

Spacecraft Clock Maintenance for MESSENGER Operations

James F. Hudson¹ and Edwin J. Colwell²

The Johns Hopkins University Applied Physics Laboratory, Laurel, MD 20723

SPACECRAFT exposed to dynamic fluctuations in temperature, velocity, and gravity present significant challenges for a high-precision timekeeping system. En route to Mercury, the MESSENGER spacecraft has endured a total of six planetary flybys, considerable variation in proximity to the Sun, and large velocity gradients. The probe completed 15 solar revolutions during the heliocentric cruise phase, with a minimum perihelion distance of 0.302 AU and a maximum aphelion distance of 1.07 AU. Its environment changed once again when the spacecraft was captured in the gravity well of Mercury and the planetary orbital phase of the mission began. The accuracy goal of the timekeeping system is to maintain ground knowledge of the spacecraft clock to within 100 ms. This paper summarizes MESSENGER's complex timekeeping system and the steps that have been taken to address the mission's unique challenges.

I. Introduction

The MErcury, Surface, Space ENvironment, GEOchemistry, and Ranging (MESSENGER) spacecraft, now in orbit about Mercury, uses two types of onboard oscillators: an oven-controlled crystal oscillator (OCXO) for precision timekeeping as well as an internal, coarse oscillator. Nominally, both oscillators run simultaneously, giving the mission operations team the option to select either. During certain operations, such as maneuvers that require the firing of thrusters, it is preferable to select the coarse oscillator. This precaution decreases the probability that spacecraft vibrations will trigger an anomaly with the OCXO that could potentially result in a reboot of the main processor (MP). For the majority of the mission, however, it is desirable to select the OCXO. The selected oscillator drives the clock that increments the mission elapsed time (MET), which provides a common frame of reference on which to base all events for the life of the mission. When commands are sequenced, they are executed at an assigned MET, and all information coming from the spacecraft is time-stamped with this MET. The unit of measurement for the MET is seconds, counted in terrestrial dynamical time (TDT). The formula for calculating TDT is expressed numerically as $TDT = UTC + 32.184 \text{ s} + (\text{leap seconds})$.¹

For the ground-system component of MESSENGER's timekeeping system, the mission operations team generates spacecraft clock data files called spacecraft clock (SCLK) kernels. One SCLK kernel is generated from each Deep Space Network support during which contact is established with the spacecraft. Telemetry and predictive ephemeris data are used to interpolate the clock drift rate and to create a linear formula to convert MET to TDT. This linear formula consists of a MET/TDT intercept and the drift rate of the oscillator (TDTRATE). These parameters are loaded to the spacecraft to support onboard calculations of TDT from MET. The spacecraft uses TDT knowledge to derive position from the onboard ephemerides. The guidance-and-control subsystem utilizes these position states to orient the spacecraft relative to Earth, the Sun, and other heavenly bodies. The SCLK kernels are also used to support ground system operations, converting time tags of telemetry data, supporting command-sequencing operations, and evaluating the performance of onboard TDT knowledge.

While analyzing timekeeping data through the planetary flybys, the mission operations team observed unpredicted trends in the clock drift-rate data. These trends showed a divergence from the nominally expected drift rate. The unusual drift-rate behavior caused a greater-than-expected error in the onboard TDT knowledge, although the calculated error was still within the specified requirements. Analysis by the MESSENGER team determined that

¹ MESSENGER guidance and control analyst, SEO, 13-N388.

² MESSENGER timekeeping analyst, SEO, 13-N384.

relativistic effects were responsible for the additional error. The changes that the spacecraft experienced in the gravity field and velocity gradients caused the oscillator to run faster or slower than the Earth reference. By using detailed knowledge of the spacecraft trajectory provided by the navigation team, the mission operations team generated new projections for the TDTRATE that factored in these relativistic effects. These new projections improved the linear modeling of the OCXO behavior onboard the spacecraft.

