CHANDRA
X-ray Center 60 Garden St., Cambridge Massachusetts 02138 USA

MEMORANDUM

Date: January 19, 2018
From: Richard J. Edgar and Gregg Germain
To: Chandra Operations Team
Subject: Chandra Radiation Events and Shutdowns in July and September 2017
Cc: MSFC Project Science, CXC Director's Office
File: `SEP1017_memo.tex`
Version: `1.3`

1 Abstract

This memo discusses the thought process that the operations team, especially the ACIS operations team, used during the very high radiation week and events in July and September, 2017.

During the July storm, we did not shut down, and the autonomous mechanisms did not trigger (but very nearly did trigger).

During the week of Sept 10 2017, we experienced two radiation-induced shutdowns, both triggered by the HRC Antico shield.

The key decision points to continue or suspend science operations are reviewed. In both cases the ACE satellite was not being tracked during strong spikes in the soft proton background, and as a result we accumulated an attenuated ACE P3 fluence of 3.85×10^9 in July and 2.63×10^9 in September.

The agreed annual budget for this quantity is 2.0×10^{10}. As of Sept 20, we have accumulated a summed fluence of about 7.7×10^9, or roughly 40% of the annual budget.

2 Introduction

Despite being near solar minimum, 2017 has been host to two very different intense radiation storms. This memo covers the properties of each, reconstructs the timelines and decision processes, and discusses the radiation recieved in terms of the single-orbit and annual budgets.

In addition, two observations were somewhat compromised by high background, and we present scalings of those background rates to what would have been seen if ACIS had been observing with full-frame and 4 or 5 chips active. In both cases, telemetry would have been saturated, meaning less effective exposure time as full frames would be dropped.

3 July 2017 Detailed Timeline

- 2017:195 **Thursday July 14, 2017** JUL1017A load is in progress.
- 2017:195:01:55 **Comm** begins (60 min)
- 2017:195:01:59 M-class solar flare in progress; ACE electrons elevated. 10 hrs 40 min til belt entry; no action. This flare was M2 at its brightest.
- 2017:196 **Saturday July 15, 2017**
- 2017:196:12:35 **Comm** begins (60 min)
- 2017:196:13:02 ACE P3 = 1000; fluence 1.5×10^7. Four or more spikes in ACE electron rates; two spikes in the GOES protons. There is a CME inbound. 38 hours to belt entry; no action.
- 2017:196:21:45 **Comm** begins (60 min)
- 2017:196:22:21 Start of obsid 19703, a 1/8 subarray S3 observation of Sgr A*, no gratings, coordinated with Spitzer.
- 2017:197 **Sunday July 16, 2017**
- 2017:197:03:30 **Comm** begins (60 min)
- 2017:197:03:35 (23:35 EDT) Acquisition of ACE signal, P3 $\sim$ 40,000, fluence $\sim 3 \times 10^8$. It's likely the orbital fluence will exceed 1×10^9. No action taken. 24 hrs to belt entry.
- 2017:197:05:00 P3 spike when the CME shock passed ACE.
- 2017:197:05:07 (01:07 EDT) MTA ACE alert (2 hr average P3 exceeds 50k). P3 fluence $\sim 6 \times 10^8$, next comm 08:15 EDT. P3 up sharply at $\sim$ 100k.
- 2017:197:05:25 (01:25 EDT) ACIS ground software reports a "Bogus P3 alert" due to a sharp rise (110k to 284k) in P3 which is real.
- 2017:197:06:00 – 09:00 ACIS threshold crossings (txings) high and rising. Ground algorithm suggests it might have triggered, but the flight instrument did not trigger.
- 2017:197:06:10 FLUMON fluence hits 1×10^9.

- 2017:197:10:20 (06:20 EDT) Fluence reads 2.8×10^9, and is probably higher due to the software's erroneous Bogus Alert.
- 2017:197:12:15 **Comm** scheduled start; OC obtained early data around 12:00Z.
- 2017:197:12:00 (08:00 EDT) Comm reveals NO SCS-107 trip. Estimate another $4 - 5 \times 10^8$ fluence on top of the 2.9×10^9 reported by the fluence monitor software.
- 2017:197 All day: Geomagnetic storming is reported.
- 2017:197:19:25 FLUMON fluence hits 3×10^9 (plus the $4 - 5 \times 10^8$ as above).
- 2017:197:19:15 **Comm** begins (60 min)
- 2017:197:19:30 P3 down to 2400.
- 2017:198:03:08 RADMON disable (from the daily loads)
- 2017:198:03:30 **Comm** begins (60 min)

4 July 2017 Discussion

The total accumulated fluence for the storm (computed from the ACE dump data) is 3.85×10^9. The P3 flux was over 10^5 (averaging about 2×10^5) for nearly four hours, after the last comm on the evening of day 196, between 1:30 and 5:30 AM local time. This is the bulk of the fluence accumulated for this storm.

The HRC Shield GOES proxy was never over about 45, on a scale where the trigger value is 245. This was a very soft storm, with ACE P3 well over 1×10^5 but the GOES proxy never near its trigger value.

ACIS telemetry did not saturate and therefore no frames were dropped, because of the fortunate choice of subarray and a single chip. In figure 3 we plot the actual event rate (all telemetered events) binned to 100 seconds, and also scaled up to 4 and 5 chips full frame, to show what might have happened during a more typical observation.

The non-x-ray background was greatly enhanced during this observation.

The small subarray also made the txings signal noisy, and so less likely to trigger, as it needs to be above the trigger level and monotone increasing for five consecutive (averaged) samples.

5 Data plots for the July 2017 storm

In figure 1, we have plotted the ACE P3 flux rate, in the usual units, which are protons s^{-1} cm^{-2} sr^{-1} MeV^{-1}, throughout the July storm. Also marked are radiation belt passages, and which instrument is in the focal position. Note that the flux was above 10^5 for several hours, and that the onset of the intense part of the storm was very sudden (each point represents a 5 minute average).

In figure 2, we present the MIT psci plot of the ACIS threshold crossings as a function of time, for the day of the storm.

In figure 3 we show the light curve for all telemetered events for obsid 19703, which was observed during the storm, starting at time 2017:196:22:21. The Level 1 event data were binned in 100 sec bins, and plotted in blue. Since this observation fortuitously used a 1/8 subarray and one chip, we had no issues with telemetry saturation. A more typical instrument setup, say four or five chips clocking full frames would have been near the VF telemetry saturation limiting count rate for the entire observation, and approached the Faint mode limit during the peak of the storm (the Faint mode telemetry saturation limit would have been exceeded twice for brief periods if 5 CCDs had been in use). The curves in figure 3 are simply scaled by the ratio of chip areas, without consideration for differing count rates due to background events in FI vs. BI CCDs.

Also plotted, in figure 4 are the GOES proton rates, from both active GOES satellites, summing the E and W detector P2 channels (about 4–9 MeV). This shows some elevated activity on day 196 and then a sudden dramatic spike above 100 flux units coinciding with the spike in ACE P3 protons. See discussion below on the behavior of the GOES P2 channel protons during the September storm, when the flux was an order of magnitide higher.

