Chapter 17
The UT1 and UTC Time Services Provided by the National Institute of Standards and Technology

Judah Levine

Abstract I will describe the network-based time services provided by the Time and Frequency Division of NIST. The network service is realized with approximately 20 servers located at many locations in the USA. The servers receive approximately 340,000 requests per second for time information in a number of different formats. The Network Time Protocol (NTP) receives approximately 98% of these requests. In addition to the servers that transmit UTC time, we also operate a single server that transmits UT1 time in NTP format. The UT1 server implements the offset between UT1 and UTC based on the data in the International Earth Rotation and Reference Systems Service (IERS) Bulletin A. The accuracy of the extrapolation is better than 4 ms; the accuracy of the time received by a user will depend on the stability of the network connection and is generally better than 8 ms. I will also describe the plan to improve the accuracy of the time service. This upgrade should be completed in the next few months.

Keywords Leap Second • NTP • TA(NIST) • UTC • UTC(NIST) • UT1

Introduction

I will describe the clock ensemble that is used to generate the NIST time scales UTC(NIST) and TA(NIST). The atomic time scale, TA(NIST), was synchronized to TAI in 1978 when it was first implemented as a real-time time scale. It has been free-running since that time. The UTC(NIST) time scale is realized as an offset from TA(NIST). It is steered toward UTC as computed by the International Bureau of Weights and Measures (BIPM). The steering is accomplished by automated frequency adjustments to the output phase stepper as I will describe below in more detail. The clocks are not steered.

J. Levine (✉)
Time and Frequency Division, NIST, Boulder, CO 80305, USA
e-mail: Judah.Levine@nist.gov

E.F. Arias et al. (eds.), *The Science of Time 2016*, Astrophysics and Space Science Proceedings 50, DOI 10.1007/978-3-319-59909-0_17

The digital time services provided by NIST respond to requests in several different formats that are received over dial-up telephone lines and the public Internet. The servers currently receive approximately 340,000 requests per second or about 3×10^{10} requests per day. About 98% of these requests are in the Network Time Protocol (NTP) format (Mills 2011; Mills 1991), and the remainder are in several other older formats, which are supported for legacy applications. The number of requests continues to increase at a rate of about 4% per month, and most of this growth is in NTP-format requests.

The NIST Time Scales

The NIST time scale (Levine 2012) uses an ensemble of commercial high-performance cesium standards and hydrogen masers. The input to the ensemble calculation is an array of time-difference measurements between one of the clocks in the ensemble, which is designated as the reference clock for the measurement process, and all of the other clocks. The time differences are measured every 12 min, and the time scale is calculated as soon as each measurement cycle is completed. The reference clock for the measurement process is chosen for its reliability and stability and has no special role in the ensemble calculation. The clock selected to be the reference clock also has no special role in generating the UTC(NIST) output signal.

The time-difference measurements are realized with a dual-mixer system (Stein et al. 1982). The output from each clock at 5 MHz is mixed with a frequency of 5 MHz–10 Hz, which is synthesized from the reference clock of the measurement process. The time difference of each clock is measured as the phase differences between the 10 Hz beat frequency generated by the mixing process and the 10 Hz frequency generated by the reference clock channel. The down-conversion of the 5 MHz input signals to 10 Hz increases the resolution of the measurement process by a factor of 5×10^5, so that a measurement resolution of 10^{-7} of a cycle at 10 Hz translates into an effective resolution of less than 1 ps at the input frequency of 5 MHz. The stability and accuracy of the time-difference data is limited by the corresponding parameters of the front-end mixer and is typically a few ps. This performance is routinely verified by measuring the time difference between a single clock connected to two measurement channels. The differential nature of the time-difference measurements attenuates any common-mode variation in the characteristics of the measurement channels.

The atomic time scale TA(NIST) is computed as the weighted average of these time differences. The weight assigned to each clock is derived from its stability, which is calculated as the RMS residuals of the time difference between the time of the clock after its deterministic parameters have been removed and the time of the ensemble averaged over the previous 30 days. This method overestimates the stability of a clock, since the clock is also a member of the ensemble that is used to evaluate its stability. This correlation is attenuated by decreasing the weight of a high-weight clock (Tavella and Thomas 1991) and also by administratively limiting the weight of any clock to 30% of the ensemble. (Most time scale algorithms have similar limits.) Since the weights of all of the clocks must sum to 100%, limiting the weight of a good clock implicitly increases the weights of clocks that have poorer stability. Therefore, the administrative limitation on the maximum weight of any clock implicitly degrades the stability of the ensemble.

