

FUSE attitude control: target recognition and fine guidance performance

Thomas B. Ake^{a,b}, H. Landis Fisher^c, Jeffrey W. Kruk^a,
Patricia K. Murphy^c, and William R. Oegerle^a

^aDept. of Physics and Astronomy, The Johns Hopkins University,
3400 N. Charles St., Baltimore, MD, USA

^bScience Programs, Computer Sciences Corporation,
Lanham-Seabrook, MD

^cThe Johns Hopkins University Applied Physics Laboratory,
Laurel, MD USA

ABSTRACT

The FUSE satellite employs innovative techniques for autonomous target acquisitions and fine pointing control. One of two Fine Error Sensors, incorporated in the optical path of the science instrument, provide the Instrument Data System computer with images, for target identification, and field star centroids, for fine pointing information to the spacecraft attitude control system. A suite of “toolbox” functions has been developed to locate stars, selected and track on “unknown” guide stars from the image, identify the star field, track preselected “known” guide stars, follow moving targets, and provide pointing optimizations to fine-tune the centering of a target.

After a maneuver to a new field, initial attitude is determined by identifying stars found in a $20' \times 20'$ image. Identification is done by matching stars with an uploaded table of up to 200 objects selected from the Hubble Space Telescope (HST) Guide Star Catalog (GSC), ranging from $V = 9$ to 13.5 mag., and typically covering a one degree field around the target. During identification, tracking is performed on unidentified stars in the image to prevent the satellite from drifting. A corrective slew is then commanded to place the target at the desired position. Tracking is then resumed on preselected guide stars. If desired, further fine alignment of the science apertures is performed by a target peakup using the FUV detectors.

We discuss the target acquisition process; end-to-end performance; and problems encountered due to the limitations of the small field of view of the FES, HST GSC errors, and stray light in the telescope baffles.

Keywords: FUSE, Satellites, Attitude control

1. INTRODUCTION

The *Far Ultraviolet Spectroscopic Explorer* (FUSE) is a low earth orbiting satellite used to obtain high resolution spectra of astronomical sources.¹ The science instrument consists of four telescopes and associated Rowland-circle spectrographs, designated as either LiF or SiC channels depending on the coating of the optical elements. Three entrance apertures for the spectrographs are cut into four focal plane assemblies (FPAs), one for each spectrograph, which can be adjusted in the grating dispersion direction to keep the selected target centered while inertial pointing is maintained by the spacecraft. Fine guidance is performed using one of two Fine Error Sensors (FES)² which view the reflective surface of one of the LiF channel FPAs to obtain an optical view of a $20' \times 20'$ field around the target. The FES provides full field images for target identification, and centroid measurements for up to six guide stars when operated as a tracker.

All mission and science operations are performed from The Johns Hopkins University (JHU) utilizing a dedicated antenna at the University of Puerto Rico Mayaguez and a commercially available backup ground station in Hawaii. Since communication passes occur for only half the orbits per day, and last about ten minutes each, the satellite was designed to operate in an semi-autonomous mode. Science observations and housekeeping activities are prescheduled

E-mail: ake@pha.jhu.edu

on the ground, and the real-time contacts are used to load and dump the onboard control systems and check on spacecraft health.

As part of in-orbit operations, the FUSE satellite must be able to recognize the celestial field of a scheduled target, perform corrective motions to place the object in the desired spectrograph aperture, and acquire and track on guide stars during the observation. These functions are performed by the FUSE Instrument Data System (IDS) using image and centroid information from an FES. The IDS flight software is a combination of a commercial off-the-shelf system, Spacecraft Command Language (SCL) by Interface and Control Systems, Inc., and customized routines written in C for FUSE, developed at the Applied Physics Laboratory of JHU.³ The FUSE guidance task development was part of the latter effort.

Because of the breadth of the types of observations to be made by FUSE and the uncertainties of in-orbit performance, the IDS flight software for target acquisition and guiding was developed as a suite of “toolbox” tasks that can be tailored for each observation. As experience was gained, most of the activities and control tables have become standardized. Many functions are accomplished using SCL flight scripts permanently installed onboard, with detailed control using table loads and commands in temporary scripts tailored for the scheduled observation.⁴

Besides onboard software, a mission planning ground system is used to create a detailed timeline of activities, interleaving slews, target acquisitions, and science exposures with orbital events such as earth occultations and crossings of the South Atlantic Anomaly (SAA). The Hubble Space Telescope (HST) Guide Star Catalog (GSC)⁵ and Digitized Sky Survey (DSS; see <http://www-gsss.stsci.edu>) are used to generate information on field stars for uploading to the IDS.

