Chandra Data Processing: Lessons Learned and Challenges Met

J. Nichols[a], C. S. Anderson[a], P. J. Mendygral[a], D. L. Morgan[a], and G. Fabbiano[a]

[a]Smithsonian Astrophysical Observatory, 60 Garden St., Cambridge, MA, USA

ABSTRACT

Six years into the mission, Chandra data processing operations has reached a stage of maturity that allows nearly complete automation as well as dynamic flexibility to accommodate future changes in mission and instrument status and constraints. We present a summary of the procedural and technical solutions that have been developed since the launch of Chandra to meet unanticipated challenges in the area of data processing. Lessons learned concerning data processing are discussed, including an explanation of the source of each problem and the Chandra team's response to the problem. Potential pitfalls that might affect future projects are also included. The user interface, data quality screening, and quicklook software developed specifically to address issues identified after launch have proved valuable in meeting the goals of low-cost, efficient, and flexible mission operations for the Chandra mission and can provide insight for future mission designs.

Keywords: Data processing, data operations, processing automation

1. INTRODUCTION

Data System Operations (DSOps) receives incoming raw telemetry from the satellite via the Chandra Ground System and processes the data to create archival science and engineering data products, performs Validation and Verification (V&V) on the final science data products, creates quicklook data and images for anomaly and acquisition problem detection, and transfers realtime data from communication passes with the satellite to numerous groups at Chandra X-ray Center (CXC) for evaluation. In addition, DSOps reprocesses the entire mission dataset every few years to make the most current calibrations and algorithms available to users of the archive. The Automated Processing (AP) system[?] that incorportates Standard Data Processing (SDP) tools[?] is used to process the data for distribution to the users.

This paper discusses a system of data quality control and processing automation that has been developed by DSOps. Lessons learned and challenges met during the course of seven years of the missions are also presented in the context of data processing operations. Overall, many solutions to our the lessons learned are easily transferable to other missions, including the front-end software we have developed which functions outside the science data processing system.

2. OPERATIONAL AUTOMATION

At the time Chandra launched, there were no tools specifically designed for tracking observations in data processing operations or for automation of routine procedures. Such tools, in general, are developed after launch, addressing the realities of flight. The obvious long-term goal of data processing was to implement maximum automation of all processing, both AP and reprocessing efforts. The standard data processing and any reprocessing efforts need to run without operator intervention as much as feasible, both to reduce the need for manual

Further author information: (Send correspondence to J.S.N.)
E-mail: jnichols@cfa.harvard.edu, Telephone: 1 617 496 7149

intervention and thus the possibility of operator error, as well as minimize resource usage and processing time. Operator error is the more costly issue, because of the time required to determine the source of the problem and the sometimes extensive cleanup. Therefore, the primary goal in the design of automation in data processing was the reduction of manual intervention. The secondary goal was the reduction of the time required to process and deliver the data to the user. These goals have been achieved and the resulting decrease in elapsed time between end of the observation and data distribution is shown in Figure 1. The figure is a bar graph with indications for average time in each 6 month interval of (1) end of observation to data receipt by DSOps, (2) data receipt by DSOps to submission to AP, (3) submission to AP to completion of Level 2 processing,[?] and (4) completion of Level 2 processing to completion of V&V and data distribution. Improvements in the time required to process data in AP (3) are due to replacement of older machines with faster ones, as well as improvements in the efficiency of the SDP code. Improvements in (2) are due to the operational automation described in this paper. Improvements in (4) are due to procedural changes and lessons learned that are described in this paper, as well as some automation tools provided by SDP.

2.1. The Operational Interface: The Telemetry Tracker

We have developed a data operations tool known as the Telemetry Tracker (TT9), which provides control and organization of nearly all aspects of data processing operations, as well as automatic initiation of quicklook and AP processing. This system serves as a front-end to AP. The system is written in Perl and may be easily adaptable to other missions. It consists of a customizable interactive GUI, daemon processes, and a database backend used to store information specifically relevant to DSOps. Management of data processing is simplified due to the dynamic display of the status of all observations in the processing system. Figure 2 shows how TT9 fits into the Chandra Data Systems operational scenario.

A key feature of TT9 is its flexibility. Because it does not have the extensive infrastructure of the AP processing system, it is easily modifiable on a short time scale to respond to changing spacecraft and instrument operation parameters and anomaly resolution support. The automation of processing provided by this system has greatly reduced the need for personnel, thus contributing to cost savings for the project. One of the more visible successes of TT9 is the near elimination of delayed raw telemetry verification prior to AP submission. For several years this stage of processing accounted for approximately 25% of the total time from the end of an observation to completion of V&V but has now been reduced to less than 8%.

TT9 is easily maintained by 1 FTE. It is well suited to use by other missions because it does not interact directly with any files except the incoming data files. Automated database queries are used to determine the state of the processing system and the processing status of each observation. Daemon processes execute these queries every few minutes. The database that tracks all of the information needed to populate the GUI is a MySQL database with tables that are easily modifiable to collect relevant information for other missions. The information is primarily organized by observation, with timestamps for each stage of processing useful for compiling metrics. In addition, observation reports of unusual acquisitions, the validation and verifiction reports, and data quality reports are stored by observation and by downlinked raw telemetry file. The evolution over time of the Chandra data processing automation, quality control, and quicklook automation is shown in Figure 3.

