Deep Impact Attitude Estimator Design And Flight Performance

Michael P. Hughes
NASA Jet Propulsion Laboratory, Pasadena, CA

Charles N. Schira
Ball Aerospace Corporation, Boulder, CO

Abstract

In July, 2005 Deep Impact successfully impacted comet Tempel-1 and imaged the crater evolution. To achieve this, an onboard autonomous navigation system had to estimate the trajectory and maneuver the Impactor spacecraft into the comet's path. A similar system, operating on the Flyby spacecraft, calculated the impact location for imaging. At the heart of each autonomous navigation system is the attitude and gyro bias estimator using star trackers and gyros. The encounter geometry and timing placed stringent demands on attitude bias magnitude and drift calling for sequenced tuning of the respective filters. This paper describes the attitude and gyro bias estimator algorithms, the estimator flight performance and particular challenges overcome by the team leading up to a spectacular encounter.

Introduction

Mission Overview

Deep Impact's basic objective was to crash an Impactor spacecraft into comet Tempel-1 and observe, via a Flyby spacecraft, the formation of an impact crater and exposed sub-surface material. In this way scientists may better understand the comet's material composition and structural properties.

The Impactor spacecraft, shown in Figure 1, carried an Impactor Targeting Sensor (ITS) telescope. The Flyby, shown in Figure 2, had two instruments: a Medium Resolution Imager (MRI) and a High Resolution Imager (HRI). The MRI was essentially identical to the ITS; the HRI contained an infrared spectrometer in addition to the visible detector. Instrument parameters are shown in Table 1.

Table 1. Deep Impact Key Instrument Parameters

	ITS	**MRI**	**HRI vis**	**HRI IR**
Camera resolution	10 ur	10 ur	4 ur	10 ur
Field of view	10 mr x 10 mr	10 mr x 10 mr	2 mr x 2 mr	2.5 mr x 10 ur
Typical exposure times	10 ms – 100 ms	10 ms – 100 ms	10 ms – 100 ms	10 ms – 100 ms

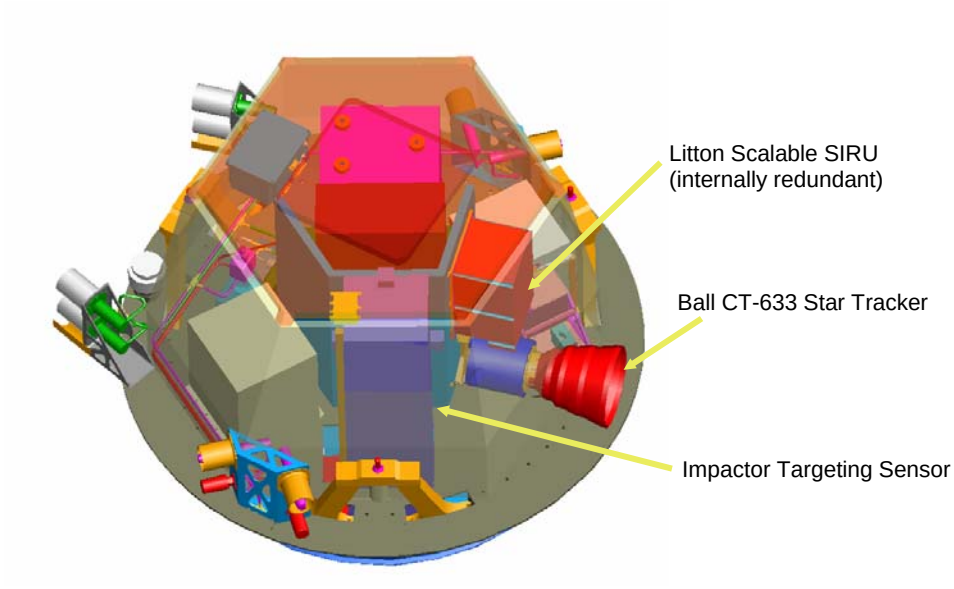

Figure 1. Deep Impact Impactor Spacecraft configuration

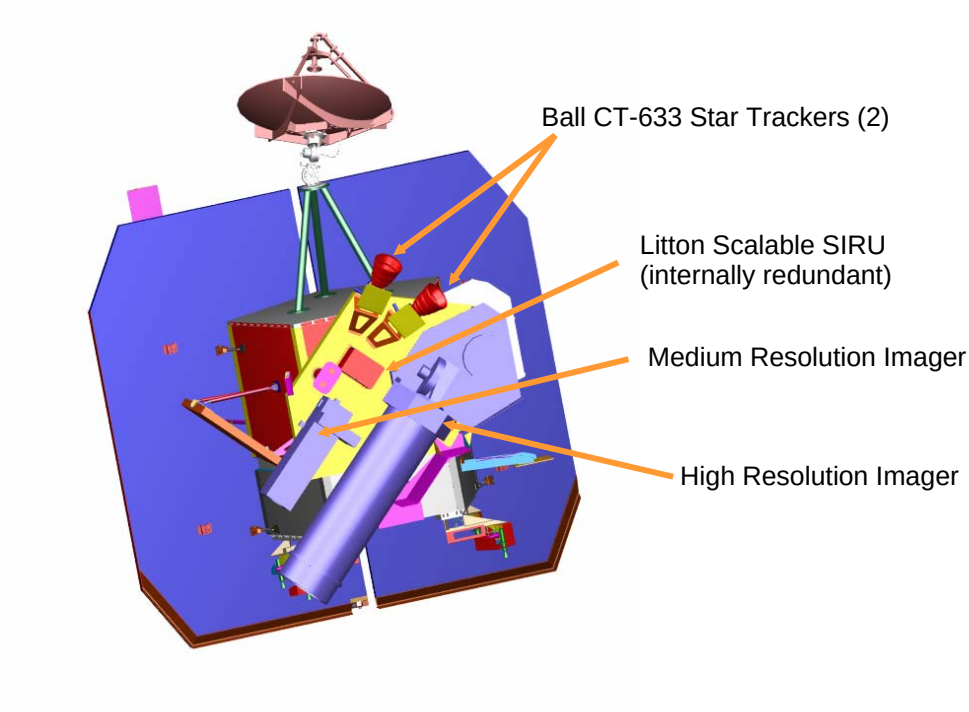

Figure 2. Deep Impact Flyby Spacecraft Configuration

Encounter Overview

The two spacecraft cruised to the comet as a Flight System, mated together and controlled by the Flyby spacecraft. A series of Trajectory Correction Maneuvers (TCMs) during cruise nominally put the flight system on a collision course with Tempel-1. Figure 3 shows a graphical timeline of the final 24 hours. A day away from comet encounter, the two spacecraft separated and the

Flyby executed a trajectory deflection burn. The Impactor, now a fully autonomous spacecraft, performed attitude control to maintain ITS imaging of Tempel-1. During the final two hours, on-board navigation algorithms processed ITS images and performed orbit determination of the Impactor's orbit. Three Impactor Targeting Maneuvers (ITMs) were performed as the Impactor autonomously self-corrected it's trajectory insure a collision in a well-lit area. Meanwhile the Flyby deflection burn put it on a 500 km closest-approach trajectory and slowed it's velocity to provide approximately 15 minutes of imaging between impact and entry into a protective shield mode.

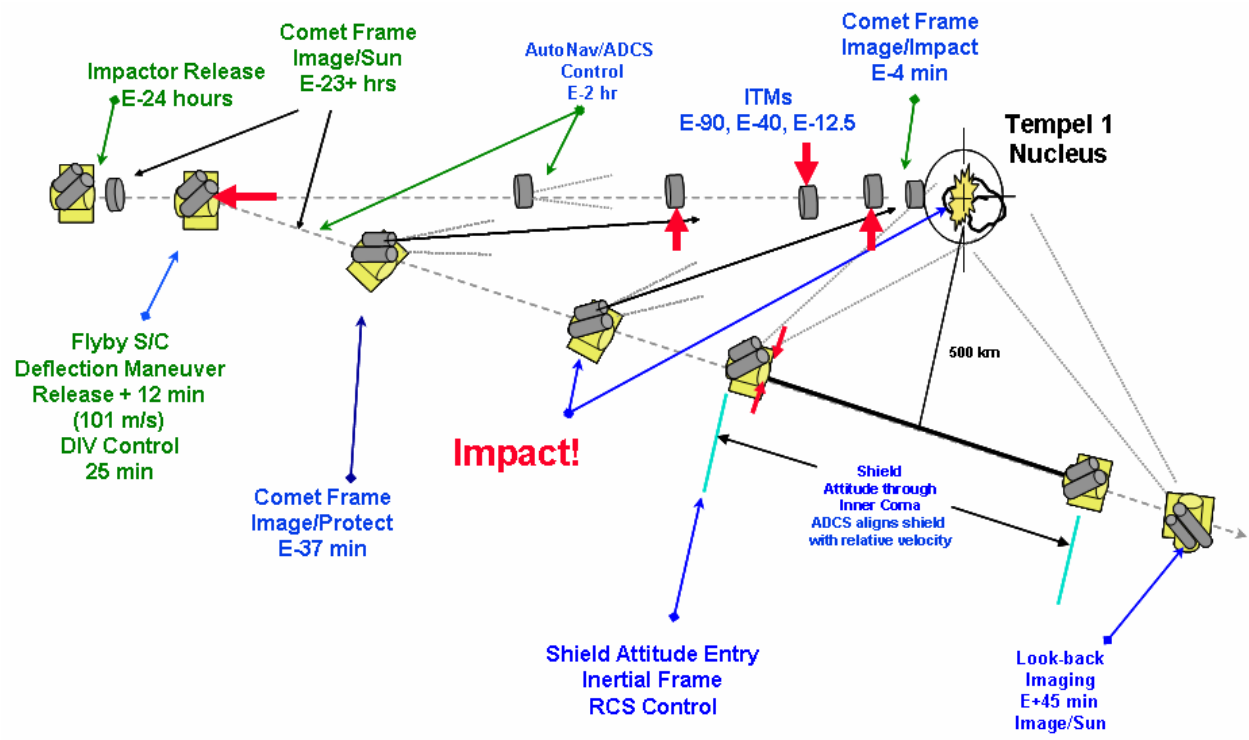

Figure 3. Encounter Overview

<u>AutoNav Overview</u>

Both spacecraft used autonomous optical navigation algorithms (AutoNav) to determine their respective orbits. The OD solutions were used to determine the desired attitude profile for imaging instrument pointing. On the Impactor, in addition to pointing, the AutoNav solutions were used to perform autonomous divert maneuvers, or Impactor Targeting Maneuvers (ITMs) [1].

Orbit determination by optical navigation (opnav) depends on availability of simultaneous observations of a near-field body for which *a priori* orbit knowledge exists and background stars that provide knowledge of inertial motion of the instrument providing the near-field body observations. When optical navigation observations are available in a single sensor, the method is known as "star-relative." For DI, the orbit geometry, relative brightness of the comet and

stars, and small telescope fields of view drove the use of "starless" opnav in which comet observations and attitude information come from separate sources, the latter provided by the Attitude Determination and Control System (ADCS) attitude estimator.

The OpNav system has heritage to many JPL deep space programs and the on-board real time implementation, known as AutoNav, has heritage to Deep Space-1.

AutoNav and ADCS Coupling

Using starless OpNav drove some mission-unique requirements on the ADCS-provided attitude knowledge. Although absolute pointing was still critical (the comet needs to be in the telescope field of view!), attitude knowledge error became especially important, and the most critical of all was the rate of change of attitude knowledge error. This parameter came to be known as attitude knowledge error drift (AKED).

AKED became important because, for any reasonably small attitude knowledge error, a constant attitude error will not adversely affect the AutoNav OD accuracy (AutoNav will still give ADCS the correct inertial pointing direction to maintain comet imaging); however, a changing attitude knowledge error will lead directly to an error in the velocity of the estimated spacecraft orbit. This velocity will not only cause errors in the desired attitude direction, but, more importantly, will introduce erroneous velocity change requests in the ITM process [2].

A simple geometric analysis reveals the coupling. Consider the Impactor spacecraft on a "miss" trajectory as shown in figure 4. (This analysis works for the Flyby as well, but a miss is, of course, desired for the Flyby.)

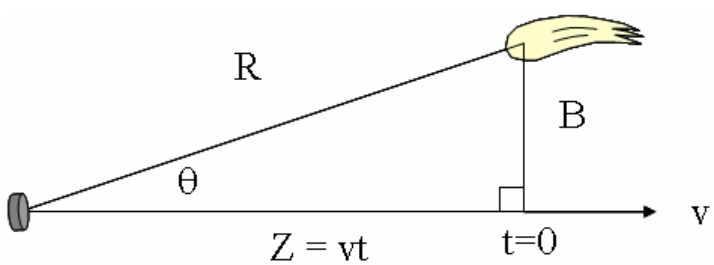

Figure 4. Encounter Relative Geometry

B = magnitude of B plane intercept vector
v = encounter relative velocity
t = seconds to B plane intercept

The angle between the comet and the B-plane intercept (aka closest approach point) is θ and is calculated by

$$\tan\theta = \frac{B}{vt} \qquad (1)$$

As the spacecraft approaches the comet, the true angle θ grows as

$$\dot{\theta} = \frac{B/v}{t^2 + (B/v)^2} \tag{2}$$

To drive B to zero for impact requires a DV of

$$DV = \frac{B}{t} \tag{3}$$

Now AutoNav calculates B and DV based upon multiple images that reveal the apparent $\dot{\theta}$. But what AutoNav sees is the true $\dot{\theta}$ corrupted by attitude knowledge error drift so it's calculations of B and DV will be in error accordingly.

$$\dot{\theta}_{ANAV} = \frac{B/v}{t^2 + (B/v)^2} + \dot{\theta}_{AKED} \tag{4}$$

For a true impact trajectory (B=0), any DV from AutoNav will be based entirely on attitude knowledge error drift and result in a miss. Conversely, a true miss can be mistaken for an collision course if the AKED approximately cancels out the true geometric rate.