II. Mission Overview

MESSENGER launched on August 3, 2004. The mission is ambitious and characterized by substantial challenges. It has joined the elite group of spacecraft that have successfully achieved orbit around another world, and it holds the distinction of the first to circle the planet closest to the Sun. The data collected will provide the science community with much needed insights regarding not only Mercury but all of the terrestrial planets. Specifically, the spacecraft's suite of instruments will explore Mercury's density, planetary core, geologic history, exosphere, and magnetic field. MESSENGER has recently gathered evidence that supports earlier observations provided by the Arecibo radio telescope and Goldstone's Very Large Array regarding the presence of ice near Mercury's poles. Finally, MESSENGER's vivid images will complete the montage that was started by Mariner 10 some thirty years ago. It will also expose hidden regions of the planet that were heretofore unobserved.

There are several features of the observatory and the mission that make MESSENGER unique. The spacecraft is distinguished in appearance because of the large sunshade that dominates its profile. This safeguard is requisite for a spacecraft operating in the intense heat present at such close proximity to the Sun. A number of autonomous fail-safes ensure that the spacecraft attitude maintains a position where its sensitive electronics are protected by the sunshade. MESSENGER is the first deep-space satellite to use a fixed phased-array antenna² in addition to a suite of secondary antennae of various gains. The eschewal of deployable or gimballed antennae reduces the potential for communication failure. The circuitous expedition en route to Mercury was designed to take advantage of a number of gravity assists in the form of planetary flybys. This method of lowering the spacecraft velocity greatly reduced the amount of propellant required to complete the mission. The Earth flyby along with the second Venus flyby also afforded the opportunity to calibrate MESSENGER's instruments, whereas the three Mercury flybys yielded useful science data. MESSENGER science data collection is extremely dependent on precise timekeeping. Specifically, the Mercury Laser Altimeter has a stringent requirement of ± 100 ms in operational knowledge of TDT onboard the spacecraft. This requirement has driven the development of the mission's complex, maintenance-intense timekeeping system.

III. Methodology

Accurate timekeeping is not just a vital component of all space missions; in fact, it has been an important part of exploration for hundreds of years. In 1714, when the British parliament passed the Longitude Act, a prodigious monetary reward was offered to anyone who could develop a precise means of determining longitude. The inefficacy of prior navigation methods severely impacted travel and trade. It was a tremendous engineering challenge, and the solution lay in timekeeping.³

MESSENGER's timekeeping process works with files whose format is defined by the Navigation and Ancillary Information Facility information system called Spacecraft Planet Instrument C-matrix Events (SPICE). SPICE is widely used by a number of space missions for its utility with solar system geometry. SPICE consists of a series of utility library routines and a number of data files called kernels. These kernels are essential to SPICE functions and contain mission-specific information. The SCLK is of primary importance and is made up of a MET, a time string of TDT that corresponds to that MET, and a clock drift rate. An updated SCLK is calculated after each coherent contact with the spacecraft. (A coherent contact is defined as any interaction with the spacecraft that includes a carrier uplink being received by the spacecraft at the same time that telemetry downlink is being received on the ground.)

Another kernel that is vital to mission operations is the Spacecraft and Planetary Ephemeris Kernel (SPK). These ephemerides describe the spacecraft's position as a function of time. As such, they are commonly referred to as a trajectory file. There is another SPK for planetary position information. A leap-second kernel is also maintained to accommodate any leap-second adjustments that occur throughout the mission. Leap second knowledge is needed to convert from normal operational time, UTC, to TDT, which is used as a basis for ephemeris time. Finally, there is a meta-kernel that acts as an index of the names and locations of the other kernels that are used as inputs to the timekeeping process. The meta-kernel is used by the SPICE library to load the appropriate kernels autonomously.

The aforementioned list of kernels is by no means exhaustive, although it does include several of those kernels that are most crucial to timekeeping operations on MESSENGER.

The computation of the spacecraft's one-way light time (OWLT) is also performed by using the SPICE library and the timekeeping ground system software. The OWLT is the amount of time required for a signal to reach the spacecraft from the ground. A telemetry packet that contains timing information is transmitted from the spacecraft numerous times throughout each real-time support. The timekeeping software analyzes the calculated OWLT and determines the clock drift rate that best fits this data collection.