6 Lessons Learned from the July 2017 storm

- The ACE P3 flux can increase rapidly and accumulate a fluence of a few $\times 10^9$ within four to five hours.
- With the present operations paradigm, we will take hits like this one from time to time.
- It's possible to second-gess the ACIS Ops decision at 23:30 EDT on July 15, but given the evidence then available, the decision is defensible.
- Software in the FLUMON Fluence Monitor process may edit out very rapid rises in the flux. The ACIS Ops person should check cron and spam folders for BOGUS P3 ALERT e-mails.
- It's possible during an intense storm to have telemetry saturation that results in high on-chip background and exposure time loss, as whole frames of data are discarded because they don't fit in the telemetry stream. This renders the data less useful than they would otherwise be. Even those frames which are telemetered to the ground will contain elevated background levels due to the high flux of charged particles passing through the detector.

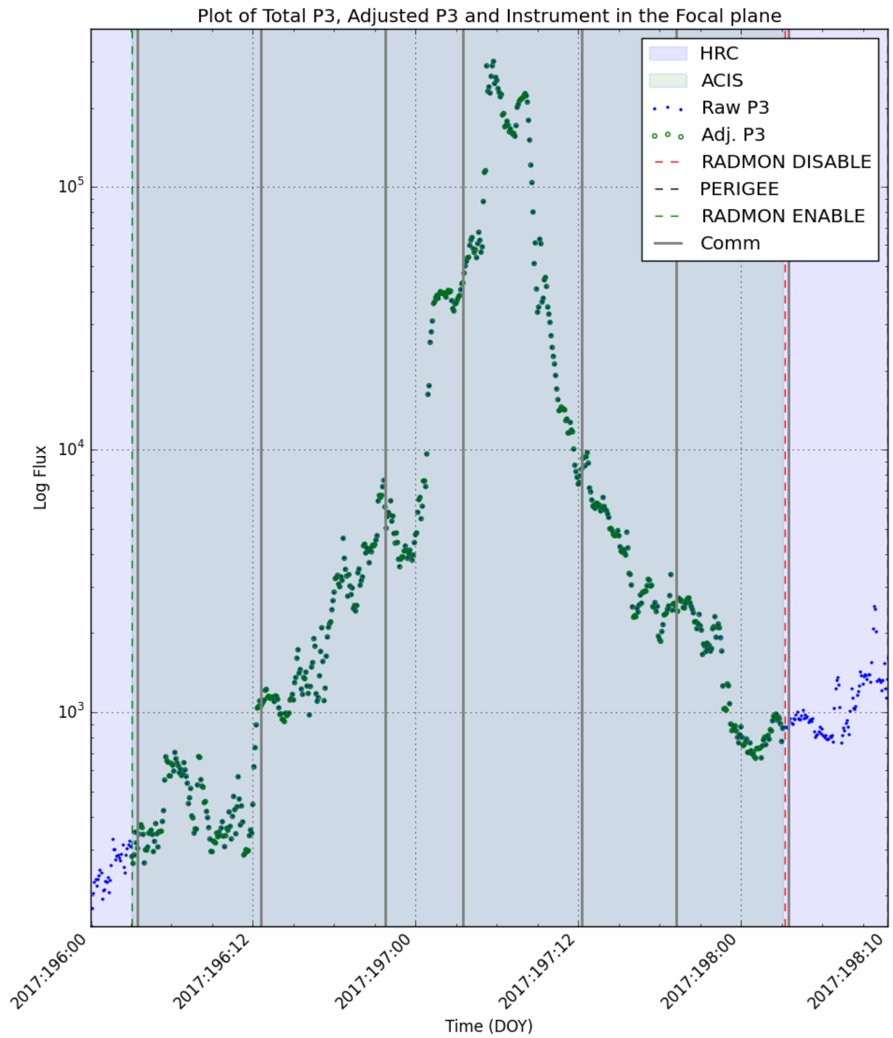

Figure 1: A plot of the ACE P3 flux during the July storm. Shading indicates the satellite configuration, with blue for times when HRC is in the focal position (radiation belt passages) and green for when ACIS is in focus. Fluence is integrated when ACIS is in focus. Grey vertical bars mark starts of scheduled DSN communications.

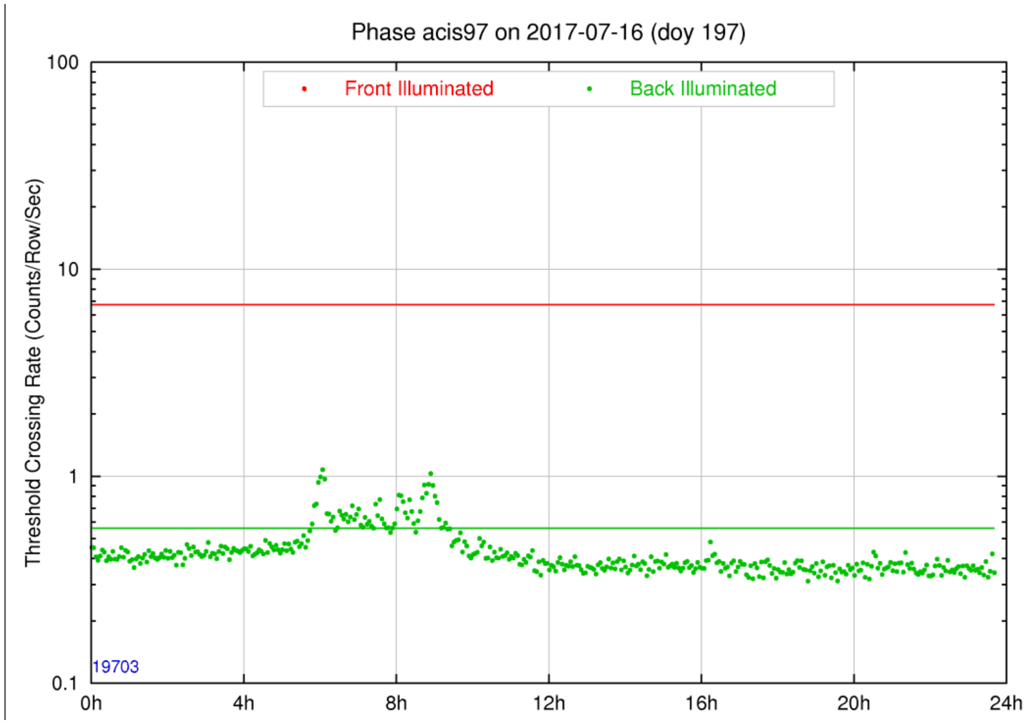

Figure 2: The txings plot for 2017:197 (July 16). Red is for FI chips, green is for BI chips. The horizontal lines are the trip thresholds for each type of chip. The txings algorithm came close to triggering an SCS-107.

