



**AAS 06-086**

## **MESSENGER IMU INTERFACE TIMING ISSUES AND IN-FLIGHT CALIBRATION RESULTS**

Daniel J. O'Shaughnessy, Robin M. Vaughan, Dave R. Haley,  
and Hongxing S. Shapiro  
The Johns Hopkins University Applied Physics Laboratory

---

### **29th ANNUAL AAS GUIDANCE AND CONTROL CONFERENCE**

---

February 4-8, 2006  
Breckenridge, Colorado

Sponsored by  
Rocky Mountain Section



AAS Publications Office, P.O. Box 28130 - San Diego, California 92198

## MESSENGER IMU INTERFACE TIMING ISSUES AND IN-FLIGHT CALIBRATION RESULTS

**Daniel J. O'Shaughnessy<sup>†</sup>, Robin M. Vaughan<sup>†</sup>, Dave R. Haley<sup>†</sup>,  
and Hongxing S. Shapiro<sup>†</sup>**

Flight operations are underway for the MESSENGER – MErcury Surface, Space ENvironment, GEochemistry, and Ranging – mission. As part of NASA's Discovery Program, MESSENGER will be the first spacecraft to observe the planet Mercury closely since the Mariner 10 flybys of the mid-1970s. The guidance and control system on MESSENGER maintains a 3-axis-stabilized spacecraft and uses (among other sensors) an inertial measurement unit (IMU) for attitude estimation, control, and autonomous cut-off of propulsive maneuvers. A larger than expected drift between the flight computer and IMU clocks was observed prior to launch and has resulted in problems with the rate and acceleration estimates in flight. The nature of the clock drift will be explained, as well as a technique for correcting the IMU time tag. Software exploiting this explained technique was loaded to the spacecraft in October 2005 and attitude performance before and after this software upload will be shown. Additionally, the latest results for in-flight alignment/sensor calibrations will be shown, both before and after the software load.

### INTRODUCTION

Flight operations are underway for the MESSENGER – MErcury Surface, Space ENvironment, GEochemistry, and Ranging – mission. As part of NASA's Discovery Program, MESSENGER will be the first spacecraft to closely observe the planet Mercury since the Mariner 10 flybys of the mid-1970s. This scientific investigation of Mercury will provide insight into the formation and evolution of all the terrestrial planets<sup>1</sup>. The spacecraft was launched on August 3, 2004, successfully completed its Earth flyby on August 2, 2005, and will make flybys of Venus and Mercury prior to orbiting the planet for one Earth-year beginning in March of 2011. MESSENGER carries a diverse suite of miniaturized science instruments to characterize the planet globally<sup>2</sup>, and the flybys assist in developing the focused science gathering of the year-long orbit phase of the mission.

Due to the extreme environment at Mercury, thermal and radiation concerns are the main factor driving the spacecraft design<sup>3</sup>. Protection from this environment is accomplished with a large sunshade mounted on the  $-Y$  body axis, which shields the spacecraft components from direct exposure to the Sun, as shown in Figure 1. Nominally, MESSENGER is flown with this sunshade centered on the sunline, and the Sun vector must remain near the  $-Y$  axis at all times while the spacecraft is in orbit about Mercury.

The MESSENGER guidance and control system maintains a 3-axis-stabilized spacecraft for the entire mission in order to satisfy the thermal and radiation constraints. This requirement is accomplished by using Teldix reaction wheels as the primary actuators for attitude control. Two Galileo Autonomous Star TRackers (A-STR) provide knowledge of inertial attitude. Adcole Sun sensors are used to provide Sun-


---

<sup>†</sup> Space Department, Mission Concept and Analysis Group, The Johns Hopkins University Applied Physics Laboratory, Laurel, Maryland.

relative attitude knowledge as a backup to the inertial sensors for spacecraft safety. Spacecraft rotation rates and translational accelerations are provided by a Northrop-Grumman (NG) S-SIRU (Scalable Space Inertial Reference Unit) inertial measurement unit (IMU) with four hemispherical resonator gyroscopes (HRGs) and Honeywell QA3000 accelerometers. The acceleration measurements are used to estimate imparted velocity during trajectory control maneuvers (TCMs).

The launch version of the G&C flight software (FSW), known as Build 4.7, uses data from the IMU in a variety of different ways. The accumulated angles are used in conjunction with the attitude estimates from the star trackers inside a Kalman filter for attitude determination. This attitude estimation is a low-rate task which operates at 1 Hz and processes 100-Hz IMU data and 10-Hz star tracker data collected over the previous second. The G&C control software uses a nonlinear control algorithm<sup>4</sup> to compute the wheel torque commands inside the high-rate software task at 50 Hz. Because the attitude estimation task produces attitude solutions at 1 Hz, the 50-Hz control task contains independent processing of the high-rate IMU data to interpolate the low-rate attitude estimation solutions. There is also 50-Hz processing of the IMU data in the guidance task in order to determine the cut-off time for TCM's by processing the accelerometer data.

Prior to launch, testing revealed that the flight computer and IMU clocks were drifting with respect to one another, and this caused periodic anomalies in the G&C software estimates of body rates and accelerations. Once on orbit, these problems persisted, and resulted in inaccurate burn termination and mis-estimation of spacecraft body rates and gyro biases<sup>5</sup>. An extensive analysis and modeling effort was undertaken shortly after launch to characterize the data interface and to update the G&C FSW to improve handling of the actual IMU data. Ground models of the clock drift were developed to simulate the characteristics of the flight IMU, and modifications were made to the flight estimation and control software blocks in order to handle the observed clock drift. Both the high- and low-rate software tasks were updated as a part of the post-launch software delivery (Build 4.9) that was uploaded to the spacecraft in October 2005.



**Figure 1. MESSENGER Spacecraft Configuration and Body Frame Definition**

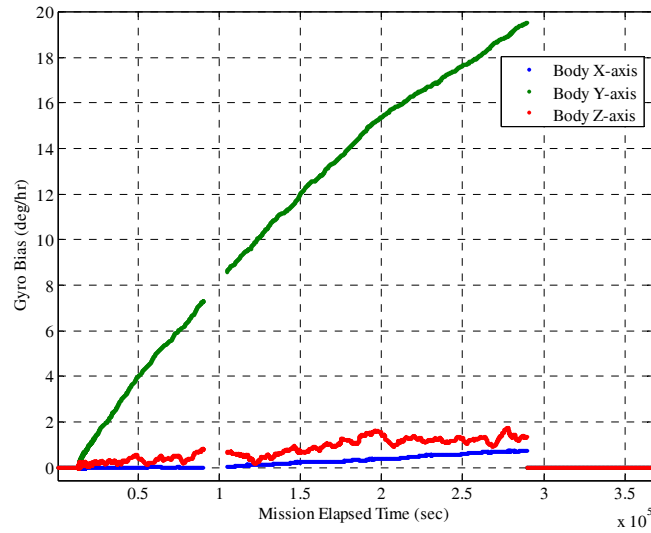
The flight software changes that made up the G&C portion of Build 4.9 included restructuring the rate and acceleration measurement processing to handle repeated and skipped measurements from the IMU, addition of a separate tracker-only estimation algorithm, estimation of star tracker alignment outside of the Kalman filter, and internal changes to the existing Kalman filter algorithm used in the nominal case where both tracker and gyro rate data are available. These changes improved attitude estimation and control performance.

The need for in-flight IMU calibration had been identified prior to launch due to problems with vendor calibration and the possibility of gyro alignment shifts that could result if inadequate welds in the gyro housing failed. Plans were made to have the spacecraft perform a series of turns while collecting high-rate raw data from both the gyros and star trackers, with the data later returned to the ground for analysis.


Filtering algorithms were used to process the data from the calibration maneuver to estimate gyro calibration parameters such as alignment and scale factor errors<sup>6</sup>. The first sequence turns was executed in September and October of 2004. A second calibration maneuver was performed in late October of 2005. Observability of the calibration parameters was maintained under different attitude constraints, and good convergence to many of the calibration parameters was achieved.

