The Evolution of the FUSE Spike Long Range Planning System

H.M. Calvani, A.F. Berman, W.P. Blair, J.R. Caplinger[a], M.N. England[a], B.A. Roberts

Department of Physics and Astronomy
Johns Hopkins University
3400 N Charles St. Baltimore, MD 21218 USA
{calvani | aberman | wpb | caplingr | england | broberts}@pha.jhu.edu

R. Hawkins, N. Ferdous, T. Krueger

Space Telescope Science Institute
3700 San Martin Drive, Baltimore, MD 21218 USA
{rhawkins | ferdous | krueger}@stsci.edu

Abstract

The *Far Ultraviolet Spectroscopic Explorer* (FUSE) satellite was launched in 1999 to perform high resolution spectroscopy of astronomical sources in the 905-1187 Å spectral region. The Long Range Planning (LRP) of all the science, calibration and engineering activities for the FUSE mission is performed using a FUSE-specific Spike scheduling software package developed at the Space Telescope Science Institute (STScI). In this paper we present a description and evolution of the FUSE Spike LRP system given the pre-launch mission assumptions, on-orbit realities, and the operational challenges encountered after mechanical failure, and subsequent modification of the attitude control system in November-December 2001. Despite the operational challenges faced throughout the mission, the FUSE Spike planning software has successfully adjusted to the dynamic set of operational constraints and has maintained the predicted pre-launch average science efficiency (~30%).

1 Introduction

The *Far Ultraviolet Spectroscopic Explorer* (FUSE) was launched into a low Earth orbit on June, 24, 1999, and was designed to perform high resolution far ultraviolet spectroscopy of a wide range of astronomical sources over a three-year prime mission (Moos et al., 2000). The FUSE science instrument consists of four co-aligned telescopes and Rowland spectrographs equipped with twin microchannel plate detectors (Sahnow et al., 2000).

FUSE is in a 756 km circular orbit with an inclination of 25° to the equator. Primary contact with the satellite is accomplished through a ground station antenna at the University of Puerto Rico, Mayaguez, which provides 6-7 daily contacts of approximately 12 minutes duration each. The mission is operated from the Johns Hopkins University (JHU) Homewood campus in Baltimore, MD. At JHU, the Mission Planning (MP) team is responsible for the planning of all the science, calibration and engineering activities onboard the observatory.

FUSE science observations begin in the form of proposals that are submitted to NASA and are reviewed in a yearly (Cycle) peer review process (Blair et al., 2002). Investigators who are awarded FUSE time then submit detailed target and observation information to JHU, where MP performs validity checks and ingests the information into the mission planning database. This information is then processed into input target files used by MP to schedule the observations at the Long Range Planning (LRP) level. Approximately 600-700 observations are ingested into the LRP system per year (Cycle), amounting to roughly 9 million seconds of observing time.

1.1 The Spike Scheduling Tool

To find suitable LRP times to schedule observations on the FUSE satellite, MP utilizes a FUSE-specific version of the Spike scheduling software developed at the Space Telescope Science Institute (STScI). Spike is a general framework for planning and scheduling, originally developed for the Hubble Space Telescope (HST) by STScI. The initial implementation of Spike had an underlying constraint propagation mechanism combined with multiple methods of search, including procedural, rule-based and Neural Network (NN) based approaches. Evaluations showed that the NN approach was clearly the best at solving large, complex scheduling problems (Adorf and Johnston 1990). As a result, Spike evolved to a Constraint Satisfaction Problem (CSP) model which

[a]Also with Computer Sciences Corporation

embodied heuristics based on the NN using a mechanism that was far more flexible and adaptable to new missions.

Throughout its evolution, Spike has used a powerful yet efficient method, *suitability functions*, to represent the wide variety of strict and preferential constraints encountered in real scheduling problems. The suitability function (is a function of time whose value) represents how desirable it is to start an activity at a given time. Suitability functions are derived from constraints, an arbitrary number which may be associated with each observation, and preference functions which indicate the degree of desirability of a particular temporal assignment. The total suitability function of an observation is the product of the suitability functions derived from its constraints (or preferences).

