From Mercury to Pluto:
A Common Approach to Mission Timekeeping

Stanley B. Cooper
The Johns Hopkins University Applied Physics Laboratory

ABSTRACT

The successful launch of the New Horizons spacecraft for a rendezvous with Pluto and Charon and the continuing progress of the MESSENGER spacecraft toward Mercury now positions mankind to unlock mysteries of our solar system from Mercury to Pluto and beyond. Both missions, though very different in concept, use the same generic timekeeping system design. This paper explores how we maintain time on these spacecraft and how we establish on the ground the correlation between spacecraft time and Earth time. It further reviews the sub-millisecond correlation accuracy that has been demonstrated for the MESSENGER mission and the time accuracy we expect to achieve for that mission at Mercury and for the New Horizons mission at Pluto-Charon.

INTRODUCTION

When the Near Earth Asteroid Rendezvous (NEAR) spacecraft launched in 1996, the correlation between spacecraft time and Earth time was determined manually in the Mission Operations Center (MOC) at The Johns Hopkins University Applied Physics Laboratory (APL). That was a tedious process, prone to error. In addition, very accurate time correlation (±20 ms) was needed to support the mission, and that was difficult to maintain with the infrequent manual computations. The time correlation process was redesigned and successfully automated in 1998, in time to support the NEAR spacecraft (later renamed NEAR Shoemaker) becoming the first man-made object to orbit an asteroid and later to land on an asteroid.

The Mercury Surface, Space ENvironment, GEochemistry, and Ranging (MESSENGER) mission to Mercury is operated for the National Aeronautics and Space Administration (NASA) by essentially the same APL team that controlled NEAR. The MESSENGER timekeeping system was modeled after the NEAR timekeeping system, but the details were substantially changed and enhancements made. The MESSENGER timekeeping system is the prototype timekeeping system for other APL missions, including New Horizons. Work began on MESSENGER timekeeping in 2000 and on New Horizons timekeeping in 2002. A third mission, the Solar Terrestrial Relations Observatory (STEREO) was also under development at that time. That provided a unique opportunity to formalize a generic timekeeping system framework that could support all three missions.

THE FLIGHT COMPONENT

Let's first consider the flight component of each of these three "core" missions. In order to correlate spacecraft time to Earth time, we need to identify an onboard time reference event to which the time of all other events on the spacecraft can be referred. For each of these missions, that reference event is the leading edge of a one-pulse-per-second (1 PPS) signal that is generated in hardware at a rate of approximately 1 Hz. Then, knowing the correlation between the 1 PPS reference event and Earth time allows us to determine the correlation between any other onboard event and Earth time.

A second element in common among the three missions is the method by which telemetry frames are downlinked from the spacecraft to Earth. Unlike previous missions such as NEAR, downlink transmission of telemetry frames is not synchronized to the 1 PPS reference event. Knowledge of the frame transmission times does not necessarily provide knowledge of 1 PPS times. However, a typical timekeeping system normally depends on knowledge of the time of transmission of downlink telemetry frames to determine the correlation between the time of the onboard reference event (1 PPS) and Earth time [1]. In order to resolve that dilemma, we have used the same fundamental clocking scheme for all the core missions. First, as usual, each mission has an onboard counter representing an integer number of seconds since some start epoch. For these missions, we have added a second counter representing sub-seconds since the 1 PPS reference

---

Author's Current Address:
The Johns Hopkins University Applied Physics Laboratory, 11100 Johns Hopkins Road, Laurel, MD 20723-6099, USA.

Manuscript received February 22, 2006; revised July 14, 2006.

