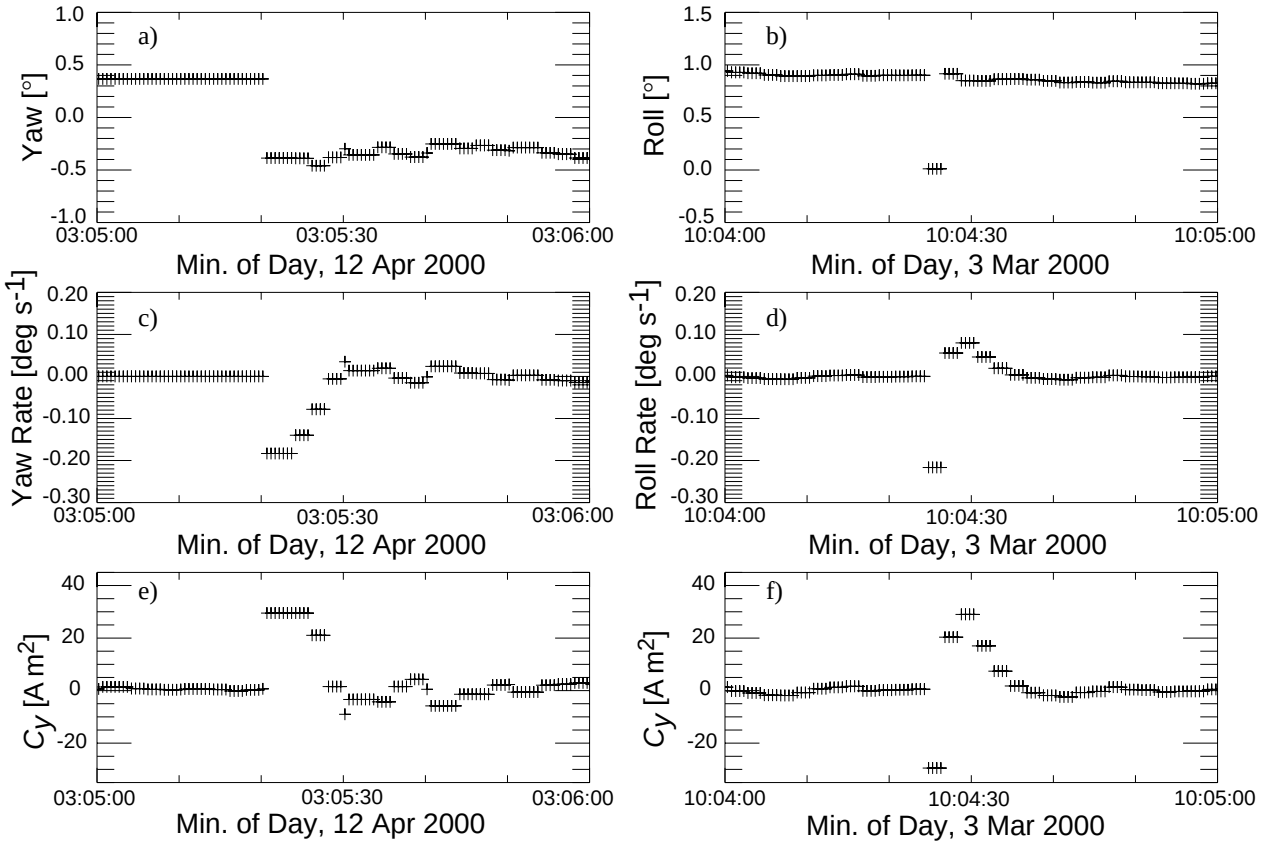

Fig. 40. Onboard computed rates and coil command responses to a jump or spike in attitude: **a)** sample jump in yaw, **b)** sample spike in roll, **c)** rate response to jump, **d)** rate response to spike, **e)** C_y response to jump, and **f)** C_y response to spike. These high data rate samples from HRPT data typically show four samples during each 2 s control cycle. Note that a couple of isolated symbols plotted in **c)** and **e)** are because of HRPT transmission noise and can be ignored.

shown along with the corresponding rates (Figs. 40c and d), and the effect on the C_y coil commands (Figs. 40e and f). The C_y commands are closely correlated with the associated computed rates, except the maximum coil command magnitude of 30 A m^2 is reached for rates over some threshold that depends on the control geometry.

The first case shown occurs after the subsolar yaw hold region, and is a common situation after the yaw hold duration was expanded. Seasonal adjustments for this condition were discussed in Sect. 2.4.5 and illustrated in Fig. 12a. Because a filter is applied on board to smooth the rate data, the discontinuity in yaw shows up in the yaw rate for several samples (Fig. 12c). The high computed rates cause maximum C_y coil commands in one direction for three control cycles (Fig. 12e).

It is noteworthy that the rate transient from the yaw jump does not always increase the nutation level, and in some cases, reduces existing nutation. In particular, during June and July there is often a noticeable level of nutation of about 0.1–0.2° amplitude at the start of the GAC span (due to regular disturbances in the back orbit), and this nutation is more often than not reduced as a result of the yaw error jump near the subsolar point. This is a reason that the further seasonal adjustments to the HS were not pursued during these months.

The second case shown (Figs. 40b, d, and f) is for a roll spike (a single outlier data point) which was, in this case, due to a DSS FOV edge anomaly (Sect. 4.1.1), but a variety of anomalies can cause such transient spikes. In the case of a spike in attitude, the rates jump first in one direction and then the other direction, but the second jump has a dominant effect overall because it lingers in the filtered rate estimates. While not a cause of major nutation events, random attitude errors due to various relatively small sensor anomalies, contribute to some of the variability in the nutation levels.

Because these transient effects on the torques typically last only three or four control cycles, the torque spike is often not seen at all in the routine GAC data, which only samples every fifth control cycle, or every 10 s. The only indication that such an event occurred is a slight change in the nutation levels. Of more concern than spikes, are attitude estimation anomalies that are sustained for more data samples. The prime examples for the OV-2 mission

are the subsolar error signatures, whose effects have already been discussed, and whose minimization was a key focus of ACS adjustments. Simulation of results from subsolar error signature is discussed next.

3.6 Simulation Results

Results from the ACS simulation tool provide some additional insight concerning on the effects of coil-generated attitude behavior. Typical motions expected of the spacecraft with nutation can be illustrated through simulations using the tool. This shows that the expected motion may be different from the onboard or ground-estimated motion in some cases, particularly because of the poor yaw observability around the subsolar region.

3.6.1 Simulation of Error Spikes

The simulation tool was used to examine the effect of single spikes in attitude. Results showed that the maximum coil commands for even a few cycles that happen following a large rate spike could initiate a small amount of nutation on the order of 0.1°. As also observed in the flight data, the rates and coil commands spiked in one direction, and then reversed direction and persisted for a few cycles. Larger rate spikes caused the coil commands to persist at high levels for longer periods.

The effects of error spikes on nutation depended on the interval between spikes. Simulation of alternating attitude values every control cycle (2 s) caused very little nutation because the effects of the steps on the rates and coil commands cancelled out. Attitude error spikes repeated at about 10 s intervals amplified the nutation, however, as each spike resulted in a net coil command in the same direction.

3.6.2 Simulation of Subsolar Effects

This section illustrates a typical example of subsolar generated nutation via simulation. It included dipole induced attitude effects and sensor biases, causing a typical subsolar error signature. An exact simulation of real flight motion was not attempted; that would have require inserting real flight data at every control cycle into the simulator—an option that was not developed. Nevertheless, the general characteristics compare very well with the flight data. The spin axis was moved off orbit normal with a simulated dipole bias, and bias errors in the HS were added to create a subsolar yaw error signature similar to that seen in flight. Both the simulated yaw and roll output and sample flight data are shown in Fig. 41, where the PC simulation of the onboard attitude computation used HS biases are consistent with those apparently affecting the flight data on 2 September 1999, so that the yaw error signature and resulting nutation disturbance at the subsolar passage were generated.

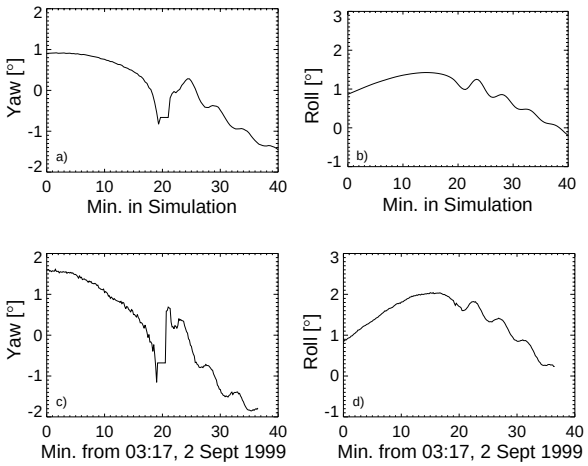

Fig. 41. Simulated estimates for **a)** yaw and **b)** roll, compared to sample GAC data for **c)** yaw and **d)** roll.

Figures 42a and b show the simulated yaw and roll *truth model* from this same simulation (as opposed to the onboard estimate of yaw and roll shown previously with the characteristic subsolar error signature). The difference of the onboard estimate from the truth model shows up mainly in yaw, because the bias primarily affects the onboard yaw estimate near the subsolar point. Figure 42c shows the path of the wheel axis in inertial space from the simulation. Note that the nutation creates a looping motion of the wheel path starting around the subsolar point, and this looping is damped out gradually toward the end of the simulated data span.

Comparison of the simulated motion near the subsolar point with typical ground attitude estimates shows how the yaw motion at the start of nutation cannot be tracked well with the current ground estimates. The poor geometry for observing yaw, and the difficulty of developing a dynamic model for modeling nutation, represent challenges for any effort to improve yaw attitude estimates in this particular region.

The simulation demonstrates that the coils and control laws play an active roll in damping any nutation that does occur. If the coils are turned off in the simulation, any nutation present continues indefinitely. In addition, any constant or slowly changing torque does not trigger nutation; thus, the effect of the residual dipole on the attitude drift does not produce disturbances. Any rapid change in the torque levels (from environmental changes or coil command changes) can, however, initiate nutation.

The simulation does not currently have as much noise in the coil commands output as the actual flight data, probably because of underestimated sensor noise. Nevertheless, the simulated attitude behavior appears to be a realistic representation of what has been observed in the flight data. Most notably, the simulation shows realistic nutation levels and nutation damping (as seen most clearly in roll).

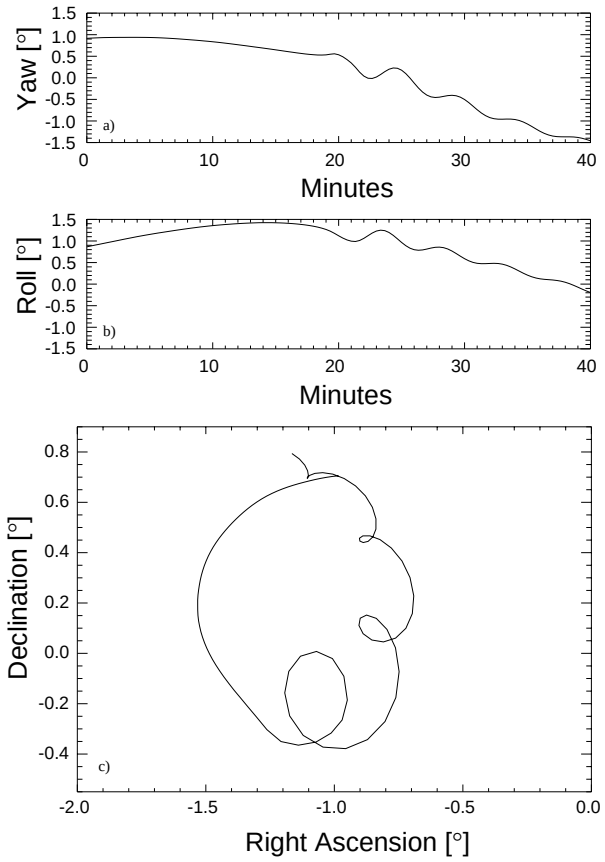

Fig. 42. Simulated attitude history truth model for **a)** yaw, **b)** roll, and **c)** spin axis path with respect to orbit normal. A dipole bias causes the large looping motion of spin axis and its offset from orbit normal, while onboard HS biases cause the nutation disturbance starting during the subsolar passage, which is slowly damped.

4. ANOMALIES

This section discusses various anomalies observed in the attitude sensor telemetry data during the mission, along with other anomalies (GPS and time tag) that affect the ACS or navigation processing. The types and instances of anomalies described are generally limited to observations made during science data collection, and therefore, should be considered representative rather than exhaustive. The following sections describe the anomalies in the DSSs, HSs, TAMs, GPS and spacecraft time tags, along with other miscellaneous anomalies.

4.1 DSS Anomalies

The three DSSs on OV-2 have shown the best overall consistency and the fewest problems of the attitude sensors. The main problem has been observed with readings near the edges of the FOV, particularly for DSS-A and -B at the gap near the subsolar point. In addition, the size of the subsolar gap increases in January and February of each year, because of obstructions on the spacecraft. These anomalies are described below.

4.1.1 DSS FOV Edge Anomalies

As described in Sect. 2, the forward and aft viewing DSSs (A and B) are mounted with their boresights angled 64° from the spacecraft zenith direction. In principle, this means that their nominal FOVs (±64°) would abut with no overlap, and there would be no interruption in coverage. In actuality, the FOV of each sensor is slightly less, with a short gap (2–3°, or less than 1 min of time) in coverage.

The DSSs have produced slightly anomalous values at the edge of the FOV with some regularity. The most common occurrence is that, in the first valid data sample adjacent to the coverage gap, the D_y angle is in error by 1–2°. The error can easily be discerned by comparing this point with the subsequent data samples in that orbit.

An example of the DSS edge anomaly is shown in a plot of the DSS-B D_y measurement angle versus latitude (Fig. 43) for a GAC scene on 14 March 2003 starting at 05:03. Note that because the spacecraft orbit is descending, the time order of the data points is from lower right to upper left. The first sample, at about −5° latitude, deviates visibly from the trend line of the subsequent points.

