

The Operation and Evolution of the *Swift* X-ray Telescope

Jamie A. Kennea^{*a}, D. N. Burrows^a, C. Pagani^c, J. E. Hill^d, J. L. Racusin^a, D. C. Morris^a, A. F. Abbey^b, A. P. Beardmore^b, S. Campana^c, G. Chincarini^c, G. Cusumano^e, P. A. Evans^b, N. Gehrels^d, O. Godet^b, T. Mineo^e, V. La Parola^e, V. Mangano^e, A. Moretti^c, J. A. Nousek^a, J. P. Osborne^b, K. L. Page^b, M. Perri^f, R. L. C. Starling^b, G. Tagliaferri^c, F. Tamburelli^f, and A. A. Wells^b.

^a Pennsylvania State University, 525 Davey Lab, University Park, PA 16802, USA

^b University of Leicester, University Road, Leicester, LE1 7RH, UK

^c INAF-Osservatorio Astronomico di Brera, Via Bianchi 46, 23807 Merate, LC, Italy

^d NASA-GSFC, Greenbelt, MD 20771, USA

^e INAF-IASF, Via U. La Malfa 153, 90146 Palermo, Italy

^f ASI-ASDC, Via G. Galilei, I-00044 Frascati, Italy

ABSTRACT

The *Swift* X-ray Telescope (XRT) is a CCD based X-ray telescope designed for localization, spectroscopy and long term light curve monitoring of Gamma-Ray Bursts and their X-ray afterglows. Since the launch of *Swift* in November 2004, the XRT has undergone significant evolution in the way it is operated. Shortly after launch there was a failure of the thermo-electric cooler on the XRT CCD, which led to the XRT team being required to devise a method of keeping the XRT CCD temperature below 50C utilizing only passive cooling by minimizing the exposure of the XRT radiator to the Earth. We present in this paper an update on how the modeling of this passive cooling method has improved in first ~1000 days since the method was devised, and the success rate of this method in day-to-day planning. We also discuss the changes to the operational modes and onboard software of the XRT. These changes include improved rapid data product generation in order to improve speed of rapid Gamma-Ray Burst response and localization to the community; changes to the way XRT observation modes are chosen in order to better fine tune data acquisition to a particular science goal; reduction of “mode switching” caused by the contamination of the CCD by Earth light or high temperature effects.

Keywords: *Swift*, XRT, Thermal, Passive Cooling, Software, X-ray, Telescope.

1. INTRODUCTION

The *Swift* Satellite¹, launched from Cape Canaveral Air Force Station on November 20th, 2004, was designed with the purpose of detection and rapid follow up of Gamma Ray Bursts (GRB). GRBs are bright transient astrophysical phenomena usually associated with the collapse of a massive star into a black hole or collision/merger of 2 neutron stars in a binary system. These GRBs require rapid follow-up in X-ray and optical due to their rapid power-law decay in both X-ray and optical/UV wavelengths.

In order to achieve both rapid detection and follow up *Swift* consists of 3 complementary instruments and a spacecraft platform capable of rapid slewing in order to get on target in usually less than 2 minutes after initial detection of a GRB. The Burst Alert Telescope² (BAT) performs this initial detection. BAT is a wide field of view Gamma-Ray/Hard X-ray telescope utilizing coded mask optics imaging sensitive in the 15-150 keV energy range. For accurate burst localization and follow-up *Swift* has two Narrow Field Instruments (NFIs): the X-ray Telescope³ (XRT), a Wolter type I grazing incidence X-ray telescope Ritchie Chrétien telescope with a micro-channel plate based photon counting detector with lenticular filters in optical and UV wavelength as well as Optical and UV Grisms for spectroscopy⁴.

The *Swift* XRT is a sensitive (135 cm² effective area), 24 arc-minute field of view X-ray telescope utilizing a 3.5 meter focal length grazing incidence Wolter type I telescope to focus X-rays onto an CCD detector. The XRT prime science objective is to perform rapid localization of GRB afterglows with an uncertainty of ~5 arc seconds, in order to improve the initial position found by the BAT which has an uncertainty of 3 arc minutes. XRT also has spectroscopic capabilities in order to characterize non-thermal emission and possible line spectra from GRBs. The high sensitivity of XRT allows for GRB afterglow follow-up for analysis of long-term GRB afterglow light-curves often for weeks after the initial GRB detection.