The CSP toolkit provides an object-oriented application-independent mechanism for implementing new telescope schedulers. Core reusable components include astronomical pointing/calculation utilities, constraint propagation mechanism and constraints. For example, adding a new constraint type is as straightforward as subclassing an existing constraint type, providing basic set–up and calculation methods, and adding the new constraint to the list of constraints in a scheduler. The core propagation mechanism takes care of the details.

The FUSE implementation of Spike uses the CSP problem solving paradigm, with variables representing observations and values corresponding to the potential scheduling times. These values are mapped to discrete time steps (user defined "quanta"), which are in turn associated with week long scheduling bins. The FUSE Long Range scheduling algorithms produce schedules by iteratively selecting observations and assigning a time quanta (and thus scheduling bin) to them. Week-long short term Mission Planning Schedules (MPSs) are produced by taking the pool of observations assigned to that week's bin and refining the assigned time and the orbital layout details to the fine grain level required to produced command uploads to the FUSE satellite.

2 FUSE Mission Planning Constraints

Before FUSE was launched, Spike was designed to compute the following orbital and spacecraft constraints in order to determine a target's visibility and hence the Spike suitabilities over the Long Range Plan:

- Beta (β) Angle Restrictions: This is the angle from the anti-sun direction to the telescope boresight. The pre-launch viewing restrictions forced FUSE to perform observations within a beta angle of $15^{\circ} < \beta < 105^{\circ}$. However, due to telescope coalignment considerations (see below) discovered after launch, most observations are currently performed in the $30^{\circ} < \beta < 95^{\circ}$ region.
- Ram Angle Avoidance: The Ram vector points in the direction of instantaneous spacecraft motion on orbit. FUSE must stay more than 20° from the Ram direction at all times to prevent mirror damage which could result from collisions of residual atmospheric particles. Currently the Ram constraint has been lowered to 10°.
- Moon Avoidance: No observations are performed when the target is within 10° of the earth's moon.

The above constraints are calculated in the Spike LRP scheduling software as *absolute constraints*. In addition, there are also *relative constraints,* user-specified special scheduling requirements that can be requested for science observations. Examples of relative constraints are observations in a monitoring over time mode (i.e. ordered and offset by a specified time or grouped within a specified time interval), phased observations of time variable sources, observations requested at a particular spacecraft orientation, and those that are to be observed at a specific requested time (i.e. to support a coordinated observation or a target of opportunity).

The goal of the initial LRP system was to use a long range scheduling function that would compute all the aforementioned constraints for a FUSE target pool and allocate potential observations to weekly-sized bins over a user specified scheduling period (for instance 1 year).

3 Initial FUSE Spike LRP Scheduling Algorithms

Given that the Spike prototype for FUSE was based on the original Spike application developed for the HST, the software incorporated a variety of built-in algorithms for automated scheduling. These included some basic search algorithms such as "*Early Greedy*", which chooses the observation that can be scheduled at the earliest time (breaking ties with other factors such as priorities and preferences) and some based on the "Repair" concept such as:

- *Earliest Least Minimum Conflicts*
- *Earliest Minimum Conflicts Deterministic*
- *Maximum Preference*
- *High Priority Maximum Preference*

An example of the repair-based algorithms, the "*Maximum*

Preference" algorithm, is based on the concept of "multi-start stochastic repair" (Johnston and Miller 1994). When constructing an automated schedule with this algorithm, a trial schedule is first made which maximizes suitabilities based on heuristics. Repair techniques are applied to eliminate conflicts and violations if possible, and, if necessary, observations are re-assigned to other scheduling times in order to produce the minimum number of conflicts. Any remaining conflicts in the schedule are resolved by removing observations or relaxing constraints until a feasible schedule remains. Interactive usage through the Spike LRP Graphical User Interface (GUI) allows inspection and modification of the automated schedule as needed.

