

On the Chronometry and Metrology of Computer Network Timescales and their Application to the Network Time Protocol^{1,2,3}

David L. Mills
Electrical Engineering Department
University of Delaware

Abstract

This paper summarizes issues in computer network timekeeping with respect to the Network Time Protocol, which is used to synchronize time in many of the hosts and gateways of the Internet. It describes the methods used to coordinate and disseminate international time services and how they are incorporated into NTP time servers. It discusses the hazards on reckoning NTP dates with the conventional civil calendar and a standard method for numbering the days. Finally, the paper describes the NTP timescale, its mapping to conventional civil time and date and its interpretation of leap seconds.

Keywords: network clock synchronization, standard time distribution, timescale coordination, leap seconds, Network Time Protocol.

1. Introduction

Each year brings exotic new applications for precision timekeeping, such as measuring the density of the universe [RAW87] and calibrating global warming [SCI91]. Recent improvements in time-transfer technology have doubled the precisions attainable about every seven years to the current regime of better than a few nanoseconds over the global range. While inexpensive, ubiquitous time transfer using a global computer network ordinarily cannot approach such precisions, accuracies of a few tens of milliseconds have been demonstrated for most paths in the Internet of today [MIL90a].

Among the requirements for ubiquitous timekeeping in a global computer networking community is to provide a standard chronometry for precision dating of present and future events. This requires a standard interpretation of computer timekeeping media with respect to conventional civil time and date as disseminated by national means. Among the problems addressed must be the rationalization with the calendar, the ambiguity of presentation and the interpretation of leap seconds.

This paper consists of an extended discussion on *computer network chronometry*, which is the precise deter-

mination of computer time and frequency relative to international standards, and on *calendar metrology*, which is the determination of conventional civil time and date according to the modern calendar. It describes the methods conventionally used to establish civil time and date and the various timescales now in use. In particular, it characterizes the Network Time Protocol (NTP) timescale relative to the Coordinated Universal Time (UTC) timescale, and establishes the precise interpretation of UTC leap seconds in NTP.

In the following discussion the terms *time*, *timescale*, *oscillator*, *clock*, *epoch*, *timestamp*, *calendar* and *date* are used in a technical sense. Strictly speaking, the *time* of an event is an abstraction which determines the ordering of events in some given frame of reference or *timescale*. An *oscillator* is a generator capable of precise frequency (relative to the given timescale) within a specified tolerance. A *clock* is an oscillator together with a counter which records the (fractional) number of cycles since being initialized with a given value at a given time. The value of the counter at any given time t is called its *epoch* and recorded as the *timestamp* $T(t)$ of that epoch. In general, epoches are not continuous and depend on the precision of the counter.

A *calendar* is a mapping from epoches in some timescale to the year-day *dates* used in everyday life. Since multiple calendars are in use today and sometimes disagree on

1. Sponsored by: Defense Advanced Research Projects Agency under NASA Ames Research Center contract number NAG 2-638 and National Science Foundation grant number NCR-89-13623.
2. Author's address: Electrical Engineering Department, University of Delaware, Newark, DE 19716; Internet mail: mills@udel.edu.
3. Reprinted from: Mills, D.L. On the chronology and metrology of computer network timescales and their application to the Network Time Protocol. *ACM Computer Communications Review* 21, 5 (October 1991), 8-17.

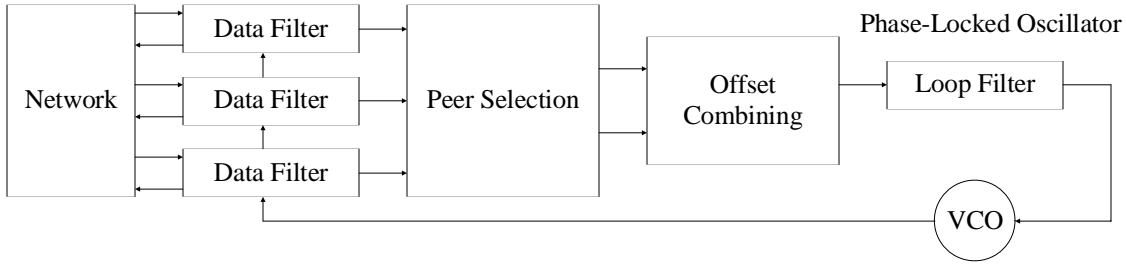


Figure 2. Network Time Protocol




Figure 1. Subnet Synchronization

the dating of the same epoches in the past, the metrology of past and present epoches is an art practiced by historians. However, the ultimate timescale for our world is based on cosmic oscillators, such as the Sun, Moon and other galactic orbiters. Since the frequencies of these oscillators are relatively unstable and not known exactly, the ultimate reference standard oscillator has been chosen by international agreement as a synthesis of many observations of an atomic transition of exquisite stability. The frequency of each heavenly and Earthbound oscillator defines a distinctive timescale, not necessarily always continuous, relative to that of the standard oscillator.