### SEEING SPIKES

MESSENGER was launched in its lowest level safing mode, Earth Acquisition (EA) mode, and remained there for several days to allow the subsystem teams to observe their systems' behavior. In this mode, the spacecraft maintains an axis at the Sun, as determined by the onboard Sun sensors, and slowly rotates around this axis at the rate of 1 revolution every 3.5 hours. During this time, it was observed that the estimated gyro bias terms from the attitude Kalman filter were growing far beyond their expected values, as shown in Figure 2. Figure 2 shows the growth in the gyro bias plotted versus the mission elapsed time (MET). MET is an integer number of seconds since the onboard oscillator started its count, which is




**Figure 2. Gyro Bias Growth as Seen in MESSENGER Flight Data from August 3-6, 2004**



**Figure 3. Slow Decrease in True Spacecraft Rate from MESSENGER Flight Data Taken August 3-6, 2004**

roughly at the launch epoch. Consequently, the spacecraft began to slow its rotation rate as the attitude control system attributed an increasing amount of the spacecraft rotation to gyro bias. This decreasing rate was independently verified through quaternion differences from the attitude solutions produced by the star trackers. Figure 3 shows the decreasing rate, where the black dashed line is the rate command. The estimated rate (in blue) continues to track the commanded rate, as it is the result of the rate from the IMU processing minus the estimated gyro bias. The true rate, which was obtained by differencing consecutive star tracker observations, was slowly decreasing as shown by the red data. Upon observation of the bias growth and the decreasing rotation rate of the spacecraft, ground controllers sent the command to disable the bias estimation at MET = 290,000. This is shown in Figures 2 and 3, as the bias estimate immediately drops to zero, and the true rotation rate begins to track the command.

After several days of ground analysis, the linear growth in the gyro bias was attributed to a time drift between the IMU and main processor (MP) clocks. Collection of the IMU high-rate data revealed a pattern in the IMU data where the IMU would cyclically go through a period of consistent measurement (a new message in every 10-ms sample window) followed at regular intervals by a noisy period of oversampled and undersampled IMU messages. These oversampled and undersampled messages result in time differences between consecutive samples of 0 and 20 ms, respectively. This irregularity in the time between consecutive measurements produces the “spike” pattern shown in Figure 4.




**Figure 4. IMU Spikes as Seen in MESSENGER Flight Data from August 3, 2004**

Although the behavior in the gyro bias, and subsequently the spacecraft rotation rate, was unexpected, the IMU spikes had been observed on the ground prior to launch. This phenomenon on MESSENGER was initially observed during the IMU integration, and though its cause wasn't fully understood at the time, it was known to be a result of the IMU clock walking relative to the spacecraft clock. It was determined that MESSENGER would launch without updating the G&C software to handle the clock walk problem, as there wasn't sufficient time in the schedule to fix the problem. Furthermore, the impact this clock walk would have on the G&C FSW was uncertain. Although spacecraft-level testing revealed the data patterns shown in Figure 4, this behavior didn't result in any adverse effects on the system performance during these tests. The gyro bias growth shown in Figure 2 could have been detected with a sufficiently long simulation, but running closed-loop simulations for many hours isn't practical at the spacecraft level. At launch there wasn't any evidence to suggest that there would be performance issues associated with the IMU clock drift. However, after the post-launch problems with the gyro bias growth, it was apparent that the G&C software needed modification to handle the IMU clock walk.

## UNDERSTANDING THE SPIKES


It was surmised during I&T that the patterns in Figure 4 resulted from the asynchronous nature of the spacecraft and IMU clocks; however, the details of this behavior weren't understood until after launch. The problems shown in Figures 2 and 3 are not internal to the IMU but are born out of the data interaction between the IMU and spacecraft flight software as well as the IMU data processing in Build 4.7. The MESSENGER IMU attitude message is comprised of four gyro measurements, four accelerometer measurements, a time tag (created by the IMU internal clock), and various health/status flags. This message is placed in the IMU's output data buffer every 10 ms, as determined by the IMU internal clock. The spacecraft bus schedule was designed so that data are extracted from the IMU data buffer every 10 ms, where the precision of this period is determined by the MIL-STD-1553(A) bus clock, the processor interrupt latency, and the main processor's master oscillator. The IMU clock is not synchronized with the MP master oscillator, and as such it is possible for the clocks to drift with respect to one another.

When the IMU clock runs slow relative to the spacecraft clock, its measure of 10 ms is slightly longer than the spacecraft's 10-ms interval. This difference results in the spacecraft (or equivalently, the spacecraft bus schedule) oversampling the IMU output buffer, resulting in duplicate measurements. This data pattern is shown in Figure 5. The red X's in Figure 5 represent the times the IMU made a message available in its output buffer termed an "IMU put," and the black lines represent spacecraft clock ticks, or times the IMU output buffer is sampled by the spacecraft bus ("IMU pulls"). Each spacecraft clock sample window contains a new IMU message (a red X), and so at each black line, the spacecraft reads the previous IMU message. If no new message is put into the IMU output buffer by between spacecraft pulls (as seen in the window indicated in Figure 5), the spacecraft reads the previous IMU message, which is a repeat of the message read at the previous pull. This situation results in missed or "stale" data. This condition occurs when the spacing between IMU puts (red X's) is greater than the spacing between the IMU pulls (black lines).



**Figure 5. IMU Data Pattern when IMU Clock Is Slower than Spacecraft Clock**

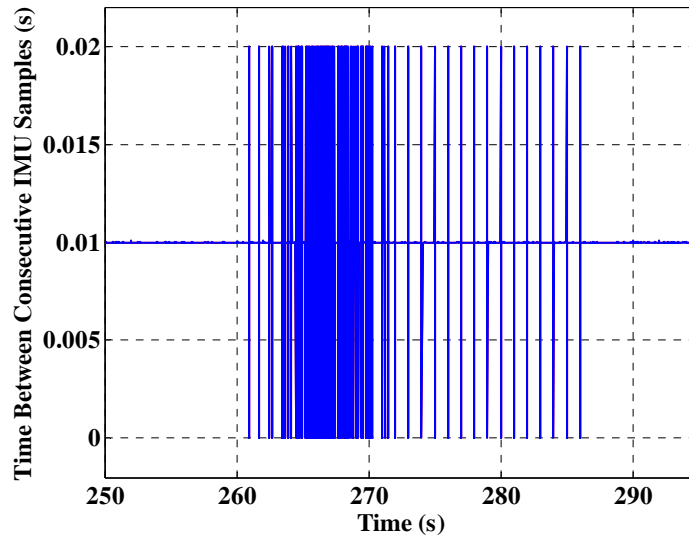
Similarly, when the IMU clock runs faster than the spacecraft clock, the IMU is undersampled, resulting in missed measurements from the IMU. This data pattern is shown in Figure 6. As before, the red X's represent the IMU puts, and the black lines represent the IMU pulls. All spacecraft clock sample windows contain a single IMU message, except one window contains a pair of IMU messages. Because the spacecraft pulls only the most recent data, one IMU message is missed. In Figure 6, the IMU puts are closer together than the IMU pulls, resulting in an extra IMU put in one of the spacecraft clock windows, as noted in the figure.



**Figure 6. IMU Data Pattern when IMU Clock Is Faster than Spacecraft Clock**

Jitter (or nonuniformity) on the clocks causes small increases and decreases in the spacing between consecutive IMU puts (IMU clock jitter) and IMU pulls (spacecraft bus clock jitter). Since the relative drift rate between the clocks is small (the specification for the IMU calls for 50 ppm, 1- $\sigma$ ) and because both clocks have some jitter, both repeated and skipped IMU measurements are readily seen in the flight data from Figure 4. The clock jitter complicates the detection of the true clock drift as both skipped and repeated messages are present in Figure 4. An expanded view of one of the timing “bursts” is shown in Figure 7. The random clock jitter causes the initial sequence of missed and repeated messages that appear to occur at random between 260 and 275 s. If neither clock had any jitter, the burst would be replaced by a single missed or repeated message that occurs once every  $\sim 200$  s, as the IMU put walks through the IMU pull as shown in Figures 5 and 6.