2.1.1. Custom processing scripts: secondary method of pipeline processing

Prior to launch, it was clear that having a means to completely process an observation outside of AP, the primary processing system, would be useful for trial run processing of problem datasets. Due to the strict requirements placed upon AP, that system is not well suited for concurrent trial run processing many observations from various time periods in the operational environment.

Manual procedures used for such trial run processing formed the basis of a series of scripts now referred to as the Custom Processing (CP) scripts that reduce manual processing of a single observation starting from raw telemetry to a single command. CP scripts are wrappers around the same pipeline tools used by AP. The CP data products are as identical to AP products as possible; however there are key differences prior to Level 1 products. CP ignores restrictions on low level data products such as overlap and versioning (see Fig 4), which allows processing observations independent of one another and makes the CP scripts ideal for trial run processing. AP adheres to all data product requirements to interface with the Chandra data archive. By disregarding all interface

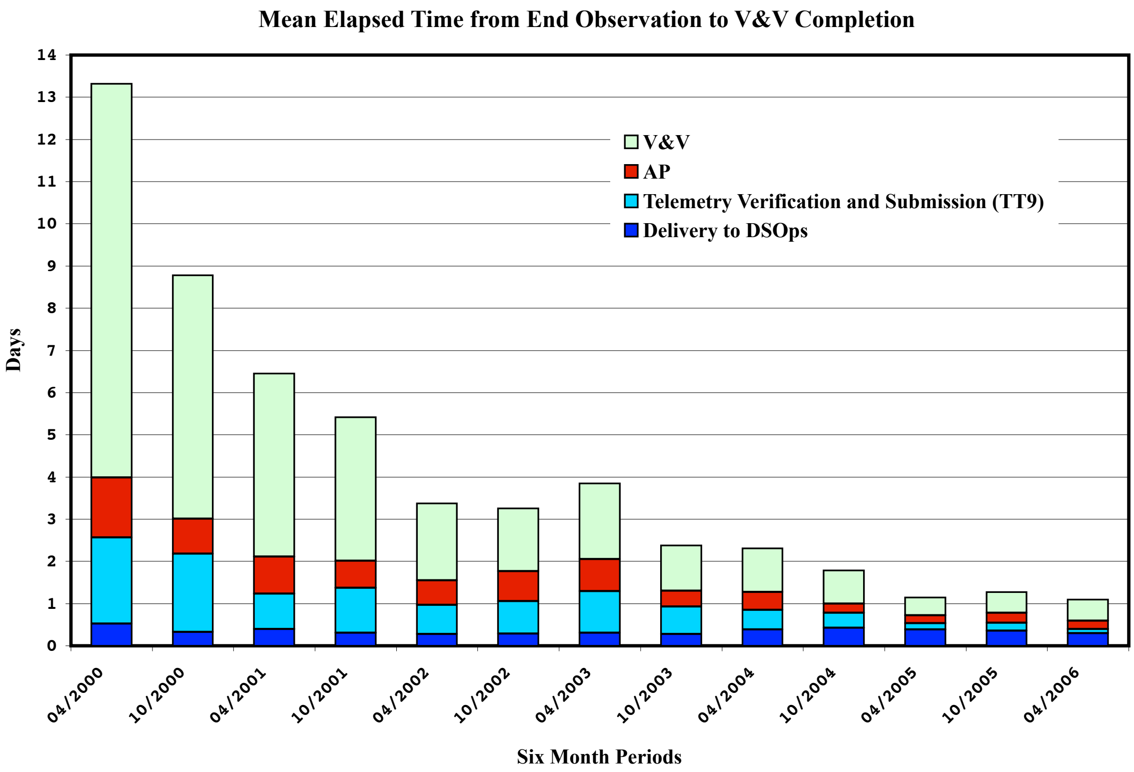

Figure 1. Improvements in automation at each stage of data processing have dramatically reduced the elapsed time from the end of an observation to data distribution. Early in the mission, an observation could be expected to be in the data processing system for weeks. Today, the majority of observations are out the following day. During the time range presented, personnel were reduced from 7 to 4 operators, which additionally attests to the increased end-to-end efficiency of Chandra data processing.

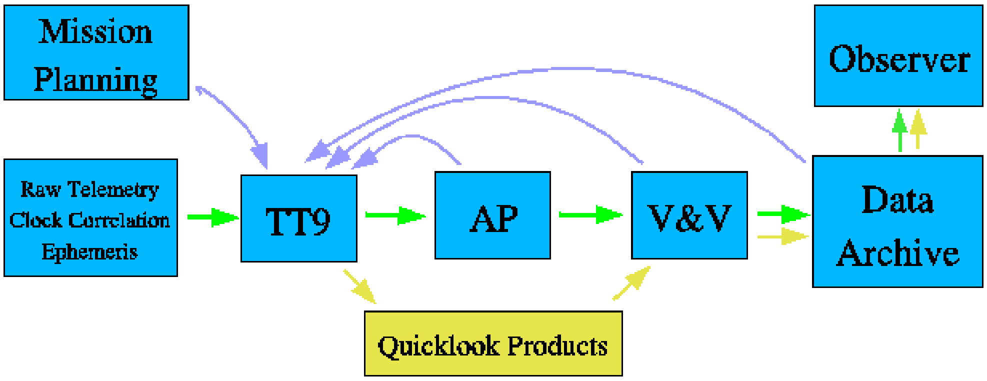

Figure 2. A high level diagram of Chandra data processing flow, showing how TT9 fits in the entire process. Quicklook products produced with Custom processing scripts can optionally be sent to V&V for distribution (lower path), but only AP products can be archived. TT9 monitors each stage of processing for end-to-end visibility (curved arrows).