In this paper the *stability* of a clock is how well it can maintain a constant frequency, the *accuracy* is how well its time compares with national standards and the *precision* is to what degree time can be resolved in a particular timekeeping system. The *time offset* of two clocks is the time difference between them, while the *frequency offset* is the frequency difference between them. In this paper reference to simply “offset” means time offset, unless indicated otherwise. The *reliability* of a timekeeping system is the fraction of the time it can be kept connected to the network and operating correctly relative to stated accuracy and stability tolerances.

In order to synchronize clocks, there must be some way to directly or indirectly compare them in time and frequency. In network architectures such as DECnet and Internet local clocks are synchronized to designated *time servers*, which are timekeeping systems belonging to a *synchronization subnet*, in which each server measures the offsets between its local clock and the clocks of its neighbor servers or *peers* in the subnet. In this paper to *synchronize frequency* means to adjust the clocks in the

subnet to run at the same frequency, to *synchronize time* means to set them to agree at a particular epoch with respect to Coordinated Universal Time (UTC), as provided by national standards, and to *synchronize clocks* means to synchronize them in both frequency and time.

2. Network Time Protocol

The Network Time Protocol (NTP) is used by Internet time servers and their peers to synchronize clocks, as well as automatically organize and maintain the time synchronization subnet itself. It is evolved from the Time Protocol [POS83] and the ICMP Timestamp Message [DOD81a], but is specifically designed for high accuracy, stability and reliability, even when used over typical Internet paths involving multiple gateways and unreliable networks. The following sections contain an overview of the procedures and algorithms used in NTP. A formal description and error analysis of the protocol is contained in [MIL90b]. A detailed description of the NTP architecture and protocols is contained in [MIL91], while a summary of operational experience and performance is contained in [MIL90a].

NTP and its implementations have evolved and proliferated in the Internet over the last decade, with NTP Version 2 now adopted as an Internet Standard (Recommended) [MIL89] and NTP Version 3 now adopted as a Proposed Standard [MIL90b]. NTP is built on the Internet Protocol (IP) [DOD81b] and User Datagram Protocol (UDP) [POS80], which provide a connectionless transport mechanism; however, it is readily adaptable to other protocol suites. The protocol can operate in several modes appropriate to different scenarios involving private workstations, public servers and various network configurations. A lightweight association-management capability, including dynamic reachability and variable poll-interval mechanisms, is used to manage state information and reduce resource requirements. Optional features include message authentication based on crypto-checksums and provisions for remote control and monitoring.

In NTP one or more primary servers synchronize directly to external reference sources such as radio clocks. Secondary time servers synchronize to the primary servers

and others in the synchronization subnet. A typical subnet is shown in Figure 1a, in which the nodes represent subnet servers, with normal level numbers determined by the hop count from the root (stratum one), and the heavy lines the active synchronization paths and direction of timing information flow. The light lines represent backup synchronization paths where timing information is exchanged, but not necessarily used to synchronize the local clocks. Figure 1b shows the same subnet, but with the line marked x out of service. The subnet has re-configured itself automatically to use backup paths, with the result that one of the servers has dropped from stratum 2 to stratum 3. In practice each NTP server synchronizes with several other servers in order to survive outages and Byzantine failures using methods similar to those described in [SHI87].

Figure 2 shows the overall organization of the NTP time-server model, which has much in common with the phase-lock methods summarized in [RAM90]. Timestamps exchanged between the server and possibly many other subnet peers are used to determine individual roundtrip delays and clock offsets, as well as provide reliable error bounds. Figure 3 shows how NTP timestamps are numbered and exchanged between peers A and B . Let $T_i, T_{i-1}, T_{i-2}, T_{i-3}$ be the values of the four most recent timestamps as shown and let

$$a = T_{i-2} - T_{i-3} \quad \text{and} \quad b = T_{i-1} - T_i.$$

If the network delays from A to B and from B to A are similar, the roundtrip delay δ and clock offset θ of B relative to A at time T_i are:

$$\delta = a - b \quad \text{and} \quad \theta = \frac{a + b}{2}.$$

Each NTP message includes the latest three timestamps T_{i-1}, T_{i-2} and T_{i-3} , while the fourth timestamp T_i is determined upon arrival of the message. Thus, both the server and the peer can independently calculate delay and offset using a single bidirectional message stream. This is a symmetric, continuously sampled, time-transfer scheme similar to those used in some digital telephone networks [LIN80]. Among its advantages are that the transmission times and received message orders are unimportant and that reliable delivery is not required.

