Operations with the new *FUSE* observatory: three-axis control with one reaction wheel

David J. Sahnow[*a], Jeffrey W. Kruk[a], Thomas B. Ake[a,b], B-G Andersson[a], Alice Berman[a], William P. Blair[a], Robert Boyer[a,b], James Caplinger[a,b], Humberto Calvani[a], Thomas Civeit[a,c],W. Van Dyke Dixon[a], Martin N. England[a,b], Mary Elizabeth Kaiser[a], Mark Kochte[a,b], H. Warren Moos[a], Bryce A. Roberts[a,d]

[a] Department of Physics & Astronomy, The Johns Hopkins University, 3400 North Charles Street, Baltimore, MD 21218
[b] Computer Sciences Corporation, 7700 Hubble Drive, Lanham-Seabrook, MD 20706
[c] Centre National d'Etudes Spatiales, 2 place Maurice Quentin, 75039 Paris Cedex 1, France
[d] Space Sciences Lab, University of California at Berkeley, 7 Gauss Way, Berkeley, CA 94720

ABSTRACT

Since its launch in 1999, the Far Ultraviolet Spectroscopic Explorer (*FUSE*) has had a profound impact on many areas of astrophysics. Although the prime scientific instrument continues to perform well, numerous hardware failures on the attitude control system, particularly those of gyroscopes and reaction wheels, have made science operations a challenge. As each new obstacle has appeared, it has been overcome, although sometimes with changes in sky coverage capability or modifications to pointing performance. The CalFUSE data pipeline has also undergone major changes to correct for a variety of instrumental effects, and to prepare for the final archiving of the data. We describe the current state of the *FUSE* satellite and the challenges of operating it with only one reaction wheel and discuss the current performance of the mission and the quality of the science data.

Keywords: FUSE, satellite operations, attitude control

1. INTRODUCTION

The *Far Ultraviolet Spectroscopic Explorer* (*FUSE*), is a NASA astrophysics mission designed to collect spectra of astronomical objects at moderately high resolution ($\lambda/\Delta\lambda \approx 20{,}000$) in the 905 – 1187 Å wavelength region. *FUSE* was launched into a 100 minute, low-earth orbit with a 25° inclination on June 24, 1999 aboard a Delta II launch vehicle. After several months of in-orbit checkout, science observations officially began in December 1999. During the three year prime mission that ended in March 2003, observing time was divided between the Principal Investigator (PI) team and Guest Investigators (GIs) from institutions around the world. Since then, in the extended mission phase, all time has been allocated to GIs who propose through NASA.

The details of the design and early performance of *FUSE* has been described previously[1,2] and will only be reviewed briefly here. The optical design consists of four coaligned telescopes, each feeding a holographically-ruled, aberration-corrected grating on a Rowland-circle mount. In order to keep the target within the spectrograph entrance apertures and to obtain high quality spectra, the satellite must maintain stable pointing in all three axes. Jitter values of no more than ±1 arcsecond in pitch, ±10 arcsecond in yaw, and ±1° in roll are required over periods of five minutes to an hour (typical exposure times) in order to obtain high-quality science data. Faint objects, which can require many thousands of seconds on a single target, require that the instrument configuration be repeatable, since such long observations may be divided up into multiple visits that may be separated by days or months.

[*] sahnow@pha.jhu.edu; phone 410-516-8503; fax 410-516-5494; http://fuse.pha.jhu.edu

The *FUSE* Attitude Control System (ACS) contains two sets of three ring-laser gyroscopes (Inertial Reference Units, or IRUs) for attitude estimate propagation. Three-axis magnetometers and coarse sun sensors provide coarse attitude information to ±2°. The required fine pointing is achieved by using a signal from a Fine Error Sensor[3] (FES) in the instrument which images a region of the sky around the spectrograph apertures. Four Reaction Wheel Assemblies (RWAs) are used to control the attitude of the satellite and manage angular momentum. Three wheels are arranged along the primary axes (pitch, yaw, and roll) of the satellite, and the fourth, or skew, wheel is oriented equidistant from the others. The skew wheel is biased to minimize zero-speed crossings on the other wheels and can serve as a substitute in case of failure of one of the other RWAs. Three magnetic torquer bars (MTBs), mounted along the prime axes, are used to control the momentum load on the wheels by acting on the earth's magnetic field. A more detailed description of the design of the ACS, along with a description of its performance early in the mission and in two-wheel mode (described below) has been presented elsewhere.[4,5]

Beginning in late 2001, the failure of spacecraft components began to affect operations of the satellite. Now, in mid-2006, the satellite is operating with only one of the four reaction wheels and two of the six gyroscopes. Loss of these hardware components has required a significant redesign of the way onboard tasks are handled, but with software upgrades and improvements in the ground mission planning system, *FUSE* is again making observations with an efficiency approaching that of its earlier days, albeit over a smaller portion of the sky.