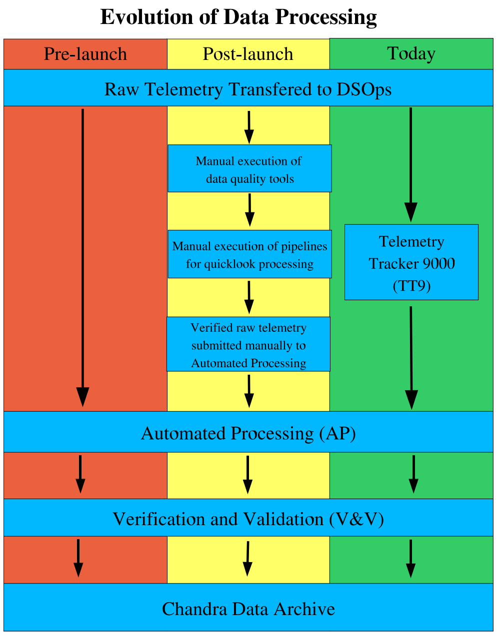

Figure 3. Evolution of data processing steps over the life of the Chandra mission. Pre-launch activites include test datasets.

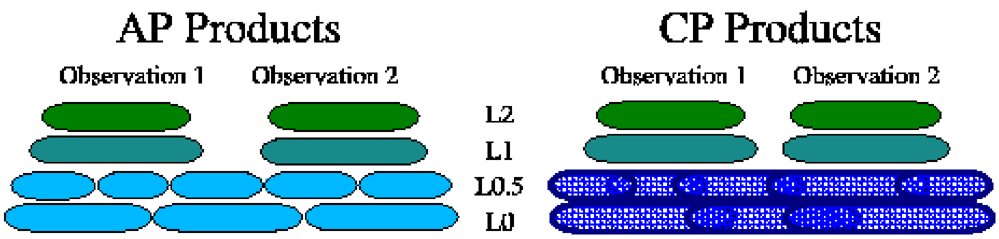

Figure 4. Comparison of AP and CP data products. AP data products at all levels are required to be contiguous and non-overlapping. CP ignores this requirement for Level 0 and 0.5 data products to provide a means of simple and efficient trial run processing. Levels 1 and 2 CP data products are identical to AP, which are distributable data products.

requirements, CP provides a secondary method that is well suited for generating quicklook data, processing an observation with incomplete raw telemetry, repeatedly processing an observation for test purposes and providing time-critical data products that may not be immediately available from AP.

Years into the mission the CP scripts remain useful for rapid processing of time-sensitive observations, a simple means to test new calibration files and pipeline changes, and for quicklook processing, which doubles as a form of raw telemetry quality control.

2.1.2. Quicklook processing

The request by the Science Operations Team for quicklook data of each observation was brought to our attention shortly after launch. We developed a very simple technique for producing quicklook data through a truncated Level 1 processing using the CP scripts. The quicklook system is simple to maintain because it lacks the infrastructure and internal tracking of the AP system. The quicklook processing system only knows about one observation at a time. Quicklook data are produced automatically as raw telemetry files become available. The result of quicklook processing is returned to TT9 as "success" or "failure" as a data quality test. For observations that span several raw telemetry files, partial quicklook data are produced as each sequential raw telemetry file arrives. These partial quicklook data allow early identification of a problem with the acquisition. The quicklook images produced from quicklook processing have been the source of identification of at least 80% of the instrument anomalies and incorrect acquisitions. Quicklook images are available within a couple of hours after the raw telemetry is received, meaning these unusual situations are caught early enough to reschedule time-critical observation or modify.

2.1.3. Incoming data quality control

Incoming raw telemetry files from the spacecraft are verified with a number of tests that identify lost data in the Ground System, which may be recoverable, double-bit corruption that is not detectable by the Ground System, data lost air-to-ground that may be recoverable by a repeat downlink from the spacecraft, and missing data in crucial packets that can prevent the observation from processing correctly. Occcasionally missing data or double-bit corruption in critical data packets require the reconstruction of the missing data in order for the observation to process. This process is manual and often requires the support of the instrument team to recreate the missing data, such as a parameter block for Advanced Camera for Imaging and Spectroscopy (ACIS) that instructs the instrument configuration for the observations. If data are missing from an ACIS bias file, the data processing group has developed an algorithm for replacing the missing bias values so that the observation is useful. Because the data quality checks identify these situations of missing critical data before the data are submitted to AP, corrected files can be delivered, tested, and submitted so AP runs without error.