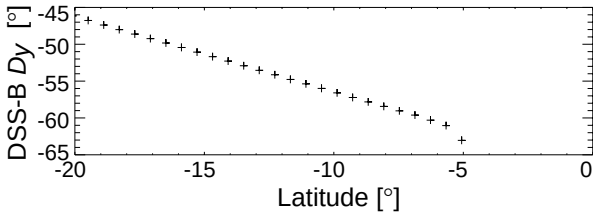

Fig. 43. DSS-B D_y samples from 14 March 2003, showing FOV edge distortion.

A second type of edge anomaly has also been observed, in which the sensor indicates a loss of sun presence for a sample near the FOV edge, and then briefly regains presence. An example of this for DSS-A is shown for a GAC scene on 11 November 2002 starting at 01:17 (Fig. 44). In this case, sun presence is lost for one point (indicated by the gap between points at the lower left corner), regained for one point, and then lost for the remainder of the scene as the sun passes out of the FOV.

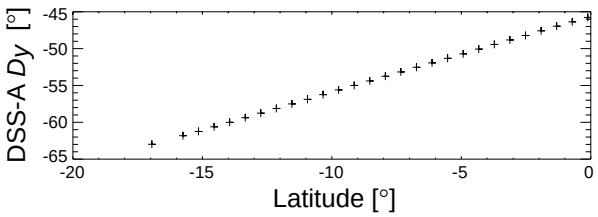

Fig. 44. DSS-A D_y samples from 11 November 2002, showing FOV edge sun presence loss.

The causes of these anomalies are unknown, although FOV edge anomalies are commonplace among optical spacecraft sensors. On OV-2, the forward and aft sensors are mounted such that the edges of the FOV are adjacent to the sides of the spacecraft, where solar arrays are mounted. This could easily lead to reflection of sunlight and distortion of the sensor output. The frequency of these anomalies has not been investigated.

This type of anomaly is easily handled in both the onboard and ground systems by setting limits on the FOV. As stated in Sect. 2.5.2, the onboard FOV has been limited for most of the mission to a radius of about 55°. The ground system limit is set to 1.7 in the tangent of D_y, equivalent to an angle of about 59.5°. This has effectively filtered nearly all occurrences of the anomaly.

4.1.2 Subsolar Gap Anomaly

The gap in the DSS coverage tends to be roughly constant over most of the year, as stated above. During January and February, however, the size of the gap increases significantly. This corresponds to the time of year when the sun angle, β, becomes positive.

The phenomenon can be seen by comparing the coverage gap at the positive β angle from January and February (Fig. 45) with the gap for the negative angles in other months. On 1 February 2000, the plot of β angle versus latitude for the valid DSS data samples shows that, at a β angle of about 6°, the gap spans more than 10° of latitude, corresponding to about 3 min in time. On 10 November, with a β angle in the range of −1 to −2°, the gap spans less than 3°, or less than 1 min.

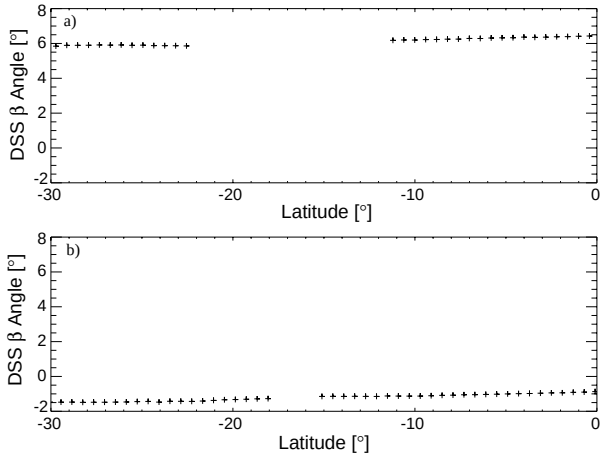

Fig. 45. The DSS β angle showing the variation in subsolar coverage gap in 2000: **a)** 1 February 2000, positive β angle; and **b)** 10 November 2000, negative β angle.

The cause of this anomaly was discovered in early 2000, when photographs of the spacecraft (taken shortly before launch) were examined. Clearly visible were thruster assemblies that protruded from the sides of the spacecraft adjacent to DSS-A and DSS-B. These thrusters were part of the avionics hardware, and were only used during the initial ascent phase of the launch. The reason why the obstruction occurs only at positive β angles is that the DSS mountings are offset toward the side of the spacecraft corresponding to the $\vec{Y}$ axis (negative orbit normal), while the thruster assemblies are approximately centered on the fore and aft spacecraft sides. The sunlight will, therefore, be obstructed when the sun direction is angled toward the orbit normal, corresponding to a positive β angle.

As previously described (Sect. 2.5.3), this anomaly actually led to the understanding of the effect of TAM data weighting in the onboard attitude estimation. It had no effect on attitude control after the sun sensor FOV was subsequently restricted and the yaw hold region was expanded.

4.2 HS Anomalies

The HSs measure the effective height of the CO_2 layer of the atmosphere at a wavelength of 14 µm. This makes them fairly sensitive to variations in the atmosphere, to other sources of infrared radiation, and to general noise effects. In the following sections, anomalies are described that have been observed in conjunction with unusually cold atmospheric conditions, moon interference, and passages through the SAA.

4.2.1 Cold Atmosphere Anomaly

In January 2001, a pattern of substantial nutation was observed for a few orbits each day in a specific time range. Starting on 16 January, one or more GAC scenes starting in the time range of 23:00–05:00 had nutation at the start of the scene (i.e., originating in the back orbit). The magnitude of the nutation often exceeded 1° peak-to-peak, and affected up to four scenes each day. This continued through 27 January, and then ended abruptly.

Although nutation had occurred occasionally coming out of the back orbit, it had not previously been observed with this magnitude and consistency over several consecutive days. Similar behavior had occurred in late 2000, from 6–9 December, for one or two orbits per day. A plot of roll angle for the first three GAC scenes on 20 January 2001 (Fig. 46) shows the consistency of the nutation coming out of the back orbit.

A review of back-orbit HS values for the affected orbits showed a pattern of anomalous values. The phase and chord samples (Fig. 47) both show anomalous values, which triggered nutation on these orbits. An attempt was made to schedule high-rate back-orbit data collection for one of the affected orbits, but the anomaly ended before the data were collected.

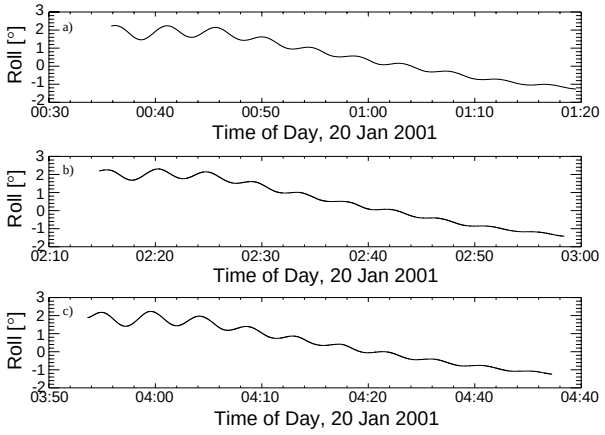

Fig. 46. Ground-computed roll angles for the first three GAC scenes on 20 January 2001, exhibit nutation triggered during the cold stratosphere anomaly in January: **a)** starting at 00:35, **b)** starting at 02:14, and **c)** starting at 03:53.

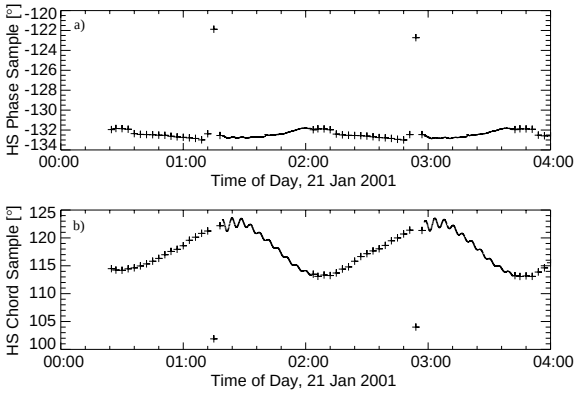

Fig. 47. HS phase and chord samples, **a)** and **b)**, respectively, from GAC and back-orbit data on 21 January 2001, showing anomalous values during the cold stratosphere anomaly.

A source of online stratospheric temperature data was identified from the Stratospheric Research Group of the Meteorological Institute, Free University Berlin. A review of this data for January 2001 showed that a series of unusually low temperatures were reported in the European Arctic region. Starting in mid-January, the temperatures dropped to values as low as −91°C. The temperature data for the first 30 d of 2001 (Fig. 48) show that the low temperatures stayed close to −85°C for the first 12 d, dropped to about −90°C by 14 January, stayed close to that value (with one exception) through 22 January, and then rose to −76°C by 29 January. Reviews of temperature data for other periods showed that the typical low values range from −60 to −80°C.

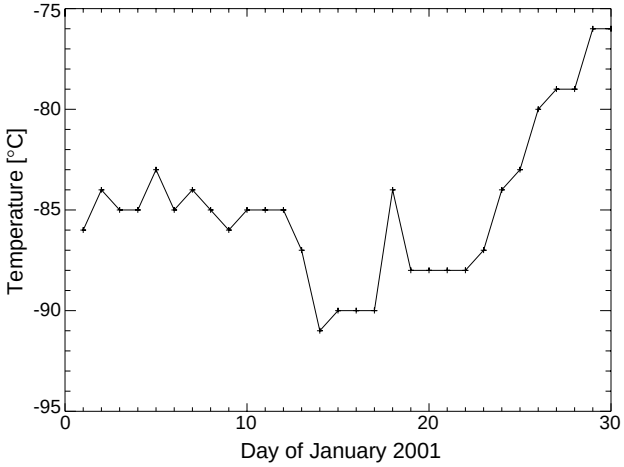

Fig. 48. Stratospheric temperature data showing a period of low temperatures from 13–23 January 2001.

Information provided by OrbImage from the HS vendor indicated that the sensor was designed to sense the Earth's horizon at 220 K, or about −53°C. Although the sensors have worked well at stratospheric temperatures well below that level, nonetheless, this information indicates that there is a limit below which the horizon crossings will not be reliably detected. In addition, it is not clear why the anomaly started 2 d after the low temperature period, and continued for a few days after the temperatures rose. One possibility is a dynamic lag between the temperatures posted by the Stratospheric Research Group and those in the uppermost CO_2 layer of the atmosphere, which affects the HSs. As stated, a similar event occurred for a few days in December 2000, and low stratospheric temperature data were also found for this period.

In December 2001, a trend of decreasing stratospheric temperatures was noted, and high-rate data collection was scheduled with the anticipated possibility of catching a similar event. Nutation was observed in a similar fashion from 17–21 December, and low stratospheric temperatures were found in this time period as well. High-rate back-orbit data collection was scheduled in this period, on 17 December. The orbits that caused the largest nutation were not the ones scheduled; however, one orbit in which a moderate amount of nutation occurred (about 0.4° peak-to-peak) was sampled in the back orbit, and gave some indication of the anomalous behavior. The HS phase and chord data from this orbit (Fig. 49) show a number of anomalous values.

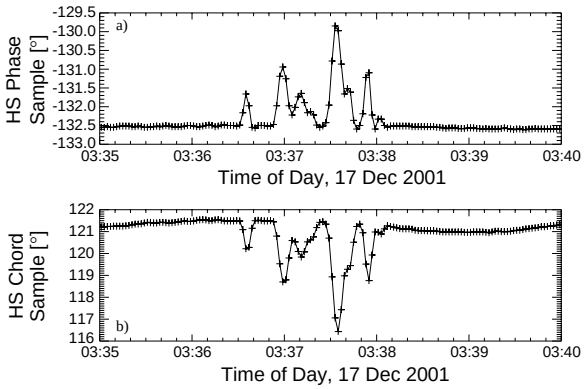

Fig. 49. HS phase and chord samples: panels **a)** and **b)**, respectively, from high-rate back-orbit data on 17 December 2001, showing a pattern of spikes.

The mechanism by which the anomaly could cause the large nutation was examined using the PC simulator. It was known that neither noisy coil activity nor large attitude errors caused high nutation. The simulation showed that a very constant large error, or a large error jumping back and forth between two values, did not trigger nutation; however, inserting large errors about every fifth sample would quickly amplify the nutation. The reason inferred was an asymmetric coil response to spikes, caused by the way rate filtering is applied. The frequency of the actual anomaly seen at high data rates showed an intermittent nature, causing peak spikes over five or six control cycles in duration. An exact simulation using the actual data was not attempted, but undoubtedly the special characteristics of the HS errors were key to causing the large nutation.

This anomaly is unavoidable, given the design of the HS and the stratospheric temperature variations. Fortunately, it has affected only a small number of orbits during the mission.

4.2.2 Moon Interference

In 1999, a series of brief roll angle excursions was noted on several consecutive orbits on 5 January. These had a magnitude of up to 0.3°, and lasted for about 1 min. They appeared initially near the start of the GAC scene for two orbits: 1) near both the start and end for one orbit, and then 2) near the end for two more orbits. The event was noted in a February 1999 review of the spacecraft performance, at which time the possibility of moon interference with the HS was raised, but not pursued.

Similar behavior was reported in July 2000, and again in November and December 2000. In each case, the events occurred a few days after a full moon. It was suspected that the moon might be viewed at, and interfering with, the horizon crossings for HS-A at that point in the lunar month. Following the December observation, this was confirmed by comparing the direction of the moon with that of the horizon crossings in the affected orbits. Examination of other periods following previous full moons showed that the roll excursions occurred regularly a few days afterward, with similar characteristics.

The typical characteristic can be seen in the roll and yaw angle for a GAC scene on 21 January 2003, starting at 12:15 (Fig. 50). This shows an instance of roll excursions at both the start and end of the scene. A small yaw excursion is also seen at the end of the scene, although this probably would not have been noticed if not for the roll excursion. Plots of the horizon in-crossing and out-crossing angles (Fig. 51) show that only the in-crossing angle is affected at the start of the scene, and the out-crossing angle at the end. This is consistent with the suspected moon interference, because the direction of the in-crossing would be close to the plane of the lunar orbit near the start of the scene, and likewise for the out-crossing.