2. HARDWARE PROBLEMS AND WORKAROUNDS

In January 2000, several of the gyroscopes first showed signs of premature aging when a telemetry warning flag tripped, signifying that the laser intensity was dropping; however, they all continued to perform nominally. In August 2000, we received the first indication of a problem with a reaction wheel, when the pitch wheel showed signs of increased static friction, or stiction. It was eventually returned to service, and it was not until 2001 that *FUSE* experienced its first permanent component failure. Since then, the loss of gyroscopes and reaction wheels has had a major effect on our ability to operate and control the satellite. Table 1 lists the hardware problems seen in the IRUs and reaction wheels since launch. The effects of losing these components on the mission are described in this section.

Table 1 Status of the *FUSE* Gyroscopes (IRU-A and IRU-B) and Reaction Wheels

Axis	IRU-A	IRU-B	Reaction Wheel Assemblies
Yaw	1/6/00 Warning flag tripped **Operating normally**	12/10/02 Warning flag tripped 7/31/03 Failed	2/16/01 Stopped; restarted in 11 days 11/25/01 Failed
Pitch	1/18/00 Warning flag tripped **Operating normally**	8/31/01 Warning flag tripped 9/28/04 Noisy / Turned off	8/4/00 Stopped; restarted in 40 days 12/10/01 Failed
Roll	4/19/00 Warning flag tripped 5/30/01 Failed	10/6/01 Warning flag tripped 5/17/05 Failed	12/17/03 Stopped; restarted in 2 hrs 12/27/04 Failed
Skew			**Operating normally**

When the first gyro failed, onboard software allowed the satellite to continue operating by switching to IRU-B, and science operations were interrupted for less than 24 hours. Similarly, the loss of the first reaction wheel had no major effect on operations, since the remaining three wheels were sufficient to fully control the satellite. Further component losses have had a more serious effect.

At present, only two of the six gyroscopes are still functioning, and both have tripped their warning flag (indicating that the laser intensity has dropped below a pre-set threshold). The time between the flag tripping and failure of a gyro varies widely, but based on the behavior of the other gyros, it is believed that these two will fail at some point during the mission. The three reaction wheels aligned with the principal axes of the satellite have also failed. Attempts to restart them have been unsuccessful, and they are now considered unrecoverable.

Each of the failed wheels exhibited a significant stiction event roughly a year prior to failure. These events were attributed to touchdown of the rotor onto the kapton tape lining the wheel housing. It was thought at the time that small bubbles of gas trapped under the tape might have grown over time until they spanned the small gap between the rotor and housing,[6] but the subsequent hard failures did not appear to be consistent with this model. The working hypothesis at present is that the combination of thermal cycling of the spacecraft structure on orbital timescales in conjunction with flexing of the standoffs on which the three wheels were mounted resulted in gradually accumulating stresses in the wheel housings. A small distortion would be sufficient to close the narrow gap between rotor and housing. In this scenario, the initial stiction events resulted from contact between rotor and the kapton tape lining, which was quickly worn away. The increasing distortion would ultimately cause the rotor to touch the housing itself, resulting in a hard failure. The remaining skew wheel is mounted in a manner different from the other three and has never shown signs of stiction. We therefore believe that it is not susceptible to the same type of failure.

Even before the failure of the first gyro, plans were developed for spacecraft operations using less than the full complement of gyros. The loss of two reaction wheels in late 2001, well before "gyroless" pointing could be implemented, significantly complicated the procedure. The failure of two wheels in 15 days meant that only two wheels remained to control the three spacecraft axes. Inertial pointing could not be maintained with the wheels alone, and initially the satellite drifted about the uncontrolled axis, meaning that no scientific observations could be made. The solution devised was to control this axis with the magnetic torquer bars; within 10 days of the failure, *FUSE* could be held in a safe, inertial pointing using that technique. To gain enough control for fine pointing and resumption of science data collection, however, took until late January 2002. The details of two-wheel mode, including a discussion of the forces on the satellite and how they are balanced, have been given previously.[4,7] By March 2004 improved planning tools and the use of offsets in roll angle had made it possible to access all points on the celestial sphere at some time during a calendar year.

Loss of the third wheel in late 2004 required that the MTBs control pointing in two axes in addition to managing momentum. Loss of the second roll gyro in 2005 meant that roll information had to be provided by the FES, which provides poorer data on that axis than in pitch and yaw. Much of 2005 was spent developing new software and operational methods to deal with these challenges.

3. ONE-WHEEL OPERATIONS

3.1 Torque authority

Since the remaining reaction wheel can control pointing along only a single axis, the MTBs are now required to provide pointing control in the two orthogonal axes. There are two major disadvantages to using magnetic control: the torquer bars generate much less torque than the reaction wheels (a maximum of about one-tenth), and their ability to provide torque is not constant – they are dependent on the instantaneous direction and strength of the local magnetic field at a given time.

Gravity gradient forces provide the largest disturbance torque on the *FUSE* satellite. The difference in the gravitational force at the two ends of the ~5 meter long satellite tends to pull the satellite into a nadir- or zenith-pointing orientation; the resultant torque can be as large as 5 milliNewton meters, and the wheel and MTBs must counteract this torque for the spacecraft to remain inertially pointed. The ability of the MTBs to provide the appropriate counteracting torque for this and other disturbances is a function of the continually varying relative orientation of the earth's magnetic field and the satellite, since the dipole field will have little effect if pointed in a direction nearly parallel to the earth's magnetic field lines. Since the earth's field and the satellite orbit are known, the availability of adequate magnetic control can be predicted as a function of time and satellite orientation; times when control is available are said to have positive "torque authority."

