THEMIS/ARTEMIS Timekeeping Conventions

THM-SOC-135

April 5, 2012

Document Revision Record

Rev.	Date	Description of Change	Revised By
-		Initial Release	
B	2011-11-17	Added details for handling of SLP and GBO data.	JWL
C	2011-12-21	Incorporated feedback from Bryce Roberts (MOC) and Harald Frey (GBO).	JWL
D	2012-01-09	Minor correction to S/C clock management details, noted leap second scheduled at the end of June 2012.	JWL
E	2012-04-05	Emmanuel cleaned up TOC; JWL changed body text font, added notes about TT2000 CDF times and how they are handled in TDAS.	EM, JWL

Distribution List

____
Jim Lewis, THEMIS/ARTEMIS Science Software Team

____
Manfred Bester, THEMIS/ARTEMIS Mission Operations Manager

____
David A. King, THEMIS/ARTEMIS Science Software Team Manager

____
Vassilis Angelopoulos, THEMIS/ARTEMIS Principal Investigator

Table of Contents

1. Introduction

1.1 Purpose

When comparing data sets from multiple sources and institutions, it is useful to express event times in a common, standardized, unambiguous time system. UTC (Coordinated Universal Time), the basis for most definitions of civil time worldwide, is a natural choice. However, the necessity of dealing with leap seconds (which are part of the definition of UTC) introduces some subtleties that must be handled carefully, in order to avoid systematic errors in converting project-specific timekeeping conventions to a common UTC representation.

In particular, the widespread use of "Unix time" for a numerical representation of UTC time strings is somewhat problematic, due to the fact that (1) the definition of "Unix time" explicitly states that leap seconds are not accounted for, and (2) there are no libraries in widespread use among the scientific software development community to support leap-second-aware calendrical computations. Many developers are not even aware of the leap second problem, and may consider "UTC", "UT", and "Unix time" as synonyms, further complicating the task of reconciling timekeeping conventions between different projects.

Owing to the difficulty of independently establishing the exact times of occurrence of phenomena measured by the THEMIS/ARTEMIS spacecraft, a direct end-to-end absolute time calibration with precision approaching one second of time is not feasible. It is therefore necessary to use more indirect methods to demonstrate the correct interpretation of the timestamps used in our data products.

This document is intended to provide a clear and correct understanding of the timekeeping conventions and procedures used by the THEMIS/ARTEMIS spacecraft and data products, and other related data sets such as the ground magnetometer network, all-sky imagers, and the spacecraft, sun and moon ephemerides.

2. Definitions

UT ("Universal Time"): A time system referenced to the earth's rotation, measured by observations of celestial bodies from the Earth's surface. UT can be considered as a modern extension of Greenwich Mean Time. There are several variants of Universal Time (UT0, UT1, UT1R, etc.), which exist to smooth out such effects as polar motion and periodic tidal and seasonal variations. The UT1 variant is the one most commonly encountered.

TAI ("International Atomic Time"): A time system referenced to the SI second and measured by atomic clocks. This time system does not use leap seconds. Since the earth's mean solar rotation period differs from an integer number of SI seconds, TAI times eventually drift out of sync with the day/night cycle. Far in the future, 12:00:00 noon (TAI) eventually will fall near 00:00:00 midnight (UTC).

UTC ("Coordinated Universal Time"): A time system referenced to the SI second and measured by atomic clocks. Unlike TAI, UTC is kept within 0.9s of UT1 by the insertion (or theoretically, deletion) of leap seconds, as needed, so that 12:00:00 noon (UTC) is always close to 12:00:00 noon Greenwich mean solar time. 1972-01-01/00:00:00 UTC corresponds to 1972-01-

01/00:00:10 TAI, and for all later times, TAI and UTC differ by an integer number of seconds. As of this writing (April 2012), there have been 24 additional positive leap seconds since then (with a 25th positive leap second to be inserted at the end of June 2012), so that TAI and UTC currently differ by 34 seconds. For example, 2011-01-01/00:00:00 UTC corresponds to 2011-01-01/00:00:34 TAI.