Jitter obscures the relative rate of the clocks, as it is difficult to see whether there are more missed or repeated messages during the random portion of the burst. By looking at the accumulated time (time tag difference) between IMU messages separated by a long interval (a thousand seconds), the clock drift is obvious. An average of the difference between IMU time tags that are 1000 s apart (by definition the time between 100,000 IMU pulls) reveals that 999.995 s have elapsed according to the IMU internal clock. This result shows that the IMU clock is running slow relative to the MP clock; furthermore, it demonstrates that the clock is well within specifications. At this drift rate, in 200 s the IMU loses 10 ms on the MP clock, resulting in a missed message. This is consistent with the spacing between spike bursts of about 200 s, as shown in Figure 4.



**Figure 7. Flight Data Sample of IMU Missed/Repeated Measurements at 100 Hz**


Note that the burst of missed/repeated measurements in Figure 7 seems to show two different stages. The early part of the burst shows a more random sequence of dropouts and repeated messages, followed by a regular pattern of repeated/missed messages. The clock walk/jitter described above explains the random sequence at the beginning of the burst. This random distribution of missed and repeated measurements is due to the random characteristic of the clock jitter.

The tail of the burst distribution is not random and was investigated further to develop a more complete understanding of the IMU-MP interaction. This tail is the portion of the burst from  $\sim 275$  to 287 s as shown in Figure 7. The portion of the burst is characterized by a uniform, repeatable sequence of a skipped message immediately followed by a stale message. This pattern repeats exactly once per second and is separated by 98 10-ms windows containing exactly one IMU message each. Three instances of this pattern are expanded in Figure 8 for clarity.

The tails of the burst are the result of an irregular sampling of the IMU in the spacecraft bus schedule. The IMU message is read by the spacecraft at slightly irregular intervals during the first 10-ms window of the spacecraft clock second. In other words, rather than pulling the IMU message from the IMU at MET.01 and MET.02, the spacecraft bus pulls the data at MET.0108 and MET.02 (where MET



represents the integer portion of the spacecraft clock time). This behavior causes the first window to be slightly longer than 10 ms and the second window to be slightly shorter than 10 ms. Because the first two “pulls” of the IMU measurements are further apart, it is more likely that a skipped IMU measurement will occur. This results in a longer time where two measurements are put out by the IMU during the first sample window. By a similar argument, because the second and third pulls of the IMU are closer together, this sample window will contain a repeated message for a longer duration. This bus schedule irregularity produces the pattern seen in Figure 8. In effect, this is similar to the clock jitter described above; however instead of the clock intervals between pulls changing at random, they are uniform and deterministic.



**Figure 8. Expanded View of the Tail of the Burst**

#### **BUILD 4.7 RESPONSE TO THE SPIKES**

The patterns shown in Figures 7 and 8 apply to the angular increment from the gyros ( $\Delta\Theta$ ) and the accumulated linear velocity increment ( $\Delta V$ ) from the accelerometers as the entire IMU message is subject to the clock drift. So at any time, the angular rate (or linear acceleration) of the spacecraft could be accurately obtained simply by taking consecutive gyro (or accelerometer) measurements and dividing them by the reported time increment between messages ( $\Delta t$ ). When the message is stale (i.e.,  $\Delta t=0$ ), the message can be ignored. When the time step is larger (an IMU message is missed), the correct rate is still obtained by dividing by the larger  $\Delta\Theta$ 's by the correspondingly larger  $\Delta t$ .

The launch version of the G&C flight software did not use the time tag from the IMU clock and instead always assumed that the time between messages was constant. The time tag passed into the G&C FSW was the spacecraft clock time at the IMU pull. Constructing the time tag in this fashion is equivalent to assuming that all measurements from the IMU were separated by exactly 10 ms. This turned out to be a critical flaw, as the drift rate of the clocks invalidated this assumption. The software did contain logic to detect and reject stale messages based on a stale count out of the IMU clock before applying the MET time tag to the message. This gave some protection against repeated measurements; however in the case of a stale message, the time tag assigned to that message was zero. This was an intended feature of the code, but it wasn't anticipated that there would be frequent stale messages, but rather it was designed to counter a failed IMU. In the end, this software logic caused the G&C FSW to reject more valid IMU data. When the time tag was set to zero the  $\Delta t$  between samples went from 10 ms to a very large negative number to a large positive number and finally back to 10 ms. These large positive and negative numbers failed a software check on  $\Delta t$ , causing both samples to be rejected. Table 1 shows this sequence.

While the MET (spacecraft clock) in Table 1 marches along at even intervals, the IMU clock counts up based on its own internal rate and produces a new IMU message once every 2500 counts (which is nominally every 10 ms). Table 1 shows this sequence for the first two messages, resulting in a nominal 0.01-s  $\Delta t$  as computed by G&C. The third IMU message is stale, so the clock counts from the internal IMU



clock don't change. This lack of change is detected in the FSW, and a time tag of 0 is assigned to the message. The G&C software differences this time tag with the previous one and computes a large negative number. This result fails the internal software check on the  $\Delta t$  between samples and the software does not use this measurement in the Kalman filter. In the fourth sample, the spacecraft gets a new IMU message that represents 20 ms of  $\Delta\Theta$ , but it rejects it because the computed  $\Delta t$  is much larger than the nominal 10 ms. In the sequence described in Table 1, there is only one stale data point, but two data points are rejected by the G&C software because they fail the limit checks on  $\Delta t$ .

**Table 1**  
**BUILD 4.7 IMU TIME TAG SEQUENCE**

| MET      | IMU Clock Counts | G&C IMU Time Tag | G&C Computed $\Delta t$ | Actual $\Delta t$ |
|----------|------------------|------------------|-------------------------|-------------------|
| 10001.00 | 25000            | 10001.00         | -                       | -                 |
| 10001.01 | 27500            | 10001.01         | 0.01                    | 0.01              |
| 10001.02 | 27500            | 0                | -10001.01               | 0                 |
| 10001.03 | 32500            | 10001.03         | 10001.03                | 0.02              |
| 10001.04 | 35000            | 10001.04         | 0.01                    | 0.01              |


Table 2 shows a similar sequence, but columns are added for the gyro angular increment and the rate computed by the G&C software during a missed measurement and a repeated measurement. As in Table 1, the sample with the stale IMU measurement and the one immediately following it are rejected due to invalid  $\Delta t$ . This action also prevents the software from computing a valid rate for those two messages. The G&C software also had no provision for retaining the previous good gyro angles, so the fourth sample is valid, but no rate can be computed. This point serves only as a reference for computing the following time steps' value of  $\Delta\Theta$ . The fifth point in the table shows that the point is not stale, so the HGC software applies a time tag which results in a  $\Delta t$  of 10 ms, even though the counts from the IMU clock indicate 20 ms passed since the last sample. This step is the missed message, and the rates computed by the software are incorrect, as they apply the correct  $\Delta\Theta$  over the incorrect  $\Delta t$ . This causes an apparent spike in the rates, and this rate spike results in the Kalman filter over-estimating the gyro bias, as shown in Figure 2.

**Table 2**  
**BUILD 4.7 IMU TIME TAG SEQUENCE**

| MET      | IMU Clock Counts | G&C IMU Time Tag | G&C Computed $\Delta t$ | Actual $\Delta t$ | Incremental Gyro Angle ( $\Delta\Theta$ ) | G&C Computed Rate                         |
|----------|------------------|------------------|-------------------------|-------------------|---|---|
| 10001.00 | 25000            | 10001.00         | -                       | -                 | $\Delta\Theta_0$                          | -   |
| 10001.01 | 27500            | 10001.01         | 0.01                    | 0.01              | $\Delta\Theta_1$                          | $\Delta\Theta_1/0.01s$                    |
| 10001.02 | 27500            | 0                | -10001.01               | 0                 | $\Delta\Theta_2$                          | -   |
| 10001.03 | 30000            | 10001.03         | 10001.03                | 0.01              | $\Delta\Theta_3$                          | -   |
| 10001.04 | 32500            | 10001.04         | 0.01                    | 0.01              | $\Delta\Theta_4$                          | -   |
| 10001.05 | 37500            | 10001.05         | 0.01                    | 0.02              | $\Delta\Theta_5$                          | $(\Delta\Theta_4 + \Delta\Theta_5)/0.01s$ |
| 10001.06 | 40000            | 10001.06         | 0.01                    | 0.01              | $\Delta\Theta_6$                          | $\Delta\Theta_6/0.01s$                    |

The processing in the high-rate control and guidance tasks is similar to the data shown in the above tables, but this task has only two IMU messages in any given software cycle. The net effect is a similar over-estimation of the spacecraft rate, as rejected measurements are ignored by the rate processing algorithm. While the estimation task drove up the bias estimate in the presence of the spikes, the rate spikes shown in Table 2 cause the controller to respond by adjusting the spacecraft rate to compensate for the apparent spike. This adjustment caused a transient response of the spacecraft during each of the burst periods. In EA mode, the spacecraft slowed its rotation rate, as onboard processing of the IMU data made the controller believe that the spacecraft was actually rotating faster than commanded. Once the IMU clock walked further, and the burst went away, the controller brought the spacecraft rotation rate back up to the commanded rate. A plot of the controller response through the burst can be seen in Figure 9. This plot was created by ground processing the high-rate IMU telemetry and represents the true spacecraft rate.