The available torque authority changes constantly; some orientations (particularly pointings in the direction toward the orbit poles) will have positive torque authority for many contiguous orbits, while others will vary over much shorter timescales. Losing torque authority means that the disturbance torques overwhelm the ability of the satellite to compensate, and fine pointing control is lost. If the loss of torque authority is brief, there may be only a small disturbance in the pointing, and the object being observed may remain in the aperture; longer losses make it more likely that the pointing cannot be held steady. Experience has shown that allowing modest losses of torque authority and scheduling extra time on an object makes available scheduling windows that could not otherwise be used. The CalFUSE pipeline can then be used to filter out the periods with large pointing errors (Section 4).

Figure 1 shows Torque Authority Contours (TACOs) on a map of the southern hemisphere for March 8, 2006 at 02:00 and 06:00 UT. The dark contours show the regions with positive torque authority for at least 85%, 90%, and 95% of the time over a three-orbit period. In both maps, the TACOs cover a relatively small region of the sky, and these regions are preferentially near the orbit poles (near ±65° declination). The map on the left shows a TACO region at higher declination (to the right in the figure) which has nearly disappeared four hours later. Over just these four hours, the shapes and positions of the TACOs vary considerably, illustrating the challenge of maintaining stable pointing for an extended time. Only targets falling inside a TACO on both figures would be considered stable for the entire period.

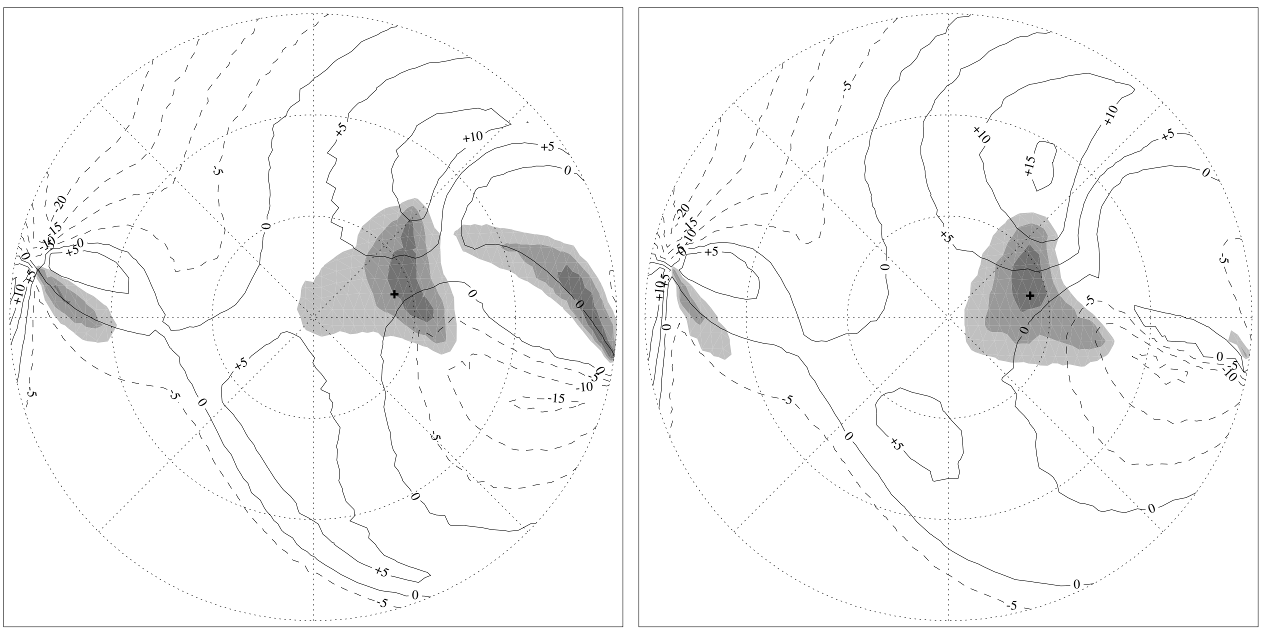

Figure 1 Two sky maps, separated by four hours, showing the southern hemisphere with torque authority contours (grey shaded regions) and momentum buildup (solid and dashed contours) overlaid. The grey contours show the regions where the torque authority is positive for 85%, 90% and 95% of the time over three orbits, when the satellite roll angle is -15°. The units of momentum buildup are N m s per orbit. In the four hours shown, the region with stable torque authority on the right side disappears, and the one near the pole changes its size and shape.

3.2 Momentum management

The original purpose of the MTBs was to manage the angular momentum of the reaction wheels. Since the external torques on the satellite do not in general average to zero over an orbit, without some means of removing the resulting angular momentum, the wheels would soon reach their limit of ±21 Newton meter seconds (Nms) and would no longer be able to provide pointing control. The dipole field generated by the MTBs can be oriented to provide torque on any two axes perpendicular to the magnetic field. In one-wheel operations, magnetic torque is required to control both axes perpendicular to the remaining wheel, so the component of magnetic torque along the wheel axis is not a free parameter. If the MTBs were used to continuously control the wheel momentum as in the original design, stable pointing of the satellite would be continually interrupted. Consequently, momentum is now managed primarily by judicious selection of targets, and automatic unloading via the MTBs is limited.