Therefore, the AKED requirements imposed by starless AutoNav were very tight. Key ADCS requirements are shown in Table 2.

Table 2. Key Deep Impact ADCS Requirements

Parameter	**Impactor Rqmt**	**Flyby Rqmt**
Absolute pointing	< 320 ur	< 61 ur
Attitude knowledge	< 100 ur	< 61 ur
Attitude knowledge error drift	< 35 ur/hr	< 45 ur/hr

Attitude Determination & Control System Overview

The DI ADCS has a high degree of heritage to several spacecraft built by Ball Aerospace, including GFO, QuikScat, MTI, QuickBirds-1 and 2, and ICESat [3]. Figure 5 shows a functional block diagram of the ADCS.

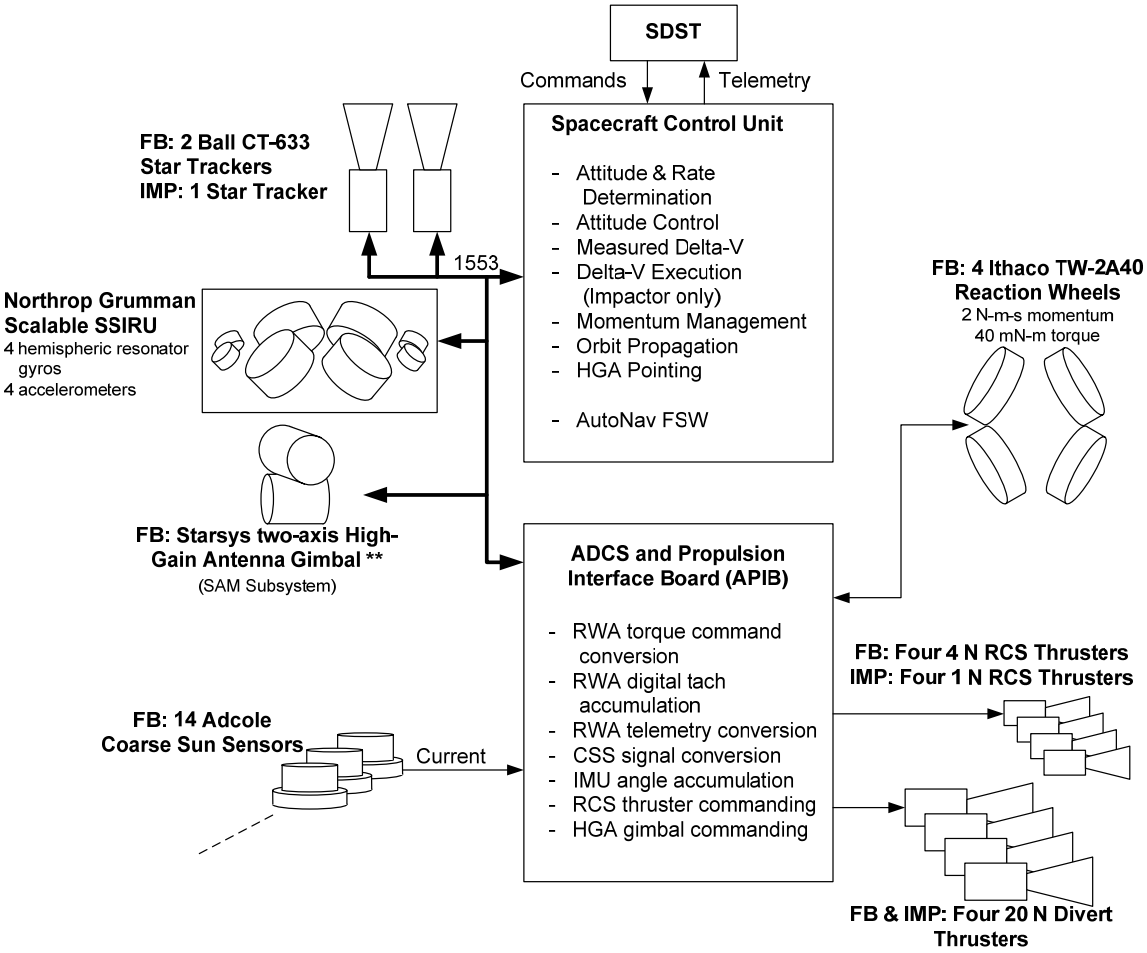

Figure 5. DI Attitude Determination and Control System Functional Block Diagram

The attitude determination sensor suite included: Ball CT-633 quaternion output star trackers (two on the Flyby, one on the Impactor); a Northrop Grumman Scalable Space Inertial Reference Unit (SSIRU) on each spacecraft. Key performance parameters for the trackers and gyros are presented in tables 3 and 4 below.

Table 3. Key Star Tracker Characteristics

Parameter	**Value**	**Units**
Model	Ball CT-633	n/a
NEA	< 73	μr, 3σ, per star
Spatial error	< 87	μr, 3σ, per star
Attitude knowledge (cross-boresight)	< 87.3	μr, 3σ

Table 4. Key Inertial Measurement Unit Characteristics

Parameter	Value	Units
Model	Northrop Grumman Scalable SIRU	n/a

Angle Random Walk	0.0001	deg/rt(hr)
Angle White Noise	0.003	arc-sec/rt-Hz
Bias stability	0.0016	deg/hr, 3σ
Quantization	0.05	arc-sec/LSB

For acquisition and safe mode use, 12 coarse sun sensors were mounted on the Flyby. Actuation was achieved using: four Ithaco TW2A40 reaction wheels on the Flyby; four RCS thrusters (4.5 N on the Flyby, 1 N on the Impactor); and four 20 N divert thrusters on each spacecraft. Flight software on both spacecraft executed on a Rad-750 based Spacecraft Control Unit (SCU).

Attitude & Gyro Bias Estimator Design

A block diagram of the DI attitude estimator algorithm is shown in Figure 6. The algorithm provides attitude and rate estimates based on optimally-mixed star tracker and gyro data. In addition, star tracker attitude measurements are used in a second-order low-pass filter to estimate gyro biases. The gyro bias filter is shown in Figure 7.

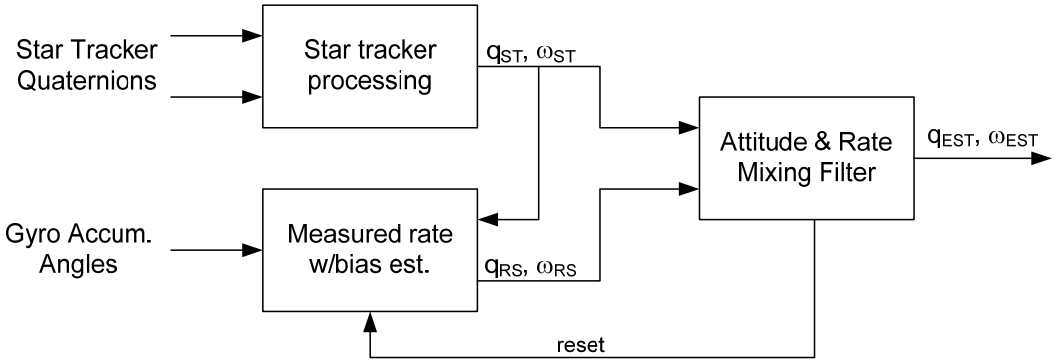

Figure 6. Attitude estimator block diagram

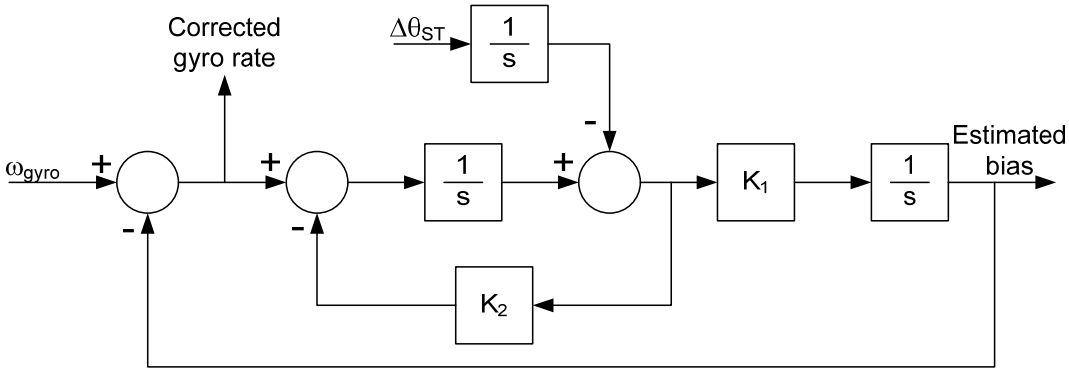

Figure 7. Gyro bias estimator block diagram

The attitude and gyro bias filter parameters were optimized for DI's mission requirements, driven primarily by the tight requirement on attitude knowledge error drift, and based on the relative performance of the star trackers and the IMU.

While many missions would optimize attitude estimation filters to minimize attitude knowledge errors, DI's AKED requirement drove the design to a mix of star tracker and gyro data that sacrificed attitude knowledge in order to reduce the effect of star tracker star field changes on attitude estimate transients.

The gyro bias estimator treats star tracker measurements as a truth reference, and any deviation between star tracker and gyro measurements is blamed on gyro bias.

Another somewhat unique aspect of the DI mission was that both the Flyby and Impactor spacecraft attitudes were almost inertially fixed for many hours leading up to Impactor release and comet encounter. This meant that stars would be fixed in constant locations on the star tracker CCDs, and any spatially correlated star tracker errors would appear as an attitude measurement biases to the attitude estimator.

The AKED requirement applied when AutoNav OD was activated – approximately the final two hours before impact on each spacecraft. Command sequences were designed to disable star tracker updates to the attitude and gyro bias estimation filters just prior to this point, and from this point forward, attitude estimates would be based solely on integrated gyro outputs. When the filters were frozen, attitude knowledge would drift from that point forward at the rate of gyro bias estimate error; thus it became critical that both filters be settled and in steady-state at the time they were disabled.

Attitude & Gyro Bias Estimator Flight Performance

Given the tight requirements on the attitude knowledge error drift, the attitude & gyro bias estimators' performance were scrutinized throughout the mission. This was done by plotting star tracker residuals and gyro bias estimates daily and for long term trends. Star tracker residuals were calculated on the ground by performing a quaternion difference between the estimated quaternion and the tracker output quaternion in the body frame. To eliminate latency errors, the tracker quaternion had to first be propagated to the estimated quaternion time tag. In the academic literature on estimation theory [4], this quantity is sometimes called "innovations process", "prediction error" or "measurement residual".

$$\mathbf{e}_{resid} = \mathbf{y} - C\hat{\mathbf{x}} \qquad (5)$$

Of course, $\mathbf{y}$ is the tracker quaternion, $\hat{\mathbf{x}}$ is the estimated attitude quaternion, C is the transformation from body to tracker coordinates. The minus sign is really a quaternion inverse multiplication. By examining star tracker residuals, transients induced by the star trackers and longer term drifts induced by the gyros or the estimation process could be separated and revealed.

Plots of flight activities were reviewed in team meetings attended by both AutoNAV and ADCS experts. Among most interesting activities were a few flight Optical Navigation tests. These provided MRI derived quaternions that could be used as an "external view" of the attitude knowledge drift.

Early Mission Flyby focus

Early on, virtually all of the flight activities were Flyby spacecraft centered. Therefore, the Impactor gyros and tracker were unpowered for much of early mission except for special instrument and calibration tests. This was done to reduce operations complexity and workload on the team. As the both spacecraft used the same gyros and tracker, it was believed that characterizing the Flyby would go a long way toward predicting systematic performance of the Impactor.

Much of the early mission was crowded with slews to have the Flyby instruments view various targets for calibration including the Earth and Moon, Jupiter, Saturn as well as certain star fields and stray light test attitudes. During many of these slews, the ADCS team noticed errors in the star tracker residuals and gyro bias estimate plots. At the time it seemed apparent that virtually all of the errors were a result of imperfect alignment tables and expected tracker lost lock conditions. As the flight alignment calibration activity had been stretched out and delayed, it was expected that gyro alignment errors during large slews would perturb the attitude estimate. As a workaround, the flight team agreed to perform slews using a heavier weighting on the star tracker output in the mix of star tracker and gyro data to reduce the estimator's divergence caused by the imperfect pre-launch alignment values. This solution seemed to work well.

Figure 8 shows gyro bias estimate trends have a stable mean with +/-100 urad/hr transients occurring over hours to days. The many spikes in the star tracker residuals plots of figure 9 were typical of what is expected when the tracker output is disabled during the many slew maneuvers. The top two plots show fairly good tracker cross boresight performance. The bottom plot shows the about boresight error is on the order of +/-500 urad.