0885/8985/06/ $17.00 © 2006 IEEE

Table 1. Core mission MET formats

Mission	iMET bits	vMET bits	vMET Resolution	vMET counts per second	Total MET bits
MESSENGER	28	20	1μs	1,000,000	48
New Horizons	32	16	20 μs	50,000	48
STEREO	32	8	3.90625 ms	256	40

event. In keeping with usual APL terminology, we call the composite counter the Mission Elapsed Time (MET). The integer seconds component of MET is denoted "iMET," and the sub-seconds ("vernier") component of MET is denoted "vMET." Both iMET and vMET are downlinked in the header of every telemetry frame. That allows us on the Earth to relate the time of the 1 PPS reference event to Earth time, as described later. Table 1 lists the pertinent characteristics of the composite MET counter for each mission, reflecting the widely divergent requirements of these missions:

Each of the core missions as well as some earlier APL missions package part of the flight electronics in a central unit called the "Integrated Electronics Module" or "IEM." Each IEM includes the Command and Data Handling (C&DH) Subsystem that controls spacecraft command execution and the flow of data between the spacecraft and the ground. Each IEM also includes the processor for the Guidance and Control (G&C) Subsystem that controls attitude and directs firing of thrusters, using a variety of sensors and control mechanisms. The New Horizons IEM takes this one step further and includes several major components of the radio frequency communications (RF) Subsystem. That extension supports implementation of a unique "non-coherent navigation" system. (The original STEREO IEM design [1] was similar to

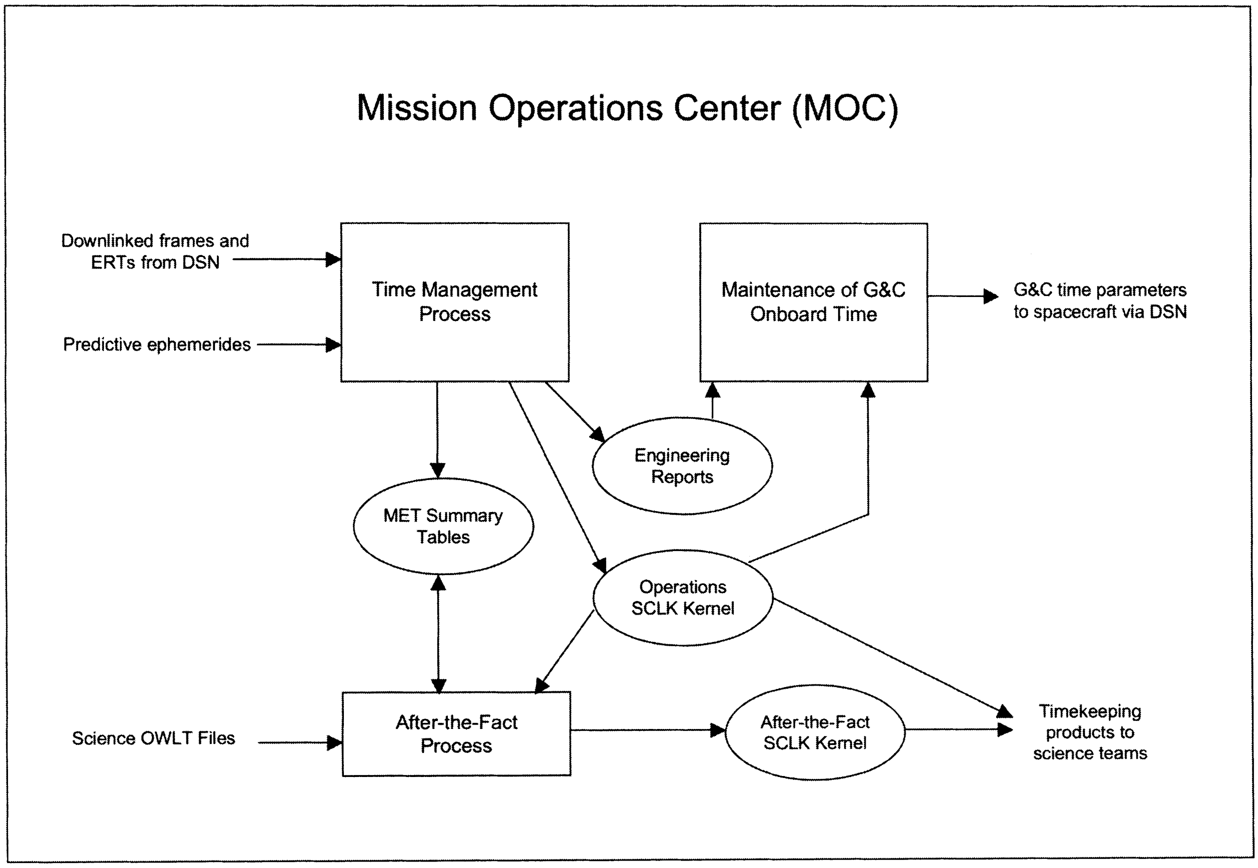

Fig. 1. Framework for the MESSENGER and NEw Horizons timekeeping ground component

