

Applying the lessons learned from HST operations to new missions

Glenn E. Miller and Peg Stanley

Space Telescope Science Institute, Baltimore, MD

ABSTRACT

In order to use the next generation of space and ground based observatories for the greatest scientific benefit, the experiences of current missions should be carefully examined to find strategies which have worked well and also to identify areas where new paradigms are needed. With the operation of the Hubble Space Telescope, the Space Telescope Science Institute pioneered the large scale application of non-traditional operations models including observation preparation tools, integrated (queue) scheduling for increased scientific return, service observing, and multi-year long-range planning. This paper discusses the key aspects of HST operations, including concepts which worked well and those which did not. We discuss how this experience can be applied to new ground- and space-based missions.

Keywords: observatory operations, service observing, integrated scheduling, scientific efficiency, long-range planning

1. INTRODUCTION

Astronomers today are fortunate to have a new generation of space- and ground-based observatories in operation and several more nearing completion. Much of this growth in capability has been brought about by technological advances in telescopes and instrumentation. In addition to this revolution in hardware, we are seeing, perhaps more subtly, a revolution in how observatories are operated. In the era of flat or declining budgets, there is a growing realization that we must break the traditional cost paradigm and that this requires new modes of operation. However these modes must be carefully evaluated against scientific performance measures. In this paper we examine HST operations to distill principles which can be applied to the next generation of astronomical observatories. Section 2 presents the case for a comprehensive science and mission operations concept to guide development. Section 3 examines HST operations in more detail and discusses how tools and procedures have evolved. Section 4 examines applications to other observatories and section 5 gives a general discussion.

2. SCIENCE AND MISSIONS OPERATIONS CONCEPT

A science and mission operations concept describes, at a high level, how an observatory conducts its science mission. The operations concept presents a holistic view of the telescope and instruments, system software support, internal data flows, external interfaces (including the observer) and staff operating the observatory. (Although the terms “mission” and “ground system” usually refer to space observatories, we use these terms to describe ground-based observatories as well.)

Team building experts have characterized the development stages of a team as “forming, storming, norming and performing”. A mission goes through very similar stages and a lot of information has been published on the norming and performing phases. This section explores some of the facets of forming and storming for a mission and the advantages of creating a comprehensive science and mission operations concept to capture the ideas, assumptions and goals that come from these early phases.

2.1 The Investment

Investing early in a comprehensive science and mission operations concept will yield a substantial pay-off in the development and operations phases of a mission. While many of the details of the mission may not be available early, it is possible to address in a comprehensive way the ground system elements, internal data flow, external interfaces and appropriate staffing and skill mix required to operate the mission. Doing so will enable trade-offs to be made between initial development costs and routine operations costs. It is important that such a concept be updated regularly as architecture decisions are made. The

concept must also be used as a blueprint for building the observatory systems at the working level. The ESO VLT data flow diagram¹ serves as an example of how an observatory's processing is built around a unified concept.

Elements of the HST ground and flight systems were developed independently over a long time span. As individual elements were designed and delivered, it was necessary to integrate them into the operational environment including getting independent systems to talk to each other. For example, the short-term science scheduling system was developed assuming that observers would specify HST observations in the format native to the scheduling system (a typical observing program would require hundreds of pages of detailed parameters). Integrating this system to the needs of the astronomical community required the development of a user-oriented observation description system (see section 3.2), as well as substantial modification to the scheduling system. A well-defined science operations concept would have helped circumvent such problems later in the mission.

The earliest versions of the science and mission concept document can be used to capture the ideas, assumptions and goals resulting from the initial mission design discussions. It is especially important to identify the areas where mission design raises questions and issues for routine operations. As the science and mission operations concepts are being developed, strong involvement by the *science staff responsible for long term science operations* is essential. The sole purpose of any science mission is to achieve its scientific objectives, not provide an engineering model or vehicle. Trade-offs made in science and mission operations must address the long-term ramifications to the science goals, and therefore the staff responsible for sustaining operations must have a strong role as architectural tradeoffs are made. One must be very careful in the early stages of a mission to avoid the pitfall of believing that operations staff can always figure out a way to do it later. While it may be true that a way will be found to achieve the original science goals, the solution may be much less effective and more costly (e.g. manually intensive) than one designed into mission at the outset.

In the case of HST, the STScI was not created until after the telescope and major ground system components had been designed and their implementation begun. Consideration of how observers would interact with the telescope and ground system were largely unspecified and had to be retro-fitted into the existing systems. Each instrument employed different operational assumptions which increased the complexity of the HST.