As shown in Figure 2, the computed delays and offsets for each peer are processed by the data-filter algorithm to reduce incidental timing noise. As described in the [MIL90b], this algorithm selects from among the last several samples the one with minimum δ and presents the associated θ as the output. The peer-selection algorithm determines from among all peers a suitable subset of peers capable of providing the most accurate and trustworthy time using principles similar to those described in [VAS88]. In NTP this is done using a cascade

Oscillator type	Stability (per day)	Drift /Aging
Hydrogen Maser	2×10^{-14}	$1 \times 10^{-12}/\text{yr}$
Cesium Beam	3×10^{-13}	$3 \times 10^{-12}/\text{yr}$
Rubidium Cell	5×10^{-12}	$3 \times 10^{-11}/\text{mo}$
Quartz (Oven)	1×10^{-9} 0-50 deg C	$1 \times 10^{-10}/\text{day}$
Quartz (Digital)	5×10^{-8} 0-60 deg C	$1 \times 10^{-9}/\text{day}$
Quartz (Analog)	5×10^{-7} 0-60 deg C	$3 \times 10^{-9}/\text{day}$
Quartz	$\sim 1 \times 10^{-6}$ per deg C	don't ask

Table 1. Characteristics of Standard Oscillators




Figure 3. Measuring Delay and Offset

of two subalgorithms, one a version of an algorithm proposed in [MAR85] and the other based on maximum likelihood principles to improve accuracy [MIL91].

The resulting offsets of this subset are first combined on a weighted-average basis using an algorithm similar to that described in [JON83] and then processed by a phase-lock loop (PLL). In the PLL the combined effects of the filtering, selection and combining operations are to produce a phase-correction term, which is processed by the loop filter to control the local clock, which functions as a voltage-controlled oscillator (VCO). The VCO furnishes the timing (phase) reference to produce the timestamps used in all timing calculations.

3. Methods for Time and Frequency Synchronization

The primary servers of a synchronization subnet must themselves synchronize to a source of standard time, such as a radio clock or telephone modem. It is of considerable interest to explore how this can be done and what accuracies can be expected. The following sections discuss issues involved in providing accurate and stable reference sources using national means of dissemination.

3.1. Time and Frequency Standards

A *primary frequency standard* is an oscillator that can maintain extremely precise frequency relative to a physical phenomenon, such as a transition in the orbital states of an electron. Presently available atomic oscillators are based on the transitions of the hydrogen, cesium and rubidium atoms. Table 1, which was compiled from published data, shows the characteristics for typical oscillators of these types compared with those for various

types of temperature-compensated and uncompensated quartz-crystal oscillators found in electronic equipment. Present practice is to operate multiple cesium-based clocks at national standards laboratories and coordinate their readings using satellite-based time-transfer methods. On the other hand, local clocks used in computing equipment almost always are designed with uncompensated quartz oscillators.

For the three atomic oscillators listed in Table 1 the Drift/Aging column shows the maximum offset from nominal standard frequency due to systematic mechanical and electrical characteristics. In the case of quartz oscillators (oven-controlled, digital-temperature-compensated, analog-temperature-compensated or uncompensated) this offset is not constant, which results in a gradual change in frequency with time, called aging. Even if a quartz oscillator is temperature compensated by some means, it must be periodically compared to a primary standard in order to maintain the highest accuracy. For all types of oscillators the stability column shows the maximum variation in frequency per day due to circuit noise and environmental factors.

As the telephone networks of the world are evolving rapidly to digital technology, consideration should be given to the methods used for synchronization in digital networks [BEL86]. A network of clocks in which each oscillator is phase-locked to a single frequency standard is called *isochronous*, while a network in which some oscillators are phase-locked to different master oscillators, but with the master oscillators closely synchronized in frequency (not necessarily phase-locked), to a single frequency standard is called *plesiochronous*. The NTP synchronization subnet operating in the Internet is a plesiochronous system in which multiple primary time servers synchronize using radio clocks and national broadcast services of different countries, which are not necessarily phase-locked to each other.