Clock correlation files and ephemeris files are required for processing Chandra data and are delivered twice a week. Each file is assessed for consistency with the previous file and with the trend for the last year. Large deviations can indicate an error in the creation of these files. Raw telemetry files, clock correlation files, and ephemeris files must pass quality control before being submitted to AP. TT9 prevents a clock correlation file that fails data quality checks from being submitted to AP until an operator overrides the TT9 hold status or alternately rejects the file in favor of a replacement.

2.1.4. Automated data submission to processing system

Automation of the submission of raw telemetry to the processing system has essentially eliminated AP latency from the Chandra data processing time metrics. This automation was achieved using a component of TT9 that queries the operational database for verified raw telemetry and monitors the overall status of the AP system. Data are submitted to AP once (1) the raw telemetry has passed quality control measures, (2) quicklook processing has completed without error and (3) the state of the AP system is nominal. At the 95% level, these conditions are satisfied and the full archival processing proceeds without operator intervention. This allows data processing to occur 24/7 although shifts are only staffed 13 hours per weekday with 1 8-hour shift on Sunday. Before raw telemetry is submitted to AP, the state of the AP system is checked by TT9. We have chosen very conservative criteria for these checks, although the criteria can be modified if necessary. The AP processing system must be idle and all logs error-free. There can be no individual pipeline in AP in an error state unless that error is flagged to be ignored by an operator. All previous automated TT9 tasks such as disk cleaning and archive file ingestion must have run successfully. If these criteria are met, the next raw telemetry files that have passed data quality checks and completed quicklook processing are automatically submitted to AP. The automation has resulted in a significant improvement in data delivery statistices for Chandra.

Our experience with the automated submission of raw telemetry to the data processing system, which has been active for more than two years, has been very positive. There have been no instances of preventable inappropriate raw telemetry submissions to AP during that time, and there have been no instances of errors in the processing that caused extra resources to resolve because of the automation.

2.1.5. Operational database

Information management and organization were identified early on as functions that needed simplification for DSOps. The TT9 system alone must track thousands of observations, input files and processes and have diagnostics available for internal calculations and criteria evaluation, as well as for operator use. Maintaining such a quantity of data in an organized and efficient manner is best handled by a database. DSOps selected the MySQL database server as the backend for TT9 and dynamic web pages. The decision to use MySQL was driven by several factors including ease of administration, quality of available APIs, cost and performance. A database server that is completely independent of the archive was preferred, allowing DSOps to maintain maximal flexibility in response to the Chandra operation situations. The database used by TT9 is relatively simple and does not require many of the features that more advanced relational databases offer. MySQL is an open source database server that is readily interfaced with a Perl DBI package, DBD::mysql, and satisfied all of our requirements. The available APIs also enable clients of many forms including CLIs, GUIs and web pages. This has allowed DSOps to maintain maximum system wide visibility and availability because operators can monitor and manage TT9 and other DSOps systems through each of the client interfaces. Other missions should consider these factors and more when choosing the database server that will best suit all of their teams' needs.

2.1.6. DSOps GUI

The primary interface for monitoring DSOps system status is the TT9 GUI. The TT9 GUI is a Perl/Tk application that contains "tabs", following a notebook theme each defined by a particular task or set of information to display. The goal of the TT9 GUI was to provide a highly visual end-to-end interface for operators to manage TT9 and other processes. Several tabs display information about telemetry files, observations, V&V and system processes. Other tabs allow an operator to control maintenance tasks for AP, initiate additional quicklook processing or modify quality control status of input files. An example of the tab for incoming observations is shown in Figure 5. An example of the tab for incoming raw telemetry files, showing the available pulldown menu for each file, is shown in Figure 6. Many tabs are simply clients of the TT9 database and execute customizeable queries of certain tables within the database. Standardized APIs define templates for new tabs and enable rapid creation of interfaces to most aspects of DSOps systems. For example, we have recently added tabs to organize the latest reprocessing effort consisting of 5 years of Chandra data. Perl/Tk was chosen primarily to simplify integrating existing Perl scripts. We found it is a well documented toolkit with numerous CPAN packages resulting in shorter time to deployment.

2.2. V&V

A relatively new responsibility of Chandra DSOps is the V&V analysis for each observation before it is release to the user or to the public. DSOps began doing the V&V task in January, 2005. Since that time, the average time required to do V&V has been reduced significantly (see Figure 1). The software is described by Evans et al., 2006.[?] This software became available about the same time that DSOps started doing V&V, and is in part responsible for the increased efficiency.

The most common reason for rejecting a processed observation in V&V is to request the removal of one guide star from the reconstructed pointing solution that has not been tracked with sufficient accuracy. Another common reason for rejecting a processed observation is the incorrect identification of the zeroth order source for a grating spectrum, which results in the spectral extraction region not including the entire dispersed spectrum.

An important function of the V&V process is to convey information directly to the user concerning acquisition problems or unusual circumstances of the observation. This is the primary means of communicating with the user on these issues. V&V reports include information about temperatures onboard that are outside the calibrated range, known software and instrument calibration issues that will be addressed in the future, reason for less than requested time, phase, and roll angle constraints that were not met due to scheduling and spacecraft issues, and complexities of the data that may require special analysis techniques.

Most of the issues with each observation are known to DSOps before the observation gets to the V&V stage, which expedites the V&V process and makes the V&V reports to the user more complete and accurate. It is sometimes necessary to contact the spacecraft engineers, mission planners, and instrument team members to resolve issues with observations.