Although a science and missions operations concept should address long-term goals, it should also be realized that the commissioning of an observatory is an extended process and not all capabilities will be available at launch or first light. Indeed, an operations concept should be used as a vehicle for phasing capabilities as operational experience is gained.

2.1 Contents of the Science and Mission Operations Concept

The science and mission operations concept needs to take a *process-oriented* view of the mission which place emphasis on the product and not on the individual functions. This view will help to ensure that artificial hand-offs and boundaries are avoided. Science observations start with the investigator planning observations and end with the investigator receiving and analyzing the data. Many missions today are being designed for service observing rather than letting the scientists take direct control of the observatory, so a science and mission operations staff is required between the scientists and the observatory. In principle, there are only two products and hand-offs required to successfully operate a mission: observation details from the scientist to operations and data distribution from operations back to the scientist. Minimizing hand-offs and interfaces will greatly simplify routine operations and reduce costs plus will have the added benefit of instilling a sense of ownership in the observatory staff that will serve to increase the quality of the product.

HST operations was developed around functions rather than products and significant effort has gone into correcting this and eliminating unnecessary hand-offs. For example, HST science planning was performed at the STScI while HST mission planning was performed at the Goddard Space Flight Center (GSFC). The STScI provided a weekly science mission specification to staff at GSFC responsible for defining the spacecraft and communications activities. STScI systems performed only rough modeling of essential factors such as onboard data management, solar array and power management and spacecraft guiding. This created the need to iterate between science planning and mission scheduling. Progress has been made on HST by combining science and mission planning process into one team, but the legacy of independently developed software systems to accomplish these functions lives on. Improvements have been made by increasing the knowledge about

solar array management and fixed head star tracker visibility into the science scheduling system and efforts will continue to address ways of improving this process.

As another example, HST observation implementation was originally organized by function: one group was responsible for user support, another for observation preparation, another for long-range planning. This led to a fragmentation of responsibility and many hand-offs. Operations were greatly improved by having one group responsible for the successful implementation of observations, which involved the operation of several different software systems.

Any new space-based mission will undoubtedly face trade-offs between performing functions in the on-board systems versus the ground system. The cost of on-board ("flight") software tends to be much higher than the cost of ground software. Flight software updates are also viewed as much riskier than ground system updates which are performed routinely. As a result, decisions may be made to include a function in the ground system even though it may be more efficient to include it in the on-board system. The future cost of an "efficiency tax" which is paid with every observation may not be considered on an equal basis with initial development and occasional flight software maintenance costs. Maximizing the overall observing efficiency should be considered a very high priority factor in all relevant operations design and architecture trade-offs. When viewed in isolation, a couple of extra minutes for a specific observatory function (e.g. guide star acquisition or instrument overhead time) may not seem large, but when multiplied by the number of times the function will occur and the push that will occur later to find a way to avoid the "tax" may turn the perceived up-front savings into a real loss. When making the design and architecture trade-offs between cost and on-orbit efficiency, care should be taken that attempting highest possible efficiency and lowest cost doesn't degrade the quality of the observations. While more data is usually better, more less-than-optimum data is a no-win situation. While there may be unforeseeable factors which degrade the quality of some observations, the design and architecture of the system and operations should not be one of them. The science and mission operations concepts should include a specification of overall observing efficiency goals and would do well to address an on-orbit cost formula for performing cost-benefit analyses.

2.2 Routine Use of the Science and Mission Operations Concept

The science and mission operations concept must be used routinely as the roadmap to operational systems development and organization. This means that as architectural decisions are made and principles change, the concept must be updated and the stakeholders in the concept must modify systems and processes.

Although the operations concept must take a comprehensive view of the observatory, it must be concise in order to be comprehensible by individuals and of practical use. In today's technology, this means that the operations concept must be a relatively small document (e.g. no more than a few hundred pages). If we project current trends in computer technology, an operations concept in the near future might include multimedia components such as interactive simulations of observatory constraints (e.g. pointing constraints) or process flow simulations.

An important aspect of the operations concept is to establish the criteria for measuring the observatory's performance (see section 5). *Quantitative performance metrics* are essential. Building in the capability to measure basic operational metrics such as the observing efficiency, process time, cycle time, work in progress, units requiring rework, failures etc. will enable more timely intervention and process improvements. Building the metrics into the system will ensure that the appropriate data are available and captured in a useful manner without placing burden on the operations staff.

