

# Chandra Mission Scheduling On-Orbit Experience

Sabina Bucher , Brent Williams, Misty Pendexter, David Balke, Chandra Flight Operations Team, Northrop Grumman Space Technology, One Space Park Redondo Beach, CA. 90278

## ABSTRACT

Scheduling observatory time to maximize both day-to-day science target integration time and the lifetime of the observatory is a formidable challenge. Furthermore, it is not a static problem. Of course, every schedule brings a new set of observations, but the boundaries of the problem change as well. As a spacecraft ages, its capabilities may degrade. As in-flight experience grows, capabilities may expand. As observing programs are completed, the needs and expectations of the science community may evolve. Changes such as these impact the rules by which a mission is scheduled. In eight years on orbit, the Chandra X-Ray Observatory Mission Planning process has adapted to meet the challenge of maximizing day-to-day and mission lifetime science return, despite a consistently evolving set of scheduling constraints. The success of the planning team has been achieved, not through the use of complex algorithms and optimization routines, but through processes and home grown tools that help individuals make smart short term and long term Mission Planning decisions. This paper walks through the processes and tools used to plan and produce mission schedules for the Chandra X-Ray Observatory. Nominal planning and scheduling, target of opportunity response, and recovery from on-board autonomous safing actions are all addressed. Evolution of tools and processes, best practices, and lessons learned are highlighted along the way.

## 1. BACKGROUND

The Chandra X-ray Observatory, one of NASA's Great Observatories, is a space-based telescope designed to observe X-ray sources. Like any telescope, Chandra must be able to point at many locations in the sky and to set up its instruments to correctly observe a wide variety of science targets. Moving safely and efficiently from target to target and ensuring that the vehicle is correctly configured when it arrives at the core of Chandra operations. It is the Mission Planners who ensure that every observation and engineering activity is scheduled in accordance with the vehicle constraints and the observation requirements requested by the user. Before exploring how this is accomplished, it is important to have a bit of background on how the vehicle works and how the parameters of the scheduling process are defined.

### 1.1. Spacecraft

The Chandra X-Ray Observatory, shown in Fig. 1, is composed of three main elements: the Spacecraft system, the Telescope system, and the Integrated Science Instrument Module (ISIM). The telescope system consists of four nested pairs of mirrors and their supports system. The ISIM houses the science instruments (SIs) and the mechanisms used to configure them for observing. The spacecraft system provides the support structure and functionality that allows the telescope and ISIM to operate as an observatory.

Chandra has two science instruments that can be placed at the focal point of the telescope, the High Resolution Camera (HRC) and the Advanced CCD Imaging Spectrometer (ACIS), and two transmission

---

Send correspondence to Sabina Hurley; E-mail: [shurley@cfa.harvard.edu](mailto:shurley@cfa.harvard.edu)

Copyright 2008 Society of Photo-Optical Instrumentation Engineers.

This paper was published in *Observatory Operations: Strategies, Processes, and Systems II*. Edited by Brissenden, Roger J.; Silva, David R. Proceedings of the SPIE, Volume 7016, pp. 70160W-70160W-12 (2008), and is made available as an electronic reprint with permission of SPIE. One print or electronic copy may be made for personal use only. Systematic or multiple reproduction, distribution to multiple locations via electronic or other means, duplication of any material in this paper for a fee or for commercial purposes, or modification of the content of the paper are prohibited.

grating arrays (OTGs) which can be flipped into the path of the X-rays to provide spectroscopy data. The SIs are sensitive detectors and susceptible to damage from the radiation environment. Therefore an on-board radiation sensor, EPHIN, is used to detect elevated radiation rates and stow the instruments when rates reach pre-determined levels. ACIS and HRC are located on the Science Instrument Module (SIM), which is translated to position the specified instrument at the focal point and also provides focus adjustments. Each observation requires a specific instrument, translation position, focus position and grating, each of which is specified in a strictly formatted Observation Request (OR).

The OR also specifies the precise location of the observation target so that the observatory can be slewed to the correct position. Pointing of the observatory is provided by the Pointing Control and Aspect Determination (PCAD) subsystem. The PCAD subsystem system has gyros, an aspect camera, sun sensors, and reaction wheels to monitor and control where the telescope is pointing at any given moment. During observations the Aspect Camera Assembly (ACA)




Fig. 1. Chandra Spacecraft

is used to track stars and fiducial lights (point sources fixed in the spacecraft body frame), which are used to provide a highly accurate pointing solution. The reaction wheels are used to orient the vehicle, maintain pointing and store accumulated angular momentum. Small thrusters are used to unload angular momentum from the reaction wheels as required. Gyros provide spacecraft rate data which is used to control the motion of the vehicle. The sun sensors provide Chandra with autonomous pointing control capability, which, when particular faults are detected, will bring Chandra to a safe attitude.

A spacecraft cannot operate without power, and Chandra's electrical power comes from two solar array wings. The energy generated is distributed by the Electrical Power Subsystem (EPS), which also provides three banks of batteries to provide power during eclipses. Chandra goes through two earth eclipse seasons annually and also experiences occasional lunar eclipses.

