CTI Correction Approaches

Catherine Grant, Leisa Townsley,
and the ACIS IPI Team

26 April 2001
revised 30 April 2001

1 Introduction

It has been shown that a CTI correction algorithm can improve spectral resolution using only the standard telemetry products. Two correction approaches currently exist. The first, from the Penn State ACIS team, (Townsley et al., 2000) was originally developed to correct for CTI effects in BI ACIS CCDs and has been under development for well over a year as a scientific data analysis tool for use with data from the (now degraded) FI detectors.[1] The PSU corrector was originally applied to reduce spatial variations in detector response, which are especially pronounced at the focal plane temperature (-110C) used early in the mission. The additional improvement in spectral resolution gained from correcting column variations was realized more recently (Townsley & Broos, 2000). The second CTI correction approach discussed here was developed by the MIT ACIS team as part of the continuing effort to understand and reduce the influence of cosmic rays and sacrificial charge on FI detector performance. One of the primary conclusions of this effort was that a (ground based) CTI corrector could significantly improve ACIS spectral resolution (Bautz et al., 2001).

These findings have motivated the CXC to consider development of a CTI correction tool that could be applied to Chandra data already in hand. One purpose of this memo is to summarize and compare the various methods used within the ACIS team to perform this correction. A second objective is to identify the significant design choices, and attendant costs and benefits, to be made if the CXC is to develop a new CTI corrector for Chandra users.

The remainder of the memo is organized as follows. We conclude this introduction with a brief comparison of the two approaches. The next two sections contain more detailed descriptions of the physical assumptions and algorithms used. (Still more detailed explanations of the algorithms and the thought behind them can be found in the memos and papers listed in section 7.) In section 4 the spectral resolution improvements of the two approaches are shown to be quite similar, and the effectiveness of the PSU grade corrector is demonstrated. Key differences in approach are discussed in section 5. A summary of design choices is presented in the last section.

We have found no significant differences in the fundamental physical ideas underlying the two approaches to CTI correction. The two approaches produce very similar improvements in spectral resolution. Both analyses include energy- and position-dependent corrections of individual pixel amplitudes for the effects of CTI. Both analyses measure and apply corrections for column-to-column variations of trap density. While both analyses identified a significant variation of trap density within each column, only the PSU analysis applies a correction for this effect. The PSU corrector also compensates for CTI-induced changes in the spatial distribution of charge within each event (i.e., grade correction.) The effect of this correction on detection efficiency is most

[1] The Penn State CTI corrector has been made available to the Chandra user community through the CXC Contributed Software Exchange. Section 8 is a list of publications benefiting from the Penn State corrector.

significant for data taken early in the mission (when the focal plane temperature was at -110C), but we believe it is a scientifically very valuable capability, especially because it can restore some flexibility in grade selection at the current operating temperature (-120C).

In both approaches the correction is applied event by event to all telemetered events. The parametrization of both the energy- and row-dependence of the charge loss function are used to adjust some of the pulseheights in the event island (BI CCDs also require column-dependent charge loss correction). The performance improvement is a result of two factors; correcting each pixel in split events separately before summing the pulseheights and correcting for differences in the charge loss in each column to account for column to column variations in the trap density. The first factor is more important at high energies, the second at lower energies. After correction the event is regraded and the PI and energy values must be recomputed using gain tables calibrated for corrected data.

Both the RMFs and ARFs must be recalibrated for use with CTI corrected data, but because the performance is more uniform, fewer regions can be used. The CTI correction models discussed here do not require any additional information about the event beyond that in a standard eventlist; this implies that the effects of sacrificial charge are included in the calibrated parameters. A significant change in the background or in the intrinsic radiation damage would require recalibrating the CTI correction parameters and the standard calibration products (gain, RMF, ARF), but a very similar recalibration would be needed even if no CTI correction were implemented.

In addition, the CTI corrector and the related calibration products are probably different for different CCD readout modes such as subarrays and continuous clocking since the change in exposure time implies a change in the effective particle background. Neither CTI corrector has been calibrated or tested yet on data in other readout modes but the methodology should be similar to that outlined here.

2 MIT CTI Correction

The MIT approach is based on a direct model in which the pulseheight values in the observed event are used to determine an additive correction. Charge is only added to the observed event island pulseheights, never removed, so that the overall charge loss is corrected, but not any change in grade due to morphing and charge trailing.

The charge loss as a function of energy is parametrized as a power law. The total charge loss for a charge packet of size Q at position x and y is then

$$\Delta Q(x, y, Q) = y\, n\, V = y\, n\, k\, Q^{\alpha}$$

where y is the number of parallel transfers through the damaged region, n is the average trap density between rows 1 and y, V is the volume of the charge cloud, k is a constant of proportionality and α is the power law index. In principle the amount of charge in the packet should be recalculated for each pixel transfer to account for the decreasing size of the charge packet, however the change is small and is currently ignored. In this formulation, Q is the original charge deposited by the photon which, for an individual X-ray photon, is unknown. A zeroth order approximation of Q is the observed charge packet, Q_{obs}. A better approximation of Q, which is implemented in the MIT corrector, is $Q_{obs} + \Delta Q_{obs}$.

This loss factor is determined separately for each of the highest pixels in the three columns of the event island. If no pixels in a column are above the split threshold, it is considered to not contain signal and is ignored. In this way the grade of the observed event is preserved.

The parameters in the charge loss model which must be predetermined are n, k and α. Since n and k are not easily separated, they are treated as a single parameter. In principle nk can be differently valued for every pixel in the array, however to date we have assumed uniform trap density along each column and are only calibrating the x or CCD column dependence. Because it is linked to a fundamental CCD parameter (the size of the charge cloud before the charge is transfered), the value of α is assumed to be the same for all nodes of I3. The model presented in Section 6.1.3 of the MIT CTI Correction memo (Bautz et al., 2001) did not include the x dependence of nk and had a fixed value of $\alpha = 0.5$.

The data used to calibrate both nk and α are from the external calibration source. Since the normalization, nk, is fit separately for each column in the CCD, a large number of photons are required; precisely how many is under investigation. Both groups are using a years worth of all the "good" external calibration source data taken at -120C which yields ≈ 18 million photons on the I3 CCD. A description of the spectral features of external calibration source data can be found in the ACIS Calibration Report.

The power-law index, α, is found by fitting the charge loss versus energy function to data from each CCD node. Because α is best determined over a broad energy range, charge loss data from ten spectral lines, kindly provided by Penn State, are used in the fit. The data and the best fit model for I3 Node 3 at -120C are shown in Figure 1. Identifications of the lines and details of the charge loss determination can be found in Townsley & Broos (2000) and Townsley & Broos (2001). The fitted values of α for each node are very similar and a weighted mean is used in the corrector for the entire CCD.

For each column's eventlist, the CTI is fit at each of the three major calibration lines (Mn-Kα, Ti-Kα and Al-K at 5.9, 4.5 and 1.5 keV, respectively). The charge versus CTI for each column is fit to the charge loss function described above ($CTI(x, Q) = \delta Q/Q = nkQ^{\alpha-1}$). The power-law index α is fixed at the value determined previously for the whole CCD, so the only free parameter is nk. The nk values can then be used directly in the corrector to calculate charge loss.