Note that the transient decrease in the rotation rate contains two distinct phases, corresponding to the two phases of the burst. The first portion shows the marked decrease in the rotation rate, followed by a rise back to the near steady-state value, appearing at ~275 and ~485 s in Figure 9. The second, more subtle phase immediately follows the first and shows a short (about 10-15 s) period at a rate slightly below the steady-state value. The first phase corresponds to the noisy part of the burst shown in Figure 7. The second phase corresponds to the tail of the burst shown in Figure 7, and expanded in Figure 8. The abscissa of the plots shown in Figures 4 and 7-9 are consistent, and taken together they show the spacecraft response as a result of the attitude controller processing the IMU timing burst. The short time scale shown in Figure 9 (several hundred seconds) is not sufficient to show the slowly decreasing rate due to the increasing bias, shown in Figure 2, which is apparent only over many hours.



**Figure 9. MESSENGER Rotation Rate in Earth Acquisition Mode on August 4, 2004**

Of course, the patterns in the flight data (Figure 8) are much more complex sequences of repeated and missed measurements than those shown in Table 1 and Table 2, but the effects are the same. It can be succinctly described in these terms: when an IMU message is repeated, the data are rejected; when an IMU message is skipped, the measurement processing causes an error in the estimated rate which *always* overestimates the rate. In the case of the Build 4.7 software, when batch processing these samples at 1 Hz inside the Kalman filter, this over-estimation of the rate caused the rate from the IMU to look artificially high. This “extra” rate was interpreted as gyro bias by the filter, which resulted in the linear growth in the gyro bias shown in Figure 2 and the attendant slow decrease in the spacecraft rotation rate from Figure 3. Similarly, the 50-Hz processing in the control algorithms overestimated the spacecraft rate during the spike period only, and the control system responded by introducing a transient decrease in the spacecraft rotation rate, as shown in Figure 9. The architecture of the G&C software had to be changed to handle the unanticipated missed and repeated IMU measurements, allowing the Kalman filter to converge to the correct gyro bias and preventing the controller from initiating unwanted excursions.

#### **IMU DATA PROCESSING CHANGES FOR BUILD 4.9**

In Build 4.7 of the flight software, the attitude estimation task was based on the assumption that there were always 100 IMU messages each second. If the IMU data were rejected for any reason, the rate measurement was replaced with star tracker rate information. Build 4.9 altered this approach to recognize that there may be several stale messages in any one second, without resulting in gaps in the accumulated angles or accumulated  $\Delta V$ . In other words, a stale or missed message does not mean any missed information. Additionally, by using the information in the IMU time tag, the correct rate/acceleration can

be computed when IMU messages are missed. These are the two fundamental changes necessary to allow the software to handle the IMU spikes.

The processing in the low-rate task (attitude estimation) requires that the IMU data have correct relative time tags in order to reconstruct the spacecraft body rates. The low-rate task also requires that the absolute time (or alternately, the MET) of the IMU measurement be correct, so that it may be aligned with the star tracker measurement for use in the Kalman filter. The requirement that the time difference in the IMU measurements be known relative to one another suggests using the IMU time tag directly. However, since the IMU measurement must also be known relative to the tracker measurement (which is tagged with a spacecraft MET), the spacecraft clock is a more appropriate reference frame for the IMU time tag. Resolution of this apparent inconsistency is the driving force behind the design for the new G&C flight software.


Because the IMU clock and the spacecraft clock are asynchronous, it is not easy to determine the IMU measurement's time tag in terms of the spacecraft MET. When the spacecraft pulls data from the IMU internal buffer, the time when that sample was made available to the spacecraft is unknown, other than that it occurred prior to the current spacecraft sample time. Looking back at Figure 6, when the IMU is sampled (at any black bar), the message the spacecraft receives is the previous red X. Anytime a new message is read by the spacecraft, it is known for certain that it occurred sometime within the last 10 ms (as measured by the spacecraft clock). However, it is impossible to know for certain when in the last 10 ms the IMU put that message out. Additionally, if the message is stale, it is known that no message was put by the IMU in the last 10 ms. Said another way, getting the same message on consecutive pulls means there was no IMU put during the last sample window. If the current pull is not the same as the previous pull, it is known that the IMU message occurred somewhere within the last 10 ms, as measured by the spacecraft clock.

From the above information, the IMU message could be tagged at the middle of the spacecraft clock sample window, leaving +/-5 ms of ambiguity on the time tag. This scheme works well for sequences of IMU measurements where messages are infrequently stale and is similar to the way the IMU data were time tagged in the launch version of the flight software. However, for sequences that include frequent missed/stale messages, this scheme is not effective, as has been demonstrated by MESSENGER's flight performance. In the case shown in Figure 6, the sample window that includes two IMU messages produces a problem. With this method, for the sample window just prior to the one with two X's, the red X is tagged up at the current spacecraft sample time, say it is MET.00. The next time the spacecraft goes to pull a message, two samples have occurred, and the spacecraft only sees the most recent. The first message in that window is missed, and the second gets read and assigned a time tag at the current MET, MET.01. Now, from the spacecraft clock's perspective, two IMU measurements occurred, separated by 10 ms, and thus the spacecraft clock tags them appropriately. But from the IMU's perspective, these messages are 20 ms apart, since the spacecraft skipped over one message. So which clock is correct, the IMU that says compute body rates over 20 ms, or the spacecraft which says compute the rates over 10 ms? From Figure 6 it is easy to see that the red X's in this example are indeed separated by nearly 20 ms. The application of time tags according to the spacecraft clock is misleading because the spacecraft has no idea that the first IMU measurement, tagged at MET.00, was really closer to MET.00-0.01. The IMU message must be time tagged with better precision than +/-5 ms otherwise, when the spacecraft misses messages, the computed rates will be wrong.

By taking the difference between the last time tag from this second and the last time tag from the previous second, the number of IMU messages in the last second can be computed. When the IMU is running normally (no hardware failures or unexpected bad data points and messages are only missed or repeated due to asynchronous clocks), the spacecraft receives 99, 100 or 101 messages in any 1000-ms window (as defined by the spacecraft clock). Said another way, the IMU will put out no fewer than 99, but no more than 101 messages in any one spacecraft clock second. When the spacecraft detects 99 or 101 messages, it is known that the IMU messages are put at nearly the same time they are pulled. The burst is the consequence of the clock boundaries walking across one another, so it is known that the clock steps are nearly aligned when the spacecraft experiences the burst. Consequently, the spacecraft can use that information to assign the IMU time tag more accurately. (Actually, the MESSENGER S-SIRU time tag increments at a rate of 250,000 counts/sec, but this time tag is stored in the IMU message as a 16-bit


unsigned integer. This procedure causes the time tag message to roll over several times each second, and this rollover must be appropriately handled by the software.)

By the above rationale, if the spacecraft computes that 101 messages were received in the last second, it is known how to time tag these messages with a high degree of precision. If the spacecraft-computed  $\Delta t$  was exactly 1000 ms (by definition, because each spacecraft clock tic is exactly 10 ms in the spacecraft clock frame), the IMU managed to count 1010 ms, as computed by the difference in the time tags between the first message in the second and the last message in the second. This outcome means that the very first message in the second was close (but just after) the prior major frame (MAF) boundary, and the last message in the second was very close (but just prior) to the current MAF boundary. (Note that MAF boundaries are coincident with increments in the MET clock.) This sequence can be seen in Figure 10, but the figure only shows 10 sample windows between MAFs for clarity.