To keep the instruments and spacecraft components operating correctly and the telescope alignment sufficiently accurate, temperatures on Chandra must be well controlled. The thermal subsystem uses several of methods to provide this control. Cooling radiators, insulation, heaters, thermostats, and reflective surfaces are used to keep each component within its specified thermal limits. As the spacecraft has aged the efficacy of the multi-layer insulation (MLI) for keeping unit temperatures stable and within limits has degraded. More specifically, the sun-side of the vehicle has warmed over time, causing exposed units on the sun side to near or exceed their maximum operating temperatures. To control temperatures, the sun exposure of these units must be limited. As the angle between the Sun and the spacecraft changes units receive more or less direct sunlight, depending on their position. So thermal control of several units on Chandra is now closely tied to sun angle.

Linking all of these subsystems and components together is the communications, control, and data management (CCDM) system. It provides transmit and receive capability through two low gain antennas, receives and processes commands from the ground and from the on-board computer (OBC), collects and distributes measurements from sensors, and stores data to two solid state recorders (SSRs). The Deep Space Network (DSN) is used to communicate with Chandra. Generally three one-hour tracks a day

provide opportunities to confirm the vehicle state of health, to uplink new command products and to dump the data stored on the SSR to the ground.

## 1.2. Constraints

Each instrument and subsystem has rules by which it must be operated. Furthermore, there are system level rules designed to protect the vehicle as a whole. Before launch these rules were documented in a constraints and restrictions document (CARD). With on-orbit experience some of the constraints in the CARD could be relaxed or eliminated. As the vehicle has aged new constraints have become necessary. Since not all vehicle constraints impact mission scheduling and since the CARD was designed for engineers, not schedulers, a second constraint document, the Mission Planning Guidelines, was created several years into the mission. This document provides the schedulers with a prioritized list of the rules and implementations they must make their scheduling decisions by. It is a living document and is frequently updated. To improve flexibility and accessibility the Mission Planning Guidelines have been recently ported from a single document into a TWiki web. The Revision Control System built into TWiki provides adequate configuration management for a controlled document. The TWiki web, hosted on the Chandra Operations Control Center (OCC) website, allows linking between the guidelines, the CARD, and reference materials and provides a widely accessible, platform independent means for viewing the guidelines. The guidelines are now the primary source for information on vehicle constraints that impact scheduling.

Although one generally thinks of constraints as flight rules for the spacecraft, they are not limited to how the vehicle is operated. In that each observation has specified parameters that determine how it is scheduled; each observation is itself a set of constraints. Some observations only define where the vehicle must point, for how long and which SI must be used. Other observations have exact timing constraints, roll angle requirements, or coordinations. Constraints that restrict how and when an observation is scheduled are specified in the same OR that defines the pointing and instrument configuration for the observation.

The combination of vehicle and observation constraints creates a constraint space, which is constantly evolving. Thermal constraints change with time and with season, the radiation zones evolve with the orbit, restrictions on use of consumables become increasingly difficult to meet with mission extension and observation constraints become more complex as observers gain experience using Chandra. To date, Chandra has had no mission impacting unit failures, but unit failure could also impact the constraint space. The table below provides a summary of the current scheduling constraints that Chandra Mission Planners work within for every schedule.

Table 1. Chandra Constraint Summary

Science	Target, Window, Phase, Roll, Coordination, and Target of Opportunity turnaround time all dictate when an observation can be scheduled and how difficult it is to schedule. Except for target, any of these can also be specified as a preference, which is less restrictive than a constraint.
Attitude	Sun position constraints, planetary and bright X-Ray source avoidance, and star quality requirements all impact when an observation can be scheduled.
Thermal	EPHIN temperatures, Propulsion Line temperatures, thruster temperatures and ACIS Power Supply temperatures all impact if and when an observation can be scheduled. The temperatures of these units are now dominated by sun exposure, which turns thermal constraints into time at sun angle (attitude) constraints.
Consumables	Minimizing number of momentum unloads and SIM moves do not currently drive scheduling, but if ignored entirely may become increasingly important
Radiation	Radiation Zones determine the time in any given orbit that can be used for observations

The mission planning system delivered at launch, the Offline System (OFLS) was designed to schedule all observations in accordance with the CARD and with constraints specified in the OR. However, very soon after launch, things started to change. Within the first few orbits of science operations, ACIS started incurring damage from low energy protons. This observed damage led to a new scheduling constraint that was not anticipated before launch, and, therefore, not incorporated into the ground system. Later followed changes to, among other items, scheduling of OTG moves, restrictions to the number of momentum unloads allowed and restrictions on the temperatures of various units. Each constraint change required an adjustment to the way the mission was scheduled. The Chandra scheduling process, therefore, by necessity, became adaptable and flexible.