The industry has agreed on a classification of clock oscillators as a function of minimum accuracy, minimum stability and other factors [ALL74]. There are three factors which determine the classification: stability, jitter and wander. Stability refers to the systematic variation of frequency with time and is synonymous with aging, drift, trends, etc. Jitter (also called timing jitter) refers to short-term variations in frequency with components greater than 10 Hz, while wander refers to long-term variations in frequency with components less than 10 Hz. The classification determines the oscillator stratum (not to be confused with the NTP stratum described previously), with the more accurate oscillators assigned the lower strata and less accurate oscillators the higher strata, as shown in Table 2.

Stratum	Min Accuracy (per day)	Min Stability (per day)
1	1×10^{-11}	not specified
2	1.6×10^{-8}	1×10^{-10}
3	4.6×10^{-6}	3.7×10^{-7}
4	3.2×10^{-5}	not specified

Table 2. Stratum Assignments

The construction, operation and maintenance of stratum-one oscillators is assumed to be consistent with national standards and often includes cesium oscillators or precision quartz oscillators synchronized to national standards. Oscillators assigned higher strata represent the stability required for interexchange systems, exchange switches and PBX systems, respectively. With respect to this classification, most computer local clocks would be assigned stratum four. These clocks are most often constructed using an uncompensated quartz oscillator with typical frequency tolerance in the order of 10^{-4} and stability of 10^{-6} per day for equipment operating continuously at room temperature.

3.2. Time and Frequency Dissemination

In order that atomic and civil time can be coordinated throughout the world, national administrations operate primary time and frequency standards and coordinate them cooperatively by observing various radio and satellite broadcasts and through occasional use of portable atomic clocks. Most seafaring nations of the world operate some sort of broadcast time service for the purpose of calibrating chronographs, which are used in conjunction with ephemeris data to determine navigational position. In many countries the service is primitive and limited to seconds-pips broadcast by marine communication stations at certain hours. For example, a chronograph error of one second represents a longitudinal position error of about 0.23 nautical mile at the Equator.

The U.S. National Institute of Standards and Technology (NIST - formerly National Bureau of Standards) operates three radio services for the dissemination of primary time and frequency information. One of these uses high-frequency (HF or CCIR band 7) transmissions from Fort Collins, CO (WWV), and Kauai, HI (WWVH). Signal propagation is usually by reflection from the upper ionospheric layers, which vary in height and density throughout the day and season and result in unpredictable amplitude and delay variations at the receiver. While these transmissions and those of Canada from Ottawa, Ontario (CHU), and other countries can be received over large areas in the western hemisphere, reliable frequency comparisons can be made only to the order of 10^{-7} and time accuracies are limited to the order of a millisecond [BLA74]. Available radio clocks which operate with

these transmissions provide accuracies to the order of ten milliseconds and are priced in the \$1,500 range.

A second service operated by NIST uses low-frequency (LF or CCIR band 5) transmissions from Boulder, CO (WWVB), and can be received over the continental U.S. and adjacent coastal areas. Signal propagation is via the lower ionospheric layers, which are relatively stable and have predictable diurnal variations in height. With appropriate receiving and averaging techniques and corrections for diurnal and seasonal propagation effects, frequency comparisons to within 10^{-11} are possible and time accuracies of from a few to 50 microseconds can be obtained [BLA74]. Some countries in western Europe operate similar services which use transmissions from Rugby, U.K. (MSF). Mainflingen, Germany (DCF77), and Allouis, France (France Inter). Available radio clocks which operate with these transmissions provide accuracies to the order of a millisecond and are priced in the \$1,500 range.

The third service operated by NIST uses ultra-high frequency (UHF or CCIR band 9) transmissions from the Geosynchronous Orbit Environmental Satellites (GOES), and can be received over most of the western hemisphere. Available radio clocks which operate with these transmissions provide accuracies to the order of a millisecond and are priced in the \$6,000 range.

Another service offered by NIST uses an ordinary telephone and modem [NBS88]. The service is intended for use by personal workstations to set clock-calendars, for example. However, in order to maintain an accuracy of a few milliseconds using an uncompensated quartz oscillator at typical limits of tolerance and without frequency compensation, it would be necessary to make a long-distance call every few minutes, which would not be suitable for a large population of clients calling on a regular basis without further redistribution.