Figure 1 also illustrates the angular momentum buildup as a function of position on the sky. The solid and dashed contours show the average increase or decrease of angular momentum with time, in units of N m s per orbit. Momentum buildup can change rapidly over a few orbits in some parts of the sky. In this example, the TACOs and momentum buildup are complementary; the momentum change inside the TACO near the pole is only a few N m s per orbit, and both positive and negative values are available. Such an arrangement makes planning observations relatively straightforward, since targets can be scheduled to control momentum buildup.

3.3 Flight software improvements

Three full software builds and several patch loads have been made over the past year to implement new ideas, improve the robustness of operations, and gradually improve satellite performance. These changes to the flight software have added: two robust safe modes (one requiring ephemeris data and one not); a completely new torque distribution algorithm, which has been tuned over the past year; an improved treatment of the cross-coupling between axes, including vector-limiting of the control and unloading torque commands; continued tuning of the controller gains and the ability to change the gains on-the-fly; and an enhanced ACS to instrument-computer interface to allow attitude information to be shared between the two. Further modifications will be made as necessary once we gain more experience operating the satellite in one-wheel mode.

3.4 Planning observations

As described above, additional constraints have been placed on the planning and scheduling process in one-wheel mode. The scheduling of observations now must now take into account the limited regions of torque authority, along with the management of momentum. As a result, only a small region of the sky is accessible for observing at any time. Considerable effort has been expended to develop planning and analysis tools to improve the predictions of pointing and improve science observing efficiency. Ever more sophisticated tools have been developed to minimize the effects of observing constraints. Improvements to the planning tools are as important as upgrades that have been made to the flight software; our new software can provide arcsecond pointing if there is enough torque available, but no control if not. It is up to mission planning to schedule observations only when the needed torque is available.

In order to successfully adapt to the new set of scheduling constraints, a series of modifications were made to the *FUSE* Long Range Plan (LRP) software: (1) The predictive tools were modified to accurately determine the time periods with torque authority for all the targets in the *FUSE* database, (2) the range of allowed roll offsets was relaxed to ±25° to increase the number and length of visibility windows, (3) the LRP oversubscription per week was increased to optimize the target sample for satisfying the torque authority and momentum management constraints when short term mission planning schedules are created, and (4) additional science and background programs were added to the target pool to improve scheduling flexibility.

Figure 2 shows the amount of observing time available during a one year period in the north celestial pole region for declinations above 45°; a plot showing declinations below -45° looks similar. All pending targets are overplotted, with the size of the symbol representing the amount of requested time. It is clear that the available observing time has a strong dependence on declination. For this reason, the call for proposals for Cycle 7 limited targets to absolute declinations above 50°.

