The FUSE detectors: on orbit use and lessons learned

David J. Sahnow*
Department of Physics and Astronomy, The Johns Hopkins University

ABSTRACT

The FUSE (Far Ultraviolet Spectroscopic Explorer) instrument includes two large-format microchannel plate detectors with double delay line anodes. The generally good detector performance has permitted the collection of scientific data with high spectral resolving power, and has enabled the observation of fainter objects than could be easily observed with previous missions in this wavelength range. As with any complex instrumentation, however, there have been numerous challenges which have arisen during the mission. We discuss the on-orbit performance of the FUSE detectors since launch, and describe some of the lessons learned. This includes a discussion of their operation on orbit and the effects that detector performance has had on the scientific data collected. The strategies taken to minimize the impact of detector anomalies on the data will also be discussed.

Keywords: Far Ultraviolet Spectroscopic Explorer, FUSE, microchannel plates, delay line

1. INTRODUCTION

The Far Ultraviolet Spectroscopic Explorer (FUSE) is a NASA-funded mission designed to obtain high resolution far ultraviolet (900 - 1200 Å) spectra of astronomical objects. FUSE was launched into a low-earth orbit on June 24, 1999**.** The prime mission consists of three years of operations, with the observing time divided between the Principal Investigator Team and the Guest Investigator (GI) Community. NASA has recently approved at least two more years of operations, with GI-only observations. Details of the FUSE mission, its science goals, and its early on-orbit performance have been given previously.[1,2]

The FUSE instrument consists of four coaligned channels. Each channel includes an off-axis parabolic telescope mirror, a holographically-ruled diffraction grating, a Focal Plane Assembly containing four spectrograph entrance apertures, and half of a large-format microchannel plate detector with a double delay line anode. The mirrors and gratings on two of the channels are coated with silicon carbide (SiC) in order to maximize the throughput at the lowest wavelengths; the optics on the other two channels are coated with lithium fluoride (LiF) over aluminum for maximum efficiency above ~1020 Å. The two LiF channels also contain a visible light Fine Error Sensor (FES) camera, which is used for fine guiding. Details of the FUSE spectrograph design are given in Green et al.[3]

This paper will describe some of the major issues that have affected the performance of the FUSE detectors during the three years since launch. It will discuss the lessons learned so that future missions utilizing similar detectors can benefit from the FUSE experience. Similar detectors are planned for future instruments, including the Cosmic Origins Spectrograph (COS) on the Hubble Space Telescope.[4]

2. THE DETECTOR AND ITS ON-ORBIT OPERATIONS

The two FUSE detectors are identical microchannel plate (MCP) detectors with double delay line anodes. The active area of each is divided into two segments, each with its own MCP z-stack, anode, and associated electronics. The < 10 mm gap between the two segments allows each to be operated independently without affecting the other. The details of the design of the FUSE detector hardware has been presented earlier,[5,6] and will not be repeated here. However, it is worth noting that the detector is by nature an analog device, with no fixed pixels. The pixel values reported by the electronics are the results of an analog-to-digital conversion of the calculated photon position; this fact has important implications for many of the effects discussed below.

---

*sahnow@pha.jhu.edu; phone: 410 516 8503; fax: 410 516 5494; Department of Physics and Astronomy, The Johns Hopkins University, Baltimore, MD 21218.

The detector Data Processing Unit (DPU) controls the detector by changing its state in response to commands that come from the ground or the instrument computer, known as the Instrument Data System (IDS). The DPU contains code which processes the incoming events and calculates the (x,y) positions. There are actually two copies of this code onboard - one in a memory area known as lower core, and the other known as upper core. The detector is also responsible for some of its own health and safety monitoring; it monitors the CRC of its own code images, for instance, and watches the count rates to make sure that they are within some allowed range. Out of specification conditions result in a diagnostic code which is sent to the ground, and if the problem is serious enough, may result in the shutdown of the unit. The FUSE mission design requires a great deal of autonomy on board, since ground station passes are infrequent in low earth orbit; this scarcity of ground contacts complicates non-standard operations, such as ramping up the high voltage.

The detector system also includes two stimulation, or “stim” lamps, which are used to illuminate the detector directly, without going through the rest of the optical system. These lamps (one for each detector) were not designed to be stable with time, and provide only a quasi-uniform illumination of the detectors. Illumination with the stim lamps show the shadows of the two 95% transmissive grids which are in front of the microchannel plates. The first of these is the plasma grid, meant to keep charged particles from reaching the MCPs. The second is a “QE grid,” which enhances the detector efficiency by collecting photoelectron from the MCP web. Both have wire spacings on the order of 1 mm in both x and y.