The frequency of TA(NIST) was set equal to the frequency of the International Atomic Time (TAI) frequency when the dual-mixer system was inaugurated, in 1978, and it has slowly drifted upward in frequency since that time. The current frequency of TA(NIST) is greater than the frequency of TAI by approximately 37 ns/day (a fractional frequency of approximately 4.5×10^{-13}).

The UTC(NIST) time scale is derived from TA(NIST) by an offset in time and frequency that is calculated administratively. The offset is a piecewise linear function that has no time steps. That is, changes to the offset equation are realized as changes in frequency only. The origin time of each of the linear steering equations is the ending time of the previous function. In other words, the origin time of each steering equation is the integral of all of the previous frequency offsets from 1978 when the system was initialized to the current time. The steering frequency has been negative for many years to remove the positive frequency offset of TA(NIST) that was described in the previous paragraph. The current and proposed future steering equations are published in the periodic Bulletin of the NIST Time and Frequency Division (NIST 2017a).

The steering adjustments are computed from the data in the BIPM Circular T (BIPM 2017a), which gives the difference between the UTC time scale computed by the BIPM and the time scales of all of the contributing laboratories (BIPM 2017b). The Circular T for any month is typically published about the 10th day of the following month. The BIPM also computes a more rapid version of UTC. The rapid version, UTCr (BIPM 2017c), provides daily estimates of the same data as in Circular T with a delay of 1 week. The results of the UTCr calculation are published on Wednesday afternoon (about 1700 UTC) for the period ending on the previous Monday. Starting in January, 2016, the data from UTCr are also used to compute a steering adjustment for UTC(NIST). This is done to improve the short-term accuracy of the UTC(NIST) time scale. The BIPM data for UTC-UTC(NIST) and UTCr-UTC(NIST) are shown in Fig. 17.1.

After each 12-min time scale computation, the software computes the offset in time and in frequency of UTC(NIST) with respect to TA(NIST) based on the appropriate steering equation and applies these parameters to a hardware phase stepper that produces the requested output signals at 5 MHz and 1 Hz. The reference signal for the phase stepper is derived from one of the clocks in the ensemble. The signal from the reference clock itself is not steered. The reference clock for the phase stepper has no special significance in the time scale algorithm, and is chosen mostly for its reliability and longevity. In general, the reference signal for the phase stepper is derived from one of the hydrogen masers in the ensemble, because these signals have the best short term frequency stability. The 5 MHz output signal from the phase stepper is fed back into the time scale measurement system, where it is