**Figure 10. IMU Measurement Sequence for 11 IMU Samples in 10 Spacecraft Clock Windows**

This overproduction of IMU data (101 time steps in 1 s) can be the result of the IMU clock running fast, or it can be the result of jitter on the clocks. In either case, it is known that the first message the IMU put out during that second was very near the previous MAF. It is important to note that this is the first message put by the IMU, not necessarily the first pull by the spacecraft. In other words, the spacecraft would have had a missed message in the first minor frame (MIF), where the first IMU message occurred very near the MAF boundary, and then the second IMU message occurred just prior to the spacecraft MIF boundary. This scheme is depicted in Figure 11. The data sequence in Figure 11 causes a missed message in the first MIF of the second, but this outcome can easily be detected by computing the difference between the first IMU time tag from the first MIF of the second and the time tag from the last MIF of the prior second. It is trivial for the software to store the time tag of the last IMU measurement from the previous second to detect accurately any missed messages during the first MIF.



**Figure 11. Missed IMU Message in the First Minor Frame**

If the number of samples in any one second is either 99 or 101 it is known that the IMU clock and spacecraft clock frames are nearly aligned. As with any single MIF, it is a straightforward computation to obtain the number of IMU messages obtained in a second by differencing the appropriate IMU clock counts. If these counts reveal that 99 samples were made available by the IMU during the last spacecraft clock second, it is known for certain (barring unusual behavior by the IMU) that the last IMU message occurred roughly 10 ms prior to the last pull. In other words, a sample was just missed at the most recent MAF boundary, as that is the only reasonable way the IMU can put out 99 measurements in one second. Likewise, if the clock counts from the IMU reveal that 101 IMU messages were made available over the last 1 s (as measured by the spacecraft clock), it is known for certain that a measurement occurred just

before the most recent MAF boundary. This is the only reasonable way that the IMU can put 101 messages in a second.

From the above argument, it is easy to detect when the clock frames are nearly aligned, and we can accurately time tag IMU measurements when this is the case. Extending the idea a bit further, the software can easily trend the time (in terms of spacecraft clock seconds) between successive periods of aligned clock frames. This defines the clock drift period, and a linear interpolation can be used to compute the time tag of IMU measurements between the times when these clock frames are aligned. This correction is based on the assumption that the drift rate between the clocks is nearly constant, which is valid on the basis of the flight data. In flight, the clock drift period has been observed to be between 180 and 220 s. However, the majority of this drift appears to be thermally driven, and as such the drift rate is nearly constant (+/-2 seconds) over short intervals (hours). A more accurate time tag is constructed by determining the number of samples the IMU put out during the last second. This tag is accomplished by obtaining the change in clock counts from the IMU over the entire second and dividing this difference by the nominal clock counts per second.

The MESSENGER IMU clock is running slow relative to the spacecraft clock, so the first indication of a burst is a second where the IMU puts 99 messages out between major frame boundaries. (A fast-running IMU clock would put 101 messages first.) By timekeeping the interval between consecutive beginnings of these bursts (always marked by a long period of 100 samples/s followed by a 99 sample second) the period of the clock drift can be calculated. Since this drift period remains nearly constant, the IMU time tag between bursts can be calculated as:

$$tt_{IMU} = MET_{IMU} + (t_{CD} - t) / t_{CD} * 0.01 \quad (1)$$

where  $tt_{IMU}$  is the IMU time tag in the spacecraft clock frame,  $MET_{IMU}$  is the time the IMU measurement was pulled in terms of the spacecraft clock,  $t_{CD}$  is the clock drift period,  $t$  is the elapsed time since the beginning of the last clock period, and all times are in seconds. By using the clock counts from the IMU, the relative drift rate of the clocks can be computed, and in turn the IMU time tag can be interpolated into the spacecraft clock frame. This procedure produces a time tag accuracy of nearly the jitter on the clocks (less than 1 ms). This improved time tag accuracy also aids in reducing the Kalman filter residuals.

This technique for interpolating the IMU time tags is useful when processing messages for the attitude estimation task, which runs at 1 Hz. This method requires some batch processing of measurements, in this case 100 IMU measurements at a time, to detect the clock walk period and to interpolate the IMU measurements inside the 10-ms spacecraft clock window. This process is not useful in the high-rate attitude control task, which runs at 50 Hz, as this task operates on two IMU measurements at a time. Recall that the attitude control task is not concerned with an absolute time tag (a time tag relative to the spacecraft clock frame, or alternately the star tracker measurement time tag), but instead requires that time tags be correct relative to the sequence of IMU measurements. For this purpose, it is enough simply to use the counts from the IMU clock to determine the time between gyro measurements. Although these time tags are given in counts, they can easily be converted to seconds by the conversion factor supplied by the vendor. So to fix the high-rate gyro and accelerometer processing, it is enough to use the IMU clock time tag when computing the body rates and linear acceleration. As long as the IMU messages are internally consistent (i.e., the change in time tags matches the change in gyro angles and accelerometer-integrated velocities), the body rates and accelerations derived from these messages will be correct.

#### **Additional Software Changes for Build 4.9**

In addition to changing the way the IMU data were processed for the new software update, the attitude estimation Kalman Filter also underwent restructuring. The previous version of the filter used body rate measurements from the star tracker in place of gyro rates when the gyro data were unavailable or rejected. In these cases, the star tracker measurement was used for propagation *and* update in the Kalman filter, which could lead to erratic filter performance. This procedure was altered as a part of Build 4.9 so that when significant gyro data have been lost or rejected, the attitude estimation task no longer runs the Kalman Filter, but instead operates in a “Star Tracker-Only Mode.” In this mode, the attitude is determined directly from the star tracker observation, and spacecraft body rate information may be obtained from the star trackers directly. If the tracker rate measurement is unavailable the spacecraft body rates are

determined from the sequence of attitude observations. This processing mode is not sophisticated, as it considers only the last valid quaternion and rate from the star trackers in determining the attitude estimate. This situation continues to be the case when both trackers are on; there is no effort to combine the measurements from the trackers to extract a more optimal attitude estimate. Persistent operation in this mode would mean that at least two gyros had failed, and the software is not intended to be tolerant of double faults. The tracker-only mode is designed for temporary periods of gyro data outage. Enhancements would be needed to meet science pointing requirements using this mode on a permanent basis in the event of an unrecoverable IMU failure.

The last significant change made to the attitude estimation task as a part of Build 4.9 was removal of the star tracker relative misalignment states from the Kalman filter. This tracker alignment estimation is easily handled external to the Kalman filter, thereby simplifying the filter formulation. Rather than carrying three filter states for the relative tracker misalignment, the misalignment is estimated by low-pass filtering the quaternion error between synchronous star tracker measurements. The low-pass filter is active only when both trackers are enabled and producing valid data. This condition is expected to hold over only a very small portion of the mission, as the operational plan is to have one star tracker enabled for normal operations. Note that this estimation is of the alignment error relative to the nominal (pre-launch optical) alignment, which is expected to be small and slowly varying, mostly due to the spacecraft thermal environment. These misalignment states aren't important in any case; the important alignment information is the star tracker-to-IMU (or more precisely, gyro) alignment. These alignment numbers are estimated on the ground via the calibration maneuvers that are discussed below.

#### BUILD 4.9 PERFORMANCE

The new spacecraft FSW which included the G&C software changes described above was loaded to the spacecraft on October 11, 2005. A main processor (MP) reset command was sent on October 24, 2005 to cause the MP to reboot on the new software. Figures 13-15 were taken from a test conducted shortly after loading the new software. The guidance algorithms were set up to command the spacecraft to perform a rotation similar to the EA-mode attitude rotation (one revolution about the Y-axis every 3.5





Figure 12. Gyro Bias Estimation Using New G&C FSW


hours). This operation mimicked the post launch operations shown in Figures 2 and 3. The new software was expected to prevent the linear growth in the gyro bias shown in Figure 2 and the transient decrease in spacecraft body rates shown in Figure 9. The new software did not show this linear growth in the bias, as can be seen in Figure 12. The bias estimates remained stable for the duration of the test (~15000 s), but are only plotted for 3000 s to illustrate the initial convergence. Evidence that the control task handled these spikes correctly is provided in Figure 13, as can be seen by noting the differences from Figure 9. This plot

only contains ~300 s for clarity, although the transient rate decreases were not seen in any portion of the test. Filter residuals on the new software were as expected, as can be seen in Figure 14. Further examination of the raw IMU telemetry on the ground revealed that the data still contained the “spikes,” but the high-rate control task (50 Hz) and low-rate attitude estimation task (1 Hz) were processing the rate information correctly.