A streamlined process has been developed by the mission operations team to facilitate the timekeeping process. A series of PERL scripts work together to calculate updated SCLK kernels based on the latest SPK. For orbital operations, when the precise oscillator is in use, a new series of kernels is generated daily. A cron job initiates the process at a regular interval, after the timekeeping data have been telemetered and received in the data archive. When the coarse oscillator is in use, the cron job is suspended and the PERL scripts are run manually, once per week.

Periodically, it is necessary to update the onboard SCLK. Whereas the creation of kernels addresses the ground-system portion of timekeeping, the spacecraft's onboard knowledge of ground time is established through loading the onboard SCLK in the form of TDTMET parameter blocks. These parameter blocks are created directly from the ground SCLKs. A ground software process makes spacecraft TDTMET command templates from the latest ground SCLK. These templates are then converted into real-time command procedures, which are executed on the ground and sent to the spacecraft. The standard process for updating parameters on the spacecraft involves loading them via command, verifying them via dumps, and then committing them to use. Once SCLKs are committed on the spacecraft, their use can be verified in normal spacecraft telemetry.

The mission operations team performs a number of quality control measures to validate and verify the accuracy of the generated SCLKs. The most basic assessment of the kernels is to trend the TDTMET drift rate and clock error on a series of plots. This procedure allows for at-a-glance evaluation of any deviation of the trend. It also gives some knowledge of when the error rate edges toward a violation of the acceptable limits outlined by mission constraints.

Additionally, once per week, a test is conducted with the spacecraft called a latch MET test. When these tests occur, a series of commands is sent along two pathways simultaneously. The first path is directly to the spacecraft. The second is routed through a ground-system receiver and converted to a Ka-band frequency. Both the ground receiver and the spacecraft will time stamp the receipt of these commands, and a comparison of the two demonstrates the offset between the two systems. Finally, because SPKs are calculated by using predicted ephemeris data, a more accurate SPK is generated after the fact, using high-fidelity position knowledge collected from spacecraft ranging data. Contrasting these after-the-fact kernels with the predicted SPK yields useful insights regarding the accuracy of the SPKs.

IV. Transition Between Oscillators

As of the start of the mission cruise phase, the OCXO was not normally used for timekeeping. Exceptions were limited intervals during which science data were collected. The OCXO was first used 24 days after launch on August 27, 2004, during postlaunch checkout. The OCXO was again selected to drive the clock in support of science collection during the Earth flyby, the second Venus flyby, and the three Mercury flybys. The operations team selected the coarse oscillator during the execution of critical trajectory-correction maneuvers such as deep-space maneuvers (DSMs), Mercury orbit insertion (MOI), and orbital correction maneuvers (OCMs). The use of the coarse oscillator during these events mitigated against any MP reset risks that can come when operating under the OCXO, which is external to the MP. The coarse oscillator was also used during the three flight software update events that occurred during the cruise phase because of the MP resets involved.

By 2009, the ground process was sufficiently automated to handle the added processing of the OCXO operations without much added impact to the regular workload of the mission operations team. At that time, the mission operations team decided to operate with the OCXO selected as default for the rest of the cruise phase, which ended with the start of orbit operations in March 2011. It was still necessary to switch to the coarse oscillator for the events mentioned above. There was one unplanned switch of the oscillator during an anomaly in September 2009 that resulted in an MP reset. This reset caused the MP to stop using the OCXO and switch to the coarse oscillator. The mission operations team switched back to the OCXO during the reset recovery operations a few days later. The OCXO is currently in use full-time during orbital operations, but the mission operations team switches over to the coarse oscillator for approximately 1 day around OCMs. The following table summarizes the oscillator usage so far during the mission.

Table 1. Oscillator Usage during the MESSENGER Mission to Date.