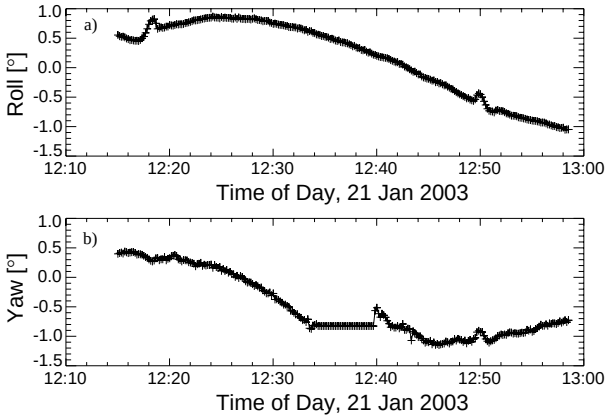

Fig. 50. Onboard angles: **a)** roll, and **b)** yaw from a GAC scene on 21 January 2003, showing effects of moon interference near the start and end of the scene.

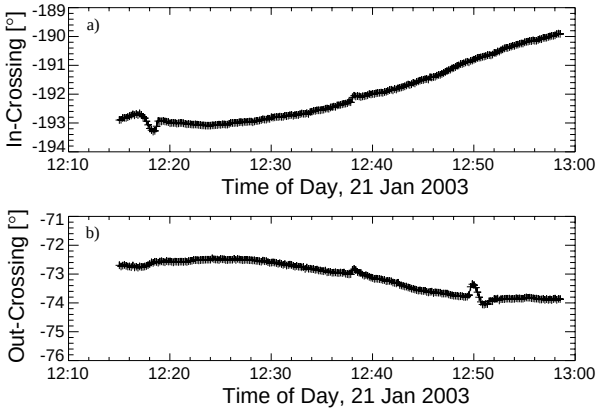

Fig. 51. HS crossing angles for this scene: **a)** in-crossing angle, with interference near the start of the scene; and **b)** out-crossing angle, showing interference near the end.

As stated above, the moon interference was confirmed by determining that the HS crossings passed near the moon

on the affected orbits. This is shown in a plot of the moon and the horizon crossings in spacecraft coordinates (Fig. 52) for the same GAC scene. The view is along the pitch ($\vec{Y}$) axis, and shows the $\vec{Z}$ versus $\vec{X}$ coordinates of the moon (solid line, clockwise motion), in-crossings (symbols on the left side of the curve) and out-crossings (symbols on the right side). This anomaly is unavoidable and has minimal effect.

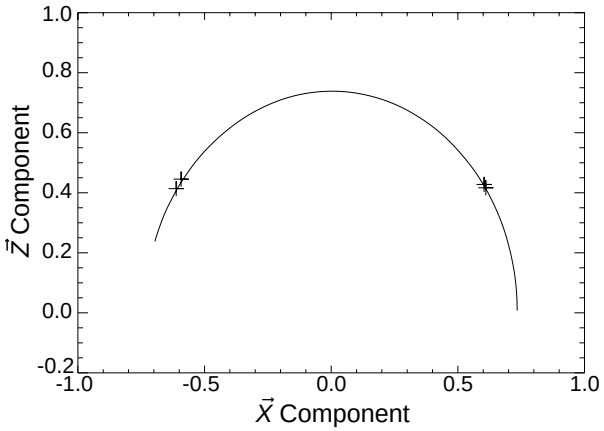

Fig. 52. Moon direction (line) and horizon crossings (+) for a GAC scene on 21 January 2003, showing the moon's proximity to the crossings.

4.2.3 HS Noise Spikes

Single-point anomalies have been observed in the onboard roll angles with some regularity throughout the mission. Although DSS FOV edge effects were suspected as a frequent cause (as described above), in some cases, these errors were ascribed to spikes in the HS data. This was noted during the 12 February 1999 spacecraft review.

After onboard control improvements all but eliminated the FOV edge effects, the frequency of single-point anomalies was reduced, but these were still seen in some orbits. In December 2000, the behavior and timing of the anomalies was investigated. These all occurred in GAC scenes collected in the time range from about 12:00–16:00, near the middle of each scene. The magnitude was typically 0.1–0.2°, but occasionally larger, up to nearly 1°.

The timing of these anomalies indicates spacecraft positions over the South Atlantic Ocean and South America. This area corresponds to the SAA region of low geomagnetic field strength (sometimes called the Brazilian Anomaly). In this region, the geomagnetic field lines curve downward toward the Earth, producing unusually large charged particle fluxes and energies at spacecraft orbital altitudes. The impact of these particles on spacecraft electronic components is a common cause of in-flight anomalies, most commonly transient events known as Single-Event Upsets (SEUs). SEUs are often manifested as elevated noise levels or more frequent anomalous values in spacecraft telemetry and instrument data. The intensity and frequency of SEUs also increases during periods of high solar activity.

A series of these anomalies was observed during an HRPT pass from the Chile station on 6 December 2000. This provided more visibility into the nature of the anomalies because of the 2 s sampling rate. A plot of the HS in-crossing and out-crossing angles versus latitude from this scene (Fig. 53) shows that the anomalies spanned multiple (three or four) data samples, and the largest effect occurred in the in-crossing angle.

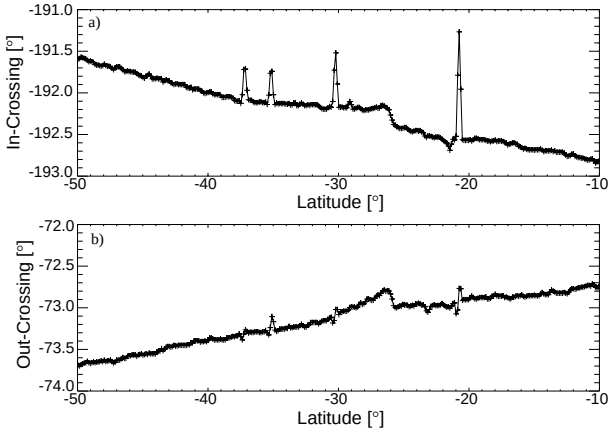

Fig. 53. HS crossing angles for a Chile HRPT station pass on 6 December 2000 showing spikes during the SAA passage: **a)** in-crossing angle, and **b)** out-crossing angle.

To verify that the anomalies were indeed associated with the SAA, all of the occurrences for one month of data (April 2003) were analyzed. As stated above, all of these were observed in the time range of 12:00–16:00. The spacecraft suborbital position at the time of each anomaly was plotted on a map (Fig. 54). This figure shows that all occurrences are in the South Atlantic region, which confirms the connection with the SAA.

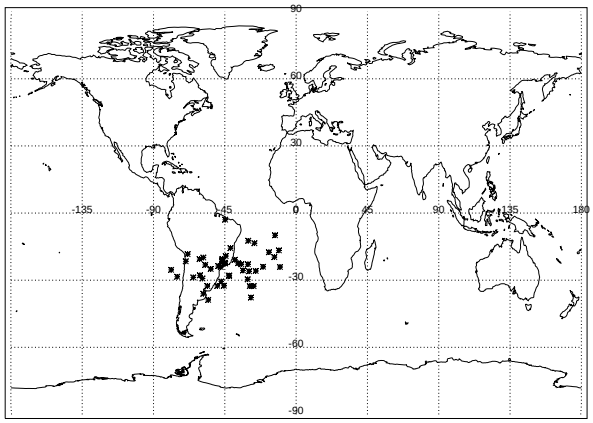

Fig. 54. Spacecraft suborbital positions for HS anomalies during April 2003.

This anomaly has had minimal effect on both attitude control and navigation processing. The anomalous points do not appear to affect the ACS, and are filtered or smoothed out during ground-attitude processing.

4.3 TAM Anomalies

The TAMs are generally the least well characterized of the attitude sensors. There are two reasons for this: TAM data are mostly used for onboard attitude determination in the back orbit, where the visibility on sensor performance is much less than during science data collection; and the data are not used for ground attitude determination, which has been a significant source of information about the sun and Earth sensors. Two anomalies have been noted in the TAM data:

- Discontinuities in the onboard calibration tables, and
- Anomalous values associated with currents generated during passages near the magnetic poles.

4.3.1 Calibration Discontinuity

The ACS applies a calibration function to the TAM output prior to use for attitude determination. The form of the calibration function changes according to the magnitude of the output. A high-order polynomial is used for output in the range of −30,000 to +30,000 nT†. Beyond this range, a simplified form is used.

This transition raises the possibility of discontinuities in the calibrated TAM output. The at-launch calibrations were continuous over this transition. When new TAM biases were uplinked on 20 October (Sect. 2.5), however, discontinuities appeared in the T_x measurements of both TAMs. It was subsequently discovered that the uplinked biases are only used in conjunction with the high-order calibration.

This phenomenon is illustrated in plots of the T_x output of both TAMs before and after calibration (Fig. 55) for a GAC scene starting at 05:15 on 20 October 1998, immediately after the biases were uplinked. The raw output (solid line) is continuous, but the calibrated output (+ symbols) shows the discontinuity in the range where the raw values are less than −30,000 nT. The T_x values typically reach this threshold on one or two orbits each day, at about 05:00.

In principle, this transition also occurs for the T_y and T_z values. It is not apparent in these axes for two reasons: 1) the T_y output never reaches the threshold because the typical magnitude does not exceed 12,000 nT, and 2) the T_z calibration was never modified with an uplinked bias, so no discontinuity is evident. This anomaly would require a modification to the onboard software to correct; however, it has a negligible effect.

† Nanoteslas.

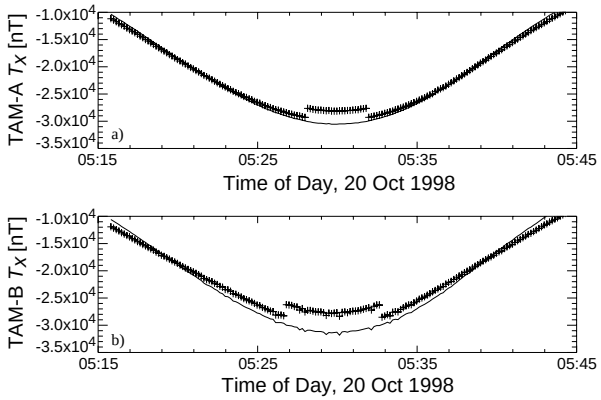

Fig. 55. TAM raw and calibrated T_x data showing calibration discontinuity, where the solid line is the raw data and the symbols are calibrated data: **a)** TAM-A T_x measurement, and **b)** TAM-B T_x measurement.

4.3.2 Birkeland Currents

The magnetic field measurements regularly show small distortions when the spacecraft passes near the Earth's auroral oval–a roughly 20° circle centered about each magnetic pole. This primarily affects the direction of the field, and appears in the T_y TAM value, which is the smallest, and therefore, most sensitive to the field direction. The distortions are visible as small changes in the TAM values in the GAC scenes; during summer of either hemisphere, the effect is more prominent, as OV-2 science data are taken closer to that hemisphere's magnetic pole. Around the equinoxes, the distortions can be seen often in both the Northern and Southern Hemisphere auroral ovals. They can occur during any orbit of the day, but data review suggests that they may occur more often in the last few orbits of the day at high northern latitudes, and in the first few orbits at the southern extremes. During periods of high solar activity, the distortions are significantly larger and tend to occur in more orbits. An example of this occurred on 24 November 2001. In this case, the distortions exceeded 2,000 nT, and were also visible in the T_x value.

These short term excursions in the TAM data were noted to occur in both TAM-A and TAM-B simultaneously in the OV-2 ACS Subsystem Review presentation in February 1999. They were, therefore, presumed to be associated with anomalies in the Earth's field, and the association with the auroral regions was speculated. Consultation with magnetosphere experts indicates these field disturbances are caused by field-aligned currents, or Birkeland currents (Schield et al. 1969) associated with aurora. This anomaly has negligible effect on attitude on board.

4.4 GPS Anomalies

The OV-2 GPS receiver and subsystem provides the orbit information for the spacecraft, which is used for onboard attitude control. The GPS also provides a precise

time reference, which is used to synchronize the spacecraft clock. The GPS data in the spacecraft telemetry are fitted to an orbit model as part of the navigation processing by the SeaWiFS Project (Patt 2002).

The GPS receiver has experienced anomalies throughout the mission. Most of these result from occurrences of a *loss of track* condition in the receiver itself. There is also an annual rollover condition near the vernal equinox, caused by a GPS subsystem software error, which has been associated with anomalies that occurred three times during the mission. The following subsections describe the initial diagnosis, analysis, and resolution of the GPS receiver problems early in the mission; the GPS resets and anomalies that have occurred during the mission; and the vernal equinox rollover problem. The spacecraft time tag anomalies resulting from the GPS resets are discussed separately in Sect. 4.5.

4.4.1 Problems and Initial Resolution

The problems with the GPS receiver started shortly after the completion of orbit raising, and were first identified by OSC in mid-September 1997. The fundamental problem was that the receiver would stop processing the signals from the GPS satellites (loss of track), resulting in orbit data gaps, and would not recover automatically. In addition, some erroneous orbit vectors were flagged as valid during these periods, causing onboard attitude errors. The GPS time signal was also degraded, causing spacecraft time tag errors.

There were two distinct problems causing the loss of track. In the first, the number of GPS satellites tracked would exceed the processing capacity, or *throughput,* of the receiver. This typically occurred when 10 or more satellites were visible. In the second, the number of satellites tracked dropped below four (the minimum for a valid solution) for an extended period, during which the propagated orbit errors would exceed the receiver's ability to recover. This condition typically happened when OV-2 was near the Earth's poles.

The initial response by OSC was to disconnect the ACS from the GPS input by disabling the automatic updating of orbit state vectors in the ACS from the GPS data. This required that orbit state vectors be regularly updated from the ground to maintain the accuracy required by the ACS. A feature in the GPS receiver was enabled to allow autonomous resets when the loss of track occurred. These changes minimized the possibility of a catastrophic spacecraft event while the problems were investigated.