Since launch the XRT has suffered two hardware anomalies. During the instrument activation phase, the power supply for the CCD thermoelectric cooler (TEC) suffered a failure, rendering the TEC inoperable. The TEC was required to keep the XRT CCD at its designed operating temperature of -100C. A second anomaly occurred in May 2005 when the XRT CCD was apparently struck by a micrometeorite, damaging the CCD causing the formation of several hot columns. For more details on this micrometeorite hit and its impact to XRT operation please refer to Abbey et al. (2005)⁵.

In this paper we discuss the evolution of the operation and on-board flight software of the Swift XRT in the light of the hardware issues, optimization of the XRT to better meet the science goals, and the changes required to meet the growing and changing role of the Swift observatory from a primarily GRB centered mission to a more general platform for multi-wavelength astronomy, both for long term monitoring and for rapid follow-up of transient events. We also update the status of the XRT CCD temperature control mechanism using passive cooling⁶, utilizing data from the first 1000 days of XRT operation.

2. CONTROLLING THE XRT CCD TEMPERATURE

2.1 UPDATED TEMPERATURE MODEL

The loss of the XRT TEC meant that the XRT CCD temperature would never reach the designed temperature of -100C. Fortunately it was determined that the XRT CCD performed exceptionally well at temperatures up to 50C. Above 50C the CCD became dominated by hot pixels and thermally dependant bias and dark current levels, making it effectively useless. It was seen during the early part of the activation of Swift that the XRT CCD temperature varied from around -70C to -45C, with rapid temperature changes of ~5-10 C in a single 96 minute orbit often observed. It was determined that these temperature variations were related to the angle of the XRT radiator relative to a line connecting *Swift* and the center of the Earth. In essence, when the XRT radiator was pointing into space, the XRT was cooler, and when the XRT radiator was pointing at the Earth, XRT was hotter. After this was determined, a method was devised to keep the XRT CCD cool by maximizing through mission planning the average angle between the radiator and the Earth. A model was devised to predict the XRT CCD temperature for a set of pre-planned targets. The initial results of this method of passive cooling were previously reported⁵.

Data from the first 1000 days of Swift operations have confirmed conclusively that the XRT CCD temperature is directly related to the so called "Earth Elevation Angle" (EEA), which is defined as the angle between the plane of the XRT radiator surface and the center of the Earth. The correlation between the day averaged value of the temperature and the day-averaged value of the EEA is shown in Figure 1. The updated equation used to derive a prediction of XRT temperature is as follows:

$$T_{CCD} = 56.62 + 0.734E + 0.006E^2 \text{ } ^\circ\text{C} \quad (1)$$

Where E is the orbit-averaged value of the Earth elevation angle. Note that this equation differs from the value given by the previous analysis in that it does not contain the component dependent on the β -angle (the angle between the current *Swift* orbital plane and the Sun) that was previously derived. The removal of this dependence was a result of better understanding of the effect after further thermal data from XRT showed that the β -angle dependence was actually a transient effect. Figure 2 shows a plot of the residuals found when subtracting the model-calculated temperature from the measured CCD temperature averaged over a 24-hour period for the first 1000 days of XRT operation. As can be seen there are periods in which the residual appears to be strongly anti-correlated with the orbital β -angle. These periods occur during summer in the North Hemisphere, however attempts to model this variation have not yielded any good results. This summer β -angle dependence is most likely caused by a higher Earth albedo during the summer. Although these excursions are roughly correlated with being in the summer, it is not possible to include an accurate estimate of these effects in the current model, however through monitoring the current temperature residuals, the planners can predict the onset of this effect, as well as being aware of approximate the times of year where this effect is a problem.

Controlling of the XRT CCD Temperature using passive cooling methods combined with careful science planning using targets optimized for XRT cooling has been highly successful with ~92% of the XRT operational time with a average CCD temperature of < 54C, and 97% below 52C.




Figure 1. An updated model for estimating the XRT CCD temperature as a function of Earth Elevation Angle.




Figure 2. Residuals from the model show in Figure 1 (black) and the β -angle (red) for the first 1000 days of XRT operation.