DATE	OSCILLATOR STATE	COMMENT
August 4, 2004	COARSE	Default state at launch
August 27, 2004	OCXO	Support early cruise testing
September 15, 2004	COARSE	End of early cruise testing
April 26, 2005	OCXO	Support Earth science and Earth flyby
September 2, 2005	COARSE	Select coarse oscillator for DSM 1
March 29, 2007	OCXO	Support Venus flyby 2
June 25, 2007	COARSE	Select coarse oscillator for DSM 2
December 21, 2007	OCXO	Support Mercury flyby 1
February 18, 2008	COARSE	Select coarse oscillator for DSM 3
August 20, 2008	OCXO	Support Mercury flyby 2
December 2, 2008	COARSE	Select coarse oscillator for DSM 4
January 26, 2009	OCXO	Support cruise operations
July 9, 2009	COARSE	Select coarse oscillator for Flight S/W Load
July 16, 2009	OCXO	Support cruise operations
September 2, 2009	COARSE	Result from anomalous processor reset
September 8, 2009	OCXO	Support Mercury flyby 3
October 12, 2009	COARSE	Select coarse oscillator for DSM 5
December 7, 2009	OCXO	Support cruise operations
March 4, 2011	COARSE	Select coarse oscillator for MOI
March 18, 2011	OCXO	Support orbit operations
June 14, 2011	COARSE	Select coarse oscillator for OCM 1
June 15, 2011	OCXO	Support orbit operations
July 25, 2011	COARSE	Select coarse oscillator for OCM 2
July 26, 2011	OCXO	Support orbit operations

Transitioning between the oscillators causes interruptions in the normal automated process that creates the SCLKs described earlier. Manual intervention is required to operate the timekeeping system through these oscillator transitions. Typically, the automated process is temporarily halted at the same time that the coarse oscillator is selected via real-time command. This switch occurs 1–2 days before the event that necessitates the transition. Afterward, the OCXO is selected to drive the onboard clock once again, and the procedure to initiate the automated processing of OCXO data can begin. The first step of this procedure is to create a clock kernel by using the last section of data retrieved from the OCXO before the transition. Another clock kernel is made for the time on the coarse oscillator. The process for constructing the kernels is the same, except in these cases the user specifies the particular time period to be used during the processing. In the automated process, any data from the entire previous day are used. The timekeeping software cannot create an SCLK using a combination of data from both oscillators. It is imperative that the time spans are set in a way to isolate the data coming from the OCXO and coarse oscillators. Because the OCXO experiences very little drift with respect to the coarse oscillator, the SCLK created with the coarse oscillator data will detail how much the clock drifted while not on the OCXO. Typically this drift can be measured in terms of tenths of seconds. The manual creation of the kernels will continue until the software can reliably calculate the drift rate of the OCXO. The first few kernels that are created for the OCXO after transition require special processing regarding the calculation of the clock rate. The software averages data up to 3 days before the start of the span; therefore, the coarse oscillator data present problems for determining an accurate clock rate during transitions. The OCXO rate will typically be billionths of a second off per second, whereas the coarse oscillator will be millionths of a second off per second. This large disparity makes it impossible to derive a meaningful rate from an analysis of both types of data in the same span. The software has a control to work around this problem. The kernel creation process can be run in a fixed-rate mode whereby the user specifies an assumed rate, and the data are all processed with the clock rate held constant. This step allows all of the historical clock data

to be placed into the active archive of the software process and also provides a precise estimate of the rate because the OCXO operates very predictably. The rate of “1” is typically used to initiate the process following the OCXO reselection. Each time the post-transition kernels are manually created, the calculated fit error of each resulting run is monitored. When the error falls below a certain threshold, the operator can be assured that older coarse oscillator data are no longer being used in the kernel processing. At this point, the automated process can be enabled again, and accurate calculate of the clock rates of the OCXO can resume.