GPS ("Global Positioning System") time: To support the precise timekeeping required for satellite-assisted navigation, each GPS satellite carries an atomic clock, which is set to "GPS time". GPS time is not adjusted for leap seconds. It was defined to match UTC in 1980, but has diverged since then. GPS time therefore has a fixed relationship to TAI, such that TAI-GPS equals 19 seconds. The GPS navigation message includes the difference between GPS time and UTC (15 seconds, at this writing), and most GPS receivers use this offset to display the time as UTC, or convert it to a user-preferred time zone. Some GPS receivers have the capability to output a very accurate pulse-per-second signal, which can be used to keep a computer's real time clock (RTC) synchronized with UTC, by measuring and compensating for the RTC drift rate.

Unix time: Unix time (also called "POSIX time", after the relevant standardization body) is defined as a count of seconds since the Unix epoch (1970-01-01/00:00:00 GMT), but not including leap seconds. A standardized group of library functions (gmtime(), gettimeofday(), mktime(), time(), etc.) exists to facilitate conversion between date+time strings and numeric time_t values. In this system, leap seconds are not unambiguously representable -- for example, the time_t values for the dates "2008-12-31/23:59:60" (the 2008 leap second) and "2009-01-01/00:00:00" are identical, even though they represent different times. Another consequence of this definition is that differencing two time_t values does not correctly represent the true elapsed time, for the case when that time interval includes one or more leap seconds. To get leap-second-aware time differences between two dates, a leap second correction must be applied, and there is no standard Unix library function that supports this. Any third-party solution would require patches from time to time, as new leap seconds are announced by IERS (the International Earth Rotation Service).

NTP (Network Time Protocol): NTP is a network protocol that allows a workstation to keep its local clock synchronized to UTC, by communicating with one or more time servers over a network. The protocol takes network propagation delays into account, and can use a pool of several time servers for improved accuracy and fault tolerance. NTP servers are characterized by their "Stratum number", with Stratum-0 referring to the reference clock itself (for example, a cesium clock, or GPS receiver), Stratum-1 referring to a server directly attached to a reference clock, Stratum-2 referring to a server that is synchronized via NTP to a Stratum-1 server, and so forth. NTP can be used to "jam" the local clock to the correct time (via the ntpdate command, for example), or it can run as a daemon, and incrementally discipline the local clock to keep it synchronized with a lower stratum NTP server or pool of servers.

TDAS (THEMIS Data Analysis Software): A suite of data processing and visualization tools, implemented in IDL, for working with data from THEMIS/ARTEMIS and other missions. The TDAS timekeeping convention is defined by the pair of functions time_string() and time_double(), which convert between string and floating-point numeric representations of dates. The underlying logic is based on Unix time, therefore leap seconds are ignored during the conversions and when performing time calculations in the floating-point domain.

CDF: "Common Data Format", a standardized binary file format and library used by various NASA missions (including THEMIS/ARTEMIS) for representing and publishing their scientific data sets. The CDF library includes several data types for representing timestamps: CDF_EPOCH (referred to as CDF_EPOCH8 in the IDL implementation), an 8-byte data type capable of representing times to millisecond precision; CDF_EPOCH16, a more precise 16-byte time representation; and in recent versions of the CDF library, TT2000, an alternate 8-byte time representation. Conversions to and from CDF_EPOCH and CDF_EPOCH16 are not leap-second-aware. Conversions to and from TT2000 do account for leap seconds.

TT2000: A CDF data type introduced in version 3.4.0 of the CDF library, which represents times as a signed count of nanoseconds elapsed since 2000-01-01 12:00:00 TT (terrestrial time). Unlike the CDF_EPOCH and CDF_EPOCH16 types, conversions between date strings and TT2000 values account for leap seconds. THEMIS does not currently use TT2000 timestamps, but TDAS has the capability to load CDFs that use TT2000 sample times (for compatibility with other missions such as RBSP). For interoperability with the rest of TDAS, TT2000 timestamps are converted to Unix times during the loading process (by subtracting out leap seconds, shifting the epoch to 1970-01-01, and converting from nanoseconds to double-precision floating point seconds). CDF's implementation of TT2000 requires a table of leap seconds that must be updated whenever IERS announces a leap second insertion. TDAS will automatically download the updated TT2000 leap second table from GSFC whenever a change is detected. Users should be aware that other libraries (in particular, SPICE/ICY) may have their own leap second tables, and may require occasional manual updates to keep in sync with the CDF TT2000 updates.