While the precise issues facing Chandra are not likely to re-occur on other missions, most spacecraft will reach a point where they must be operated by a new and changing set of rules. This paper will walk through the Chandra scheduling process, highlighting lessons learned and best practices that could aid other missions in preparing for the inevitable parameter space change caused by vehicle age, on-orbit experience, environmental exposure, and changing resources and expectations.

## **2. Nominal Mission Scheduling Process**

### **2.1. Process Overview and Goals**

Chandra mission scheduling is performed on a week by week basis. Each weekly schedule includes all of the maneuver and reconfiguration commanding required to perform each observation and to safely stow the instruments when passage through the Van Allen belts makes the radiation environment too hostile to observe. Responsibility for mission scheduling is split between Science Operations Team (SOT) and Flight Operations Team (FOT) Mission Planners. FOT Mission Planning performs the detailed level scheduling that turns an unordered list of observations into a weekly observing schedule and finally command loads. This detailed level scheduling will be the focus of this paper.

The detailed scheduling process aims to generate optimal schedules for Chandra that meet all spacecraft health and safety constraints, meet science constraints of the observations and respect preferences of observers as much as possible. Chandra mission scheduling takes a heuristic approach to solving the problem of optimizing observing schedules. Schedules are manually ordered using the follow list of priorities: (1) Spacecraft Safety, (2) Science Goals, (3) Maximize Time for Science, (4) Minimize use of consumables. Chandra is fortunate in that at their current rate of use consumables will last beyond the planned duration of the mission, so use of consumables can be prioritized below science objectives.

### **2.2. Process Description**

Every year, once the call for proposals is complete and the peer review process recommends targets for observation, approved observations are translated into ORs. The ORs are laid out into a yearly observing plan, binned by week. This yearly plan is called the Long Term Schedule (LTS). Layout of the LTS is the first step toward generating efficient weekly schedules. First, the observations that are short and have no observing constraints are pulled out into the “pool”. Pool targets are not assigned to any given week, but are instead used to fill in around constrained and long observations, improving efficiency. All remaining observations are placed into a particular week of the year, based on science constraints specified in the OR and on availability of the specified target within vehicle constraints during that week.

Each week the targets scheduled in the LTS are put into an OR list that is delivered to the Flight Operations Team (FOT) Mission Planners. FOT Mission Planning then performs the detailed scheduling that turns an unordered list of observations into a weekly observing schedule and finally command loads. The mission scheduling software provided to the FOT Mission Planners at launch was designed as a “black box” that could ingest ORs and Engineering Requests (ERs) and output efficient mission schedules. An optimization routine with knowledge of spacecraft and observation constraints would do most of the work of mission

scheduling. In practice, the optimization routine did not generally produce desirable results. Furthermore, on-orbit experience quickly led to new scheduling constraints that were not anticipated before launch, and, therefore, not incorporated into the optimization routine. So, while each new constraint was incorporated into the scheduling software, building safe and efficient schedules required manual intervention.

Unfortunately, a lack of visibility into the software made manually assembling a useable mission schedule an exercise in trial and error and thus an extremely inefficient process. To remedy the problem, the FOT designed and developed a stand alone, GUI based, tool, the ORViewer, that reads in observation and attitude requests and displays them in several frames of reference. Fig. 2 provides a snap of the ORViewer GUI. The displays allow the scheduler to visualize the attitudes required for the week and to assemble a sequence of maneuvers and dwells (an attitude profile) through an interactive timeline view. An efficiency view allows the scheduler to immediately assess the impact of schedule changes on efficiency. The ORViewer also allows adding in ERs, which schedule activities required for spacecraft maintenance and other non-science related activities. As requested, the ORViewer pushes the attitude profile to a second, home grown tool, the Maneuver Constraint Checker (MCC). MCC displays the attitude profile with respect to various scheduling constraints, providing visual feedback of the impact of schedule changes on each constraint. The schedulers move back and forth between ORViewer and MCC until they have assembled an attitude profile that meets all vehicle constraints and efficiently moves through all of the observations planned for the week. These two tools remain a very important part of Chandra mission scheduling processing and ordering the OR list using them is first step in assembling a schedule.