No attempt has yet been made to calibrate the MIT approach for data on CCDs other than I3, at focal plane temperatures other than -120C or to expand the correction to include serial CTI effects important in modeling BI CCDs. Calibrating the model parameters, nk and α, for other CCDs and focal plane temperatures would require the same data and analysis as described above and should present no difficulties. To expand the model to include serial CTI effects would require calibrating both the row- and column-dependence of the charge loss separately, however in principle the same formalism would apply. In the first step the charge loss in the serial direction would be calibrated and removed, then the charge loss in the parallel direction. Some iteration would likely be required in the beginning to balance the serial and parallel charge loss but the eventual implementation would be very similar to that for FI CCDs.

3 PSU CTI Correction

The Penn State approach is an event reconstruction model in which a true event is hypothesized, degraded by a CTI model and compared to the observed event. This process is iterated until the

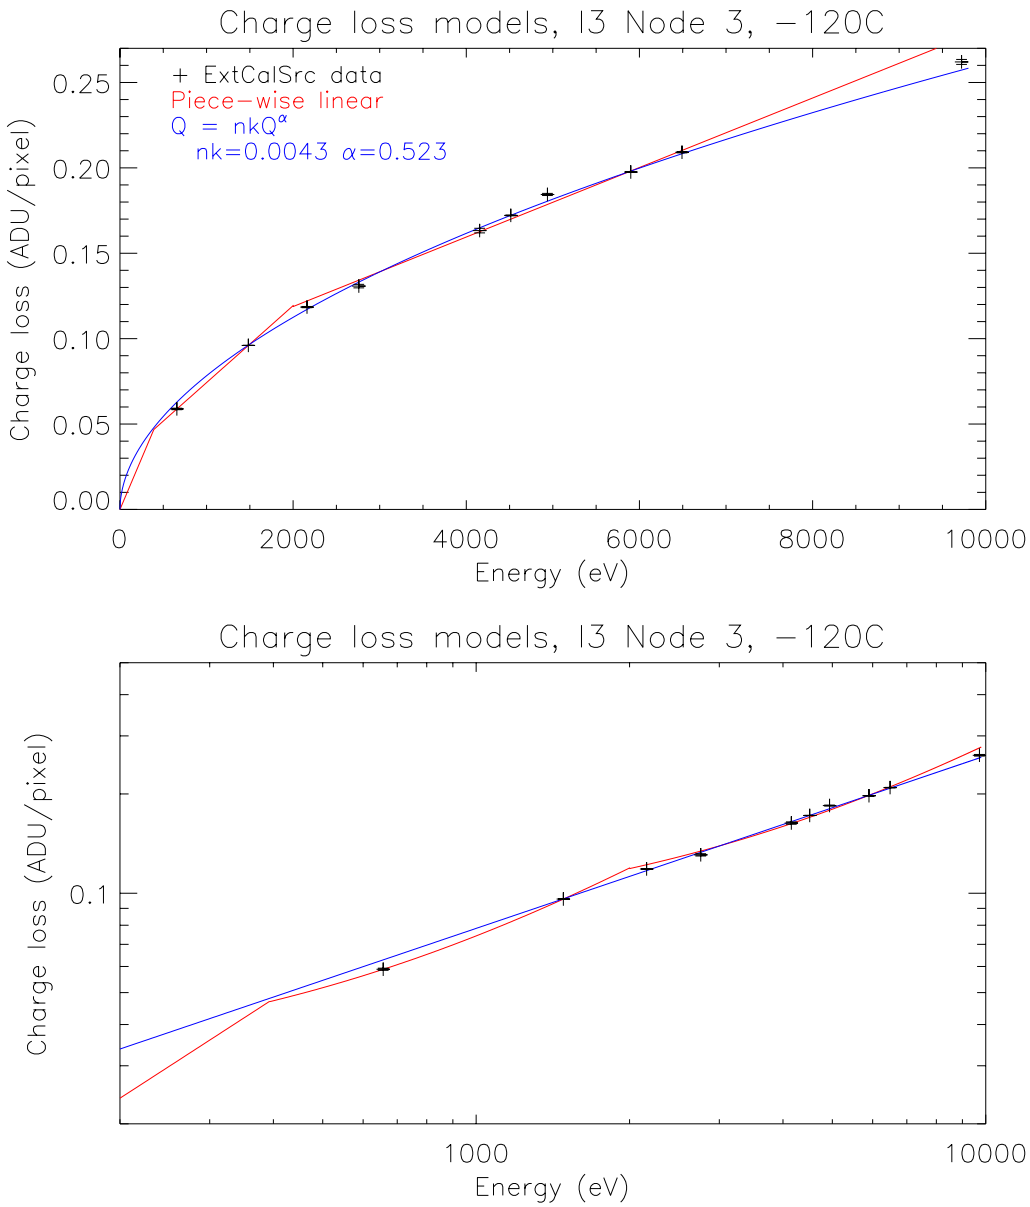

Figure 1: Comparison of the piece-wise linear (red) and power law (blue) charge loss models used by the Penn State and MIT CTI correctors in linear (top) and logarithmic (bottom) plots for a focal plane temperature of -120C. This is an example fit; other nodes and devices can be better or worse fit by either model. Data used to calibrate this relationship are also shown (courtesy of the PSU CTI corrector group).

modeled and observed events agree. Both charge loss and charge trailing are modeled. In this way both the charge loss and the grade morphing are corrected.

The charge loss as a function of energy is parametrized as a piece-wise linear function which is required to go through the origin. The total charge loss for an isolated charge packet is then

$$\Delta Q(y, Q) = y\, L(Q)$$

where Q and y are as before and $L(Q)$ is the charge loss per pixel transfer as a function of the charge in the pixel. At -120C, three lines are required to fit $L(Q)$. Figure 1 compares the current Penn State and MIT charge loss models for I3 Node 3 at -120C. It is not surprising that the two models have converged to be very similar since they both use the same calibration data. $L(Q)$ is calibrated separately for each CCD node.

Initially the Penn State corrector included a power-law parametrization of the charge loss. The switch to a piece-wise linear function was made because the power-law systematically underfit the charge loss at very high energies for both calibration and astrophysical data at -110C. This problem seems to be much less important for data at -120C.

The Penn State corrector also incorporates a model of the charge emitted into the first trailing pixel. This is again a piece-wise linear function that is required to pass through the origin and it is calibrated once for each CCD. The total charge trailed by an isolated charge packet is

$$\Delta Q_t(y, Q) = y\,T(L)$$

where ΔQ_t is the charge re-emitted into the trailing pixel and $T(L)$ is the charge trailed per pixel transfer as a function of $L(Q)$, the charge loss per pixel transfer defined above. Two lines are needed to fit $T(L)$ to insure that the function does go through zero. Figure 2 shows the data and the fit to the charge trailing function for all nodes of I3.