3. THEMIS/ARTEMIS Spacecraft Clock Management

THEMIS and ARTEMIS have adopted an "epoch + offset" representation for timekeeping onboard each of the five spacecraft. The THEMIS epoch is defined as 2001-01-01/00:00:00 UTC. Two leap seconds have occurred since the THEMIS epoch, at the end of the years 2005 and 2008, with a third leap second to be inserted at the end of June 2012.

Each spacecraft has a MET (Mission Elapsed Time) counter, which is initialized to 0 whenever the BAU (Bus Avionics Unit) is restarted, and counts seconds and subseconds since the BAU restart. A second quantity, the so-called UTC offset (telemetry mnemonic UTCO), represents the time interval between the THEMIS epoch and the most recent BAU restart. The UTCO value can be adjusted via commanding from the THEMIS MOC (Mission Operations Center), to compensate for any drift in the MET counter.

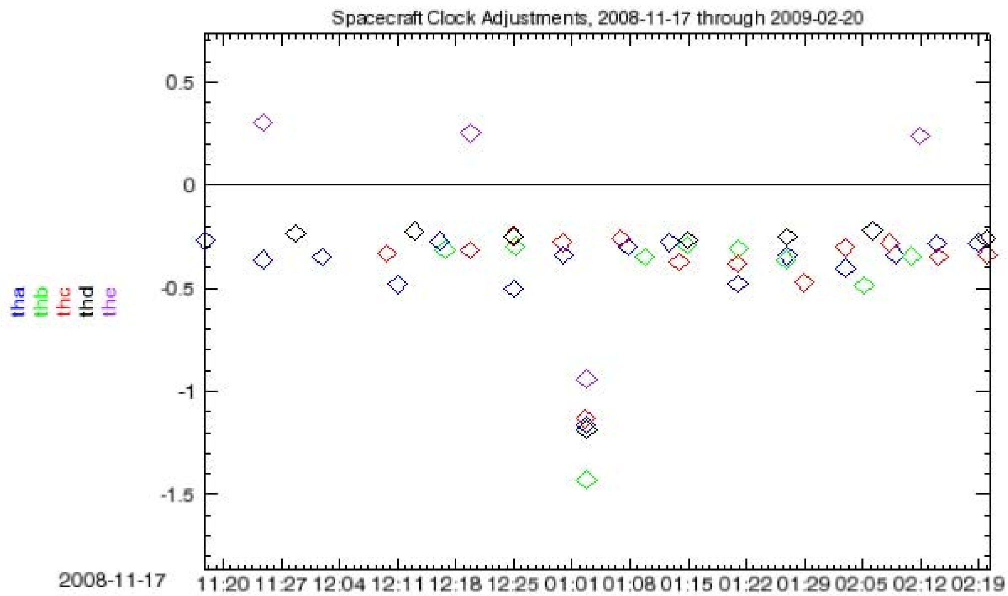

Figure 1: UTCO adjustments around 2008 leap second

A 1 PPS clock signal is generated by the BAU, with the pulses denoting the times when the MET subsecond field is zero; this is referred to as the "one second mark". A message is sent from the BAU to the IDPU (Instrument Data Processing Unit) once per second, including the spacecraft time (UTCO + MET) at the most recent one second mark; this value is recorded in the apid 0x404 housekeeping packets (telemetry mnemonic ONESECMARK). The IDPU uses this information to timestamp the packet and frame headers it generates. The packet timestamps and ONESECMARK values are used by the L0->L1 telemetry processing, along with knowledge of IDPU FSW (flight software) interrupt times and instrument configuration, to calculate accurate timestamps for each data point.