3. IMPROVING RAPID BURST RESPONSE

3.1 The problem of finding a rapid X-ray GRB afterglow position

The XRT provides rapid localization of the X-ray afterglow to the community by taking an image of the field and performing on-board centroiding of the brightest X-ray source in the imaged field of view. This is achieved by taking a series of short Image Mode exposures, which consist of raw CCD exposures with no event recognition performed, on which a centroiding algorithm is run on-board to attempt to find the brightest source in the field of view. The sequence of Image Mode data taken consists of one 0.1s followed by up to two 2.5s exposures if no source is found. The results of this centroiding are downlinked rapidly to the ground through the use of Tracking and Data Relay Satellite System (TDRSS) network of satellites and distributed to the astronomical community via The Gamma-ray Bursts Coordinates Network⁸ (GCN) within seconds of the localization occurring. The rapidly downlinked data include a "Position Message" containing the RA, Dec and positional error of the bright source and a "Postage Stamp", which is a 64x64 pixel image centered on the source, that can be used to determine if the source is a real astrophysical object or a detector artifact such as a cosmic ray event.

Unfortunately the X-ray brightness of GRBs are often too faint to be localized on-board using Image Mode data (for an explanation of XRT observing modes see Hill et al, 2005⁷) either because the GRB X-ray afterglow is initially faint, or because the X-ray afterglow fades too quickly to be caught in a bright state at the time Image Mode data is taken. The inability to find a rapid position for a GRB has led to the reporting of GRB afterglow positions to be delayed by in the worst case up to 8-10 hours, depending on the time until the next ground station downlink, and the speed at which the data can be processed and analyzed.

In this section we discuss the changes made to the XRT on-board software that add new data products downlinked through TDRSS that allows the much more rapid localization of a GRB.

3.2 "THRESHPIX" data

In order to increase the likelihood of detecting the presence of a GRB in the XRT Image Mode data, "THRESHPIX" data were included in the standard XRT data products sent out immediately through TDRSS. Implemented in January 2006 via a change to the XRT flight software, these new THRESHPIX data products consist of a list of pixels coordinates and their DN values. Only pixels with DN values above a specified threshold level are telemetered, and the threshold level being set so as to minimize the amount of background noise downlinked, as TDRSS telemetry needs to be kept to a minimum level at all times.

Upon being telemetered to the ground these pixel lists are then reconstructed into a full 600x600 image of the XRT field of view with an astronomical coordinate system attached, and sent out as FITS files through the GCN system. These reconstructed THRESHPIX images can be visually inspected for the presence of any bright X-ray afterglow of the GRB. Typically if the XRT is unable to centroid it takes a sequence of 3 images, one 0.1s and 2 2.5s exposures. With the THRESHPIX data all 3 of these images are downlinked, allowing the data analyst to combine the images to maximize likelihood of finding the GRB X-ray afterglow. This data mode also acts as a vital diagnostic tool for cases when the XRT on-board centroiding algorithm centers on a source in the field of view that is clearly not an astrophysical source (e.g. a cosmic ray).

With normal Image Mode data, it is required that GRB X-ray afterglow brightness is a minimum of ~8 counts per second in order for the on-board software to be able to successfully find a position. These THRESHPIX data allow for the XRT team to localize, using ground analysis, moderately bright X-ray afterglows for which the brightness is not high enough to be found using the centroid algorithm (e.g. ~2-8 counts per second). Typically utilizing these THRESHPIX data the XRT team has been able to report a prompt XRT afterglow position approximately 5 minutes after a settling on a GRB, greatly increasing the response time compared to previously having to wait until a full data downlink through Malindi.

3.3 "SPER" data

Some GRB X-ray afterglows proved to be too faint even to be detected in the ~5.1s of Image Mode data downlinked through TDRSS. These faint bursts are often some of the most scientifically critical for which to report a rapid position, for example in the case of a short GRB, where the source is often faint and decays extremely rapidly. In this case a rapid position from XRT may be necessary for getting a good optical image from a ground based telescope, and possibly a red-shift measurement to estimate the distance to the GRB.

In order to obtain a rapid position it was decided to downlink the Photon Counting (PC) mode data through TDRSS. Photon Counting is the most commonly utilized mode of XRT observation for faint sources. This mode consists of a series of 2.5s CCD exposures, which are searched for X-ray events. Each event consists of a 3x3 array of DN values and events are assigned a grade based on the distribution of DN across this 3x3 array (for a full description of PC mode see Hill et al, 2005⁷).

PC mode events are normally telemetered in the science data downlinked through the Malindi ground station, located on the East Coast of Kenya. At the best the gap between these Malindi passes will be a single 96 minute orbit, although longer gaps of 3 hours and up to 8 hours (referred to as the "Malindi Gap") occur daily due to the orbital inclination. PC mode events can be reconstructed into images, light-curves and spectra of X-ray sources.