The U.S. Department of Defense is developing the Global Positioning System (GPS) for worldwide precision navigation. This system will eventually provide 24-hour worldwide coverage using a constellation of 21 satellites in 12-hour orbits. For time-transfer applications GPS has a potential accuracy in the order of a few nanoseconds; however, various considerations of defense policy may limit accuracy to hundreds of nanoseconds [VAN84]. Available radio clocks provide accuracies to 100 ns and are priced in the \$10,000 range.

There are several other navigation systems such as LORAN-C, operated by the U.S. Coast Guard and agencies of other countries, OMEGA, operated by the U.S. Navy, and various communication systems using very-low-frequency (VLF or CCIR band 4) transmissions. These systems can provide accuracies in the order from less than one to about 50 microseconds; however, the broad-

cast formats of these systems are not well suited for standard time distribution and the receivers have to be initialized with position information before use.

Note that not all transmission formats used by the various radio broadcast services [NBS79] and not all currently available radio clocks include provisions for year information and leap-second warning. This information must be determined from other sources. NTP includes provisions to distribute advance warnings of leap seconds using the leap-indicator bits described in the NTP specification. The protocol is designed so that these bits can be set manually or by the radio clocks at the primary time servers and then automatically distributed throughout the synchronization subnet to all other time servers.

4. Calendar Metrology

In the simplest terms a calendar system is a method to rationalize sightings of the Sun and Moon with certain religious and agricultural seasons which recur at characteristic frequencies. The various metrologies of the calendar amount to systems to assign day-numbers and sometimes day-names which determine the periodicity of these holy and secular oscillators. One might ask whether a (suitably disambiguated) NTP timestamp documenting the adventures of Alexander the Great merits the precision implied in its format. The following sections consider this issue in the light of the Western civil calendar and the establishment of a standard day-numbering plan.

4.1. Evolution of the Calendar

The calendar systems used in the ancient world reflect the agricultural, political and ritual needs characteristic of the societies in which they flourished. Astronomical observations to establish the winter and summer solstices were in use three to four millennia ago [CAL86]. By the 14th century BC the Shang Chinese had established the solar year as 365.25 days and the lunar month as 29.5 days. The lunisolar calendar, in which the ritual month is based on the Moon and the agricultural year on the Sun, was used throughout the ancient Near East (except Egypt) and Greece from the third millennium BC. Early calendars used either thirteen lunar months of 28 days or twelve alternating lunar months of 29 and 30 days and haphazard means to reconcile the 354/364-day lunar year with the 365-day vague solar year.

The ancient Egyptian lunisolar calendar had twelve 30-day lunar months, but was guided by the seasonal appearance of the star Sirius (Sothis). In order to reconcile this calendar with the solar year, a civil calendar was invented by adding five intercalary days for a total of 365 days. However, in time it was observed that the civil year was about one-fourth day shorter than the actual solar year and thus would precess relative to it over a 1460-

year cycle called the Sothic cycle. Along with the Shang Chinese, the ancient Egyptians had thus established the solar year at 365.25 days, or within about 11 minutes of the present determination. In 432 BC, about a century after the Chinese had done so, the Greek astronomer Meton calculated there were 110 lunar months of 29 days and 125 lunar months of 30 days for a total of 235 lunar months in 6940 solar days, or just over 19 years. The 19-year cycle, called the Metonic cycle, established the lunar month at 29.532 solar days, or within about two minutes of the present determination.

The Roman republican calendar was based on a lunar year and by 50 BC was eight weeks out of step with the solar year. Julius Caesar invited the Alexandrian astronomer Sosigenes to redesign the calendar, which led to the adoption in 46 BC of the Julian calendar. This calendar is based on a year of 365 days with an intercalary day inserted every four years. However, for the first 36 years an intercalary day was mistakenly inserted every three years instead of every four. The result was 12 intercalary days instead of nine, and a series of corrections that was not complete until 8 AD.

The seven-day Sumerian week was introduced only in the fourth century AD by Emperor Constantine I. During the Roman era a 15-year census cycle, called the Indiction cycle, was instituted for taxation purposes. The sequence of day-names for consecutive occurrences of a particular day of the year does not recur for 28 years, called the solar cycle. Thus, the least common multiple of the 28-year solar cycle, 19-year Metonic cycle and 15-year Indiction cycle results in a grand 7980-year supercycle called the Julian Era, which began in 4713 BC. A particular combination of the day of the week, day of the year, phase of the Moon and round of the census will recur beginning in 3268 AD.