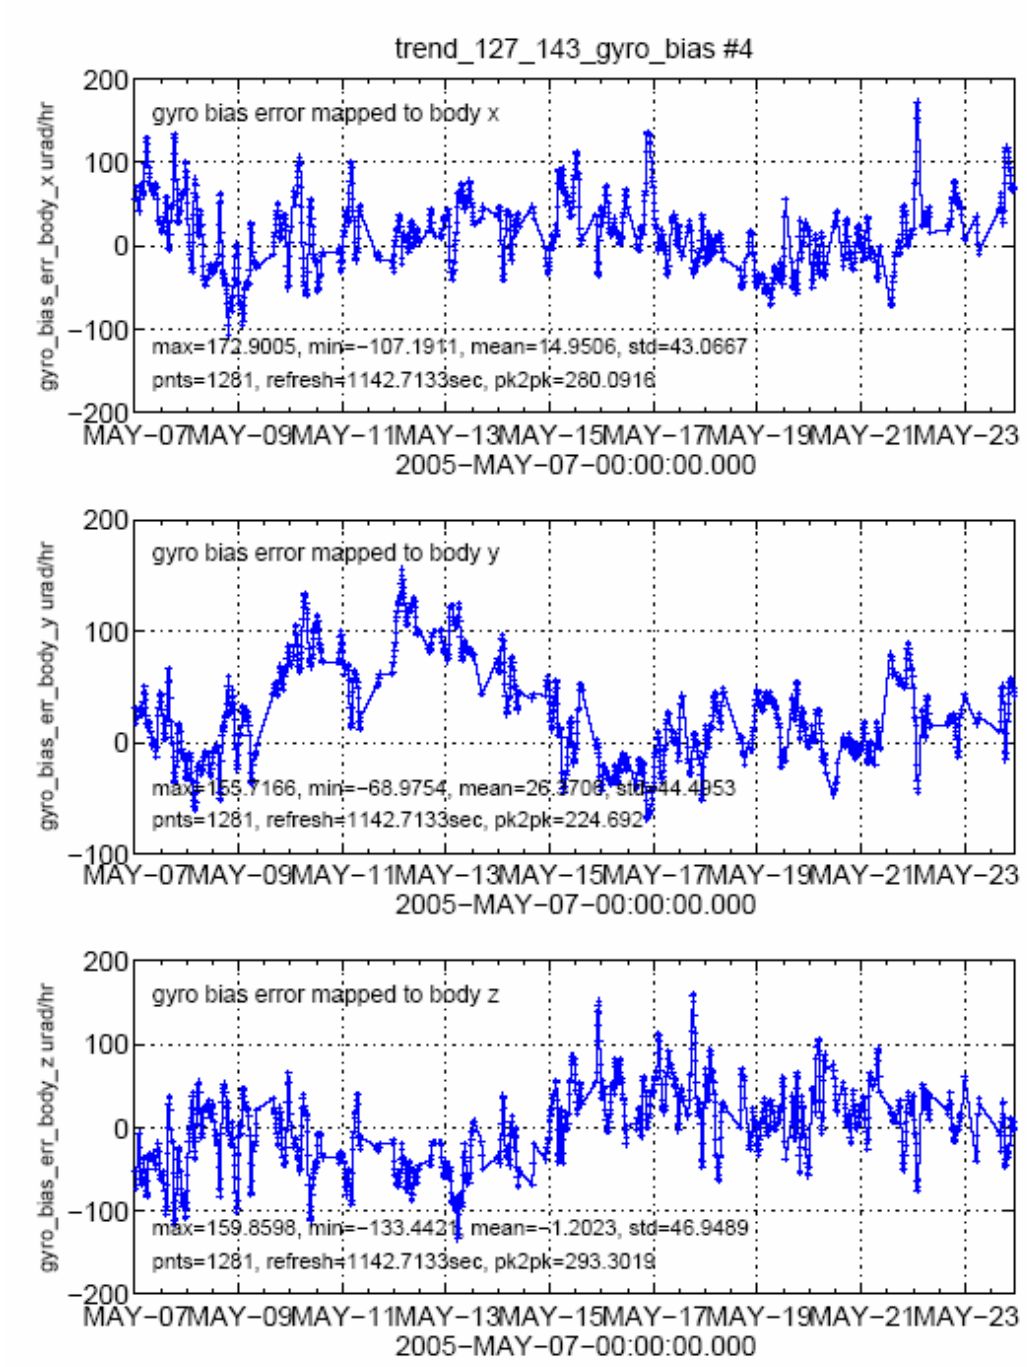

Figure 8. Flyby Gyro Bias Estimates Mapped to Spacecraft Body frame

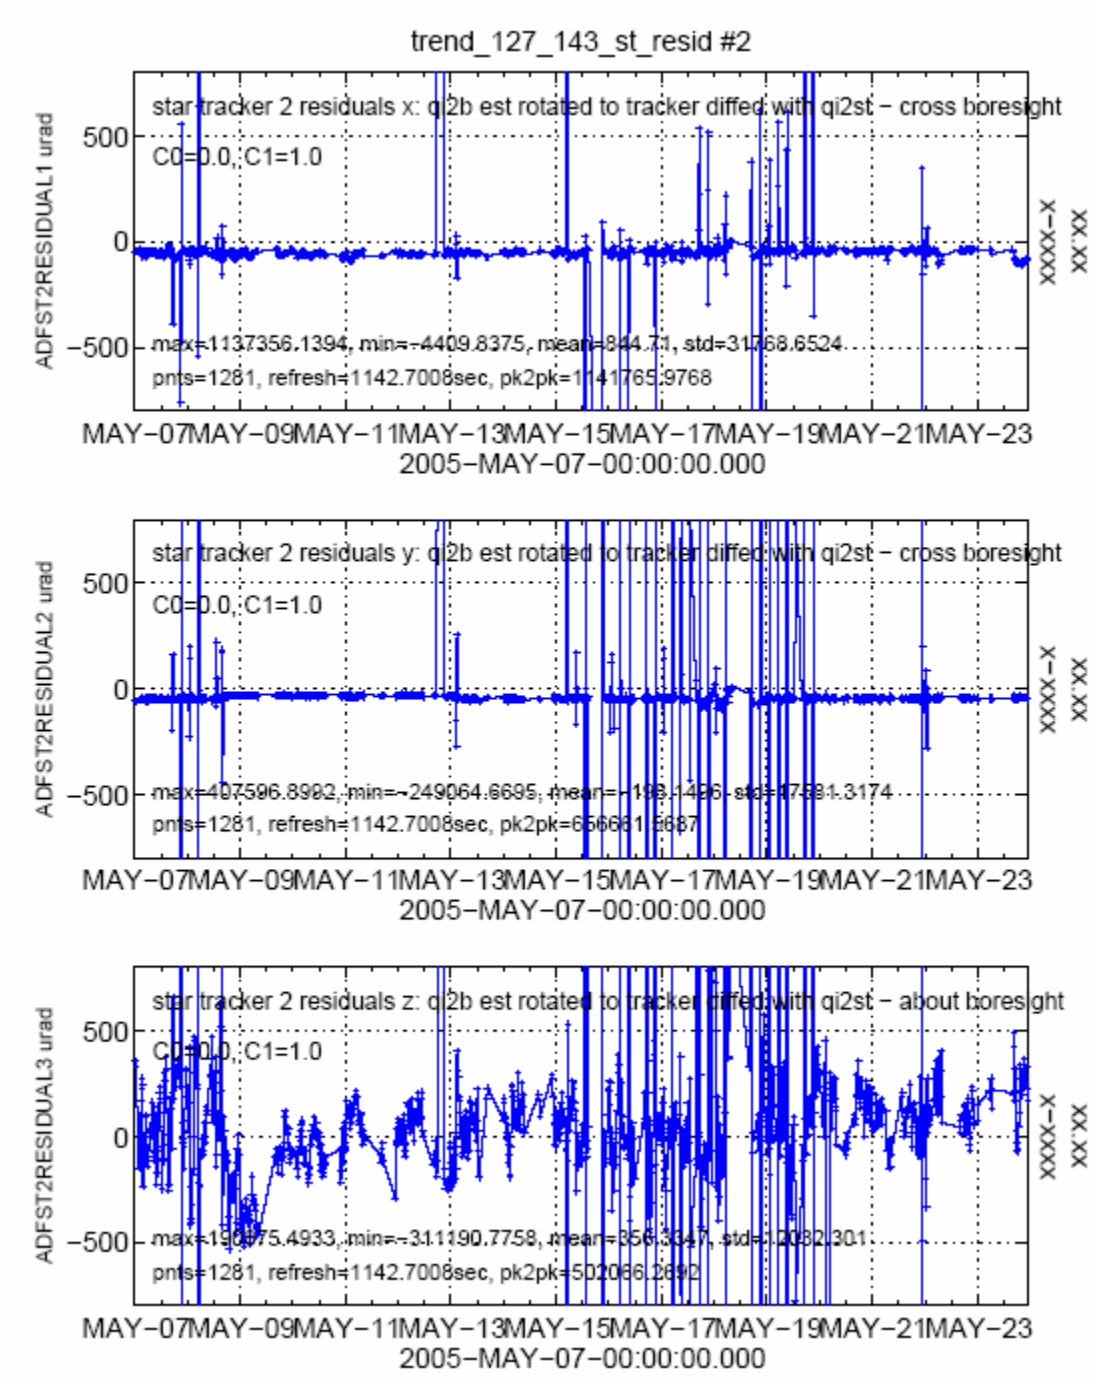

Figure 9. Flyby ST2 residuals Mapped to Tracker frame

<u>Flyby True Gyro Bias Revealed</u>

On April 23, 2005 a Flyby Optical Navigation test was performed. This test involved taking many images from the MRI telescope at a nearly inertial attitude and processing them on the ground. The ground process produced an inertial quaternion for each image, and thus can be thought of as a high performance star tracker. The MRI quaternions could then be used to create

MRI residuals using the same algorithm already in routine use for calculating the tracker residuals. During the planning of this activity, it was decided to disable the tracker output to both estimators as this was how we planned to execute the last 2 ½ hours of encounter. In this way we could see how the gyros are really behaving when uncoupled from the tracker noises. Two features were noticed.

First, the true gyro biases seemed to vary between 20 and 50 urad cross MRI boresight during the 1:40 minute test. (About boresight measurements are corrupted by MRI geometry in a similar way to trackers.) For the Flyby, this was good news in that we had lots of confidence from tests and analysis that this would support encounter. However, the peak slope was about 75 urad/hr – exceeding the 45 urad/hr specification. This caused some concern and did not recur in an identical test. This is shown in figure 10.

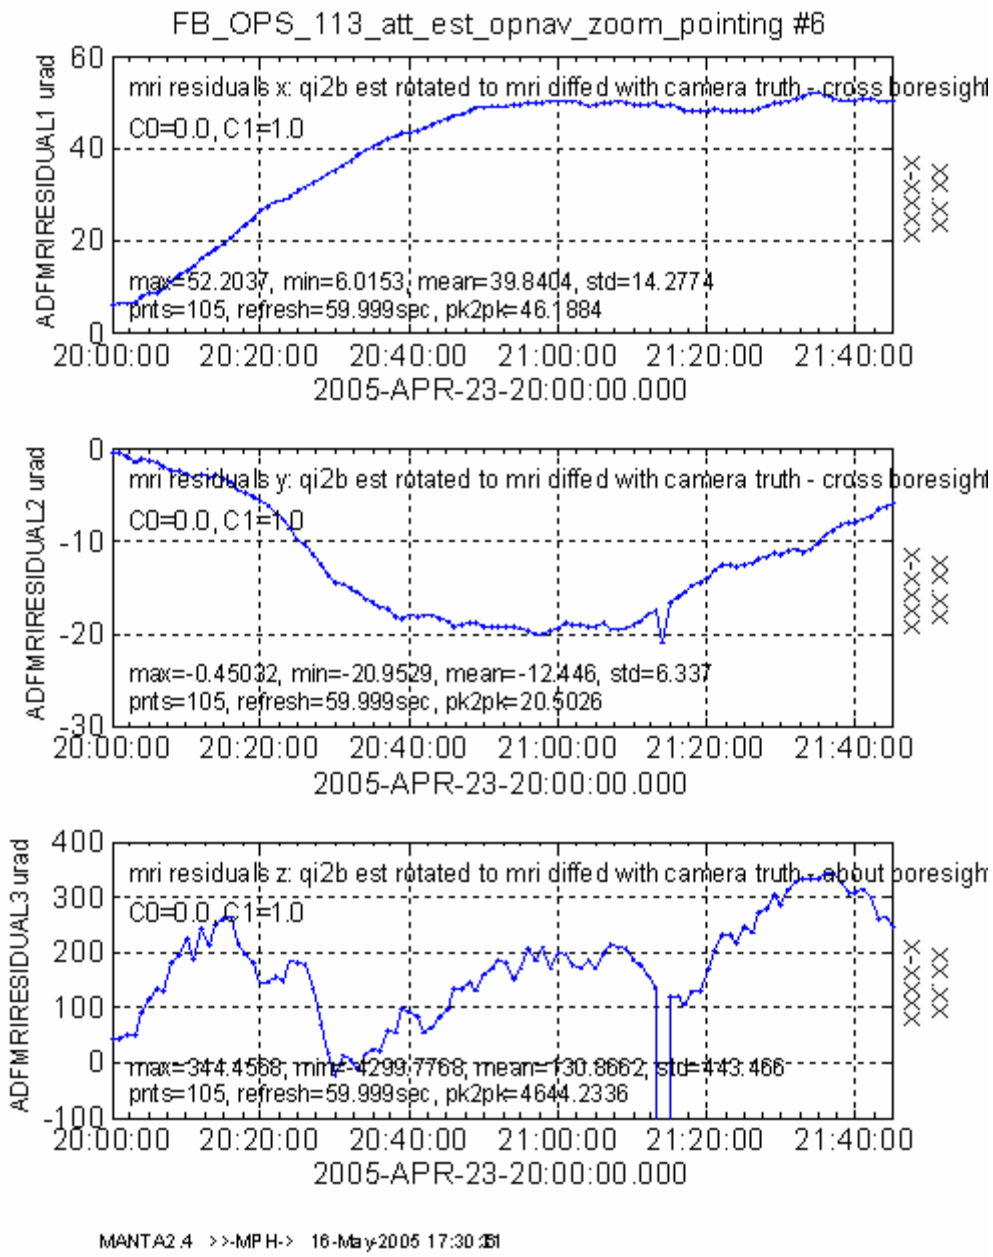

Figure 10. MRI Residuals from OpNav Test Revealing True Gyro Bias Performance

The next feature, shown in figure 11, was that there appeared to be significant step changes in tracker output. At the time, the cause of the steps was unknown but later, after intensive investigation, it was learned that step changes occur when the tracker had changed its star set. These steps affected the tracker about boresight axis the most. This seemed to happen somewhat randomly and was not expected. Why the star tracker changed star sets in this way is still not fully understood, although it was conjectured to be related to the space weather or possibly gamma rays. Regardless, it was not a great concern for the Flyby as there are nominally two

trackers combined to produce the measurement. In a sense, the two tracker solution "forgives the sins" of either tracker.

