EUVE Telemetry Processing and Filtering for Autonomous Satellite Instrument Monitoring

M. Eckert, C. Smith, F. Kronberg, F. Girouard, A. Hopkins,
L. Wong, P. Ringrose, B. Stroozas, and R. F. Malina
Center for EUV Astrophysics, 2150 Kittredge Street,
University of California, Berkeley, CA 94720-5030, USA

ABSTRACT

A strategy for addressing the complexity of problem identification and notification by autonomous telemetry monitoring software is discussed. The Extreme Ultraviolet Explorer (EUVE) satellite's science operations center (ESOC) is completing a transition to autonomous operations. Originally staffed by two people, twenty-four hours every day, the ESOC is nearing the end of a phased transition to unstaffed monitoring of the science payload health. To develop criteria for the implementation of autonomous operations we first identified and analyzed potential risk areas. These risk areas were then considered in light of a fully staffed operations model, and in several reduced staffing models. By understanding the accepted risk in the nominal, fully staffed model, we could define what criteria to use in comparing the effectiveness of reduced staff models. The state of the scientific instrument package for EUVE is evaluated by a rule-based telemetry processing software package. In the fully automated implementation, anomalous states are characterized in three tiers: critical to immediate instrument health and safety, non-critical to immediate instrument health and safety, and affecting science data only. Each state requires specific action on the part of the engineering staff, and the response time is determined by the tier. The strategy for implementing this prioritized, autonomous instrument monitoring and paging system is presented. We have experienced a variety of problems in our implementation of this strategy, many of which we have overcome. Problems addressed include: dealing with data dropouts, determining if instrument knowledge is current, reducing the number of times personnel are paged for a single problem, prohibiting redundant notification of known problems, delaying notification of problems for instrument states that do not jeopardize the immediate health of the instrument, assuring a response to problems in a timely manner by engineering staff, and communicating problems and response status among responsible personnel.

Keywords: EUVE, autonomous operations, autonomous telemetry monitoring, RTworks

1. INTRODUCTION

The Extreme Ultraviolet Explorer (EUVE) satellite is an orbiting telescope facility (Bowyer & Malina 1991). The science payload is operated from the Center for Extreme Ultraviolet Astrophysics (CEA) at the University of California, Berkeley (UCB). The Explorer Platform is operated by the Flight Operations Team (FOT), at Goddard Space Flight Center (GSFC) in Greenbelt, MD. Since launch this NASA contract has been served by Loral AeroSys, now part of Lockheed Martin Space Mission Systems. EUVE launched aboard a Delta II rocket in June, 1992, and remains functional beyond its planned 3.5 year mission. It is currently operating in an "extended mission" mode under a reduced budget but with virtually uncompromised science returns.

Prior to launch the decision was made to monitor the EUVE science instrument with staff controllers at CEA in the same way that other missions were: 24 hours a day, 7 days a week. We required two people, generally a staff controller and a student controller-aide, be available in the EUVE Science Operations Center (ESOC) at all times to monitor telemetry as it arrived.

To reduce EUVE's operating costs and to increase the mission's competitiveness, CEA decided to create an autonomous monitoring system for the payload. A discussion of the factors that went into this decision can be found in Malina (1994) and Morgan & Malina (1995).

2. TRANSITION FROM FULLY STAFFED TO AUTONOMOUS OPERATIONS

The transition from fully-staffed to autonomous operations was planned in phases. We had a long transition phase from fully staffed to single shift operations, then operated in a single shift mode. We are in the final transition phase to autonomous, or "zero-shift," operations. The transition to zero shifts is being carried out in three phases: first, eliminate all routine human monitoring of the instrument; second, eliminate scheduled weekend and holiday support; third, eliminate or automate all remaining routine console support tasks.

2.1 from THREE SHIFTS TO ONE SHIFT

Applying our experience in the fully staffed model, we summarized everything that the operators routinely did during their shifts. This produced a list of duties that had to be addressed in the light of first a single, and then a zero shift model. To begin, all routine tasks done during the night shift were moved to the day shift. Then, once the rule-based, automated monitoring software (Eworks) was developed (see section 4.1) and running in the ESOC, the operators moved out of the ESOC for the night shifts. As problems arose, the staff returned to the ESOC to investigate. This test period was an important phase in

Eworks development. Bugs were found and fixed, the operators gained confidence in the software, and, as software issues arose, new software was requested. Once Eworks and other software were fully tested, and the operations-related nighttime workload was reduced to responding to pages (via electronic pager) generated by Eworks, a review was held for groups from UCB and GSFC, and the official transition to single shift operations was approved.