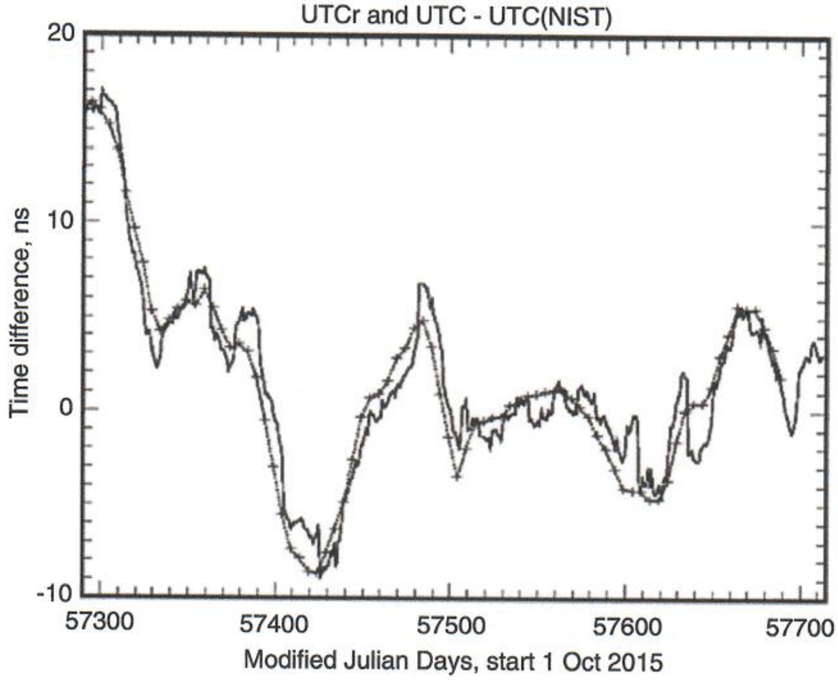

Fig. 17.1 The time differences in ns between the UTC and rapid UTCr time scales computed by the BIPM and UTC(NIST). The *dotted line* with the symbols shows the UTC time difference, and the *solid line* is the UTCr time difference. The *X* axis is in Modified Julian Days, and the plot shows the time differences from October 2015 through November 2016

measured along with all of the other clocks in the ensemble. However, the signal from the phase stepper is not included in the ensemble calculation, and its time difference with respect to the reference clock of the measurement system is used to verify that the requested steering offsets have been correctly applied. In normal operation, the difference between the physical output of the phase stepper and the equation that defines UTC(NIST) for that epoch is a few ps and is well characterized as white phase noise.

The output signals from the phase stepper are used as the reference signal for all of the NIST time services. They are also used to control the link between NIST and the BIPM through PTB (the Physikalisch-Technische Bundesanstalt, which is the National Metrology Institute of Germany). This link is used by the BIPM to compute the values for UTC-UTC(NIST) that are reported in the monthly Circular T and in the more rapid UTCr data. The link between NIST and PTB is implemented using the two-way message exchange protocol, and this method is also used to support the NIST digital time services. I will describe this method in the next section.

judah.levine@nist.gov

The Two-Way Method

Estimating the delay in transmitting messages from the source to the user is usually the most important and most difficult aspect of any time distribution system. The two-way method, in which the one-way delay is estimated as one-half of a measurement of the round-trip value, is widely used for this purpose.

There are a number of different realizations of the two-way method. In the version that is generally used to measure the delay in packet-switched networks, the local station transmits a time query to the remote system, and the remote system replies. The systems at both ends transmit continuously in the satellite implementation of the protocol. In both cases, the two transmission time stamps and the two received time stamps are combined to estimate the transmission delay and the time difference between the clocks at the two stations. The details of the calculation are described in the following paragraph.

Let T_{1s} be the time when system 1 transmits a message to system 2 as measured by the clock in system 1. System 2 receives the message at time T_{2r} and replies at time T_{2s}, where both of these times are measured by the clock in system 2. This second message is received back at system 1 at time T_{1r}, again measured by the clock in system 1. The round-trip travel time, as measured by system 1, is given by

$$\Delta = (T_{1r} - T_{1s}) - (T_{2s} - T_{2r}) \tag{17.1}$$

The first term is the total elapsed time for the message exchange, and the second term is the delay between when the second system received the message and when it replied. For systems that transmit continuously, the messages in both directions are combined so that the second term is positive. If the path delay in one direction is d, then the message transmitted at time T_{1s} arrives at the second system at time $T_{1s} + d$ as measured by the clock in system 1. The time of system 2 at that instant is T_{2r}, so that the time difference between systems 1 and 2 is estimated by

$$\delta t_{12} = (T_{1s} + d) - T_{2r} \tag{17.2}$$

If the path delay is symmetrical, then $d = \Delta/2$, and the time difference between systems 1 and 2 is given by

$$\delta t_{12} = \frac{T_{1s} + T_{1r}}{2} - \frac{T_{2s} + T_{2r}}{2} \tag{17.3}$$

The accuracy of the two-way method depends on the assumption that the delay is symmetric—that the inbound and outbound delays are equal. The inbound and outbound delays are comprised of two contributions: the delays through the station hardware at each end point and the delay through the channel itself. The delays through the station hardware can be measured in principle, and any asymmetries can be calibrated and removed administratively. The uncorrelated variations in the inbound and outbound delays in the station hardware must be minimized as much as possible. A residual admittance to fluctuations in the ambient temperature is a common problem, and variations in the ambient temperature must be minimized for that reason.