Each detector segment has six counters which are continuously updated whenever the detector is powered on. These are the Fast Event (FEC), Digitized Event (DEC), Active Image (AIC), LiF, SiC, and Autonomous Shutdown (ASC) counters. The first two measure the total number of counts which are detected and processed over the entire segment. The remaining counters are each controlled by a ground-programmable mask which can be used to include or exclude regions of the detector from being considered. The AIC mask is used to define which regions of the detector data are included in the data sent from the detector to the IDS; photons landing in excluded regions are discarded. The SiC and LiF counters are used primarily for peakup purposes; each of these is used to count only the events in a certain region (those arriving from the appropriate SiC or LiF channel aperture). The ASC - also known as the SAA (South Atlantic Anomaly) counter - is used to monitor a part of the segment that does not have any spectrum falling on it, and thus is a measure of the detector background. If the count rate in this region reaches a ground-programmable level, the segment’s internal SAA protection reduces the voltage, since the instrument may be mistakenly passing through the SAA with the voltage at its full value. This protection is also sometimes tripped by large event bursts (section 3.2).

Events that are not excluded by the AIC mask are sent to the IDS as 32 bit words: 14 bits for x, 10 for y, 5 for pulse height, 1 for detector, 1 for segment, and 1 to identify the word as a photon event. The IDS then saves the data in one of two formats: as a photon list (also called time tag or TTAG mode) or a spectral image (also known as histogram or HIST mode). In TTAG mode, the entire 32 bit word is saved, and time markers are inserted periodically (at 1 second intervals by default). In HIST mode, a two dimensional image of the detector active area is made, and only the number of counts that fall in each pixel are saved; the pulse height and timing information are discarded. A further set of masks, known as SIA tables, is provided by the IDS and is used to map regions of the detector to regions in IDS memory. Since HIST images are usually binned by 8 in y, and only the regions around the primary aperture are saved, taking data in that mode can save memory when the count rate is above ~2500 counts per second. Science data that is saved by the IDS is later downlinked during a ground station pass, and reassembled into a list of incident photons. Next, the CalFUSE pipeline[7] is run to convert these events into calibrated spectra for scientific use. CalFUSE is also responsible for removing the artifacts introduced by the detector.

During normal operations the high voltage for a given segment toggles between two voltage levels: ‘full’ and ‘SAA.’ The full level is the level at which the gain is high enough to collect counts, and is adjusted from time to time to compensate for gain sag (see section 3.9). The SAA level is a much lower, safe level, which is low enough to prevent incident photons from generating events; this level was chosen to be low enough so that the large number of counts generated when passing through the SAA have very low gain and do not affect the lifetime of the MCPs. Since there is no shutter on the instrument, the only way to stop photons from being detected is to drop the high voltage from full to SAA level. A single detector command changes the voltage level between full and SAA (or vice versa). The full voltage level has been changed three times since launch, and future adjustments are likely to be made every 6-12 months in order to correct for gain sag, as described below.

Even though the detector has a built in high count rate protection (typically called ‘SAA’ protection), that protection is not used routinely, and passages through the SAA are managed by the IDS using an onboard SAA manager. The SAA manager maintains a list of times for SAA passages, and sends an ‘SAA Reduce’ or ‘SAA Resume’ command as appropriate. This commanding is done asynchronously with the onboard scripts that control the data collection; thus problems with those scripts do not affect the health and safety of the detector.

Since late 2001, an onboard ‘occultation’ manager has also been in place. The occultation manager is used by the observation scripts to drop the high voltage to the SAA level when the target is occulted by the earth, and also at most other times that no exposures are being saved. This was implemented in order to minimize the gain sag due to exposure of bright airglow lines, which continually illuminate the detectors (see section 3.9). This extra protection is necessary because the system is designed so that the detector collects data constantly, independent of whether the IDS is collecting an exposure; the detector continually collects data (no matter what the high voltage level), and the only difference is that during an exposure, the IDS saves the collected data rather than ignores it.

Most of the time, the detector high voltage toggles between full and SAA high voltage levels. Exceptions include the annual Leonid meteor shower, when the high voltage was turned off for safety reasons, and other on-board maintenance which requires the IDS to be shut down. Some single event upsets, or SEUs, which are described in more detail in section 3.1, also turn the high voltage completely off, and a set of commands must then be sent from the ground in order to return the voltage to normal levels.

3. ON-ORBIT CONCERNS

This section describes some of the most important issues that affect the quality of data from the FUSE detectors. Many of these effects are now addressed by the CalFUSE pipeline, with varying degrees of success.