2.2 from ONE SHIFT TO ZERO SHIFTs

After completing the transition to single shift operations, we began working toward an unstaffed, or zero-shift, model. The job descriptions of the controllers had changed, and they became known as Anomaly Coordinators for the ESOC (ACE). The ACEs again reassessed their operations-related duties to determine the operational changes necessary to accomplish zero-shift operations. The transition from human monitoring of the instrument during the day shift with the automated software was an easy one and required little development or re-evaluation. We reevaluated what would be required to expand the unstaffed time from night shifts to weekends. Software was written or expanded to make this transition. In the early stages of this implementation the ACE was required to log-in on weekends to verify that everything was working as expected. As software testing progressed and new software was implemented and confidence increased, the requirement to log-in on weekends was dropped.

As of this writing we have completed the transition to autonomous operations for weekends and holidays. ACEs and students still work in the ESOC on workdays. Most, but not all, of the routine duties have been eliminated. We have a target of late 1996 for the transition to continuous, autonomous operations. Some routine maintenance tasks will always require human presence, as well as a few tasks that pose very difficult problems for automation and will probably not be addressed in this mission.

3. Risk Evaluation

A risk evaluation group (REG) formed to discuss the risks of changing over to an autonomous monitoring system. Risk was divided into two broad categories: (1) not receiving telemetry, and (2) telemetry indicating a problem with the instrument. The REG identified the risk areas and considered them in the light of the fully-staffed model as well as at several different levels of reduced staffing. These results were then used to prioritize tasks at each step and to aid in focusing the work.

3.1 Receiving Telemetry

Before the REG formed, we realized a risk is always involved in being out of touch with the instrument. In the fully-staffed model we were routinely out of contact for 70 minutes of every orbit, and that was an accepted risk. Occasionally problems with ground systems force us to be out of contact for extended periods. In light of this, a contingency procedure was written to define what to do when the ESOC is, or expects to be, out of touch with the payload for an unacceptable length of time. The "20 hour rule" attempted to balance the risk of an unmonitored payload with the risks of commanding without feedback. The rule states that if we are out of contact with the payload for more than 20 hours the high voltage power to the detectors will be turned off. This rule defined the basic risk already accepted and became the basis for many decisions for the implementation of automation.

Initially, we used two timeout periods to signal that we were not receiving telemetry, one at 6 hours and another at 16. The shorter period gave ample time to work problems, and the 16 hour period warned that time was running out and we should be determining when a realtime contact was available to command the detector high voltage be turned off. Once we were operating with this system, we found that if the 6 hour timeout expired after the day shift, little support was available to solve the problem before morning. If it occurred during working hours, the problem was usually apparent long before the timer tripped. We decided to eliminate the 6 hour timer, and rely on Eworks solely for emergency notification that we had been out of touch for too long and needed to plan to turn the instrument off.

3.2 Instrument problems

All of the risks and problems that could be determined by analyzing the telemetry were organized into three tiers: concerns critical to immediate instrument health, concerns non-critical to immediate instrument health, and concerns that effected science data only. The results from this system formed the criteria for what the expert system needed to do. Rules were written for Eworks that addressed the critical issues first.

The risks and problems have been revisited as we approached each new phase of the project to ensure that we have considered their implications with respect to the next phase.

4. IMPLEMENTING AUTONOMOUS OPERATIONS

The automation project was divided into three main parts: an expert telemetry monitoring system (Eworks) (Lewis et al. 1995; Wong et al. 1996), a paging system to automatically contact the ACE (Lewis et al. 1995), and a system to monitor the ground systems (Abedini & Malina 1994).

In addition to the systems themselves, the operations staff revamped the procedures that governed how they operate. Also, the way in which they communicated within the ESOC, with the rest of CEA, and with GSFC had to change (see Kronberg et al. 1995).