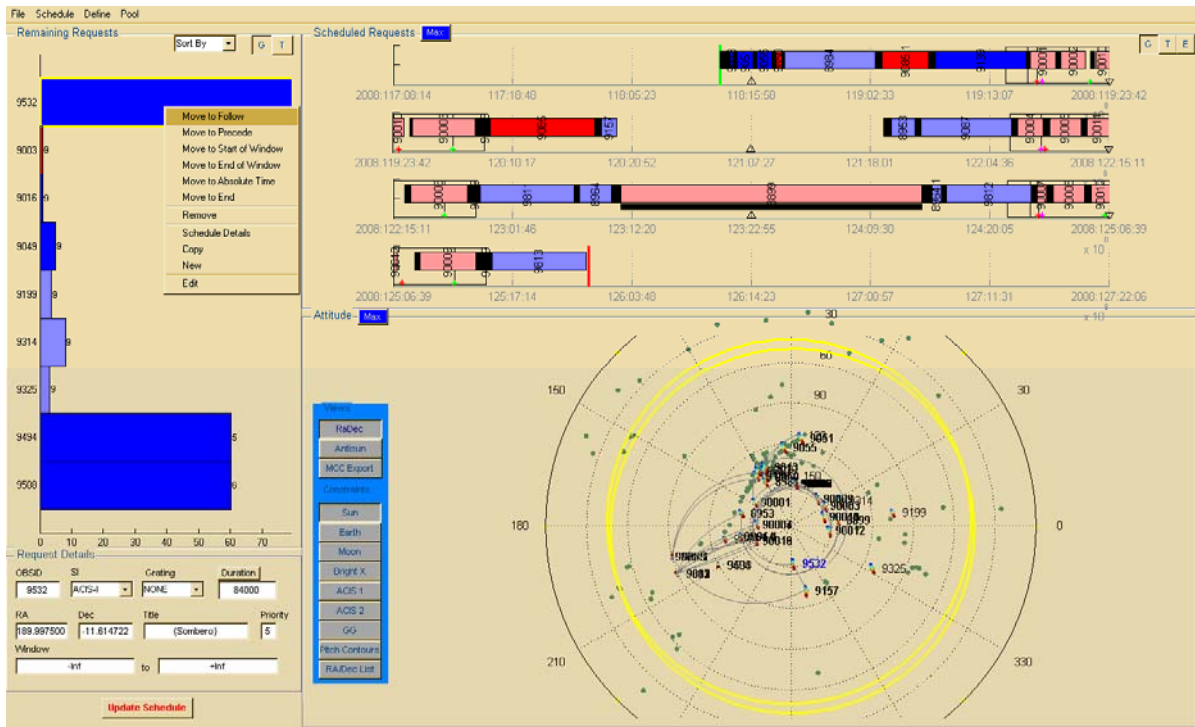


Fig. 2. ORViewer scheduling tool graphical interface. Each panel can be toggled to display different views of the schedule. Constraint displays can be toggled on and off to provide maximum capability without visual clutter. Scheduling is achieved through context menus in the timeline and remaining request panels.

Once the ORs and ERs are assembled into a weekly schedule they are linked together and put into the OFLS scheduling routine. Linking puts the observations into a single chain, so that the optimization routine's algorithm only has one solution; the one designed by the scheduler. It may seem odd to pass an optimization routine a solved problem, but the routine is central to the design of the OFLS scheduling software and cannot be disabled. The OFLS scheduling software currently provides the only algorithm with knowledge of all of the exact timing and commanding requirements used to build safe mission schedules. So, by necessity, it is used to build command products from an assembled schedule. Unfortunately, linking

observations together is not guaranteed to force the optimization routine to schedule observations as requested. The mission schedulers, therefore, do a fair amount of troubleshooting, trying to get the optimization routine to return a schedule that matches the observing plan laid out in the ORViewer.

Once the OFLS schedule is successfully replicating the observing plan laid out in the ORViewer, the schedule products are passed to an external algorithm for star selection. The OFLS scheduling software identifies candidate guide stars for observations before scheduling them, and only schedules observations its algorithm determines have adequate stars. With the original settings, this algorithm disallowed many observations. Clearly, approved observations need to be observed and should only be disallowed as a last resort, so the star selection parameters were relaxed to allow observations to schedule. The relaxation had the unfortunate side effect of producing catalogs with insufficient stars for accurate aspect reconstruction and, sometimes, for adequate attitude reference. Initially, unacceptable catalogs were edited by hand, with ACA engineers hand selecting desirable stars and mission planners editing them into command products. Over time the process was automated into tools and scripts. Finally, a working group of PCAD and ACA Engineers and Mission Planners defined a star selection algorithm that shifted the way star selection was performed. It assumed that an adequate catalog would be found for all observations, and accepted the risk of a re-schedule if one could not be. The algorithm was coded into a stand alone tool, the Simple AXAF User Select Acquisition and Guide star Editor (SAUSAGE), that would be used to replace every star catalog used in mission schedules. SAUSAGE also provided a GUI interface for viewing and editing star catalogs. The tool is used to this day and observations are very rarely rescheduled due to poor quality star catalogs.

Once stars have been replaced the schedule products are used to generate command products, which can be uplinked to the spacecraft. The scheduling software and the command load generation software have a single interface, the Detailed Operations Timeline (DOT). The DOT is simply a formatted text file, so it can be easily modified. The DOT is what allows replacement of star catalogs after the fact and provides a means for mission schedulers to correct small scheduling problems by hand. While hand editing schedules is not desirable, it provides ultimate flexibility. This flexibility has proven to be a necessity in efficiently generating mission schedules. The command load generation software processes the DOT and uses predefined command sequences to generate command loads.