3.1. Single Event Upsets

Soon after the detectors were powered on, the operations team discovered that the CRC values in code memory were spontaneously changing during some SAA passages. Since the SAA is a region containing a large number of energetic charged particles, it was soon realized that these changes were radiation-induced bit flips; subsequent analysis showed that the memory chips used were not radiation hardened enough. Early in the mission, these Single Event Upsets (SEUs) were occurring at the rate of about one event every 1½ days. Although most of these changes were benign, there was concern that an important part of memory could be affected (such as that controlling the high voltage values, for example), and put the health of a detector at risk.

A large effort was therefore mounted to minimize the effects by reloading the DPU code whenever an SEU occurred. Because two copies of the code are normally stored in RAM simultaneously, this was usually as simple as stopping execution of the corrupt code, swapping to the region containing the uncorrupted copy, replacing the bad code from the ground, and then continuing on. Since the code loads could only be done during ground station passes, however, there were times when both cores would be corrupted before a fix could be made, and the safest course of action was to reboot the DPU and ramp up the high voltage again. This had a major impact on efficiency early in the mission. By late 1999, IDS code was uploaded to allow the DPU code to be stored on board (in the more radiation-hardened IDS memory) and loaded autonomously whenever a CRC change was detected. This greatly improved the likelihood that an SEU would be repaired before another occurred, and decreased the number of reboots necessary. In addition, the memory area included in the CRC was decreased to include only the part of memory which contained the code image, rather than the entire region of RAM. Despite these improvements, slightly less often than once per month an SEU still occurs which causes the detector to reboot due to the watchdog timer timing out, or the IDS noting a problem with telemetry.

One other change made after the discovery of the susceptible memory was that variables controlling the SAA count rate limits, the SAA high voltage levels, and the current limits were refreshed after each pass through the SAA to minimize the chance that they will become corrupted. No change has ever been noted in these memory regions.

3.2. Event Bursts

Event bursts are a still unexplained feature of all four segments of the FUSE detectors. A typical burst consists of a temporary increase in counts; the total time can be anywhere from several seconds to several minutes, and the increase in count rate can be as large as tens of thousands of counts per second. A typical burst seen early in the mission created a narrow streak across the active image of the detector; sometimes these streaks also had scalloped edges. Later, bursts which nearly uniformly filled the active area of a segment began to appear along with the original ones. A given burst can appear in one or more segments, and one or more detectors, but it is usually much stronger on one segment. See Fig. 1 for an example of a set of bursts which is visible on multiple segments. During the 1800 seconds shown, first a burst occurs on segments 2A and 2B, with a stronger signal on 2B. A few seconds after the burst ends, side 1 shows a large burst on both segments. About 1000 seconds later, side 1 has another burst on both segments.

Since the pulse height characteristics of bursts are similar to those of normal photons, they do not appear to be *caused* by the detectors, but rather just *detected* by them. Bursts often repeat on time scales of exactly one orbit, which implies they are related to the observing geometry. There also seems to be some long term or seasonal dependencies, since the number seen varies significantly throughout the year.

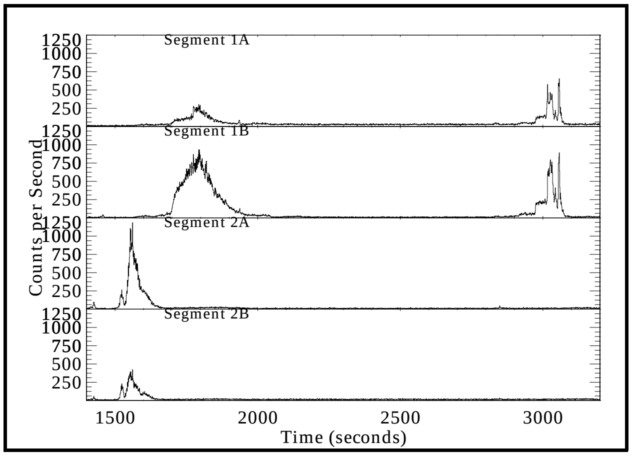

Figure 1 Multiple bursts on all four segments are seen in this figure. Note that bursts are often correlated between segments of the same detector.

Bursts have only two effects on the collected data. A large burst can substantially raise the background for a short time during an exposure, so the CalFUSE pipeline now excludes time periods containing bursts from TTAG data during processing; this normally has only a small effect on the data quality, since the bursts are short compared to the exposure length. The other effect is that a large burst can trip the SAA count rate protection on the detector; the large increase in counts is similar to what would be seen if the instrument went through the SAA with detectors at full voltage, so the detector drops the high voltage to the safe SAA level, and ground intervention is required to return it to full voltage with the present threshold values. These shutdowns now occur on the average of once per month; the high voltage is usually restored during the next ground station pass.