2. FUSE TARGET ACQUISITIONS

FUSE observations are scheduled one to two weeks in advance of their execution. At that time, stars expected in the FES field of view (FOV) are examined by a mission planner for the targets being observed. The 200 brightest objects in the HST GSC within 30' of a target are requested using the electronic HST guide star interface, as is an image from the DSS for direct comparison with the catalogued stars. The GSC is not a complete magnitude-limited survey, but rather was designed to provide a uniform coverage of potential guide stars. It contains positions and magnitudes of 15 million stars from $7 < V < 16$, with relative positions good to 0.2–0.8" and magnitudes to ~ 0.3 mag. The magnitude range attainable with the FES, and its field of view, matches well the brightness and spatial distribution characteristics of stars in the GSC. By displaying the DSS image with the GSC entries, the mission planner can decide which stars appear to be best for tracking. Up to six can be selected (five if the target is to be measured as part of the acquisition; see below). Criteria include brightness, isolation from contaminating sources, and arrangement in the FOV. The 200 field stars are placed in a star table for use by the onboard star identification process (Sect. 3.2), and a list of the stars for guiding is prepared for command generation.

Observation scripts are created through an automated tool, and are uplinked to the IDS several orbits before the scheduled execution time. These scripts contain the calls to perform slews, initiate target acquisitions, command guidance control, and obtain exposures for the science program. Activities are executed either at an absolute time, such as at the end of an earth occultation for the target, or in relative time after a specific event has occurred, such as starting an exposure as soon as the target is acquired.

A target acquisition sequence is started in relative time after a slew completes or in absolute time when the target reappears after an earth occultation or SAA passage. It begins with a one-second, unfiltered FES exposure of the field. Slew errors and drifts are small enough that even for the longest slews and occultation periods the target is in the FES FOV. Images are normally taken in 2×2 binned mode to reduce the transfer time to the IDS across the data bus. It takes about 15 seconds to read out and transfer the image to the IDS.

The rest of the acquisition then proceeds in the semi-autonomous, relative-time mode. The IDS processes the FES image, determining the locations and intensities of stars in the field. Positions are determined to ~ 0.1 pixel, or 0.5" in 2×2 binning. Objects brighter than $V \sim 8.5$ are usually saturated in the peak, but stars can be detected down to $V \sim 14$. The IDS then chooses stars from this list to track on temporarily. The FES is commanded into centroid mode, such that small subimages, usually 25" or 40" on a side, are taken around the stars. For each one, the FES removes an average background based on the perimeter pixels of the subimage, corrects for cosmic-ray hits by looking for illuminated pixels not seen in the previous one, then computes a flux-weighted position of the center of light. The centroid information is made available to the IDS at one-second intervals for attitude computations.²

While in this “unknown star” guidance mode, the IDS attempts to identify the field based on the uploaded star table of 200 objects. This can take from several seconds to over two minutes depending on the complexity of the field. Unknown tracking prevents the spacecraft from drifting and allows the ACS to settle and update drift rates. Once the field is identified, a transition is made to “known” tracking mode. In this activity, a corrective slew is performed based on the star identification calculation. This takes out the residual error from the initial maneuver to the target or drift after an occultation and places the target in the appropriate aperture for the observation. FES subarrays are then placed at the predicted positions of the guide stars selected on the ground for known tracking. Another FES image is not necessary. Once in known tracking, FUSE is ready to start the science exposure.

The sequence of taking an FES image, achieving unknown tracking, performing star identification, and completing known tracking can take as little as 3 minutes. Settling after the initial maneuver is scaled by the slew length, (settle time in seconds = $3.33 \times$ slew length in degrees), so the initial acquisition of a field can take up to 11 minutes, but is typically about 6. Tracking can be maintained when the spacecraft passes through the SAA, but the task is idled during occultation. Acquisitions after earth occultation require no settling time. Spacecraft drifts during occultation are typically 1–2 arcminutes, which is larger than the FES subimage size. Thus it is not possible to simply lock up on the guide stars since they will have moved out of the subimages. The variation in acquisition times is one reason that relative time commanding is used for subsequent activities.