A plan was developed by OSC and presented on 21 October 1997. The proposed solution had three aspects:

1) Implement ACS software patches, to filter out invalid orbit vectors and autonomously reset the GPS receiver;
2) Adjust the receiver mask angle, to minimize the throughput problems; and
3) Increase the interval at which the spacecraft clock was synchronized to the GPS time, to reduce the likelihood of GPS time errors propagating to the other spacecraft time tags.

The orbit filter was based on limit checks of the orbit height and velocity, and a measure of the GPS solution accuracy (referred to as "dilution of precision," or DOP). OSC proposed to develop and test the required software patches by late November or early December 1997. In the interim, the ACS was operated using uplinked orbit state vectors, which were generated by OSC from postprocessed GPS data.

The software patches were uplinked and verified on 14 and 15 December, 1997. The SeaWiFS instrument was powered down and stowed for this event to minimize the risk of damage, and science data collection was resumed early on 16 December. Uplinking of orbit state vectors continued until 30 January 1998, while monitoring of the GPS filtering and reset logic was performed. At that time, the GPS was reconnected to the ACS, and autonomous updating of the state vectors was enabled.

4.4.2 Reset and Orbit Anomalies

GPS receiver resets have occurred regularly throughout the mission since the software patches were uplinked in December 1997. They are observed during science data collection at the rate of one or two per month. The events are evidenced in the spacecraft data stream by multiple indications:

a. The spacecraft minor frame time tag interval varies from the nominal value;
b. The number of GPS satellites decreases to less than four; and
c. The GPS height and velocity telemetry values deviate from the normal behavior.

The Mission Operations Time Tag Anomalies Web page (Sect. 4.5) allows the times of these events to be easily determined.

A reset during the first GAC scene of 29 January 2002 demonstrates typical behavior, as shown in plots of the minor frame time tag error, number of satellites tracked, GPS height, and GPS velocity (Fig. 56). The time tag error accumulates stepwise to a total of 1 s, and is corrected in a single step about 16 min later. The number of satellites drops to zero twice, but actually appears normal for much of the anomaly. The GPS height and velocity show significant discontinuities in the middle of the period, compared to the (mostly) continuous behavior during the rest of the scene. The minor frame time tag is not actually affected by the anomalous GPS output until a few minutes afterward, and the time tag error correction occurs well after the end of the GPS anomaly.

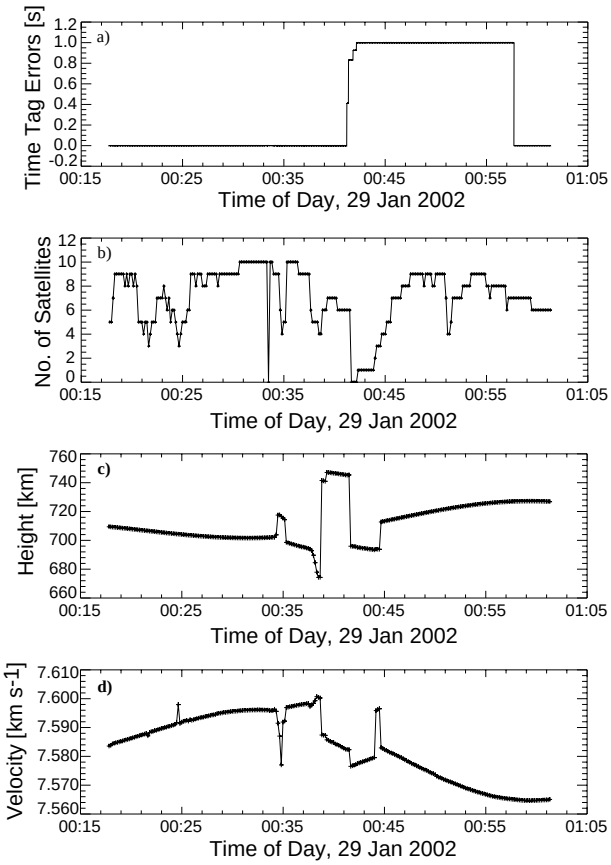

Fig. 56. Time tag errors and GPS telemetry data during the GPS reset on 29 January 2002: **a)** minor frame time tag errors, **b)** number of GPS satellites tracked, **c)** GPS height, and **d)** GPS velocity.

The GPS height represents a single dimension (the radial component) of the GPS orbit position measurements. A more complete understanding of the errors during the anomaly is gained by examining the errors in three dimensions. This can be accomplished by comparing the GPS orbit vectors in the spacecraft telemetry with the fitted orbit used for navigation processing (Patt 2002). The differences in the Earth-centered rotating (ECR) vectors are converted to along-track and cross-track components (Fig. 57). This comparison shows that the largest errors (nearly 200 km) appear in the along-track direction, several times larger than those in the height and cross-track directions.

In most cases, the orbit position errors caused by the resets have little or no effect. The onboard attitude control is minimally affected, and as stated previously, the ground navigation processing uses an orbit model fitted to the GPS vectors, which effectively filters out the erroneous values. The minor frame time tag errors do cause navigation errors, as described in Sect. 4.5.

On several occasions, however, GPS anomalies have resulted in unusual behavior of the ACS, ranging from brief transient motion to large instabilities. The most extreme instance of this was on 25 January 1999. On this date, an OV-2 safe haven event occurred when the ACS despin mode was triggered. OrbImage determined that this was caused by an anomalous GPS orbit vector that was not detected by the filter and triggered unstable behavior in the ACS. An orbit vector was uplinked, the GPS receiver was reset manually, and normal operation was resumed early on 26 January. This event caused a loss of more than one day of data.

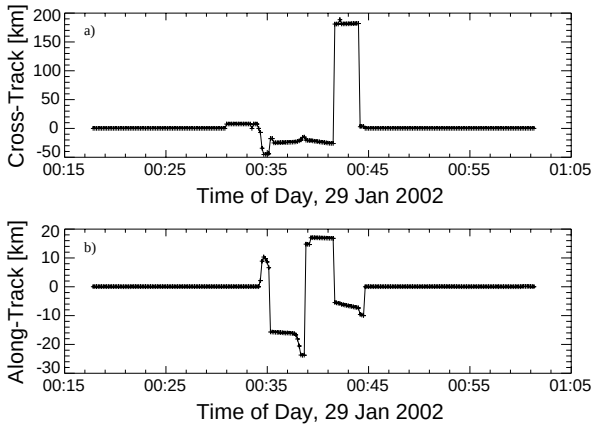

Fig. 57. Computed GPS position errors during the GPS reset on 29 January 2002: **a)** along-track error, and **b)** cross-track error.

On a few occasions, the GPS track-loss condition has not been automatically corrected by the onboard resets, and a manual reset had to be commanded by the OrbImage controller. This happened when the number of satellites tracked remained below the minimum for an extended period, or when the onboard propagated orbit diverged from the actual position. In one event on 20 July 2002, the onboard propagated orbit brought the apparent spacecraft position within 100 km of the Earth's surface. This resulted in a roll excursion of over 4°, and required several hours to recover normal attitude. A similar event (with a smaller roll excursion) occurred on 20 July 2000 at about the same time of day, but no cause for this coincidence has been identified.

The documented instances of unusual ACS behavior associated with GPS anomalies are summarized in Table 14. The time given for each event includes the total recovery time, i.e., return to normal attitude. A number of other factors are also given for each event:

- Whether or not an autonomous GPS reset occurred (Y or N, respectively) or if a manual reset was commanded (M);
- The nature of the GPS anomaly (i.e., either an erroneous GPS vector that did not get filtered and caused errors in the onboard orbit, or an extended loss of GPS satellite tracking with a significant accumulation of error in the onboard orbit); and
- The ACS behavior observed.

Table 14. ACS events associated with GPS anomalies, where "Y" indicates an autonomous reset, "M" a manual reset, and "N" no reset.

Year	*SDY*	*Time*	*Reset*	*GPS Event*	*ACS Response*
1998	355	13:35–14:15	Y	Bad vector	Roll and pitch motion
1998	364	00:00–05:00	M	Track loss	Pitch and nutation
1999	25–26	18:12–20:34†	M	Bad vector	ACS Despin—Safe Haven
2000	201	16:50–18:46	Y	Bad vector	Roll and pitch motion
2000	208	10:00–18:27	M	Track loss	Pitch drift
2001	55	22:35–22:50	N	Bad vector	Nutation and pitch motion
2001	345	13:50–14:00	N	Bad vector	Yaw error
2002	29	00:37–00:42	Y	Bad vector	Roll and pitch motion
2002	161	21:24–21:45	Y	Bad vector	Pitch motion
2002	201	16:40–00:00	Y	Track loss	Pitch and spin axis motion
2003	349	00:15–00:45	Y	Bad vector	Roll and pitch motion

† The ACS event in 1999 spanned two days, starting at 18:12 on SDY 25 and ended at 20:34 on SDY 26.

In addition to these very large GPS and onboard orbit anomalies, there are smaller errors that occur regularly, but have minor effects. Instances of these were noticed in January 2002, and motivated a more systematic comparison of the onboard orbit with the GPS data, which was posted on the telemetry SOH Web plots early in 2002. The events that are visible in the onboard attitude affect pitch only, by 0.2–0.4° over 10 min or less, occurring once or twice a month. During these events, the onboard orbit is in error by 20–50 km in the along-track direction, apparently due to a few bad GPS data points that get picked up by the onboard filtered orbit. Because of the error in the onboard estimate of position, the errors in the inertial-to-local vertical coordinate transformation corrupt the DSS data usage for the onboard attitude. The effect is relatively small, and the ground navigation processing is not affected because it filters the GPS data independently. Errors of even smaller magnitude occur in the onboard orbit—spikes of typically 1–2 km, occasionally as high as 8 km, which get corrected within a minute—but these have very minimal effects in the onboard attitude.

4.4.3 Vernal Equinox Rollover Anomaly

A special case of GPS anomaly can occur near the vernal equinox each year. It is caused by a coding error which prevents successful resetting of the GPS receiver for a period each year around the equinox. On one occasion, in 2000, this caused an extended GPS outage and significantly disrupted navigation processing.

The GPS subsystem code computes the right ascension of the ascending node (RAAN) from the orbit state vector, relative to the Greenwich Sidereal Time (GST) at the start of the GPS week (Sunday at 00:00). Both the RAAN and the GST are computed in the range $(-\pi, \pi)$. Each year, around the vernal equinox, the RAAN and the GST both roll over from π to $-\pi$. Although the nominal orbit is sun synchronous, the actual RAAN drifts slowly with respect to GST, and therefore, these two values roll over at different times. More importantly, the reference GST is only computed for the beginning of the GPS week, so the effective date of the GST rollover can be nearly 7 d after the nominal rollover. During the OV-2 mission, the RAAN has always rolled over before the reference GST.

The difference between the RAAN and GST is also required to be in the range $(-\pi, \pi)$. The problem arises in the code during reset events when this check is performed; if the difference falls outside of this range, it is corrected by π instead of 2π. This results in an incorrect calculation of the relative RAAN and prevents successful resetting of the GPS receiver. The conditions for this occur each year, between the RAAN rollover and that of the reference GST. Table 15 shows when these conditions have occurred in each year of the mission (projected through 2004). The table shows the day and time of the RAAN rollover and the actual GST rollover, as well as the next GPS week rollover, which is the effective rollover time for the reference GST. The conditions for unsuccessful resetting of the GPS receiver, therefore, have existed each year of the mission, from the date and time of the RAAN rollover to the next GPS-week rollover. In 1998, this period exceeded 7 d, and in 2004, it is projected to be about 10 d. The period will continue to increase in future years as the RAAN rollover continues to drift earlier, unless an orbit adjustment is done to reverse the drift rate.

There are two known occasions where GPS anomalies occurred during these periods, and resets were unsuccessful: in 2000 and 2001. On SDY 84 (24 March) of 2000, an autonomous GPS reset at 03:37 failed. OrbImage noted the loss of GPS tracking during the next pass, and commanded a manual reset, which was also unsuccessful. Several subsequent manual resets on SDYs 84 and 85 also failed. The GPS receiver was not tracking at any time during this period. Finally, on SDY 86 at 18:24, a reset

Table 15. GPS vernal equinox rollover conditions.

Year	*RAAN Rollover*		*GST Rollover*		*GPS Week Rollover*	
	SDY	*Time*	*SDY*	*Time*	*SDY*	*Time*
1998	80	18:49	81	17:09	88	00:00
1999	81	13:06	81	23:00	87	00:00
2000	82	01:43	82	04:50	86	00:00
2001	81	02:44	81	10:39	84	00:00
2002	80	14:02	81	16:28	83	00:00
2003	79	08:42	81	22:16	82	00:00
2004	77	16:00†	82	04:04	88	00:00

† Estimated time.

was successfully commanded and tracking resumed. As indicated in the table, this was after the rollover of the GPS week and the reference GST, although at that time, the cause of the unsuccessful resets was not yet understood.

The problems resulting from the extended GPS outage were twofold. First, the GPS orbit vectors were not useful for navigation processing; and second, the time tag reference from the GPS was not available. The lack of GPS vectors to fit to the orbit model in the navigation software (Patt 2002) meant that the orbit was propagated from the last set of GPS data through the outage (approximately 2.5 d). By the end of this period, the propagated orbit had accumulated several kilometers of error, primarily due to underestimation of the atmospheric drag. The drag coefficient in the model was adjusted manually to minimize the accumulated error during the outage, and most of the data were successfully navigated.

The time code problem manifested itself as multiple small shifts in the time code during the outage. Subsequent analysis of the navigation results with the corrected orbit showed that most of these did not significantly affect the accuracy. There were several orbits, however, where large errors remained in the time code, and could not be reliably estimated. The affected period was from about 06:00–18:15 on SDY 85, spanning eight orbits. The data from this period were excluded from further processing.

Following the anomaly, a review of the code uncovered the problem in calculating the relative RAAN at the rollover. This also allowed all of the potential problem periods to be determined as in Table 15. The table shows that the potential for an extended anomaly existed in both 1998 and 1999. A GPS track loss event occurred in 1999, on SDY 84 from about 02:00–14:00, and was accompanied by pitch angles increasing to about 4°. In this instance, the problem corrected itself prior to the week rollover, although the mechanism for this was not understood. In 2001, a BCR reset on SDY 83 at 04:03 caused the GPS receiver to be power cycled. This occurred during the problem period for that year, so the GPS receiver did not automatically resume tracking. Fortunately, the problem was understood by then, and occurred close to the GPS week rollover. A stored command was uplinked to command a GPS reset after 00:00 on SDY 84, which restored tracking.