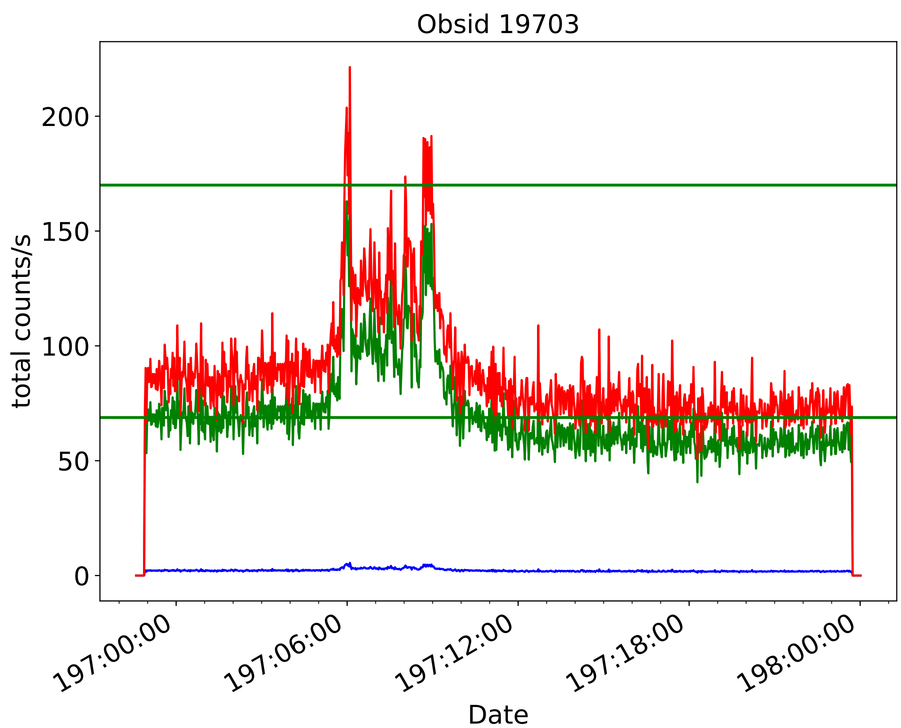

Figure 3: Light curve for all telemetered events for obsid 19703. Measured count rates are in blue. Horizontal lines are the telemetry saturation limit for VF and F modes, at 68.8 and 170.2 counts per second. Time is in days from start of observation. Also plotted are rates scaled to 4 full-frame chips (green) and 5 full-frame chips (red). Common observing configurations would have saturated telemetry during this storm.

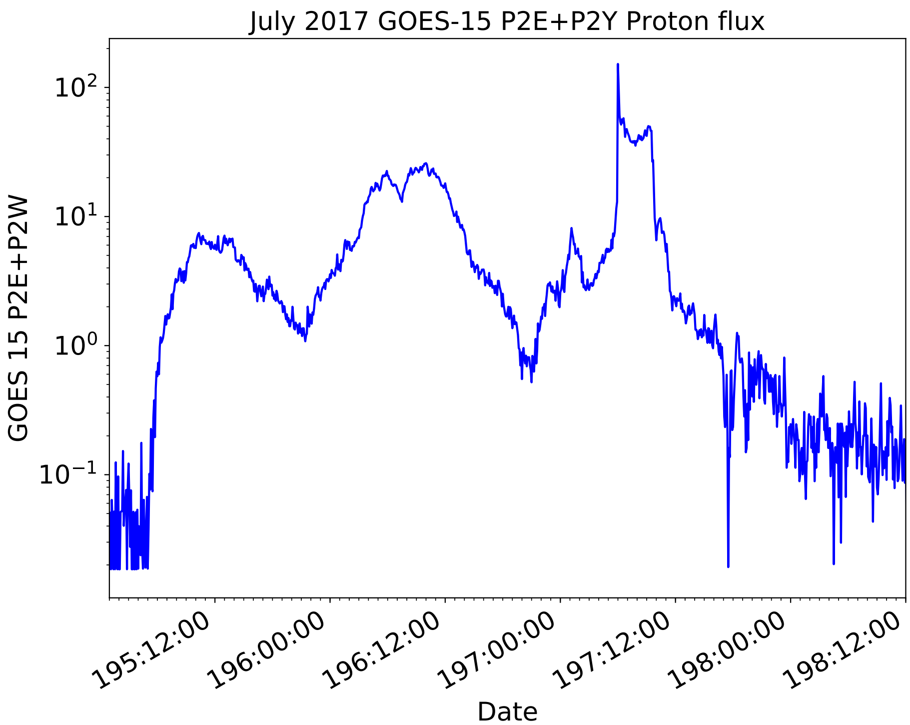

Figure 4: Proton rates from GOES P2 channels (about 4–9 MeV) for the UT days July 14–17, 2017 (2017:195–198). The spike above 100 flux units coincides closely with the maximum in the ACE P3 channel (see figure 1). There follows a plateau with readings of about 40 units for several hours.

7 September 2017 Detailed Timeline

- 2017:247 **Monday, September 4, 2017** SEP0417A in progress.
- 2017:247:20:35 M5+ flare in progress (SWPC alert). CME observed.
- 2017:247:23:15 **Comm** begins; P3$\sim$ 100 but increasing. 12 hr to belt entry.
- 2017:248 **Tuesday, September 5, 2017**
- 2017:248:00:45 **Comm** begins (60 min)
- 2017:248:03:30 **Comm** begins (60 min)
- 2017:248:03:33 MTA GOES alert: over 1/3 P4GM threshold
- 2017:248:07:15 SWPC alert: 10 MeV flux over 100 pfu
- 2017:248:12:45 **Comm** begins (60 min)
- 2017:248:16:00 Summary:
  - GOES rates turned over; steady at an elevated rate
  - ACE proton spectrum very hard; expect higher rates soon
  - ISWA model predicts shock passage at about 2017:251:00:00
  - Expect P3 to rise going into the shock; might fall off after.
  - Recall we took over 3×10^9 in one orbit in July
  - Possible opportunity for long CTI run if there's a shutdown.
- 2017:248:19:20 **Comm** begins (60 min)
- 2017:249 **Wednesday, September 6, 2017**
- 2017:249:01:35 **Comm** begins (60 min)
- 2017:249:09:04 X2.2 flare (SWPC alert)
- 2017:249:11:30 **Comm** begins (60 min)
- 2017:249:11:56 X9.3 flare (SWPC alert)
- 2017:249:11:58 Isolated spike to 248 in HRC shield rate; no shutdown. Recall 245 is the trigger, but it needs 10+ consecutive samples.
- 2017:249:19:00 **Comm** begins (60 min)

- 2017:249:19:00 ACE P3 fluence $\sim 2.3 \times 10^8$. HRC shield rate ~ 59. A halo CME was observed after the X9.3 flare.
- 2017:249 Note the ISWA model with this CME in it removed the wave from the Sep 05 flare.
- 2017:249:23:45 10 MeV flux over 100 pfu (SWPC alert)
- 2017:249:23:48 Shock passes SOHO, DSCOVR at the sun-earth L1 point.
- 2017:249:23:48 Summary:
  - ACE is LOS, but at last report ACE P3 = 18k and rising;
  - GOES protons spike
  - nothing in the HRC proxy.
  - ACE P3 fluence is at least 6.0×10^8
  - We projected, given P3 $\sim$ 23k, fluence would be $\sim 1.6 \times 10^9$ by the 07:30 EDT comm Thurs 7 Sep
  - Shock passage suggests P3 might go down
  - Therefore, continue!
  - Obsid 18997 (1/4 subarray, S2,S3,S4) flirting with TM saturation.
  - txings could have triggered (above threshold, but slightly too noisy to be monotone increasing).
- 2017:250 **Thursday, September 7, 2017**
- 2017:250:02:41:28 HRC shield triggers SCS-107 (out of comm)
- 2017:250:03:30 **Comm** begins (60 min)
- 2017:250:03:37:33 Comm shows SCS-107 had run; telecon called. Target Friday night Sep 8 (day 251) for return to science.
- 2017:250:10:15 M5+ flare in progress (SWPC alert)
- 2017:250:14:36 X1.3 flare (SWPC alert)
- 2017:250:22:00 Shock from X9.3 flare passed.
- 2017:251 **Friday, September 8, 2017**
- 2017:251:07:42 MTA alarm for P3 spike (spurious?)
- 2017:251:07:46 M5+ flare in progress (SWPC alarm)

