Kalman Filtering in Chandra Aspect Determination

Roger Hain, Thomas Aldcroft, Robert Cameron, Mark Cresitello-Dittmar, M. Karovska
Harvard-Smithsonian Center for Astrophysics

The ability of the Chandra X-Ray Observatory to achieve unprecedented image resolution is due, in part, to the ability to accurately reconstruct the spacecraft attitude history. This is done with a Kalman filter and Rauch-Tung-Striebel (RTS) smoother, which are key components of the overall aspect solution software. The Kalman filter/RTS smoother work by combining data from star position measurements, which are accurate over the long term but individually noisy, and spacecraft rate information from on-board gyroscopes, which are very accurate over the short-term, but are subject to drifts in the bias rate over longer time scales. The strengths of these two measurement sources are complementary. The gyro rate data minimizes the effects of noise from the star measurements, and the long-term accuracy of the star data provides a high-fidelity estimate of the gyro bias drift. Analysis of flight data, through comparison of observed guide star position with expected position and examination of the reconstructed X-ray image point spread function, supports the conclusion that performance goals (1.0 arcsecond mean aspect error, 0.5 arcsecond aspect error spread diameter) were met.

An enhanced kalman filter module which will function with degraded or limited sensor data is currently being studied. Advanced estimation techniques to increase robustness of the algorithm may be able to maintain aspect accuracy in the case of a problem with on-board aspect hardware, such as gyro or aspect camera failure

This schematic illustrates the basic flow of the Kalman filter design for Chandra aspect determination. The main purpose of the filter is to maintain an accurate estimate of the spacecraft angular position, or attitude, at any point in time. Together with fiducial light data for estimating spacecraft flex, this attitude estimate is used to accurately determine the source in the celestial sphere of X-rays detected by the science instruments.

The Filter estimates two main quantities: error in current attitude estimate, and gyroscope drift rate. Each of these quantities is estimated about three axes, resulting in six estimated quantities referred to as the filter "state". The schematic to the right illustrates the basic flow of the filter. Beginning with an initial estimate of the state, and initial statistics which characterize how well the state estimates are known (the "covariance"), the filter iterates through two steps.

First, as shown in the upper left box, gyro data is received. The state estimate of gyro bias is subtracted from the raw gyro data, and the spacecraft attitude is updated according to the angular motion derived from the corrected gyro data. The error statistics associated with the state are also updated, since uncertainty in gyro bias will grow over time, and attitude errors due to rate uncertainty will also increase. In this way the state and attitude estimate are maintained and propagated forward in time until a star camera measurement is available.

The second step is shown in the upper right box. A measurement (the location of a guide star) is received from the aspect camera and incorporated into the filter's state estimates and error statistics. The filter mathematically compares the new attitude estimate based on the current measurement with the expected attitude from the prior propagation step. The error statistics of the expected attitude and the expected error of the measurement (the noise from the star camera) are taken into account, and new estimates of attitude error and gyro bias drift are calculated. Since more data leads to a more accurate attitude estimate, the error statistics are updated to reflect more accurate attitude knowledge, and the cycle begins again with additional gyro data.

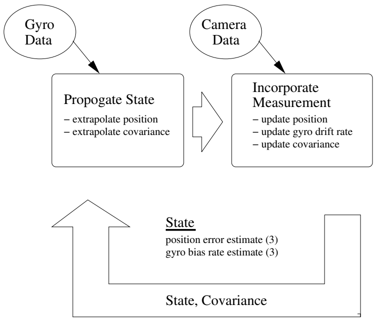

The data shown on the right demonstrates the quality of the aspect solution with current Chandra data. The top two charts show the difference between a known guide star position and the aspect software's best estimate of that guide star's position over a 50,000 second observation. Both top charts show the same data, but the one to the right has fixed axes at +/- 1 arcsec in each direction, and the star has been centered at (0,0). While the absolute position of the star is quite good in this data (near 0.1 arcsecond radius/axis), calibration accuracy has a significant effect on absolute position accuracy. The recentered image on the right more clearly reflects the contribution of the Kalman filter and smoother in particular to the quality of the aspect solution.

The data contributing to this top right image is plotted in the bottom two charts to the right. The middle chart shows the difference between the known star location and the best estimate of star location (adjusted for zero-mean) along the spacecraft x-axis as a function of time. The bottom chart shows the same quantity for the spacecraft y-axis. The calculated root mean square (RMS) values for this data are shown on the chart, and are 0.05 arcseconds and 0.04 arcseconds, respectively. This is equivalent to a 0.13 arcsecond RMS diameter, significantly better than the image reconstruction requirement of 0.5 arcsecond RMS diameter.

Since other errors can contribute to overall image reconstruction, the charts below expand on an examination of aspect quality by calculating X-ray image aspect reconstruction. The examination of the guide star image reconstruction, however, makes clear the ability of the Kalman filter to provide a highly accurate spacecraft attitude estimate on which to build overall aspect accuracy.