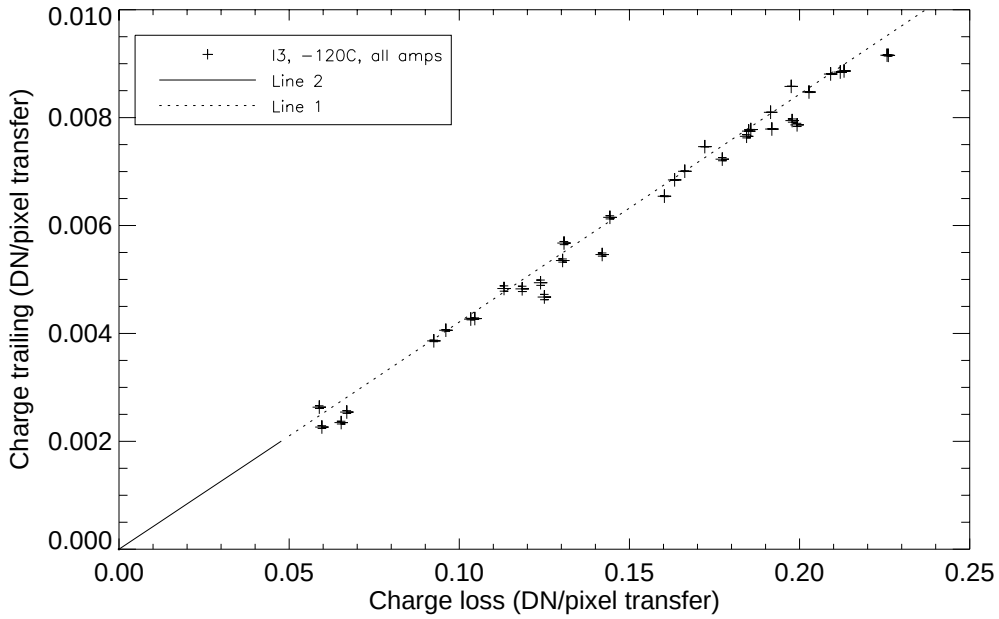

Figure 2: The piecewise-linear model of parallel charge trailing per pixel transfer for I3 (all nodes) at -120C. The Au-L line produced unreliable results and is not included.

In addition to correcting the amplitude for the charge lost to an event, the Penn State model attempts to restore the entire original event island. This is accomplished by a shielding model that compares the amount of charge in a pixel (Q_y) to that in the preceding pixel in the event island (Q_{y-1}) when computing the charge lost and trailed. If the preceding pixel is smaller ($Q_{y-1} < Q_y$), the effective charge lost is the difference between the calculated charge lost to the pixel and the

preceding pixel

$$\Delta Q^{eff}(y, Q_y) = \Delta Q(y, Q_y) - \Delta Q(y-1, Q_{y-1})$$

If the preceding charge is larger ($Q_{y-1} > Q_y$), the effective charge trailed into the pixel is the difference between the calculated charge trailed for the pixel and the preceding pixel

$$\Delta Q_t^{eff}(y, Q_y) = \Delta Q_t(y-1, Q_{y-1}) - \Delta Q_t(y, Q_y)$$

The shielding model is better described in Townsley & Broos (2001). One assumption made in this 'self-shielding' model is that trap time constants are longer than the pixel transfer time (40 μs) which is consistent with our knowledge of the trap parameters at -120C.

The Penn State CTI corrector works in the forward direction to reconstruct the event. For each observed event, a hypothesis is made as to the true event, i.e. the event that would have been obtained if no CTI effects were present. In the present incarnation of the code, the initial guess at the true event is the observed event itself. The charge lost and trailed are computed for each pixel above the split threshold in the true event as described above. The CTI-degraded true event is compared to the observed event and the difference between them (observed - CTI degraded true) is added to the true event. This iterative adjustment continues until the CTI-degraded true event and the observed event differ by less than 0.1 ADU. The observed event is then replaced by the true event.

To correct for pulseheight variations along both columns and rows a deviation map is applied via a two-dimensional map with an energy-dependent amplitude

$$D(x, y, P) = D_0(x, y)(1 + gP)$$

where D_0 is a two-dimensional map of deviations at 0 pulseheight, x, y and P are the position and summed pulseheight of the event, and g is the amplitude to pulseheight scaling of the deviations. Note that the deviation map is applied to the center pixel pulseheight of the event, and not the surrounding pixels in the event island. Deviation maps are constructed from CTI-corrected data using maps of the pulseheight variation after CTI correction, binned by 1 pixel in x and by 16 pixels in y, of the three strongest lines in the external calibration source. The deviation maps at each spectral line are scaled to zero pulseheight in a number of averaging and smoothing steps. The process of creating the deviation map is further described in Townsley & Broos (2001).

4 Performance comparison

Our primary figure of merit for judging any CTI mitigation approach is the full width at half maximum (FWHM) of spectral features at the top of the CCD (far from the readout). On the I3 CCD this region corresponds to the ACIS I-array aimpoint. Except where noted, only the G02346 events are used in this comparison. The line width versus row is calculated for the three strongest lines of the external calibration source; Al-K at 1.5 keV, Ti-Kα at 4.5 keV, and Mn-Kα at 5.9 keV. Each row bin includes 32 rows to minimize the pulseheight differences across the bin. FWHM values in eV are shown in Figure 3 and use the local gain in each bin. Using the local gain slightly overestimates the FWHM of data without any pulseheight correction relative to corrected data. The + symbols correspond to the performance of I3 with no CTI correction which is equivalent to the performance with current CXC calibration products. The * symbols are the MIT CTI correction reported in Bautz et al., (2001) which did not include correction for variations of trap

density between columns. The diamonds and squares indicate the current MIT and PSU CTI corrector performance. Also shown in Figure 3 are the theoretical limits on the performance of I3 (as discussed in Section 8 of Bautz et al., (2001)) and the current performance of the BI CCD S3 (without CTI correction).

Figure 3 is also a nice demonstration of the energy dependence of the two parts of the correction by comparing the earlier MIT results that did not include charge loss variations in columns to the later results. The higher energy Mn-Kα line receives most of its improvement from the pulseheight correction of the event island while the lower energy Al-K line receives much more benefit from the correction of column to column variations.

The spectral resolution performance of both CTI correction approaches is very similar. Figure 4 is an expanded version of the previous figure which emphasizes the differences between the models. While the performance of both correctors is very similar at the top of the CCD, the PSU corrector performs slightly better on the lower part of the CCD (rows $\sim$ 100-500). We believe this is because the current MIT model does not include any variations in trap density along each column even though such variations are known to exist (Section 4.1 of Bautz et al., 2001).

An important feature of CTI corrected data is that the gain or energy scaling should be constant with position within a CCD node as it is for undamaged CCDs. Correcting for deviations along columns as well as rows is required to produce this. Both CTI correction approaches have small residual variations in pulseheight versus row. Without correction for variations along columns the deviations are of order 20-30 eV; including deviation corrections reduces this to 10 eV or less.