This clock management scheme assumes the existence of an accurate time source on the ground, against which the required UTCO values can be measured and implemented. The fact that this process is implemented on a Unix workstation raises the question of whether the UTCO adjustments take leap seconds into account, or are performed in a non-leap-second-aware fashion using the native Unix libraries. An examination of the UTCO corrections performed near the time of the 2008 leap second shows that leap seconds do not appear to be accounted for in this process. See Figure 1, which shows the magnitude (in seconds) and sign of the adjustments performed on all five spacecraft for several months bracketing the 2008 leap second. Each spacecraft displays a rather stable clock drift rate over this period. But in each case, the first clock adjustment of 2009 shows that the value of UTCO was adjusted to about one second earlier in time than would be expected based on the observed clock drift. The only reasonable explanation for this is that the 2008 leap second was not accounted for when the ground and onboard clocks were compared for the first time in 2009, and the discrepancy due to the leap second was absorbed into the UTCO correction. The calculated clock drift rate for this time interval is also "out of family" for each probe. When the clock drift is recalculated after removing the leap second from the UTCO adjustment, the result is much closer to the nominal value for each probe. See Table 1 below, showing the history of UTCO corrections for THEMIS-A bracketing the 2008 leap second.

Date	delta UTCO (sec)	Interval (days)	Raw Drift (sec/day)	Leap Corrected Drift
2008-11-24/18:16:56	-0.36070251	6.9139773	-0.052170046	
2008-12-01/20:38:00	-0.34988403	7.0979589	-0.049293612	
2008-12-11/00:35:32	-0.48193359	9.1649481	-0.052584432	
2008-12-16/00:31:31	-0.27279663	4.9972191	-0.054589688	
2008-12-24/23:51:51	-0.50161743	8.9724479	-0.055906419	
2008-12-30/23:10:51	-0.34136963	5.9715238	-0.057166251	
2009-01-02/16:15:01	-1.1615601	2.7112365	-0.42842447	-0.059589069
2009-01-07/18:30:01	-0.29821777	5.0937465	-0.058545860	
2009-01-12/17:15:57	-0.27891541	4.9485616	-0.056362925	
2009-01-20/22:51:48	-0.47744751	8.2332353	-0.057990267	
2009-01-26/22:37:40	-0.34400940	5.9901812	-0.057428880	
2009-02-02/22:20:52	-0.40330505	6.9883286	-0.057711232	
2009-02-08/22:04:07	-0.33775330	5.9883758	-0.056401487	
2009-02-13/22:03:22	-0.28263855	4.9994758	-0.056533637	
2009-02-18/21:41:26	-0.27934265	4.9847653	-0.056039279	

Table 1: THEMIS-A clock corrections, correction baseline intervals, and drift rates, showing the anomalous behavior around the 2008 leap second.

Bryce Roberts, from the THEMIS/ARTEMIS mission operations team, has provided the following detailed description of the spacecraft clock management strategy:

> Before each frame is transmitted, the time tag in the frame secondary header is filled in with the spacecraft's current time (MET + UTCO). When the first bit of this frame reaches the antenna, the frame is annotated with an earth receive timestamp, as received from a highly accurate GPS receiver-based UTC clock. Despite being a GPS-based time source, the time code processor converts from GPS standard to true UTC by introducing leap seconds, as needed. The UTC timestamp from this clock is inserted into the earth received timestamp field of the SMEX frame, using the "PB5" format (in which the date is broken down into fields for truncated Julian day, seconds, and milliseconds). (This is the step which determines whether the UTCO values are leap-second-aware, or not. Note that the UNIX time of workstations slaved to UTC by NTP is irrelevant, since the ground receive time comes only from a GPS clock's time code processor.)
>
> After removing the electronic delays between the time tagging and actual transmission -- which was characterized on the ground before launch -- as well as the one way light travel time, we compare the delta between spacecraft and ground time, assuming that ground time is "truth". If the delta is large enough, we make a step function adjustment in the s/c clock to reduce the delta to 0. (This comparison requires a conversion of the PB5 earth receive timestamp to an offset, in seconds, from the THEMIS epoch. The treatment of leap seconds in this conversion determines whether the THEMIS timestamps account for leap seconds, or not.)
>
> The mission requirement is to keep all THEMIS clocks to within 1 second of each other, which can be met or exceeded if all spacecraft clocks are within +/- 0.5 seconds of UTC. In practice, we usually adjust the clock when the delta reaches about 0.33 seconds, to add margin against the possibility of a spacecraft going outside the 0.5 tolerance before the next available ground station pass.