3.1 Post Launch Scheduling Challenges

The Spike auto scheduling algorithms were conceived prior to launch assuming that FUSE would be relatively free to slew about the sky and change hemispheres (i.e. cross the orbit plane) whenever desired. Constraining observations to a restricted beta angle for extended time periods was considered a possibility, but it was not an assumed observational constraint. However, during the FUSE In-Orbit-Checkout (IOC) phase (July 1999 - December 1999) it was discovered that the four separate optical channels were experiencing large misalignments due to thermal variations which were driven by beta, pole angle (angle between the current pointing and the pole of the orbit plane), and hemisphere (north or south) changes.

It was observed that targets within $50^\circ < \beta < 90^\circ$ generally do not require channel alignment checks/corrections provided that the pole angle is within 20° degrees of the previous channel alignment activity. Observations with $\beta < 30^\circ$ were prohibited, and frequent alignment corrections were required for $30^\circ < \beta < 50^\circ$ and for $95^\circ > \beta > 90^\circ$. However, in general, the thermally induced motions were characterized and determined to be repeatable, permitting predictive motions to be used in order to maintain approximate channel alignment (Blair et al., 2002). But the discovery of channel misalignments following orbital plane crossing slews added a strong preference for limiting the target pool to one hemisphere or the other.

Since the initial LRP auto scheduling algorithms were not designed with these scheduling constraints in mind, these FUSE constraints rendered the basic auto schedulers unusable. Consequently, during Cycle 1 of FUSE operations (1 December 1999 - 30 November 2000), the MP team had to manually select and schedule targets in the allowed beta and pole angle regions, and organize observations in terms of hemisphere campaigns in order to minimize the impact of channel misalignments. The manual LRP process consisted of using software developed by MP to independently calculate the visibility and time urgency of FUSE observations. In this manner, observations were selected for the generation of MPSs. But manually building a long term schedule proved to be a tedious and time-consuming process which required a lot of bookkeeping in order to track all the observations which requested special scheduling requirements throughout the observing cycle. Once new scheduling requirements for Spike were finalized (in Fall 2000), Spike software developers at STScI began to work towards the development of new auto scheduling algorithms that would support the new operational demands of the FUSE mission.

4 Improvements to the FUSE Spike LRP

The original Spike algorithms were fairly flexible, fast, and had the ability to produce robust schedules. The main drawback of these algorithms was that the mechanism they used for making choices applied the choice strategies sequentially, until a choice was made, and thus did not allow MP to give a combined weight to multiple scheduling criteria (i.e. beta and pole angle change restrictions).

To face the new scheduling challenges learned during IOC, the FUSE Spike LRP software was modified to include two new automated scheduling algorithms: The Criteria Scheduler and the Campaign Scheduler. These algorithms were fully integrated into the FUSE Spike LRP in March 2001 and successfully satisfied the new operational constraints.

4.1 The Criteria Scheduler

The Criteria Scheduler uses scoring functions that allow MP to dynamically control the role of different criteria in the scheduling process. The algorithm works by combining the score for each criterion in a multiplicative fashion. Each criterion is implemented so that it returns a score in the range [0,1]. Each criterion also has an associated weight assigned to it, which is then taken into account when the scores are combined. The overall score for each observation is calculated as (1 - weight × (1-score)).

Hence, to support the new FUSE scheduling constraints, the Criteria Scheduler was configured with a beta angle criterion that would preferentially score different beta angle regions in the sky (i.e. a score of 0.5 for $30^\circ < \beta < 50^\circ$ and for $50^\circ < \beta < 90^\circ$). Minimum beta and pole angle change criteria that would also preferentially assign scores to minimize changes in beta and pole angles were also included, thus decreasing the need for channel alignment activities by clustering targets in beta and pole space.