3.3. Current Spikes

Current spikes, also known as “crackles” are short-duration current spikes seen in the auxiliary power and high voltage power supplies. A crackle can be as short as several milliseconds (the time sampling is 1 ms), or longer than 60 ms. The DPU code monitors these currents, and shuts off the high voltage if the current stays above a ground-programmable threshold for more than a specified time (originally 20 ms, now 60 ms). Because the shutdowns occur so quickly, it appears that there is no long term effect on the detector. There is an effect on efficiency, however, since observations taken on a segment shut down by a crackle are lost until the high voltage is returned to normal via ground command, which can take as long as 24 hours, depending on the timing of the ground station passes. Crackles seem to have a similar ‘seasonal’ variation to the bursts, and, in fact, the two seem to be somewhat correlated. Because crackles are relatively rare, however (recently, on the order of one every 60 days), it has been difficult to prove this correlation.

3.4. Background Events

The intrinsic detector background is very low on all four segments. The total rate per segment is on the order of 10 counts per second during most orbits. The only exception is when the satellite passes near the SAA (but not close enough to fall within the boundaries we have defined), or through other localized, temporary regions of high background. Because even these cases normally raise the background by only $\lesssim$50 counts, and the fact that a spectrum covers only a fraction of the more than 16 million pixels on the detector, even in these cases, the effect on the spectrum is extremely small, and is usually negligible compared to other (non-detector) effects such as scattered light.

3.5. Quantum Detection Efficiency

On orbit, it is not possible to measure the quantum detection efficiency (QDE, or QE) separately from the mirror reflectivities and grating efficiencies, but it appears that there has been very little degradation since launch. The total measured change in effective area has been 5 - 30%, depending on channel and segment, and since there does not seem to be much of a correlation between channels which illuminate the same segment (e.g. LiF1A and SiC1A), most or all of that decrease is likely due to coating degradation on the mirrors and gratings. The flux from the stim lamps, which illuminate the detectors directly, varies with time, and therefore cannot be used to show QDE variation; in addition, most of its photons are out of band.

3.6. Dead Time

Detector dead time has been measured on orbit by measuring the fraction of stim pulse events which are recorded as a function of total count rate. An error in the algorithm used on board resulted in this number being higher than anticipated; the dead time is 10% at ~8000 counts per second for each segment. The dead time also has a component due to the IDS, since the IDS can process a maximum of 8000 counts per second in TTAG mode, or 32,000 counts per second in HIST mode. The larger than expected dead time has only a minor effect on the data quality, however, since most FUSE observations have relatively low count rates.

3.7. Thermal Stability

A consequence of the analog nature of the FUSE detectors is that the x and y pixel scales change as the temperature changes; both a shift and stretch of the pixel scale are seen in the data. During a typical orbit, the detector thermal environment varies by less than 1° C, but between different targets at different orientations with respect to the sun and earth, these variations can be as large as several degrees. Since the changes in scale are on the order of 5 pixels (over the entire detector) per degree, this can result in errors of ~10 pixels over the entire length of the detector. This effect is corrected for in the CalFUSE pipeline by measuring the positions of two electronic stimulation (or “stim”) pulses which are introduced into the detector electronics at a known position at the beginning and end of each exposure. The position is measured in both x and y in order to determine how much of a change in format has occurred. There are two problems with this approach: the y position of the two stim pulses for a given segment are located so close together, that no stretch information is available in that direction; and the stim pulses are turned on only for 60 seconds at the beginning and end of each exposure, so no temporal adjustments can be made.

3.8. Geometric distortions

A major consequence of the analog nature of the detector is that the pixel scale is not constant, i.e. not all pixels are the same size. The variations are found on multiple spatial scales, and can cause variations in the x width of pixels of a few percent. In the y direction, large distortions are visible by eye, giving the spectra a wavy appearance in the raw data. Distortion measurements were made on the ground before detector delivery, but the detector electronics settings were adjusted after that, limiting their usefulness. Since the effect of x distortions is simply a change in the wavelength scale, the CalFUSE pipeline does not attempt to correct for it, and instead includes it in the wavelength scale. Future versions of the pipeline will separate these two effects by using the detector grid wires to determine the distortion. The y distortion correction already uses the grid wires.

3.9. Gain sag

As microchannel plate detectors are exposed to photons, the gain (the number of electrons per incident photon) drops. If the exposure were spatially uniform, simply raising the high voltage would be enough to return the gain to its original value. The FUSE optical design, however, does not give uniform illumination of the detectors. Each segment has six spectra falling on it, three from a SiC channel, and three from a LiF channel. In addition, these spectra are not constant as a function of wavelength. Although most objects observed by FUSE have continua with absorption lines, the details vary from object to object, and other effects, such as the location of the object in the slit, affects the exact pattern of photons striking the detector. In addition, bright emission lines from the earth’s atmosphere continually illuminate the detector from all slits, since there is no shutter in the system. And finally, there are regions of the detector that are not illuminated by any slit; their only exposure is to scattered light and background events.