The accuracy of the pointing in the FUSE apertures is ultimately dependent on the position of the target as provided by the observer compared to the guide star frame. It is possible to refine the pointing by measuring the position of the target itself with the FES, as long as it is not too bright, too faint, or too extended. In this case, the target is not placed in the aperture when going to known tracking, but is put at a reference point about one arcminute away from the high resolution slit. During a ten second integration, the locations of the target and guide stars are all measured, and an offset position of the target is calculated. All subsequent slews take into account this offset, even after a reacquisition after earth occultation. Before the science exposure is started, a slew is performed to move the target from the reference point to the observing slit. The only special scheduling requirement is that the same guide stars be visible in the FOV for the science observation as for the target correction measurement.

A final pointing optimization can be made if the medium resolution (MDRS; $4'' \times 20''$) or high resolution (HIRS; $1.25'' \times 20''$) aperture is used. If the target is sufficiently bright in the FUV channels (~ 10 counts s^{-1} in each channel), the spacecraft can be scanned perpendicular to the narrow dimension and the location of the peak flux can be measured for each individual channel. Nine steps are made (every $3''$ for the MDRS; $1''$ for HIRS), with ten second integrations at each point. For the LiF1 channel, which is the FES side currently used for guiding, the spacecraft is maneuvered to the spot with the highest counts. For the other channels, the FPAs are moved so that the slits are aligned with the LiF1 location. The peakup sequence takes about 8 minutes to execute, and is most necessary when HIRS observations are made because of the channel motions that occur over an orbital period.⁶ Peakups in the low resolution aperture (LWRS; $30'' \times 30''$) are not performed since the FES subimages cannot be made large enough to keep the guide stars inside during the step and dwell sequence.

3. FLIGHT SOFTWARE DESIGN

As part of the toolbox approach to the flight software, various uploadable control tables were designed to be used by the IDS. Table 1 lists the tables used for target acquisition and fine guidance. The majority of these do not have to be updated for each observation, but can be modified for special observation cases.

Except for centroid computations performed by the FES when in tracking mode, all attitude calculations are performed by the IDS. Coordinate systems and transformations in the instrument frame are defined by the IDS tables. The IDS removes optical distortions in the FES image, computes the spacecraft attitude with a quaternion estimator, QUEST,⁷ for both star identification and tracking purposes, and commands the FES to compute centroids for guide stars. Centroid acceptance checks are performed once per second while guiding. Measured quaternions, covariances, and configuration flags are sent to the spacecraft attitude control system (ACS) once per second for fine pointing management. In return, the ACS provides status messages back to the IDS when events such as slews are completed.

3.1. Image Processing

The image processing task repairs an FES image using a ground generated list of bad pixels and determines the positions and brightnesses of celestial objects in it. The list is used by the IDS for two purposes. First, it provides

Table 1. IDS Tables for Acquisition and Guidance

IDS Table	Purpose	Update Interval
bad_pixel_map	FES image repair	Infrequent
opt_dist_params	FES distortion correction	Infrequent
img_proc_params	Image processing control	Infrequent
star_id_params	Criteria for pattern matching	Infrequent
att_est_par_si	QUEST criteria for star ID	Infrequent
cent_verif_param	Centroid acceptance criteria for tracking	Infrequent
fes_parameters	Boresight/plate scale definition	Infrequent
sub_image_params	FES command parameters for guidance	Infrequent
att_est_par_guid	QUEST criteria for guidance	Infrequent
star_table	Field star positions from the GSC	Observation Specific
peakup_config_table	Step and dwell values	Observation Specific

candidates for the unknown tracking mode. The IDS will select the 3–6 stars brightest stars that appear in different quadrants of the image. FES subarrays 10 pixels wide are placed around the positions and “faux” measured quaternions are generated. The ACS is informed that the pointing error from the commanded one is not a true estimate and FUSE is held to keep the unidentified stars at their detected positions. Secondly, the list is input to the star identification task (Sect. 3.2) for determining the attitude at the time the image was taken. The actual position of FUSE is then computed and the ACS is informed to slew to take out the true error when going to known tracking.

Several parameters are available to control the processing (Table 2). Currently, no bad pixel processing is required and object detection is performed on the raw image. In addition, there is no explicit cosmic-ray rejection algorithm, unlike for the FES centroids. This then drives all target acquisitions to be scheduled outside of the SAA.