**Figure 13. Spacecraft Y-Axis Rate in Earth Acquisition-like Rotation Using Build 4.9**


The revised Kalman filter was not without some unexpected behavior, however. The prior version of the flight software had an error in the Kalman filter formulation, which caused the process noise to be an order of magnitude too large. These formulation errors masked several minor problems that were ironed out during the testing of the new FSW. One such consequence of the large process noise was that the Build 4.7 software would accept any star tracker measurement into the filter despite its noise (covariance) condition. So the only time the Kalman filter rejected a measurement on the launch version of software was when the tracker data were flagged invalid or the tracker was off.



**Figure 14. Kalman Filter Residuals Using Build 4.9**



The new software fixed this problem with the process noise, but an unexpected consequence was the rejection of a high percentage of star tracker measurements by the filter. (A  $\chi^2$  test is performed in the software prior to running the Kalman filter to determine if a measurement is to be accepted into the state estimate. If the measurement is inconsistent with the state, based on the process noise and measurement covariance, the measurement is rejected and the filter is not run that second. If this condition persists for some parameterized amount of time, the filter is reset.) It was expected that the filter would accept nearly every star tracker measurement, as the threshold for the  $\chi^2$  test was set very conservatively. However, as can be seen in Figure 15, the software was rejecting nearly 50% of the measurements. The pertinent enumeration of the status flag is: 7 indicates both IMU and star tracker data were good, and 6 indicates the IMU data were good, but the star tracker measurement was rejected due to the  $\chi^2$  test. This figure shows the status flag toggling between 6 and 7, and rejecting roughly 50% of the star tracker measurements.



**Figure 15. Estimation Status Flag Using Build 4.9**

After analysis of the star tracker measurement data, the problem was traced to the measurement covariance reported by the star tracker. The diagonal of the covariance matrix is output as a part of the normal star tracker measurement message and is used directly in the Kalman filter. This number is output as a 16-bit integer and is a scaled, quantized representation of the star tracker estimate of the true measurement covariance. This scaling and quantization can result in zeros being reported for the cross-boresight covariance terms, which indicate no measurement uncertainty on these axes (a near impossibility). It is these measurements that are rejected by the filter due to the misreported covariance.

This problem could be considered a cosmetic issue, as the rejection of the tracker measurements does not adversely affect the filter performance. When these measurements are rejected, the filter continues to propagate the state based on the good gyro data and will continue to do so until a good tracker measurement is available. If this condition persists for 50 s, the filter is reset, which is undesirable but unlikely, since the tracker tends to report a zero covariance about half the time. Consecutive instances of a zero covariance are common, but 50 consecutive occurrences is a very-low-probability event. Upon filter reset, the tracker measurements are forced into the filter (no  $\chi^2$  test), so a filter reset is not a flight concern until precise attitude determination is required during the science observations that will occur at Mercury. Fortunately, this covariance matrix can be manipulated via flight parameters for scale factor and bias for each diagonal term in the matrix. These parameters will be adjusted as a part of the calibration maneuver parameter load to be completed sometime in early 2006 to prevent the Kalman filter from seeing zeros in the covariance matrix. This fix has been verified on the ground to eliminate the  $\chi^2$  test failures.

#### **ATTITUDE SENSOR CALIBRATION**

As mentioned previously, a need to perform a calibration of the IMU had been identified prior to launch due to problems with vendor calibration and the possibility of gyro alignment shifts. The calibration

is accomplished by executing a series of turns while collecting high-rate data from both the gyros and star trackers, with the data later returned to the ground for analysis. These data are then processed via an Alignment Kalman Filter<sup>6</sup>, which results in updated gyro alignments, as well as information about gyro biases and scale factor errors. Current operation plans call for the IMU calibration to be repeated approximately every 6 months throughout cruise, and on-board calibration parameters will be updated as needed should the ground solutions indicate significant changes as the IMU ages.

The MESSENGER S-SIRU is internally redundant and can be operated from either of two Processor and Power Supply Modules (PPSMs). Normally PPSMA is used with PPSMB serving as a backup in the event of a problem with PPSMA. Separate calibrations were performed for each PPSM by NG prior to installation the spacecraft. The in-flight calibrations also operate the IMU from each PPSM in order to detect changes from the ground calibration values.

To date, the spacecraft has executed three sets of IMU calibration maneuvers as shown in Table 3. The first attempt to perform calibration maneuvers on September 13, 2004, was aborted prematurely by an unexpected demotion to safe mode that was unrelated to any G&C activity. The second attempt to execute these same turns successfully ran to completion on October 6, 2004. The 2004 calibration maneuvers were executed while the spacecraft was near 1 AU, and thermal constraints were less stringent, resulting in a wider array of available attitudes. The second set of calibration turns was performed in late October 2005 using the new flight software, Build 4.9. These sequences were executed with the spacecraft between 0.78 and 0.75 AU from the Sun, which imposed the constraint that the spacecraft must maintain pointing of its sunshade to within roughly 15° of the sunline. This constraint forced the calibration turn sequence to be redesigned to maintain observability of the interested parameters.

**Table 3**  
**IN-FLIGHT IMU CALIBRATION EXERCISES**

| Date & Time        | Comments  |
|--------------------|---|
| September 13, 2004 | Aborted during Y axis turns on PPSMA; turns using PPSMB not performed   |
| October 6, 2004    | Completed all turns with IMU on PPSMA and PPSMB                         |
| October 27, 2005   | Completed all turns with IMU on PPSMB                                   |
| October 28, 2005   | Completed all turns with IMU on PPSMA; same turn sequence as October 27 |

There were sufficient data to perform the ground filter calculations from each of these tests, although the test of September 13, 2004, only yielded solutions for operation on PPSMA. The results for each exercise are compared in the following sections.

### **Calibration Maneuver Design**

Between launch and March 2005, MESSENGER was oriented with the back of the spacecraft facing the Sun and the spacecraft remained outside of 1 AU from the Sun. The thermal control system was able to tolerate attitudes with the Sun offset from the +Y axis beyond the normal 15° constraint. The first IMU calibration turn sequence took advantage of these conditions by commanding turns of nearly 30° about the body axes, resulting in relatively large changes in the Sun azimuth and elevation angles. The spacecraft was initially commanded to an attitude that placed the Sun at an azimuth of 170° and elevation of 10°, where the azimuth and elevation conventions are defined in Figure 1. Then the series of turns shown in Table 4 were executed. This turn sequence was repeated from the same starting attitude three times: once with only star tracker 1 on, next with only star tracker 2 on and finally with both star trackers on. The star tracker alternation was necessary because the thermal control system cannot maintain sufficiently low temperatures for both trackers operating simultaneously when the sunshade is not facing the Sun. The test was designed to minimize the period when both trackers were powered on while still providing some data for tracker alignment computation. The three repeats of the turn sequence were, in turn, executed twice: first with the IMU controlled by PPSMA and then with the IMU controlled by PPSMB. The Sun azimuth varied between 165° and 195° while the elevation varied between +15° and -15° during each set of body axis rotations. The rotation rate was just under 0.3°/s.