2.3. Special Processing for Problem Cases

A very small percentage of observations fail V&V and must be reprocessed with customized techniques. A configurable version of AP, called Special AP (SAP[?]) is used to modify parameters in the processing to obtain the best possible science data products. CXC Data Systems also maintains a database of data issues and processing issues, called the ISSUES database, that is linked to the processing status of each dataset. The processing status is web-available (http://asc.harvard.edu/soft/op/op_pst.html) to all users. Data products produced by AP and SAP are versioned in the archive to allow multiple reprocessings and partial reprocessings while maintaining a hierarchical system in the archive that delivers that latest version of each product to the user.

3. LESSONS LEARNED

3.1. Information on the status of the system and on metrics should be in a centrally located place.

Once AP was delivered to DSOps for use in data processing, we found that TT9 would be useful for more than just raw telemetry verification and quicklook processing. There was no single location providing end-to-end visibility into AP. To determine what observations were planned, awaiting AP, awaiting V&V, distributed to the user, or in an error state, we had to search for status data in several locations. Problem situations could go undetected for hours or even days if this investigation was not done frequently. We developed TT9 modules that monitor for key data products from each stage of data processing and update the operational database with completion times for each observation. Time of completion is stored providing metrics for each stage of data processing. The end-to-end status is then retrieved and displayed by the TT9 GUI for each raw telemetry file and the observations contained as they progress through the data processing. Color codes in the TT9 GUI are used to highlight error and warning conditions. The GUI allows evaluation of all processing activities at a glance at any time and interactive control, even from remote locations. The metrics are used along with monitoring mission planning products to estimate the completion time of planned observations. This is useful for time critical observations where observers request the expected time of data distribution.