V. Performance During Flybys

The first time that the OCXO supported the collection of science data was during the Earth flyby, which occurred on August 2, 2005. The OCXO was selected almost 100 days before the flyby. This interval of OCXO operations gave the mission operations team ample time to ready the timekeeping system for the flyby. The OCXO was selected on April 26, 2005, and the drift rate for the clock was calculated on May 2, 2005. When the OCXO was selected, a drift rate that had been determined before launch was loaded to the spacecraft. This onboard drift rate proved to be adequate for operations because it closely matched the drift rate trends leading up to the Earth flyby. As mentioned above, the onboard OCXO drift data, or the “onboard clock kernel,” are used to calculate TDT from the spacecraft clock MET. This TDT information is then used to interpret onboard ephemerides, which in turn tell the flight software how to locate the spacecraft and potential planetary targets. The error in TDT calculation comes from the difference between the actual OCXO drift rate and the drift rate stored in the onboard clock kernel. Large differences between the onboard drift parameter and the actual drift will create large growth in the TDT calculation error. Figure 1 shows the TDT calculation error during the period leading up to the Earth flyby in 2005.




Figure 1. TDT calculation error during the period leading up to the Earth flyby in 2005.

As may be seen in Figure 1, the onboard TDT error remains well within the 100 ms in-flight requirement. The mission operations team was prepared to update the onboard clock kernel a few days before the flyby, but no update was necessary because of the solid performance of the OCXO.

Approximately 30 days after the Earth flyby activity, the coarse oscillator was selected, which removed the OCXO from control of the clock. The OCXO was selected again on March 29, 2007, and put into use to support the

second Venus flyby activity. The first drift-rate calculation for the OCXO occurred 4 days later, which allowed the older coarse clock data to work their way out of the timekeeping data process. A clock kernel was then made daily from that day forward, in contrast to the weekly generation that occurred while on the coarse oscillator. The clock kernel that was generated on April 12, 2007, was then loaded to the spacecraft to more accurately calculate onboard TDT from MET. Two more kernels were loaded to the spacecraft during the time period leading up to the second Venus flyby, which occurred on June 5, 2007. The updates were not needed to maintain TDT calculation accuracy; instead, they were uplinked to practice the loading process leading up to the flyby. This exercise improved the operational readiness for the upcoming Mercury flybys. The error behavior leading up to the flyby is detailed in Figure 2. The error did grow slightly over time, and updates to the onboard kernel brought the growing error back down to near-zero levels.




Figure 2, TDT calculation error during the period leading up to the second Venus flyby in 2007.

The growth in the error was somewhat larger than previously observed, so the rate behavior of the OCXO leading up to the second Venus flyby was further scrutinized. Once several rates were calculated, a new trend was observed. The sinusoidal-shaped trend could not be explained by any thermal changes or other hardware behavior. Eventually, the lead timekeeping engineer found that the trend matched relativistic effects caused by large changes in the spacecraft velocity and gravity gradient. These gradients had not been as pronounced earlier in the mission because of more benign heliocentric orbit geometry. A plot of the drift rate behavior is shown in Figure 3.




Figure 3, OCXO drift rate during the period leading up to the second Venus flyby in 2007.

The blue dots in Figure 3 represent solutions of the OCXO drift rate taken from the daily-generated clock kernels. The sinusoidal shape is apparent in the curve before the second Venus flyby. The purple line is the projected behavior of the clock drift rate accounting solely for the relativistic effects of velocity and gravity changes. The clock drift rate, TDTRATE, is the number of TDT seconds per second that is ticked off by the OCXO. It is simply a measure of how well the clock can account for a true second. If the rate is >1 , the OCXO is running slow, i.e., there are more TDT seconds for every MET second. If the rate is <1 , the OCXO is running fast, i.e., there are fewer TDT seconds per MET second. The plot shows a relative peak at about April 20. This peak indicates that the OCXO is slowing down during this time. The peak also corresponds to a relative solar distance minimum in the spacecraft's heliocentric orbit. At this solar distance minimum the spacecraft is at a relative maximum in velocity magnitude with respect to the Sun. This follows from relativity theory, by which time advances more slowly as velocity increases. The purple line on the plot is derived from Einstein's equations using the spacecraft trajectory file as an input and taking into account the spacecraft's velocity and its position in the gravity fields of the Sun and Venus. These calculations result in a projection of how the clock rate should change, accounting for its motion and relativistic effects. The measured clock rates signified by the blue line match up to the purple line projection very well, as the plot shows. This result demonstrates that relativistic effects, if not included, can introduce a substantial source of error to the system. The plot also shows how the system can respond to this error. The red line represents the rate onboard the spacecraft that is used to calculate TDT. That parameter can be manipulated to minimize the error growth. The plot shows how the parameterized rate was updated to follow the actual changing rate of the clock. In addition to the drift rate, the TDT for a particular MET in the most recent past is also loaded to the spacecraft, and these values effectively reset the TDT calculation error to zero. Figure 3 shows how the accuracy is increased each time a new kernel is loaded onboard.