Once the schedule has been made into command loads, the schedule and command products are packaged for review. The review package is ingested by a suite of independently developed and maintained set of Perl scripts, Backstop, that reverse engineer the command loads into command mnemonics and time-tags. The resulting report is processed and reviewed by each subsystem, both science instrument teams and science mission planning. Each reviewer reports back to the FOT Mission Planning Manager identifying any problems detected in the loads or approving the load set. Once review comments have been received from every review team, problems are addressed by rescheduling or by editing the DOT and new command loads are generated. The new command loads go through the same review process until no problems are reported. Once every reviewer has approved the load set they are approved by the FOT Mission Planning Manager and Flight Director. Approved command loads are saved into a database, which enables the spacecraft controllers to uplink them to the spacecraft.

The scheduling process, from receipt of the OR list to command load uplink, takes approximately two and a half weeks to complete. The end-to-end process is summarized in Fig. 3. While the process may seem manual and labor intensive, it is highly effective and generally well controlled. An average of 93% of available science time is used for observations. New constraints or requests are incorporated quickly. Finally, when questions or problems arise the Mission Planner responsible for the week knows the details of every decision made in scheduling that week and can, therefore, respond more quickly and accurately.

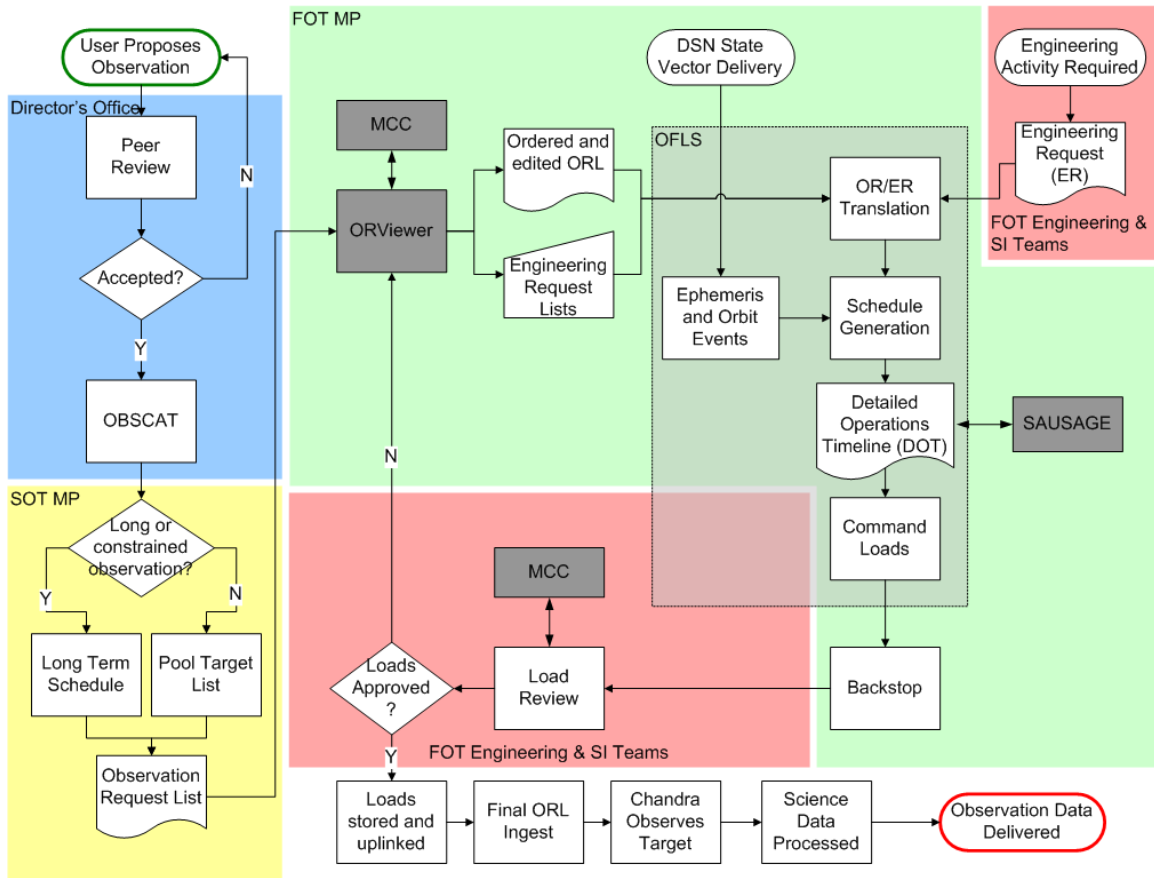


Fig. 3. End to end Chandra mission scheduling process

### 2.3. Process Controls:

The two metrics that best monitor the performance of the scheduling process are: (1) effort required to generate a schedule, and (2) the efficiency of the schedule produced. While it is difficult to get accurate measurements of effort, efficiency is a simple calculation. The Chandra FOT computes and tracks schedule efficiency on a weekly basis and uses more qualitative measures of effort required to generate a schedule. Not surprisingly, effort and efficiency tend to track together. As schedules become harder to assemble, they also become less efficient. As Fig. 4 shows, schedule efficiency has a rather large standard deviation, but the mean was fairly consistent until 2004. In late 2004, efforts to protect the EPHIN radiation detector from getting too hot and the propulsion lines from getting too cold forced splitting observations into pieces with large maneuvers between them. The scheduling problem had become over-constrained and, as a result, the effort to schedule observations increased significantly and the availability of science targets and efficiency of schedules suffered. Clearly, the scheduling process was no longer well controlled.