Along with the routine acquisition of performance metrics, it is equally important to *seek and heed user feedback*. The STScI has approached user feedback in several ways including an official, formal Space Telescope Users' Committee (STUC), smaller, focused user working groups and user surveys. Each serves a different purpose and are most beneficial when the right forum addresses the right question. It is vital that such communication involve the user as a "stakeholder" and not simply as the customer.

HST's formal STUC meets twice per year and primarily focuses on the "big picture" questions and provides general "pulse of the community" feedback. They bring items of general concern to the STScI and help formulate the areas in which more in-depth examination and interaction with the community may be useful.

Focused user working groups have been beneficial when in-depth, detailed user feedback and responses are needed. Most recently, the STScI chartered an HST Parallel Science Working Group. Members were interested in optimizing the scientific return of the HST parallel science program and produced a cohesive set of HST parallel observing programs intended to populate the HST archive with large, uniform data sets.

User surveys can potentially reach more users, but can be very time-consuming to prepare well. Questions must be carefully worded to elicit meaningful (and perhaps semi-quantitative) responses on the topics of interest. There is also a selection effect in that users who are generally pleased often do not respond to a survey so responses can be skewed towards those with unresolved issues. While there will be little hesitation to tell you what you're doing wrong, it takes more effort to frame questions that allow the users to tell you what you're doing right (and you're sure to be doing something right) or to suggest new approaches. For example, a question which asks "what additional observation preparation tools would you like to see" may not be as useful as one which asks the users to prioritize a number of potential tools.

3. SCIENCE OPERATIONS

In this section, we describe HST science operations in more detail and highlight areas which are potentially useful for other observatories.

3.1 Proposal Development

HST observing is divided into two stages. Phase I is the competition for observing time, while Phase II is the specification of the detailed observing program. Two overriding considerations motivate this two-phased approach:

- Time on HST is competitively awarded and selection is based primarily on scientific merit. Phase I is focused on the scientific justification, while Phase II deals with the fundamentally different questions related to specifying and executing the observations. For example, detailed target acquisition strategies, dithering patterns or exact exposure times are not relevant to selection.
- Competition for HST time is intense, with an oversubscription ratio (requested time/available time) of six or more. Since the high oversubscription ratio means that most proposers are not awarded time, the Phase I process was carefully constructed to place the minimum overhead on the observer. The proposal asks only for information which is actually used in selection and the mechanics of the Phase I process is simple. The proposal is prepared in LaTeX format, a medium which is known to most astronomers. Embedded keywords are read by a computer program to populate the proposal tracking database. Postscript figures can be easily included and proposals can be electronically submitted.

HST's Phase I process is similar to that of other observatories^{2,3,4} (e.g. ESO, NOAO) and this type of commonality should be strongly encouraged in future missions, especially in view of the trend towards multi-wavelength observations. Future trends in proposal selection are discussed by Moller, et. al.⁵

3.2 Observing Program Development

The complexities of the telescope and ground system along with the constraints imposed by a low earth orbit require HST to be operated in a service mode where the observer is not in direct, real-time control of the telescope. In addition, observations from different programs are interleaved to optimize the individual science and overall observatory efficiency. The primary interface for observers using HST is RPS2⁶ (Remote Proposal Submission, second generation). The input is a text file in keyword-value format (e.g. `config: WFPC2` or `aperture: PC1-FIX`). The STScI provides general observers with a customized template file, based on the approved program.

The proposal file defines 3 kinds of entities: targets, visits and scheduling requirements on visits. A visit is a series of consecutive exposures, with overheads, on a target and may last for more than one orbital visibility period. Scheduling requirements can be imposed for scientific reasons and may affect the scheduling of a particular visit (e.g. between two dates or at a specific orientation on the sky) or may link two or more visits (e.g. maximum time separation, same orientation). This file can be edited with any text editor, or with the graphical Proposal Editor (PED)⁷. Figure 1 shows a PED screen for editing astronomical targets. This interface guides the user in entering the correct information and checks input for errors. Also included is a help system which links to the HST documentation on the WWW.