Although downlinking PC mode data through TDRSS instead of Malindi provides near-instant X-ray data, there are significant limitations on the available bandwidth as the maximum allowed TDRSS downlink bandwidth for *Swift* is 256 bytes per second. In order to downlink PC data through TDRSS, the PC mode data production rate needs to be significantly smaller than normal. The requirement to minimize telemetry is further complicated by the presence of Earth Limb contamination, where reflected light from the Earth shines on the XRT CCD, which can cause significantly higher count rates in PC mode also, with count rates up going up to 1-2000 counts per second in some cases. It is clear that it would not be feasible to telemeter raw PC mode data through TDRSS, without significantly delaying the arrival time of other TDRSS products generated by the other instruments on-board *Swift* that are being produced in the aftermath of a GRB detection.

In order to reduce the PC mode telemetry to be suitable for TDRSS downlinking, the following filtering was performed: Firstly, the downlinked PC mode events were limited to the central 200x200 pixels of the chip, representing an approximately 8 arc minute field of view. Given that the error on BAT localization is typically 3 arc minutes, and that the error in the slew is of the order of 1-2 arc minutes, this means that the majority of GRBs will land inside of this 200x200 window. However this window only represents 1/9th of the total area of the CCD so is significant reduction in the amount of data telemetered, especially reducing the Earth limb contamination that is primarily at the edges of the CCD.

Earth limb however can still be a problem, as particularly bad cases can create low energy events across the entire field of view. In order to reduce the number of these events telemetered, two steps are taken. Firstly all events below 200DN (approximately 0.5 keV) are not telemetered. Secondly only events that are entirely contained within a single pixel (known as "Grade 0") are telemetered. This works as an effective method of removing any remaining Earth limb events (which usually have high grades due to not being real X-ray events), although this does have the effect of throwing out approximately 10-15% of "good" X-ray events. However the resultant reduction in telemetry and high signal to noise ratio achieved by this filtering make for extremely high data quality. This data packet is referred to as the "Single Pixel Event Report" (SPER), due to the fact that only single pixel events are telemetered.

The final step in reducing the telemetry size is to remove empty frame headers, which are generated when no X-ray events are found within a PC mode frame. As PC mode frames can sometimes be dropped, or the XRT can switch in WT mode if a burst undergoes bright flaring activity that requires a faster readout mode to avoid pile-up, these empty frame headers are required to accurately calculate the exposure time for a combined PC mode dataset. As the primary requirement for SPER data are to localize the burst, an accurate exposure time is not necessary, however it does mean that SPER data are not useable to calculate accurate light-curves, spectra or flux measurements for GRBs.

Figure 3 shows a typical light-curve for a moderately bright GRB, and the sequence of events are marked on the plot. In this case the XRT took 114 seconds to settle on the BAT determined position and to start taking image mode data. In this case, the XRT did not centroid and began taking WT mode data, as the burst was brighter than 2 counts/second. At T₀+135s the afterglow had faded enough for the XRT to transition into PC mode, at this point SPER data starts to be downlinked through TDRSS. At T₀+450s a watchdog timer is triggered at the GCN, at which point the combined SPER packets are processed into a fits file and sent to the XRT team via a GCN notice. This fits file can then be run through a centroiding algorithm in order to find a localization of the GRB (inset in Figure 3). The XRT team member can then send a manual GCN XRT position alert out based on the position found from the SPER data. Typically this happens approximately 10 minutes after the burst occurs, although this is now being automated allowing for an even quicker response. SPER data are continually downlinked throughout the first orbit of observations only, allowing the position to be refined.




Figure 3. XRT Light-curve of GRB070808. The vertical dashed line represents the time at which the XRT took a 2.5s image of the field and attempted to produce a rapid position through centroid. The points on the light-curve represent the X-ray afterglow brightness in data in WT mode (circles) and PC mode (triangles). The dot-dash box represents the time period in which the first batch of SPER data are being taken. (Inset) Image formed from SPER data on GRB 070808. The error circle plotted on this image is the calculated centroid position and 90% confidence level error circle calculated from the data taking the location of hot columns into account.

The development of these SPER data has resulted XRT being able to determine X-ray GRB afterglow positions at approximately $\sim T_0 + 10$ min for almost all bursts in which an immediate slew is triggered by BAT. Also in cases where the slew is delayed by a Spacecraft constraint, the sensitivity of the SPER data allows positions to usually be found within 10 minutes of the slew beginning, despite the burst not being observable for up to an hour after the initial detection. Also as these PC mode data typically has more counts in a source than the Image Mode data, positions determined with this method typically can have errors as small as 3.5 arc seconds (90% confidence) compared to ~ 5 arc seconds for on-board centroiding. See <http://www.swift.ac.uk/sper.php> for a description of the accuracy achieved by this method.