A new turn series was designed for the October 2005 and all future IMU calibrations that can be executed within the normal Sun keep-in (SKI) bounds for the shade-to-Sun orientation. These bounds constrain the Sun direction to stay within approximately +15° and -15° in both azimuth and elevation (a

region centered on the  $-Y$  axis or sunshade normal). The turn sequence for the IMU calibration was designed to keep Sun azimuth and elevation angles within  $\pm 12^\circ$ , comfortably within the true SKI limits. The starting attitude places the Sun along the  $-Y$  axis ( $0^\circ$  azimuth and elevation). Then rotations are

**Table 4**  
**2004 IMU CALIBRATION TURN SEQUENCE**

| Axis | Turn Size  |
|------|------------|
| +Y   | $90^\circ$ |
| +X   | $28^\circ$ |
| -X   | $28^\circ$ |
| -Y   | $90^\circ$ |
| -Z   | $26^\circ$ |
| +Z   | $26^\circ$ |

performed about the X, Y, and Z axes and about an axis with non-zero X, Y and Z components as shown in Table 5. The rotation rate was set at just under  $0.3^\circ/\text{s}$ , as had been used for the 2004 calibration turns. Star tracker 2 was turned on and generated attitude solutions along with star tracker 1 throughout the entire turn sequence. Both star trackers could be operated simultaneously without a significant temperature increase since the spacecraft had its sunshade facing the Sun. For operational convenience, the two executions of this turn sequence with the IMU controlled from different PPSMs were performed on separate days with the PPSMB session done first.

**Table 5**  
**2005 SKI-CONSTRAINED IMU CALIBRATION TURN SEQUENCE**

| Axis                              | Turn Size                                | Comment   |
|-----------------------------------|--|---|
| +X, -X, +X, -X                    | $10^\circ, 20^\circ, 20^\circ, 10^\circ$ | Sun elevation varies between + and $-12^\circ$ (with overshoot) |
| +Z, -Z, +Z, -Z                    | $10^\circ, 20^\circ, 20^\circ, 10^\circ$ | Sun azimuth varies between + and $-12^\circ$ (with overshoot)   |
| +Y, -Y, +Y                        | $360^\circ, 720^\circ, 360^\circ$        | Sun stays at AZ and EL near $0^\circ$                           |
| +X, +Z                            | $10^\circ, 10^\circ$                     | Sun ends at azimuth and elevation of $-10^\circ$                |
| AZ, EL= $-10^\circ$ (3-axis turn) | $360^\circ$                              | Sun stays at azimuth and elevation of $-10^\circ$               |
| -Z, -X                            | $10^\circ, 10^\circ$                     | Sun ends at azimuth and elevation of $0^\circ$                  |


### IMU/Star Tracker Time Synchronization

It is necessary to properly synchronize the IMU data with the measurements obtained by the star trackers. This synchronization is necessary to compensate for the varying internal hardware delays in the IMU and star trackers, as well as the previously discussed imprecision on the IMU time tag. Additionally, because the IMU/bus clock drift rate is nearly constant over the calibration maneuver, it is possible to create more accurate time tags by accounting for this drift rate on the ground, following similar processing as used in the new flight software.

Once the IMU time tags are corrected for clock drift, the IMU data must be aligned with the tracker data to minimize the calibration filter residuals and improve the overall calibration. The procedure is simply to calculate the rates from the IMU data and the rates from each of the ASTR measurements, and shift the IMU rates in time until they are aligned with those computed from the star tracker observations. Because the angular acceleration capabilities of MESSENGER are limited, the data latency between star tracker and gyro data is difficult to discern under normal spacecraft slews. However, structural vibrations excited during the calibration maneuver proved useful for visualizing the time correction applied to the AST time tags. This can be seen in Figure 16, which depicts several seconds of the X-axis spacecraft rate for clarity. The other axes show similar characteristics for the duration of the maneuver, but this time sample's dominant structural activity is most readily apparent on the X-axis.

The value of the time correction applied to the star tracker time tag computed for the calibration maneuver was  $-0.1$  s for both October 27 and 28. This shifts the red curve in Figure 16 to the left by  $0.1$  s,

and aligns the data from the two sensors. As expected, the values for both trackers on both days showed nearly identical shifts relative to the IMU data. This is the case because, unlike the IMU, the star trackers are synchronized with the spacecraft bus schedule and produce measurements at nearly identical times, and their delays are determined by fixed delays internal to the hardware devices and the spacecraft flight software version itself (e.g., processing delays, bus delays). These values computed for the delays were further verified as small perturbations from these numbers used in the calibration filter resulted in larger filter residuals.



**Figure 16. Relative Time Shift in IMU/Star Tracker Data During Calibration on October 27, 2005**

#### Star Tracker Alignment

The calibration filter does not estimate an alignment for AST1 as the AST1 frame is used as the fiducial reference. It does estimate the AST2 misalignment, which is an error on the ground-supplied AST1-to-AST2 nominal transformation. Both trackers are synchronized with the spacecraft bus (and thus one another), so the AST measurements are assumed to occur at the same time (which should be accurate to within several  $\mu\text{s}$ ). This allows the calibration filter misalignment estimate to be compared with the more direct computation of the average difference between synchronous tracker measurements. Said another way, the relative tracker alignment can easily be obtained from the mean quaternion error between synchronous tracker observations. This quaternion error is defined as the vector part of the quaternion  $q_e$ :

$$q_e = q_{\text{err}}(q_{\text{ast1}}, q_{\text{ast2}}) \quad (2)$$

where  $q_{\text{err}}$  is the quaternion function defined by the quaternion product of the conjugate of the star tracker 1 quaternion,  $q_{\text{ast1}}$ , times the star tracker 2 quaternion,  $q_{\text{ast2}}$ . A plot of the misalignment between both trackers is shown in Figure 17 obtained from this  $q_{\text{err}}$  function. Note that for convenience, the star tracker quaternions are defined to be the nominal inertial-to-body quaternions and include the nominal optical alignments relative to the spacecraft body frame.

The calibration maneuver was conducted with both trackers on and enabled, which is an unusual thermal situation. The gradual slopes in Figure 17 appear to be due to thermal conditions at the star trackers during this maneuver. The mean values from Figure 17 compare favorably with the estimates from the calibration filter, from the results of prior calibrations shown in Table 6. Similar results were obtained for the calibration maneuver on October 28. Table 6 encapsulates the different calibration results, which rotates the body coordinates plotted in Figure 15 into star tracker 2 coordinates (the output format of the ground calibration filter). Note that the calibrations done in 2004 show numbers that differ from the 2005 calibrations by nearly 40 arcsec, (root squared sum, or RSS). This is most likely due to different thermal conditions on the spacecraft, as the respective calibration maneuvers were done in entirely different attitudes and Sun-relative inertial positions.




Figure 17. Relative Misalignment Between AST2 and AST1 Telemetry in Spacecraft Body Coordinates on October 27, 2005

Table 6  
STAR TRACKER RELATIVE ALIGNMENT CALIBRATION IN AST2 COORDINATES

| Date               | Method    | X-Axis Alignment Error (arcsec) | Y-Axis Alignment Error (arcsec) | Z-Axis Alignment Error (arcsec) |
|--------------------|-----------|---------------------------------|---------------------------------|---------------------------------|
| September 13, 2004 | $q_{err}$ | 1308.4                          | -608.9                          | -426.0                          |
| September 13, 2004 | KF        | 1309.0                          | -608.7                          | -426.6                          |
| October 6, 2004    | $q_{err}$ | 1318.1                          | -613.4                          | -432.9                          |
| October 6, 2004    | KF        | 1318.1                          | -613.4                          | -431.6                          |
| October 27, 2005   | $q_{err}$ | 1333.4                          | -640.9                          | -428.3                          |
| October 27, 2005   | KF        | 1334.3                          | -641.0                          | -428.1                          |
| October 28, 2005   | $q_{err}$ | 1332.1                          | -640.2                          | -426.7                          |
| October 28, 2005   | KF        | 1333.0                          | -640.3                          | -425.9                          |

### Gyro Alignments

The alignment estimate produced by the calibration filter includes orthogonal and non-orthogonal alignments internal to the IMU and the external orthogonal misalignment relative to AST1. These alignments are estimated in the calibration filter as errors from the nominal alignment matrix. The estimated values from the calibration filter depend on the initial nominal alignment, although the final transformation matrix is independent of the initial nominal alignment matrix. For consistency with prior calibrations, the pre-launch optical alignments are used for the initial IMU alignment, so current calibration filter results can be compared with prior calibrations directly. Note that these preliminary alignments are used in the ground calibration filter, not onboard for the maneuver itself. Onboard alignment transformations are handled in the G&C flight software, and as such, do not impact the telemetry used for the ground calibration.