the New Horizons IEM but was later changed to a configuration functionally closer to the MESSENGER IEM.) New Horizons and MESSENGER each include two IEMs, one primary, one backup, while each of the two STEREO spacecraft includes a single IEM. Another element in common among the timekeeping systems of the three missions is the use of a common approach to timekeeping system testing during ground "Integration and Test" (I&T) of each spacecraft. Each IEM provides the same three timekeeping test signals to an I&T testbed, allowing the performance and functionality of each timekeeping system to be characterized accurately.

From this point on, the design of the flight component of the STEREO timekeeping system diverges from MESSENGER and New Horizons to comply with substantially different requirements. MESSENGER and New Horizons distribute iMET to all the science instruments. For these missions, we call iMET the "Spacecraft Clock." STEREO instead distributes an onboard estimate of UTC to the science instruments, and we call that UTC estimate the STEREO Spacecraft Clock. For all three missions, the Spacecraft Clock is also used for controlling the time of execution of commands.

THE GROUND COMPONENT

The ground components of the timekeeping systems for the three missions all follow the same framework, although STEREO does not use all the elements defined for MESSENGER and New Horizons. This similarity has allowed substantial reuse of the ground software written to implement the MESSENGER timekeeping prototype. The remainder of this paper focuses on the implementation of MESSENGER and New Horizons, since that is of primary interest in this discussion. Figure 1 is an overview of the elements of the timekeeping system ground component for those two missions.

Referring to Figure 1, each downlink telemetry frame is received by a station of NASA's world-wide Deep Space Network (DSN). The receiving DSN station attaches to each frame an "Earth Received Time" or "ERT" that is the time of receipt in UTC at the station antenna of a reference point in the frame. Subtracting the time it took the frame to travel from the spacecraft to the station gives the time at which the frame reference point was transmitted from the spacecraft antenna. That travel time is commonly called "one-way-light-time" or "OWLT." The core missions are all designed so that the MET placed in the frame header represents the approximate time that the reference point of the previous frame left the spacecraft antenna. Using the sub-seconds component vMET of MET allows us to determine the time of the last 1 PPS reference edge prior to that frame, when vMET was set to zero. That gives us a correlation between iMET, representing the time of the 1 PPS reference event, and UTC. That correlation is saved in a file called the Operations SCLK ("Spacecraft Clock") Kernel.

The three core missions are the first APL missions to use a new DSN capability called turbo coding. For turbo coding, the frame time reference point is located at the leading edge of the first frame bit following the "attached sync marker" or "ASM" that precedes every frame. The current DSN implementation of decoding turbo coded frames has a deterministic bias in ERT as shown in Figure 2 that is fairly large at low downlink bit rates. The solid curve is that deterministic bias for R = 1/6 turbo code, and the dashed curves are the 3σ upper and lower bounds on statistical ERT errors. All errors are with respect to the local station clock, which is kept within microseconds of UTC (NIST), the UTC estimate distributed by the National Institute of Standards and Technology. The next deployment of DSN turbo decoding will eliminate the deterministic bias and leave us with the much smaller statistical ERT errors. Figure 3 provides a more detailed view of the 3σ bounds on the statistical ERT errors.