To process the image, the IDS divides it into a grid of 33×33 blocks, each 16 pixels on a side. The routine makes a single pass through the image, attempting to detect stars block by block. Block processing improves the local background determination, particularly in areas of varying contrast in the image. For each block, the mean and RMS intensity are computed, which provide the background subtraction value and detection level criterion for each pixel. Each pixel in the block is examined and a threshold test is applied to determine if the counts are significant. A pixel is considered to be “illuminated” if it is at least $1\text{-}\sigma$ and 100 counts above the background. For most images, the background $\sigma \sim 35$, so the 100 ADU minimum becomes the determining criterion in selecting illuminated pixels.

Table 2. Image Processing Parameters

Parameter	Function	Value
min_star_size	Minimum valid star size	4 pixels
max_star_size	Maximum valid star size	40 pixels
min_star_dn	Minimum peak brightness	100 ADU
max_star_dn	Maximum peak brightness	32000 ADU
nreport	Number of stars returned	30
block_col	Processing block column dimension	16 pixels
block_row	Processing block row dimension	16 pixels
t_alpha1	Initial threshold	1σ
t_alpha2	Final thresholding parameter	4σ

In high background regions, $\sigma > 100$, so the $1\text{-}\sigma$ threshold prevails. If the pixel passes the threshold test, a check is made of its nearest neighbors. Contiguous illuminated pixels become an object, and the flux-weighted centroid, peak and total intensity, and object size are determined.

When succeeding blocks are processed, information on bordering pixels from previous blocks is retained for continuity. An object spanning any block boundary is characterized by pixels from all contributing blocks. As the blocks are processed, the IDS tracks the 250 brightest objects, with fainter ones being dropped from the list if the field is exceptionally crowded.

After all blocks have been examined, a final pass is made through the object list, and only those that meet the brightness and size criteria specified by the image processing table are retained. At least 4 pixels must contribute to the object; an upper limit of 40 has been found eliminate extended objects and very bright stars. We have found that a lower brightness threshold of $4\text{-}\sigma$ helps to remove artifacts such as diffraction spikes and column bleed when bright stars are present in the image.

A final sorting by intensity takes place based on statistics of the entire image. Although 250 stars are retained by the IDS, only the brightest 30 are allowed to be reported. This number more closely matches the quantity of field stars typically available from the GSC in the FES FOV (Sect. 3.2). If too many stars are returned, a false pattern match with the star field can result.

Figure 1 shows the HST DSS and FUSE FES images for the globular cluster, NGC 6723. Objects in the GSC for this field are shown on the DSS image. The FES image shows the objects that were returned by the image processing task. Stars near the center of the cluster, which would not be useful for either star identification or tracking, were not returned due to the crowded nature of the region.

3.2. Star Identification

The star identification routine uses the objects found by image processing (Sect. 3.1) to determine the FUSE attitude at the time an FES image was taken. The star field is recognized using a pattern matching technique based on the separations and position angles between objects compared to those in the uplinked field star table. It does not explicitly rely on star brightnesses, i.e., no comparison is made on catalogued *vs.* measured intensity, but since the field star table and image processing results are both sorted in order of decreasing intensity, matches are first attempted on the brightest stars in the image.

Star identification occurs in two main steps and uses the parameters shown in Table 3. First, angular separations of pairs of objects in the field star table are compared with imaged star separations until a potential match is found. Only stars separated by at least $60''$ are considered. The separation must agree to within $2''$, a fairly stringent requirement. The star identification process is computationally expensive and this tight criterion prevents an undue number of false search paths from being followed. Once a potential match is made, an attitude estimate is computed from the pair and a check is performed that the computed roll angle is within tolerance (5°). This tolerance is large since FUSE has a relaxed requirement on spacecraft roll because it is not an imaging spectrograph. Since there is no knowledge of which star of the candidate pair is which, there is a 180° ambiguity in the roll that must be checked as well.