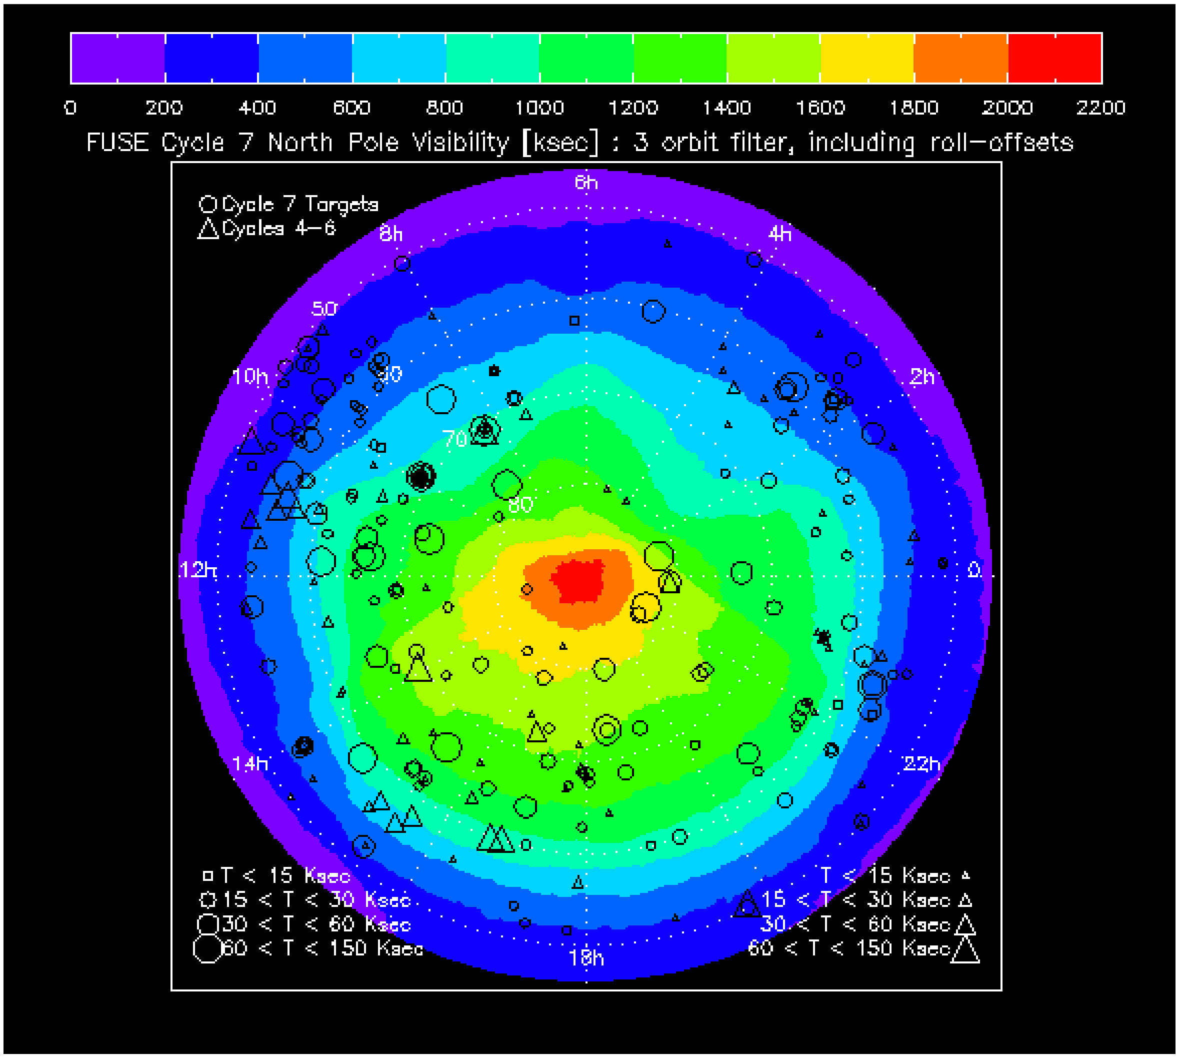

Figure 2 One year sky visibility in the northern hemisphere.

The figure shows the total amount of observing time available during the year, but it does not show how the visibility windows are distributed in time. For a typical one-week period, only a small subset of these targets can be observed, since the region of positive torque authority on the sky typically changes position, shape, and size on timescales of hours. In addition, preexisting constraints for the instrument, such as solar avoidance, further limit the available targets on any given day. Thus, the larger the pool of targets, the more efficient is our use of observing time.

Using the results of the LRP, the Mission Planning Schedules (MPSs) are developed. A typical one-week MPS includes observations of ~8-10 unique targets. The goal is to plan as much observing time as possible in regions of high torque authority, while keeping the wheel momentum well away from saturation. Simultaneously optimizing these disparate requirements often requires compromises: observing a high priority target which is in a region of positive momentum buildup, for example, often must be followed by a less-desirable target with negative buildup to avoid saturating the wheel.

4. THE CALFUSE PIPELINE

The CalFUSE science data pipeline processes raw science and engineering data to create spectra. The latest version (3.2) has recently been described in detail by Dixon et al.[8], so this section will provide only a brief summary, focusing on changes made to account for some of the effects described above.

As originally designed, the CalFUSE pipeline assumed that the instrument was stable during an observation. Early data showed that thermally-induced motions of the mirrors and gratings, along with changes in the detector format as a function of temperature and count rate, made this assumption incorrect, and no simple changes to the software would permit these effects to be included. As a result, the pipeline was completely redesigned beginning in 2002 so that most of these effects could be included, and in doing so, the design became more flexible. A fundamental design change was that some instrument and spacecraft engineering data, such as detector count rates, high voltage levels, aperture positions, and pointing information were incorporated to track changing conditions during an exposure. As a result of these improvements, very few modifications were necessary when the change to one-wheel mode observing occurred. The pipeline had already been modified to compensate for minor pointing variations (jitter), and this modification also easily handled larger pointing excursions due to short losses of torque authority.

Because of these changes the science data quality is, in many cases, the same now as it was before the loss of the roll reaction wheel. Tests on selected data sets show that the spectral resolution has not changed, despite the fact that the pointing jitter has increased and much larger pointing excursions are common. One effect of our less-stable pointing is that observations of a target may now consist of a larger number of shorter exposures than before, so it may be necessary to combine more exposures to achieve the same signal-to-noise ratio. Since long *FUSE* observations have always required the observer to combine exposures in this way, the tools for doing so were already available.[9]

All science data taken earlier in the mission are now being reprocessed and re-archived at MAST[10] in order to take advantage of these improvements. This reprocessing, which is expected to be the last full reprocessing of the data, should be completed by the end of 2006. Calibration files that reflect the current properties of the instrument will continue to be provided for the life of the mission.

5. CURRENT PERFORMANCE AND FUTURE PLANS

A significant amount of development and testing of the one-wheel control system occurred in the spring and summer of 2005, allowing *FUSE* to return to full science operations on November 1, 2005. This section describes the performance of the satellite as of April 2006.

5.1 Pointing stability

The pointing stability of the satellite is crucial for obtaining high-quality spectra in several ways. In the largest spectrograph aperture (30 × 30 arcsec) stable pointing ensures the highest possible spectral resolution. In the smaller apertures (4 × 20 arcsec and 1.25 × 20 arcsec) stable pointing increases the instrument throughput. The recent improvements to the CalFUSE pipeline can correct time-tag (low count rate) data for target motions within the aperture. However, photons lost when a target drifts out of an aperture can never be recovered. Early in the mission, pointing stability was ~0.3 arcsec rms, better than the pre-launch specification of 0.5 arcsec. In two-wheel operations, this increased to ~0.5 arcsec, with the largest errors along the weak axis (45° to the dispersion axis). The loss of the third reaction wheel did not appreciably decrease stability during times of good torque authority, but instead added additional periods of much larger (>30 arcsec) pointing excursions.