Since the two correction approaches yield very similar results for our primary figure of merit, it is worthwhile to examine differences in other event properties. The Penn State corrector changes the distribution of charge in the event island and, therefore, the event grade while the MIT corrector does not. Figure 5 shows the ratio of counts in the Mn-Kα line at 5.9 keV for corrected events to uncorrected events for the two approaches (G02346 events). As expected the MIT corrector does not change relative numbers of G02346 events. The increase in this ratio at high row numbers for the Penn State corrector is a result of grade de-morphing in the event reconstruction. The additional events are real events whose grade was morphed into "bad" grades (G157) by CTI degradation. While at -120C the difference is small ($\sim$ 1% at the top of the CCD), at -110C the increase in quantum efficiency at the top of the CCD by recovery of morphed grades is much larger, $\sim$ 10% at 5.9 keV (Townsley & Broos 2000).

Figure 6 further demonstrates which grades are affected by CTI-induced grade morphing and how well the Penn State event reconstruction recovers the original event. For an undamaged detector with no charge-transfer inefficiency, the fraction of events with any particular shape should be constant as a function of position (under the assumption that the input spectrum is the same at all positions). Figure 6 shows the ratio of events in specific grades to the total number of events over all energies in the external calibration source spectrum. Grade ratios and therefore grade morphing are energy dependent, so these results will vary for spectrally harder or softer sources. Figure 6 highlights the y-dependent structure of some interesting grades. In particular for the unreconstructed data, the number of single pixel events (G0s) decreases with row number while the number of upwardly singly split events (some G2s) increases with row number because of charge trailing from the center pixel into the following pixel. The PSU corrected data has a nearly flat distribution of G0 and upper split G2s. Smaller effects are seen for other grades. For data at a focal plane temperature of -110C or with higher CTI, the grade morphing effects can be much larger.

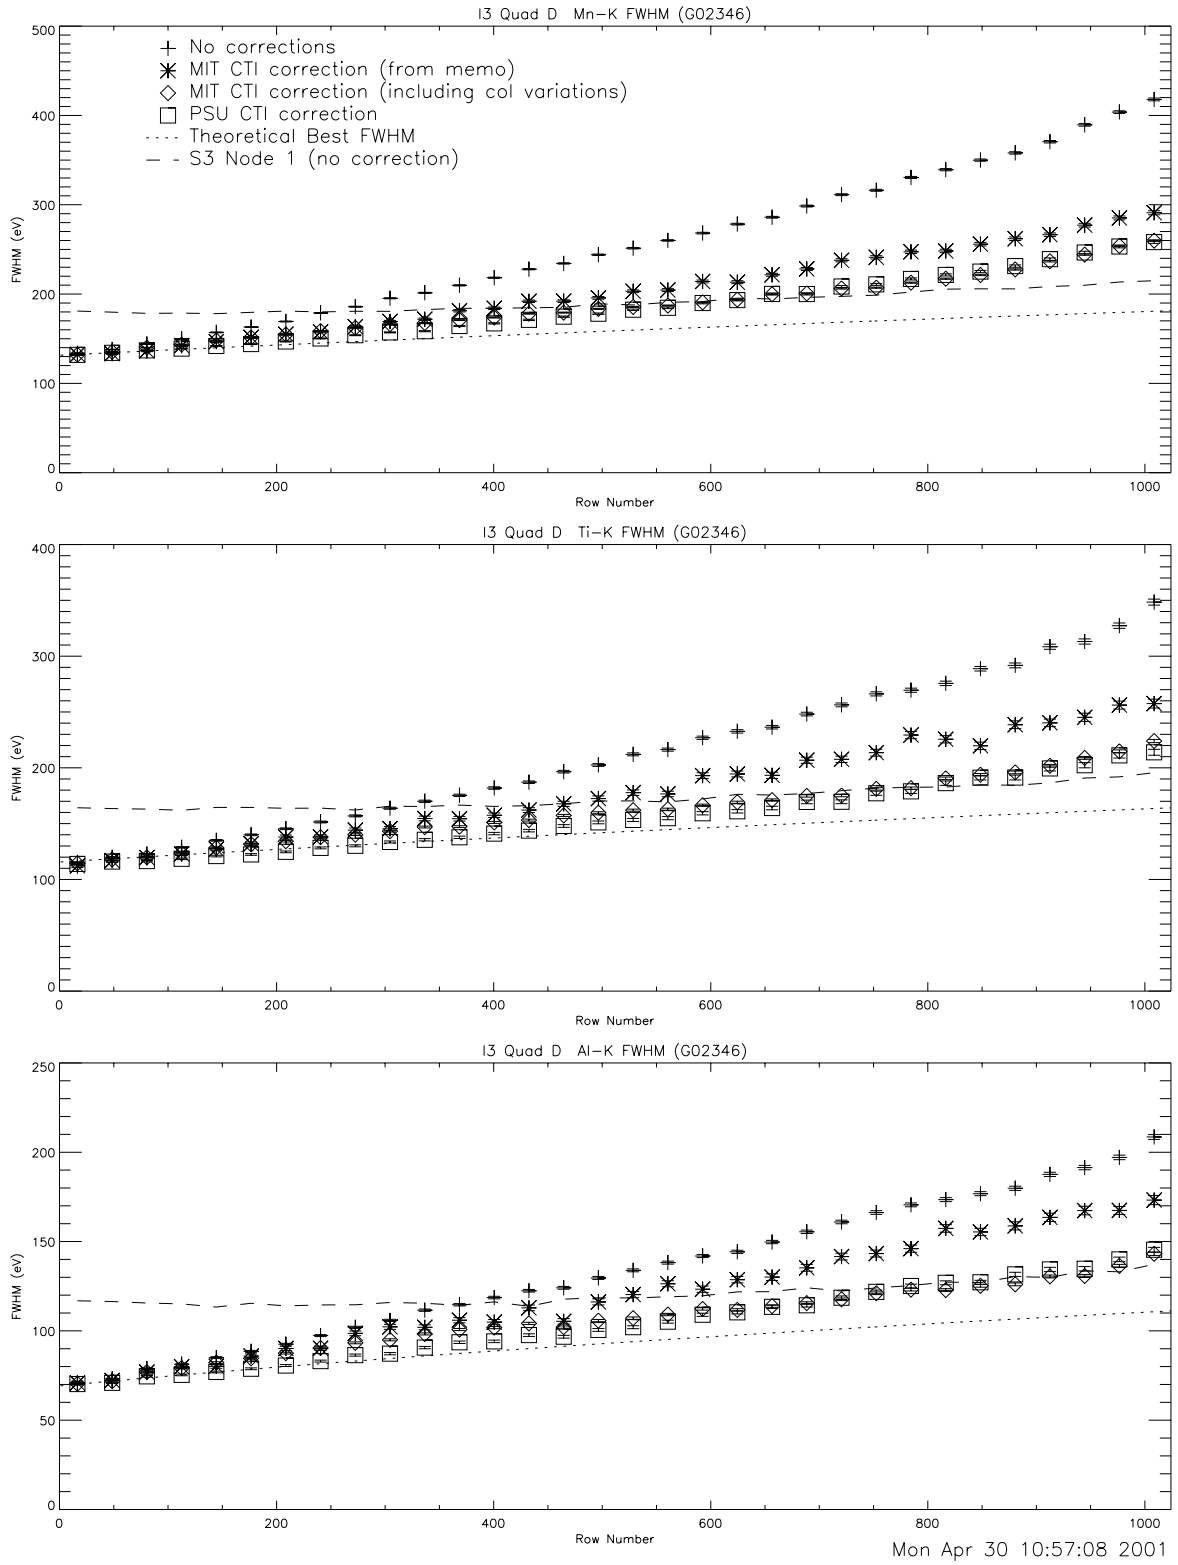

Figure 3: FWHM versus row at 5.9, 4.5 and 1.5 keV. Data corrected using the current MIT (diamonds) and PSU (squares) CTI correctors perform very similarly. Also shown are uncorrected data (+), the MIT CTI correction reported earlier (*), the theoretical limits on I3 performance and the current performance of S3 (without CTI correction). The error bars shown are from counting statistics and are therefore very small.