Leap seconds are applied as if the UTC second 23:59:59 was "held" for 2 seconds. This occurs instantaneously and automatically in the time code processors in the MOC that are responsible for doing the earth receive time annotation, but the spacecraft doesn't have a similar, instantaneous knowledge of the leap second -- it keeps free counting. Consequently, during the first ground station pass after a leap second is applied, there suddenly appears to be an increase of 1 second in the clock delta. Thus, the controllers make an effort to perform clock adjusts soon after the leap seconds are applied, and try to do so on all spacecraft as close to simultaneously as possible. However, it is true that in the brief time between the application of the leap second and the successful clock correlation of the 5th THEMIS probe, the probe-to-probe clock correlation is somewhat indeterminate.

For the most accurate comparison of THEMIS/ARTEMIS probe data timestamps with other data sets, any ground based leap second processing should take into account not only that the leap second occurred, but exactly when each of the free-running spacecraft clocks was re-slaved to UTC. We can provide the precise times at which the change was effected by looking into our engineering logs.

4. THEMIS/ARTEMIS Ephemeris Timekeeping Conventions

The THEMIS MOC is responsible for producing and maintaining predicted and definitive ephemerides for the five spacecraft, including (among other quantities) probe positions and velocities sampled at one-minute time intervals. They are published as ASCII files, with UTC timestamps expressed as Year/DOY HH:MM:SS.SSS. The software used to produce these files is patched, when necessary, to keep track of leap seconds as they are announced by IERS.

These ASCII files are ingested by the THEMIS SOC automated processing routines, and incorporated into the binary L1 STATE CDFs. The STATE processing code uses the IDL time_double() function to convert the UTC time strings from the ASCII ephemeris into Unix-style double precision values (unleaped seconds since the Unix epoch).

Since the conversion to Unix time is not leap-second-aware, but the underlying ephemeris is, we would expect to see some sort of anomaly corresponding to the 2008 leap second. The ASCII ephemerides are sampled at 2008/366 23:59:00 and 2009/001 00:00:00 UTC. The true time interval between these samples is 61 seconds, due to the leap second at 2008/366 23:59:60 UTC. However, differencing the Unix timestamps for these data points, as recorded in the STATE CDFs, yields a calculated time interval of 60 seconds.

This effect can be seen by extrapolating each position+velocity sample 60 seconds forward, to the next ephemeris sample Unix time, and comparing the extrapolated position to the next value recorded in the ephemeris. See Figure 2 (centered on 2009-01-01/00:00:00, showing leap second) and Figure 3 (centered on 2008-01-01/00:00:00, no leap second) for such an analysis performed by Harald Frey.