The predicted periods for this problem will continue to be monitored as long as OV-2 operation continues. As indicated in the table, the predicted period in 2004 is 10 d, which suggests a higher than usual likelihood of a GPS anomaly in this period.

4.5 Time Tag Anomalies

Throughout the mission, the time tags from various OV-2 subsystems have exhibited a variety of anomalies. The effects of these on attitude control vary widely; many anomalies have no apparent effect, while a few have resulted in spacecraft safe haven events; time tag anomalies also have a significant effect on navigation processing by the SeaWiFS Project. As with other types of spacecraft anomalies, the visibility of time tag anomalies is largely limited to the GAC data collection period. The SeaWiFS Mission Operations Web site includes a page for time tag anomalies, with links to listings of anomalies that were generated automatically by scanning the spacecraft telemetry. This procedure was implemented in September 2000. In the following section, the spacecraft time tags that are pertinent to the ACS are summarized, followed by a description of each type of anomaly.

4.5.1 Spacecraft Time Tags

The OV-2 spacecraft data stream includes a number of time tags for various subsystems. There are three types of time tags that can affect either attitude control or navigation accuracy: the spacecraft minor frame time tag, the GPS time tag, and the ACS time tags. Each of these is described below.

The data stream contains a time tag in the header for each level-0-to-1A data record, or minor frame. (A level-0 minor frame contains either one LAC or five GAC scan lines, along with associated spacecraft and instrument telemetry.) The minor frame time tag is stored as a 40-bit word, and consists of a day counter and a milliseconds-of-day counter. The 13 most-significant bits store the day

counter, with a reference date of 13 January 1993 (Astronomical Julian Day 2,449,001). This counter has a range of 8,192 d, and will roll over on 19 June 2015. The time of day is stored in the 27 least significant bits, and represents milliseconds since 00:00, with a range of 0–86,399,999.

The GPS subsystem produces a time tag for each new GPS orbit sample. These samples are updated in the spacecraft telemetry every 10 s. The time tag consists of a series of integers which store the year, month, day, hour, minute, second, and fraction-of-second (in nanoseconds). This fraction is almost always less than .001 s, so the GPS samples typically occur *on the second* to within 1 ms. In addition, leap second adjustments to UTC are automatically incorporated to the GPS time tag within 15 min of occurrence; there has been one leap second—on 1 January 1999—since the OV-2 launch.

The time system used by the GPS satellites is actually based on weeks and time of week, where the GPS week starts at 00:00 on Sunday, and the week counter is referenced to 6 January 1980. The week counter rolls over every 1,023 weeks, which has occurred once, on 22 August 1999. This system of GPS time does not appear in the OV-2 GPS data stream, but the GPS week is used as the basis for ACS time tags.

The ACS subsystem maintains a reference time. This time is stored as an 8-byte, floating point, monotonically increasing counter with a resolution of 1 s. The reference date and time is the same as for the GPS week: 6 January 1980 at 00:00. Note that like the GPS subsystem time tag, the ACS reference time is corrected for leap seconds.

Each attitude sensor data sample includes a time tag. These time tags are stored as 4-byte integers that represent milliseconds of the week, where the week starts at 00:00 on Sunday (same as the GPS week); the range is 0–604,799,999. The ACS fields are updated every 2 s; in the GAC data, where SOH telemetry is stored every 10 s, one of every five ACS samples is transmitted; in the LAC, with telemetry stored twice per second, each ACS sample is repeated four times.

The management of the time tags in the different subsystems is complex. A Real Time Clock (RTC) in each of the onboard processors (the SCM, PSM, and GIM) is used to keep time, and is synchronized to the GPS time by the onboard software. A separate clock in the ACE is used for the attitude sensor time tags. These clocks are checked and synchronized to the GPS time by background tasks that execute typically every minute. In addition, there is a GPS Pulse-Per-Second (PPS) signal that triggers various tasks in the processors. The details of how the various timing anomalies occur within the system are not well understood.

4.5.2 GPS Reset Time Anomaly

The GPS resets described in Sect. 4.4 also result in errors in the GPS time tag, and through the time management logic on the spacecraft, these errors propagate into the minor frame time tag and the ACS time tags. The errors appear at the start of the reset event as a series of fractional-second steps, spread over about 1 min. In some cases, the steps cancel out each other within a few seconds. More frequently, they accumulate to a total of 1 s; this error remains for a period ranging from 15–20 min, and is corrected in a single step at the end of the event. In these cases, because the event duration is a large fraction of the GAC scene length, it is common to observe only the start or end of an event.

The effects of these anomalies are twofold: first, the GPS data are not useful because of the time tag errors (and also because of the other associated errors in the GPS data described in Sect. 4.4); and second, the spacecraft time management logic causes these errors to propagate to the minor frame and ACS time tags, affecting both attitude control and navigation accuracy. Since January 1998, the main problem resulting from the resets has been the navigation errors resulting from the time shifts. The response has been to flag the GPS reset periods for navigation failures during data processing (Patt et al. 2003).

4.5.3 30 s Anomaly

The minor frame time tag has exhibited an anomaly throughout the mission with the following characteristics: the time tag is shifted exactly 1 s; it is corrected 30–35 s later; no other time tags on the spacecraft are affected. The cause of this phenomenon, referred to as the "30-second anomaly," is unknown.

The occurrence of these anomalies was not actually documented until early 2001, when they were observed by the SeaWiFS Calibration and Validation Quality Control team. That the anomalies went unobserved for over three years is most likely because of their brevity and (at first) infrequency, which limited the effect on data quality. Following their discovery, the mission data set was analyzed to locate all occurrences and post them on the Mission Operations Time Tag Anomalies Web page. There were 17 occurrences in 1998, 24 in 1999, 59 in 2000, 87 in 2001, and 64 in 2002. The timing has been very irregular; in some cases, multiple anomalies have occurred in a single GAC scene, while at other times, none were observed for weeks. Geographic and temporal analyses have revealed no discernible pattern. Figure 58 shows the occurrences for the years 1998–2002; there have been two *double* (overlapping) anomalies that resulted in 2 s of error.

These anomalies only affect navigation processing. In 2002, the capability was incorporated to detect and correct the anomalies during data processing (Patt et al. 2003), which was possible because of the consistency and limited extent of the anomalies (i.e., only the minor frame time tag is affected).

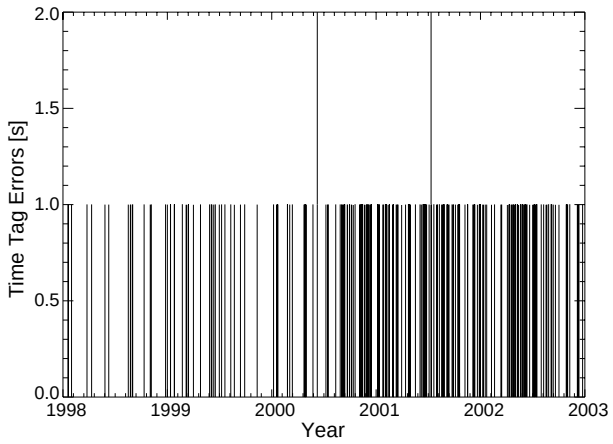

Fig. 58. Occurrences of 30 s minor frame time tag errors for 1998–2002.

4.5.4 Weekly Rollover Anomaly

There are two anomalies that affect all of the attitude sensor time tags, but no other spacecraft time tags (including the ACS reference time tags). The first of these occurs at the end of the GPS week, when the sensor time tags roll over. As stated above, this would normally occur at 00:00 each Saturday. In fact, the sensor time tags roll over 6 s before midnight, and increment at the normal rate for a short period afterward; therefore, the time tags are in error during this period by 6 s. The errors remain for a short time (10–60 s) and then are corrected when the ACE clock is resynchronized to GPS time, which is done once every minute. (The duration follows the same pattern as the 65,536 s anomaly described in the following section.)

This anomaly was first observed in December 1998, at which time it was discovered that spikes in the onboard roll and pitch angles occurred at the weekly rollover. Further investigation revealed the existence of the weekly rollover time-tag anomaly, and that this had been occurring throughout the mission. That this was only observed after more than a year of operation was most likely because of its minimal effect (at that time) and the number of more significant issues being investigated. The effects on the onboard attitude angles were observed to be transient, they were visible only when the rollover occurred during GAC data collection, and the actual spacecraft motion during the events was small. In navigation processing, the anomalous time tags were easily detected and rejected; thus, this anomaly received little attention at that time.

In late July and early August 2000, analysis of potential causes of nutation brought renewed attention to the weekly rollover anomaly. Specifically, roll disturbances were noted at the 29–30 July and 5–6 August rollover events. Shortly afterward, a spacecraft safe haven event occurred at the 26–27 August rollover, in the back orbit. Another safe haven occurred at the 9–10 September rollover, and was found to include an ACS mode change from fine pointing to despin modes. Further analysis of the 5–6 August event, which happened to occur during an HRPT pass, revealed large attitude angle and rate oscillations that could trigger a transition to despin mode.

This behavior is illustrated in plots of the time anomaly, pitch angle, and pitch rate across the day transition (Fig. 59). The first plot shows the difference between the attitude sensor (in this case, HS-A) and minor frame time tags, corrected for the day rollover. This difference is normally negative, because the sensor samples are stored in telemetry 2–4 s after they are collected. Shortly before the day transition, the difference jumps 6 s and becomes positive because of the anomaly, and is corrected approximately 50 s later. During the anomaly, the onboard computed pitch angle and rate both show large oscillations caused by the sensor time tag error.

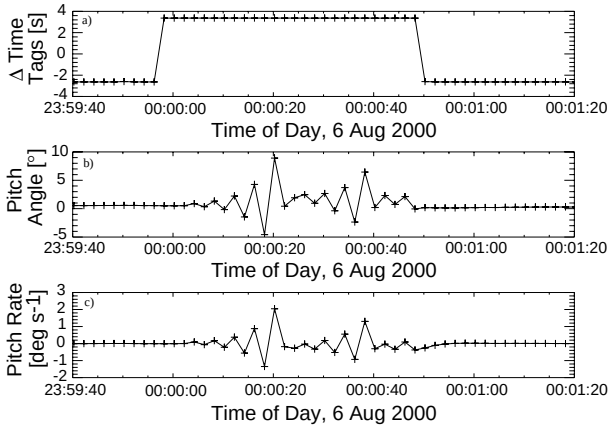

Fig. 59. Weekly rollover on 6 August 2000, showing effect on sensor time tags, pitch angles, and pitch rates: **a)** difference between HS and minor frame time tags, **b)** onboard pitch angles, and **c)** onboard pitch rates.

The cause of these oscillations was traced to the time tag error effect on sensor time alignment calculations in the onboard code. The function of this code was to align the sensor readings to the ACS reference time for the onboard attitude calculations. Because of the attitude rate estimate noise, however, large corrections were erroneously added, which in turn caused diverging rate errors during the anomaly period. A modification was proposed to an onboard table to correct the problem by eliminating this correction if the differences between the attitude sensor time tags and ACS reference time exceeded 3 s. The modification was uplinked on 22 September, and analysis of the following weekly rollover (23–24 September) using high-rate back-orbit telemetry showed that the oscillations were no longer occurring. No subsequent problems have been associated with the weekly rollover.

4.5.5 65,536 s Anomaly

The second type of attitude sensor time tag anomaly occurs each time the time tags reach an exact multiple of

65,536 s, i.e., 2^{16}. At these times, the time tags are reset to zero, incremented at the normal rate for a short time (10–60 s), and then corrected. An occasional variant of this anomaly occurs, in which the time tags are reduced by exactly 65,536 s, not to zero. The time tags have large errors during these periods. Because the sensor time tags are referenced to the start of the GPS week, these anomalies occur at precisely the same time each week. The days and times of occurrence are listed in Table 16; the cause of the anomaly is unknown.

Table 16. Days and times of 65,536 s time tag multiples.

Multiple of 65,536	*Day of the Week*	*Time of Day*
1	Sunday	18:12:16
2	Monday	12:24:32
3	Tuesday	06:36:48
4	Wednesday	00:49:04
5	Wednesday	19:01:20
6	Thursday	13:13:36
7	Friday	07:25:52
8	Saturday	01:38:08
9	Saturday	19:50:24

This anomaly was first reported in October 1998, but was not characterized until May 1999. At that time, it was proposed that the anomaly occurred at every multiple of 65,536 s each week, and had been throughout the mission. When the Time Tag Anomalies Web page was added in September 2000, the systematic occurrence of this anomaly was confirmed. Like the weekly rollover anomaly, this was not believed to affect attitude control, and was easily handled in navigation processing.

It was not until much later that this anomaly was observed to affect ACS performance when DSS data were not available. In 2002, anomalous pitch and roll values were observed in the DSS coverage gap, and were sometimes accompanied by unusual pitch motion. These events were determined to be coincident with the 65,536 s anomalies. The pitch and roll errors were, in fact, sign reversals on the angles, which occurred when the anomaly coincided with the DSS gap. When the DSS FOV limits were reduced (Sect. 2.5.2), the coverage gap during GAC data collection was increased, along with the incidence of time anomalies in the gap. Further analysis using back-orbit data showed that the sign reversals resulted any time the anomaly occurred in the absence of DSS data. The anomaly affected yaw in the back orbit as well, typically by almost 180°. Yaw is not affected in the subsolar region because the DSS gap is inside the yaw hold region.

The onboard attitude errors in GAC scenes would be inconsequential, except for the pitch excursions that sometimes result. The size of the excursions is typically a few tenths of a degree, but occasionally is larger. In 2002, on 13 July there was an excursion of 1.5°, and on 14 August, an excursion of 2.2°. The size of the excursion depends on two factors: the magnitude of the computed pitch angle, which determines the error resulting from the sign reversal; and the duration of the time anomaly. The rapid motion during these excursions is not accurately tracked during navigation processing, and the time tag errors cause the attitude sensor data to be rejected during the anomalies. Both of these factors result in larger than usual navigation errors, up to a few pixels during the anomaly.