By 1545 the discrepancy in the Julian year relative to the solar year had accumulated to ten days. In February 1582, following a plan originally devised by the astronomer Luigi Lilio, Pope Gregory XIII issued a papal bull which decreed that the day following Thursday, 4 October 1582 on the old (Julian) calendar would be Friday, 15 October 1582 on the new (Gregorian) calendar and, furthermore, that the solar year would henceforth consist of 365.2425 days [MOY82]. In order to realize the new value, only those centennial years divisible by 400 would be leap years, while the remaining centennial years would not, making the new year within about 24 seconds of the actual solar year at that time. Since the beginning of the Common Era and prior to 1991 there were 475 intercalary days inserted in the Julian calendar, but 14 of these were removed in the Gregorian calendar. While the Gregorian calendar is in use throughout most of the world today, some countries did not adopt it until early in the twentieth century.

While it remains a fascinating field for time historians, the above narrative provides conclusive evidence that conjugating calendar dates of significant events and assigning NTP timestamps to them is approximate at best. In principle, reliable dating of such events requires only an accurate count of the days relative to some globally alarming event, such as a comet passage or supernova explosion; however, only historically persistent and politically stable societies, such as the ancient Chinese and Egyptian, and especially the ancient Maya [MOR83], possessed the means and will to do so.

4.2. The Modified Julian Day System

In order to measure the span of the universe or the decay of the proton, it is necessary to have a standard day-numbering plan. Accordingly, the International Astronomical Union has adopted the use of the standard second and Julian Day Number (JDN) to date cosmological events and related phenomena. The standard day consists of 86,400 standard seconds, where time is expressed as a fraction of the whole day, and the standard year consists of 365.25 standard days.

In the scheme devised in 1583 by the French scholar Joseph Justus Scaliger and named after his father, Julius Caesar Scaliger, JDN 0.0 corresponds to 12^h (noon) on the first day of the Julian Era, 1 January 4713 BC, and runs in days and fractions since then. The Modified Julian Date (MJD), which is sometimes used to represent dates near our own era in conventional time and with fewer digits, is defined as $MJD = JD - 2,400,000.5$. While the JDN timescale is based on an absolute day numbering, it is not uniform relative to the absolute (atomic) timescale, since the Earth rotational period is not constant and has increased over ten seconds from 45 BC to the present.

For historical purposes the years prior to the Common Era (BC) are reckoned according to the Julian calendar, while the years of the Common Era (AD) are reckoned according to the Gregorian calendar. Since Monday, 1 January 1 AD in the Gregorian calendar corresponds to Monday, 3 January 1 in the Julian calendar [DER90], JDN 1,721,426.0 corresponds to 12^h on the first day of the Common Era, 1 January 1 AD. Following the convention that our century began at 0^h on 1 January 1900, at which time the solar year was already 12^h old, that eclectic instant corresponds to MJD 15,020.0 and the origin of all existing Internet timescales, including that of NTP.

5. Timescale Chronometry

International timescales are based on atomic clocks of exquisite stability compared to the astronomical clocks we normally live by. Occasionally, there is a need to correct the conventional civil timescale to account for

unpredictable wobbles in Earth rotation and related phenomena. The following sections discuss these issues with respect to the NTP timescale; in particular, how it handles leap seconds.

5.1. Determination of Frequency

For many years the most important use of time and frequency information was for worldwide navigation and space science, which depend on astronomical observations of the Sun, Moon and stars [TIM86]. Sidereal time is based on the transit of stars across the celestial meridian of an observer. The mean sidereal day is 23 hours, 56 minutes and 4.09 seconds, but varies about ± 30 ms throughout the year due to polar wandering and orbit variations. Ephemeris time is based on tables with which a standard time interval such as the tropical year - one complete revolution of the Earth around the Sun - can be determined through observations of the Sun, Moon and planets.

In 1958 the standard second was defined as $1/31,556,925.9747$ of the tropical year that began this century. On this scale the tropical year in 1900 was 365.2421987 days and the lunar month - one complete revolution of the Moon around the Earth - was 29.53059 days; however, the actual tropical year can be determined only to an accuracy of about 50 ms and the mean value has been increasing by about 5.3 ms per year. Without correction for this factor, an NTP timestamp would be in error over 26 seconds at the dawn of recorded history.

