Astronomical Data Analysis Software and Systems XII
ASP Conference Series, Vol. 295, 2003
H. E. Payne, R. I. Jedrzejewski, and R. N. Hook, eds.

Chandra Monitoring, Trending, and Response

Scott J. Wolk, Bradley D. Spitzbart and Takashi Isobe

Harvard-Smithsonian Center for Astrophysics, 60 Garden St., Cambridge, MA 02138

Abstract. The *Chandra* X-ray Observatory was launched in July, 1999 and has yielded extraordinary scientific results. As part of *Chandra's* Science Operations Team, the primary goal of Monitoring and Trends Analysis (MTA) is to provide tools for effective decision making leading to the most efficient production of quality science output from the observatory. MTA tools, products, and services include real-time monitoring and alert generation for the most mission critical components, long term trending of all spacecraft systems, detailed analysis of various subsystems for life expectancy or anomaly resolution, and the creation and maintenance of a large SQL database of relevant information. This is accomplished through the use of a wide variety of input data sources and flexible, accessible programming and analysis techniques.

1. Introduction

The Monitoring and Trends Analysis (MTA) subdivision within the science operations team (SOT) of the *Chandra X-ray Center* (CXC) is charged with providing an overview of telescope performance as it affects the science quality and efficiency of the observatory. The group often serves as a clearinghouse of data and analysis tools for *Chandra* with the engineers, instrument experts, and calibration scientists. The MTA tasks make fervent use of the World Wide Web. We maintain thousands of dynamic web pages as well as a similar number updated on a daily basis. All the data and many on-line tools can be accessed through our home page `http://cxc.harvard.edu/mta/sot.html` which has links to all of the products described below, as well as to many other MTA and CXC resources.

2. Inputs

The MTA system is designed to use a variety of interchangeable inputs. New data arrives at the CXC approximately every eight hours. Data are stored onboard and dumped during ground supports, roughly three times a day. During real-time contacts data feeds are sent directly to the CXC. Dumped data generally arrives within a few hours.

For raw telemetry decommutation we use ACORN (Wolk et al. 2000). ACORN is capable of reading from both the real-time telemetry stream or from several types of archived dump files. *Chandra* telemetry is coded in over 11,000 MSIDs (mnemonic string identifiers). Each spacecraft meter, sensor, thermistor, boolean value, etc. can be identified and tracked with an unique MSID. ACORN decodes the telemetry stream and provides times, MSIDs, and values to either the standard output or to a tab-delimited file, for other tools to use as input.

Input data is also obtained from CXC standard processing pipeline products. We frequently use files from all levels of available processing (see Plummer et al. (2001) for a full description of pipeline data products.) All of these standard products are easily accessible from the *Chandra* data archive in FITS format.

In addition to *Chandra* data, the MTA system gathers data from outside sources, most notably the NOAA GOES and ACE missions, using lynx and anonymous ftp commands.

3. Processes

The standard MTA data processing pipeline is run as part of standard *Chandra* automated processing (Plummer et al. 2001). For more customized applications, the focus of the current effort is to create simple data products which can be massaged to allow intuitive visualization. Our programming tools of choice are UNIX shell scripts, Perl, IDL, and HTML. These often are used to write wrapper tools around *Chandra's* suite of data analysis programs (CIAO[1]). Programming is approached with the intent of eventual automation. We rely heavily on the UNIX time daemon (cron) to autonomously run jobs at various times of the day and night, updating data files and web pages and monitoring processing status and telemetry. MTA's crontab consists of over 50 periodic tasks. Jobs are divided among three UNIX machines, all running Solaris 5.8. The main real-time analysis and standard processing occurs on a Sun Ultra10/440, with a completely independent real-time data flow on an UltraE450 for redundancy. A separate UltraE450 handles daily tasks and individual cron jobs.

4. Outputs

MTA data and analyses are provided to the community in three forms: Time- and mission-critical alerts are sent via e-mail to pagers, standard and custom presentations are posted on the world wide web, and all monitored values are archived in a database.

4.1. Alerts

We have created a number of e-mail aliases to which alert messages can be sent when spacecraft state violations or other problems are detected. Data are monitored in real-time and dump data is processed on receipt in a near real-time mode.

Real-time Alerts During each real-time support, MTA runs Perl scripts which create a dynamic web page known as the *Chandra* Snapshot (see Sect. 4.2.). The

---

[1] `http://cxc.harvard.edu/ciao/documents.html`

Perl code incorporates selected limit checks to color code the display, indicating any state violations. In addition, these limit checks will generate alerts if certain persistent conditions are found. Once a message is sent, a semaphore is created which prevents further alerts for the same violation. This semaphore is autonomously removed when the condition subsides for three minutes. Similar alerts are triggered if limits are exceeded on other spacecraft which monitor the radiation environment.