This can be seen by comparing the attitude angles computed on board with those from the ground navigation processing for the 14 August event at 19:01:20, which occurred during an HRPT pass (Fig. 60). The onboard values show the sign reversals in both roll and pitch at the start of the anomaly. The pitch angle begins to diverge, indicated by both the onboard value (becoming more negative) and the ground value (more positive). In contrast, the roll angles show little change during the anomaly. At the end of the anomaly, the onboard values are restored to their correct sign; the roll angle shows a small transient, and the pitch angle reverses direction and recovers over the next few minutes.

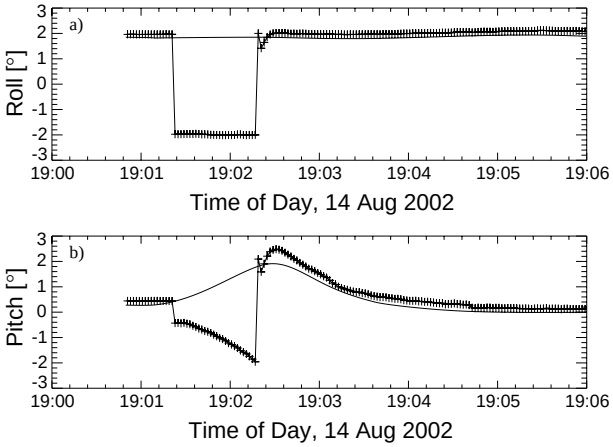

Fig. 60. Onboard (+) and ground-computed (line) attitudes during the 65,536 s anomaly on 14 August 2002: **a)** onboard and ground-roll angles, and **b)** onboard and ground-pitch angles.

The true magnitude of the pitch excursion can be seen in the onboard value shortly after the anomaly; the ground-computed value tends to damp the motion, and underestimates the peak by a few tenths of a degree (equivalent to a few pixels of navigation error). The ACS commands the momentum wheel to correct the apparent error resulting from the sign reversal, producing a rapid response. It is suspected that because the apparent (onboard) pitch appears to move in the opposite direction from the true (ground) value, divergence results. By contrast, the roll angle appears not to change during the anomaly, because the control authority of the magnetic coils is small compared to that of the momentum wheel, and the control disturbance is minimal.

One other aspect of this anomaly that merits discussion is the timing of the occurrences relative to the GAC scenes. The GAC data collection is scheduled each orbit, and the anomalies are synchronous with the week, so occurrences in the GAC data are purely coincidental. On average, less than half of the anomalies will be observed because the GAC periods are less than half of the orbit. (The GAC periods were 40 min at the start of the mission, increased to 43 min 40 s in July 2000). There are two coincidences of the OV-2 orbital period, however, which produce a pattern of anomaly occurrences:

1. The orbit period has been nearly synchronous with the week throughout the mission. When data collection started in September 1997, the period was 5,940 s, and 102 orbits exceeded 1 week by 18 min. As the orbit decayed, the period gradually decreased, and by mid-December 2001 was 5,929.4 s, with exactly 102 orbits per week. Since then, the orbit decay has continued, with a corresponding decrease in the period.
2. The interval between anomalies, i.e., 65,536 s, is slightly longer than 11 orbits; at the start of the mission, the difference was slightly more than 3 min, and as of December 2001, about 5 min.

The effect of these is twofold. First, for most of the mission, weekly events have appeared in nearly the same GAC scenes in consecutive weeks, and their timing relative to the start of the scene has shifted slowly from week to week. Second, the 65,536 s anomalies have tended to appear in multiple GAC scenes in most weeks, and their timing relative to the start of the scene has moved a few minutes for successive occurrences within the same week. The resultant pattern of occurrences is shown in Table 17, which lists every observation of the anomaly in 2002 (as determined from the Mission Operations Time Tag Anomalies Web page). These are organized by week and the 65,536 s multiple within the week; note that the weekly rollover anomaly is also shown as the zeroth multiple. Each entry in the table indicates the duration of the anomaly in GAC telemetry frames (10 s intervals). Occurrences that coincided with the DSS coverage gap are indicated by asterisks.

The anomaly duration also shows a pattern, with a given weekly event slowly increasing from week to week. In some cases, the duration reaches a peak (6 frames) and then disappears entirely for a few weeks (e.g., the fifth multiple, starting the week of 3 March) or decreases to 1 frame (e.g., the seventh multiple, from 17–24 November). The missing entries within the patterns are believed to be occurrences that were too brief to be observed at the 10 s sampling rate. The cause for the variations in duration is believed to be slow drift in the timing of the task that resynchronizes the ACE clock every minute. It has been shown that, in a given week, all of the anomaly corrections occurred at approximately the same point in the minute (given the 10 s sampling rate in the GAC data). Over the year, the timing of the corrections drifted slowly (about 2 s per week). This is also supported by the fact that shifts in the durations occurred after the SCM reset on 26 September, which had the effect of putting the corrections exactly on the minute. The timing then returned to a slow change in duration for subsequent weeks.

4.5.6 SCM Reset Time Anomaly

The OV-2 spacecraft has experienced four resets of the SCM. The first of these occurred in conjunction with the first spacecraft safe haven event, from 13–19 October 1997. The others occurred on 30 June 1998, 26 September 2002, and 21 February 2003. This particular type of onboard computer reset does not put the spacecraft into safe haven mode, or cause loss of imaging data. It does, however, cause errors of several seconds in both the spacecraft minor frame time tag and the ACS time tags, and resets ACS tables (including sensor biases) to default values. The combination of these results in both unusual spacecraft attitude motion, in all three axes, and large navigation processing errors, of approximately 50–60 km.

The sequence of events was different for each event. The 1997 reset was commanded on 17 October, as part of the spacecraft recovery operations during the week starting 13 October. Following the resumption of science data collection on 19 October, the time tags were found to be in error by an unknown amount during two consecutive orbits, but it is unclear whether this was associated with the SCM reset or some other aspect of the safe haven events.

In 1998, the reset condition was observed on 30 July, at the start of a GAC scene (17:09). The time tags were resynchronized to the GPS time at about 16:30 on 1 August; this reduced the spacecraft motion. At about 03:31 on 2 August, the ACS tables were reloaded, which restored normal spacecraft attitude control. The data prior to the time resynchronization (15 orbits) were found to have time tag errors, and were rejected from further processing because of the large navigation errors which resulted; the magnitude of the errors was not determined.

In 2002, the reset occurred on 26 September at 14:22:40, during GAC data collection. The telemetry stopped updating for the remainder of the GAC scene. Telemetry was restored at the start of the next GAC scene, but the spacecraft attitude was affected. The ACS tables were reloaded during two successive nighttime passes at 04:21 and 05:59 on 27 September, which eliminated most of the anomalous attitude motion. The spacecraft time tags were resynchronized almost a day later, during a pass at 03:26 on 28 September, which restored normal spacecraft attitude control. Subsequent analysis showed that the spacecraft minor frame and ACS time tags appeared to be in error by exactly −8 s from the start of the anomaly to the resynchronization, although the GPS time tags were not affected. A modification was made to the navigation

Table 17. A listing of most of the 65,536 s time tag anomalies observed in 2002. A *boxed* entry (e.g., [1]) indicates an anomaly during a DSS coverage gap. Multiples of 65,536 s are given above the days of the week.

Week Starting		0	1	2	3	4	5	6	7	8	9
Date	SDY	Sun.	Sun.	Mon.	Tue.	Wed.	Wed.	Thurs.	Fri.	Sat.	Sat.
06 January	6								6	5	4
13 January	13									5	4
20 January	20								1	5	4
27 January	27							3	2	6	4
03 February	34							4	2		5
10 February	41						5	4	3	1	[5]
17 February	48						6	4	2	1	[5]
24 February	55					1	6	4	3	[1]	6
03 March	62					2		5	3	[1]	6
10 March	69				3	1		5	[3]	1	5
17 March	76			5	3	2		[5]	3	1	6
24 March	83			5	4	2		[5]	3	2	
31 March	90		1	5	4	[3]		5	4	2	
07 April	97	4	2	6	[4]	2	1	6	4		
14 April	104	3	1	6	[4]	3	1	3			
21 April	111	3	1	[6]	4	3	1				
28 April	118	4	[1]	6	5	3					
05 May	125	[5]	3	6							
12 May	132	5	2	6	1						
19 May	139	5	3	5							
26 May	146	5	2								
02 June	153	5									
09 June	160										1
16 June	167										3
23 June	174									5	4
30 June	181									5	4
07 July	188							3	1	6	[4]
14 July	195							3	2	[6]	5
21 July	202					1	6	3	2		5
28 July	209				2	1	6	3	[2]		5
04 August	216				3	1	6	[4]	3		5
11 August	223			4	3	2	[5]	4	3	1	
18 August	230		6	5	2	[1]	6	4	2	1	
25 August	237	2	6	4	[3]	1		4	3		
01 September	244	2		5	[4]	2		4			
08 September	251	3	1	[5]	4	2	1				
15 September	258	3		5	4	2					
22 September	265	[3]		6	4						
29 September	272	6	4	2							
06 October	279	6	4								
13 October	286	5	4								
20 October	293										
27 October	300										3
03 November	307								2	5	4
10 November	314							1	7	5	4
17 November	321						3	2	6	5	[3]
24 November	328					5	4	3	1	5	3
01 December	335				1	5	4	[3]	1	5	4
08 December	342		4	3	1	5	4	2	1	5	

processing software to adjust all of the affected time tags. This largely restored the navigation accuracy, although the period prior to the reloading of the ACS tables had some degradation caused by the unusual attitude motion. The data immediately following the reset was not recoverable because of the loss of telemetry.

In 2003, the reset occurred on 21 February at 14:10:21, during GAC data collection. The telemetry stopped updating for the remainder of the GAC scene. Telemetry was restored at the start of the next GAC scene, but the spacecraft attitude was affected. The ACS tables were reloaded and the spacecraft time tags were resynchronized during a back-orbit pass at 04:07 on 22 February, which restored normal spacecraft attitude control. The spacecraft minor frame and ACS time tags were found to be in error by −8 s as before, and the navigation processing software was modified to adjust the affected time tags for this period as well. This largely restored the navigation accuracy, with some degradation caused by the unusual attitude motion. The data immediately following the reset were not recoverable because of the loss of telemetry.

4.5.7 Time Synchronization Anomaly

Shortly after the October 1997 safe haven event, the spacecraft experienced a problem with the synchronization of the spacecraft time. Starting on 22 October at about 01:00, until 24 October at about 04:00, the spacecraft minor frame time tags had random shifts of up to several seconds. The time shift was constant within each GAC scene, but varied randomly from orbit to orbit. In these scenes, the minor frame time tag was observed to be earlier than the attitude sensor time tags (an impossibility), by as much as 14 s. The sensor and GPS time tags appeared to be correct. A few scenes during this period were unaffected, but most had to be excluded from further processing because of the large navigation errors that resulted. The problem was corrected by OSC during a midnight pass on 24 October.

4.6 Miscellaneous Anomalies

There have been anomalies associated with changes of the encryption keys, and the momentum wheel, which have affected attitude control. These are described in the following sections.

4.6.1 Encryption Key Change Anomaly

The SeaWiFS science data are encrypted for direct broadcast, to protect the commercial data rights of OrbImage. The encryption keys are provided to owners of direct broadcast licenses (as well as to the SeaWiFS Project) to allow the data to be decrypted after receipt. The keys are entered by the station owners into a ground preprocessor, which decrypts the data prior to level-0-to-1A conversion. The keys are updated every two weeks, at 00:00 on alternate Wednesdays, an operation which requires powering off the instrument. If the nominal time falls within a scheduled GAC scene, the key change is delayed to the start of the next back-orbit period.

In 2000, large nutation was observed at the start of the first GAC scene of 8 March (a Wednesday). This was connected with a nearly identical event on 12 January (also a Wednesday). In both cases, the magnitude was roughly 1° peak-to-peak, and no other orbits were affected.

The timing of the event suggested a connection with the key changes. Further investigation showed that this had also occurred on a number of Wednesdays in 1998 and 1999. The suspected mechanism was the powering off and on of the SeaWiFS instrument, which includes spinning the rotating telescope down and then up again. The instrument has a momentum compensator that counteracts the telescope angular momentum during operation. The angular acceleration of the telescope after it is powered on apparently imparts a torque about the telescope rotation axis, which is perpendicular to the spacecraft angular momentum.

The question was why the nutation occurred for a small number of key changes (e.g., five events out of 26 key changes in 1999). The answer appeared to be the timing of the change relative to the start of the GAC scene, which depends on the start time of the scene. The two events in 2000 occurred in GAC scenes starting at 00:17 on 12 January and 00:11 on 8 March. The timings of the other affected scenes were also within 30 min of midnight. There was a strong correlation between the magnitude of the nutation and the start times of the scenes, with the largest nutation occurring for scenes that started soonest after midnight. The other key changes had either occurred for cases where the first GAC scene of the day started more than 30 min after midnight, or had been delayed to after a GAC scene, which ensured that they occurred more than 50 min before the start of the next scene. This allowed the nutation to be damped out by the control system prior to the start of the scene.

This information was presented to OrbImage, and a request was made to perform all future key changes within the first 20 min of a back-orbit period, to allow the nutation to be damped prior to the start of the GAC scene. The request was accepted and the procedure was implemented starting with the first key change in April 2000. Since that time, little or no nutation has been observed at key change times.

4.6.2 Momentum Wheel Anomaly

On 4 September 2001, transient pitch motion was observed in two successive GAC scenes. At about 01:39, a pitch excursion of about −1° occurred over a few minutes; at about 03:14, motion of about 0.41° was observed. The timing of the motion did not correlate with any known cause.

An investigation showed changes in the momentum wheel speed at these times; however, the timing and direction of the speed changes suggested that they were the cause of the pitch excursions, rather than being a response to them. The momentum wheel current and voltage showed stepwise changes at the times of the two speed changes. The current, which normally averages about 0.15 A, spiked to over 0.7 A, and then stabilized at about 0.4 A. It decreased slightly over the remainder of the scene. At the start of the next scene, the current was about 0.25 A, and then dropped to the normal range. The voltage showed corresponding, but less drastic, changes jumping from an average of about 3 V to about 4 V in the first scene, decreasing to about 3.5 V at the start of the second scene, and then dropping back to 3 V in the second scene.