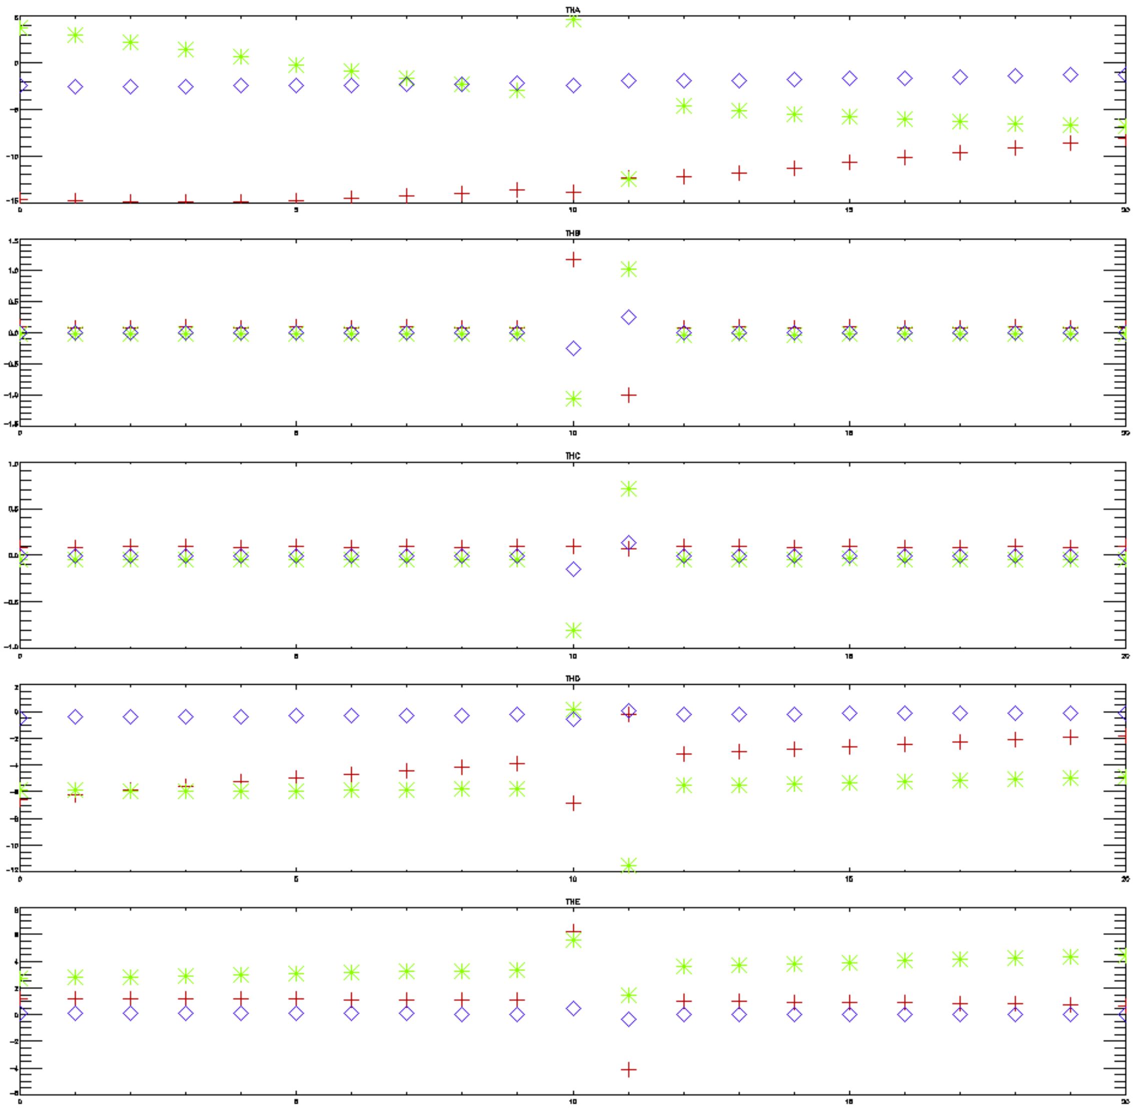

Figure 2: Differences between extrapolated and recorded probe positions, for time interval containing 2008 leap second