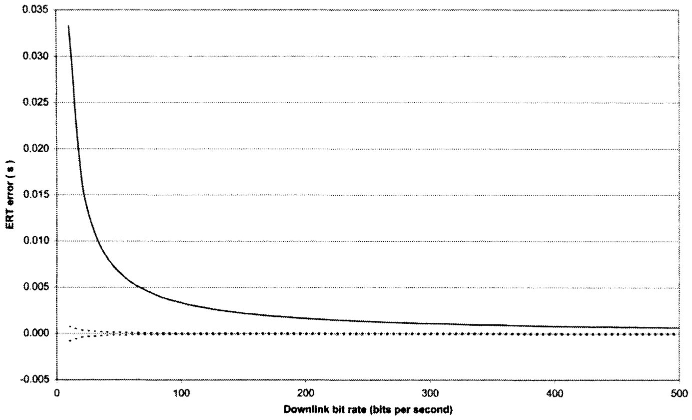

Fig. 2. Bias in DSN Earth Received Time (ERT) for R=1/6 turbo code

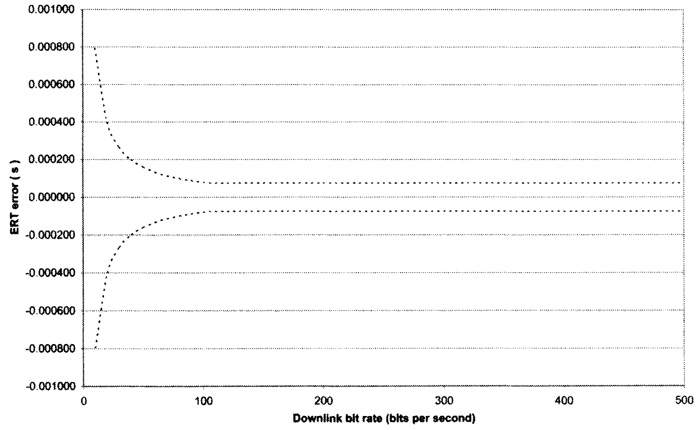

Fig. 3. 3-sigma error in DSN Earth Received Time

The downlink telemetry from DSN is transmitted to the Mission Operations Center (MOC) at APL where it is processed by the automated timekeeping system ground software, identified as the "Time Management Process" in Figure 1. That software first filters the downlink frames to ensure the METs and ERTs are consistent with each other and with previously-processed frame times. These filters were first

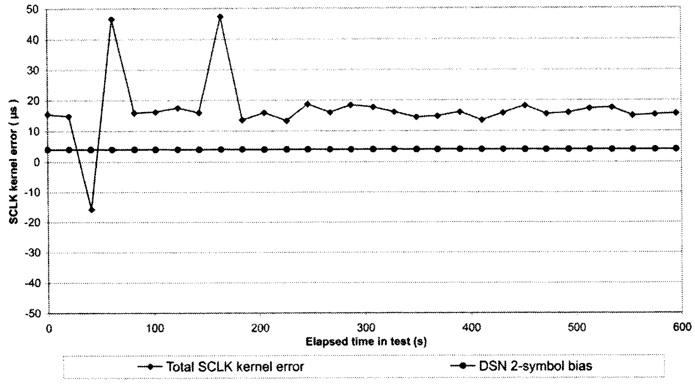

Fig. 4. Operations SCLK Kernel error vs. elapsed time, 5/9/2005

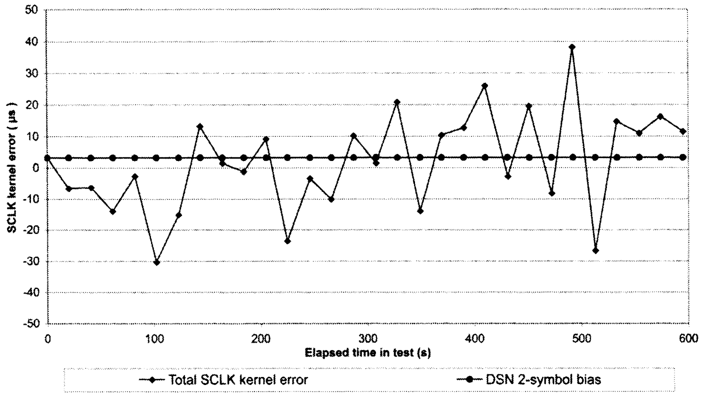

Fig. 5. Operations SCLK Kernel error vs. elapsed time, 8/9/2005

applied on the NEAR mission and have been adapted for our set of core missions. It then correlates frame iMET to Earth time using the time system called "Terrestrial Dynamical Time" (TDT or TT). The TT system uses the same standard atomic second (the "SI" second) used by UTC but does not include leap seconds, and is trivially convertible to UTC. For a given number of leap seconds, TT has a fixed offset with respect to UTC so any drift rate is identical whether expressed in terms of TT or in terms of UTC.