Spike then calculates the net multiplicative score for each observation as a function of time for each time (quanta) within a schedulable window, and schedules the observation at the time where the maximum score is obtained. In the case where tied scores are obtained, Spike would schedule the observation at the earliest of the tied times.

4.2 The Campaign Scheduler

The Campaign Scheduler was developed to minimize the impact of channel misalignments resulting from hemisphere-crossing slews. The algorithm works in conjunction with the Criteria Scheduler to schedule observations in the form of hemisphere campaigns of a user specified length.

The Campaign Scheduler is a three-phase scheduling algorithm. First, Spike generates a constrained observation schedule by selecting all the observations which are constrained by either limited visibility, or special scheduling requirements, and locks them in the scheduling timeline. During the second (repair) phase, Spike makes modifications to the campaings created in phase 1 to make them fit, as close as possible, the user specified campaign requirements. These changes include modifying (or removing) extremely short campaings or extending the campaign lengths to fill gaps etc. Lastly, the algorithm fills in the rest of the schedule using the Criteria Scheduler, clustering targets in beta and pole space within the campaign. For the second phase of the process, a north-south criterion is also incorporated with a score that depends on whether or not the time being evaluated for the observation is in the same hemisphere as in the campaign.

5 FUSE LRP Performance

Experience during the Cycle 1 (manual) LRP scheduling process showed that hemisphere campaigns had an average duration of 21 ± 7 days. The Campaign Scheduler allows a large tuning flexibility; but after thoroughly exploring the Campaign and criteria parameters through LRP simulations, it was determined that the nominal campaign length arrived at in Cycle 1 in fact worked quite well with the automated one year generated LRP.

Given the combination of competing absolute and relative constraints, occasionally some observations disrupt the nominal campaign lengths, and conflicts are created. In general, this only occurs when a target with very limited visibility and/or a specific timing constraint is present, and a sufficiently large target pool is not available in that particular part of the sky to support a hemisphere campaign. When such conflicts do arise, the dynamic and interactive nature of the Spike LRP provides flexibility to modify the Criteria Scheduler or manually re-assign (if possible) observations to other scheduling windows. Hence MP can have real control over the choices Spike makes during the generation of the FUSE LRP. Changing hemispheres for a single observation or relatively short number of observations is not disallowed, but these reduce the observing efficiency and disrupt the scheduling process.

The first year of FUSE science operations (1 December 1999 - 30 November 2000) accomplished a respectable average science efficiency of 27.2% - the actual on-target science exposure time as a percentage of the wall clock time. This is slightly higher than the predicted 26% despite the additional operational restrictions discovered during IOC. But without the use of the Spike LRP it proved to be an inefficient, tedious process which required a heroic effort on the part of MP. The development and incorporation of the Campaign Scheduler into the Spike LRP simplified tremendously the FUSE scheduling process and provided a robust planning solution to the new scheduling constraints. Beginning on March 2001, the recently modified LRP system was generating a modest 28% science efficiency three months into the beginning of Cycle 2 operations (1 December 2000 - 30 November 2001).

6 New Operational Challenges: Reaction Wheel Failure

In November-December of 2001, FUSE lost two of the four reaction wheels in the spacecraft, presenting an enormous challenge to FUSE mission operations.

FUSE is three-axis stabilized satellite that requires sub-arcsecond pointing stability. The satellite uses reaction wheels to control slew motions between targets and to maintain fine pointing control during observations. FUSE has four reaction wheels, three that are used for each of the body axes (yaw, pitch, roll) and a skew wheel that is used mainly for redundancy. A minimum of three wheels are needed to successfully operate the attitude control system (ACS) and achieve the required pointing stability and accuracy.