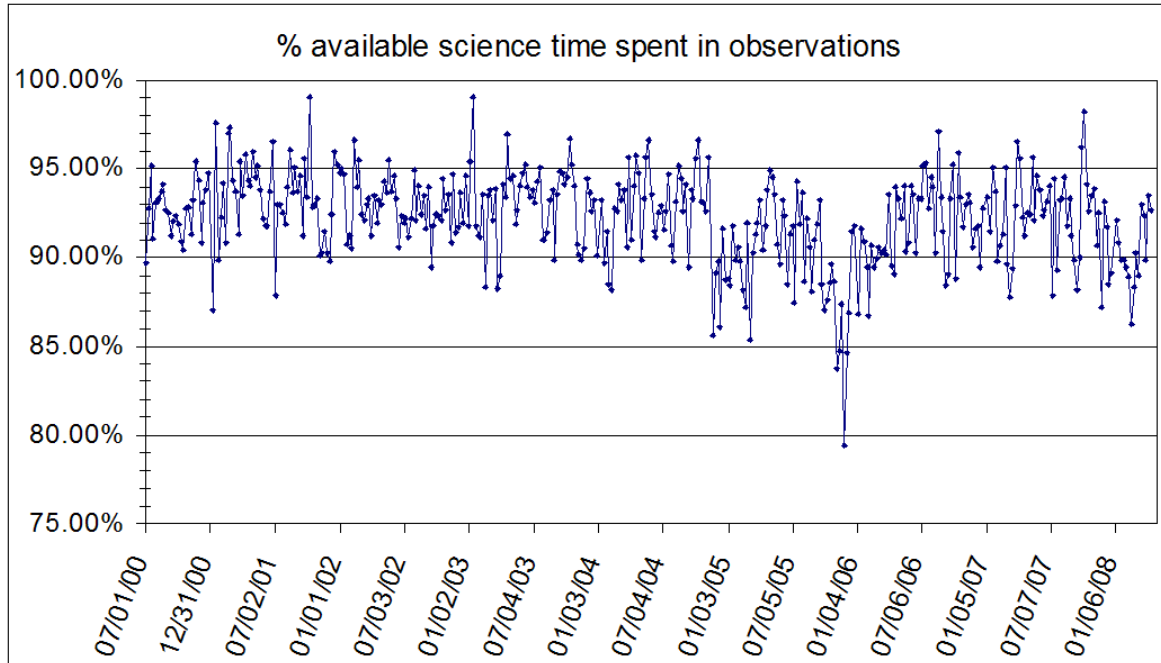


Fig. 4. Weekly scheduling efficiency. Calculated by dividing time spent observing science targets by time available for observing. Week to week variation largely due to OR list profile. Trends generally due to constraint migration.

The decline in scheduling ability prompted cross-team analysis and discussion, which led to relaxation of the EPHIN temperature limits, which, in turn, allowed scheduling ability to recover. This cross-team group was the precursor to a new working group, the Mission Planning and Constraints Working Group (MPCWG). The mission of the MPCWG is to minimize the impact of constraints while protecting the safety of the vehicle. The group has representatives from mission planning, engineering, both SIs and project management. Since 2005, the group has been actively and effectively managing constraints and optimizing scheduling, which has kept scheduling efficiency near early mission values.

While it is intuitive that if a constraint is relaxed, scheduling will get easier, it is not simple to discern how much impact a given relaxation will have. It may be that all of the work put into relaxing a constraint is for naught, because another constraint overwhelms scheduling in that part of the constraint space. To resolve this problem, the MPCWG developed a tool that can run all of the models used in scheduling and display the impact of constraint changes. Given a set of constraints and starting conditions, the tool generates a maximum allowed dwell time at every attitude in the sky. The user can set constraint thresholds, model parameters or output dates and regenerate maximum dwell times. Maximum dwell time correlates well to scheduling efficiency and target availability, since shorter allowed dwells forces splitting long observations into parts, which decreases target availability, and requires more maneuvering, which reduces efficiency. So measuring maximum allowed dwell time gives insight into scheduling ability, and generating a delta chart allows the user to see the impact of a constraint change. The tool can analyze the impact of given constraint over time, of a constraint relaxation or of a new constraint. This constraint integration tool has been fully incorporated into the Chandra constraint management process and is used to assess the impact of constraints over time and inform decisions regarding constraint relaxation versus scheduling efficiency. It also provides dwell time data that is now fed into the process of laying out the yearly observing plan. This predictive data maximizes the probability that all of the observations chosen for a given week will be able to be scheduled once that week arrives. Fig. 5 provides a sample output from the constraint integration routine.