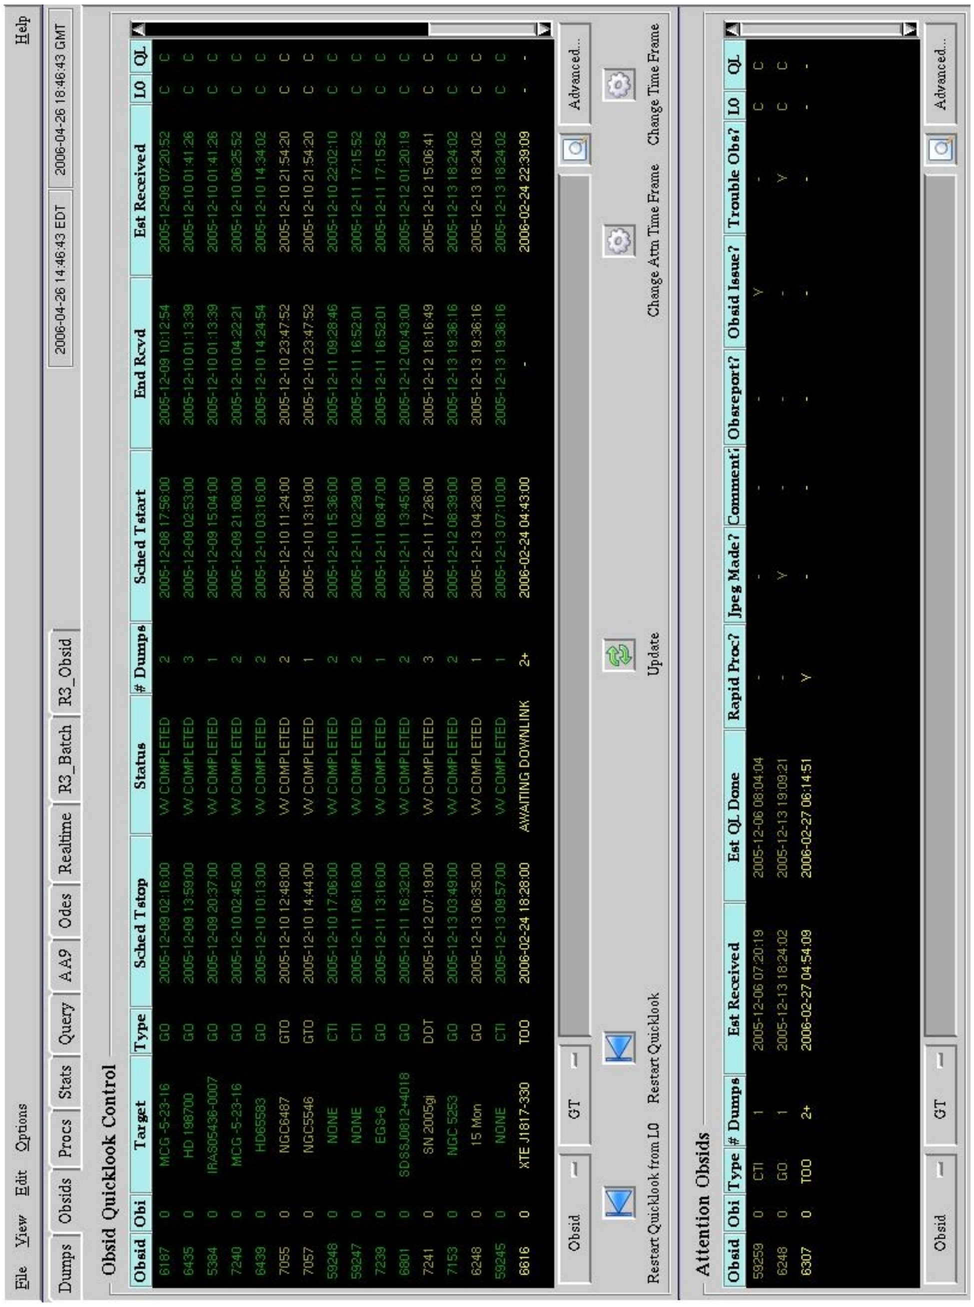

Figure 5. Observation ID (Obsid) TT9 GUI tab showing one observation per row in order of acquisition. Timestamps for each stage of data processing such as quicklook completion and submission to AP are available. Time metrics from the operations database and mission planning data are used to calculate columns such as estimated receipt of raw telemetry (Est Received). The status column tracks each observation from planned to V&V completion. There are 50 viewable and sortable columns that can be displayed in the Obsid tab.

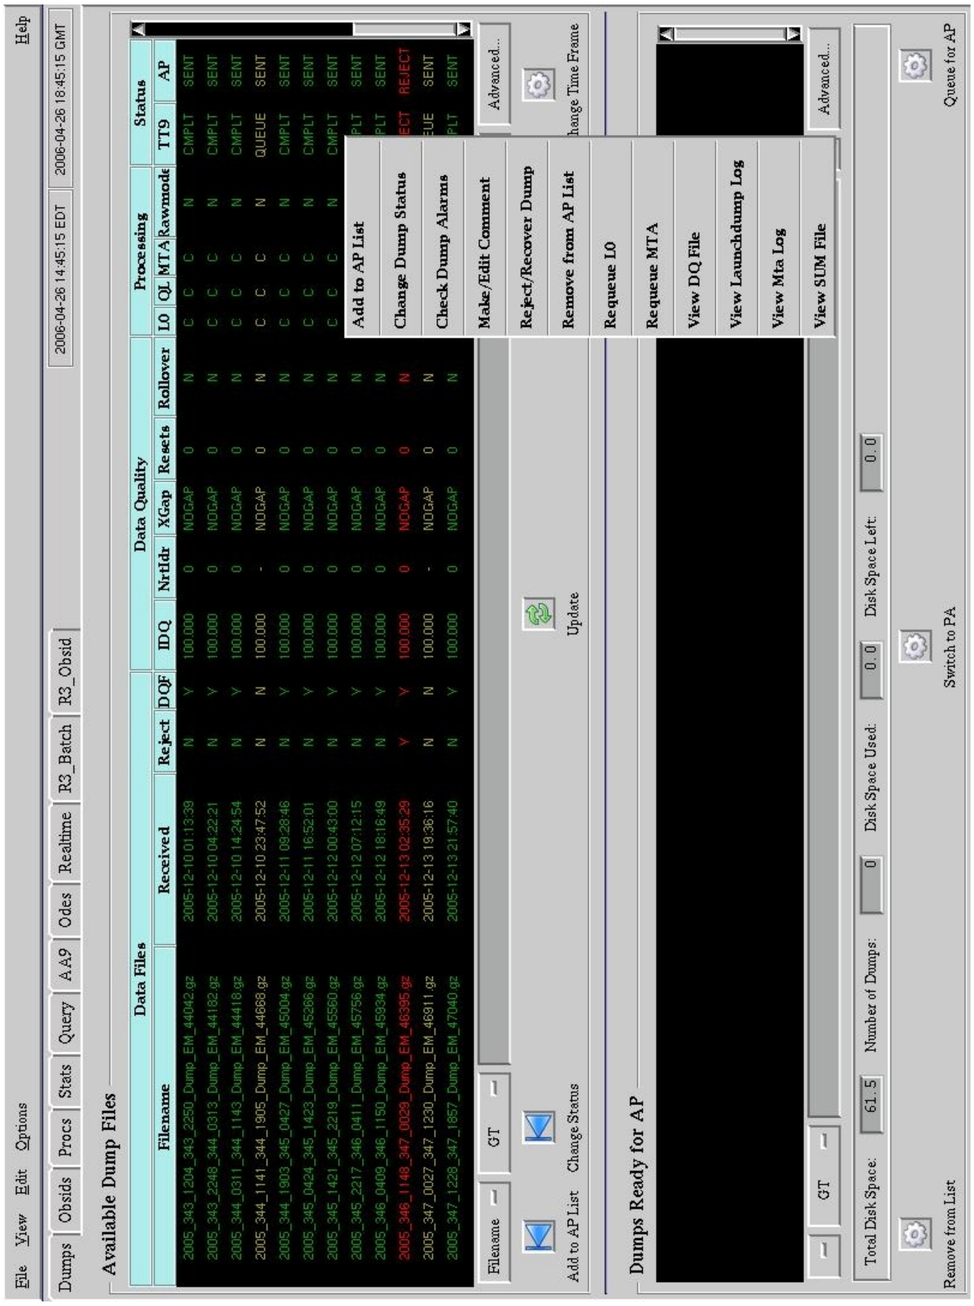

Figure 6. Raw Telemetry (Dumps) TT9 GUI tab showing each raw telemetry file in order of receipt and the data quality status of the file. The pull-down menu is available by highlighting a row in the GUI and right-clicking to view more information or take action on the selected row. Other available tabs can be seen at the top of the figure, such as the "R3_Batch" tab used to organize input of large raw telemetry batches to an instance of AP configured for reprocessing, a second instance of the Obsid tab customized for tracking observations going through reprocessing, and the AA9 tab tailored for monitoring and controlling AP maintenance. Currently, the TT9 GUI contains 16 customized tabs for DSOps.