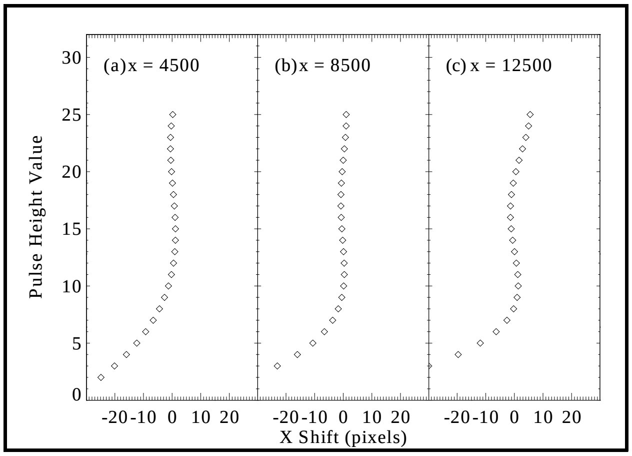

Figure 2 X shift (walk) change with pulse height for 3 positions in x along segment 2A. Below a pulse height of ~10, there can be a significant error in the reported x position of an event.

As a result, the integrated illumination pattern on the detectors is quite complicated. Some regions, such as the bright airglow line at H I Lyman-β (1026 Å) have received tens of thousands of counts per pixel since the beginning of the mission, while certain background regions have received only a tiny fraction of that. Because of these great differences in exposure, there is a great difference in modal gain across the detector, even though it started out roughly constant. Thus, if the voltage were to be raised such that the regions with the most gain sag were returned to the value at launch, other regions would have a gain that was too high (possibly leading to a higher background, for example). It has therefore been necessary to leave some of the most gain-sagged regions below their optimal value in order to keep the majority of the detector at a reasonable value. The most serious effect due to the low gain regions are 'walk' effects - the x position of a photon is a function of the gain. The walk variations are negligible above pulse height values of ~10, but below that level, it has a strong variation with position (primarily x) on the detector and gain. Fig. 2 shows the shift in position of an event as a function of pulse height value at three positions on segment 2A, including one near the bright O I 989 Å airglow line at x≈ 2500. The figure shows that if the pulse height is below about 10, depending on detector x position, there can be severe errors in the reported x position of the photon.

Because the walk can seriously distort the data, a walk correction module has been added to the CalFUSE pipeline. Since essentially no walk data was collected before launch at the final electronics settings, the data for this correction had to be determined after launch, by taking stim lamp exposures at a variety of voltage settings, and measuring how the positions of the grid wires changed as a function of x and pulse height (walk is independent of y, since the y position is determined by charge division, rather than time delay). After collecting many hours of lamp data, a correction was obtained for each x position on the detector. One of the first steps of the TTAG calibration pipeline now adjusts the x position of each photon, based on its position and the pulse height value. For HIST data, where no pulse height values are available, the correction routine must estimate the gain for a given pixel (based on the x position on the detector, and the amount of exposure that region of the detector has seen since launch), and applies a mean (and therefore much less satisfactory) correction. Walk effects are constant with time (except for the fact that gain changes with time) for a given set of electronics adjustment, so we continue to make these low gain stim lamp observations in order to refine the walk correction. Fig. 3 shows the two dimensional image of a gain-sagged O I 989 Å line (taken shortly before the high voltage was raised), along with a walk-corrected version of the same data. The corrected spot is more symmetric, and has more of the light concentrated in a few pixels.

Although the gain sag effect was expected, it turned out to be much more serious than expected, because of a combination of the way the detector electronics were trimmed, the gain properties of the MCPs, and the fact that the highest effective area occurs near the edges of segments, where the walk effects are the largest. Once it was realized that this was a serious problem, detector operations were changed so that the detector voltage is lowered to the SAA level most times when no exposures are being made. Thus, we minimize the exposure to the bright airglow lines.

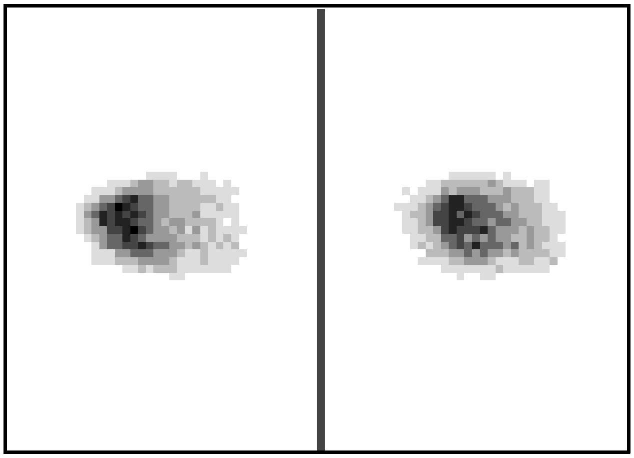

Figure 3 A raw (left) and walk-corrected (right) O I 989 Å airglow line. The walk corrected spot is more symmetric and somewhat narrower.