- 2017:251:19:45 **Comm** begins (60 min)
- 2017:251:19:45 CAP 1432 (long CTI run) started
- 2017:252 **Saturday, September 9, 2017**
- 2017:252:01:45 GO for uplink of SEP0917C loads
- 2017:252:02:00 **Comm** begins (60 min). Loads uplinked.
- 2017:253 **Sunday, September 10, 2017**
- 2017:253:11:40 **Comm** begins (60 min).
- 2017:253:16:06 X8.2 flare (SWPC alert)
- 2017:253:16:25 100 MeV flux over 1 pfu
- 2017:253:17:05 10 MeV flux over 100 pfu
- 2017:253:17:18 MTA alert for GOES protons
- 2017:253:18:23 MTA alert for GOES protons
- 2017:253:18:25 **Comm** begins (60 min).
- 2017:253:18:31 HRC proxy rising rapidly; HRC shield rate and ACIS txings rates also rising rapidly at RADMON disable, but neither triggers SCS-107. Estimates are both were minutes from triggering. Huge rates of Solar Energetic Particles; half the front-illuminated chip pixels are above the event threshold during belt-entry CTI run (obsid 49789).
- 2017:253:18:46 10 MeV flux over 1,000 pfu (SWPC alert)
- 2017:254 **Monday, September 11, 2017**
- 2017:254:03:30 One hour **comm** in the radiation zone. The GOES proxy was still extremely high. This comm could have provided an opportunity for a manual shut-down.
- 2017:254:07:51:39 SCS-107 tripped by HRC shield upon belt exit (plus 10 minutes) Safing included a stuck science thread active bit requiring warmbooting the ACIS BEP. In addition, the HRC team is concerned about possible damage to the HRC Anticoincidence Shield photomultiplier tubes.
- 2017:254:11:30 **Comm** begins (60 min).
- 2017:254:19:55 **Comm** begins (60 min).

- 2017:254:20:02 CAP 1433 to warm boot the ACIS BEP

  The GOES proxy for the HRC shield stayed very high for days.

- 2017:256 **Wednesday, September 13, 2017**

- 2017:256:17:30 HRC proxy 1.2 times limit; P3 $\sim$ 10k; both very flat. Estimate annual fluence may be $\sim 8 \times 10^9$; nearly half the annual budget of 2.0×10^{10}.

- 2017:256:18:10 **Comm** begins (60 min). SEP1317B loads uplinked.

- 2017:256:21:23 Resume science with SEP1317B. HRC shield rate at trip level; ACIS txings slightly over but steady (bright source in CC/Graded mode)

8 September 2017 Discussion

The series of eruptions from Active Region 2673 presented the operations staff with a unique challenge, both in terms of the intensity of the radiation storm, and its spectral properties. This region formed quickly near the centerline of the solar disk, and within the week plus before rotating off the disk emitted multiple M-class flares, and three X-class flares, two of which rank as the strongest solar flares in a decade or more.

In figure 5, we again plot the ACE P3 flux data. Also plotted are the two SCS-107 transitions (vertical red bars), radiation zone passages (vertical dashed lines), and which instrument was in the focal plane. Blue dots are 5-minute averages of the full P3 flux, while green dots represent the attenuated flux (weighted with zero for HRC in the focal plane, 0.2 for HETG inserted, and 0.5 for LETG inserted). The integral of this attenuated flux is the fluence which we monitor.

In figure 6, we show the HRC anticoincidence shield rate in uncalibrated units, where 245 is the trigger level. Four events are apparent:

- The momentary excursion at 249:11:58 which did not cause a shutdown; these were the prompt protons from the X9.3 flare 2 minutes earlier.

- A strong rise to the trigger level at 250:02:41, which did trigger SCS-107.

- A strong rise just prior to belt entry at 253:18:31. This would probably have triggered (if ACIS txings didn't beat it to the punch), except that RADMON was disabled for a regularly scheduled belt entry before that could happen.

- Upon exit from that radiation zone, the shield rate was above the trigger value as soon as it was powered up. This resulted in another SCS-107 shutdown at 254:07:51.

The spike in the HRC shield rate observed at 249:11:58 could have been taken as a warning of things to come. Some flares emit prompt Solar Energetic Particles (SEPs), and some do not. Shock-accelerated protons and other nuclei may not be able to escape from the flare region on the sun, because of overlying magnetic field structures. Thus if one strong flare from a given active region emits SEPs, one may expect a higher probability that future flares from the same active region will do so as well. Also any super-thermal perticles lingering in the area from a previous flare can be accelerated to much higher energies by the shock of a subsequent flare.

One further risk of repeated CMEs is that the second shock wave will overtake the first one, and the resulting converging flows can efficiently accelerate energetic particles, which should arrive with the inter-shock zone of the two CMEs.

While the catalog of previous Chandra radiation shutdowns contains two ”harder” events from the 2001 time frame, this storm was unusually hard and pervasive. For example, the Solar Energetic Particles caused auroras seen on Mars when they hit that planet's atmosphere.

In figure 7, we show the HRC Anticoincidence Shield proxy based on GOES proton rates. This plot is available in near real time on the CXC Replan Central page. Depending on the spectrum of the protons in a storm, it can closely approximate the actual HRC shield rate. Note that the trigger threshold is 245, and the maximum value is around 8000, a factor of about 30 over the trigger threshold.

Also plotted, in figure 8, are the GOES proton rates, from both active GOES satellites, summing the E and W detector P2 channels (about 4–9 MeV). This shows some elevated activity starting on day 248 and then sudden dramatic spikes above 1000 flux units coinciding with the spike in ACE P3 protons. After returning to background levels on day 253, there was another dramatic rise to a few hundred flux units in day 254, followed by a very slow decrease over the following days. The lack of detailed correspondence between GOES P2 and ACE P3 during the second phase of the storm, days 254-258, shows that these proton detectors are measuring different phenomena.

In the panels of figures 9 through 11, we show the ACIS threshold crossings during the September 2017 storm.

Late on day 2017:249, the txings went above the RADMON trip threshold, but were slightly too noisy to qualify for five samples of monotone increasing values. This came very close to tripping, and the version of the ground algorithm in use at the time did in fact show it triggering.

On day 253, the hard proton storm arrived coincident with the regularly scheduled RADMON disable. Both the ACIS txings and HRC anticoincidence plots show values rising rapidly toward trip thresholds, but the RADMON disable happened first. Coming out of the belts on day 254, the txings were still well above threshold, but decreasing, during the ECS measurement. ACIS was still taking bias frames when the HRC antico shield triggered the RADMON process at the start of the first science observation.

On day 256, science was resumed. The HRC proxy was at threshold and decreasing very slowly; the txings plots show the same thing for the FI chips. The first science target (obsid 20203) was a CC mode observation of a bright source, and the S3 txings plot shows factor of two variations at random, so did not trigger even above threshold.

In figure 12 we show the light curve for all telemetered events for obsid 18997, which was observed during the storm, starting at time 2017:249:06:01. The Level 1 event data were binned in 100 sec bins, and plotted in blue. Since this observation fortuitously used a 1/4 subarray and three chips, we flirted with telemetry saturation only once, about 23 hours into the observation. A more typical instrument setup, say four or five chips clocking full frames would have been above the VF telemetry saturation limiting count rate for the entire observation, and violated the Faint mode one by large factors during the peak of the storm. The curves in figure 12 are simply scaled by the ratio of chip areas, without consideration for differing count rates due to background events in FI vs. BI CCDs.