The next major events after the second Venus flyby were the three Mercury flybys, which occurred on January 14, 2008, October 6, 2008, and September 29, 2009. The first two Mercury flybys performed very similarly to the second Venus flyby event with respect to the timekeeping system. The variation of the clock rate was slightly more pronounced for the later events because of more severe relativistic effects. The changing and ever-tightening geometry of the spacecraft trajectory accounted for the increased magnitude of the relativistic effects. The mission operations team considered taking the relativistic projections and using them to "lead" the expected error in order to

minimize it, but a more conservative approach was chosen for the first two Mercury flybys. The strategy remained one of updating the onboard clock kernel with the most recent ground kernel in order to reset the growing error back to zero and match the changing clock rate as closely as possible. Figure 4 shows the error performance for the second Mercury flyby.

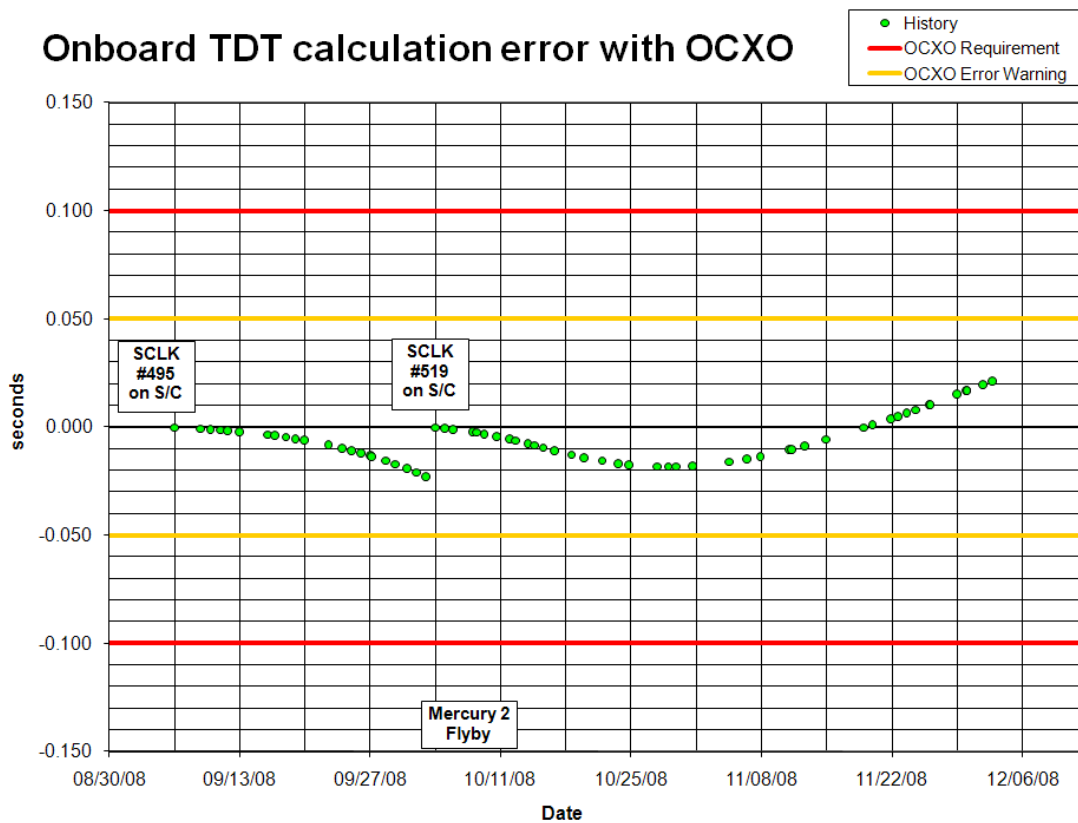


Figure 4, TDT calculation error during the period leading up to the second Mercury flyby in 2008.