The symmetry of the channel delay is a more difficult problem, because it is usually not under the control of the operators at either station. The symmetry requirement is not too difficult to realize with simple physical circuits, but it is much more difficult to do so when the circuit is realized as a packet-based communications channel with intervening routers, switches, and gateways. The queueing delays in these elements often depend on the traffic load, especially when that load is asymmetric, as is often the case with systems that provide web-based services or streaming audio or video information. Even under the best of circumstances, an asymmetry of a few percent of the total round-trip delay is quite common in packet-switched networks, and this asymmetry often sets the limit on the accuracy of any distribution method that uses packet-switched networks such as the public Internet. Since the round-trip delay is often on the order of 100 ms in a packet-switched network, these networks typically provide accuracy that is often not much better than a few milliseconds. This limit has nothing to do with the format of the messages or the protocol that is used to transfer the information, and there is no advantage of one protocol over another to the extent that the accuracy is limited by this consideration.

It is possible to mitigate the impact of the asymmetric delay contributions of routers, switches, and gateways in the communications channel by special purpose hardware at each one of these network elements (IEEE 2008). The special-purpose hardware bridges the network element and measures the latency delay of a message passing through the device. The details of how the latency is measured vary somewhat from one implementation to another, but the important point is that this technique removes one of the primary limitations to the accuracy of packet-based time distribution. Unfortunately, these extra devices are not present in the general public Internet, so that the improvement in accuracy is not usually possible when the public Internet is used for time distribution.

Although the one-way delay in a link between two ground stations that uses a synchronous communications satellite link is quite large, the asymmetry is very small, so that time comparisons that use this method can realize sub-microsecond accuracy. The stability and reciprocity of the delay through the satellite transponder and the ground station equipment often limit the accuracy of the comparisons.

The NIST Time Servers

The Time and Frequency Division of NIST operates a number of network-based time servers that are synchronized to the NIST clock ensemble. The servers that are located at the NIST facilities in Colorado are synchronized by a direct hard-wired connection between the 1 pulse-per-second signal that realizes UTC(NIST) and the server hardware. The accuracy of the time at the server is a few microseconds. The servers that are located at other locations (see the Internet Time Service list at http://tf.nist.gov for a list of the locations) are synchronized by a telephone connection between the server and the NIST clock ensemble in Boulder. The connection uses the ACTS protocol (Levine et al. 1989), which transmits time over dial-up telephone lines. The delay in the telephone connection is measured by using the two-way protocol described above, and the accuracy of the time at the server is typically a few milliseconds.

The servers respond to time requests received over the public Internet in a number of different formats. The most common request is in the Network Time Protocol (NTP) format. This format implements the two-way method for estimating the transmission delay, and the accuracy of the messages is limited by the considerations that I discussed in the previous section.

There are three types of NTP servers. Most of the NIST servers are synchronized to UTC(NIST) and respond to requests by using the standard NTP message format. The second type of server also transmits time data based on UTC(NIST) but includes support for message authentication based on the MD5 (Rivest 2011) or SHA1 (National Institute of Standards and Technology 2015) hash codes. There are currently approximately 400 users of this service, and each one has a unique key. The key is combined with the standard NTP request and reply packets, and the hash function of the combination is computed by the sender and appended to the message. The receiver computes the hash function of the message and compares the computed result with the transmitted version. The message is accepted only if the computed hash function agrees with the one in the message. Since the key parameter is a secret known only to NIST and the single registered user, the authentication process ensures that the inbound and outbound messages were not sent by a rogue third party and were not modified in transmit. The authentication process does not affect the accuracy of the time transmitted in this way, since the accuracy is limited by symmetry of the path and is not affected by the contents of the message.

The third type of NTP server transmits UT1 time in the standard NTP format. The UT1 time is computed by adding the difference between UT1 and UTC published by the International Earth Rotation and Reference Service (IERS) in Bulletin A (IERS 2017). The bulletin is published every week and contains estimates of the UT1-UTC time difference for about a year in the future. The accuracy of the prediction for 30 days in the future is approximately 4 ms, and this is comparable to the accuracy of the NTP time transfer process as I have described above. Therefore, the values used by the time server are updated every few weeks, and the longer-term predictions in each bulletin are never used.