9 Lessons Learned from the September 2017 Storm

- The ISWA Heliosphere model actually removed the first CME before it arrived, when the 2nd CME was added. It's not clear why this was done; possible human error, or knowledge on the part of the providers that the model might not work well with a second shock catching up to the first one. ACIS Ops people should be aware.
- We could consider altering the ACIS txings parameters during CTI measurements. If they had the sense of triggering whenever the rate got to be extremely large, whether it was increasing or not, it could have triggered either at the entrance or exit CTI run on days 253-254. This would require HRC Antico Shield high voltage to be turned off (as it is now), but the RADMON left enabled until the end of the CTI measurement.
- The HRC Mission Planning guidelines are under review to prevent a situation such as the one on day 2017:254.
- If the HRC proxy is above about 3 times the limit during a perigee passage as it was here, we or the HRC ops team can suggest (to hrcdude, for example) that we shutdown manually. We should be prepared to support the HRC team in any way they need during such a discussion.

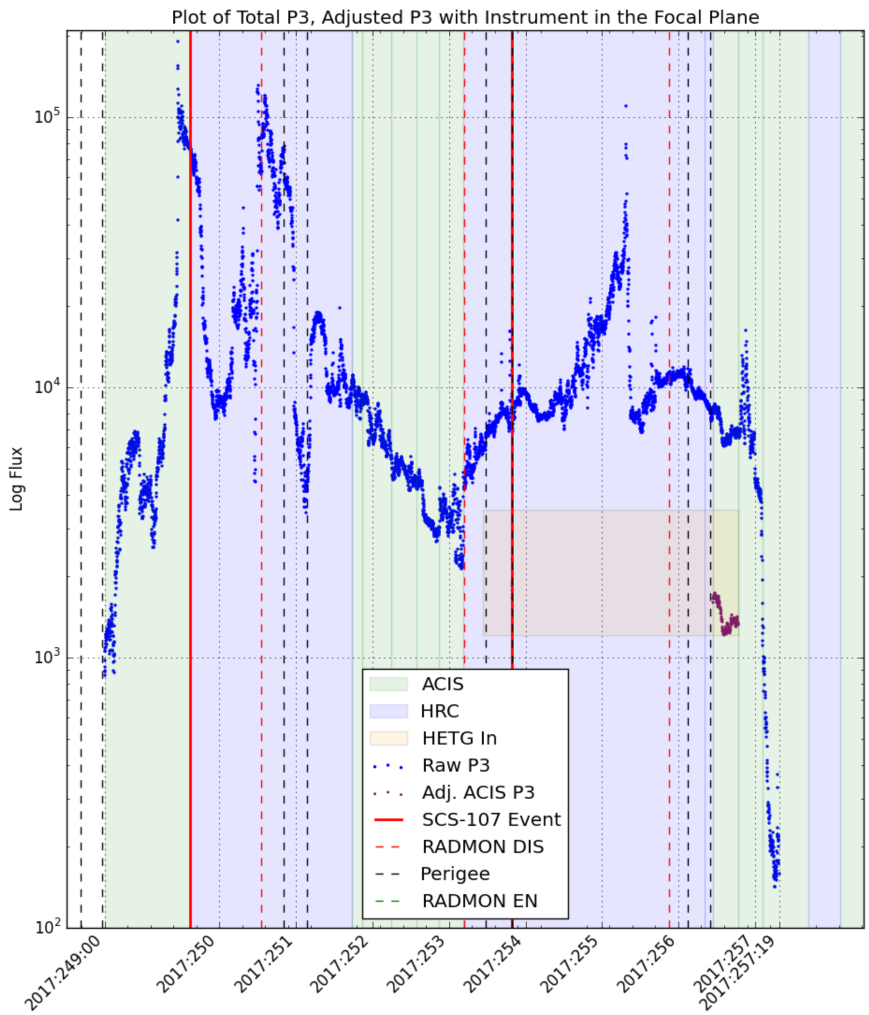

Figure 5: A plot of the ACE P3 flux during the September storm. Shading indicates the satellite configuration, with green for times that ACIS is in the focal position, blue for HRC (typically shutdowns and radiation belt passages), and orange when the HETG is inserted. Blue data points are the actual P3 rates, and the orange ones on day 256 are adujsted for the factor of 5 attenuation due to the HETG.

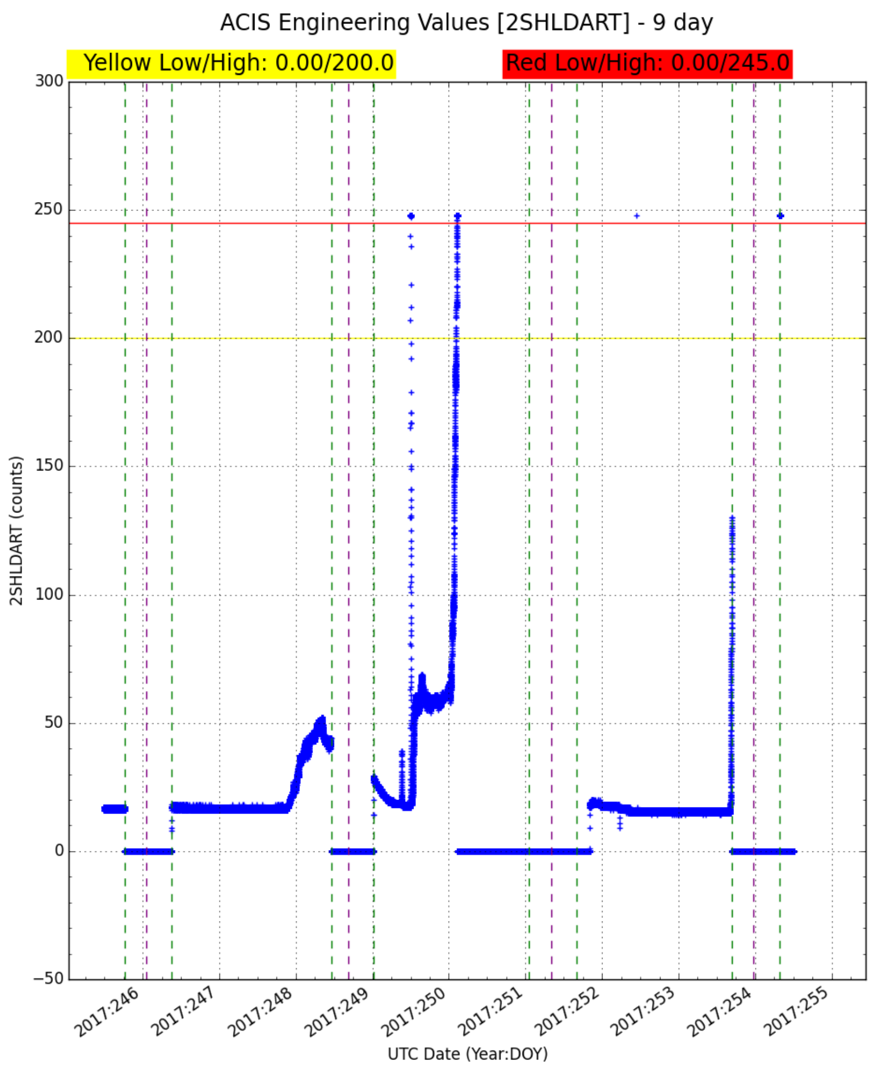

Figure 6: Plot of the HRC Anticoincidence shield rate, showing two SCS-107 shutdowns and two near misses. Note a few points above the red line on day 2017:254, just prior to the second shutdown.