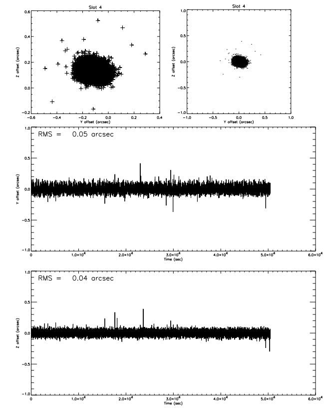

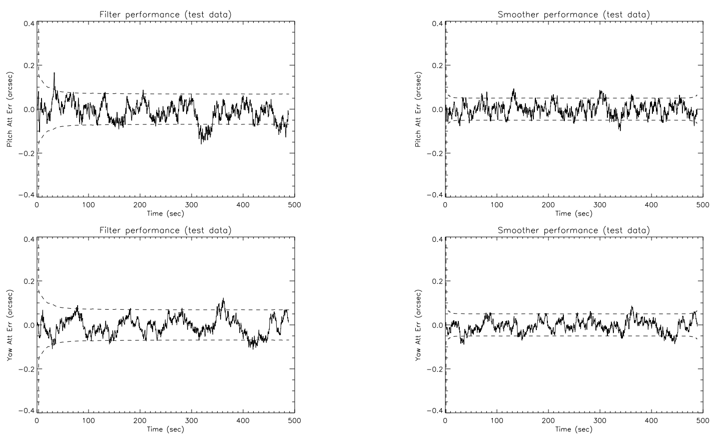

The two charts shown above on the left show the aspect pipeline Kalman filter performance on test data, not flight data, and are shown primarily to illustrate the behavior over time of the filter. In the upper left graph, the solid line shows actual error in pitch attitude estimate for a 500 second nominal dither observation. The dashed lines represent the root mean square (RMS) value of the covariance for pitch error (i.e., the filter's own estimate of how accurate approximately 2/3's of yaw estimates will be). Note that in this plot and the similar plot below it showing error in yaw estimate, the error in each axis quickly converges to a steady state value. The two charts shown above on the right demonstrate the improvement in yaw and pitch estimate error with the addition of a smoothing algorithm. The use of a smoother is possible in the aspect determination system since it is not necessary to calculate the best estimate in real time (as it is, for example, in the on-board flight software, which uses only a Kalman filter). The smoother allows better overall estimates by utilizing state and measurement data from the entire time period to improve the estimate at each point in time. The filter alone is not able to use future data to estimate any given point. The improvements with a smoother are smaller overall attitude error and much improved performance in the very early part of the data, as seen by the much smaller RMS values. .

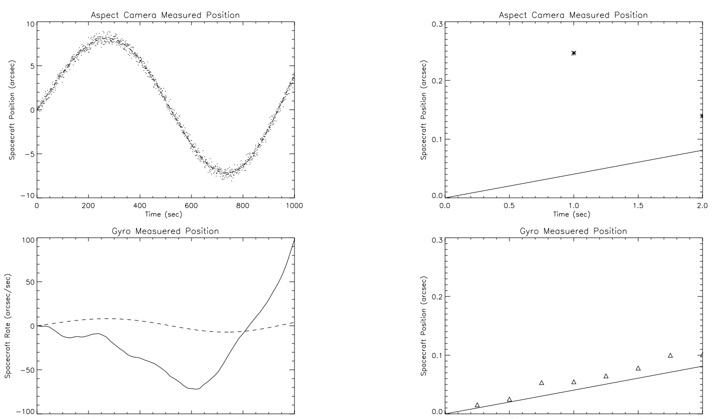

The two charts shown above on the left are shown to highlight a key feature of the Kalman filter - the ability to optimally combine (in the least-squares sense) data which is in one case accurate over long time periods but noisy in the short term, and data which has the opposite characteristics of being accurate over short time intervals, but degrades over time. The top left chart shows a characteristic change in attitude for a single spacecraft axis, similar to the kind of motion typically seen during dither on Chandra. Overlaid on top are estimates of the axis' position derived from aspect camera type data. The estimates generally track this axis well over time. Contrast this with the lower left chart, showing the same axis motion tracked only with gyro data. As the unknown gyro bias drifts over time, the rate error grows as a random walk, and diverges significantly as time increases. These two charts demonstrate the superior ability of the camera data to provide long term accuracy. In contrast, the two charts above on the right show the same data, but only for the first two seconds. In this case, the star camera data, shown at the top right, is much more in error than the gyro estimates shown below. The reason for this is simply that for such a short time period, the gyro bias rate has not drifted much, causing the gyro data to be relatively accurate. One key strength of the Kalman filter is the ability to use the short term accuracy of the gyro data to reduce attitude error from noise on star camera data, and at the same time use the long term accuracy of the camera data to estimate gyro bias drift rate, resulting in a good quality aspect solution over all time intervals.