The Time Management Process places the above MET – TT correlation into the MET Summary Tables, a collection of timekeeping results modeled after the NEAR MET Summary Table. Additional information obtained from downlinked data packets (that are contained in downlinked frames) is also placed into the MET Summary Tables. Those Tables are then used to update the Operations SCLK Kernel and various engineering reports.

The accuracy of the MESSENGER Operations SCLK Kernel MET – TT correlations during a flyby of the Earth in 2005 was measured by a unique in-flight timekeeping test system, and the results are shown in Figures 4 and 5. For these tests, the DSN ERT deterministic bias contributed only 3 to 4 µs to the total SCLK kernel error. The Operations SCLK Kernel is updated with an interpolation technique that improves the MET – TT correlation for past observations, and these figures are illustrative of those "past" SCLK kernel correlation accuracies.

An experiment was conducted on May 27, 2005 (during the Earth flyby period), with the Mercury Laser Altimeter (MLA) on the MESSENGER spacecraft, that provided confirmation of Operations SCLK Kernel accuracy. The MLA instrument team was able to determine that the "past" correlation error in the Operations SCLK Kernel was +349 µs during that experiment [2]. That particular error was largely due to a +366 µs bias at that time in the Operations SCLK Kernel caused by a large DSN ERT deterministic bias from a low-rate downlink of telemetry the previous day. The roughly –17 µs residual error is consistent with Figures 4 and 5. Note that the –17 µs residual includes the small uncertainty in timing of the MLA instrument.

Such small time correlation errors are not surprising close to Earth where the environment is fairly benign for timekeeping and where good knowledge of spacecraft ephemeris provides very accurate values for OWLT. Our navigation partner, KinetX, has committed to providing spacecraft ephemerides with an error no worse than ±400 km throughout the mission, including while MESSENGER is on orbit around Mercury. Since OWLT is computed from the spacecraft ephemeris, that means the upper bound in OWLT uncertainty exceeds 1 ms. One requirement for this mission is that we must be able to correlate MLA data times to UTC to better than 1 ms at Mercury, so the Operations SCLK Kernel will not support that requirement. Referring back to Figure 1, we are planning to provide a much more accurate file called the "After-the-Fact SCLK Kernel" that will be used exclusively to support analysis of science data. That Kernel is designed to achieve sub-millisecond accuracy. The software to produce the After-the-Fact SCLK Kernel has not yet been commissioned, so that product is not yet available. Two methods will be used to improve the accuracy over what is available with the Operations SCLK Kernel. The primary improvement will be the use of a more accurate OWLT supplied by KinetX from reconstructed ephemerides. The second improvement will block the use of telemetry frames that were downlinked at lower data rates to avoid the larger statistical ERT errors shown in Figure 3, because the mission accuracy requirement cannot be satisfied with those larger ERT errors. The first delivery of the After-the-Fact SCLK Kernel will include the Earth flyby period, and the low-rate-blocking method will be used that one time to block the large DSN ERT deterministic biases at low downlink rates illustrated in Figure 2. It will be interesting to learn how that improved SCLK kernel will affect the results reported by the MLA instrument team for the May 27 test. Once the ERT deterministic bias is removed from the DSN system in late 2006, we will no longer need to use this work-around to achieve high correlation accuracy. New Horizons will similarly require an After-the-Fact SCLK Kernel to achieve the 10-ms science data time correlation accuracy requirement at Pluto-Charon but, again, the software to produce that Kernel has not yet been commissioned. New Horizons is strictly a "flyby" mission, and it will not block low-downlink-rate telemetry since that telemetry may be required to obtain sufficient timekeeping data to maintain the required "after-the-fact" mission accuracy during the limited flyby window. That detail is reflected in the error budget for the New Horizons After-the-Fact SCLK Kernel. The MESSENGER mission is a bit more forgiving in this regard since it will be in orbit around Mercury for an Earth year, and there will be multiple opportunities for gathering science data in the unlikely event that too much downlink timekeeping data is lost during low-rate downlinks.