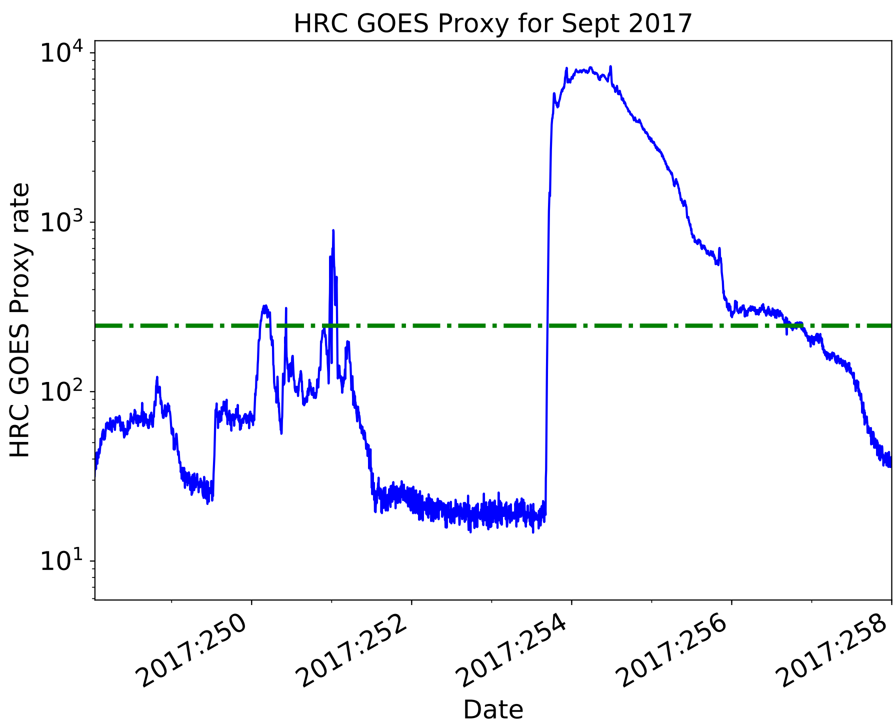

Figure 7: Plot of the HRC Anticoincidence shield proxy, based on GOES proton fluxes. This is always available (there's a plot on the Replan Central page), and closely approximates the actual HRC Antico Shield rate. The trigger threshold for the HRC Antico Sheild rate is 245, plotted in green.

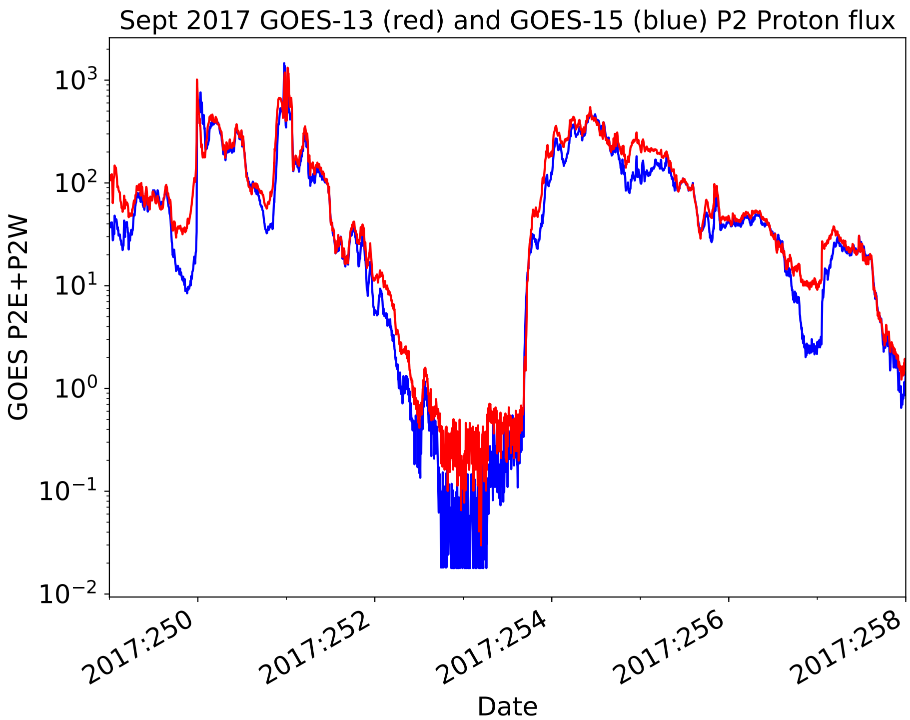

Figure 8: Proton rates from GOES P2 channels (about 4–9 MeV) for the UT days September 5–15, 2017 (2017:248–258). The spikes above 1000 flux units coincides closely with the maximum in the ACE P3 channel (see figure 5). The second sharp rise on day 254 accompanies a less dramatic rise in ACE P3.