Table 3. Star Identification Parameters

Parameter	Function	Value
min_pair_ang	Minimum pair separation for match	$60''$
max_pair_ang	Maximum pair separation for match	$28'$
pair_ang_tol	Separation tolerance for first pair	$2''$
addl_mtch_tol	Separation tolerance for other pairs	$4''$
max_roll	Maximum position angle tolerance	5°
min_mch	Minimum number of matches	4
des_mch	Desired number of matches	12

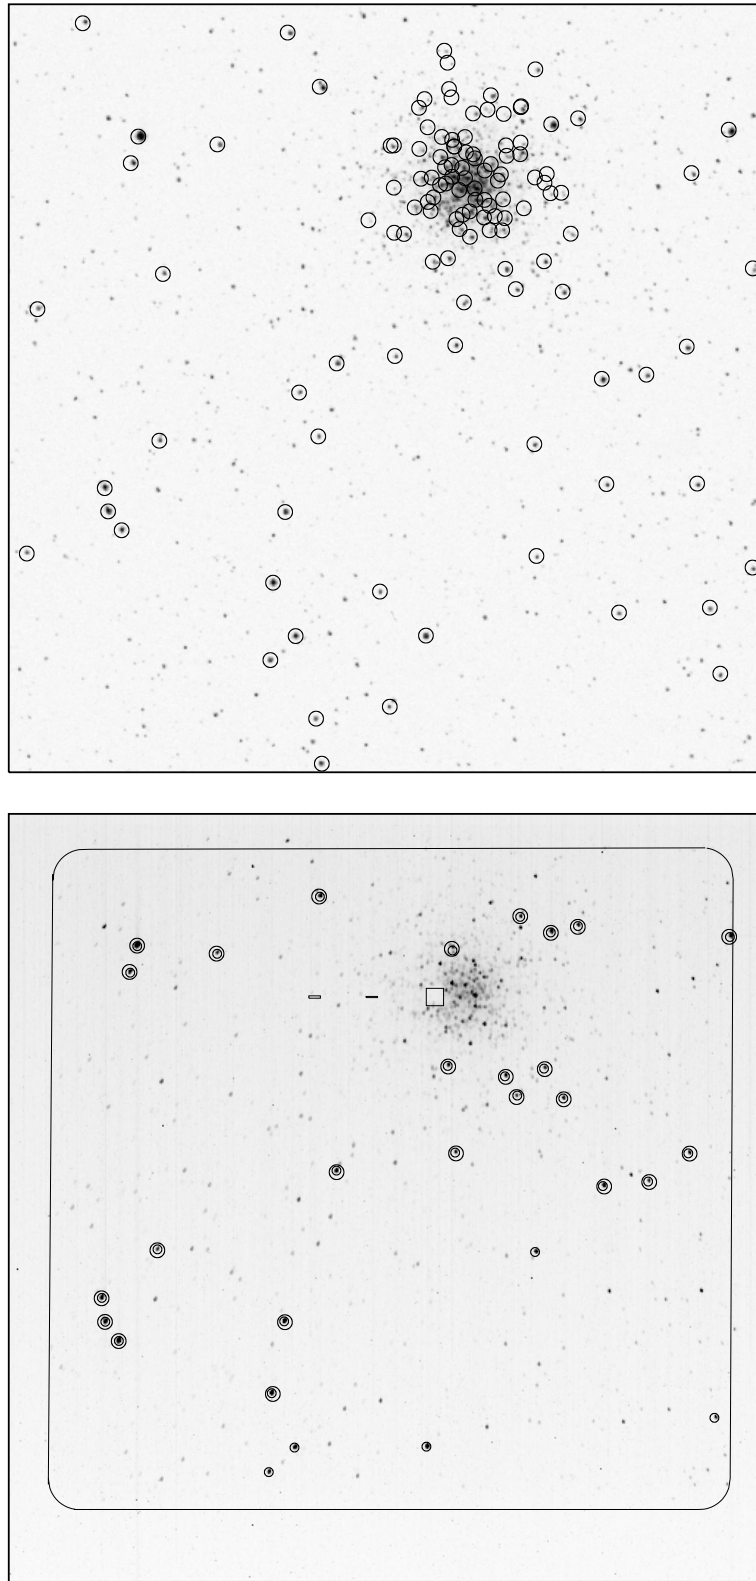


Figure 1. Top: DSS image of the globular cluster NGC 6723. Objects in the HST GSC are shown (*circles*). Bottom: FUSE FES image of NGC 6723. Objects found by the image processing task are shown as small circles; those that passed the pattern matching criteria with the GSC are noted with larger circles. Also shown are the locations of the FPA field stop (inner frame) and the three sciences apertures (LWRs on the right; HIRS, middle; MDRS, left).