Figure 4 shows that the error starts off near zero but slowly grows negative because the clock rate is moving away from the rate used by the spacecraft to calculate TDT. The onboard kernel was updated on October 3, 3 days before the second Mercury flyby, and the plot shows how the new kernel minimized the error and slowed its growth by more closely matching the actual clock rate. Only one onboard update was needed to keep the TDT calculation error well within requirements. The performance during the first Mercury flyby was nearly identical in terms of error management.

The mission operations team decided to be a little more aggressive in managing the onboard kernel for the third Mercury flyby. A relativistic projection was used to create a custom-made kernel that, in turn, created an error profile that was optimized for the third Mercury flyby event. An onboard kernel was created that produced an error profile that went through zero right at the time of the event. A simple model was produced that predicted what the error profile would be by using a certain onboard kernel and assuming that the relativistic clock rate projection was “truth.” The model started with a known TDT–MET pair and then marched forward by using discrete time steps. At each step, the error accumulated to that time because of the difference between the proposed clock rate and the current “truth” clock rate was calculated from the model. The “truth” clock rate was taken directly from the relativistic projection. A few resulting error profile predictions for the third Mercury flyby from the model are shown Figure 5.

Onboard TDT calculation error with OCXO




Figure 5. TDT calculation error during the period leading up to the third Mercury flyby in 2009.

The red, blue, and purple lines in Figure 5 represent three candidate profiles that originated from three different proposed rates to use on the spacecraft for the third Mercury flyby. The process for choosing the ideal rate involved using the model to find which rate centered the error at or near zero during the third Mercury flyby encounter time. In the case above, the blue line shows zero error right at the time of the third Mercury flyby. The rate used to generate the blue profile was the onboard rate eventually used for the encounter. The green dots show the actual error performance that resulted from the activity. The observations closely followed the prediction, but there was a slight offset shown here in the plot. Note that this plot has a smaller error scale than the previous plots. The offset is due to the amount of error accrued on the clock between the time of the MET–TDT pair used to start the model and the time that the load was activated on the spacecraft. The mission operations team was successful in using the relativity projections to guide the selection of the onboard rate and lowering the TDT calculation error considerably.

VI. Performance During Orbit

The next milestone after the third Mercury flyby was MOI and the start of the long-awaited orbital mission phase. The coarse oscillator was used during MOI, and then the OCXO was quickly put into operation once MESSENGER achieved orbit about Mercury. The planned trajectory for the orbital phase of the mission was used to generate a new relativistic clock-drift profile. This new profile was, in turn, used to select an optimal clock rate to store onboard for clock operations. Figure 6 shows the selected clock rate being implemented on May 12, 55 days after orbit insertion.