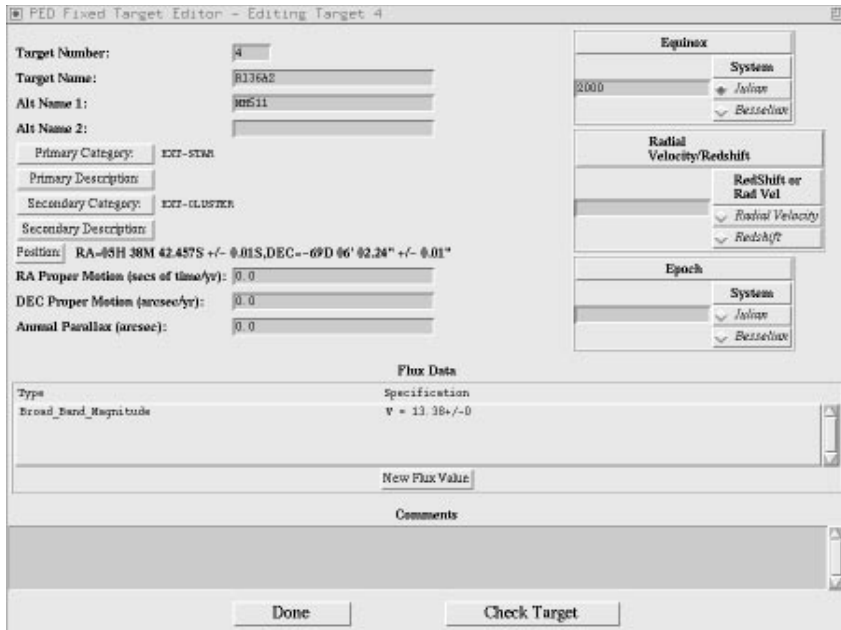


Figure 1 - The fixed target editor screen from the RPS2 Proposal Editor. This guides the observer in entering target coordinates, keywords, magnitudes, etc. and checks entries for errors. Other screens are used to enter solar system targets, visits, exposures, scheduling requirements, etc.

The RPS2 software performs several functions. It checks the proposal file for legal keywords and values, transforms the high-level observation description into the detailed parameters needed for the planning and scheduling software and calculates the basic observing constraints to determine when the observations can be scheduled during the observing cycle. Output is displayed graphically and is highly effective in showing the astronomer how the observations will be structured and their scheduling opportunities. Potential problems are highlighted by diagnostic messages and explanations. Figure 2 shows sample RPS2 output.

Before RPS2, STScI provided astronomers with software called the Remote Proposal Submission System (RPSS). RPSS checked only for some syntactical errors. The proposal was submitted to the STScI, which used other in-house software systems to check the complete syntax, feasibility and schedulability of proposals. This led to several problems being seen in most programs:

- Schedulability - interacting constraints specified by observers frequently prevented observations from being scheduled
- Opacity - The HST ground system, not the observer, divided observations into visits. Since this was done after submission, the observer had no clue how the program was structured.
- Efficiency - Because of overheads involved in executing observations (guide star acquisition, target acquisition, mechanism motion times, buffer readouts, etc.) large gaps in observing time were all too common.

Correcting these problems resulted in many iterations between STScI staff and observers. RPS2 bundled the existing in-house software systems in an interface that observers could use and provided them with the information needed to make correct and efficient observations.

RPSS was the first electronic proposal submission system (operations began in 1986!), and advances in computer technology paved the way for the evolution to RPS2 in 1994. The key technologies included graphical rather than textual displays (e.g. X-windows), ubiquitous network access (the Internet), scripting languages to “glue” different software components (e.g. Tcl), “artificial intelligence” software, and the tremendous increase in computer processing power.




Figure 2 - Detail of RPS2 output. The upper panel shows the approximate structure of the observations during an HST orbital period. (Immediately obvious in this case is 22 minutes of unused observation time.) The lower panel shows the observing constraints on the visit, including Sun and Moon avoidance, target visibility duration, a user-specified absolute orientation and the combination of all constraints.

A number of principles guided the development of observing tools:

- The system must provide astronomers with the information needed to make scientific tradeoffs - the impact of a choice needs to be apparent and shown in a meaningful way.
- Astronomers should work with familiar, astronomical concepts.
- The system should provide reasonable defaults for parameters, but allow seasoned users to specify full details.
- Observations should be built with appropriate flexibility, including flexibility of when they can execute due to changes in observing conditions.
- Astronomers should be provided with the same software tools used by the observatory staff, rather than a separate software system.
- The transformation of a proposal from a high-level, astronomically oriented observing description to detailed planning and scheduling parameters can be automated. This allows similar observations to be performed with proven methods and allows observations to be quickly modified if the observatory performance changes or if new observing methods are implemented.