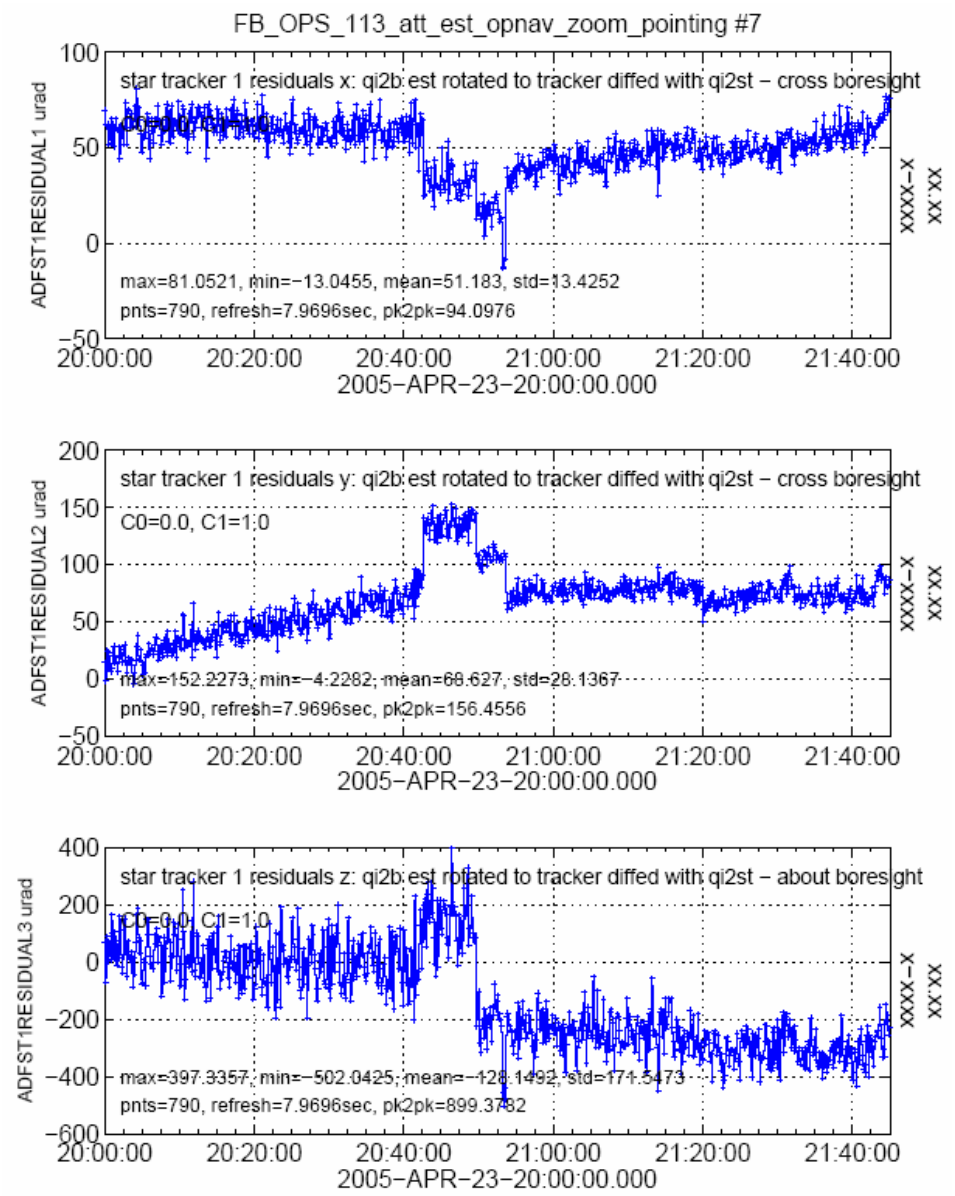

Figure 11. Tracker 2 Residuals from OpNav Test Showing Steps due to Star Set Changes

The two features, gyro drift and star tracker steps, were two separate phenomena. We know this because the filters were disabled during the April 23 test and thus there was no data coupling between the tracker and gyros. It seems coincidental that they were both revealed in the same 1 hour 40 minute period.

<u>Shift to Impactor Focus</u>

Although the step changes in the above figure 11 were not a concern for the Flyby, they were a big concern for the Impactor for three reasons. First, the Impactor is inherently more susceptible to tracker errors due to its single tracker solution. Second, the impact geometry drove tighter performance for attitude knowledge error drift. Third the Impactor tracker's about boresight axis was the cross boresight axis for the ITS instrument—the same axis AutoNav depended on.

Unfortunately, during the early mission the Impactor was commanded for only a few short checkout and calibration activities. The rest of the time it remained powered in a "sleep state". Planned operations called for waking the Impactor two weeks before encounter. Therefore, there were precious few opportunities to gather extended periods of star tracker, gyro and attitude estimator performance telemetry for performance evaluation. During early May, 2005, the big question was "is the Impactor performance good enough for encounter?"

About the same time, a suite of tests was being executed on the testbed to stress the encounter system with various parameter errors to examine robustness. These robustness tests revealed a very disturbing effect when the attitude knowledge error drift was perturbed by a large amount. The effect was called the ITM zig-zag. What would happen is AutoNav would see the large drift, mistake it for a large miss distance and command a large Impactor Targeting Maneuver to correct it. This happened even when the Impactor was on a simulated collision course, ie. it would command the Impactor to miss the comet! The two later ITM burns would be similarly affected. In some tests, the Impactor propellant tank was emptied as a result of the ITM zig-zag.

Some ITM zig-zag was basically expected behavior when the signal to noise ratio was bad, i.e., when the true targeting error was small and the spacecraft was far from the comet as in the case of ITM-1. (Robustness tests started with a perfect trajectory, so any non-zero ITM-1 would make the trajectory worse, and any noise in the system would cause a non-zero ITM-1). If we had it to do over again, we probably only would allow ITM-1 to be performed if there was an indication of a bad release from the Flyby.

Working the Solution

Following these revelations, a tiger team was formed to work the solution. This team included attitude and tracker experts from JPL and Ball Aerospace. The team decided to power on the Impactor ADCS early for a set of evaluation and tuning operations. After evaluating the star tracker hardware telemetry of all three units, the team declared them all healthy. The tracker software was looked at closely and it was decided to load a patch to command break track on stars in the corners, the regions susceptible to the worst calibration errors.

One month before encounter, the daily plots revealed series of very large tracker spikes accompanied by large gyro bias estimate errors. One of these events is shown in figures 12 through 14 below. These plots show a near 4 mrad spike in the tracker output causing a change in gyro bias estimate (with a 7 hour time constant) of about 150 urad/hr. This occurred when three of the five stars being tracked were declared invalid by the tracker.