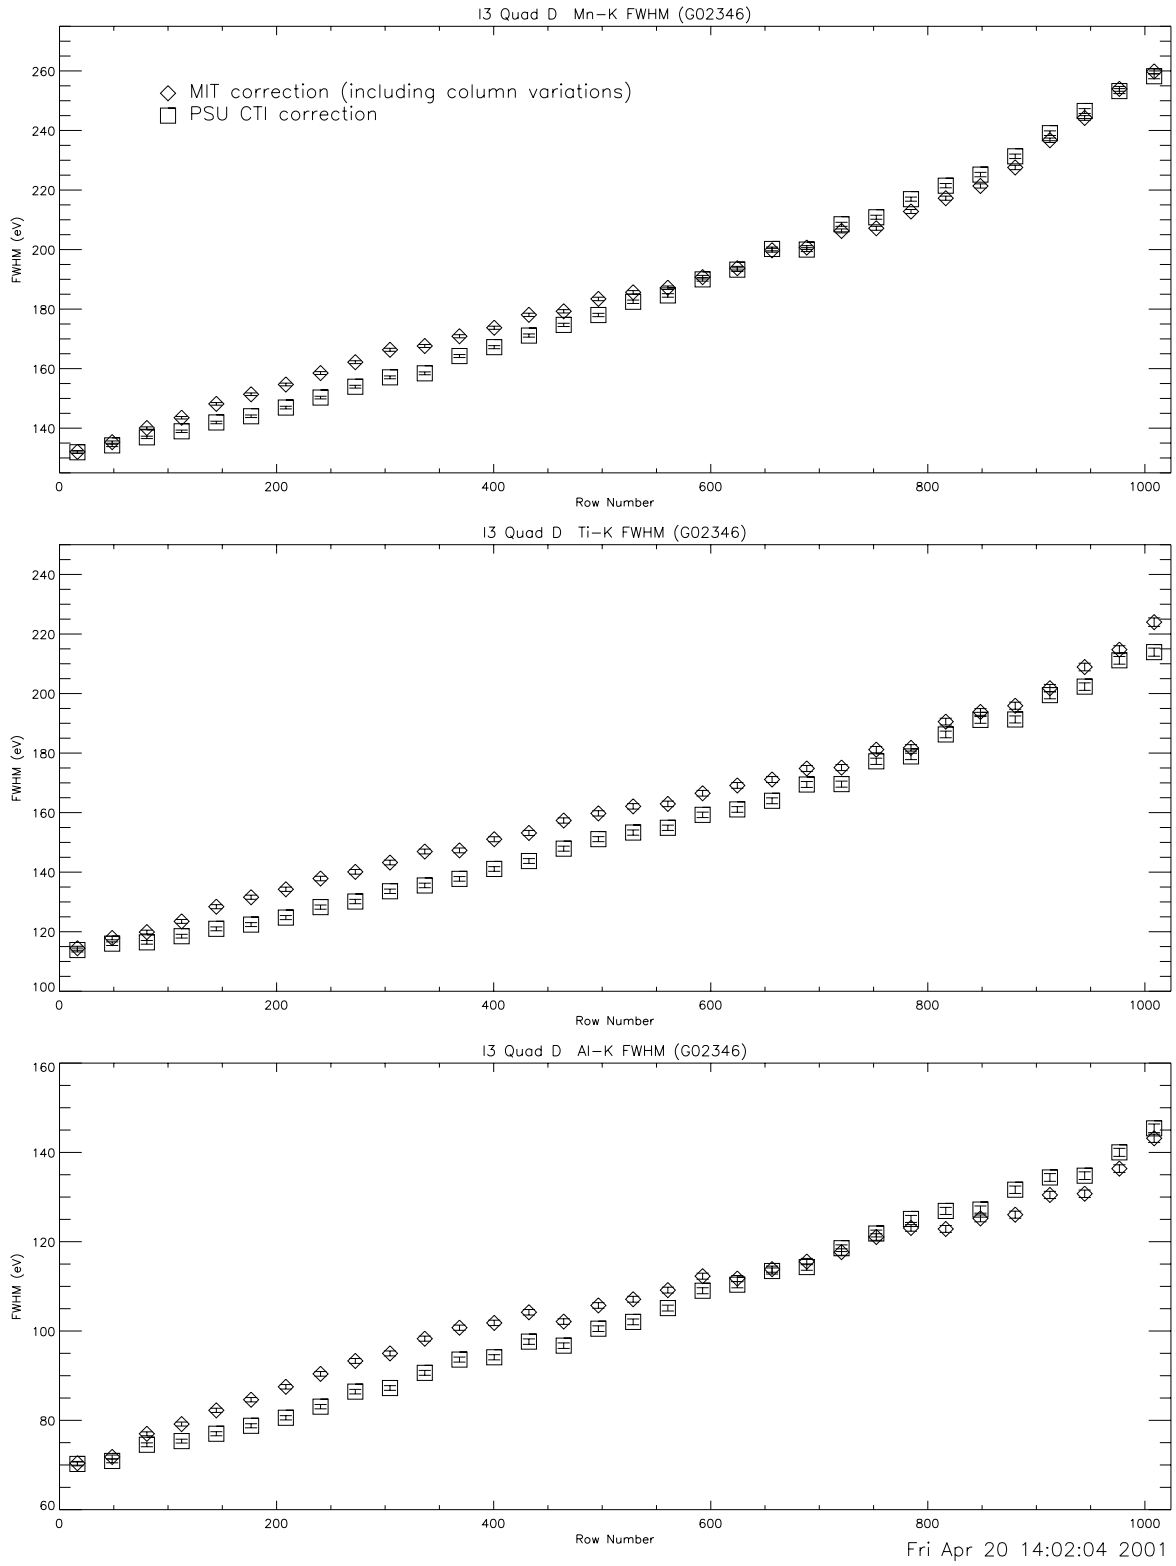

Figure 4: FWHM versus row at 5.9, 4.5 and 1.5 keV. This is the same data as Figure 3 with the y-axis range decreased to emphasize the differences in the MIT and PSU CTI corrector performance. The error bars shown are from counting statistics and are therefore very small.