Access to HST observation status is provided via the WWW (<http://www.stsci.edu/public/propinfo.html>), see Figure 3. Information provided this way is:

- **Timely** - Information about a proposal is generated on demand as each request is made. There is no built-in delay due to caching or periodic (e.g. daily) updates.
- **Accurate** - Information is generated from the operational systems so it reflects the same information available to STScI operations staff.

Understandable - Database fields and values are converted to meaningful terms and WWW links are provided to definitions and explanations.

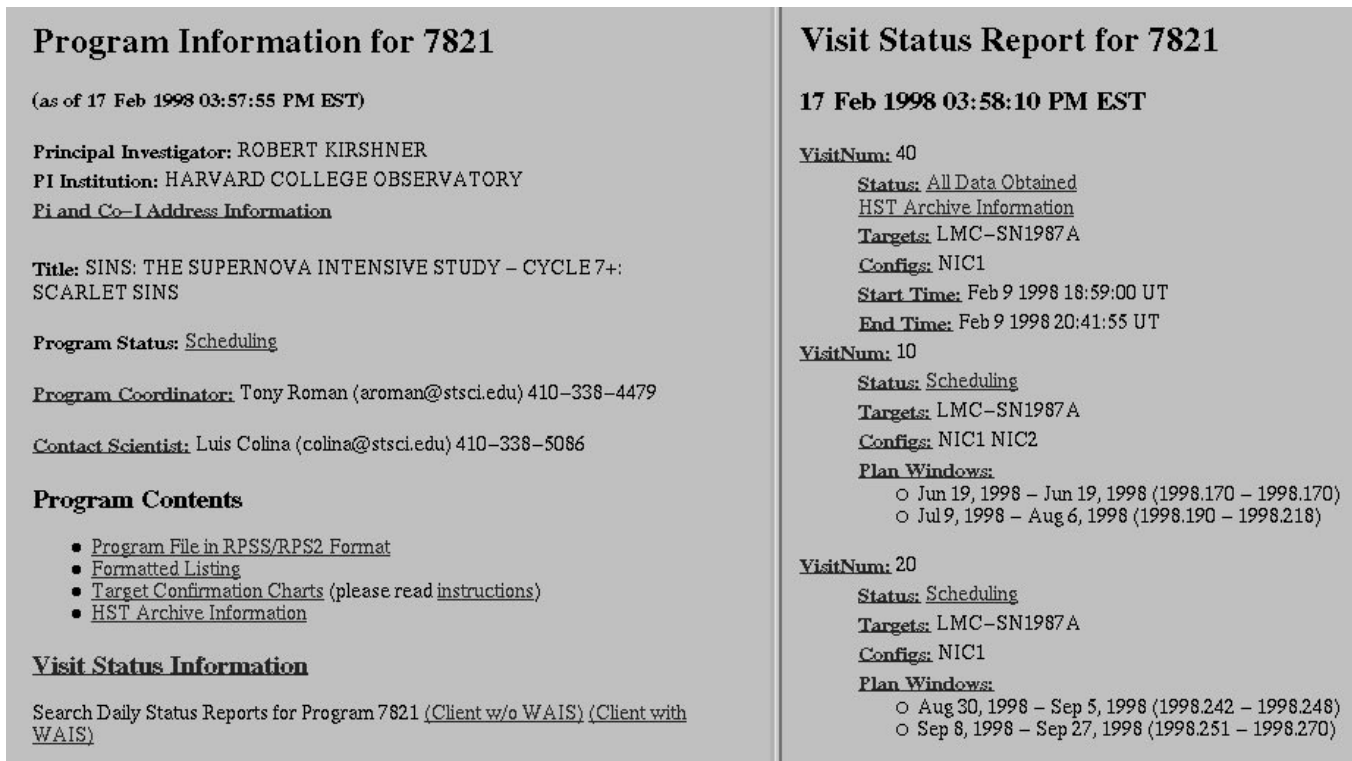


Figure 3 - HST Program Information WWW Page. The left panel shows the program summary information and allows the observer to determine the current status of the program, change their mail or email address, find STScI support staff contact information, view the current copy of the proposal, view target confirmation charts and access the Hubble Data Archive. The right panel shows the Visit Status Information link which gives the status of each visit in the program. The first visit has been executed and is available in the Archive, while the plan windows for the remaining visits are indicated.