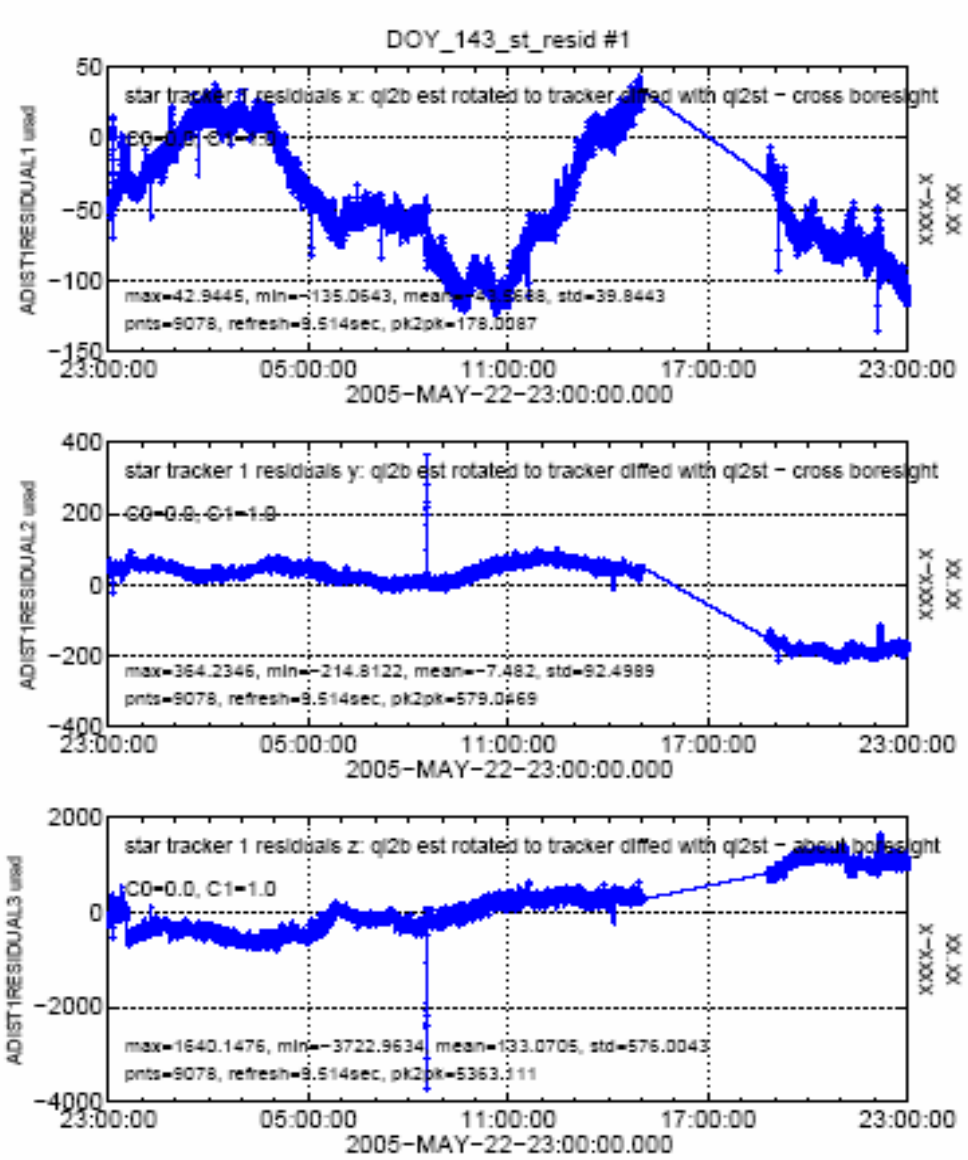

Figure 12. Impactor Star Tracker Residuals Day Of Year 143

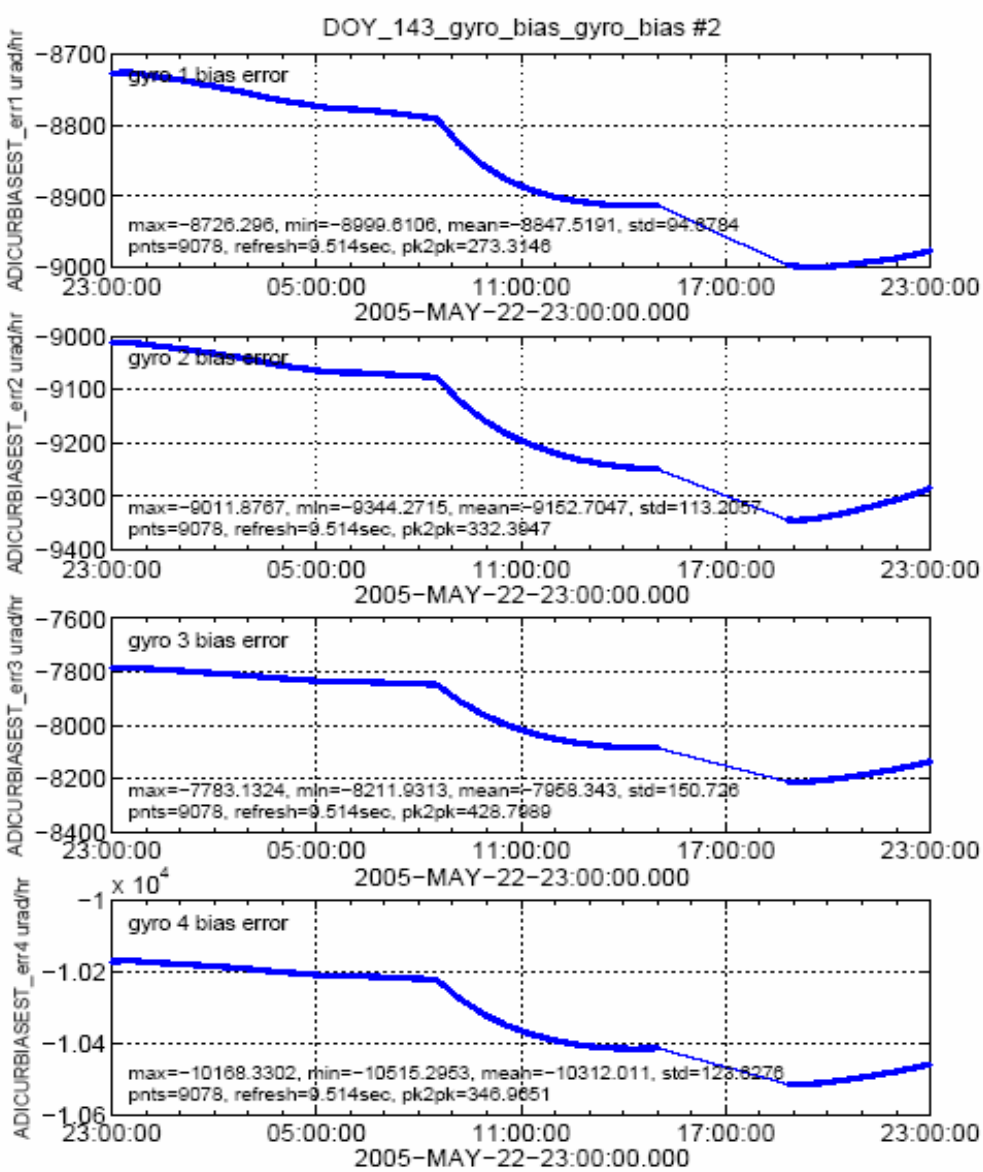

Figure 13. Gyro Bias Estimates, Day Of Year 143

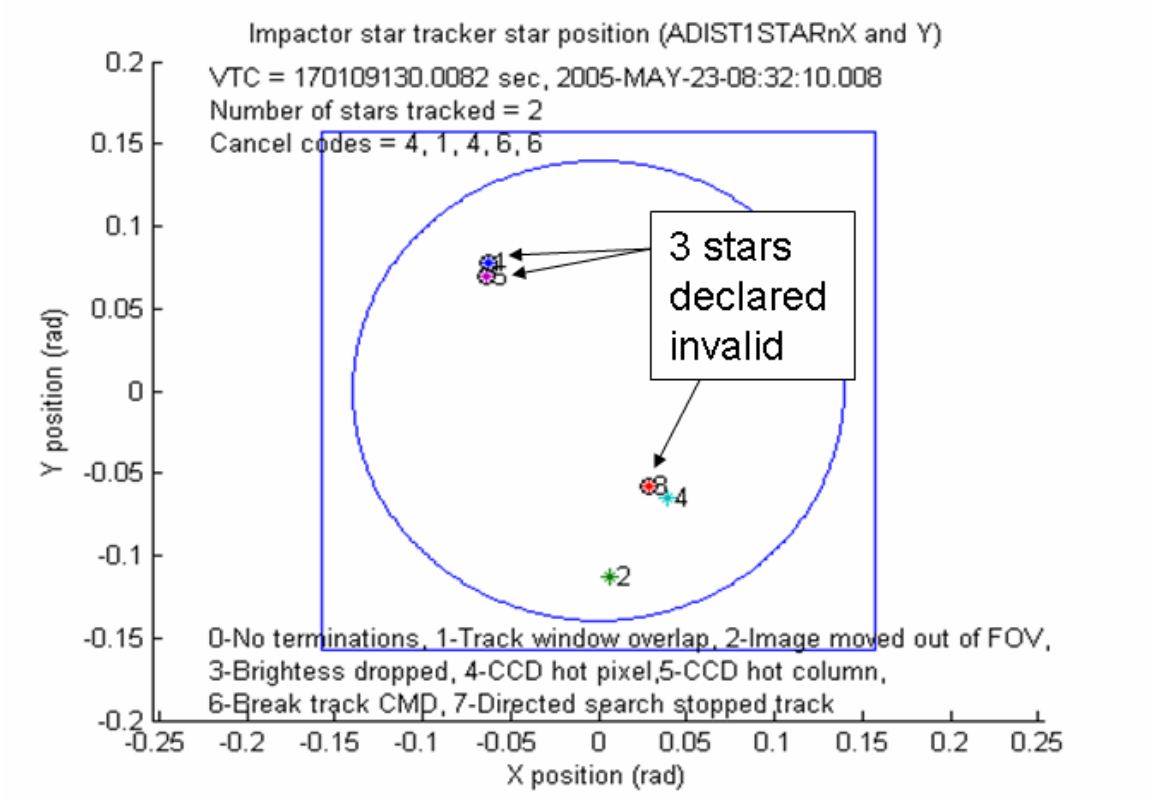

Figure 14. Star Tracker Field, Day Of Year 143, Stars 2 and 4 are used in solution

It was clear that if one of these transients occurred in the last few hours before impact that it could result in an unrecoverable ITM zig-zag.

Actions Leading up to Encounter

Further work on the star tracker software yielded a simple, yet ingenious patch. The patch dubbed "DI Combo", circularized the FOV to improve calibration, removed a few questionable near neighbor stars at the encounter attitude, and dropped any stars that were identified invalid. These efforts greatly improved performance during star set changes, but not by enough to eliminate the risk. What ADCS needed was 24 hours at the Impactor encounter attitude to see how the system performed with the tracker patch operating. With the Impactor still attached to the Flyby, this was not possible due to the heavy demands on the comet pointed Flyby spacecraft instruments for ongoing critical science and calibrations.

It was obvious to us that the gyro bias estimator time constant had to be changed, but it was not as obvious how much. Both the magnitude of tracker steps *and* the true underlying gyro bias drives time constant selection. A long time constant allowed the system to ignore tracker errors, but also the tracker inherent accuracy. After the shock of separation it was expected that the gyros might undergo a true change in bias, so this favored a short time constant to allow quick convergence to the true bias. But once the estimated bias converged, the system required smooth performance to prepare for AutoNav operations. It was thought that a sequenced approach would suffice, ie. a) separate, b) converge with short time constant, c) switch to long time

constant. But if the long time constant gain switch occurred right after a spike, the situation could result in "locking in" on an erroneous bias and ultimately loss of mission. But at the same time it was not prudent to rely on a commandable link from the ground to Flyby to Impactor after separation for any contingency operation, so options were limited. Ultimately the tiger team decided to separate with a 1 hour time constant and sequence a switch to a 4 hour time constant at separation plus 8 hours. In case of anomaly, a set of contingency commands were prepared to select several gains between 30 minutes and infinity (estimator disabled) time constants.

The 4 hour time constant was supported by analysis of the flight data and knowledge of the specified gyro noises. The time constant cannot be selected based on Kalman filter theory because the step errors from the star tracker cannot be modeled as white noises. However, because the gyro bias estimator is a linear system, the principle of superposition applies. Therefore we can superimpose the outputs of the noise driven system with the step driven system and pick the time constant corresponding to the minimum error. Figures 15 and 16 illustrate this as a block diagram and with the actual performance data.

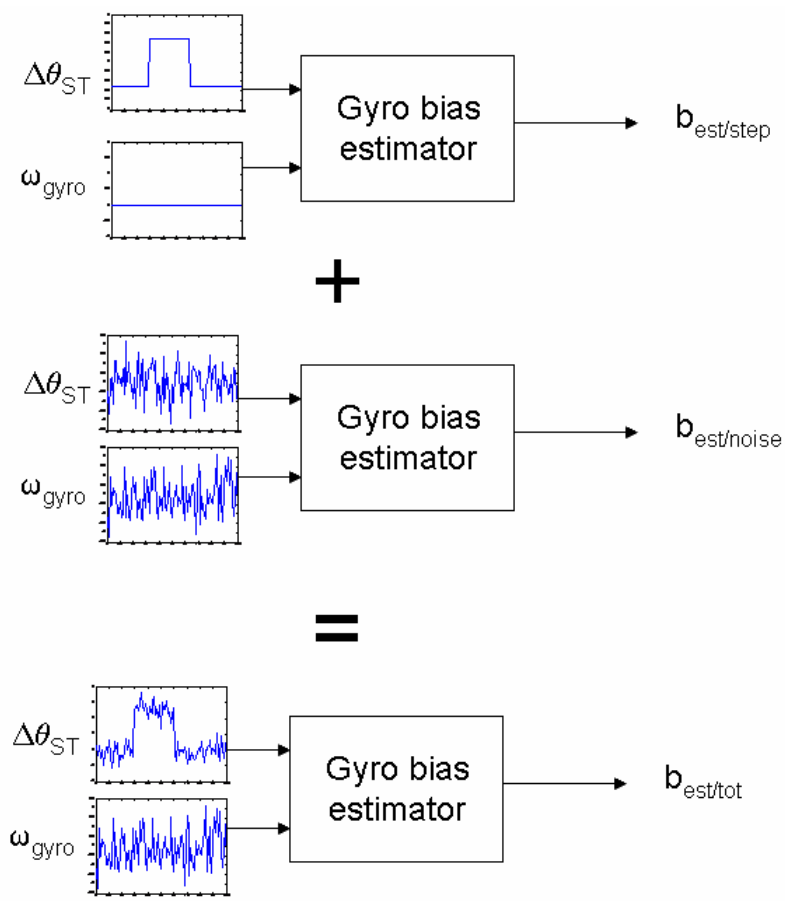

Figure 15. Block Diagram illustrating principle of superposition for gyro bias estimate error sources

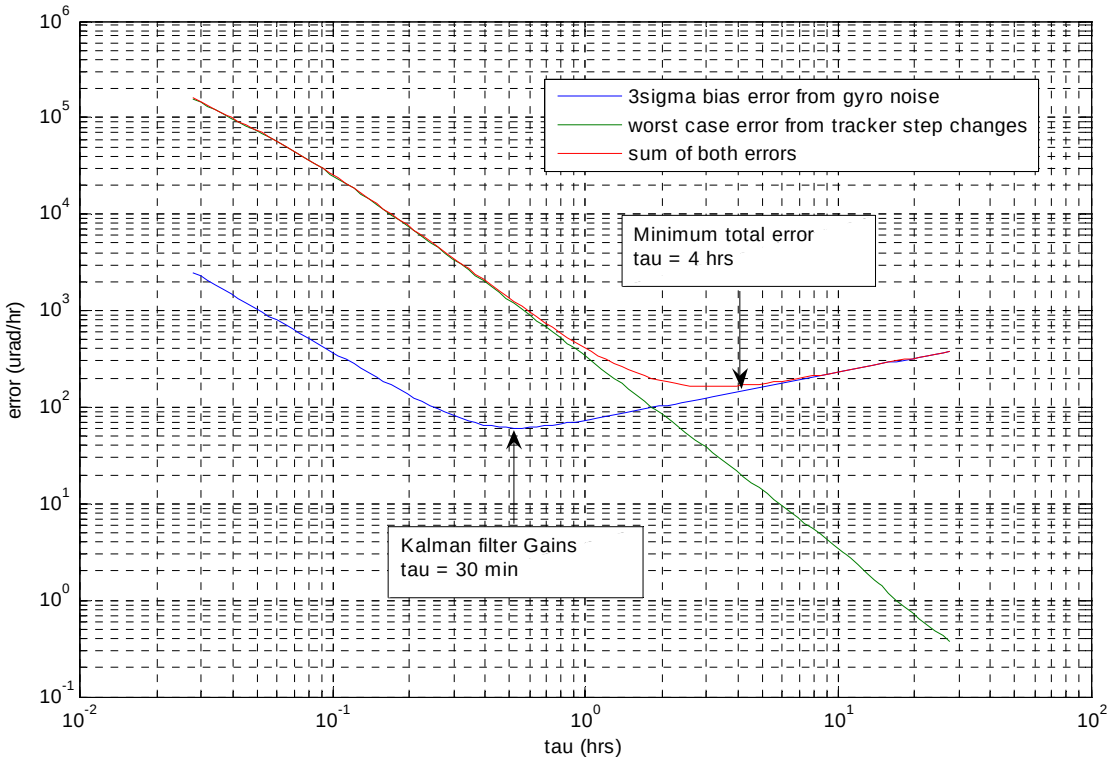

Figure 16. Analysis Supporting 4 hour time constant selection

Attitude & Gyro Bias Estimator Encounter Day Performance

<u>Separation Plus 15 hours</u>

Following a nominal separation on July 3, 2005 at 06:00:00, the Impactor bias estimator with a 1 hour time constant converged to the true bias. Our fears that the true bias would change dramatically were not founded. The bias estimates converged to values near those telemetered 24 hours before separation. Now the Impactor was free flying under thruster control. The deadband rates induced interpolation errors in our post processing method that appeared as 100 urad noise. This is shown in figures 17 and 18.