The anomaly can be seen in plots of the pitch angle, momentum wheel speed, and current for the two GAC scenes (Fig. 61). The pitch angle shows the excursions near the start of each scene, superimposed on the normal orbital variation. The momentum wheel speed shows the spikes at the start of each pitch excursion. The wheel current shows the initial spike and elevated current in the first scene, and the sudden recovery in the second scene.

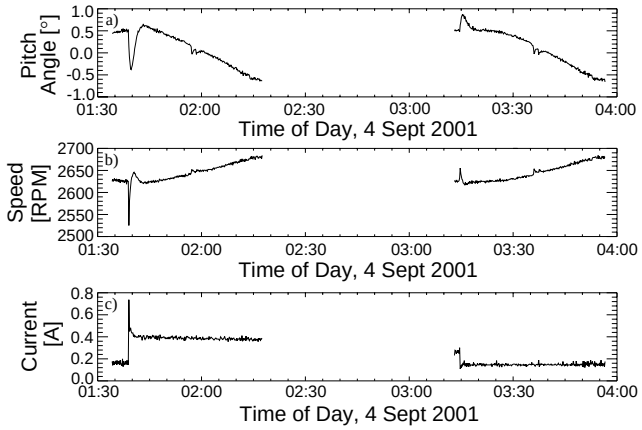

Fig. 61. Pitch and momentum wheel behavior during two consecutive GAC scenes spanning momentum wheel anomaly of 4 September 2001: **a)** onboard pitch angle, **b)** momentum wheel speed, and **c)** momentum wheel current.

This incident was reported to OrbImage, but no specific cause was found. The characteristics and sequence of events suggest that a particle may have entered one of the momentum wheel bearings. This would have caused increased friction and the sudden drop in wheel speed, in turn causing the pitch excursion. The wheel control electronics responded by increasing the voltage to return the speed to normal, which also increased the current. The slow decrease in current after the initial event indicates a gradual decrease in bearing friction. At the start of the second scene, the particle would have worked itself out of the bearing, which reduced the friction, increased the wheel speed, and produced the second pitch excursion. As before, the wheel control electronics reduced the voltage to return the speed to normal.

It appears to have been coincidental that the start and end of this anomaly was observed during GAC scenes. There were no apparent long-term consequences, and no other momentum wheel anomalies like this have been observed during the mission.

5. SUMMARY

The pointing performance for the OV-2 spacecraft carrying the SeaWiFS instrument has been described. Issues have been discussed regarding the data collection, the onboard attitude calculation, the attitude stability and control, and a wide variety of anomalies. This section summarizes key results and lessons learned that could be useful in planning for future missions.

5.1 Data Rates

The OV-2 data collection appropriately emphasized the sunlit side of the orbit, however, the severe limitation on data in the back-orbit spans proved to be a significant frustration for analyzing various anomalies. For future missions with a similar design, it is recommended that careful thought be given to including a higher rate data subset collected in the back orbit for regular monitoring purposes.

Engineering data samples at 3 min intervals are not adequate for analyzing anomalous control events. All of the engineering data need not be sampled at the high rate, but key items could be included at the high rate and still represent just a tiny fraction of the science data stream. These items could include the onboard attitude and a few mode flags.

Additional flexibility in the telemetry could also allow for options to monitor particular items decided upon after launch, as is done in other onboard systems (e.g., the Tropical Rainfall Measuring Mission). Some flexibility in the telemetry was available that made possible the occasional high rate data collections in the back orbit. An effort to re-allocate the data for more regular monitoring, however, was not practical. Another possibility to support anomaly analysis would be a *most recent data* buffer, which is saved to show engineering data leading up to any safe haven event.

5.2 Onboard Attitude

The onboard attitude performed adequately after various adjustments were made. The onboard attitude did not need to meet stringent accuracy requirements, because the data were reprocessed on the ground for SeaWiFS Project data products. Nonetheless, it proved desirable to avoid errors that could trigger nutation and other relatively rapid motions of the spacecraft.

The primary issue for the onboard attitude was the single-frame solution sensitivity to periods of poor geometry with the sun and Earth vectors around the subsolar points. The potential errors upon deployment were underestimated prior to launch, and sensor calibration plans were not in place. A simple method for rough bias estimation based on the subsolar region error signature, described previously, was developed and used soon after launch to facilitate adjustments during the first month of science data collection. This removed the largest disturbances; however, disturbances generated by errors in the onboard attitude estimates at the subsolar point were still a concern. Further adjustments proved useful seasonally, particularly with the start of single-string operation. Widening of the yaw hold region eventually proved to be a key adjustment for minimizing the subsolar error effects.

A secondary issue was the poorly calibrated TAMs and the use of an old (reference year 1980) magnetic field model onboard. This was not as critical because it only affected attitude in the Earth shadow period, and in the brief sun coverage gap at the subsolar point within the yaw hold region. It contributed, however, to fairly large errors in the back orbit, which sometimes triggered nutation at the start of the science data collection. Experience demonstrates that the most recent field model is important because of the secular drift in the Earth's dipole. Other missions have updated the models on board.

5.3 Control

The main goals of ACS adjustments were to limit nutation and pitch motion disturbances, as these were harder to track in the ground attitude estimation. The initial on-orbit errors gave nutation amplitudes of 1–2°, but after initial adjustments, the amplitudes were generally less than 0.5° except during various anomalies. The coil commands for roll–yaw steering were very sensitive to subsolar region estimation errors, and this sensitivity was reduced by control feedback reductions. Improvements in the onboard attitude biases and expansion of the yaw hold region also reduced the generation of disturbances from the coil commands. The coil commands could still be very noisy, especially in the back orbit, but that was not a problem.

The spacecraft spin axis average position stayed about 1–2° from orbit normal, after the control adjustments were completed, with a twice-per-orbit motion of the spin axis. The average position and motion of the spin axis shifted with a BCR switch in December 1998. It was inferred that a residual magnetic dipole bias was probably the main cause of this motion. A proposed method to reduce the spin axis drift and position it closer to orbit normal could not be implemented because of a coding error in the coil calibrations that could not be patched. Pitch motions have been generally small and slow enough to be tracked in ground software, except during the tilt change and during various anomalies.

5.4 Anomalies

A variety of sensor data and onboard computational anomalies were presented. These generally have only minor or infrequent consequences. There are, however, a number of anomalies that occasionally cause notable attitude disturbances:

a) Cold atmosphere anomaly—resulted in nutation because of errors detecting the Earth horizon;

b) GPS anomalies—caused errors in the onboard orbit and time tag, and occasional data loss;

c) Weekly rollover time tag anomaly—caused two safe haven modes before being fixed;

d) 65,536 s anomaly—resulted in occasional pitch disturbances of greater than 1°;

e) SCM reset time anomaly—time tag errors and other effects were observed following a processor reset;

f) Early mission time synchronization anomaly—caused loss of data because of time tag errors; and

g) Encryption key change anomaly—large nutation followed key changes before the key change schedule update.

Other anomalies were not significant sources of attitude disturbance or data loss, but were noted in this document for completeness in documenting the ACS performance.

5.5 Programmatic Support Observations

The ACS support for OV-2 was different than most NASA missions because of its unique programmatic support. The SeaWiFS data were acquired at NASA for the SeaWiFS Project Office under a unique data buy arrangement, with the spacecraft developed as a commercial investment by OSC and handed over to OrbImage for operation. The mission was beset by cost overruns and launch delays, and then after launch, it was not a strong commercial success. After orbit raising and checkout, OrbImage could call on limited support from OSC engineers, but the support available was less than for typical NASA missions, which have civil service and contractor support.

SeaWiFS Project staff worked in cooperation with OSC and OrbImage engineers on critical issues as needed. The Project had an interest in adjustments that optimized control to improve the data products, and OrbImage had an interest in satisfying NASA and their commercial customers, although available resources on both sides limited the level of effort. A good working relationship provided mutual benefits. OrbImage provided the SeaWiFS staff with the ACS simulation code and helped track down information about the spacecraft, as needed. The SeaWiFS staff made recommendations on ACS adjustments and provided analysis of anomalies.

OSC and OrbImage did not have a plan in place at launch for sensor calibration efforts. This was an oversight

in comparison with most NASA missions that GSFC supports, but requirements were not defined for this within the context of the commercial development effort. The SeaWiFS staff were able to provide assistance with this when problems were realized after launch. The effort was more ad hoc than formal, however, so initial on-orbit tuning was not as thorough as for most NASA missions. On-orbit adjustments benefited further from the analysis done for improved navigation in the SeaWiFS Project data products.

The OrbImage support priorities and budget constraints also led to limited support for changes to the onboard software. For example, the spacecraft simulator, FlatSat, could not be maintained past the first two years of the mission. This meant that some desirable changes could not be implemented (e.g., a coil bias correction). Numerous onboard software adjustments and table changes were not loaded into permanent memory onboard, as is usually done for other missions; this meant many changes had to be reloaded multiple times, after SCM resets. Some known software problems (e.g., the problems in the background software task for the oblateness correction) were set aside and not fixed. Fortunately, no problems that jeopardized mission success were left unaddressed, but there remained added risk with certain types of tools and types of support not readily available.

The automated Web-based display of engineering data at the SeaWiFS Project Office proved to be very useful to SeaWiFS staff. The Web displays could be accessed from any browser, so support could be provided from remote locations, and a wide variety of telemetry items could be checked if questions arose. It helped in the detection of anomalies and in the diagnoses and review of the spacecraft performance. Easy links were provided for scanning through the data and comparing performance from orbit to orbit and day to day. Selected long term trends were automatically updated. Additional plots of derived parameters and data summary files (e.g., orbit error plots and time tag error summaries) were added as needed for further analysis. Further use of this approach is recommended in future missions.

Overall, it should be summarized that despite the spacecraft support limitations noted above, the scientific mission from SeaWiFS has been a great success, and the ACS was tuned adequately to provide superior data products with very few and limited data outages.

Glossary

ACE Attitude Control Electronics
ACE-A Primary ACE box.
ACE-B Backup, or redundant, ACE box.
ACS Attitude Control System
AOS Acquisition of Signal

BCR Battery Charge Regulator

DOP Dilution of Precision
DSS Digital Sun Sensor
DSS-A Digital Sun Sensor-A, the front-mounted sensor.
DSS-B Digital Sun Sensor-B, the back-mounted sensor.
DSS-C Digital Sun Sensor-C, the top-mounted sensor.

ECR Earth-Centered Rotating
EEPROM Electronically Erasable Programmable Read-Only Memory

FlatSat Not an acronym, but a laboratory-bench spacecraft simulator—a shortened way of saying "Flat Satellite."
FOV Field of View

GAC Global Area Coverage
GIM Ground Interface Model
GPS Global Positioning System
GSFC Goddard Space Flight Center
GST Greenwich Sidereal Time

HRPT High Resolution Picture Transmission
HS Horizon Scanner
HS-A Horizon Scanner-A
HS-B Horizon Scanner-B

IGRF International Geomagnetic Reference Field

LAC Local Area Coverage
LOS Loss of Signal

NASA National Aeronautics and Space Administration

OSC Orbital Sciences Corporation
OV-2 OrbView-2, the satellite platform on which the SeaWiFS instrument is flown.

PC Personal Computer
PPS Pulse Per Second
PSM Payload Support Module

RAAN Right Ascension of Ascending Node
RPO Revolutions Per Orbit
RTC Real Time Clock

SAA South Atlantic Anomaly
SCM Spacecraft Control Module
SDY Sequential Day of the Year
SeaWiFS Sea-viewing Wide Field-of-view Sensor
SEU Single-Event Upset
SOH State of Health

TAM Three-Axis Magnetometer
TAM-A Three-Axis Magnetometer, A
TAM-B Three-Axis Magnetometer, B

UTC Coordinated Universal Time

Symbols

$\vec{B}$ Geomagnetic field vector.

$\vec{C}$ Coil commands as proportional to the vector cross product of $\vec{B}$.
C_x Coil command along $\vec{X}$.
C_y Coil command along $\vec{Y}$.
C_z Coil command along $\vec{Z}$.

D DSS validity check parameter.
D_x Angle of the sun line from the DSS y–z plane.
D_y Angle of the sun line from the DSS x–z plane.

$\vec{N}$ Desired torque vector.
N_w Momentum wheel torque.
N_x Roll torque.
N_y Pitch torque.
N_z Yaw torque.

S_x SeaWiFS base reference frame x coordinate.
S_y SeaWiFS base reference frame y coordinate.
S_z SeaWiFS base reference frame z coordinate.

$\vec{T}$ TAM-measured field (averaged over TAM-A and TAM-B when both are in use).
T_x TAM measurement along $\vec{X}$.
T_y TAM measurement along $\vec{Y}$.
T_z TAM measurement along $\vec{Z}$.

$\vec{V}$ Primary reference vector.

$\vec{W}$ Secondary reference vector.

x The first axis of an orthogonal triad.
$\vec{X}$ The x, or roll, axis in ACS coordinates.

y The second axis of an orthogonal triad.
$\vec{Y}$ The y, or pitch, axis in ACS coordinates.

z The third axis of an orthogonal triad.
$\vec{Z}$ The z, or yaw, axis in ACS coordinates.

β Sun elevation angle from the orbit plane.

δ_W Angular error in $\vec{W}$.
δ_x Roll error.
δ_y Pitch error.
δ_z Yaw error.

Γ Separation angle between two vectors.

η Error in rotation angle.

ω_O Orbit rotation rate.
ω_w Momentum wheel speed.
ω_x Roll rate.
ω_y Pitch rate.
ω_z Yaw rate.

References

Beard, R.M., and M. Plett, 1978: "Dynamic Motion of a Symmetric Dual-Spin Spacecraft." In: Wertz, J.R., *Spacecraft Attitude Determination and Control,* D. Reidel Publishing Company, Dordrecht, Holland, 536–539.

Chobotov, V.A., 1991: *Spacecraft Attitude Dynamics and Control,* Krieger Publishing Co., Malabar, Florida, 33–40.