Near Real-time Alerts We have developed a separate customizable Perl-based package called *config_mon*. This software acts on the spacecraft playback data when it arrives at the CXC (a few hours after the completion of each communications pass). This data contains the record of spacecraft state for the time period since the previous data dump. The values are reviewed and compared against as-planned values and operational limits using output products from mission planning and a limits database. When violations are found, alerts are sent. *Config_mon* currently monitors science instrument position, focus position, pointing, gratings positions, wheel rates and particular temperatures of concern.

4.2. World Wide Web

Our main vehicle for data dissemination is the world wide web. We maintain a large suite of dynamic web pages presenting real-time data feeds, standard processing displays, customized studies, and weekly and monthly reports. To the extent possible these pages are updated automatically. We are also experimenting with emerging WAP (Wireless Application Protocol). Spitzbart et al. (2003) have a complete report on this aspect of the project.

Real-time Web Pages Real-time data is viewed through a variety of web pages. Each one is run using a dedicated ACORN feed and underlying Perl code to format the ASCII output and color code particular items of interest. The *Chandra* Snapshot[2] provides easy access to the most relevant information from the current telemetry. Other real-time displays cover over 1000 additional MSIDs and data from the science instruments.

Standard Web Pages Spacecraft subsystem monitoring pages are produced each day as part of the standard data processing pipeline. Plots and statistics are displayed for each mnemonic and values are highlighted according to a green-yellow-red color scheme. These plots are reviewed daily by the SOT and a summary of violations or other concerns are reported to the project each week. We provide quick-look images and statistics of all observations. Certain calibration observations are further processed in specialized pipelines. The flexibility of MTA tools and data allows for the timely creation and presentation of customized studies as called for by various teams in response to current spacecraft needs or anomalies. These have included details on the radiation environment, spacecraft mechanisms and instrument performance.

---

[2] `http://cxc.harvard.edu/cgi-gen/mta/Snap/snap.cgi`

4.3. MTA Databases

At the end of the standard data processing pipeline, a five minute average and standard deviation is computed for each monitored MSID. This is ingested into an SQL database. Currently there are eight databases and 43 individual tables, divided by subsystem (Wolk et al. 2002). The DataSeeker (Overbeck et al. 2002) is used to extract and merge tables from the MTA databases. This tool is available with either a web interface or command line mode, which makes it convenient for first-time users or incorporation into automated scripts. DataSeeker seamlessly merges data keying on time. This allows users to cross-correlate data to find trends relating to temperatures, attitude, power consumption, etc. in addition to temporal trends. Another important feature of the DataSeeker is the ability to incorporate non-SQL tables. Easily generated RDB files can be merged with existing SQL database tables. This has proven valuable for rapid implementation of new tables for which the need had not been foreseen.

In practice, we call the DataSeeker via an automated trending script. The script provides plots and statistics for all the monitored MSIDs. The system attempts to predict the next six months' behavior by performing simple fitting to data and extrapolating. Past and predicted future limit violations are highlighted.

5. Conclusions

The main lesson learned from the MTA experience is that simple, uniform access to data is paramount. Unfortunately today's spacecraft and instruments, with their programmable telemetry, do not lend themselves to uniformity. What we have done is to impose uniformity on the ground data such that separate systems can be analyzed and displayed as a unit using fairly simple scripts. The next phases feature expanded databases which will include more higher level data products as well as more sophisticated trending tools. We will continue to respond to and attempt to anticipate spacecraft issues.

References

Overbeck, R.S. et al. 2002, in ASP Conf. Ser., Vol. 281, Astronomical Data Analysis Software and Systems XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP), 449

Plummer, D.A. et al. 2001, in ASP Conf. Ser., Vol. 238, Astronomical Data Analysis Software and Systems X, ed. F. R. Harnden, Jr., F. A. Primini, & H. E. Payne (San Francisco: ASP), 475

Spitzbart, B.D. et al. 2003, this volume, 162

Wolk, S.J. et al. 2000, in ASP Conf. Ser., Vol. 216, Astronomical Data Analysis Software and Systems IX, ed. N. Manset, C. Veillet, & D. Crabtree (San Francisco: ASP), 453

Wolk, S.J. et al. 2002, in ASP Conf. Ser., Vol. 281, Astronomical Data Analysis Software and Systems XI, ed. D. A. Bohlender, D. Durand, & T. H. Handley (San Francisco: ASP), 341