There are currently 24 registered users of this service. The service was originally limited to registered users so that the general user, who may not understand the difference between UTC and UT1, will not connect to the server by mistake and receive a time message that can be "wrong" by up to 0.9 s. The server was converted to open access in September 2016. It currently receives requests from about 900 unique network addresses.

In addition to responding to requests in the NTP format, the servers also respond to requests in the DAYTIME (TFC 1983) and TIME protocols (Postel 1983). These are simpler protocols that do not support the full two-way message exchange, and they are used when the client hardware is relatively simple and when the highest accuracy is not required. The NIST implementation of the DAYTIME protocol includes advance notice of the transitions to and from daylight saving time based on the US model and also provides advance notice of leap seconds. The daylight saving time information is particularly useful for older systems that may not have any other source for this information or that may have the older, incorrect US transition dates. The TIME protocol is the simplest of all of the message formats and is supported for legacy applications only. We do not recommend it for new applications because it has poor error-checking capabilities, supports no method for estimating the network delay, and provides none of the ancillary information of the DAYTIME format.

The NIST time servers receive approximately 340,000 requests per second for time in the various formats, and almost all of these requests (>98%) are in the NTP format. The number of requests has been increasing at a rate of approximately 4% per month, and this rate of increase shows no sign of moderating. The increase in the number of requests since 2004 is shown in Fig. 17.2. The figure shows that the number of requests in the DAYTIME and TIME formats has been decreasing since 2012, but there are still approximately 10^8 requests per day in these formats or about 1000 requests per second.

Traceability

The concept of traceability is important for all time services. A time measurement is traceable if there is an unbroken chain of measurements from the end user back to a national or international standard of time (Taylor and Kuyatt 1994). Each link in the measurement chain must be characterized by a delay and an uncertainty. In the USA, traceability can be satisfied by a chain that leads back to UTC(NIST) or UTC (USNO).

In addition to the technical traceability requirement of the previous paragraph, some users are required to have legal traceability. In addition to the technical requirements for traceability, legal traceability adds requirements that can confirm technical traceability in a legal adversary proceeding. The details of how legal traceability is satisfied depend on the end user but generally include properly maintained and certified log files and other documentations that can confirm the proper operation of the synchronization process. A log file that contains only error information may be inadequate, because it can be completely empty when the system is working properly, and it can be hard to distinguish this case from a system that is totally inoperative.

Although a time signal may be generated by a traceable source such as a GPS satellite or a NIST-operated time server, the traceability of these signals generally