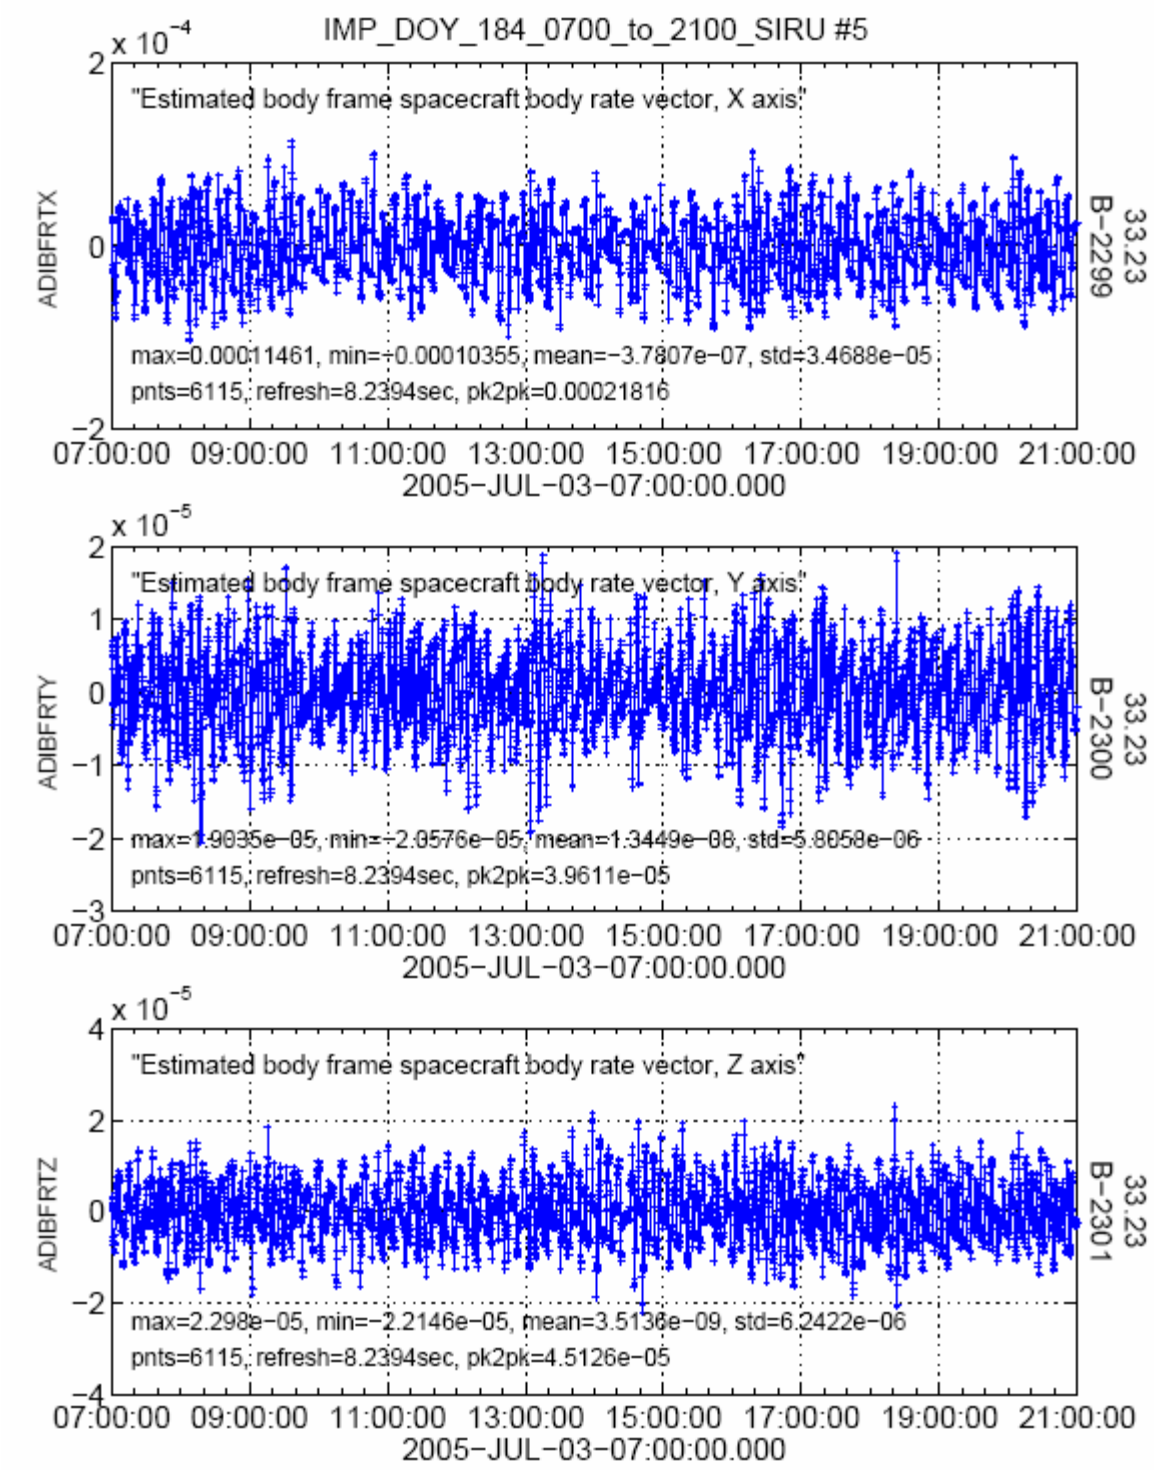

Figure 17. Post separation Impactor Body Rates showing thruster deadbanding

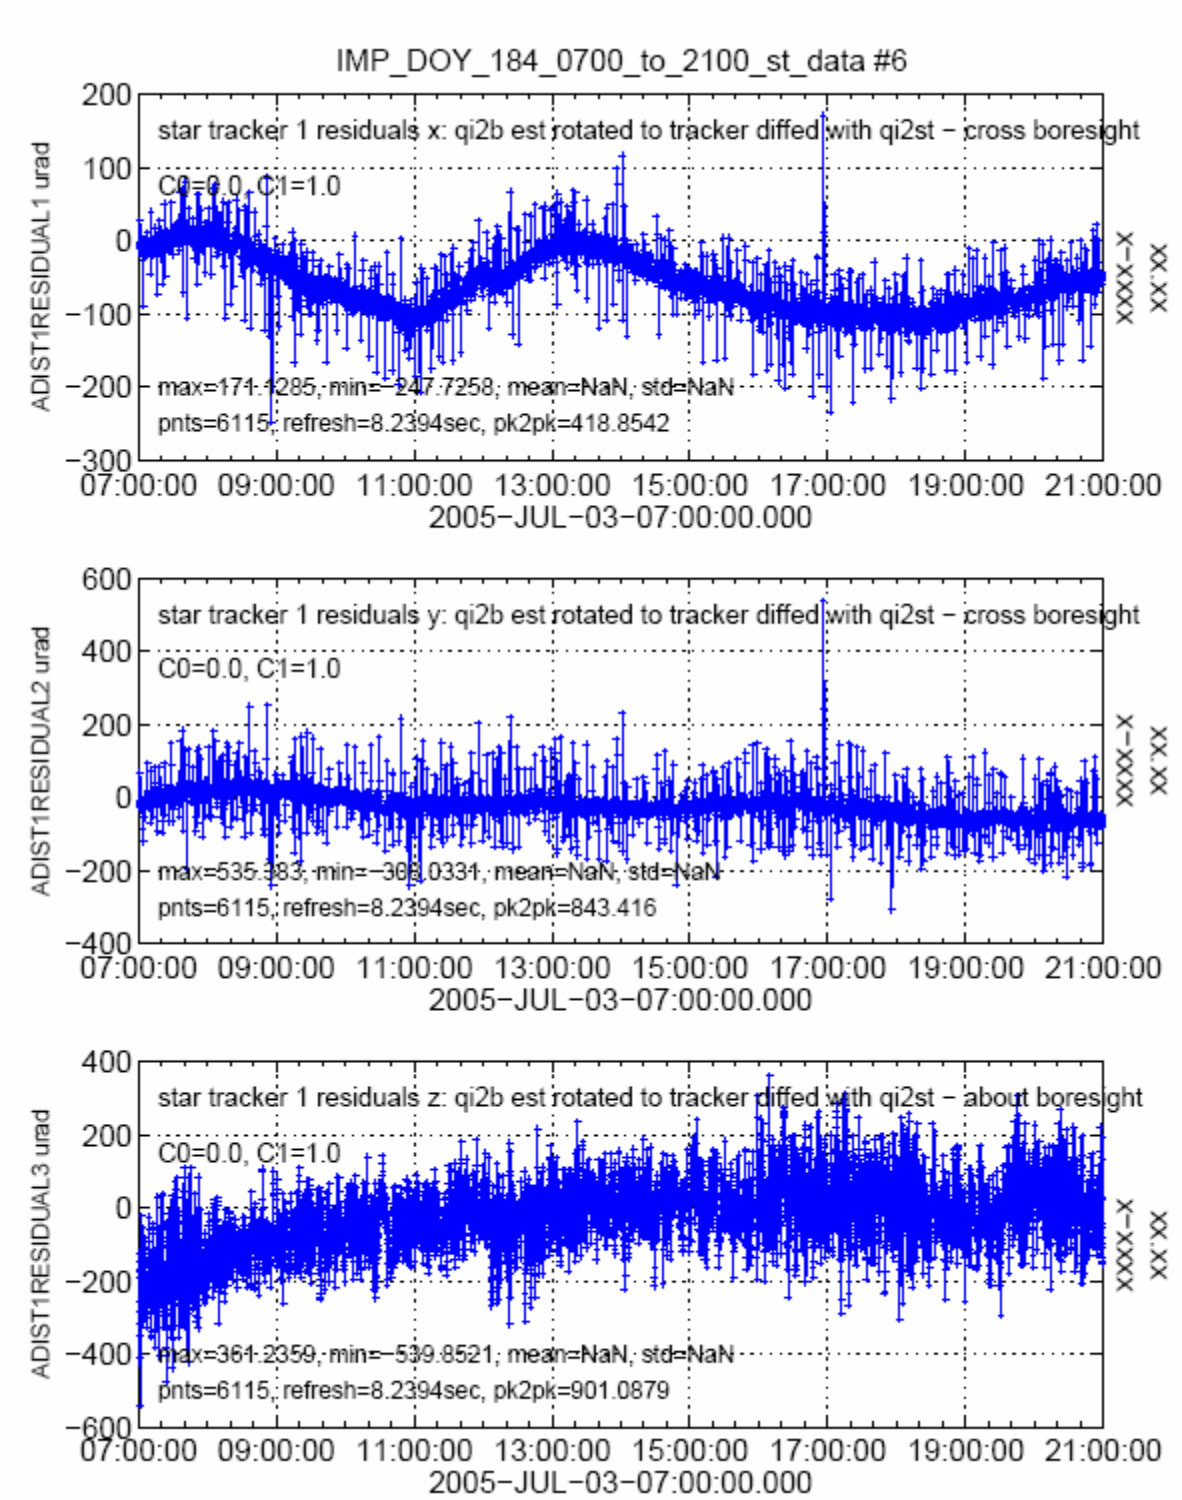

Figure 18. Star Tracker Residuals showing 100 urad noise induced by post-process interpolation error

The probability of tracker star set changes seemed to go down at the encounter attitude. This is no doubt because there were several magnitude 2 and 3 stars in the field of view. It was known

weeks previously that there were brighter stars in the encounter FOV, but it was not known that this fact could be relied upon to improve performance. All available stars in the Impactor encounter star field is shown in figure 19 below. The stars tracked one hour after separation are shown in figure 20.

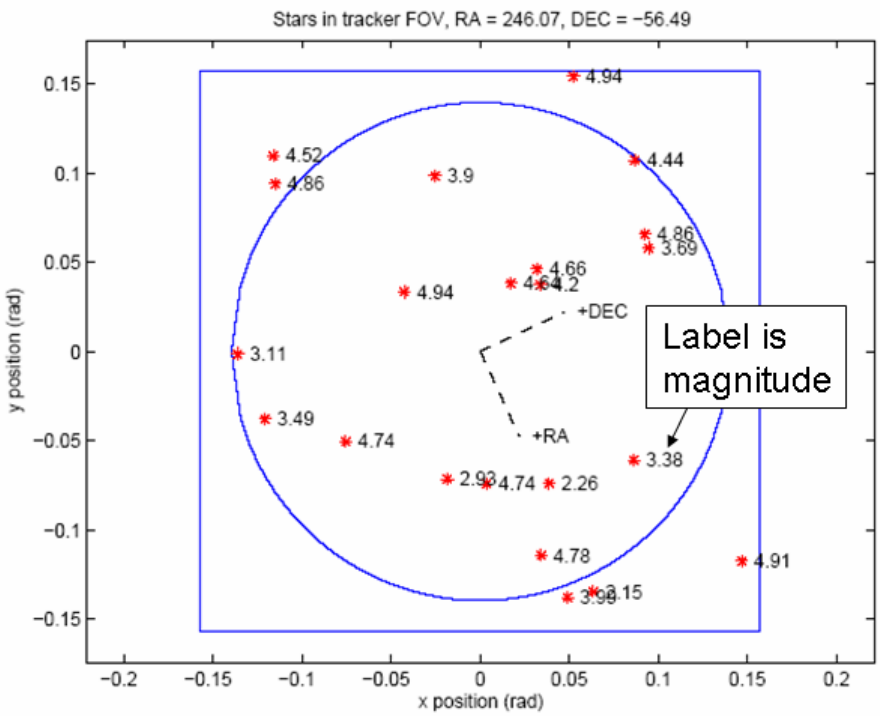

Figure 19. Available stars for Impactor star tracker at encounter

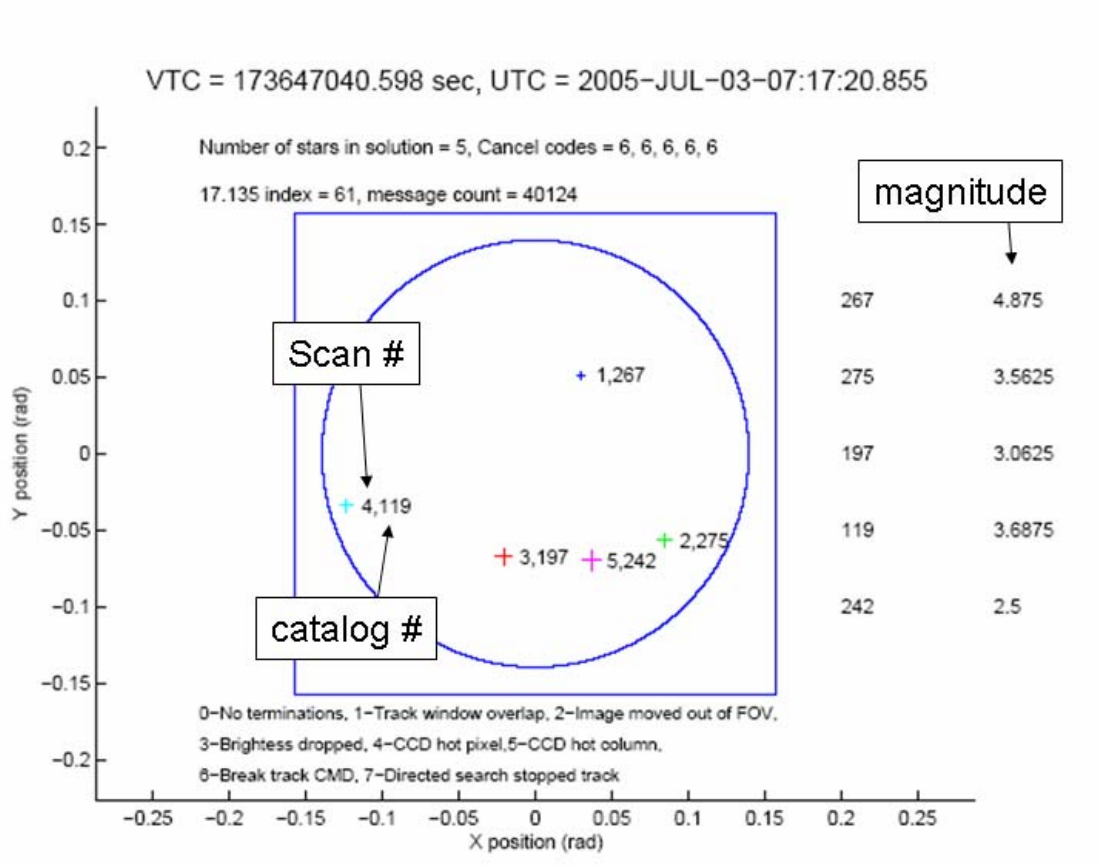

Figure 20. Impactor Encounter Attitude Star Field showing Group of magnitude 2 and 3 stars

After an 8 hour period for gyro bias estimator convergence (with the one hour time constant), the pre-loaded sequence switched to the 4 hour time constant. This had the desired effect of smoothing out the transients for encounter.