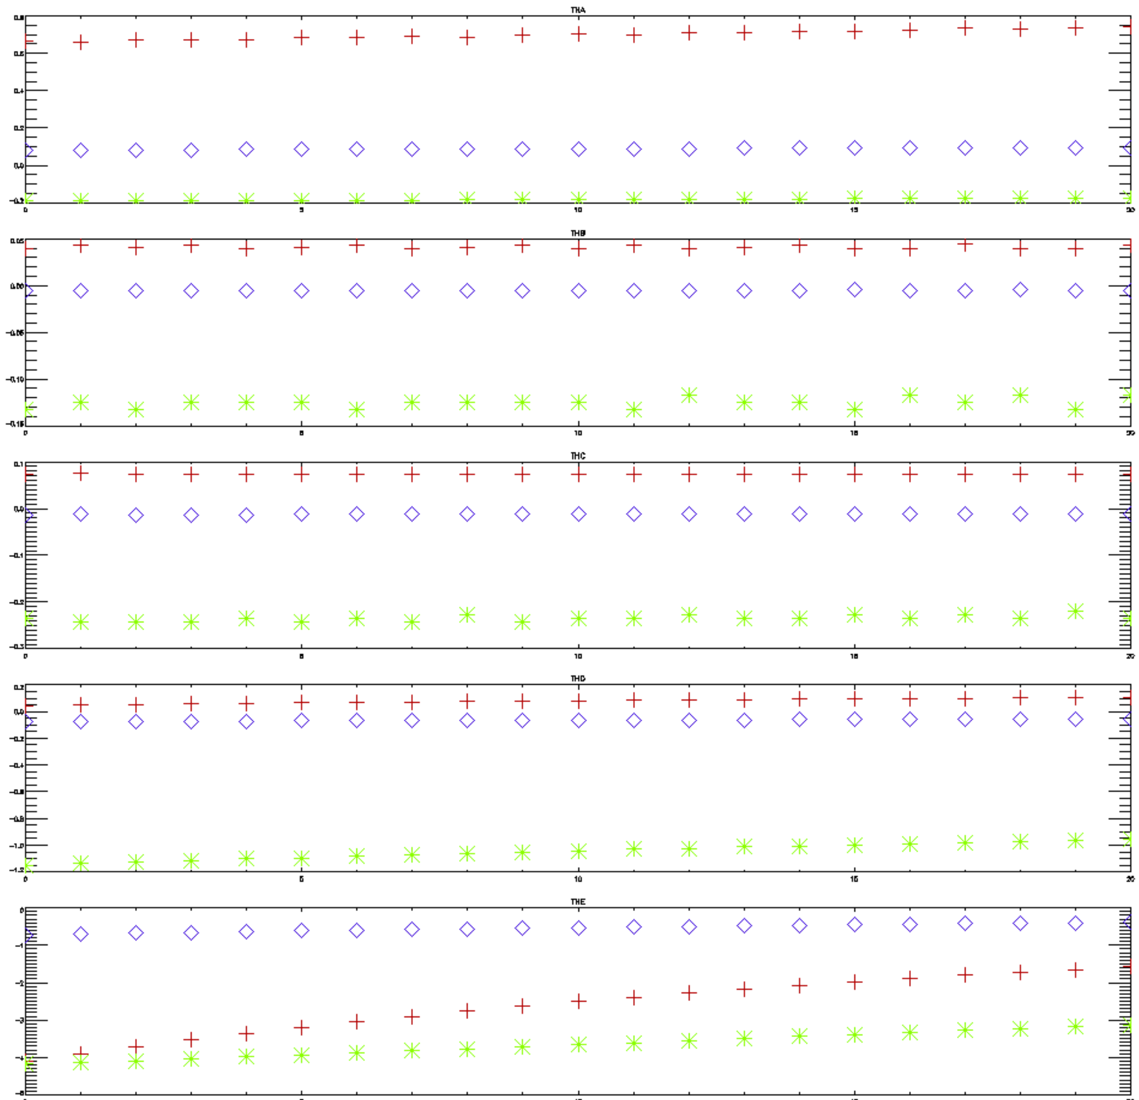

Figure 3: Differences between extrapolated and recorded probe positions, for time interval without leap second.

5. Solar and Lunar Ephemeris Timekeeping Conventions

The THEMIS project maintains a set of Level 1 SLP CDFs, containing solar and lunar ephemeris data used by TDAS for certain coordinate transforms. The UTC timestamps from these ephemerides are converted to double-precision Unix times for storage in the L1 SLP CDFs. Since the orbits of the moon and earth are well-understood, and not subject to disturbances due to maneuvers, the SLP files can be generated en masse for times extending out to year 2020. The sun and moon data is generated in IDL, using the SPICE/ICY libraries. These libraries use data files, called "kernels", to specify (among other things), a table of leap second offsets relative to a particular reference date. At this writing, we are using the kernel file naif0009.tls, which contains all leap second offsets from 1972 to 2009. In the future, whenever additional leap seconds are inserted, we will update the leap second kernel and reprocess the L1 SLP data set to account for the new leap second. The next leap second insertion is scheduled for the end of June 2012.