4.1 EWORKS IMPLEMENTATION

From the time of the decision to move to autonomous operations for night coverage, the software engineers were given six months to design and implement a system. They selected RTworks, an AI software package from Talarian Corporation, for the basis of our expert system, Eworks. They also sought guidance from groups at NASA's Ames Research Center and the Jet Propulsion Lab (JPL) with experience in designing and managing expert systems.

The software engineers designed a system that would provide the ability to do trend analysis, as well as to monitor the items that the REG had decided were critical to payload health and safety. They also planned a graphical user interface (GUI) for the controllers to use in diagnosing problems. They took their design before an internal review board at CEA. The board made a convincing argument that trend analysis was not an integral part of assuring health and safety, and it was removed from the plan. It was also decided that the GUI to be used in problem diagnosis was not critical and would be very resource-intensive to develop. They decided to develop only a fractional part of the interface.

Discussions were held with groups from Ames and JPL to discuss how the operators' knowledge could be transferred to the programmers and then used to generate the knowledge base rules. None of the methods discussed, including digraphs, decision trees, and face-to-face conversation, were appealing. In the end, flowcharts were chosen for the simplicity and general familiarity. The flowcharts were used both to design the knowledge base and to serve as a reference for the operators on how the knowledge base behaves. Once the flowcharts were agreed upon by the operators and programmers, they were taken to REG for verification and approval.

During the period of flowchart design and review, the software engineers worked on low level code for Eworks, converting our data into a representation that RTworks could use. They encountered several difficult problems. RTworks was designed to work with a continuous data stream. Our realtime telemetry is not continuous; there are data dropouts and corrupted data are received. RTworks also had a problem with the design of our telemetry stream. All of our engineering monitors share a small number of telemetry slots and appear in those slots at different rates, varying between monitors, from every frame (1.024 seconds) to every 128 frames. We needed to know whether or not the instrument knowledge was current. RTworks passed telemetry values along to the inference engine to

be processed by the rules only when the values changed. If data were missing, the value did not change. We need to know whether all the instrument knowledge is current. In RTworks, rules are only queued when values change. No rules are fired if any monitor that it references is "unknown". In our implementation, every time a frame arrives its time is checked to see if there was a gap. If data are missing, we generate a message that sets all the monitors that were expected in the gap to "unknown"; then processing resumes. Once ways were found to address all the problems and the rules were implemented, the testing process began.

4.2 PAGING SYSTEM

Notifying the ACE of a problem detected by Eworks was critical. In order to assure this notification, we needed a paging system that would page the ACE persistently until someone responded.

The paging system consists of programs that manage problem files, and user configurable parameters that control who is paged and how often, in a tiered structure. First the on-call ACE is paged, (currently every 10 minutes for 3 hours), then, if a timeout period expires with no response, backup personnel (currently 5 people) are called. Two timeout periods provide the capability of paging different people at each level. At this writing, a problem escalating beyond the on-call ACE has only occurred once, when the primary beeper's battery was dead. Paging for a given problem stops only when someone acknowledges the page by logging onto the system and moving the problem file(s) to a different directory.

4.3 GROUND SYSTEMS MONITORING

We decided to improve the reliability of our network, rather than invest a large effort in automating the monitoring of the various components in the telemetry path. The most frequent problem resulted from the network locking up when one of the file-servers failed. The addition of a redundant array of independent disks (RAID) and some reorganization of the network itself largely alleviated this problem. A workstation isolated as much as possible from the rest of the network was also added. It monitors whether Eworks has processed data and has its own modem. If any of the ground systems required to get the telemetry to CEA fail, or some critical part of our network fails, Eworks will fail to process data. If Eworks fails to process data, the stand-alone system will notify the ACE that data have not been processed.

4.4 TESTING

A working version of the expert system was installed on the ESOC network and was run in parallel with the existing systems. During the initial stages of testing, the software engineers would come to the ESOC to see how Eworks was running. They checked the validity of reported problems, and determined how to eliminate "false" problems. Sometimes this involved meetings between the programmers and controllers to discuss and change the flowcharts. As the programmers' confidence in Eworks increased, they also began testing the paging software in the operations environment. Initially a software engineer was paged for problems, first during working hours only, and eventually 24 hours a day. When the number of pages was reduced to a small number the software was reconfigured to page the operations staff. This changeover occurred while the night shifts were staffed.