3.3 Planning and Scheduling

The STScI uses a two-tier hierarchical approach to HST scheduling:

- Each week a detailed (~1 sec resolution) short-term schedule is prepared and binary command loads for the HST computers are generated.
- The long-range plan covers a several year interval. Visits on the long-range plan are assigned plan windows⁸, which are a subset of the observing constraint windows. Typically they are 8 weeks long. The size of a plan window is a trade-off between competing factors: Larger windows allow the short-term scheduler to achieve higher efficiency through a large pool of observations and provide freedom to move observations by a few weeks as needed for flexibility. Plan windows must be small enough to provide observers with a meaningful estimate of when the observation will be executed and also allow STScI staff to schedule their observation implementation work.

The long-range plan has been an invaluable tool for the HST mission. It provides observers with an estimate of when their observations will be executed and therefore allows the observer to perform an independent check on the schedule and also to plan their work. Similarly it allows STScI staff to plan work in support of HST observers. Long-range planning has also played a critical role in supporting overall observatory operations. It has allowed the STScI to plan for the introduction of new instruments during servicing missions, and assess the impact and recovery from changes to telescope and instrument performance. For HST, the long-range plan allows a balance of observing time between different instruments (e.g. currently 40-50% NICMOS due to the limited lifetime, 25% STIS to ensure use of this new instrument). The analog for other observatories is the allocation of time to national partners. Rather than allocating blocks of time per nation, an integrated long range plan allows programmatic balancing, optimization of individual science observations and overall observatory efficiency.

For many reasons an observatory's schedule is not static - observers make changes in programs in response to scientific discoveries; instrument characteristics can change with time; improved ways of operating and calibrating instruments are found, etc. Thus the execution of the observations must be monitored and plans must be updated to accommodate changes. The HST long-range plan is updated daily.

4. Applications to New Missions

The thesis of this paper is that the next generation of space and ground based missions must learn from the experiences of current missions in order to provide higher quality scientific observations without ever increasing budgets. Traditionally this has been done for hardware-related development, but in the areas of science observations and observer support there has been much less reuse of concepts and software between observatories. In this section we highlight some of the applications of the STScI's HST techniques to other missions. Odea, et. al.¹² discuss the HST model for user support by the Contact Scientist/Program Coordinator team.

Spike⁹ is a general planning and scheduling toolkit that was developed by the STScI for HST scheduling. Spike is the component of RPS2 that calculates visit schedulability and is also used for the HST long-range plan. Spike was designed with generality and reuse in mind, and has been used in several other observatories as shown in the table below. In addition, the German Space Operations Center has used Spike to schedule space station MIR experiments and Interval Logic Corporation is using Spike for semiconductor manufacturing. Spike is being used for prototype mission studies for SIRTf and SIM.

Observatory	Status
HST	Hubble Space Telescope. Used operationally for long-range planning and observation development.
ESO VLT	ESO's Very Large Telescope. Spike is scheduling engine used in VLT observation system. First operational use on NTT in February 1998.
FUSE	Far Ultraviolet Spectroscopic Explorer. Used for long-range planning and short-term scheduling. Launch in 1998.
Subaru	Japanese National Observatory 8m telescope. Spike is the scheduling engine.
AXAF	Advanced X-ray Astrophysics Facility. Spike is used for long-range scheduling.
EUVE	Extreme Ultraviolet Explorer. Spike used for long-range scheduling.
ASCA	Advanced Satellite for Cosmology and Astrophysics. Spike is used for long-range scheduling.
Rossi XTE	X-ray Timing Explorer. Spike is used for long-range scheduling.

Table 1 - Observatories which use the Spike planning and scheduling toolkit

OPUS¹⁰ is the pipeline data processing system currently operational at the STScI. It converts the incoming telemetry from HST into standard data products. OPUS software is based on a blackboard architecture and, as a result, is highly robust and portable. OPUS replaces the original HST pipeline system which was developed a decade earlier and included architectural components that were unmaintainable (e.g. VAX VMS dependence, brittle interprocess communications). OPUS was designed as a general distributed pipeline system which allows easy integration of customized data reduction applications.

Johnson¹¹ describes how the Far Ultraviolet Spectroscopic Explorer (FUSE) mission at Johns Hopkins University is using STScI tools to support FUSE within strict budgetary and schedule limitations. JHU and STScI are adapting the STScI's Phase I proposal processing software, Spike planning and scheduling software and OPUS pipeline data processing system to support FUSE's launch in late 1998. Being an "Explorer" class mission, FUSE demonstrates that many of the concepts and tools used for HST can apply to smaller missions as well.