6. THEMIS GBO Clock Management

The THEMIS network of ground magnetometers and all-sky imagers are generally located in remote areas, far from sources of electromagnetic interference and light pollution. Maintaining continuous, reliable network access for NTP timekeeping is generally not feasible under such conditions. Therefore, each GBO station is equipped with a GPS receiver to provide an accurate UTC time reference. Some of the details are documented here:

http://www-ssc.igpp.ucla.edu/personnel/russell/papers/themis_ground-based_magnetometers.pdf

Here is an excerpt with some details about the GBO clock management:

> Timing accuracy is obtained by combining the time of day accuracy of the Acutime-2000 GPS receiver with the long-term stability of a Vectron Series TC-400 temperature compensated crystal oscillator (TCXO). The combination of these two time sources allows for synchronization of data collection between magnetometers anywhere in the world. The Acutime 2000 GPS provides a 1 Hz strobe that is synchronized to occur at the zero millisecond boundary of every 1 second tic. When sufficient GPS satellite signals are present, this 1-second strobe is used to time tag the data processing logic within the magnetometer. The local oscillator derived from the Vectron TCXO is also periodically resynchronized to the GPS 1-second strobe. Each time this occurs, a measurement is obtained of the drift rate of the TCXO. These data are collected as a part of the overall GMAC data log stream. In the event that the GPS satellite signal is lost, the TCXO will free run allowing seamless data collection for an extended period of time. Without GPS satellite signal, the TCXO will be expected to drift at a rate of less than 0.2 seconds per day. This drift rate limits the duration in which synchronization of data sampling between stations can be maintained. The drift rate of the TCXO for each station is however well-known, due to the continual measurement of its drift relative to the GPS 1-second strobe. Therefore, time knowledge can be corrected to an uncertainty of less than 10 msec per day.

The GMAG and ASI data are locally time stamped by UTC time, and the conversion to CDF time is done at Berkeley. The clearest indication of this synchronization was a jump in time in the THEMIS magnetometer data 11-13 seconds after the leap second insertion on 2008-12-31. Magnetometer data are recorded twice a second and written to RMD-files every 5 seconds. The RMD files of stations CHBG, INUV, KAPU, KIAN, KUUJ, MCGR, SNAP, TPAS, and WHIT all show the time jump during the first minute in the 2009-01-01 files.

THEMIS also makes available CDF data sets derived from other magnetometer networks. Timestamps are usually received as text strings, interpreted as UTC, and converted to numeric values using the TDAS time_double() function, which follows the Unix "no leap seconds" convention.

7. THEMIS/ARTEMIS L1 and L2 CDF Timekeeping Conventions

The telemetry received from the THEMIS/ARTEMIS probes is formatted as CCSDS source packets, with a single timestamp located in the packet header. Our L0 data products consist of sequences of packets, grouped by application ID and UTC date, more or less exactly as they were produced on the spacecraft. Most packets contain multiple samples, and during the L0->L1 processing, individual sample times are calculated using the packet header time, the position of each sample within the packet, any internal delays or time offsets present in that instrument, and the sample interval (derived from instrument configuration bits in the packet header).

As described above, the spacecraft timestamps are maintained as a count of seconds and subseconds since the THEMIS epoch (2001-01-01/00:00:00 UTC), not accounting for leap seconds. These timestamps are converted to Unix times by adding a constant offset of 978307200 seconds – the Unix time corresponding to the THEMIS epoch – and written into the L1 CDFs. The L2 CDFs (calibrated, converted to engineering units, and transformed to geophysically relevant coordinate systems) follow the same convention.

Therefore, with a few caveats, the timestamps in the THEMIS/ARTEMIS L1 and L2 products are to be interpreted as Unix times, and the Unix conventions should be followed when converting the numeric timestamps to UTC date strings or other time bases.

Caveats: The L1 and L2 CDF times are not corrected for clock drift between UTCO resets. A backward adjustment of UTCO occurring while a data packet is in progress can lead to non-monotonic sample times in the L1 data. For time intervals between a leap second insertion and the subsequent UTCO adjustment, the L1 sample times will be one second ahead of true Unix time (since the MET counter runs continuously through the leap second). The next UTCO adjustment will absorb the leap second, and after that point the L1 sample times will again correspond to Unix times.