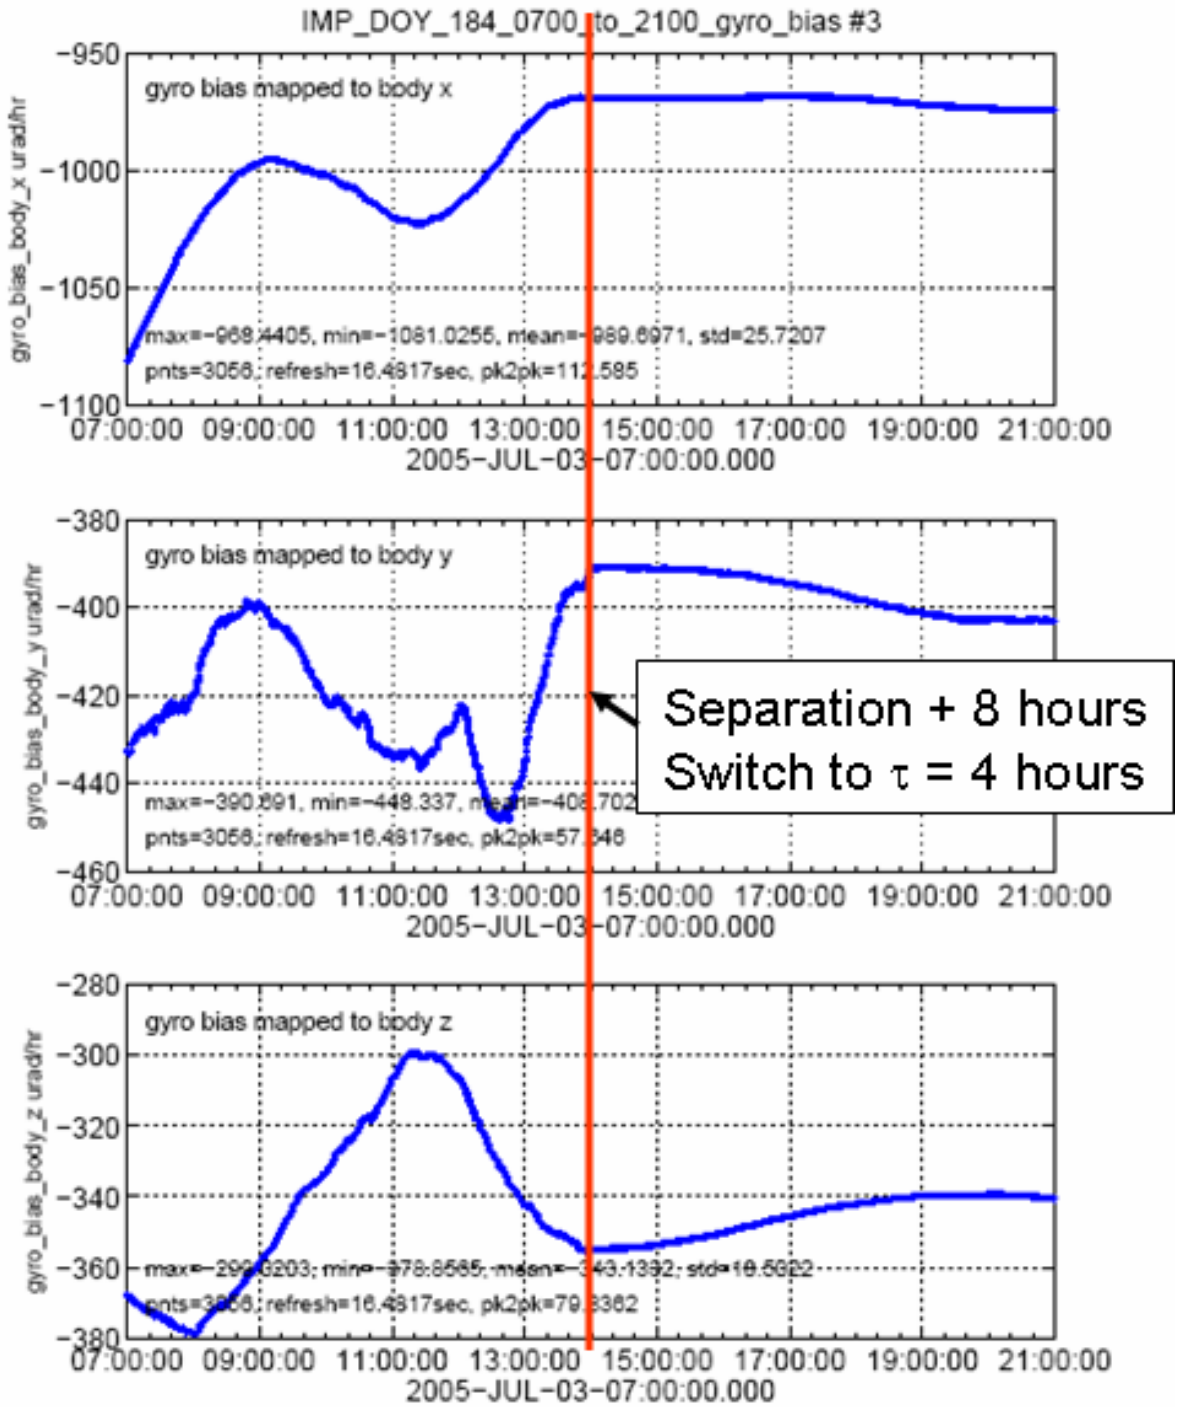

Figure 21. Impactor Gyro Bias Estimates showing switch to 4 hour time constant

<u>Impact Minus 7 hours</u>

The final hours before impact were full of activity. The gyro bias estimator, having converged for 12 hours, could be disabled 2 ½ hours before separation. Now, whatever bias error existed

would appear to AutoNav as a true miss. Fortunately, the performance was very good and resulted in only about 20 to 40 urad/hr bias error.

The three ITM burns occurred during the last two hours. They are easily seen as spikes in the star tracker residual plots in figure 22. The spikes are caused by expected tracker lost lock conditions during the fast rolls to/from burn attitude.