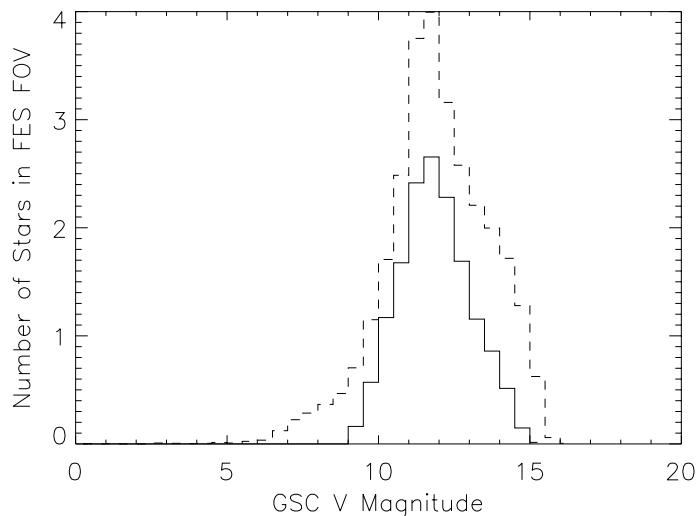


Figure 2. The average distribution of stars available in an FES field of view from uploaded star tables (*dotted line*) and those that are matched by the star identification process (*solid line*). The turnover in the magnitude of available stars for $V \geq 12$ is due to the limitation of the star table size and selection criteria used (see text, Section 3.2). Roughly two-thirds of available stars are found with the image processing and star identification parameters in Tables 2 and 3.

If the attitude passes the roll check, the second step begins. Using the attitude estimate, the field stars are converted into predicted image positions and a new search is performed on the imaged stars. The additional sources must be within $4''$ of the predicted positions. This criterion is larger than the initial match to allow for errors in the roll determination. If a sufficient number of matches is found, a new attitude solution is computed, including convergence tests, using the QUEST algorithm. If the attitude is not validated, the search begins with the first step again for another star pair.

At least four stars must be found for the star identification to be considered successful. Since this criterion is rather loose, the routine continues to search for other solutions until one is found that matches either at least 12 stars or all the stars returned by image processing.

The entire process typically takes from 15–60 seconds onboard. The second step occurs very rapidly once the correct pattern is found. If the routine does not complete within two minutes, it is stopped and an attitude determination failure is declared. This minimizes the chances that a false identification will be made based on only the faintest stars in the field star and image processing tables. The FES image in Fig. 1 displays the stars identified for NGC 6723. A few stars near the edges of the FES were not matched, likely due to residual errors in the optical distortion calibration.

Figure 2 shows the distribution by magnitude of GSC stars in the FES FOV for all observations from March–June 2000. The average number of stars with $V \leq 13.5$ is 23. The available stars artificially peak at $V = 12$ since only 200 objects can be loaded in the star table, which is sized to cover a one degree area around the target (a 3×3 raster of FES images). The star table is populated in order of increasing magnitude, so fainter GSC objects are preferentially left out. Figure 2 also shows the average number of stars identified. About two-thirds of the GSC stars are typically found. Figure 3 illustrates the distribution of the number matches typically made. While 15 stars are found on average, the range covers nearly the full allowable parameter space (4–30 stars).

The photometric calibration of FES intensities measured on the 2×2 acquisition images and based on GSC magnitudes is shown in Fig. 4 (left). The large scatter is due to inaccuracies in the GSC (~ 0.3 mag.⁵), the difference in responses of the unfiltered FES and the filters used for the GSC, stellar variability, and to some extent, the location of the stars on the subpixel scale of the FES. The FES intensities are typically larger due to its additional




Figure 3. The distribution of the number of stars identified in FES images from March–June 2000. The IDS image processing task returns up to 30 objects from an FES image (Table 2), which sets the maximum number of matches. The lower limit is 4 (Table 3).