3.10. Y Blooming

An unexplained property of the detector is that the y scale is a function of count rate. As the count rate changes, the pixel positions of two spectra, for example, move further apart in pixel space. This peculiarity was not recognized for some time after launch, and made application of the bad pixel masks (see section 3.11) problematic. This effect appears to be independent of x, and has now been measured by measuring the positions of dead spots as a function of count rate. A correction factor is applied in the CalFUSE pipeline so that all data is in the same (count rate = 0) reference frame. Fig. 4 shows how the pixel scale changes at 3 locations in y on segment 1B.

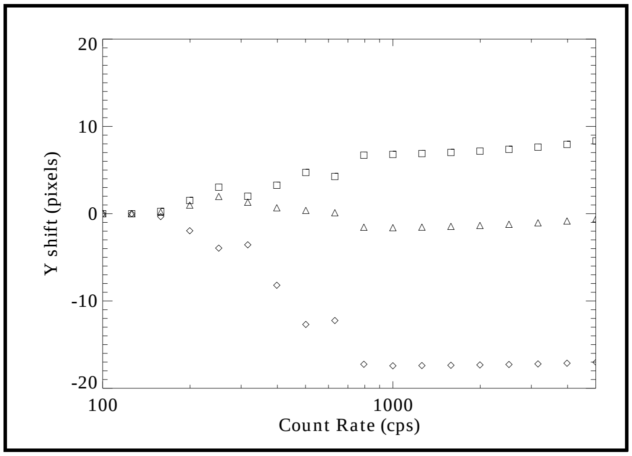

Figure 4 The y shift of the data as a function of count rate at three different y positions on segment 1B. The squares are for a position near the top of the active area, the triangles are in the middle, and the diamonds are near the bottom.

3.11. Flat Field and Fixed Pattern Noise

Even three years after turn-on, it is still difficult to determine an appropriate flat field for the FUSE detectors. Part of the problem is the large number of non-detector effects, such as the change in position of the spectra on the detector due to grating motions, which complicate the data analysis. No good ground flat fields were obtained because the temperature was not stable during spectrograph integration and test. Several segments also show an obvious moiré pattern[8] at a level of up to 15%. Numerous dead spots are visible on all segments, and photons falling in those regions of the detector are lost or compromised. A bad pixel mask has been generated, and it is used to exclude all data that falls in those regions during the calibration process. The flux must be adjusted appropriately, of course, to compensate for the lost area.

3.12. The Worm

The feature known as the "worm" (Fig. 5) is due to an unfortunate interaction between the horizontal focus of the spectrograph and the wires in the QE grid. The location of this focus in front of the detector varies with wavelength; it roughly coincides with the grid on the long wavelength end of the LiF channels, so the effect is most pronounced there. At other wavelengths and in the SiC channels, it is also present, but often at a level which is low enough to be unnoticeable. The worm can attenuate on the order of half of the photons at some wavelengths, and it moves significantly across the detector as the target moves in the slit.

At present, the CalFUSE pipeline does not correct for the worm, but recent calibration observations have been obtained during which a bright target was purposely trailed along the slit in order to determine how the worm varies. The goal is to construct a worm map as a function of position and wavelength, so that the science data can be corrected appropriately.

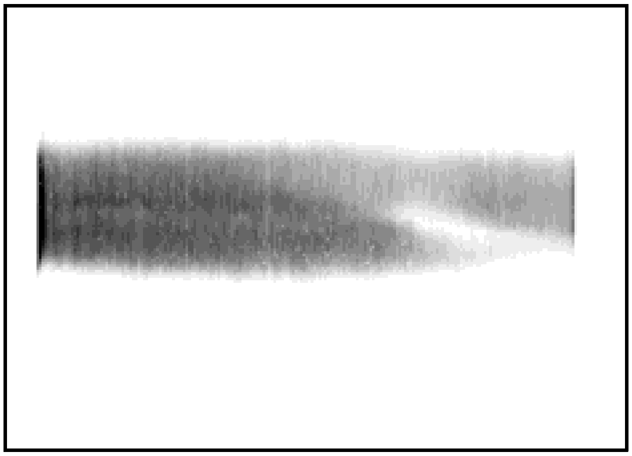

Figure 5 The worm in the LWRS aperture in LiF1B. This data has been corrected for y distortion and binned in the x direction for display purposes..