A comparison of the alignment estimates for each gyro with those provided in (1) show the current alignment estimates are consistent with prior data sets and between PPSMs as expected. The misalignments are provided in terms of orthogonal angular offsets ( $du$  and  $dv$ ) from the nominal optically aligned axes. These alignments are mutually orthogonal and are in a plane normal to the nominal optical alignment axis for a particular gyro<sup>6</sup>. Table 7 contains the alignment data for each gyro from each calibration maneuver. Although the power supply electronics (PPSMA or PPSMB) should not impact the alignments, they may slightly alter the calibrated values due to their potential impact on scale factor and

gyro bias. The values in Table 5 show standard deviations between 15 and 40 arcsec, for each respective error coordinate, and a maximum RSS standard deviation of 46 arcsec. These errors are small, and may be driven by the thermal conditions at the IMU, as well as small time tag imprecision on the star tracker and IMU measurements.

**Table 7**  
**IMU ALIGNMENT CALIBRATION RESULTS IN GYRO ERROR COORDINATES**

| Date               | Gyro A (arcsec) |      | Gyro B (arcsec) |      | Gyro C (arcsec) |      | Gyro D (arcsec) |      |
|--------------------|-----------------|------|-----------------|------|-----------------|------|-----------------|------|
|                    | du              | dv   | du              | dv   | du              | dv   | du              | dv   |
| September 13, 2004 | 480             | -272 | 438             | 308  | 839             | 266  | 839             | -302 |
| October 6, 2004    | 309             | 793  | -278            | 453  | -313            | 463  | 225             | 849  |
| October 6, 2004    | 432             | 330  | 817             | 329  | 790             | -223 | 497             | -306 |
| October 27, 2005   | -303            | 473  | -298            | 465  | 301             | 837  | 330             | 819  |
| October 28, 2005   | 812             | 328  | 796             | -260 | 511             | -252 | 473             | 318  |

**Gyro Bias**

The gyro bias estimates from the calibrations on Build 4.9 show minor differences from the results obtained with the calibrations on Build 4.7. All calibrations show a higher gyro bias on Gyro A (about 0.4 deg/hr on PPSMB compared with 0.55 deg/hr on PPSMA) than the other three gyros. This is a known characteristic of this particular gyro that was seen in vendor testing, and is within the vendor specification. The remaining three gyros show bias of +/-0.05 deg/hr, which is consistent with the manufacturer specifications. The numbers presented in Table 8 are also consistent with the onboard Kalman filter estimates of gyro bias.

**Table 8**  
**GYRO BIAS ESTIMATES FROM CALIBRATION FILTER**

| Date               | Gyro A (deg/hr) | Gyro B (deg/hr) | Gyro C (deg/hr) | Gyro D (deg/hr) |
|--------------------|-----------------|-----------------|-----------------|-----------------|
| September 13, 2004 | 0.554           | -0.043          | 0.04            | -0.033          |
| October 6, 2004    | 0.562           | -0.044          | 0.048           | -0.037          |
| October 6, 2004    | 0.371           | -0.027          | 0.044           | -0.034          |
| October 27, 2005   | 0.358           | -0.064          | 0.091           | -0.054          |
| October 28, 2005   | 0.533           | -0.071          | 0.061           | -0.052          |

**Gyro Scale Factor**

The convergence of the symmetric scale factor estimate in the calibration filter is very sensitive to the relative IMU/ASTR synchronization. Any uncompensated time shift in either the star tracker or IMU data can cause significant changes in the resulting scale factor errors. The estimates obtained from the calibrations on October 27 and 28 are significantly different from prior calibrations, and are larger than the beginning-of-life specification of +/-200 ppm 1- $\sigma$ . The differences from prior calibrations could be the result of an improved IMU time tag for the current calibration. All these results are provided in Table 9. Another point is that the prior version of flight software reset the IMU clock time tags at 1 Hz. By doing so, the time between consecutive gyro measurements was difficult to obtain in some cases, thereby “muddl(ing) the effect on scale factor<sup>6</sup>.” These values will not be updated on the basis of the calibration, since they could have a significant impact on the attitude estimation, and the calibration values are not obtained with a high level of confidence. This policy was also adopted on the TIMED mission, as the scale factor from calibration filter runs for that mission exhibited similar behavior<sup>7</sup>. Future calibrations may show improved consistency, since the IMU time tags should be more accurate for future maneuvers.

The asymmetric scale factor is theoretically zero, but it was a part of the Kalman filter formulation<sup>6</sup> and thus was estimated as a part of this investigation. Prior calibrations estimated these values

**Table 9**  
**SYMMETRIC SCALE FACTOR ESTIMATES FROM CALIBRATION FILTER**

| Date               | Gyro A (ppm) | Gyro B (ppm) | Gyro C (ppm) | Gyro D (ppm) |
|--------------------|--------------|--------------|--------------|--------------|
| September 13, 2004 | 209.11       | 357.1        | 480.35       | 303.6        |
| October 6, 2004    | 208.68       | 381.74       | 467.09       | 312.78       |
| October 6, 2004    | 282.63       | 532.07       | 554.07       | 346.98       |
| October 27, 2005   | 85.26        | 291.55       | 345.71       | 669.83       |
| October 28, 2005   | 318.61       | 28.07        | 572.9        | 353.19       |

to be smaller than +/-50 ppm, and although the estimates from the current calibration differed from prior estimates, they remained below 50 ppm. This value is considered negligible, although it is above the specification of +/-7 ppm. As with the other gyro parameters, attempts to estimate the asymmetric scale factor showed great variability with small shifts in the gyro/star tracker time tags, as well as increased variability in the symmetric scale factor estimates. These numbers are not obtained with a high level of confidence, but there is currently no provision to use an asymmetric scale factor in the G&C flight software and the numbers are estimated for informational purposes only.

### CONCLUSIONS

Changes were made to the MESSENGER G&C flight software to handle the drift between the IMU and spacecraft clocks. This drift caused intermittent periods of missed and repeated IMU messages and is handled in the new software by a redesign of the IMU time tag computation. Other changes were made to the attitude-estimation Kalman filter to simplify the processing and improve performance, including removal of rate and misalignment states from the filter and the addition of a "star tracker only" mode of operation. These software changes were loaded to the spacecraft in October 2005 and have improved attitude estimation and control performance.

Flight calibration of gyro and star tracker parameters have been conducted at several different times, under varying attitude constraints. Despite these constraints, observability of the sensor parameters has been maintained, and many of the parameters show consistent convergence to the same values over the duration of the mission. Future upload of these estimated parameters are expected to improve attitude estimation performance further.

### ACKNOWLEDGEMENT

The work described in this paper was performed at The Johns Hopkins University Applied Physics Laboratory, under contract NAS5-97271 with the National Aeronautics and Space Administration Discovery Program Office.

### REFERENCES

1. S. C. Solomon et al., "The MESSENGER Mission to Mercury: Scientific Objectives and Implementation," *Planetary and Space Science*, Vol 46, Issues 14-15, pp. 1445-1465, December 2001.
2. R. E. Gold et al., "The MESSENGER Mission to Mercury: Scientific Payload," *Planetary and Space Science*, Vol 46, Issues 14-15, pp 1467-1479, December 2001.
3. A. G. Santo et al., "The MESSENGER Mission to Mercury: Spacecraft and Mission Design," *Planetary and Space Science*, Vol 46, Issues 14-15, pp. 1481-1500, December 2001.
4. B. Wie et al., "Rapid Multitarget Acquisition and Pointing Control of Agile Spacecraft," *Journal of Guidance, Control, and Dynamics*, Vol. 25, No. 1, pp. 96-104, January-February 2002.
5. R. M. Vaughan et al., "MESSENGER Guidance and Control System Performance During Initial Operations," AAS Guidance and Control Conference, AAS-05083, Breckenridge, CO, February 4-8 2005.
6. M. E. Pittelkau, "Calibration and Attitude Determination with Redundant Inertial Measurements Units," *AIAA Journal of Guidance, Control, and Dynamics*, Vol 28, Number 4, pp. 743-752, July-August 2005.
7. M. E. Pittelkau and W. F. Dellinger, "Attitude Sensor Alignment and Calibration for the TIMED Spacecraft," Paper AAS 03-153, 13th AAS/AIAA Space Flight Mechanics Conference, Ponce, Puerto Rico, February 9-13, 2003.