In addition to testing Eworks in normal operations, all of the data from previous anomalies were "run through" the expert system to verify that it would catch all of the problems previously experienced. The expert system was also run on data from a variety of engineering tests to verify that it would recognize the abnormal instrument configuration. The final test phase was to generate telemetry with problems designed to trigger individual rules. Once this testing was completed, a presentation was put together for a review panel from CEA and from GSFC. The transition was approved and we proceeded to single shift operations as planned.

5. PROBLEMS ENCOUNTERED

Implementation of this autonomous system brought many problems to light. Problems expected from the beginning of the planning process required that procedures be written to detail the ACE's responses in the unstaffed model. Some problems are inherent in the system design. Here we discuss a subset of the problems encountered and how we resolved or worked around them.

5.1 REDUCING THE NUMBER OF PAGES AND PROHIBITING REDUNDANT PAGES

The problem of paging frequency and redundancy was recognized as soon as Eworks went into operation and has yet to be adequately addressed. Eworks generates a problem file every time one of its rules is triggered. All of the real problems that we have experienced are continuous in nature and persist until rectified. This means that, in its current version, Eworks will continue to generate problem files, and therefore pages, each time we receive telemetry. When a major problem occurs, like the main instrument processor (CDP) resetting, 20 to 30 problems are generated every time Eworks receives telemetry. Over a

period of hours this can result in the creation of hundreds of problem files, which the ACE must check to verify that no unrelated problems have occurred. One solution to this problem would be to increase the complexity of Eworks rules, such that if the rule fired, indicating a lack of power to the CDP (and therefore the entire instrument was off), no other rules would be triggered. Another solution might be to keep track of what rules had been used to generate problems and prevent generating more from the same rule. These solutions would require major modifications to Eworks, and it is advantageous to keep Eworks stable. Currently we use a method that we call page screening to reduce the number of pages. Once a problem is diagnosed, and the ACE determines that it will be ongoing, future pages that exactly duplicate problems that we have seen can be screened. When Eworks generates a problem and page screening is on, each new problem is compared with those being screened, if it matches no new problem file or subsequent page is generated.

5.2 DELAYED NOTIFICATION OF NON-CRITICAL PROBLEMS

When we made the transition to single-shift operations, we decided that responding to minor problems could be left until the day shift. For instance, a given temperature being slightly above normal is not a threat for a period of hours, it might become a serious problem if left unresolved for days. In a zero-shift scenario there would no longer be someone in the ESOC every day. We did not want to increase the number of pages with the transition to zero shifts, and we could not ignore these problems. By design, Eworks categorizes problems into two levels, "warnings" which are logged and page an ACE and "alerts" which are logged, but do not page an ACE. A program written to go through all of Eworks' logs to look for alerts automatically runs once a day and bundles all of the alert messages into one problem file, which generates a page. This way we are notified of non-critical problems daily.

Limit transitions for which no Eworks rules were written are handled in a similar way. For example, the detector high voltage is ramped down to about half its operating voltage during orbital day, protecting the detector from high photon counts. The high voltage violates yellow limits during orbital day and the condition is normal. No Eworks rule triggers when this voltage is in the yellow range. The voltage could be neither red nor at either usual value. This situation could be a problem in the zero-shift staffing model. A new module was written and added to our normal limit-checking package. The module looks at limit transitions files and compares each transition with a filter file which contains acceptable ranges for out-of-limits values. The filter file is user-controllable and can be updated as the situation changes. Software similar to that which checks for Eworks alerts, checks daily for unexpected limit transitions and generates one page for all of the unexpected transitions found.