judah.levine@nist.gov

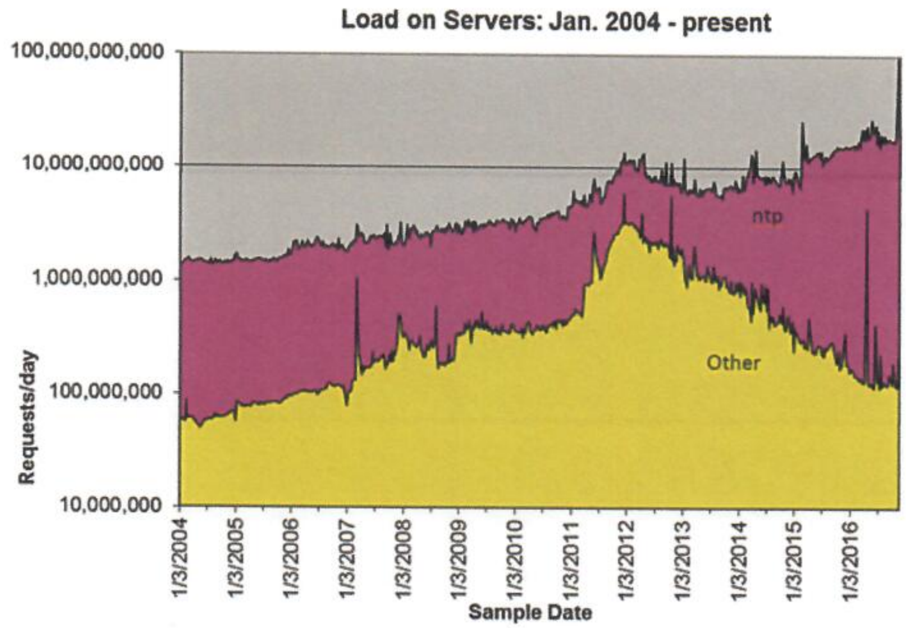

Fig. 17.2 The number of time requests received by all of the NIST time servers from January 2004 to November 2016. The logarithm of the number of requests as a function of time is displayed. The *lower curve*, labeled "other," is the number of requests in the TIME and DAYTIME formats, while the *upper curve* is the total number of requests, so that the difference is the number of requests in NTP format

does not extend to un-calibrated and unmonitored user equipment. That is, the traceability of a signal from a GPS satellite ends with the signal in space by default, and the default traceability of the time from a NIST time server ends at the server portal. The messages transmitted by third-party servers are generally not traceable.

The Leap Second Problem

The current definitions of the length of the second in both the UTC and TAI time scales are based on the same value derived from the frequency of a hyperfine transition in the ground state of cesium 133 (BIPM 2006). The length of the second that results from this definition is somewhat shorter than the length of the second of the UT1 time scale, which is derived from the rate of rotation of the Earth. The difference between the two time scales has a systematic contribution whose magnitude is currently somewhat less than 1 s per year. This systematic increase is caused by tidal friction and a number of other, smaller effects. (Nelson et al. 2001) There is also an additional contribution that is irregular and whose magnitude cannot be accurately characterized either deterministically or stochastically. The stochastic contribution to the length of the UT1 second could reverse the sign of the UTC-UT1 difference in principle, but this has not happened to date and is unlikely to happen in the future. The UTC time scale is used as a proxy for UT1 time in many applications ranging from every day timekeeping to the timing of events in astronomy, and maintaining the equivalence between UT1 and UTC is an important consideration in timekeeping.

In order to prevent the UTC and UT1 time scales from diverging, an extra "leap" second is added to UTC whenever the magnitude of the difference between UTC and UT1 approaches 0.9 s (NIST 2017b). (As I mentioned in the previous paragraph, it could be necessary in principle to drop a second in UTC to preserve this equivalence, but this is unlikely in the foreseeable future.) This extra second is added after 23:59:59 UTC, usually on the last day of June or December. (The last days of other months are possible in principle but have never been used.) The name of this extra second is 23:59:60 UTC. The leap second is defined with respect to UTC, and it therefore occurs at different local times at different locations. For example, it occurs late in the afternoon in California and during the following morning in Asia and Australia. The second after the leap second is named 00:00:00 UTC of the next day.

Computer clocks and most digital systems keep time internally as the number of seconds that have elapsed since an origin time, which can differ from one system to another. The conversion between the count of the elapsed seconds and the time in the conventional year-month-day hour-minute-second representation is computed by the operating software whenever the time is requested. Additional offsets for the extra day associated with leap years, for conversion to the local time zone, and for daylight saving time are also added if necessary. The various offsets are either strict constants (as for the offset of the local time zone) or can be computed algorithmically (as for the offsets for a leap year or for the local version of daylight saving time). Both leap days and the daylight saving time transitions introduce discontinuities between a time interval computed from the count of the elapsed seconds maintained by the system internally and a simple computation of the same time interval using the hour-minute-second representation. This discontinuity can be particularly serious during the Fall transition away from daylight saving time, since the local time moves backward by 1 h during the transition and the same time stamps will be transmitted twice. There is often no indication whether the time reported by the system during this transition period is the old daylight time or the new standard time. The impact of these changes is mitigated by the fact that the change is made at 2 a.m. Sunday local time.