5. DISCUSSION

The overall observing concept developed by the STScI for HST emphasizes scientific quality, user involvement and efficient operations. Effective use of HST required the large-scale application of several new operations strategies including:

- graphical, astronomer-oriented observation development tools
- service observing
- interleaved (queue) scheduling
- multi-year long-range planning
- electronic proposal submission

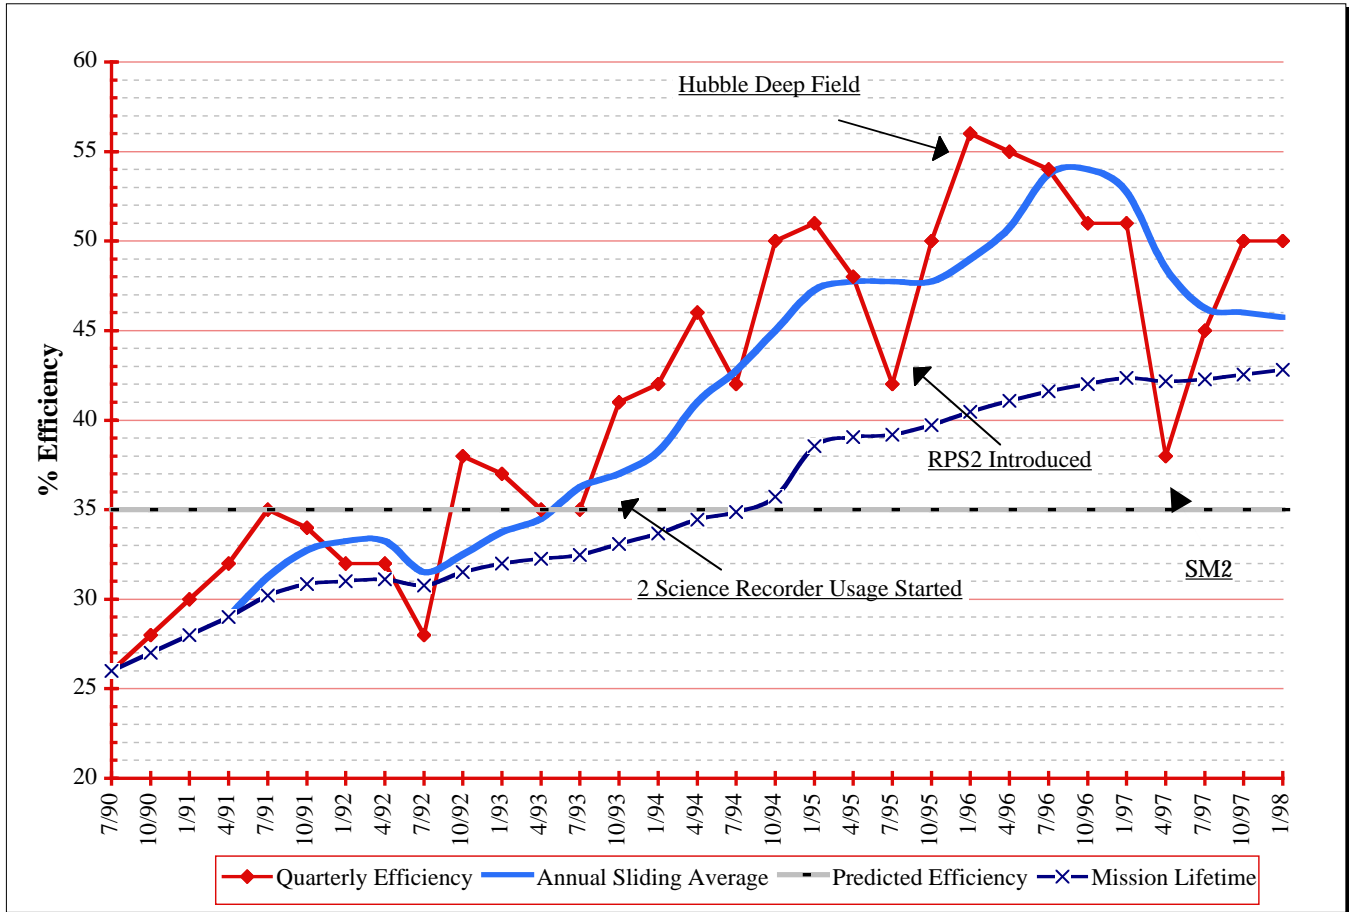


Figure 4 - HST Observing efficiency. The history of HST observing efficiency. The diamonds show the efficiency each quarter, while an annual sliding average is given by the smooth line. The global efficiency over the life of the mission is shown by the Xs. For reference, the line at 35% efficiency is a pre-launch prediction of the maximum *sustainable* average efficiency. Only the efficiency of prime observations is shown and observations made in parallel provide a further increase in science data.