5.3 ASSURING TIMELY RESPONSE TO ANOMALIES

The problem of timely response has been discussed since the inception of autonomous operations. It is solved by policy and the configurable paging software. An ACE has been on-call since the moment we stopped staffing the ESOC continuously. One ACE is designated as the person on-call and the duty rotates. The primary beeper is physically passed from ACE to ACE, so the phone number to contact the ACE on duty remains constant. The main phone number for the ESOC is forwarded to a commercial message center that offers callers the option of paging whoever is on-call. In addition, each ACE carries a beeper assigned to them individually. A written policy defines the response time for the on-call person, as well as the time limit for when the paging system changes from calling only the on-call person to calling everyone. The policy also contains directions about how to acknowledge that someone is responding. This prevents multiple ACE's from heading in to the ESOC if everyone is paged and one person has already responded. The problem files that generate pages are not moved via remote login unless the ACE is certain that the problem presents no threat to the instrument. If a serious problem occurs, the problem file is left active. If something were to prevent the ACE's arrival at the ESOC, others are paged when the timeout period expires.

5.4 COMMUNICATING PROBLEMS AND RESPONSE STATUS

When the ESOC was staffed continuously, we engaged in a ritual called shift-changeover that required that two people talk to each other about the status. Current information was also kept, in various forms, in the ESOC. When we stopped continuous staffing the need to communicate with others was a problem. We had used paper logbooks during full staffing, and when we started leaving the ESOC, we switched to on-line logging and our "logbooks" became accessible outside of the ESOC. Other information was placed on-line when possible, unfortunately some information only comes to us in paper form, and remains inaccessible except in the ESOC. Part of the ACE policy defines how to hand problems over from one person to another. Problems that require a continuous presence in the ESOC are only passed from ACE to ACE in person. Less severe problems are handed over simply by carefully recording in the logbook what happened, what has been done, and what remains to be done. Each work day, an ACE is responsible for the instrument and reads the logbook to determine whether there are problems.

6. CONCLUSION

CEA has accomplished most of the transition from fully-staffed to autonomous operations for the EUVE science instrument. The transition relied on a phased approach to solve difficult problems at all stages of planning and development. The crucial decision to concentrate solely on instrument health and safety enabled us to obtain the goal in a

reasonable period of time. The existence of a software package that permitted use a lot of existing, working code was also paramount.

The operations group has many projects they would like to see implemented in the future. That the funding or time will be available to accomplish them is unlikely, but we can envision: An automated commanding system that would be set in motion by the submission of a science plan; improvements to the expert system so that it could verify that the instrument was configured correctly at all times; and a sophisticated network watchdog system.

ACKNOWLEDGMENTS

REFERENCES

Abedini, A. and Malina, R. F. "Designing an Autonomous Environment for Mission Critical Operation of the EUVE Satellite," Third International Symposium on Space M1fission Operations and Ground Data Systems 1994, Part 1, GSFC, Greenbelt, MI, 1994, 541

Bowyer, S. and Malina, R. F. "The EUVE Mission," Extreme Ultraviolet Astronomy, Pergamon Press, New York, NY, 1991, 397-408

Kronberg, F., Ringrose, P., Losik, L., Biroscak, D., and Malina, R. F. "Re-engineering the EUVE Payload Operations Information Flow Process to Support Autonomous Monitoring of Payload Telemetry," Re-engineering Telemetry, XXXI, 1995, 286-294

Lewis, M. et al. "Lessons Learned from the Introduction of Autonomous Monitoring to the EUVE Science Operations Center," 1995 Goddard Conference of Space Applications of Artificial Intelligence and Emerging Information Technologies," NASA CP-3296, GSFC, Greenbelt, MI, 1995, 229-235

Malina, R.F. "Low-Cost Operations Approaches and Innovative Technology Testbedding at the EUVE Science Operations Center," SAG #614, 45th International Astronautical Congress, IAA Symposium on Small Satellite Missions," Jerusalem, Israel, October, 1994

Morgan, T. and Malina, R. F. "Advances in Autonomous Operations for the EUVE Science Payload and Spacecraft," <u>Robotic Telescopes: Current Capabilities, Present Developments, and Future Prospects for Automated Astronomy</u>, ASP, Provo, UT, 1995

Wong, L., Kronberg, F., Hopkins, A., Machi, F., and Eastham, P. "Development and Deployment of a Rule-Based Expert System for Autonomous Satellite Monitoring," SAG #713, Proceedings of Fifth Annual Conference on Astronomical Data Analysis Software and Systems, Tucson, AZ, October, 1995

RTworks, Talarian Corporation, 444 Castro Street, Suite 140, Mountain View, CA 94041, (415)965-8050