Hashmall, J.A., and J. Sedlak, 1997: The use of magnetometers for accurate attitude determination. *12th International Symp. Space Flight Dynamics,* Darmstadt, Germany, 179–184.

Headrick, D., 1978: "Momentum Bias Control Systems." In: Wertz, J.R., *Spacecraft Attitude Determination and Control,* D. Reidel Publishing Company, Dordrecht, Holland, 601–603.

Patt, F.S., 2002: Navigation Algorithms for the SeaWiFS Mission. *NASA Tech. Memo. 2002–206892, Vol. 16,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 17 pp.

Patt, F.S., and S. Bilanow, 1999: Attitude sensor alignment for OrbView-2 using island targets. Proc. 1999 Flight Mechanics Symp., *NASA Conf. Pub. 1999–209986,* NASA Goddard Space Flight Center, Greenbelt, Maryland, 479–493.

Patt, F.S., and S. Bilanow, 2001: Horizon scanner triggering height Analysis for OrbView-2. Proc. 2001 Flight Mechanics Symp., *NASA Conf. Pub. 2001–209235,* NASA Goddard Space Flight Center, Greenbelt, Maryland, 559–573.

Patt, F.S., B.A. Franz., W.D. Robinson, and J. Gales, 2003: "Level-1a and Level-3 Processing Changes." In: Patt, F.S., R.A. Barnes, R.E. Eplee, Jr., B.A. Franz, W.D. Robinson, G.C. Feldman, S.W. Bailey, J. Gales, P.J. Werdell, M. Wang, R. Frouin, R.P. Stumpf, R.A. Arnone, R.W. Gould, Jr., P.M. Martinolich, V. Ransibrahmanakul, J.E. O'Reilly, and J.A. Yoder, Algorithm Updates for the Fourth SeaWiFS Data Reprocessing, *NASA Tech. Memo. 2003–206892, Vol. 22,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 41–45.

Schield, M.A., J.W. Freeman, and A.J. Dessler, 1969: A source for field-aligned currents at auroral latitudes. *J. Geophys. Res.,* **74,** 247–256.

Shuster, M.D., and S.D. Oh, 1981: Three-axis attitude determination from vector observations, *J. Guidance and Control,* **4,** 70–77.

The SeaWiFS Postlaunch Technical Report Series

Vol. 1

Johnson, B.C., J.B. Fowler, and C.L. Cromer, 1998: The SeaWiFS Transfer Radiometer (SXR). *NASA Tech. Memo. 1998–206892, Vol. 1,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 58 pp.

Vol. 2

Aiken, J., D.G. Cummings, S.W. Gibb, N.W. Rees, R. Wooddd-Walker, E.M.S. Woodward, J. Woolfenden, S.B. Hooker, J-F. Berthon, C.D. Dempsey, D.J. Suggett, P. Wood, C. Donlon, N. González-Benítez, I. Huskin, M. Quevedo, R. Barciela-Fernandez, C. de Vargas, and C. McKee, 1998: AMT-5 Cruise Report. *NASA Tech. Memo. 1998–206892, Vol. 2,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 113 pp.

Vol. 3

Hooker, S.B., G. Zibordi, G. Lazin, and S. McLean, 1999: The SeaBOARR-98 Field Campaign. *NASA Tech. Memo. 1999–206892, Vol. 3,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 40 pp.

Vol. 4

Johnson, B.C., E.A. Early, R.E. Eplee, Jr., R.A. Barnes, and R.T. Caffrey, 1999: The 1997 Prelaunch Radiometric Calibration of SeaWiFS. *NASA Tech. Memo. 1999–206892, Vol. 4,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 51 pp.

Vol. 5

Barnes, R.A., R.E. Eplee, Jr., S.F. Biggar, K.J. Thome, E.F. Zalewski, P.N. Slater, and A.W. Holmes 1999: The SeaWiFS Solar Radiation-Based Calibration and the Transfer-to-Orbit Experiment. *NASA Tech. Memo. 1999–206892, Vol. 5,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, 28 pp.

Vol. 6

Firestone, E.R., and S.B. Hooker, 2000: SeaWiFS Postlaunch Technical Report Series Cumulative Index: Volumes 1–5. *NASA Tech. Memo. 2000–206892, Vol. 6,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 14 pp.

Vol. 7

Johnson, B.C., H.W. Yoon, S.S. Bruce, P-S. Shaw, A. Thompson, S.B. Hooker, R.E. Eplee, Jr., R.A. Barnes, S. Maritorena, and J.L. Mueller, 1999: The Fifth SeaWiFS Intercalibration Round-Robin Experiment (SIRREX-5), July 1996. *NASA Tech. Memo. 1999–206892, Vol. 7,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, 75 pp.

Vol. 8

Hooker, S.B., and G. Lazin, 2000: The SeaBOARR-99 Field Campaign. *NASA Tech. Memo. 2000–206892, Vol. 8,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, 46 pp.

Vol. 9

McClain, C.R., E.J. Ainsworth, R.A. Barnes, R.E. Eplee, Jr., F.S. Patt, W.D. Robinson, M. Wang, and S.W. Bailey, 2000: SeaWiFS Postlaunch Calibration and Validation Analyses, Part 1. *NASA Tech. Memo. 2000–206892, Vol. 9,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, 82 pp.

Vol. 10

McClain, C.R., R.A. Barnes, R.E. Eplee, Jr., B.A. Franz, N.C. Hsu, F.S. Patt, C.M. Pietras, W.D. Robinson, B.D. Schieber, G.M. Schmidt, M. Wang, S.W. Bailey, and P.J. Werdell, 2000: SeaWiFS Postlaunch Calibration and Validation Analyses, Part 2. *NASA Tech. Memo. 2000–206892, Vol. 10,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, 57 pp.

Vol. 11

O'Reilly, J.E., S. Maritorena, M.C. O'Brien, D.A. Siegel, D. Toole, D. Menzies, R.C. Smith, J.L. Mueller, B.G. Mitchell, M. Kahru, F.P. Chavez, P. Strutton, G.F. Cota, S.B. Hooker, C.R. McClain, K.L. Carder, F. Müller-Karger, L. Harding, A. Magnuson, D. Phinney, G.F. Moore, J. Aiken, K.R. Arrigo, R. Letelier, and M. Culver 2000: SeaWiFS Postlaunch Calibration and Validation Analyses, Part 3. *NASA Tech. Memo. 2000–206892, Vol. 11,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, 49 pp.

Vol. 12

Firestone, E.R., and S.B. Hooker, 2000: SeaWiFS Postlaunch Technical Report Series Cumulative Index: Volumes 1–11. *NASA Tech. Memo. 2000–206892, Vol. 12,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 24 pp.

Vol. 13

Hooker, S.B., G. Zibordi, J-F. Berthon, S.W. Bailey, and C.M. Pietras, 2000: The SeaWiFS Photometer Revision for Incident Surface Measurement (SeaPRISM) Field Commissioning. *NASA Tech. Memo. 2000–206892, Vol. 13,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 24 pp.

Vol. 14

Hooker, S.B., H. Claustre, J. Ras, L. Van Heukelem, J-F. Berthon, C. Targa, D. van der Linde, R. Barlow, and H. Sessions, 2000: The First SeaWiFS HPLC Analysis Round-Robin Experiment (SeaHARRE-1). *NASA Tech. Memo. 2000–206892, Vol. 14,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 42 pp.

Vol. 15

Hooker, S.B., G. Zibordi, J-F. Berthon, D. D'Alimonte, S. Maritorena, S. McLean, and J. Sildam, 2001: Results of the Second SeaWiFS Data Analysis Round Robin, March 2000 (DARR-00). *NASA Tech. Memo. 2001–206892, Vol. 15,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 71 pp.

Vol. 16

Patt, F.S., 2002: Navigation Algorithms for the SeaWiFS Mission. *NASA Tech. Memo. 2002–206892, Vol. 16,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 17 pp.

Vol. 17

Hooker, S.B., S. McLean, J. Sherman, M. Small, G. Lazin, G. Zibordi, and J.W. Brown, 2002: The Seventh SeaWiFS Intercalibration Round-Robin Experiment (SIRREX-7), March 1999. *NASA Tech. Memo. 2002–206892, Vol. 17,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 69 pp.

Vol. 18

Firestone, E.R., and S.B. Hooker, 2002: SeaWiFS Postlaunch Technical Report Series Cumulative Index: Volumes 1–17. *NASA Tech. Memo. 2002–206892, Vol. 18,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 28 pp.

Vol. 19

Zibordi, G., J-F. Berthon, J.P. Doyle, S. Grossi, D. van der Linde, C. Targa, and L. Alberotanza 2002: Coastal Atmosphere and Sea Time Series (CoASTS), Part 1: A Tower-Based Long-Term Measurement Program. *NASA Tech. Memo. 2002–206892, Vol. 19,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 29 pp.

Vol. 20

Berthon, J-F., G. Zibordi, J.P. Doyle, S. Grossi, D. van der Linde, and C. Targa, 2002: Coastal Atmosphere and Sea Time Series (CoASTS), Part 2: Data Analysis. *NASA Tech. Memo. 2002–206892, Vol. 20,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 25 pp.

Vol. 21

Zibordi, G., D. D'Alimonte, D. van der Linde, J-F. Berthon, S.B. Hooker, J.L. Mueller, G. Lazin, and S. McLean, 2002: The Eighth SeaWiFS Intercalibration Round-Robin Experiment (SIRREX-8), September–December 2001. *NASA Tech. Memo. 2002–206892, Vol. 21,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 39 pp.

Vol. 22

Patt, F.S., R.A. Barnes, R.E. Eplee, Jr., B.A. Franz, W.D. Robinson, G.C. Feldman, S.W. Bailey J. Gales, P.J. Werdell, M. Wang, R. Frouin, R.P. Stumpf, R.A. Arnone, R.W. Gould, Jr., P.M. Martinolich, V. Ransibrahmanakul, J.E. O'Reilly, and J.A. Yoder, 2003: Algorithm Updates for the Fourth SeaWiFS Data Reprocessing, *NASA Tech. Memo. 2003–206892, Vol. 22,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 74 pp.

Vol. 23

Hooker, S.B., G. Zibordi, J-F. Berthon, D. D'Alimonte, D. van der Linde, and J.W. Brown, 2003: Tower-Perturbation Measurements in Above-Water Radiometry. *NASA Tech. Memo. 2003–206892, Vol. 23,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 35 pp.

Vol. 24

Firestone, E.R., and S.B. Hooker, 2003: SeaWiFS Postlaunch Technical Report Series Cumulative Index: Volumes 1–23. *NASA Tech. Memo. 2003–206892, Vol. 24,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 35.

Vol. 25

Doyle, J.P., S.B. Hooker, G. Zibordi, and D. van der Linde, 2003: Validation of an In-Water, Tower-Shading Correction Scheme. *NASA Tech. Memo. 2003–206892, Vol. 25,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 32 pp.

Vol. 26

Zibordi, G., D. D'Alimonte, D. van der Linde, S.B. Hooker, and J.W. Brown, 2003: New Laboratory Methods for Characterizing the Immersion Factors of Irradiance Sensors. *NASA Tech. Memo. 2003–206892, Vol. 26,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 34 pp.

Vol. 27

Barlow, R., H. Sessions, N. Silulwane, H. Engel, S.B. Hooker, J. Aiken, J. Fishwick, V. Vicente, A. Morel, M. Chami, J. Ras, S. Bernard, M. Pfaff, J.W. Brown, and A. Fawcett, 2003: BENCAL Cruise Report. *NASA Tech. Memo. 2003–206892, Vol. 27,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 64 pp.

Vol. 28

Bilanow, S., and F.S. Patt, 2004: Pointing Performance for the SeaWiFS Mission. *NASA Tech. Memo. 2004–206892, Vol. 28,* S.B. Hooker and E.R. Firestone, Eds., NASA Goddard Space Flight Center, Greenbelt, Maryland, 63 pp.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	March 2004	Technical Memorandum

4. TITLE AND SUBTITLE
SeaWiFS Postlaunch Technical Report Series
Volume 28: Pointing Performance for the SeaWiFS Mission

5. FUNDING NUMBERS

6. AUTHOR Stephen Bilanow and Frederick S. Patt

Series Editors: Stanford B. Hooker and Elaine R. Firestone

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Laboratory for Hydrospheric Processes
Goddard Space Flight Center
Greenbelt, Maryland 20771

8. PERFORMING ORGANIZATION REPORT NUMBER
2003-01910-0

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Washington, D.C. 20546–0001

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
TM—2004-206892, Vol. 28

11. SUPPLEMENTARY NOTES

E.R. Firestone, S. Bilanow, and F.S. Patt: Science Applications International Corporation, Beltsville, Maryland

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified–Unlimited
Subject Category 48
Report is available from the Center for AeroSpace Information (CASI),
7121 Standard Drive, Hanover, MD 21076–1320; (301)621-0390

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*

The onboard pointing performance of the OrbView-2 (\OVtwo) spacecraft for the first five years of the SeaWiFS mission is presented. Adjustments to the onboard attitude control system (ACS) since launch are described, and various issues and anomalies regarding the performance are discussed. Overall, this relatively low-cost spacecraft has performed quite effectively after various in-flight adjustments, however, a variety of sensor and computational anomalies have caused occasional minor pointing disturbances. Many of these disturbances have implications for the navigation processing performed for the science data by the Sea-viewing Wide Field-of-view Sensor (SeaWiFS) Project at the NASA Goddard Space Flight Center. Possible further adjustments to the OrbView-2 ACS have been investigated which could lead to improved pointing performance, and conclusions from these analyses are presented. Some of the various sensor and software anomalies are fairly well understood, but some others remain puzzling. Particularly vexing are various timing anomalies resulting from the way the clocks for three separate onboard processors are synchronized to each other and Global Positioning System (GPS) time. The pattern of occurrence of some of the sensor anomalies could merit further review and trending, which may be useful to monitor for any degradation in performance as themission continues.

14. SUBJECT TERMS
SeaWiFS, Oceanography, Mission Operations, Pointing Performance, Yaw, Hold, Pitch, Flight, Stability

15. NUMBER OF PAGES
63

16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified	Unclassified	Unlimited