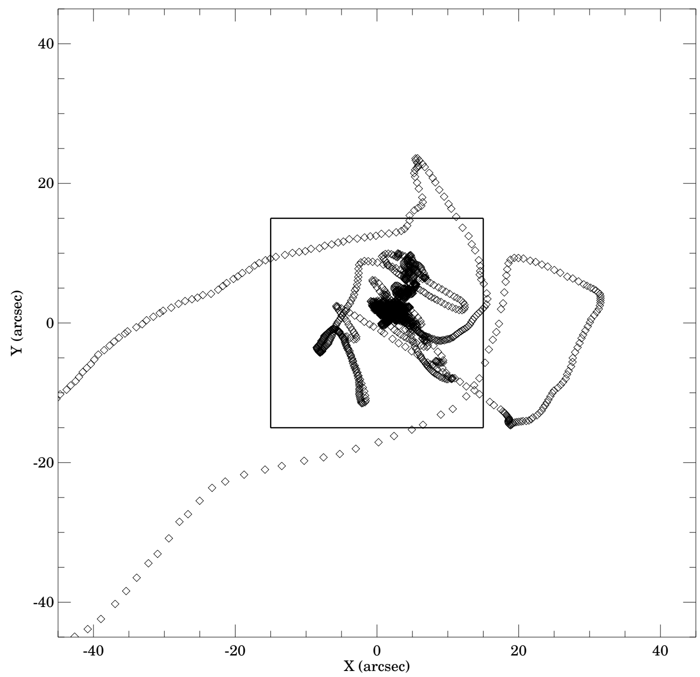

Figure 3 Pointing jitter during an exposure taken in March 2006. The symbols show the position of the boresight at one second intervals. Although the target remains well within the 30 × 30 arcsecond aperture for most of the exposure, there are several pointing excursions which bring it outside the aperture.

Figure 3 shows an example of the pointing stability of a science exposure made near the south pole of the orbit during the time covered by the sky map in the left panel of Figure 1. The symbols show the position of the instrument boresight once per second. The size of the 30 × 30 arcsecond aperture is overlaid. Although the target remains inside the aperture for most of the exposure, there are several excursions which bring it well outside.

In Figure 4 data from the same exposure is displayed as a function of time for 1000 seconds around a pointing excursion. This figure shows the pointing jitter, the count rate of the guiding channel, the angle between the magnetic field and the skew axis, and the momentum. As expected, the count rate drops to near zero during the excursion because the target leaves the aperture. This corresponds to a time when the angle between the skew axis and the magnetic field is close to 90° and magnetic control becomes difficult; the vertical lines on the angle plot mark the times when torque authority was predicted to be lost. The momentum plot shows both the predicted (top) and measured (bottom) wheel momentum. Although the shape is similar, the prediction has been shifted by 12 N m s so that it fits on the same plot. This example was chosen to highlight a period of bad torque authority; it should be noted that many exposures show stable pointing for their full duration.

5.2 Momentum predictions

Momentum management via the mission planning process introduces complications for satellite operations, primarily due to the fact that real-time conditions are not used. In particular, the wheel momentum at any time is rarely exactly what was assumed when the MPS was planned. This difference can build up through several mechanisms. For example, during a slew the gyroscopic torque is sensitive to the magnitude and sign of the initial momentum. When the gyroscopic torque is the main driver of the momentum variation during a slew, differences between the planned and actual momentum at the beginning of the slew can lead to significant differences by the slew's end.

Another complication is due to the influence of the angle between the magnetic field and the axis of the skew wheel. When the magnetic dipole axis is nearly parallel (within ~5°) to the earth's magnetic field (or alternatively, the skew axis is nearly perpendicular to it), little magnetic control is available, and thus wheel usage is typically at a maximum. Predicting the momentum variation during such periods is difficult because it also corresponds to torque authority gaps. Consequently, a small change in the spacecraft attitude (~1°) due to a loss of torque authority can change both the magnitude (e.g. 4 Nms within 20 min) and the sign of the variation.

To avoid wheel saturation, which would cause the satellite to lose pointing control, the momentum is monitored by the mission operations team. If the actual momentum strays too far from the prediction, a manual intervention is made to temporarily point the satellite towards a location that will return the momentum back towards the predicted value. These interruptions from the planned timeline typically last for 3 – 7 orbits and change the wheel momentum by 10-20 Nms.

During these interruptions no science data is collected, so when possible, they are made during background observations or other periods when they are likely to have the smallest effect on the *FUSE* science program. If a Guest Investigator target must be interrupted, it is rescheduled for a later time.