The block labeled "Maintenance of G&C Onboard Time" in Figure 1 represents the combined human/software function of ensuring that the TT estimate that is used onboard by the G&C Subsystem is maintained within the accuracy requirements of the mission. MESSENGER and New Horizons use exactly the same approach both in onboard computation of TT and in determination of the time parameters that need to be uploaded to the spacecraft to enable that computation. First, the accuracy of onboard TT is monitored in the Mission Operations Center via the Operations Time Engineering File (a.k.a. Time History File) to ensure that the computation of TT satisfies requirements. The Operations Time Engineering File is one of the engineering reports generated by the automated Time Management Process. Whenever necessary, new G&C time parameters are extracted from the information in the most recent time record of the Operations SCLK Kernel and used to create a command that will execute on the spacecraft to load the new parameters into the G&C Subsystem. Those parameters include an iMET and corresponding TT and predicted rate of change of TT with respect to iMET. Onboard, the G&C Subsystem receives iMET from the C&DH Subsystem and uses that and the new time parameters to compute the onboard estimate TT(S) of Earth time, including compensation for predicted iMET drift.

As implied above, the Operations Time Engineering File reports to the MOC analyst the estimated error in TT(S). That estimate is determined very simply by using the Operations SCLK Kernel to map the iMET value that was used onboard for computing a particular TT(S) value to a corresponding estimate TT(G) of the time of that same reference event (1 PPS reference edge). The difference between TT(S) and TT(G) is the estimated error in TT(S), provided the mapping from iMET to TT(G) is sufficiently accurate.

Whenever the Time Management Process updates the Operations SCLK Kernel, it also adds a new time record to the Operations Time Engineering File. That record contains a number of parameters of interest to the Operations staff. It includes, for example, the estimated current error in the predicted correlation between iMET and TT(G), the estimated drift rate of the MET counter, the designation of the receiving DSN station and the estimated OWLT from the spacecraft to that DSN station, the approximate downlink bit rate and the designation of the IEM that provided the timekeeping information that was used to create the new time record. Another interesting product generated automatically by the Time Management Process is the Backup IEM Time Engineering File. That file provides a comparison between the iMETs in the primary and backup IEMs to aid Operations staff in monitoring the performance of the MET in the backup IEM.

ERROR BUDGETS

Each mission has multiple time error budgets to deal with several categories of time accuracy requirements. There are requirements on the accuracy of the onboard estimate TT(S) used by the G&C Subsystem, requirements on the "quick-look" accuracy achieved from the Operations SCLK Kernel for predictions of future MET – TT correlations, and requirements on the "after-the-fact" accuracy for correlation between science data times and TT. We will examine the error budgets for correlation between science data times and TT using the After-the-Fact SCLK Kernel.

The After-the-Fact SCLK Kernel supports correlation of past values of iMET with TT and does not provide any prediction of future correlations. All MET – TT correlations are based on simple linear interpolation between time correlation records in the SCLK Kernel. In the error budgets shown below, the measurement uncertainty U_t is the upper bound on the uncertainty in the correlation between a particular iMET received in the header of a downlink frame and the corresponding TT computed from ERT, OWLT, and vMET. These error budgets apply only after the oscillator that drives the MET counters has stabilized; other budgets apply prior to oscillator stabilization.

The tightest accuracy requirement for MESSENGER is that the correlation between MLA data times and TT should be no worse than 1 ms. Here is the MESSENGER worst-case error budget for the After-the-Fact SCLK Kernel with the MLA after the oscillator has stabilized:

- 11 μs allowance for MLA instrument time uncertainties;
- 10 μs to 12 μs allowance for uncertainty due to distribution of 1 PPS to MLA;
- 130 μs bound on interpolation error due to measurement uncertainty U_t;
- 173 μs bound on interpolation error due to aging (7 days after oscillator turn-on);
- 90 μs bound on interpolation error due to background solar proton radiation;
- 44 μs bound on interpolation error due to relativistic effects on oscillator;
- 242 μs bound on interpolation error due to oscillator temperature/voltage effects; and
- 2 μs bound on computation error due to precision of SCLK Kernel change rate.