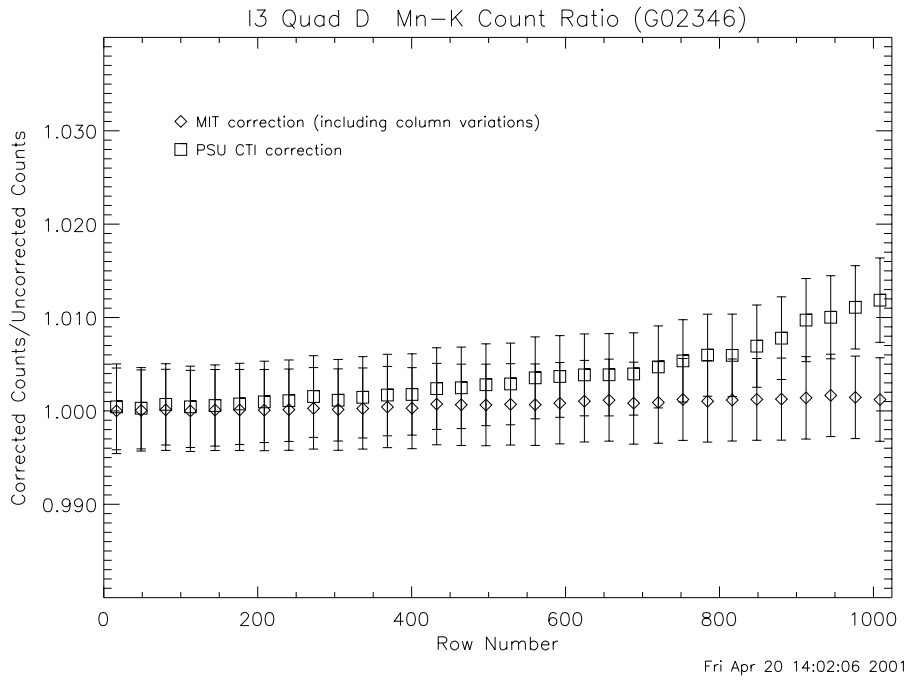

Figure 5: Ratio of corrected to uncorrected counts at 5.9 keV (G02346 events).

5 Differences in the Two Approaches

As discussed above, the two CTI correction approaches have converged to similar forms because they both attempt to model the same processes, but there remain a few differences in the execution and parametrization of the models. These differences are as follows:

- Parametrizing the shape of the charge loss versus energy function as a power-law or a piecewise linear function,
- Including position-dependent pulseheight variations in the physical model of charge loss or as an empirical correction to the center pixel pulseheight,
- Modeling both charge loss and charge trailing to reconstruct the undamaged event grades, and
- Applying the correction directly or as a forward event reconstruction model.

Since our primary figure of merit, spectral resolution at high row number, is very similar for both models at -120C, other considerations may be important in discriminating between choices.

5.1 Parametrization of charge loss

At a focal plane temperature of -120C, the differences between the two parametrizations of the charge loss function are insignificant; while the requirement to pass through the origin is more natural for the power-law, the advantage is slight. What may be most important is the adaptability of the charge loss model to conditions not directly discussed here, particularly that at -110C. Townsley