The method of timekeeping in a digital system is therefore divided into two cooperating processes: a low-level application that receives and counts the periodic signals generated by the hardware and a second higher-level application that "knows" how to convert this count to the conventional representation of hours-minutes-seconds. In general, the partition between the low-level and high-level processes is arbitrary because the application need not be concerned with the intermediate values transmitted between them. However, leap seconds are inserted on an irregular and unpredictable basis, and the times of both future and past leap seconds cannot be computed algorithmically. Therefore, there is no simple way of incorporating leap seconds into this deterministic pair of cooperating processes.

judah.levine@nist.gov

The solution used by the NIST time servers and by many other systems is to stop the hardware clock for 1 s at 23:59:59 UTC on the day of a leap second. The result is to use the time 23:59:59 UTC twice—once when that time occurs and a second time during the leap second. Neither the system applications nor the standard communications protocols have any provision for indicating that the second 23:59:59 time value is actually a leap second and should be assigned the time 23:59:60. This solution has a number of undesirable side effects. The fact that there are 2 s with the same name introduces confusion and can produce time stamps that appear to violate causality: 23:59:59.2 during the leap second actually occurred after 23:59:59.5 during the first second with this time stamp value, but the time stamps themselves indicate the opposite order. (The Fall transition from daylight saving time to standard time has the same problem with time stamps between 2 a.m. and 3 a.m.)

This method of repeating 23:59:59 is probably the closest approximation to the definition of the leap second on systems that have no method of indicating the leap second event to applications and inserting a special leap second indicator into communications protocols. The NIST DAYTIME protocol transmits time in the hour-minute-second format and transmits the leap second time correctly as 23:59:60.

Although the method I have just described has undesirable side effects, there are other implementations that are worse. For example, a number of network time-service appliances add the leap second to the first second of the new day. That is, the time stamps in the vicinity of a leap second are 23:59:58, 23:59:59, 00:00:00, and 00:00:00. The long-term behavior of this method is correct, but it adds the leap second to the wrong day and has a transient error of 1 s with respect to the definition. Other systems implement the leap second as a frequency adjustment over some time interval before (or after) the leap second (https://googleblog.blogspot.com/2011/09/time-technology-and-leaping-seconds.html). This method introduces a varying time error on the order of 0.5 s with respect to the definition and a frequency error that depends on the interval over which the leap second is inserted. Since the interval for amortizing the leap second is not defined in any standard, different implementations of this method may produce time stamps in the vicinity of a leap second that disagree both with each other and with respect to UTC.

Finally, it is clear that all of these solutions introduce steps in both time interval and frequency in the vicinity of a leap second and that a real-time physical process does not stop during the leap second. The time scales of navigation systems do not use leap seconds for this reason, although each navigation system time typically includes the number of leap seconds that have been inserted from the start of the leap second system in 1972 to the instant that the navigation system time scale was initialized. Each navigation system time will therefore have a different offset from UTC for this reason, and all of these offsets will change whenever a new leap second is inserted. This multiplicity of time offsets from UTC is undesirable because it can be a source of confusion when time stamps are derived from the transmissions of navigation satellites without a clear indication of which time scale was used.

A number of proposals have been advanced for addressing the leap second problem, but a detailed discussion of them is outside of the scope of this document, which is focused on the NIST time services.

Improvements to the Service

The Time and Frequency Division of NIST is addressing the continued increase in the number of requests for time services by upgrading the capacity of the servers and the network connections. We plan to have a number of very high capacity sites operating in the next few months. In addition, we are adding reference clocks to each of these sites to improve the accuracy of the time at the servers. We expect that the reference clocks at the remote sites will be within 10–20 ns of UTC(NIST), and we will consider offering services that can transmit the time to a user's system with an end-point accuracy of 100 ns or better.

Summary and Conclusions

The NIST Time and Frequency Division operates an ensemble of time servers that respond to requests in a number of different standard formats. The servers receive approximately 340,000 requests per second, most of which are in the NTP format. The time of the servers is directly traceable to the UTC(NIST) time scale, which is maintained by an ensemble of atomic clocks and hydrogen masers located at the NIST laboratory in Boulder, Colorado.

NIST operates a single server that transmits UT1 time in the standard NTP format. The internal time of the server is synchronized to UTC(NIST), and the difference between UT1 and UTC is added to the NTP response just before it is transmitted to the system that requested it. The UT1-UTC time difference is derived from the International Earth Rotation and Reference System Service (IERS) Bulletin A, and the value of the UT1-UTC difference added to each message is calculated by a simple linear interpolation of the previous and following values in Bulletin A. The server is open access and currently receives requests from about 900 unique network addresses.