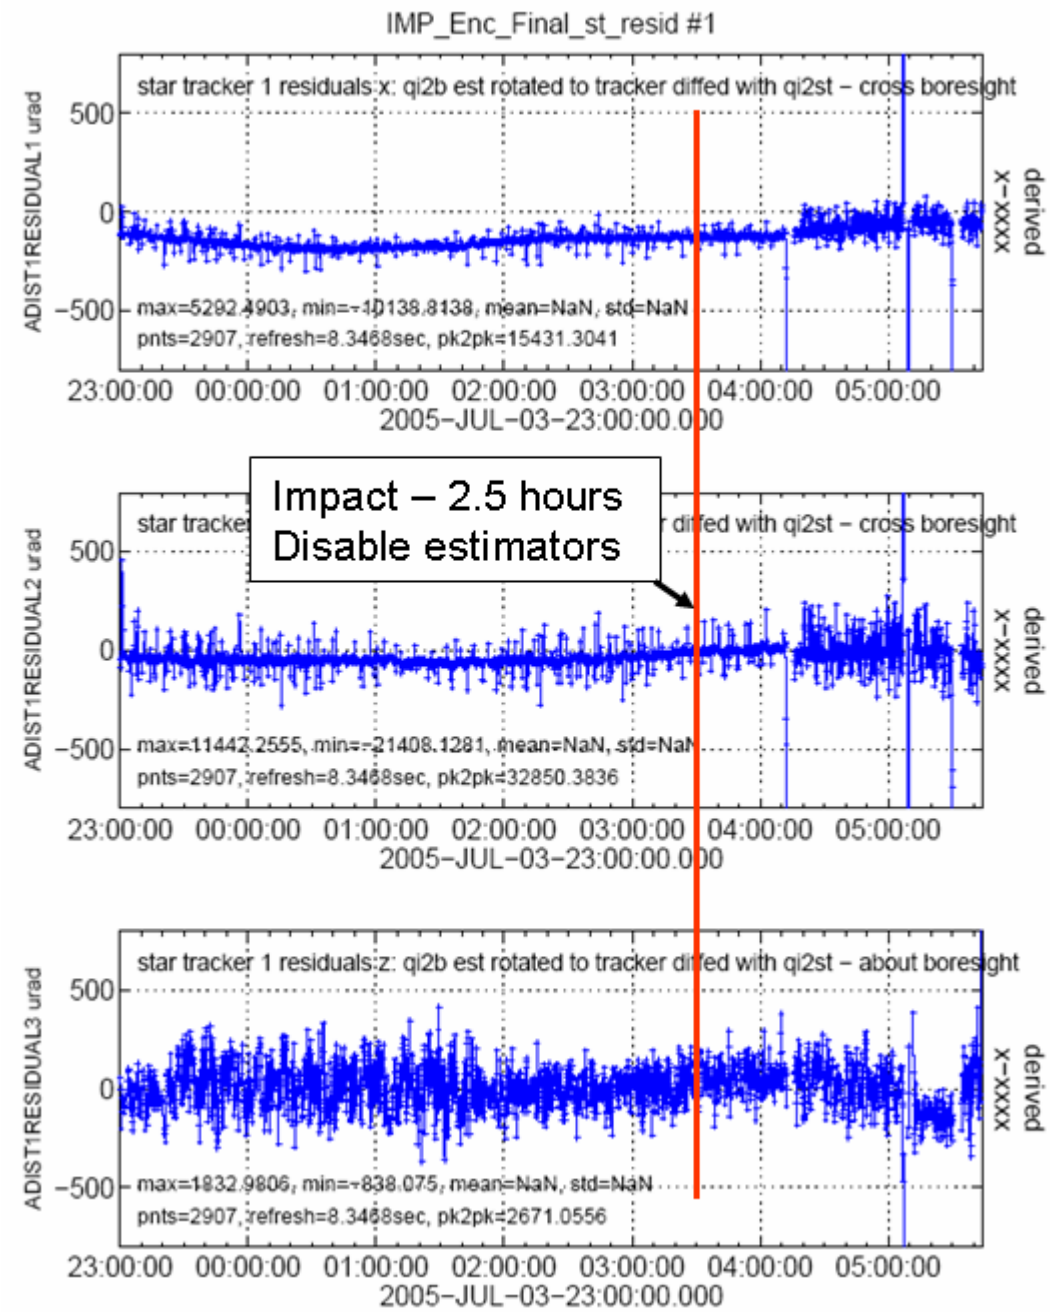

Figure 22. Star Tracker Residuals showing three ITM burns in last two hours

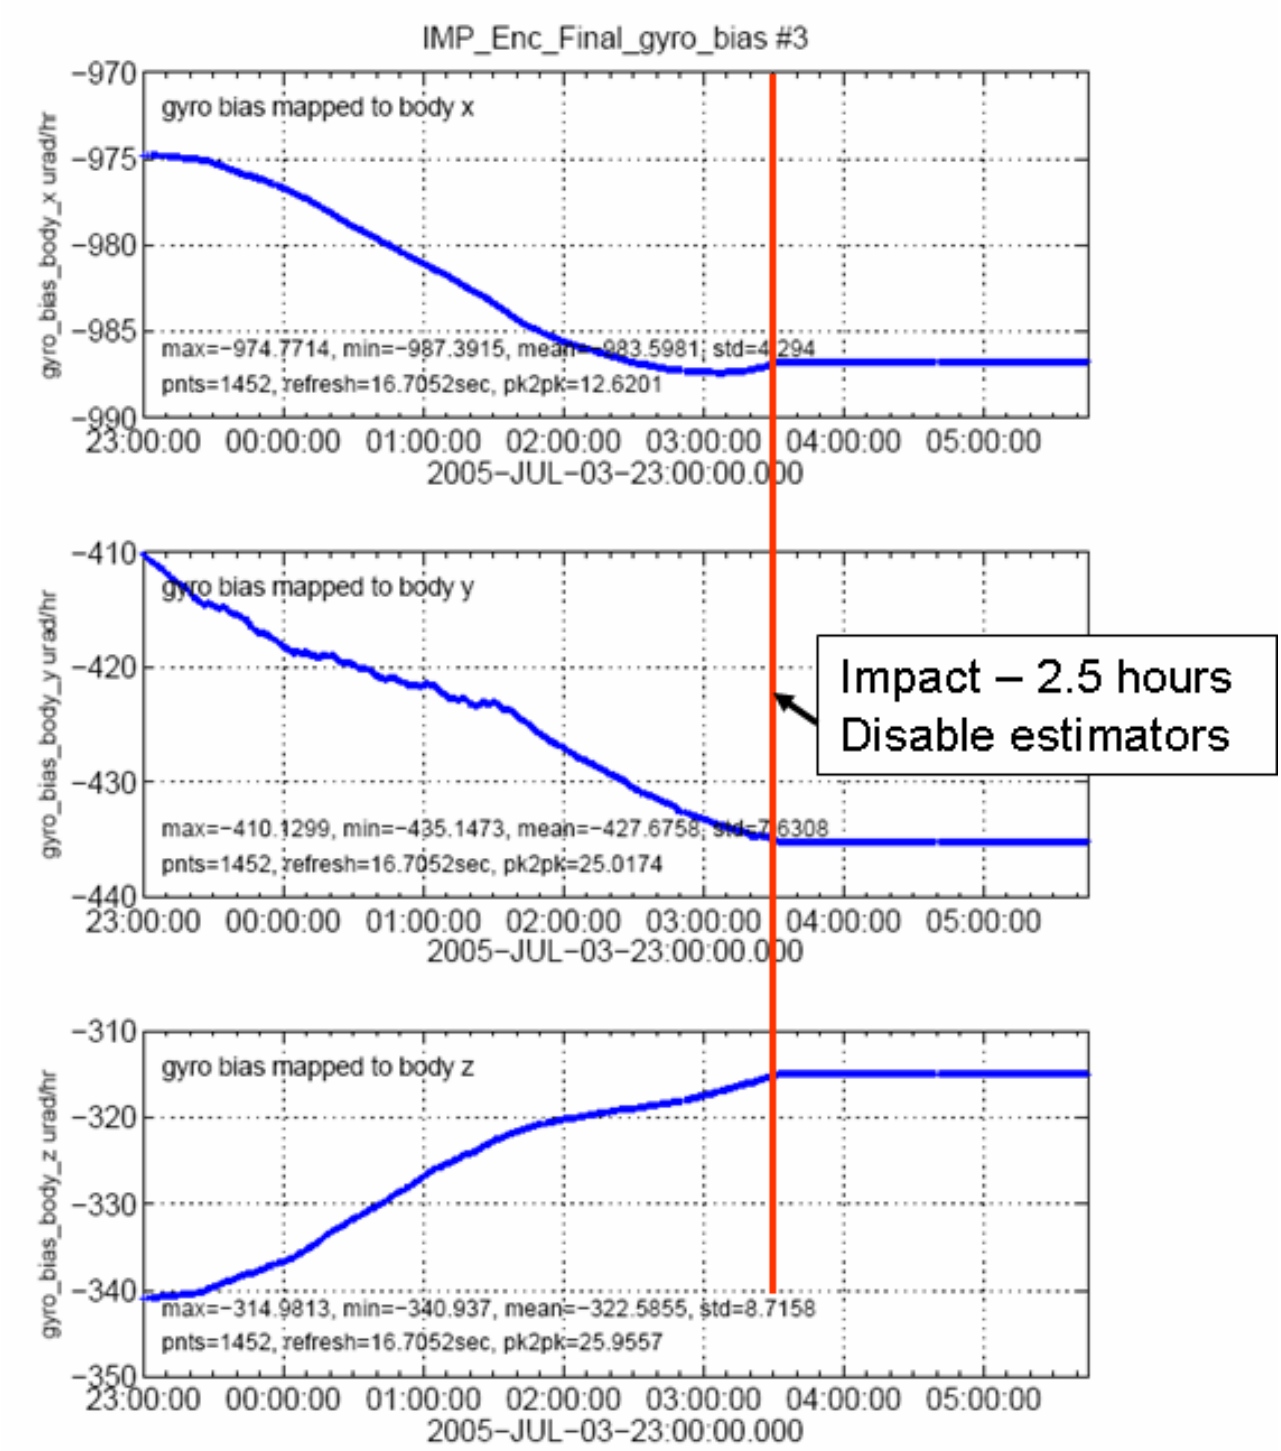

Figure 23. Gyro Bias Estimates showing disable at 2.5 hours before impact

After ITM 2, the Impactor was on a collision course. Still, the star tracker held mostly to the brighter stars. The stars in the field curiously made the image of a smile as shown in figure 24. This image made the flight team smile as well!

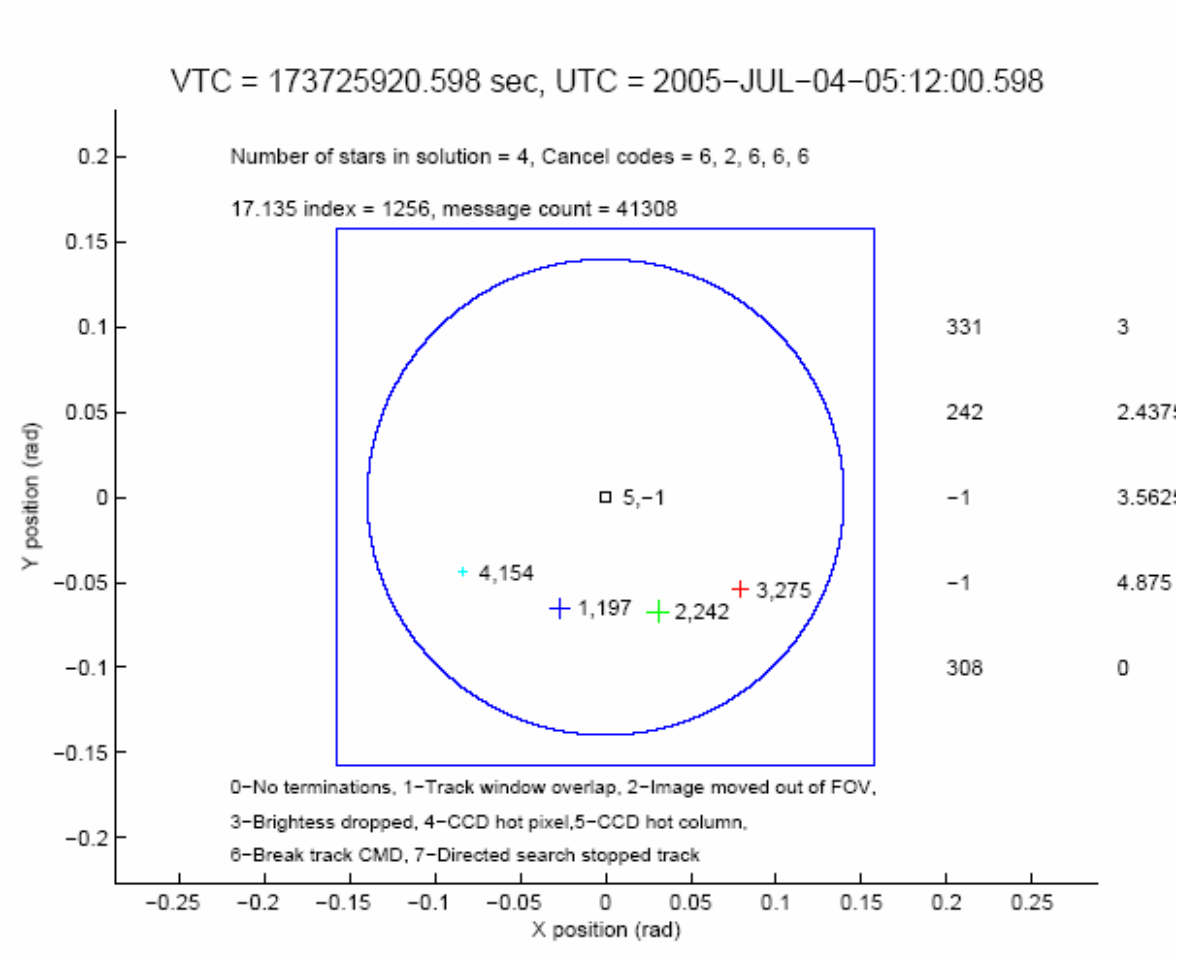

Figure 24. Impactor Star field showing star "smile"

The last frame of telemetry from the star tracker in figure 25 shows that the stars tracked changed significantly, except for one – 3rd magnitude star 197. Perhaps the bright coma obscured the stars. It was no matter as by then the tracker had done its job and was now destined to become part of Tempel 1.

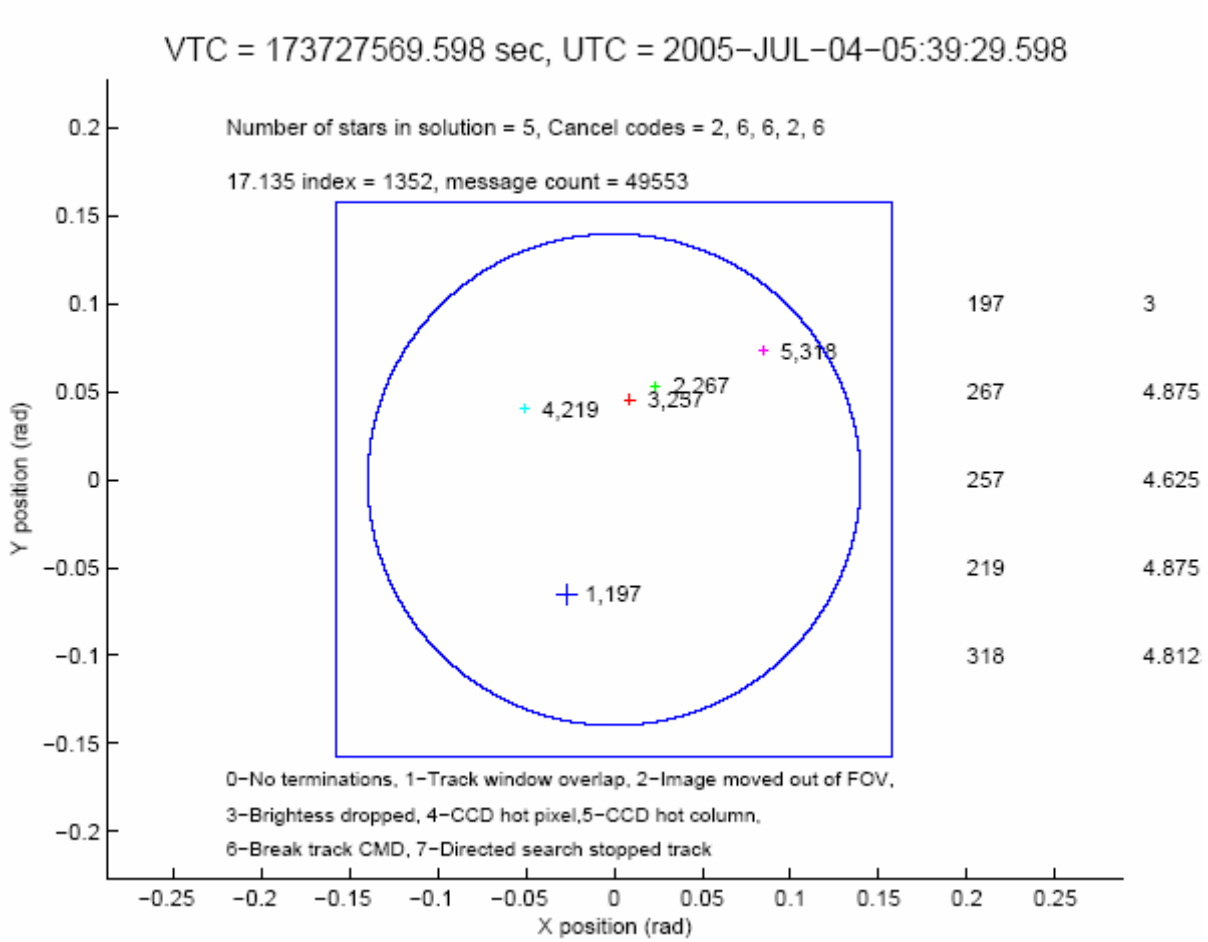

Figure 25. Last Telemetry Frame Tracker Stars

Conclusion

The Deep Impact mission was demanding with respect to both ADCS requirements and its relatively short, very busy mission timeline. The encounter design was complex and ambitious.

An existing set of attitude and gyro bias estimator algorithms was customized via parameter tables to meet stringent requirements imposed by starless AutoNav. Throughout the mission, the team evaluated daily and long term trends of attitude determination performance. Flyby spacecraft performance was acceptable for encounter, but Impactor performance caused a great deal of concern. After extensively evaluating Impactor star tracker performance, the flight team made critical adjustments to the system, including patching the star tracker software and sequencing gyro bias estimator gain changes.

The result was a successful impact by the Impactor spacecraft of comet Tempel-1 and a wealth of images returned by both the Impactor and Flyby spacecraft. The success of the DI ADCS was a result of an enormous effort by many hard-working individuals. Special thanks for the work described in this paper are due to Lew Kendall, Bill Trochman Doug Wiemer, Steve Collins,

Fred Smith, Steve Waydo, Jim Alexander, Dan Kubitschek, Nick Mastrodemos, Tom Bank, Meredith Larson, Stu Gray, Miguel San Martin, Tony Vanelli, and Shadan Ardelan.

References

[1] Kendall, Lewis, "Turn and burn – Deep Impact Imactor Autonomous Delta-V Maneuver Execution," Paper AAS 06-033, Proceedings of the Annual AAS Guidance and Control Conference, 2006

[2] Kubitschek, Dan, et. Al., "Deep Impact Autonomous Navigation: The Trials of Targeting the Unknown," Paper AAS 06-081, Advances in the Astronomical Sciences, Proceedings of the Annual AAS Guidance and Control Conference, 2006

[3] Wiemer, Doug, "Attitude Determination and Control for the Global Imaging System 2000," AAS 98-012

[4] Mendel, Jerry M., "Lessons in Estimation Theory for Signal Processing, Communications and Control," Prentice Hall Signal Processing Series, 1995