4. OVERALL PERFORMANCE

The previous section focused on the problems encountered in orbit, but overall, the detectors are performing very well. After three years of nearly continuous on-orbit operation, there are no major signs of degradation except for gain sag. The system is very robust; we have also been very cautious and have been willing to temporarily allow lower observing efficiency in order to ensure a long and vigorous life. None of the issues described above are insurmountable now that they are understood, although a number of them did cause interruptions or inefficiencies in data collection or analysis while they were initially being investigated. Resources early in the mission were focused on understanding and compensating for the thermal motions of the optics, and detailed analysis of the detector performance was delayed until more data had been collected. Since the thermal motions cause the spectra to move on the detector, the detector data obtained are more complicated, and in many cases must be adjusted for this fact. The worm (section 3.12), for example, is dependent on the y position of the spectrum on the detector. Since this changes (differently for each channel) during an exposure, it is much more difficult to correct for than if the gratings and mirrors did not move with changing thermal conditions.

As discussed in the previous section, the QDE remains at or near the preflight level, and all other detector properties appear to be stable over time. Event bursts, crackles and similar complications can usually be dealt with by simply removing the bad time periods, which result in only a small reduction in net observing time. The major complication to the data analysis is instead the analog nature of the detector; understanding the properties and how they vary on orbital time scales has been a major focus of the Science Data Processing team since launch. As we continue to improve our understanding, the need for additional calibration data becomes apparent. An example of this is the fact that we now collect stim lamp exposure data in both HIST and TTAG modes. Early in the mission, it was collected in HIST mode only, since the high count rate made this the most efficient. But it was later realized that the pulse height information was crucial for understanding the properties of the detector on both large and small scales, so TTAG observations were also begun. Other special purpose calibration observations are also scheduled as the need arises.

The detector DPU code has been upgraded twice since launch (and once before launch, but after detector delivery). The first upgrade was done primarily to more efficiently respond to SEUs, and the second simply changed a hard-coded threshold value. It is likely that no further updates will be made, since the current version of the code appears flexible enough to handle future operations.

5. LESSONS LEARNED

With three years of detector operations behind us, what has been learned? The main lesson is that the flexibility of the detector system has been its most important asset when dealing with the unexpected. This adaptability has allowed a relatively quick response to the unforseen problems encountered on orbit. The clearest example of this is the response to the SEU problem early in the mission. The redundant, dual-core DPU code design allowed an operational workaround almost immediately, and allowed operations to continue while a long term solution was developed.

One area in which more flexibility would have been helpful is that not all detector parameters can be modified from the ground. The fact that the crackle time threshold was hard-coded meant that a new version of DPU code was required to decrease the crackle frequency, and any changes to the code is a potential source of errors.

Many of the complications encountered in the detector are due to its analog nature. We now understand that the pixel scale is a function of (at least) temperature, count rate, and pulse height, for example. The first of these was known before launch, but the implications were not understood; the prelaunch calibration data, for example, was not taken with the detector temperature well-controlled. As a result, the pixel scale varied during exposures, limiting the usefulness of the data. The y blooming was unanticipated, and was only discovered after a large amount of data had been collected and analyzed. Although we were aware before launch that detector gain was a limited-lifetime item, the full implications of this were not understood. No reliable data on the MCP lifetime was available, and no modeling was done to determine how much charge would be extracted from the MCPs over the life of the mission. For the first two years, the detector was at full voltage and collecting data for nearly the entire time that the satellite was not in the SAA. This simplified the operations scenario, but resulted in serious gain sag at some locations, and required a complicated retrofitting of the occultation manager onto the system.

For all of the pixel scale effects, it was extremely helpful to have as much information in the raw data as possible. In TTAG mode, corrections for these effects can be applied long after the data is taken, and improvements to the CalFUSE pipeline are currently underway to take advantage of some our recent understanding of how the system performs. HIST data is much more difficult to correct.

Some unintended detector benefits due to the optical design are the fact that there are multiple spectra on each segment, and all wavelength regions are covered at least twice. This allows the separation of the spectral features from the detector effects.

A recent change to the CalFUSE pipeline makes use of some housekeeping data when calibrating the science data. The initial design precluded this, since the two types of data followed distinct processing paths. The recent loss of several reaction wheels prompted this change in order to compensate for larger pointing variations. Now that the framework is in place to allow this, it should be a simple matter to include other housekeeping data, such as the counters in order to calculate a time-variable dead time and y blooming correction, both of which can change dramatically during an exposure if there are bursts.