NIST also operates three time servers that authenticate the time stamps by computing the MD5 or SHA1 hash of the combination of a secret key and the standard time message. Each registered user of the authenticated service is given a unique key. There are currently approximately 400 users of this service. The authentication process ensures that the message originated from NIST and was not modified in transit. It does not affect the accuracy of the time service, which is limited by the symmetry of the network connection between the client and the server.

judah.levine@nist.gov

The time servers implement a leap second by transmitting the time corresponding to 23:59:59 a second time during the leap second event. This method is used because neither the system nor the NTP message format can represent the correct name of the leap second, which is 23:59:60. The DAYTIME format uses the hour-minute-second representation and transmits the correct representation of the leap second. There is no discontinuity in the time transmitted by the UT1 time server across the leap second.

References

BIPM, *Le Systeme International d'Unites*, 8th edn. (BIPM, Sevres, 2006)

BIPM, Published monthly by the International Bureau of weights and measures and available from the BIPM web site at ftp://ftp2.bipm.org/pub/tai/publication/cirt/ (2017a)

BIPM, Details and data for all contributing laboratories are available on the BIPM web page at www.bipm.org (2017b)

BIPM, The weekly UTCr data for all of the contributing laboratories are published in tables on the BIPM web page at: ftp://ftp2.bipm.org/pub/tai/publication/utcr. All of the data for each contributing laboratory is at: ftp://ftp2.bipm.org/pub/tai/publication/utclab (2017c)

IEEE, IEEE Standard for a Precision clock synchronization protocol for networked measurement and control systems, IEEE Standard 1588, available from: webstore.ansi.org. 2008

IERS, International Earth Rotation and Reference Service, http://datacenter.iers.org/eop/-/somos/5Rgv/latest/6 (2017)

J. Levine, The statistical modeling of atomic clocks and the design of time scales. Rev. Sci. Instrum. **83**, 012201–012228 (2012)

J. Levine, M. Weiss, D.D. Davis, D.W. Allan, D.B. Sullivan, The NIST automated computer time service. J. Res. NIST **94**, 311–322 (1989)

D.L. Mills, RFC 1305, available from: www.ietf.org/rfc/rfc1305.txt. 1991

D.L. Mills, *Computer Network Time Synchronization*, 2nd edn. (CRC Press, Boca Raton, 2011), p. 64

National Institute of Standards and Technology, Federal Information Processing Standard FIPS Pub 180-4, 2015

R.A. Nelson, D.D. McCarthy, S. Malys, J. Levine, B. Guinot, H.F. Fliegel, R.L. Beard, T.R. Bartholomew, The leap second: Its history and possible future. Metrologia **38**, 509–529 (2001)

NIST, The Steering data is available online at http://www.nist.gov/pml/div688/grp50/nistat1.cfm (2017a)

NIST, Leap Seconds are described at www.nist.gov/pml/div688/leapseconds.cfm. See also: www.itu.int/en/ITU-R/Pages/default.aspx (2017b)

J. Postel, The Time Protocol, RFC 866, www.rfc-editor.org/info/rfc866, 1983

R. Rivest, The MD5 Message-Digest Algorithm, Request for comments RFC-1321, 1992. Updated in request for comments RFC-6151, 2011. The documents are available from www.ietf.org/rfc, 2011

S. Stein, D. Glaze, J. Levine, J. Gray, D. Hilliard, D. Howe, L. Erb, Performance of an automated high accuracy phase measurement system, in *Proceedings of the 36th Annual Symposium on Frequency Control*, Fort Monmouth, New Jersey, 1982. Reprinted in NIST Technical Note 1337, available from the publications list at tf.nist.gov

P. Tavella, C. Thomas, Comparative study of time scale algorithms. Metrologia **28**, 57–63 (1991)

B.N. Taylor, C.E. Kuyatt, *Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results, NIST Technical Note 1297* (U. S. Government Printing Office, Washington, D.C., 1994)

TFC, The NIST version of the DAYTIME protocol is described on internet time service page at: tf.nist.gov. See also, J. Postel, www.tfc-editor.org/info/rfc867, 1983

judah.levine@nist.gov