4. SELECTION OF XRT MODES

4.1 Mode Switching

In its regular "Auto" state, the XRT picks one of four observing modes: PC, WT, Low-rate Photodiode (LrPD) and Piled-up Photodiode (PuPD). These four modes respectively have increasingly faster read out times, although with faster read out modes comes the loss of imaging capability (WT mode only has 1-D imaging capability, LrPD and PuPD have no imaging at all). In order to decide which mode to use, the XRT measures the count rate. This is done primarily to avoid pile up, although also the faster read outs allow for better timing resolution for bright sources. Note that since the

micrometeorite damage to the CCD, LrPD and PUPD modes have been disabled due to the high bias caused by the hot columns, so as of May 2005, only PC and WT modes are current used by XRT.

Unfortunately this auto mode selection is prone to several issues that were not predicted before XRT launch. Firstly the presence of Earth Limb contamination causes a large number of events to be seen around the corners of the CCD at the beginning and end of the orbit for targets at or near the declination of the orbit poles (± 69 degrees). Another issue is that of the varying CCD temperature, which can cause a large number of background counts and the appearance of transient hot pixels at CCD temperatures of greater than 50C.

These effects are generally limited to affecting PC mode data, and this leads to a phenomenon known as "mode switching" in which a high count rate seen in PC mode causes the XRT to switch the mode to WT, and then the count rate seen in WT is low due to the differences in the way that event recognition works in WT mode to PC mode leading to less false events from hot pixels being detected, causing the XRT to switch back to PC mode. As there is a dead time associated with switching observing modes, getting into this cycle of mode switching can be an extremely inefficient mode of observation, especially when the source in the field of view is faint, meaning that the WT frames taken are also not useful for analysis.

4.2 On-board Count Rate Calculation

In order to combat this mode switching several strategies have been put in place since launch.

Firstly as the Earth Limb contamination is generally concentrated in the corners of the CCD, the window on the CCD that is utilized to calculate the count rate was reduced from the full PC mode window, to the central 200x200 window on the CCD. This region is where a point source is most likely to fall, although in some cases sources have fallen outside of this window due to a combination of the error in the BAT position of a GRB and the error in *Swift* pointing combining. Luckily this is a very rare occurrence.

Secondly, in order to further remove Earth Limb contamination from the count rate calculation, as well as low energy background events caused by the CCD being too hot, the threshold at which X-ray events are included in the XRT count rate was raised from 80DN (~ 0.2 keV) to 200DN (~ 0.5 keV).

However, these changes, although moderately effective, do not remove the most extreme cases of mode switching, such as where the Earth Limb contamination is bright and located on the central part of the CCD. In order to further reduce the effect of mode switching, a software modification was made to the algorithm that is used to calculate XRT count rate. This modification utilized a new selection criteria to determine if a XRT photon event should be considered "good", which was that an event was only included in the count rate if the central pixel DN was higher than the split event DN (the total DN seen in the 8 pixels of the 3x3 event array other than the central pixel). As most real X-ray events will be highly centrally concentrated, whereas events due to Earth Limb contamination and elevated detected backgrounds tend to be more spread out over the 3x3 event array, this method proved to be highly effective at removing the effect of Earth limb to a point where it no longer causes any mode switching to occur. Furthermore tests show that it dramatically reduces mode switching at high CCD temperatures allowing observations to occur at a much higher temperature (51C compared to 54C prior to the change at the time of the change).

4.3 Calibration State

Although auto selection of the XRT mode has its advantages in burst observations, with the problems of mode switching discussed earlier and the requirement to obtain observations in a fixed mode regardless of source brightness (e.g. calibration observations) it became desirable to find a method of fixing the XRT into a set mode for a particular observation target. The initial solution to this problem was to place commands to set the mode in an Auxiliary Timeline file (ATF). This file consists of commands that get sent to the XRT at designated times. In order to fix an observation in WT mode the ATF would contain commands to put the XRT into Manual State, and then begin taking data in the required mode, timed to start at the beginning of the observation period for the source. At the end of the observation the XRT would be commanded to stop the exposures and return to Auto State.