Fig. 5. Maximum allowed dwell times by Ra and Dec for late 2007. Plot generated with constraint integration tool, which also outputs delimited text files for processing and import into other tools.

### 3. Fast Turn-Around Scheduling

#### 3.1. Targets of Opportunity

Target of Opportunity (TOO) observations can have requested turnaround times ranging from hours to months. Any TOO with a requested turnaround time greater than three weeks is incorporated into the nominal weekly scheduling process. For observations requested to occur within three weeks, extra effort is required to schedule them. If the load set where the TOO will be placed has not yet been uplinked to the spacecraft, the Mission Planner uses the ORViewer to incorporate the TOO into their observing plan and goes through the nominal scheduling process at an accelerated pace. The scheduler's goal is to disturb as little of the original observing plan as possible, which minimizes the effort required to incorporate the TOO and the chance for introducing errors. However, when inserting a TOO observation, something must be removed from the schedule. ORs with time constraints or those coordinated with other observatories are removed only as a last resort. So the scheduler must balance keeping the schedule intact with ensuring that only unconstrained observations are removed. This is generally achieved through close coordination between flight team schedulers and science mission planners.

If the requested turnaround time places the TOO observation within the week running on the spacecraft, the running loads must be interrupted and replaced with a new observing plan. Fast turn-around scheduling that replaces the running loads is referred to as a "re-plan". In order to provide fast turnaround times, flight and science team schedulers are on-call 24 hours a day. An automated email system sends messages to pagers and cell phones alerting the planners to a fast turnaround re-plan. Science mission planning immediately starts to coordinate with the Director's Office, science instrument teams and observer to finalize the parameters of the observation as quickly as possible. While this happens, flight mission planning takes the information they have (generally target and turn-around time) and generates a straw-man schedule.

Re-plans must take into account the requested turn around time, maneuvers required to reach the new target, spacecraft constraints, communications times and the time required to build and review a new set of products. Due to on-board memory allocation the weekly loads are divided into smaller load segments, generally lasting approximately a day and a half. Re-planning a running load segment requires stopping commanding, while it is running, and then re-planning it with new commanding. Once the running load segment has been stopped, there is no on-board commanding until a new load has been activated. Due to the slight increase in risk and the increased complexity of uplink, interrupting a running load segment is avoided whenever possible. To allow viewing all of the items that must be considered when re-planning in a single easy to use format, a graphical timeline of mission schedules is provided on the Chandra OCC website. The graphical timeline is used to give a snapshot view of the current and planned configurations of Chandra and the communications opportunities for uplink. The snapshot is used to find the first and last available uplink opportunities, load segment breaks, and to identify the lowest impact spot to place the TOO that still meets the requested turnaround. All of these items are considered when selecting the load interrupt time.

Once the load interrupt time is identified, the flight scheduler begins to use the ORViewer and OFLS to generate the replacement TOO schedule. Scheduling usually begins before the observation parameters have been finalized. The planner uses a stand-in OR, filling in as much information is available at the time, to begin the scheduling process. Once all the parameters have been set, a final OR is delivered. Generally flight mission planning simply needs to re-run the software with the final parameters and perform the last few steps of the scheduling process when the final OR is delivered. If flight mission planning were to wait until all of the OR parameters were finalized to begin scheduling, the process would be delayed by hours. Once the re-planned schedule is complete it goes through the same review process as nominal weekly loads, again on an accelerated schedule. All reviewers are also on-call 24 hours a day and email messages are sent to pagers and cell phones indicating that a fast turn around review is required. Once the loads are approved, a Command Action Procedure, which directs the console operators to replace the running loads with the TOO loads, is generated, reviewed and executed.

### **3.2. Safing Action Response**

Chandra is equipped with a cleverly designed safing subsystem. The flight software monitors for survival threatening conditions instead of specific unit failures. This system provides robust protection against unit failures, environmental changes and ground based errors. When the vehicle detects a survival threatening condition it takes a proportional response to safe the vehicle. Most safing actions safe the SIs and stop the running loads. Others add attitude safing by acquiring bright stars for guidance or by aligning the vehicle normal to the sun vector. The most drastic safing action, Safe Mode, swaps all core spacecraft units and uses a backup processor to bring the vehicle normal to the sun vector. Any safing action that stops the mission loads, requires mission planning response. A similar paging system to the one used for TOOs is used to alert engineers and mission planners when a safing action has occurred.

Once the root cause for the safing action has been identified, any necessary recovery actions have been completed and the engineering team has confirmed that it is safe for the science mission to resume, mission planning begins to re-plan the weekly schedule. Generally, safing actions will cause some science observations to be missed. The re-plan must take skipped observations into account. First of all, when a safing action occurs and observations are missed, the configuration of the vehicle is not as expected by the original schedule. Continuity between the vehicle and the mission schedule must be re-established; therefore, the original plan cannot simply be restarted. Additionally, if an observation with a coordination or time constraint is missed all efforts should be made to reschedule that observation within its required window if safely possible. Rescheduling constrained observations, however, must be balanced with ease of resumption. As with TOO scheduling, the scheduler will aim to minimize changes to the original plan to minimize the effort required for the re-plan and the risk of introducing errors. With all these complicating factors, the graphical timeline, used to re-plan schedules for TOOs, is also a valuable resource in planning science resumption loads. Once the resumption plan is laid out and approved, and continuity is accounted for, science resumption scheduling follows the nominal scheduling process, at an accelerated rate.