Of the three heavenly oscillators readily apparent to ancient mariners and astronomers - the Earth rotation about its axis, the Earth revolution around the Sun and the Moon revolution around the Earth - none of the three have the intrinsic stability, relative to modern technology, to serve as a standard reference oscillator. In 1967 the standard second was redefined as “9,192,631,770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the cesium-133 atom” [ALL74]. Prior to 0^h 1 January 1972 national frequency standards were occasionally corrected to agree with astronomical observations. Since then the time and frequency standards of the world have been based on International Atomic Time (TAI), which is defined and maintained using multiple cesium-beam oscillators to an accuracy of a few parts in 10^{13} , or a few tens of nanoseconds per day. Note that, while this provides an extraordinarily precise timescale, it does not necessarily agree with conventional solar time and may not in fact even be absolutely uniform, unless subtle atomic conspiracies can be ruled out [RAW87].

5.2. Determination of Time and Leap Seconds

The International Bureau of Weights and Measures (IBWM) uses astronomical observations provided by the

UTC Date	MJD	NTP Time	Offset
01 Jan 72	41,317	2,272,060,800	0
30 Jun 72	41,498	2,287,785,600	1
31 Dec 72	41,682	2,303,683,200	2
31 Dec 73	42,047	2,335,219,200	3
31 Dec 74	42,412	2,366,755,200	4
31 Dec 75	42,777	2,398,291,200	5
31 Dec 76	43,143	2,429,913,600	6
31 Dec 77	43,508	2,461,449,600	7
31 Dec 78	43,873	2,492,985,600	8
31 Dec 79	44,238	2,524,521,600	9
30 Jun 81	44,785	2,571,782,400	10
30 Jun 82	45,150	2,603,318,400	11
30 Jun 83	45,515	2,634,854,400	12
30 Jun 85	46,246	2,698,012,800	13
31 Dec 87	47,160	2,776,982,400	14
31 Dec 89	47,891	2,840,140,800	15
31 Dec 90	48,256	2,871,590,400	16

Table 3. Table of Leap-Second Insertions

U.S. Naval Observatory and other observatories to determine UTC. Starting from apparent mean solar time as observed, the UT0 timescale is determined using corrections for Earth orbit and inclination (the Equation of Time, as used by sundials), the UT1 (navigator’s) timescale by adding corrections for polar migration and the UT2 timescale by adding corrections for known periodicity variations. While standard frequencies are based on TAI, conventional civil time is based on UT1, which is presently slowing relative to TAI by a fraction of a second per year. When the magnitude of correction approaches 0.7 second, a leap second is inserted or deleted in the TAI timescale on the last day of June or December.

For the most precise coordination and timestamping of events since 1972, it is necessary to know when leap seconds were implemented in UTC and how the seconds are numbered. As specified in CCIR Report 517, which is reproduced in [BLA74], a leap second is inserted following second 23:59:59 on the last day of June or December and becomes second 23:59:60 of that day. A leap second would be deleted by omitting second 23:59:59 on one of these days, although this has never happened. Leap seconds were inserted prior to 1 January 1991 on the occasions listed in Table 3 (courtesy U.S. Naval Observatory). Published IBWM corrections consist not only of leap seconds, which result in step discontinuities relative to TAI, but 100-ms UT1 adjustments called DUT1, which provide increased accuracy for navigation and space science.

Note that the NTP time column actually shows the epoch following the last second of the day given in the UTC date and MJD columns (except for the first line), which is the precise epoch of insertion. The offset column

shows the cumulative seconds offset between the TAI timescale and the UTC timescale; that is, the number of seconds to add to the TAI clock in order to maintain nominal agreement with the UTC clock. Finally, note that the epoch of insertion is relative to the timescale immediately prior to that epoch; e.g., the epoch of the 31 December 1990 insertion is determined on the timescale in effect following the 31 December 1989 insertion, which means the actual insertion relative to the TAI clock was fifteen seconds later than the apparent time on the UTC timescale.

The UTC timescale thus ticks in standard (atomic) seconds and was set to the value 0^h MJD 41,317.0 at the epoch determined by astronomical observation to be 0^h on 1 January 1972 according to the Gregorian calendar; that is, the inaugural tick of the UTC Era. In fact, the inaugural tick which synchronized the cosmic oscillators, Julian clock, UTC clock and Gregorian calendar forevermore was displaced about ten seconds from the civil clock then in use, while the GPS clock is ahead of the UTC clock by seven seconds in early 1991. Subsequently, the UTC clock has marched forward relative to the Julian clock exactly one second on scheduled occasions at monumental epoches embedded in the institutional memory of our civilization. Note in passing that leap-second adjustments affect the number of seconds per day and thus the number of seconds per year. Apparently, should we choose to worry about it, the UTC clock, Julian clock and various cosmic clocks will inexorably drift apart with time until rationalized by some future papal bull.