The major disadvantage of this method is that it puts the XRT into Manual State for fixed intervals. When in manual state the XRT is unable to execute any automated mode changes, which means if a GRB is detected while the XRT is in Manual State, it will not go through its usual automated sequence of burst location, and will not send out any data packets through TDRSS until it returns to Auto State at the time specified in the ATF. This is a major impact in the effectiveness of the XRT to rapidly localize GRBs.

In order to get around this, a change was made to the XRT flight software to allow for a new "Calibration State". This Calibration State allowed the XRT mode to be set as one of the arguments in the Pre-Planned Science Timeline file (PPST). This file contains the information as to what targets *Swift* will observe. The major advantage of this is that it does not require switching to Manual State to set the mode, and if a GRB occurs, the XRT will execute its normal sequence of events, including sending burst information through TDRSS.

Although this state was primarily designed for use in Calibration observations, after its introduction in May 2007, it rapidly became the default mode of operation for non-GRB targets to always pick the XRT mode, apart from in extreme cases where the source was expected to vary. This had the benefit of completely bypassing the issue of mode switching, as well as giving the scientist the ability to choose which mode is best for their observations, rather than having the mode picked for them based on the source brightness. This is particularly important in cases where the scientist requires high timing resolution data of a source that is faint, or if they require positional information on a source that would normally be too bright for PC mode operations. As *Swift* becomes a more general platform for multi-wavelength observing, the ability to pick the XRT mode will much better suit the needs of Guest Observers than the previous Autonomous selection of observing modes.

5. CONCLUSIONS

The *Swift* X-ray Telescope has had to evolve due to several operational challenges since *Swift* was launched in November 2004. The major impact on XRT was the loss of the TEC, which meant that the XRT temperature was no longer electronically controlled to remain at a constant -100C. The requirement to keep the XRT temperature below 50C presented a significant challenge to the science planners, however that challenge has been well met and the XRT temperature has remained below an average 54C for the majority of the observing time. Analysis of data taken over 1000 days of XRT operation has allowed further refinements to be made to the model of XRT CCD temperature, and understanding of the 2nd order temperature effects apparently caused by seasonal variations in the Earth albedo.

Significant advances in the ability of XRT to perform rapid burst localization have been achieved through the increased use of TDRSS telemetry to downlink more detailed data during the first orbit after a GRB is detected. These changes have allowed the XRT team to report accurate locations on all but the very faintest bursts to the wider GRB community usually within 10 minutes of the BAT detecting the burst.

Software changes to how XRT selects its observing modes have greatly reduced dead-time by mostly removing the issue of mode switching. The ability of observers to now pick the mode they wish the XRT to collect data using the new calibration state will be increasingly important as the focus of *Swift's* science program moves away from GRBs towards more general observing programs.

ACKNOWLEDGEMENTS

This work is supported at Penn State by NASA contract NAS5-00136; at the University of Leicester by the Particle Physics and Astronomy Research Council on grant numbers PPA/G/S/00524 and PPA/Z/S/2003/00507; and at OAB by funding from ASI on grant number I/011/07/0. We gratefully acknowledge the contributions of dozens of members of the XRT team at PSU, UL, OAB, GSFC, ASDC, and our subcontractors, who helped make this instrument possible.

REFERENCES

- ¹ N. Gehrels et al., "The Swift Gamma-Ray Burst Mission," *The Astrophysical Journal* **611**, 1005-1020, 2004.
- ² S. D. Barthelmy et al., "Burst Alert Telescope (BAT) on the Swift MIDEX mission," *Space Science Review* **120**, p. in press, 2005.
- ³ D. N. Burrows et al., "The Swift X-Ray Telescope," *Space Science Reviews* **120**, p. in press, 2005.
- ⁴ P. W. Roming et al., "The Swift Ultra-Violet/Optical telescope," *Space Science Review* **120**, p. in press, 2005.
- ⁵ A. F. Abbey et al., "Micrometeorite Damage to CCDs in XMM-Newton and Swift and its Significance for Future X-ray Missions," *Proceedings of 'The X-ray Universe 2005', El Escorial, Madrid, Spain, 2005*
- ⁶ J. A. Kennea et al., "Controlling the Swift XRT CCD Temperature with Passive Cooling," *Proc. SPIE* **5898**, 589816-1, 2005.
- ⁷ J. E. Hill et al., "Readout modes and automated operation of the Swift X-ray Telescope," *Proc. SPIE* **5165**, pp. 217-231, 2004.
- ⁸ S. D. Barthelmy et al., "The GRB Coordinates Network (GCN): A Status Report," *Bulletin of the American Astronomical Society*, Vol. 31, p.718