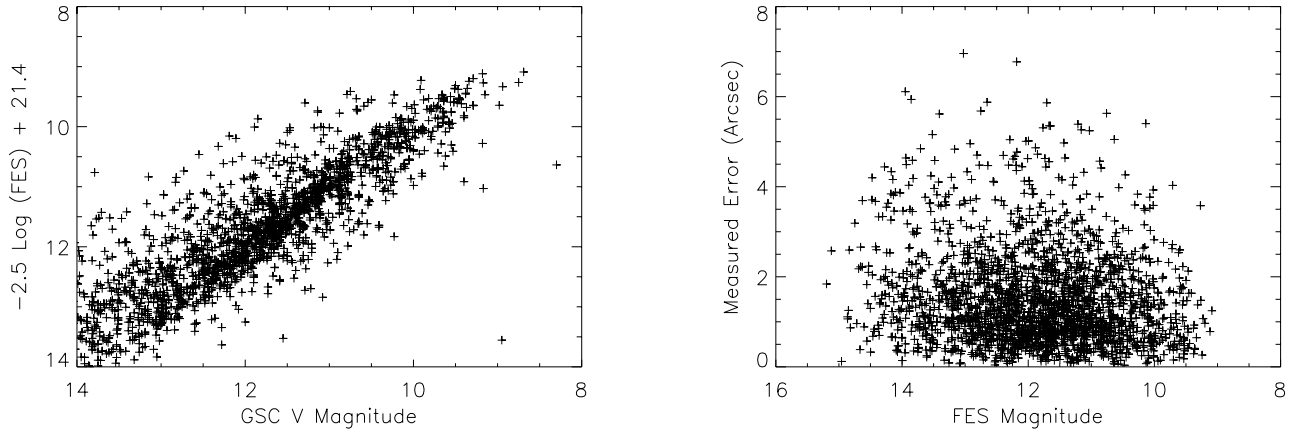


Figure 4. Left: Comparison of HST GSC magnitudes and those determined from FES 2×2 binned images. Right: Distribution of measured positional errors by magnitude.

red response.² Positional errors are independent of brightness (Fig. 4, right), a result unexpected based on prelaunch considerations. For the FUSE quaternion estimator, variances used for weighting positional accuracies are scaled from the intensity in a simplification of the algorithm. True variances in measured position would provide more accurate weighting of stars with poor positions.

The scatter in Fig. 4 illustrates why object intensity is not an explicit criterion in the star identification method used by FUSE and why up to six guide stars are chosen for tracking, particularly if only faint guide stars are present in the field. It is not uncommon for the FES measurements to be 0.75 mag. off from the GSC values, i.e., a factor of 2 in intensity. This has led to more conservative guide star selection criteria. For fields that are judged to be

susceptible to star identification failure either due to sparseness or faintness of available field stars, an offset field is selected for target acquisition.

A minor consequence of the scatter in the photometric calibration is that the star identification process can take much longer to complete when the field is dominated by stars of nearly the same brightness, as can occur in the Galactic plane or in the Large and Small Magallenic Clouds. Since image processing returns the table of detected objects sorted by intensity, the table in these cases is effectively unsorted. The star identification task must work through more false matches than in other fields. In a few instances, it has run longer than the two minute time out allowed on-board, and the acquisition failed.

4. POINTING AND GUIDANCE PERFORMANCE

The FUSE ACS is found to respond rapidly and smoothly to guidance requests from the IDS. Measured quaternions pass through a Kalman filter, along with magnetometer and gyro data, to produce a jitter of $\sim 0.35''$, well within the requirement to keep a target centered in the HIRS aperture. Drifts during periods without centroid updates are also small ($\sim 0.025'' \text{ s}^{-1}$). This is particularly fortuitous because stray light in the telescope baffles during orbital noon and low earth limb angles² raises the background level in the FES centroids so that fainter guide stars are temporarily lost. The spacecraft pointing is stable enough that the stars reappear in the subimages after noon passage without having to perform another reacquisition sequence.

Early in the FUSE mission, target acquisitions were failing regularly in the transition from unknown to known guidance mode. It was soon realized that residual misalignment of the FES and gyroscopes frames, positional errors in the GSC, and the small lever arm in determining spacecraft roll with the FES conspired to cause significant differences in the attitude determination from star identification and the found positions of the guide stars after performing the known tracking error nulling slew. The differences were large enough to place the stars at the edge, or sometimes outside, the commanded FES subarrays. Even if the stars were inside, the ACS would begin to ignore the pointing data from the IDS if the residuals were large. In that case, the stars would eventually drift out of the subimages and the IDS ceased to generate measured quaternions.

Figure 5 illustrates the pointing error (i.e., the RMS summed translational errors in X and Y) and error about the roll axis as a function of the number of identified stars contributing to the attitude solution. While the pointing error is much smaller compared to the individual star measurements (Fig. 4), the average increases as the number of stars available decreases. The error is still small enough that the target can be placed in the spectrograph aperture. The roll error is more problematical. While roll errors are relatively unimportant after the target is centered in an aperture, they are transformed into additional pointing inaccuracies during slews. This was the cause of the failed unknown to known transitions.