In November and December of 2001 two reaction wheels suffered permanent mechanical failures, leaving the spacecraft stable in only two axes and halting science operations. However, within seven weeks, engineers developed and installed new flight software to control the satellite in all three axis, using a hybrid of the two remaining reaction wheels and magnetic torque bars (MTBs) acting against the geomagnetic field to compensate for the third axis.

It was determined that although the two-wheel MTB scheme can in some cases provide nearly the same pointing accuracy and stability as before the wheel failures, magnetic torque is not strong enough to fully cancel external disturbance torques at all spacecraft attitudes. Moreover, the direction and magnitude of the geomagnetic field and gravity gradient vectors vary on orbital and daily timescales, hence adding a high level of scheduling complexity to mission operations.

The implementation of the two-wheel MTB scheme added a new absolute constraint to FUSE MP operations, *torque authority* (TA) - defined as the margin of control provided by the MTBs beyond what is needed to counteract gravity gradient induced torques. To direct the science planning process, a ground-based predictive model was develop such that observations would be scheduled only at stable attitudes (attitudes with sufficient TA). The flexibility in the Spike LRP software allowed for rapid incorporation of this new constraint into the scheduling system. This new constraint is conveyed to Spike as series of absolute constraint time windows.

6.1 Post-Reaction Wheel Failure LRP Performance

Initial predictions in late January 2002 indicated that FUSE would have stable observing regions around each of the orbit poles (CVZs) which, with orbital precession included, would allow access to ~45% of the sky for science observations. Empirical testing and modeling increased the effective sky availability to greater than 75% by late-July 2002. Relaxing the Ram constraint to 10° and improving the MTB management scheme by the beginning of Cycle 4 (1 April 2003 - 31 March 2004), provided good long term sky visibility with an effective sky coverage of ~ 95% (Figure 1).

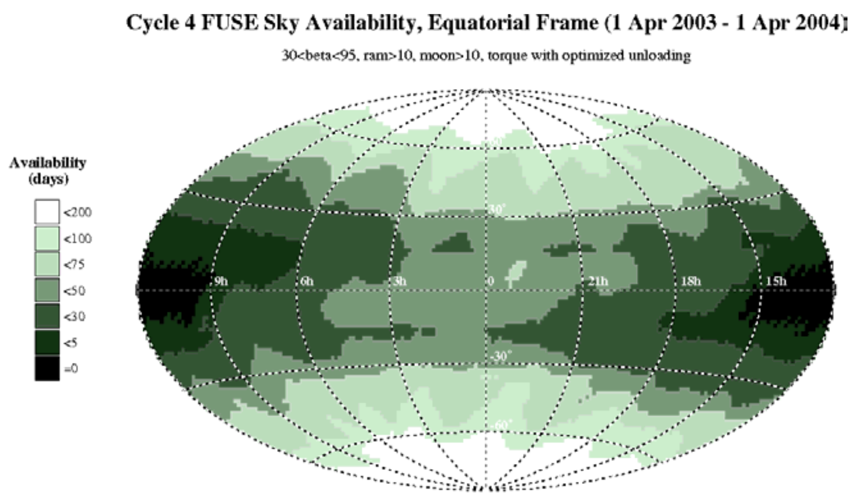

Figure 1. Cumulative sky availability during Cycle 4. The scale at left shows the integrated days of visibility. Targets located in light regions of the plot are nearly unconstrained in terms of visibility and scheduling, while targets in the dark green regions require careful scheduling.

As shown in Figure 1, during a one year time frame only a small part of the sky (black region) is not observable by FUSE. But the stability on multi-orbit timescales is essential. As a general rule, targets near the CVZ achieve good multi-orbit TA, but targets near the equator have more limited visibility (Figure 2).