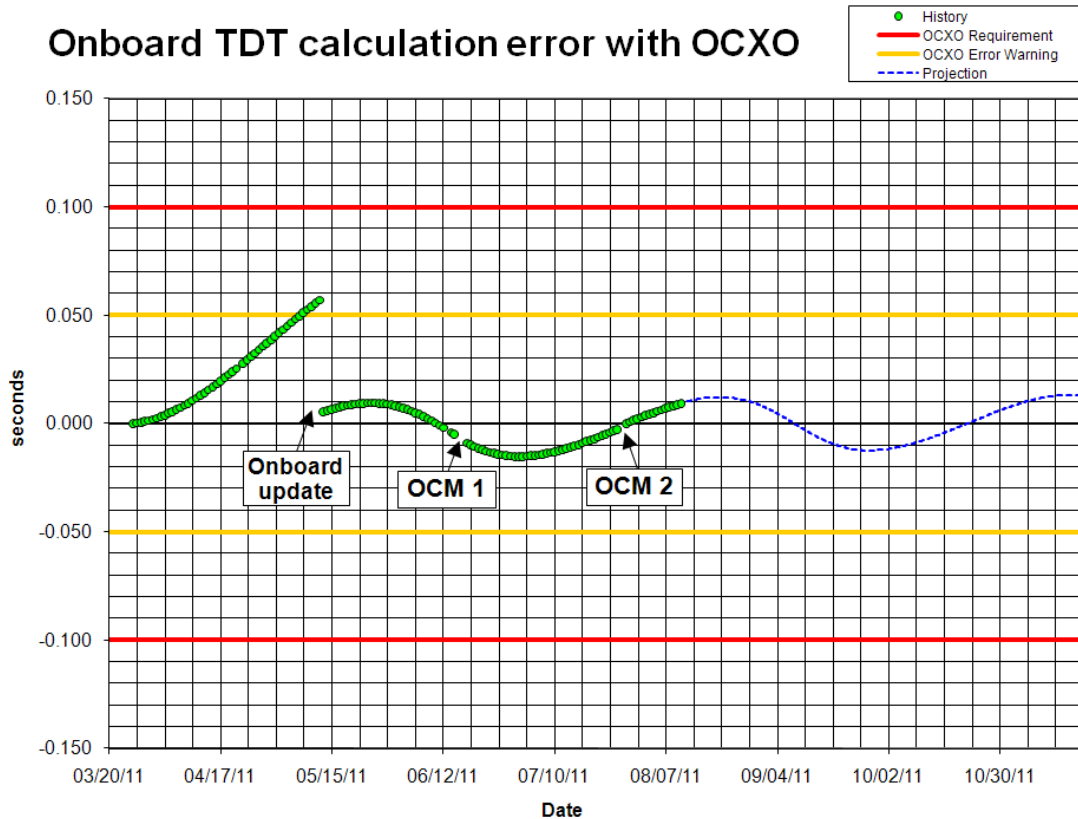


Figure 6, TDT calculation error in Mercury orbit near the time of selection of the OCXO on May 12, 2011.

At the beginning of the orbital phase in late March, the OCXO operations were started by loading a calculated clock rate to the spacecraft. This rate was later updated to correct the growing error, and the rate selected for that update centered the relativistic error around zero. This rate is applicable for the rest of the mission unless there is a major change in the shape of the orbit. The blue dashed line shows the error prediction that comes from the modeling of the relativistic error. The measured error values, represented by the green dots, match this prediction very well.

The timekeeping operations have been going smoothly in orbit but are periodically interrupted by OCMs. Two such interruptions have occurred in orbit so far (OCM 1 and OCM 2), and they are noted on the plot. These events cause the operations to fall back to the coarse oscillator for a duration of slightly more than 1 day. The gaps in the error history curve in Figure 6 show the time that operations were on the coarse oscillator. The mission operations team has streamlined the process of returning to the OCXO for orbital operations. Once an OCM burn is complete, the OCXO is reselected and the data from the reselection are processed the next day. The resulting SCLKs give a measure of how much the clock drifted while not on the OCXO. That same day, a new clock kernel is loaded to the spacecraft. The updated kernel will still have the same clock rate as before, but the accompanying MET-TDT pair is updated to compensate for the coarse oscillator drift. This update effectively translates the large error accrued while on the coarse oscillator back to a smaller error located on the predicted OCXO error curve. The plot shows that operations can easily bridge the gap in OCXO operations and return quickly to the optimal performance curve. Knowledge of the relativistic impact on error and the dependable performance of the OCXO allow the mission operations team to operate the clock with minimal error well within mission requirements.

References

- ¹Cooper, S. B., "From Mercury to Pluto: A Common Approach to Mission Timekeeping," *IEEE Aerospace and Electronic Systems Magazine*, Vol. 21, Oct. 2006, pp. 18–23.
- ²Wallis, R. E., and Cheng, S., "Phased-Array Antenna System for the MESSENGER Deep Space Mission," *2001 IEEE Aerospace Conference Proceedings*, Vol. 1, 2001, pp. 1-41–1-49.
- ³Sobel, D., *Longitude: The True Story of a Lone Genius Who Solved the Greatest Scientific Problem of His Time*, Fourth Estate Ltd., London, 1998.