Some of the on-orbit concerns are a consequence of the fact that FUSE is a modest-size mission. A small team works best with a relatively simple operations scenario, since only a limited number of people are available (and knowledgeable about the relevant system) to handle a particular problem. Having the flight operations team be heavily involved in the construction, integration and test of the subsystems was crucial to the on-orbit efficiency, since it provided experience with the instrument systems before launch. Similarly, colocating science analysis and mission operations has proven invaluable in identifying and solving problems. A small operations team also means that operations may be overly cautious, such as occurred early in the mission when the SEUs were discovered. Although this caused the efficiency to suffer, it was prudent to understand the problem before possibly putting the detector at risk. The small team also explains why it has taken several years to recognize and address all of the issues discussed above.

On the other hand, the small staff and informal working environment facilitates a quick response to problems and provides the ability to schedule calibration observations on very short notice under some circumstances. Although last minute changes to the mission planning schedules can be a headache for the Mission Planners, this can often allow the diagnosis of a problem to occur much more quickly than would otherwise be the case.

Having a backup detector (and someone to run it) also proved useful on numerous occasions, particularly early in the mission when changes to the DPU code were being tested.

Although the FUSE design contains only two detectors, in practice, each segment must be treated separately for nearly all calibration issues. Each segment has different properties, and the electronics and voltages on each are tuned differently. The advantages of having the two segments in one mechanical housing are mainly optical (both segments are on the same Rowland circle) rather than operational. The only commonality is that the plasma grid wires seen in the stim lamp exposures is the same physical grid in the two segments of one detector, and the spectrum that falls on one segment continues smoothly (except for the gap between segments) on to the other. Because of all of these differences, we have often not taken advantage of times when it is advantageous to handle each two-segment detector as one unit. Choices were made to treat each segment individually (in the pipeline, for example), which generally makes sense, but doesn’t use all the information available.

Although the detector blemishes have made the interpretation of the scientific data more complicated, many of these features have turned out to be useful for calibration purposes. For instance, the grid wires have been used to calculate the walk correction and the distortion; the dead spots have been used to measure the blooming as a function of count rate.

As is probably the case with most space-based instruments, having more time on the ground for calibration would have been very helpful. Some of the complexities discovered after launch may have been found this way, but many would still have remained undiscovered. Ground calibration time is always in short supply because of schedule constraints; thus the time available must be used efficiently. Taking data under conditions as close to on-orbit as possible can be very illuminating, but there will always be limitations on how well this can be simulated.

Finally, it is important to not underestimate the complexity of the data. We initially assumed most parameters would be stable, but now we believe that essentially everything is varying: either with time, position, temperature, or pulse height. Although the current CalFUSE pipeline has been expanded to handle many of these effects, it has became clear that it is not efficient to continue making modifications in this manner, and it may make more sense to redesign it from scratch. Being able to start over when necessary is an important lesson.

Additional information about the FUSE mission can be found at *http://fuse.pha.jhu.edu*.

ACKNOWLEDGMENTS

The author wishes to acknowledge the contributions of the members of the FUSE Science Data Processing group, who have contributed to some of the analysis described in this work. This work is based on data obtained by the NASA-CNES-CSA FUSE mission operated by the Johns Hopkins University. Financial support has been provided by NASA contract NAS5-32985.

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," *ApJ* **538**, pp. L7-L11, 2000.

3. J. C. Green, E. Wilkinson, and S. D. Friedman, "The design of the Far Ultraviolet Spectroscopic Explorer spectrograph," *Proc. SPIE* 2283, pp. 12-19, 1994.

4. J. B. McPhate, O. H. W. Siegmund, G. Gaines, J. V. Vallerga, and J. Hull, "The Cosmic Origins Spectrograph FUV Detector," *Proc. SPIE* **4139**, pp. 25-33, 2000.

5. O. H. W. Siegmund, M. Gummin, G. Gaines, G. Naletto, J. Stock, R. Raffanti, J. Hull, A. Abiad, T. Rodriguez-Bell, T. Magoncelli, P. Jelinsky, W. Donakowski, and K. Kromer, "Performance of the double delay line microchannel plate detectors for the Far Ultraviolet Spectroscopic Explorer," *Proc. SPIE* **3114**, pp. 283-294, 1997.

6. D. J. Sahnow, M. A. Gummin, G. A. Gaines, A. W. Fullerton, M. E. Kaiser, and O. H. W. Siegmund, "On-orbit performance of the double delay line detectors for the Far Ultraviolet Spectroscopic Explorer," *Proc. SPIE* **4139**, pp. 149-162, 2000.

7. "The CalFUSE Pipeline Reference Guide," *http://fuse.pha.jhu.edu/analysis/pipeline_reference.html*.

8. A. S. Tremsin, O. H. W. Siegmund, M. A. Gummin, P. M. Jelinsky, and J. M. Stock, "Electronic and optical moiré interference with microchannel plates: artifacts and benefits," *Appl. Opt.* **38**, pp. 2240-2248, 1999.