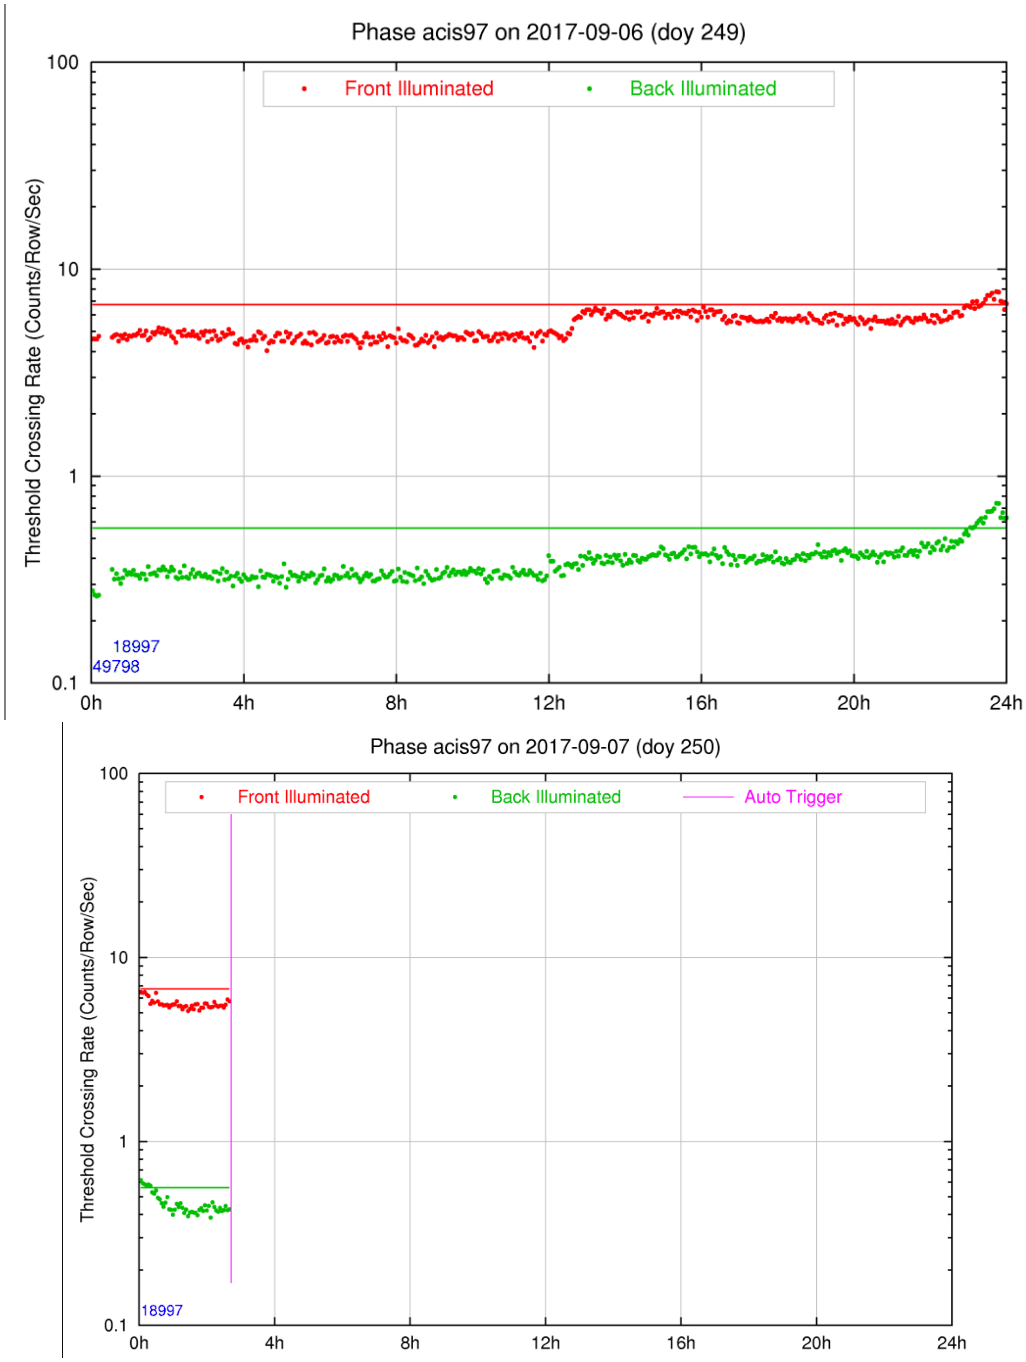

Figure 9: ACIS txings for 2017 days 249, 250. Note at 249:23 or so the txings are above threshold, but slightly too noisy to be monotone increasing for five samples, so RADMON did not trigger.

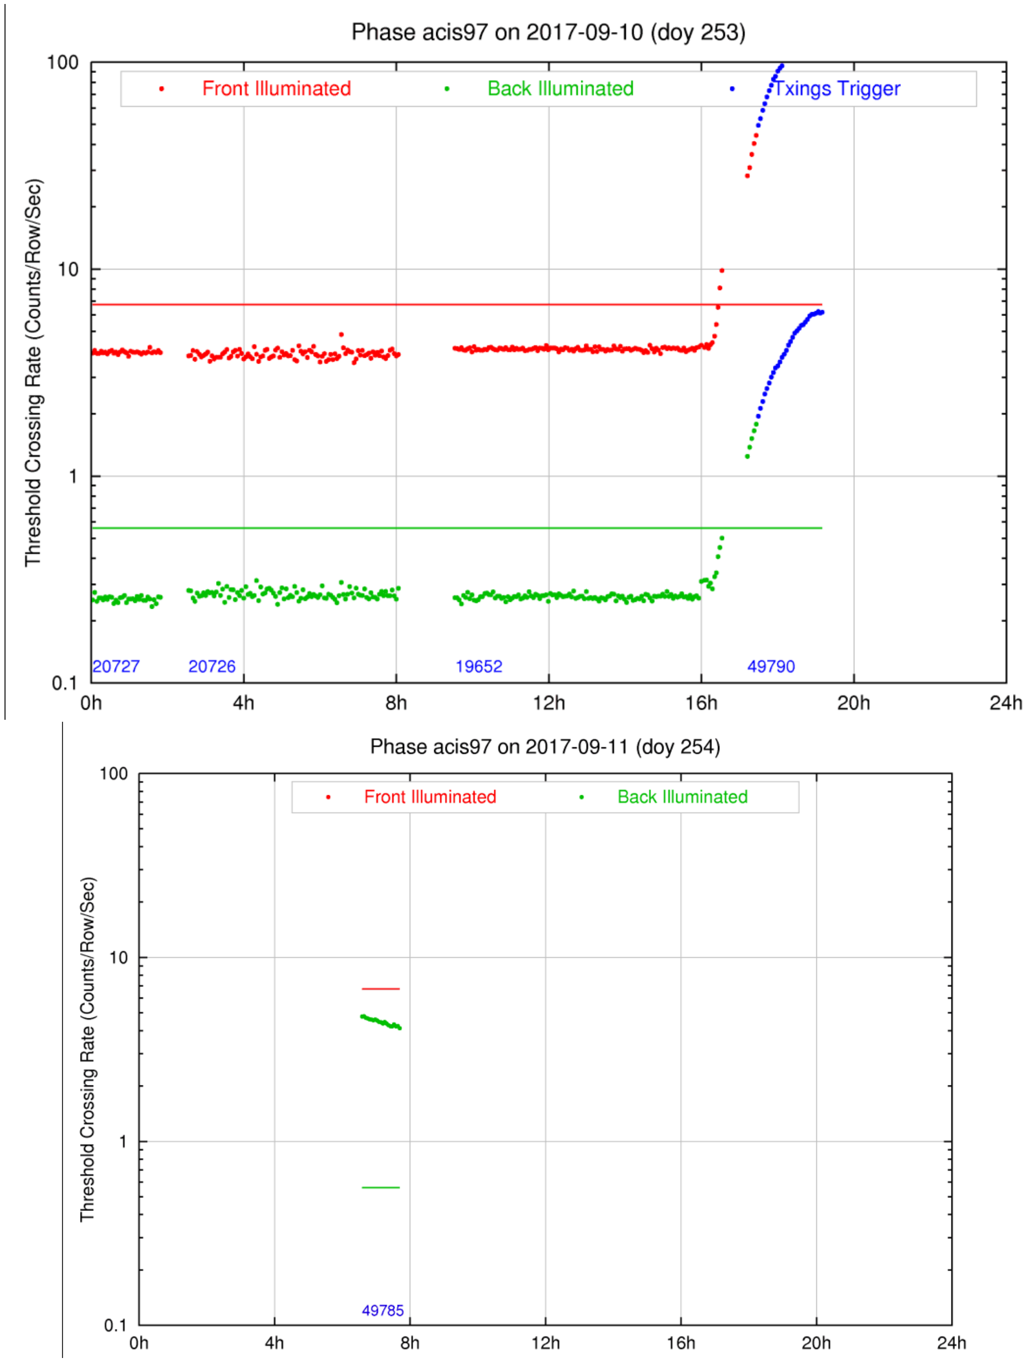

Figure 10: ACIS txings for 2017 days 253, 254. Note the sharp rise on 253 just as Radmon was disabled by the daily loads to fly into the belts. The txings algorithm would have triggered just minutes later. On day 254 during the CTI run, txings are still an order of magnitude above trigger levels, but declining. FI curve is off the chart.