The need for information management for data processing suggests systems solutions capable of providing a flexible and easily administered database for operations as well as personnel skilled with databases. Having this powerful tool early in a mission will encourage end-to-end visibility of operations and greatly increase efficiency. Each team should either include or have available personnel that are proficient in SQL and database application design. Such staff will be invaluable resources to the entire mission and have the skills needed to help maintain a well organized and highly productive operational group.

Lesson Learned: Centrally located, realtime information concerning status of all aspects of the system prevents many operator errors and allows full communications. Such a system should be prototyped prior to launch. Flexible, easily maintained databases are essential for organizing and tracking the large amounts of information that will hopefully flow through any mission data processing group. Having the database tools in place early in the mission would have saved Chandra DSOps considerable effort in back filling the database.

3.2. Proactive rather than reactive

We quickly found that we needed to be fully informed about spacecraft and instrument issues as they occurred. For example, if the spacecraft autonomously safes the instruments due to high radiation during an observation, it is important for us to know about that before we process so we can determine if any science data were collected, which dictates how the observation is processed. Also, if the first onboard Solid State Recorder fills up and there is an autonomous switch to the second SSR, the data will be delivered to us out of time order as it is retrieved from each SSR. We then are alerted to the need to reorder the data dumps into time order for submission to AP. There are several other operational issues that can affect how the data are processed, including unusual sequences in the load.

Lesson Learned: It is better to be in a proactive role and be prepared for unusual processing situations than to discover them empirically. We have ensured that we will know about such issues beforehand by maintaining communication with mission planning, spacecraft operations, and ground operations. In general simply being included on e-mail lists for these groups and also attending some meetings is sufficient. Co-location with these groups has greatly facilitated the information flow.

3.3. A flexible front-end to the standard data processing software system facilitates post-launch adjustments.

Requirements for the software were firmly in place before launch and were based on the best information available at the time. However, as with all missions, some pre-launch assumptions are not appropriate when the mission flies. In the case of Chandra, the pre-launch assumption at data from the spacecraft would always arrive from the Ground System in time order and be the best quality data obtainable proved to be incorrect on a few occasions. Sometimes a better version of the data becomes available after the first delivery because (1) the data are re-dumped from the spacecraft, (2) the second channel at JPL DSN is brought in and is of higher quality or (3) the Ground System is able to fill a data gap with data from a realtime data stream. As noted above, if the first SSR fills up and data are recorded on the second SSR, the data will arrive out of order. We realized that we needed an automated method of sorting the incoming data into time order and detecting when we needed to wait for missing data before submitting the raw telemetry to AP. TT9 was enhanced to determine if data are missing between raw telemetry files and hold submissions until all data are available. TT9 also identifies resent data and rejects the original version. These capabilities were critical to the goal of full automation of the data processing.

Lesson Learned: The lesson learned is that pre-launch assumptions necessary for design of the processing software should be understood to be assumptions. A flexible operational interface can accommodate these unexpected situations.

3.4. Scientifically motivated and creative personnel prove to be very productive and more permanent than entry-level IT personnel for data processing

Data processing operators at many missions are often selected for their computer expertise, and are generally rather low salary personnel. However, it is a perennial problem for missions that these operators tend to stay with the project for only a short time, resulting in large turnover and significant training time for new personnel.

Most of our data processing operators were selected based on an Astronomy or Astrophysics degree with very high grade point average. Most have intentions of continuing to graduate school in Astrophysics or another technical field in the future. Their interest in the science that is being done with Chandra and the chance to work with Chandra data have served to keep them interested and motivated in a job that is often is viewed as unchallenging and repetitive for IT personnel. Strong science majors have plenty of good software skills to develop needed scripts, databases, and software design, but they also have the opportunity to learn more about their chosen field and enhance their prospects for graduate school. In fact, 80% of data processing operators at Chandra have been there since launch.

Lesson Learned: Personnel with an interest in the science goals of the mission are more dedicated and motivated than IT personnel.

3.5. ISSUES database

Categorizing data and processing issues and organizing the information in a publicly available database has been a benefit to productive user interfaces. Issues are normally recognized during data quality screening and in the data processing system. Each issue is given a number and a short explanation. Each observation that is affected by the issue is then linked to the issue in the database.

Lesson Learned: The implementation of web-based access to processing status and associated information on data and processing issues has simplified enormously the user interaction for Chandra Data Systems. The number of queries from users was reduced greatly as soon as this web-based reporting was implemented. The ISSUES database is also an important source of metrics that are provided to NASA.