This budget adds up to ±704 μs, leaving a margin of almost 30% of the ±1 ms accuracy requirement. The error budget that applies when the oscillator has not stabilized leaves a smaller margin that still exceeds 20%.

The tightest accuracy requirement for New Horizons involves the New Horizons Radio Science EXperiment (REX), the Multispectral Visible Imaging Camera (MVIC), the Ultraviolet Imaging Spectrometer (ALICE), and the LOng Range Reconnaissance Imager (LORRI). The correlation between data times of the REX, MVIC, ALICE, and LORRI instruments and TT should be no worse than 10 ms at Pluto-Charon and at the farther Kuiper Belt Objects (KBOs). Here is the New Horizons worst case error budget at Pluto-Charon for the After-the-Fact SCLK Kernel with the ALICE instrument. Note that the larger error budget in this case is expressed in terms of milliseconds rather than microseconds:

- 2.5 ms allowance for ALICE instrument time uncertainties;
- 1.2 ms bound on interpolation error due to measurement uncertainty U_1;
- 0.1 ms bound on interpolation error due to aging and background neutron radiation;
- 0.1 ms presumed bound on interpolation error due to relativistic effects on oscillator;
- ~ 0 ms approximate interpolation error due to oscillator temperature/voltage/loading effects; and
- ~ 0 ms computation error due to precision of SCLK Kernel change rate.

This budget adds up to 3.9 ms, leaving a margin of 6.1 ms at Pluto-Charon. The error budget at the KBOs may total several milliseconds more but will still be well under the 10-ms requirement. The details of the ephemeris knowledge (and hence OWLT uncertainty) at the KBOs are still being worked and may be better than anticipated.

CONCLUSION

A generic timekeeping system framework has been developed at APL and used successfully to support design, implementation, test, and in-flight operation of the mission timekeeping systems for deep space missions. To date, three missions have benefited from this approach from design through the test phase. Two of these missions have been launched and the timekeeping systems for these have, so far, met all expectations. In particular, in-flight testing of the timekeeping system for the MESSENGER mission to Mercury has demonstrated accuracy in "past" correlation between MET and TT (or UTC) of the order of tens of microseconds during an Earth flyby. We expect larger errors for this mission near Mercury and for the New Horizons mission when it reaches Pluto-Charon, but we do expect that all mission timekeeping accuracy requirements will be satisfied.

ACKNOWLEDGEMENTS

A large number of individuals in several organizations contributed to the development of the generic timekeeping system described herein. Although there are too many to mention, a few who made particularly important contributions to this development deserve to be acknowledged. These include: Rich Conde, Steve Jaskulek, Joanna Mellert, Mike Reid, and many others at APL; Bobby Williams and his navigation team at KinetX; and Rich Benson, Jeff Berner, Dan Finnerty, and Al Hewitt of the Deep Space Network component of NASA's Jet Propulsion Laboratory (JPL).

REFERENCES

[1] S.B. Cooper,
Principles of Timekeeping for the NEAR and STEREO Spacecraft,
NASA Scientific and Technical Information (STI) Report,
NASA/CR-2001-209988, pp. 4-5 and 13-16, 2001.

[2] D.E. Smith, M.T. Zuber, X. Sun, G.A. Neumann, J.F. Cavanaugh,
J.F. McGarry and T.W. Zagwodzki,
Two-Way Laser Link over Interplanetary Distance,
Science, 311, p. 53, 2006.

[3] K. Hurley, T. Cline, E. Mazets, R. Aptekar, S. Golenetskii,
D. Frederiks, D. Frail, S. Kulkarni, J. Trombka, T. McClanahan,
R. Starr and J. Goldsten,
Interplanetary Network Localization of GRB 991208 and the
Discovery of Its Afterglow,
The Astrophysical Journal, 534,
pp. L-23-L-25, 2000 May 1.