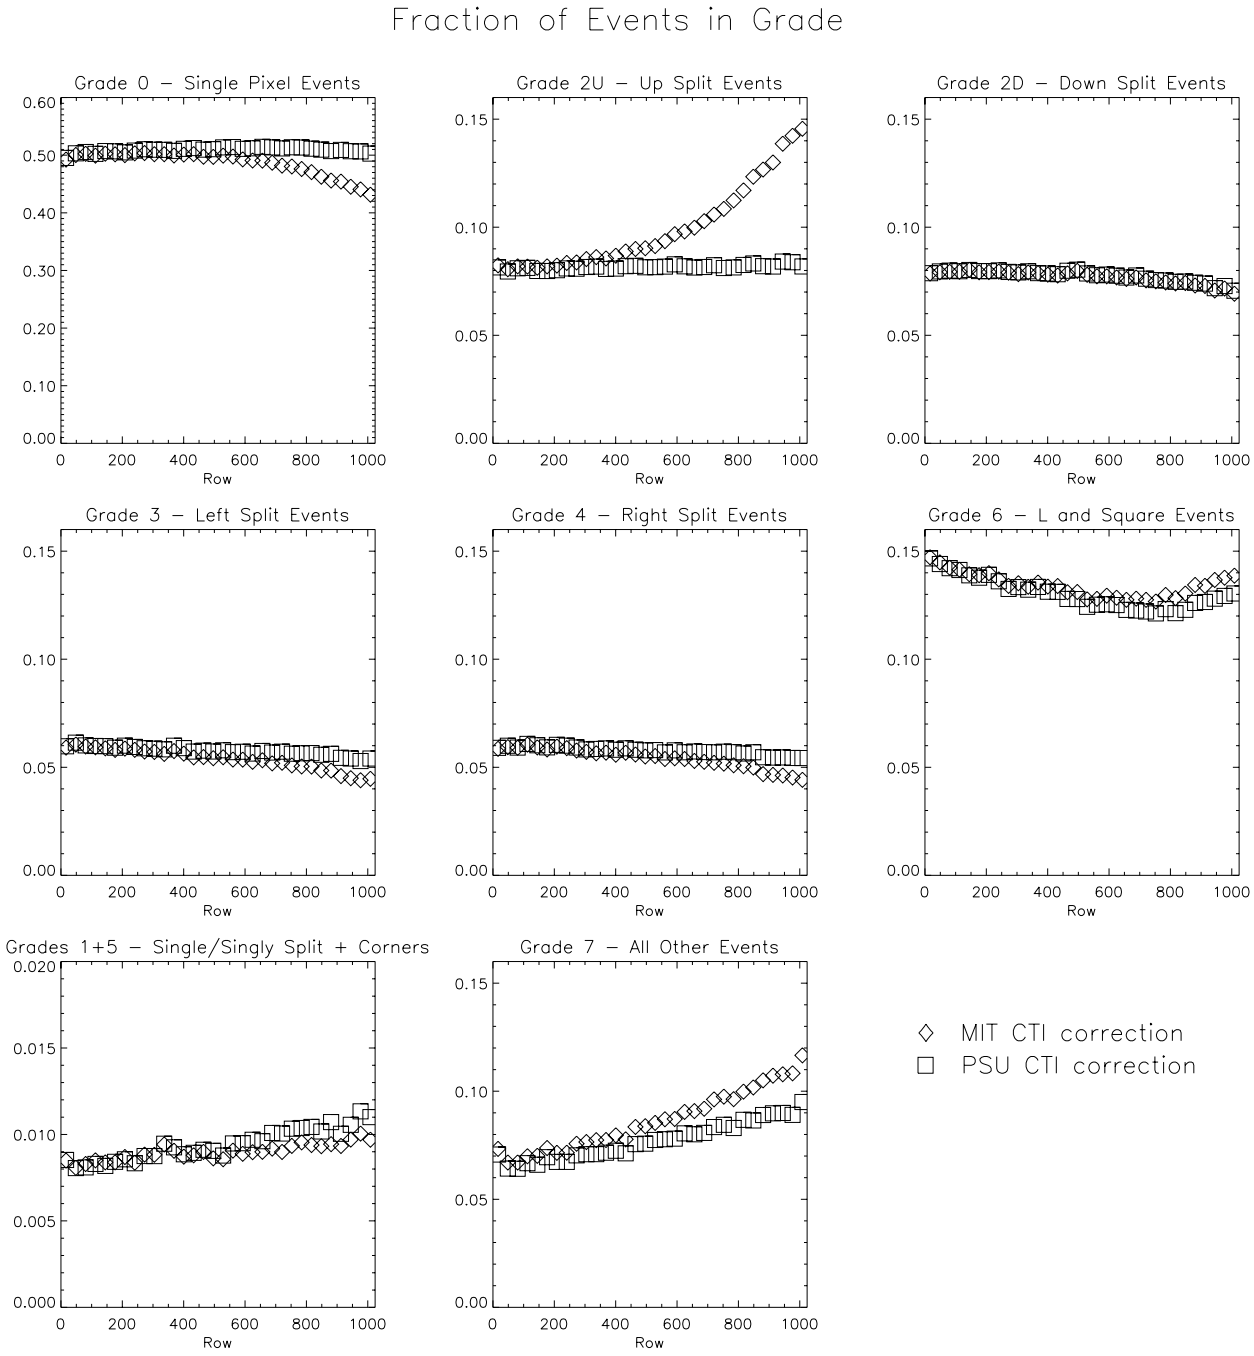

Figure 6: Ratio of events in various grades to total telemetered events as a function of row. Note that the y-axis range is different for the G0 and G1+G5 events. The MIT corrector does not change the observed event grade.

& Broos (2000) found that high energy data at -110C, both calibration source and astrophysical, could not be well fit by any power law. An example fit at -110C is shown in Figure 7; other nodes and devices have better or worse fits at high energies. If a CTI corrector is to be applied to all ACIS data, it must be accurate at all useful focal plane temperatures. In addition, future radiation damage may cause data at -120C to behave more like that at -110C. The best model for charge

loss at -110C might be a power law that transitions to a linear fit at high energies. These are strong arguments for a CTI corrector that is flexible enough to accommodate different charge loss parametrizations as CCD conditions change and our calibration at high energies improves.

5.2 Correction for position-dependent pulseheight variations

Much of the benefit of the CTI correctors is a result of the correction of spatial variations in the CTI in each column. Figure 4 demonstrates that trap variations *along* each column cannot be ignored for the best possible performance at all positions. This does not preclude the MIT paradigm of including the variations in the model trap density, but the details of calibrating such a model remain to be seen. We would like, as much as possible, to ground the CTI corrector in physical models of the charge loss. If the pulseheight deviations are a result of the non-uniformity of trap density then the energy dependence of the charge loss model should be included in the deviation correction, either explicitly in the charge loss correction or in the energy-dependence of the deviation map.

5.3 Reconstructing the Event Grades

Reconstructing the event island has a number of benefits all related to the restoration of event grades. The grade morphing by CTI causes some X-ray events to placed into event grades that will be thrown out in standard analyses; many of these events can be recovered by reconstructing the original event grade. This improves the effective quantum efficiency of the detector, especially at high row number and at high energies. The loss of QE due to grade morphing increases with CTI so data taken at -110C or future data with additional radiation damage would most benefit.

A second benefit of grade reconstruction is the wide-field grade "regularization". By removing the position-dependent effects of CTI on event grades, a particular grade designation means the same thing across the focal plane. Additional more specialized grade filtering can then be done uniformly for an extended object or a wide field observation. Here are two specific examples that could benefit from additional grade filtering that would only be valid with position-independent event grades. If an astronomer with an extended source wanted to get better spectral resolution than that with G02346 events, she can use G0 events. Additional grade filtering beyond G02346 can also be done to improve background rejection (for an example see Brandt et al., (2001)). Neither of these could be fairly done without first regularizing the event grades.

5.4 Direct or forward modeling

The performance of a simple CTI corrector does not seem to be linked to the direction in which the charge loss model is applied. A forward model can be directly incorporated into a CCD simulator as it is at Penn State, and future advances in modeling pile-up could be easily incorporated into a forward CTI corrector. LYNX, a spectral fitting tool under development at Penn State, incorporates a forward modeling approach including MARX and the Penn State CCD simulator. A forward CTI corrector could be easily included in such a tool.