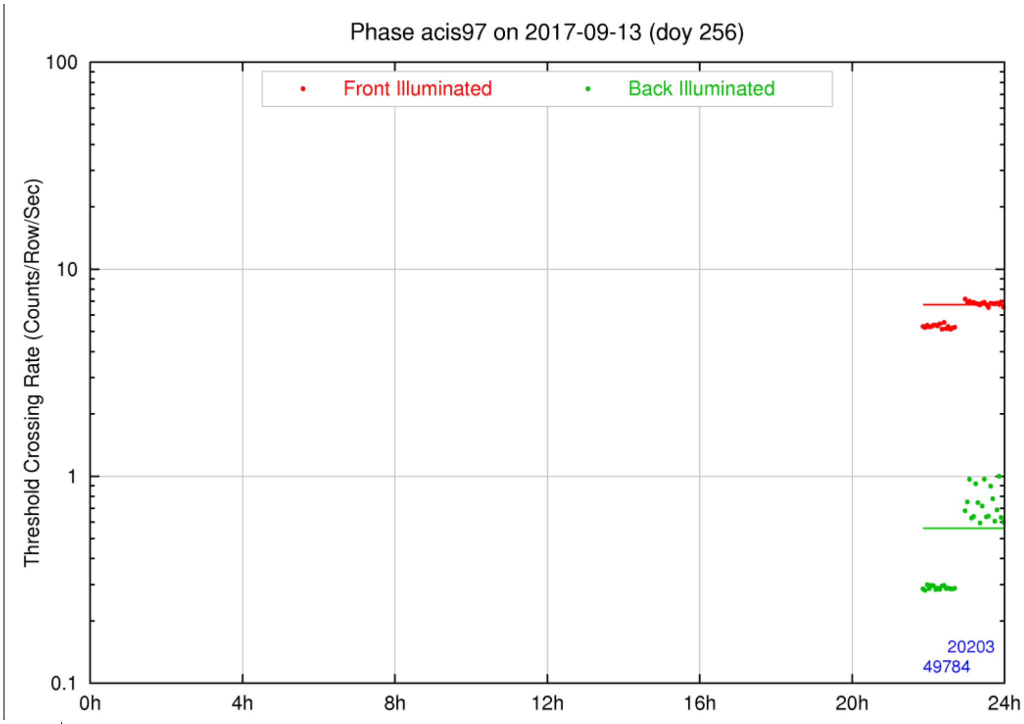

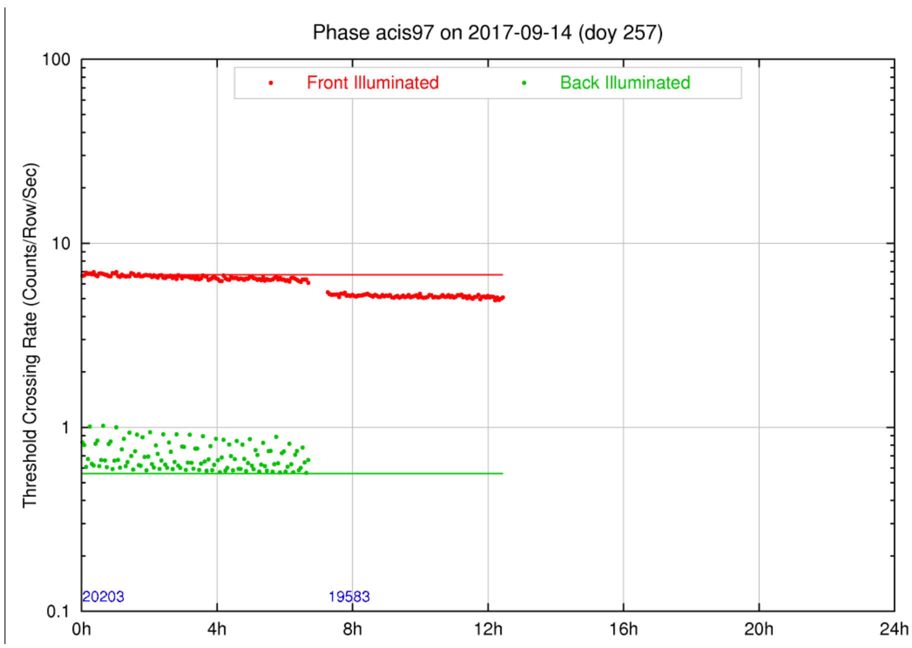

Figure 11: ACIS txings for 2017 days 256, 257. Return to science was ordered late on day 256 with rates right at the thresholds, but (very slowly) declining. The noisy BI behavior is expected for bright sources observed in the CC mode.

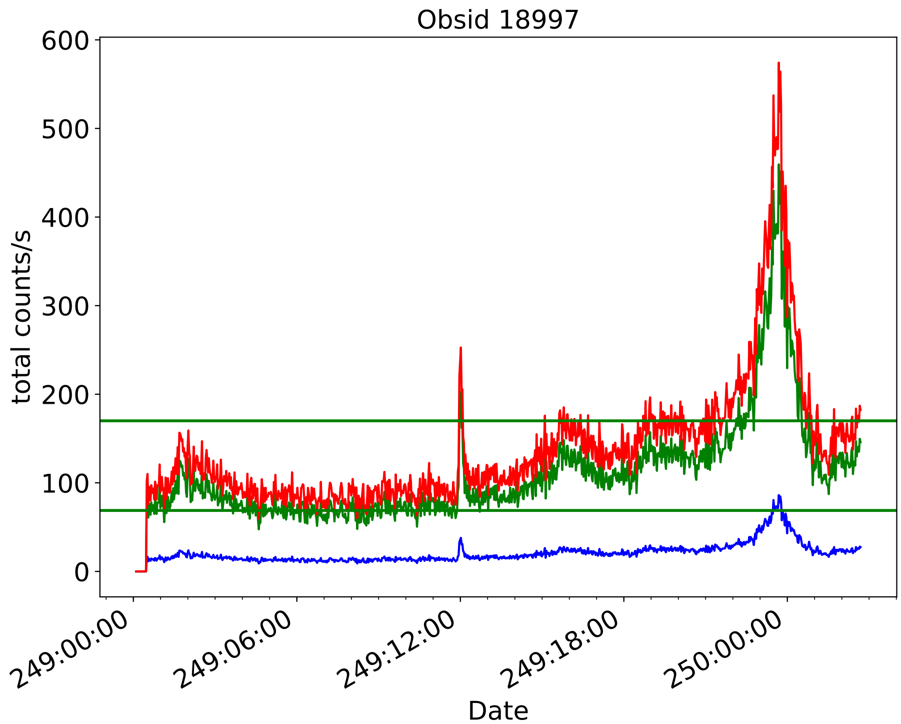

Figure 12: Light curve for all telemetered events for obsid 18997. Measured count rates are in blue. Horizontal lines are the telemetry saturation limit for VF and F modes, at 68.8 and 170.2 counts per second. Time is in days from start of observation. Also plotted are rates scaled to 4 full-frame chips (green) and 5 full-frame chips (red). Common observing configurations would have saturated telemetry by large factors during this storm.

10 Resources

In this section I list unusual sources of some of the data I have used.

The July ACE P3 data were annoyingly too new to be in the final ACE archive, and too old to be in the quick-look or processed dump data set. However, we managed to find several recent years' worth of data, plus this year-to-date, at this url:

`ftp://mussel.srl.caltech.edu/pub/ace/browse/`

The data are in HDF4 format, which can be read or converted to ascii data by use of a program hdf-4.2.13 which I downloaded and installed. One member of that suite is hdp, which I used thus:

```
$bin/hdp dumpvd -v -i 0 ACE_BROWSE_2017-001_to_current.HDF > dumpvd.txt
```

It was then possible to convert the text file into an RDB table and reformat it to recover the date string. The script I wrote for this purpose can be found here:

`/home/edgar/acis/ops_memos/SEP1017/ACE_BROWSE`.

Also of interest are a little python notebook John ZuHone put together to plot back GOES proxy data. This is the result of a calculation based on real-time GOES proton flux data, and the code that supplies the plots for Replan Central also stores the data. I have adapted the notebook for use with command-line python, here:

`/home/edgar/acis/ops_memos/SEP1017/python/plot_hrc_proxy`

The data are stored in HDF5 format here:

`/proj/sot/ska/data/arc/hrc_shield.h5`.