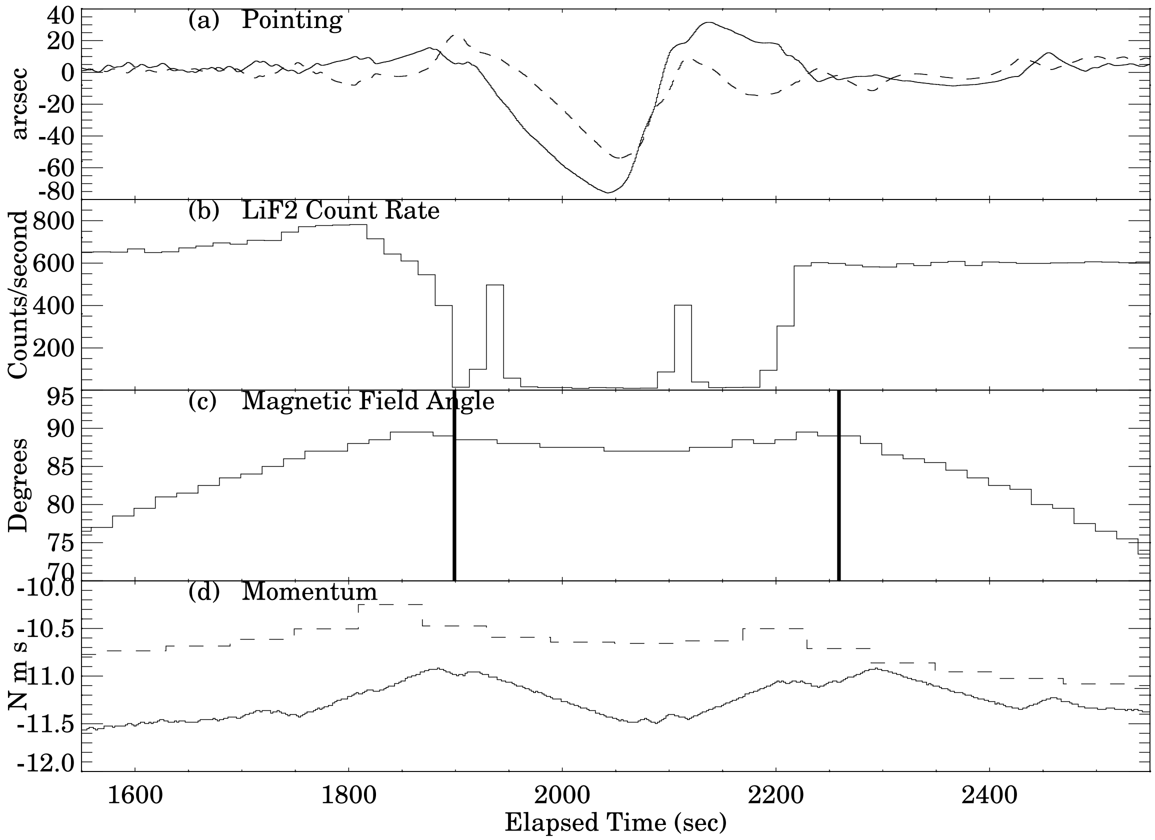

Figure 4 (a) x (solid) and y (dashed) pointing jitter as a function of time for 1000 seconds during the exposure shown in the previous figure; (b) the count rate on the LiF2 (guiding channel) for the same time; (c) the angle between the skew axis and the earth's magnetic field and the predicted times when torque authority was lost (vertical lines). (d) The predicted (dashed) and measured (solid) momentum on the skew wheel. Although the shape is similar, the prediction has been shifted in order to fit on the same plot.

One of our near-term goals is to decrease the frequency of these momentum interventions. This will be possible as we gain more experience with the current flight software and spend more time analyzing its performance. We have developed a large number of analysis tools to assist us in understanding what happens onboard. This should enable us to exercise more options during the planning process, such as using 180° rolls to provide improved opportunities for managing the momentum in real time.

5.3 Observing efficiency and instrument performance

The average science efficiency during the prime mission was approximately 30%. Automation of science operations and other improvements before the last reaction wheel failure had led to a slight increase in this number during the past several years of extended mission operations. Much of 2005 was spent testing new methods and procedures, but since science operations were officially restarted at the beginning of November, there has been a steady improvement (Figure 5) in science observing time, along with a major decrease in the time spent in non-science safe modes.

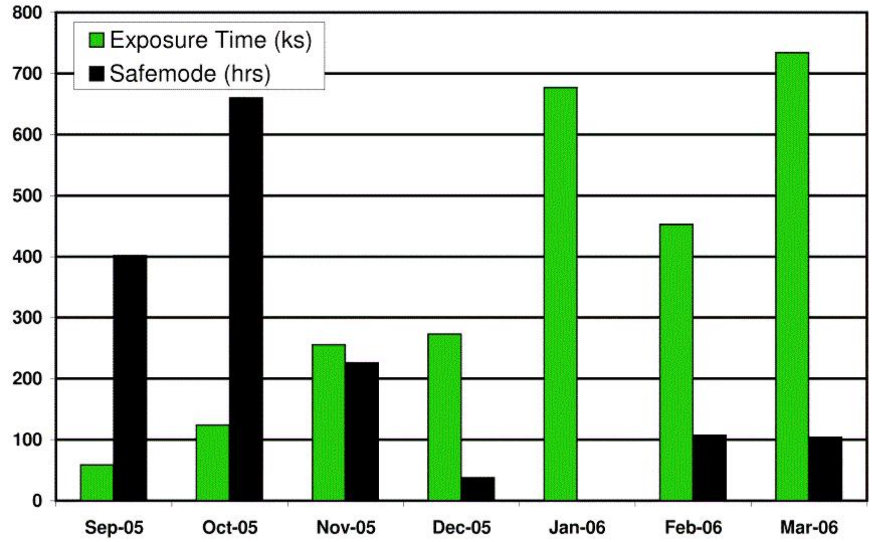

Figure 5 On-target observing time and time in satellite safe modes since the fall of 2005. Since the resumption of science operations in November 2005, the observing efficiency has increased dramatically, and is approaching the value obtained earlier in the mission.