The proof of these concepts is illustrated in Figure 4, which shows HST's observing efficiency as a function of time. Pre-launch estimates predicted the maximum sustainable average efficiency would be 35%. HST operations exceeded this over four years ago and has continued to make gains. The amount of on-board science data storage was doubled when two science recorder usage was started in 1993 and this removed a significant bottleneck imposed on operations. (Installation of a solid state recorder in the second servicing mission further increased the on-board storage capacity.) The introduction of RPS2 in 1995 allowed observers to optimize individual science observations while retaining scheduling flexibility. A peak in

efficiency was achieved with the Hubble Deep Field which used the HST's continuous orbital viewing zone to take extended exposures uninterrupted by earth occultation. Long-range planning continues to use the continuous viewing zone for observations whenever possible. The dip in efficiency in early 1997 is due to the inevitable down time of the second servicing mission.

The increase in HST observing efficiency was done without degrading the high quality of HST observations. One measure of the quality of HST observations is the frequency of Hubble Observation Problem Reports (HOPRs). A HOPR is filed when an observer (or STScI staff member) detects a problem in the resulting data. The HOPR rate for HST observations is low, about 5%, with about half of the failures due to causes which are beyond direct control of the user or observatory staff (e.g. binary guide star, instrument anomaly). HOPRs are continuously monitored to identify and address systemic problems in observational procedures and execution.

As discussed in section 4, many of the lessons learned from HST and other current missions are applicable to new missions, regardless of scope and whether they are space- or ground-based. Emphasis on scientific quality, user involvement and efficient operations are the linchpins for an observatory. Modes such as service observing or interleaved scheduling may be appropriate, depending on the mission. Providing users with the proper tools to make informed decisions on the scientific observations is essential. A long-range plan can be a valuable tool for the observers, observatory staff and overall observatory management.

REFERENCES

1. P. Grosbol and M. Peron, "The VLT Data Flow Concept", *Astronomical Data Analysis Software and Systems VI*, G. Hunt and H. Payne, ed., pp. 22-25, 1997.
2. H.E. Payne, D.J. Asson, "Electronic Submission of HST Phase I Observing Proposals", *Astronomical Data Analysis Software and Systems IV*, ASP Conference Series, Vol. 77, R.A. Shaw, H.E. Payne, and J.J.E. Hayes ed., pp. 65-67, 1995.
3. B. Blacker, G. Miller, M. Urry, H. Payne, and D. Asson, "Electronic Submission of HST Phase I Proposals", this volume.
4. A. Chavan and M. Albrecht, "An Information System for Proposal Submission and Handling", *Astronomical Data Analysis Software and Systems IV*, ASP Conference Series, Vol. 77, R.A. Shaw, H.E. Payne, and J.J.E. Hayes ed., pp. 58-61, 1995.
5. P. Moller, G. Miller, M. Urry, and B. Blacker, "Proposal Solicitation and Selection for the 21st Century", this volume.
6. A. Bose, G. Miller, A. Gerb, "An Interactive Tool to Aid in Proposal Preparation for the Hubble Space Telescope", *Telescope Control Systems*, SPIE Proceedings 2479, P. Wallace ed., pp. 434-444, 1995.
7. D. Asson, A. Bose, A. Krueger, "A Tcl/Tk-based, intelligent graphical editor for preparing HST programs", *Astronomical Data Analysis Software and Systems V*, ASP Conference Series, Vol. 101, 1996, G. Jacoby and J. Barnes ed., pp. 447-450, 1996.
8. M. Giuliano, "Achieving Stable Observing Schedules in an Unstable World", Mark Giuliano, *Astronomical Data Analysis Software and Systems VII*, ASP Conference Series in press.
9. M. Johnston and G. Miller, "Spike: Intelligent Scheduling of Hubble Space Telescope Observations", *Intelligent Scheduling*, M. Fox and M. Zweben ed., Morgan-Kaufmann Publishers, pp. 391-422, 1994.
10. J. Rose, "OPUS: A Flexible Pipeline Data Processing environment", this volume
11. C. Johnson, "Modification of an Existing Planning and Scheduling System to a New Mission", International Workshop on Planning and Scheduling for Space Exploration and Science, NASA-JPL, 1997.
12. C. Odea, A. Koratkar and K. Peterson, "User Support: Lessons Learned from HST", this volume.