When the mission loads are stopped, the planned attitude profile stops. As mentioned in Section 1, thermal control of several units on Chandra is now based on attitude. If the loads are stopped at the wrong attitude, the units can quickly get too hot or too cold. Additionally, planets can cross in front of the star tracker and angular momentum can accumulate quickly when an attitude is held for an unplanned period of time. If science cannot be resumed in time to avoid negative consequences of dwelling at an unplanned attitude a “maneuver only” load must be generated and uplinked. Maneuver only loads contain only the commanding required to perform a maneuver to a safe attitude and can be created far more quickly than science resumption loads. Maneuver only loads follow a slimmed down version of the nominal scheduling process.

To date, the most common safing action, by far, is radiation safing. When the radiation detector senses unacceptable rates or when a decision to safe the SIs is made on the ground, the loads are stopped and the SIs stowed. This also stops the planned attitude profile. At the levels protected against by the radiation detector, radiation only poses a threat to the SIs. The vehicle could continue maneuvering without concern. Continuing the planned attitude profile would avoid the negative consequences of remaining at an unplanned attitude. It would allow for more rapid resumption of science by eliminating the need to maneuver from the attitude of the stopped observation to the first target in the resumption plan. To allow maneuvers to continue without science commanding, engineering and science loads must be separated, but run concurrently. When a safety issue threatens the SIs, but not the rest of the vehicle, only the science loads can be stopped. If a safety issue threatens both, then both sets can be stopped. In anticipation of the next solar maximum, the Chandra FOT is investigating changing the ground and flight software to allow this separation.

### **3.3. Anomaly Response**

To date, the anomalies on Chandra have been detected well before they caused a safing monitor limit to be reached. Several anomalies have, however, required changes to mission schedules. If the running loads must be reworked to mitigate or work around an anomaly, the TOO scheduling process, without the effort of generating an OR for a new observation, is used to re-plan observations and engineering activities. If the loads cannot be reworked in-time to avoid undesirable reconfigurations, then they will be stopped and the SIs stowed. In this case, the safing response re-plan process is followed.

When an anomaly requires a long term change to the scheduling process, the MPCWG coordinates the required changes. New constraints often start out very conservative. The conservatism combined with the learning curve for schedulers working with a new rule and the lack of automated tools to help adhere to the new constraint can cause a temporary decline in scheduling efficiency and increase in scheduling effort. The MPCWG works to mitigate the decline, by investigating safe ways to relax the new constraint and methods to help the planners work within it.

## **4. Conclusions**

The mission scheduling process requires balancing many factors, such as, observation constraints, observer preferences, maneuver size, momentum accumulation and unloading, mechanism motions, thermal management and attitude restrictions. Such a complex problem leads people to want to develop an algorithm that does all of the work, a “black box” scheduling system that ingests requests and outputs a schedule with no manual intervention. However, such a system simply cannot support rapid changes to constraints and priorities and cannot answer questions when problems arise. Maintaining mission scheduling capability in the face of evolving constraints and expectations requires a method that allows specifying how requests are scheduled. To prevent such specifications from being trail and error, the system must provide the user with the information required to make scheduling decisions. The Chandra Mission Planners use visual representations of the schedule and the constraints to assemble a set of maneuvers and dwells that achieves the scientific goals of the mission, without compromising vehicle safety, in the most efficient manner possible. Optimization routines can be very helpful in solving the complex problem of mission scheduling, but they cannot be the only answer. A scheduling algorithm cannot work with engineers as constraints evolve, and cannot rapidly change its priorities and rules the way the human brain can. Adaptive mission scheduling requires a balance of automated tools and human input.

Providing the scheduler with all of the information that goes into scheduling and using optimization routines to aid, not replace, the scheduler allows mission scheduling to evolve as the vehicle ages.

Like any program, Chandra has experienced problems that were unanticipated before launch. To date, a majority of the problems experienced have been successfully mitigated by modifying mission scheduling. However, at one time these mitigations over-constrained the scheduling problem and caused a drastic reduction in science time efficiency. By introducing scheduling process controls and a cross-team group tasked with balancing spacecraft safety and scheduling capability the Chandra mission recovered this lost efficiency. New tools that anticipate constraint impacts allow proactive investigation and planning and have successfully kept scheduling efficiency high, despite several new constraints.

The tools and techniques developed during the operational portion of the Chandra mission have allowed the mission scheduling process to adapt to new constraints and changing expectations. An environment that fosters innovation empowered the FOT to develop these new tools and processes and continues to encourage improvements to this day. Constant innovation has expanded the resources available to schedulers, keeping the scheduling process in control, despite increasing complexity.