The fact that observations are currently constrained to high declinations has had a positive effect on these numbers, since observing near the orbit pole means that occultations of targets by the earth are short or nonexistent. The stability of alignment between the four optical channels also seems to be similar to what it was before 2005, despite a decrease in scheduled alignment activities. This may also be due to our pointing constraints, since limiting the angular separation between targets (a consequence of limiting pointing within a TACO) minimizes the thermal changes to the instrument, and thus helps maintain channel alignment.

Recent calibration measurements have shown that no unexpected changes in the performance of the science instrument have occurred. The effective area for the long-wavelength (LiF) channels has remained constant, and the response of the short-wavelength (SiC) channels has been slowly decreasing as before. Spectral resolution is unchanged, even in exposures where modest losses of torque authority occur.

5.4 The future

A large number of changes have been made to flight software, planning tools, and analysis tools during the past year. Major modifications are finished for now, and the next step is to gain a better understanding of the actual on-orbit pointing performance and make minor adjustments to the system. Since it has been just a short time since our last major code load, limited engineering data under restricted conditions has been collected so far. More experience with the real performance will allow the development of more reliable predictive tools so that we will be better able to plan for (or avoid) periods of bad pointing.

Our performance was still improving in two-wheel mode when the last wheel was lost; we expect that similar ongoing improvements will be made for quite some time in one-wheel mode. A few of the priorities for the near future are: (1) the investigation of different slew algorithms to provide more flexibility and improved sky coverage, (2) the continued improvement of mission planning tools for both long range and short term scheduling, and (3) the further study of making observations with the satellite roll angle flipped by 180°, which would significantly change the angle between the earth's magnetic field and the skew axis at any given time and help in the management of the momentum profile.

The *FUSE* mission has been a tremendous scientific success, having obtained over 3700 observations on more than 2300 unique objects or pointing positions on the sky. Through the end of 2005, the mission has generated over 53 million seconds of on-target science data. As of March 2006, 355 scientific papers have been published in the refereed literature, and the number climbs with each passing month. As we approach seven years into the mission, the observing community continues to propose cutting-edge scientific programs that use FUSE for investigations that were not even conceived of at the time of launch.

ACKNOWLEDGEMENTS

We thank the Orbital Sciences Corporation for their important contributions to the development of the algorithms and software used in one-wheel mode, and the *FUSE* Mission Operations Team from Honeywell Technology Solutions, Inc. for their support during the testing and implementation phases. This work is supported by NASA contract NAS5-32985 to the Johns Hopkins University.

REFERENCES

1 H. W. Moos et al., "Overview of the *Far Ultraviolet Spectroscopic Explorer* Mission," *ApJ* **538**, pp. L1 – L6, 2000.

2 D. J. Sahnow et al., "On-orbit Performance of the *Far Ultraviolet Spectroscopic Explorer* Satellite," *ApJ*, **538**, pp. L7– L12, 2000.

3 J. W. Kruk, P. Chayer, J. Hutchings, C. Morbey, and R. Murowinski, "*FUSE* Fine Error Sensor optical performance," *Proc. SPIE* **4139**, pp. 163-174, 2000.

4 J. W. Kruk, B. F. Class, D. Rovner, J. Westphal, T. B. Ake, H. Warren Moos, B. Roberts, and L. Fisher, "FUSE In-Orbit Attitude Control with Two Reaction Wheels and No Gyroscopes," *Proc. SPIE* **4854**, pp. 274-285, 2002.

5 W. P. Blair, J. W. Kruk, H. W. Moos, and W. R. Oegerle, "Operations with the FUSE observatory," *Proc. SPIE* **4854**, pp. 241-250, 2003.

6 B. Bialke and G. Dorsey, "FUSE Reaction Wheel Torque Anomaly Resolution," Advances in the Astronautical Sciences **107**, pp. 441-458, 2001.

7 B. A. Roberts, J. W. Kruk, T. B. Ake, T. S. Englar, B. F. Class, and D. M. Rovner, "Three-axis Attitude Control with Two Reaction Wheels and Magnetic Torquer Bars," AIAA Guidance, Navigation, and Control Conference, Providence, RI, Aug. 2004.

8 Dixon et al., "CalFUSE v3: A Data-Reduction Pipeline for the *Far Ultraviolet Spectroscopic Explorer*," in preparation.

9 *http://fuse.pha.jhu.edu/analysis/analysis.html*.

10 *http://archive.stsci.edu/fuse/*.