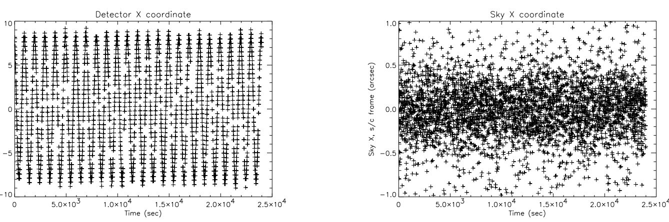

To further quantify overall aspect quality, a 50,000 second observation was examined. A point source was identified in this observation, and a circular region around the point source with a radius of approximately 1 arcsecond was isolated. The chart above on the left shows the detector X coordinate of the events from this region for the first 20,000 seconds of the observation. The 1000 second period dither pattern is clearly visible. The chart to the right shows the aspect corrected X coordinate for the same 20,000 seconds, resulting from sky coordinates transformed into the spacecraft reference frame. The effect of this transformation is to make any dither frequency remaining in the sky coordinates after aspect processing more clearly visible as a single frequency oscillation. While the overall aspect accuracy of the data is good, it is difficult to see whether any residual oscillation at the dither frequency is present with the background noise still present. The charts below show an attempt to further process the data to identify oscillations at the dither frequency by reducing background noise.

Residual aspect errors after processing are expected to be most visible near the dither frequency. The transformed sky coordinates shown in the above right chart do not show any visible oscillations, but the amount of noise is large compared to potentially small amplitude oscillations. One way to try to remove noise due to background events and make small amplitude oscillations more visible is to phase modulate the data at the dither frequency. This is done by resetting event times to "time mod(dither period)". This technique was used to create the data shown in the top charts at the right. Time only goes from 0 to the dither period, and events later than an integer number of dither periods are wrapped back to start at time 0 again. In this way, any residual aspect error at the dither frequency is preserved in the data, but random noise can be reduced by averaging. The top charts to the right show phase modulated data for the sky coordinates transformed to the spacecraft X axis (left chart) and spacecraft Y axis (right chart).

The two charts in the middle at the right show similar data for the X and Y spacecraft axes. In this case the data has been binned on 50 second intervals, with the intention of reducing the noise. Ideally, if no residual oscillation were present, it would be expected that the binned data points would become smaller and smaller as the bin size becomes larger, without showing any noticeable constant offset. If, on the other hand, a clear sinusoidal signal were present, this should become apparent in the structure of the binned data points. Although even the worst case binned data point is very small (less than 0.05 arcseconds), some slight evidence of a residual signal may be present in the Y coordinate charts. The binned data shows a disproportionate number of positive data points in the first half of the dither interval, and a disproportionate number of negative data points in the second half. Although difficult to see, it can also be argued that a small oscillation is present in the top right chart, showing the data before binning.

The two charts on the bottom at the right show the power spectrum of the original data in the spacecraft frame, before phase modulation and binning. Again it is difficult to see a strong trend, but a small increase in power seems present in two points with periods greater than 800 seconds, and one point with a period near 600 seconds. Further investigation is needed to determine whether this small oscillation represents a real trend or is simply an artifact of this data. But it is clear that any oscillations near the dither frequencies are very small, and that the Kalman filter and smoother have minimized these residual oscillations to well below the required accuracy.

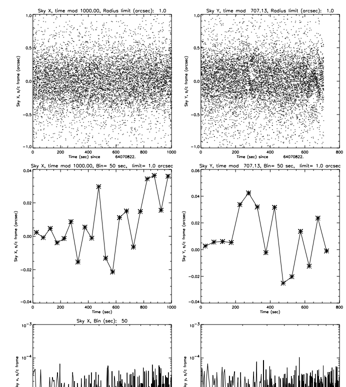

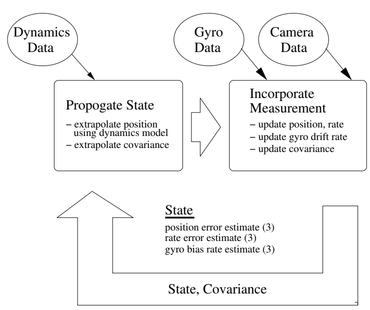

Future Work

Future extensions of the Kalman filtering algorithm for Chandra aspect determination may include the ability to continue processing in the event of limited hardware failure. In particular, a dynamics model based Kalman filter is being examined in an effort to continue to produce accurate aspect solutions with gyro data degradation. Such a filter might be designed as shown at the left, where gyro data is considered to be a measurement input, just as star camera data is currently used. The propagation equations would not be based on integration of gyro rate data but rather on a dynamics model of the Chandra spacecraft. New information, such as reaction wheel torque, would be used in conjunction with the equations of motion of the spacecraft to predict spacecraft attitude. The dynamics model would provide the accurate short term data to substitute for degraded gyro information.