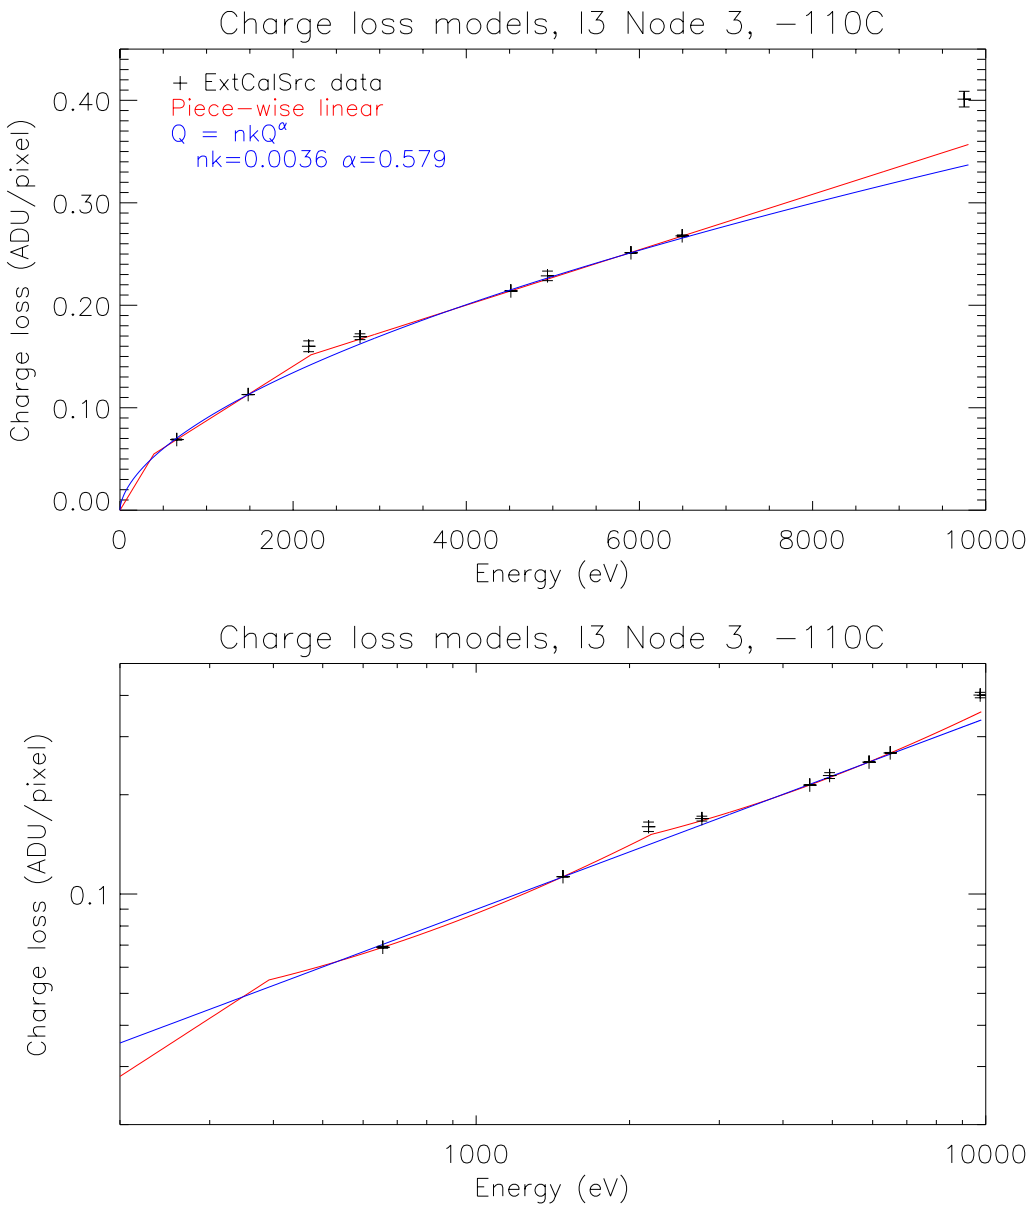

Figure 7: Comparison of the piece-wise linear (red) and power law (blue) charge loss models used by the Penn State and MIT CTI correctors in linear (top) and logarithmic (bottom) plots for a focal plane temperature of -110C. This is an example fit; other nodes and devices can be better or worse fit at high energies. Data used to calibrate this relationship are also shown (courtesy of the PSU CTI corrector group).

6 Preliminary ACIS Team Recommendations

This section outlines the major options available to the CXC in implementing a CTI correcting software tool. We also list the tradeoffs involved in these choices. There are two fundamental event parameters that can be corrected for CTI: event amplitude and event grade. For each event parameter, the possible correction choices are listed, from the simplest to the most complex. A related recommendation for more flexible calibration products is also discussed.

6.1 Event amplitude correction

Choices: The possible choices for event amplitude correction are:

a) "Uniform" trap density for each CCD node

b) Column-to-column variable trap density

c) Trap variation within a column.

Tradeoffs & implementation issues:

- The finer the spatial scale of the correction, the greater the improvement in response (see section 4), and the more difficult the calibration task.
- An implementation must choose how to represent the charge loss vs energy relationship (power law, piecewise linear, other.)
- Another implementation issue is how to model the spatial variations within column (phenomenological deviation map vs trap density map.)
- We have a general concern about calibration accuracy at E $<$ 1 keV (The Mn-L complex is the only calibration source line feature; this is weak and complicated). At the same time, the correction is expected to become a progressively larger fraction of pulse height at progressively lower energy, so the accuracy of the correction is correspondingly more important. Moreover, we know that our physical models do NOT accurately predict the actual spectral resolution after CTI correction, presumably because precursor events are not properly modeled. Any accurate "model" of the CCD will therefore contain *ad hoc* corrections to spectral resolution Of course, this is true whether we apply CTI corrections or not. This ignorance of the physics implies that we must rely more heavily on empirical measurements to obtain an accurate response model. We probably need more validation with astrophysical sources such as E0102. Of course, this is true whether we apply CTI corrections or not.

6.2 Event grade correction

Choices: The possible choices for event grade correction are:

a) none

b) PSU method

c) PSU method with slightly modified trailing model

Tradeoffs & implementation issues

- Grade correction makes quantum efficiency and grade distributions more uniform at -110C and allows more flexible use of grades (and thus, for example, better background rejection for selected observations) at -120C. The cost is some additional complexity in the correction code, additional calibration requirements and, possibly, additional computation time.
- An implementation issue: Is forward modeling required/appropriate?
- Another implementation issue: would nine corrected pulseheights be saved, or just revised grades?
- Should we consider developing corrector in two generations (first amplitude only, second amplitude and grade)?

6.3 Response matrix generation

The addition of CTI correction to the tools available to the Chandra user community, while beneficial scientifically, doubles the number of basic ACIS calibration products. The current paradigm of response matrix generation in which the response is parametrized in the FEF file allows for much flexibility in rebinning and interpolating the products but is expensive in terms of the manpower required to produce the calibration. We would like to recommend that the CXC consider an alternate method, that of producing RMFs and ARFs directly from simulated event lists. A demonstration of this method can be found in Broos & Townsley (2001). This method requires more CPU time to produce simulations at many energies but much less human intervention. The benefit of this approach is that the same simulated eventlists can be used to produce response matrixes for any subset of the events, such as filtering by position or event grade, and for either CTI corrected or uncorrected events.

7 References

M. Bautz, P. Ford, C. Grant, S. Kissel, B. LaMarr, & G. Prigozhin 2001, "ACIS CTI Correction," http://space.mit.edu/ACIS/iacis.html

P. Broos & L. Townsley 2001, "Notes on RMF Generation", ftp://ftp.astro.psu.edu/pub/townsley/rmf.ps

L. Townsley, P. Broos, G. Garmire, & J. Nousek 2000, "Mitigating Charge Transfer Inefficiency in the Chandra X-ray Observatory's ACIS Instrument," ApJ 534, L139

L. Townsley & P. Broos 2000, "The PSU CTI Corrector: New results for I3, -110C," http://www.astro.psu.edu/users/townsley/cti/I3memo/

L. Townsley & P. Broos 2001, "Modeling Charge Transfer Inefficiency in ACIS," ftp://ftp.astro.psu.edu/pub/townsley/method_memo.ps

8 Projects which have benefited from the PSU CTI Corrector

Brandt, W.N. et al. 2001, "The Chandra Deep Survey of the Hubble Deep Field North Area. IV. An Ultradeep Image of the HDF-N", AJ in press (astro-ph/0102411)

Burrows, D.N. et al. 2000, "The X-ray Remnant of SN 1987A", ApJ 543, L149

Garmire, G. et al. 2000, "Chandra X-Ray Observatory Study of the Orion Nebula Cluster and BN/KL Region", AJ 120, 1426

Griffiths, R. et al. 2000, "Discovery of Hot Gas in the Core of M82", Science, 290, 1325

Ho, L.C. et al. 2001, "Detection of Nuclear X-ray Sources in Nearby Galaxies with Chandra", ApJ 549, L51

Hornschemeier, A.E. et al. 2000, "X-Ray Sources in the Hubble Deep Field Detected by Chandra", ApJ 541, 49

Hornschemeier, A.E. et al. 2001, "The Chandra Deep Survey of the Hubble Deep Field North Area II. Results from the Caltech Faint Field Galaxy Redshift Survey Area", ApJ in press (astro-ph/0101494)

Imanishi, K., Koyama, K., & Tsuboi, Y. 2001, "Chandra Observation of the Rho Oph Cloud", ApJ in press (astro-ph/0104190)

Maeda, Y. et al. 2001, "A Chandra Study of Sgr A East: A Supernova Remnant Regulating the Activity of our Galactic Center?", ApJ in press (astro-ph/0102183)

Tsuboi, Y. et al. 2001, "Discovery of X-rays from Class 0 Protostar Candidates in OMC-3", ApJ 554 in press (astro-ph/0103401)

Tsunemi, H., et al. 2001, "Improvement of the Spatial Resolution of ACIS Using Split Pixel Events", ApJ submitted