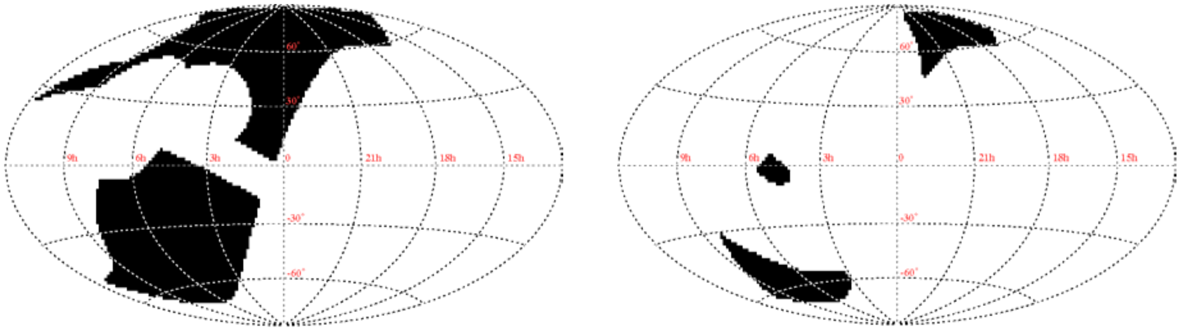

Figure 2. Visibility (dark region) for 3 January 2003. Left panel shows regions of stability for 1000 seconds; these regions shrink (right panel) if a day-long observation is desired.

In terms of the schedulability of the Spike LRP, a large fraction of the sky ($|\delta| \geq 40^{\circ}$), satisfies Ram, Beta and Moon absolute constraints and is also in or near the CVZ at some time during the year. But, for targets with $|\delta| \leq 40^{\circ}$ the visibility windows are significantly reduced due the TA constraint (Figure 3).

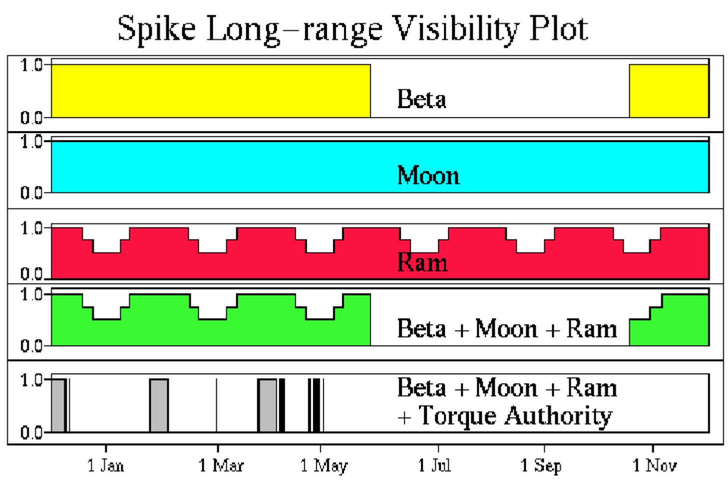

Figure 3. Spike target constraint plot (preference vs. time) for a hypothetical target near $\delta \sim 40^{\circ}$ over the course of one year. Under pre-failure conditions (Beta, Ram and Moon), the target is available for more than half the year (4th line down). The target visibility drops severely when the TA absolute constraint is added.

Targets which have a total visibility in a year ≤ 3 weeks are defined as constrained either by visibility and/or relative constraints (however not all targets with $|\delta| \leq 40°$ have less than three weeks of visibility). Note that the target pool in Cycles 1-3 operations had not been selected with knowledge of the TA constraint. As a result, the amount of challenging targets to schedule in the LRP dramatically increased (by ~ 46%) as a result of the reaction wheel anomaly.

In order to improve the scheduling efficiency, efforts were made to increase the pool of targets at higher declinations with survey (Observatory) programs, which by definition did not contain any special relative constraints. These Survey programs alleviated the LRP scheduling process by increasing the number of observations that could be used as "fillers" between the observation of targets at lower declinations. Recommendations were also made for the Cycle 4-5 proposal selection process to limit the number of constrained observations that would be accepted (no more than 35-40 observations per Cycle).