A change was made to the interface between the IDS and ACS software to eliminate this problem. During the unknown to known slew, the ACS software is now reinitialized so that the first measured quaternion from the known guide stars causes a small correction maneuver that recenters the stars into the subimages. Figure 6 illustrates the response of the pointing to this slew. In this case, a $500''$ roll difference caused a $5''$ error in pointing. The realignment slew takes about 10 seconds to execute and another 40 seconds for the error in Y to damp out. The jitter afterwards is what is typically seen with FUSE.

5. SUMMARY

FUSE target acquisitions and fine pointing are performing well. The toolbox approach has allowed most in-orbit operational difficulties to be overcome through modification of values in the IDS control tables. Other problems have been solved by changes in procedures. Only two significant items were unforeseen based on prelaunch expectations:

- Due to stray light in the telescope baffles, acquisitions need to be scheduled near orbit night. Tracking on faint stars may be temporarily lost near orbit noon.
- Due to the small FES FOV and positional errors in the HST GSC, spacecraft roll is not well-determined whenever the set of stars in a field used for attitude determination is changed. Transitioning from unknown to known tracking requires a small maneuver.

A few design changes have been identified, but are not required for the mission:




Figure 5. The accuracy of determining pointing and roll is a function of the number of stars used in the attitude solution. Spacecraft roll is much less accurate due to the small FOV of the FES and errors inherent in the HST GSC. While roll errors are relatively unimportant after the target is centered in an aperture, the consequences when transitioning from unknown to known tracking caused failed acquisitions early in the FUSE mission (see text).




Figure 6. Differences in the attitude estimate from the star identification process and the smaller population of guide stars requires an instantaneous slew to be performed when the guide star centroids are first obtained. The pointing stabilizes in less than a minute to the nominal $0.35''$ jitter.

- Positional errors of stars are unrelated to star brightness. Measured variances would provide better weighting in the quaternion estimator for stars with poor positions.
- Some of the inefficiency in scheduling target acquisitions due to stray light considerations could be reduced by performing them in the SAA. This would require a cosmic-ray rejection technique to be developed for the image processing task.

The innovative techniques used for target acquisition and fine guidance control have allowed FUSE to operate in a semi-autonomous manner as desired.

ACKNOWLEDGMENTS

We thank the developers of the FUSE attitude and guidance systems for a well-designed target acquisition and pointing control system, particularly the Orbital Sciences Corporation and the Johns Hopkins University Applied Physics Laboratory. FUSE is a joint NASA-CNES-CSA mission operated by The Johns Hopkins University. Financial support to U.S. participants has been provided by NASA contract NAS5-32985.

REFERENCES

1. D. J. Sahnou, H. W. Moos, S. D. Friedman, W. P. Blair, S. J. Conard, J. W. Kruk, E. M. Murphy, W. R. Oegerle, and T. B. Ake, "The far ultraviolet spectroscopic explorer: 1 year in orbit," *Proc. SPIE* **4139**, 2000.
2. J. W. Kruk, P. Chayer, J. Hutchings, C. Morbey, and R. Murowinski, "Fuse fine error sensor optical performance," *Proc. SPIE* **4139**, 2000.
3. B. K. Heggstad and R. C. Moore, "The far ultraviolet spectroscopic explorer (fuse) instrument data system," *Proc. 18th Digital Avionics Systems Conference*, 1999.
4. D. A. Artis, L. J. Frank, and B. K. Heggstad, "Scripted operations in the far ultraviolet spectroscopic explorer flight software," *51st Congress of the International Astronautical Federation*, 2000.
5. B. M. Lasker, C. R. Sturch, B. J. McLean, J. L. Russell, H. Jenkner, and M. M. Shara, "The guide star catalog i - astronomical foundations and image processing," *Astronomical Journal* **99**, pp. 2019–2058, 1990.
6. S. J. Conard, R. H. Barkhouser, J. P. Evans, S. D. Friedman, J. W. Kruk, H. W. Moos, R. G. Ohl, and D. J. Sahnou *Proc. SPIE* **4139**, 2000.
7. M. D. Shuster and S. D. Oh, "Three-axis attitude determination from vector observations," *Journal of Guidance and Control* **4**, pp. 70–77, 1981.