3.6. Operational records

We have kept a file for each observation since launch that has anything unusual in its acquisition or data. These records have been invaluable for mass reprocessing of the data because no one can remember every unusual or anomalous situation. As it turns out, no other group has keep such complete records. Some groups collect information about acquisition problems, but data problems as they affect AP are not recorded uniformly by any other group. Data processing operations should take this responsibility and not rely on any other group to do so. In the first years of the mission, I assumed other groups kept such information, but that was later revealed to be wrong. These observations reports are used by V&V, so that the information ultimately ends up in the archive.

Lesson Learned: The lesson learned is that data processing should keep full records of any unusual information about each obsid. Don't expect others to do that.

3.7. Make generous estimates of hardware requirements and have the resources in place at launch

We had at least one occasion when standard data processing of current data from the satellite fell behind solely due to the lack of sufficient disk space. Estimates of necessary hardware resources should not only include expected usage per observation, but also include generous margins for unexpected situations and growth of software system. It takes some time to purchase and receive new hardware, and it normally doesn't arrive in time to get out of a tight spot.

Lesson Learned: Purchase hardware with a significant margin beyond the calculated average needs. Hardware should accommodate the peak rate of processing, not the average.

3.8. V&V

It is very efficient for the staff members who are responsible for quicklook processing and AP processing to do V&V, because information learned about each raw telemetry file and each observation is often directly applicable to the V&V analysis. Because DSOps personnel do V&V, they have normally already learned about anything unusual concerning an observation. Therefore, V&V analysis goes very quickly.

Lesson Learned: It is most efficient for staff members who normally investigate problems with data, observation acquisition, scheduling, and processing to also perform V&V analysis. This approach eliminates duplicate independent investigations and facilitates to the delivery of the data to the user.

4. CHALLENGES MET

4.1. Improvement in data distribution time and staff reduction

To achieve the goal of full automation, several problems had to be solved.

- script standard tasks into one-line commands
- automate quicklook processing
- develop algorithms for time-ordering incoming data dumps, detecting when a dump is missing, and evaluating data quality
- develop notification system for operators
- develop operational GUI to allow full visibility into system and to serve as an interactive interface with TT9 system
- develop database to support GUI and maintain searchable information
- design method of determining state of AP system
- develop method of queuing and executing tasks based on status of AP system

We developed a simple system at first, writing scripts to kick off sequential pipelines automatically from a single command line. This system allowed us to process data with exactly the same tools as the archival processing system, producing quicklook data through Level 1,[?] or Level 2 if required. The use of a single command line system greatly reduced operator error, which can be very costly in resources, and later enabled full automation of quicklook processing. The system also provided a means of detecting data with corruption or other problems that would prevent the processing in AP. Complete automation was achieved with a system of scripts that monitors the state of the processing (idle, processing, error) of AP and kicks off various functions based on the detected state of the system. For example, when AP is idle, processed data are ingested into the archive and disk cleaning activities are initiated. The next data dump is then sent for processing. The automation system is controlled with the operational GUI, allowing an operator to include or exclude any of the tasks from automatic execution via a simple right click. This is seldom necessary. The introduction of the processing automation significantly reduced the elapsed time between the end of the observation and data distribution to the user. The personnel requirements have also been reduced due to the automation.

4.2. Notifications

To reduce the amount of time needed to actively monitor the processing, both for quicklook processing and AP processing, pager alerts have been set up to inform personnel of situations that need immediate attention to keep the system running. Pager alerts are activated if raw telemetry fails Level 0[?] processing. In this case, the data may be corrupted and immediate response can allow a re-dump of the data from the satellite on the next pass. Pager alerts can also be activated for specific observations, informing the operators that an observation has finished Level 2 processing and is ready for V&V. This type of alert is used for fast processing requests approved by the project for time-critical data distribution. Also, any interruption in the flow of realtime data causes a pager alert to minimize any loss of realtime data flow to other groups in CXC. The notification system was critical to achieving automation with limited staff.

5. CONCLUSION

The Chandra Data Processing Operations has reduced the time required to deliver science data to the users from an average of 13 days in 2000 to an average of a little more than 1 day in 2005. The improvement is due to several factors, including the implementation of a front-end system of data quality checks, quicklook production, and automatic control of the data submission and archiving process. The front-end system developed to achieve the automation, coupled with a database of observation-specific information and an interactive GUI for system visibility and control, is potentially adaptable to other missions due to a simple design and minimal interaction with mission-specific files. V&V is done by DSOps personnel who are familiar with any acquisition problems and data problems for each observation.

5.1. Acknowledgments

We acknowledge the dedicated and expert mission-long contributions of Jennifer Lauer and Beth Sundheim to the Chandra Data Processing Operations group. Support for Chandra data processing is provided by the National Aeronautics and Space Administration through the Chandra X-ray Center, which is operated by the Smithsonian Astrophysical Observatory under NASA contract NAS 8-03060.

REFERENCES

1. D. Plummer, J. Grier, and S. Masters, "The Chandra Automated Processing System: Challenges, Design Enhancements, and Lessons Learned," in *(this proceedings)*, 2006.
2. I. Evans, M. Cresitello-Dittmar, S. Doe, J. Evans, G. Fabbiano, G. Germain, K. Glotfelty, D. Plummer, and P. Zografou, "The Chandra X-ray Observatory Data Processing System," in *(this proceedings)*, 2006.