The limited number (and size) of visibility constrained targets (mainly un-observed observations from Cycle 1-3) rendered many targets as unschedulable in the LRP, since the requested observing time did not fit in the individual visibility windows calculated by Spike. This problem was solved by splitting long observations in order to fit shorter visibility windows. This technique worked very well and significantly reduced the number of affected observations. Another side effect of the larger number of constrained targets in the schedule was the frequency increase hemisphere campaigns. Tuning the campaign length in the Campaign Scheduler algorithm to 14 ± 7 days satisfactorily fit the new constraints optimizing overall stability and the observing efficiency (Figure 4) in the Spike LRP.

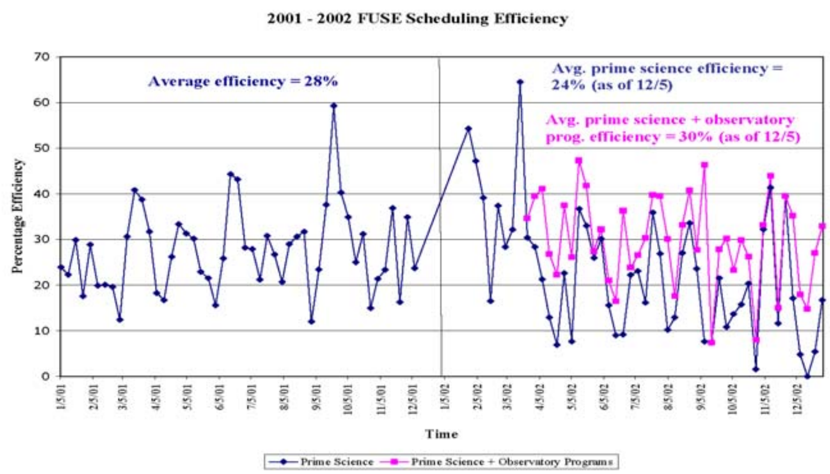

Figure 4. Scheduling efficiency for 2001-2002. The vertical black line indicates the time of the second reaction wheel failure. Note that even without the addition of Observatory programs in early 2002, the observing efficiency did not significantly changed as a result of the reaction wheel failures.

7 Conclusion

The FUSE Spike LRP system has successfully adapted to the dynamic set of operational constraints of the FUSE mission. From the channel misalignment problems encountered during IOC to the addition of the TA absolute constraint, the Spike CSP toolkit model's domain independent, object oriented approach for representing constraint types allowed the rapid incorporation of new constraints into the MP scheduling process. Spike has provided a robust solution for the long-term planning of FUSE observations and despite all the scheduling challenges faced, it has steadily maintained the pre-launch observing efficiency.

Acknowledgments

We wish to thank the FUSE Operations and Science teams and the Spike development staff at the STScI for assistance in the development efforts described in this paper. The FUSE mission and this work are both supported by NASA Contract NAS5-32985 to the Johns Hopkins University.

References

Adorf, M. H., Johnston, D. M., "A discrete stochastic neural network algorithm for constraint satisfaction problems." *Proceedings of the International Joint Conference on Neural Networks*. 1990.

Blair, P. W., Kruk, W. J., Moos, W. H., Oegerle, R. W., "Operations with the FUSE Observatory," *Proc. SPIE*, 4854, 2002.

Johnston, D. M., Miller, G., "Spike: Intelligent Scheduling of Hubble Space Telescope Observations." M. Zweben and M. Fox., eds. *Intelligent Scheduling*. Morgan-Kaufmanm, pp 391-422; 1994.

Moos W. H., et. al., "Overview of the Far Ultraviolet Spectroscopic Explorer Mission," *Astrophysical Journal*, 538, L1-L6, 2000.

Sahnow, D., et al., "On-orbit Performance of the Far Ultraviolet Spectroscopic Explorer," *Astrophysical Journal*, 538, L7-L12, 2000.