5.3. The NTP Timescale and Reckoning with UTC

In NTP epoches are determined by copying the current value of the local clock to a timestamp variable when some significant event, such as the arrival of a message, occurs. Since NTP timestamps are cherished data and, in fact, represent the main product of the protocol, a special timestamp format has been established. An NTP timestamp is represented as a 64-bit unsigned fixed-point number, in seconds relative to 0^h on 1 January 1900. The integer part is in the first 32 bits and the fraction part in the last 32 bits. The precision of this representation is about 232 picoseconds, which should be adequate for even the most exotic requirements. Should NTP be in use when this 64-bit field overflows some time in 2036, external means will be necessary to qualify time relative to 1900 and time relative to 2036 (and other multiples of 136 years). Timestamped data requiring such qualification will be so precious that appropriate means should be readily available.

Since a large body of today's protocol architects expect to continue plying their trade as of 2036, some consid-

UTC		NTP	
hours	seconds	kiloseconds	seconds
31 Dec 90 23:59	:59	2,871,590	,399 +
(leap) 23:59	:60	2,871,590	,400 +
1 Jan 91 00:00	:00	2,871,590	,400
00:00	:01	2,871,590	,401

Figure 4. Comparison of UTC and NTP Timescales at Leap

eration as to the choice of format and origin of the NTP timescale may be in order. It is of course simple to mitigate the rollover problem simply by redefining the origin as, for example 2,272,060,800 or 1 January 1972, which results in a continuous timescale through all but the last two years of the next century. This does not, of course, change the interpretation of the NTP timestamp itself, just its representation in conventional civil time. Also, as will become evident shortly, in order to establish the most precise representation with respect to conventional civil time, it is necessary to retain a record of leap-second insertions, which cannot be predicted far in advance. Having surmounted that requirement, it does not seem odious to record the 136-year NTP epoch for such cherished timestamps upon the election of a pope, for example.

The NTP timescale is based on the UTC timescale, but not necessarily always coincident with it. At 0^h on 1 January 1972 (MJD 41,317.0), the first tick of the UTC Era, the NTP clock was set to 2,272,060,800, representing the number of standard seconds since 0^h on 1 January 1900 (MJD 15,020.0). The insertion of leap seconds in UTC and subsequently into NTP does not affect the UTC or NTP oscillator, only the conversion to conventional civil UTC time. However, since the only institutional memory available to NTP are the UTC broadcast services, the NTP timescale is in effect reset to UTC as each timecode is received. Thus, when a leap second is inserted in UTC and subsequently in NTP, knowledge of all previous leap seconds is lost.

Another way to describe this is to say there are as many NTP timescales as historic leap seconds. In effect, a new timescale is established after each new leap second. Thus, all previous leap seconds, not to mention the apparent origin of the timescale itself, lurch backward one second as each new timescale is established. If a clock synchronized to NTP in 1991 was used to establish the UTC epoch of an event that occurred in early 1972 without correction, the event would appear sixteen seconds late relative to UTC. However, NTP primary time servers resolve the epoch using the broadcast timecode,

so that the NTP clock is set to the broadcast value on the current timescale. As a result, for the most precise determination of epoch relative to the historic UTC clock, the user must subtract from the apparent NTP epoch the offsets shown in Table 2 at the relative epoches shown. This is a feature of almost all present day time-distribution mechanisms.

The chronometry involved can be illustrated with the help of Figure 4, which shows the details of seconds numbering just before, during and after the last scheduled leap insertion at 23:59:59 on 31 December 1990. Notice the NTP leap bits are set on the day prior to insertion, as indicated by the “+” symbols on the figure. Since this makes the day one second longer than usual, the NTP day rollover will not occur until the end of the first occurrence of second 400. The UTC time conversion routines must notice the apparent time and the leap bits and handle the timescale conversions accordingly. Immediately after the leap insertion both timescales resume ticking the seconds as if the leap had never happened. The chronometric correspondence between the UTC and NTP timescales continues, but NTP has forgotten about all past leap insertions. In NTP chronometric determination of UTC time intervals spanning leap seconds will thus be in error, unless the exact times of insertion are known.

It is possible that individual systems may use internal data formats other than the NTP timestamp format; however, a persuasive argument exists to use a two-part representation, one part for whole days (MJD or some fixed offset from it) and the other for the seconds (or some scaled value, such as milliseconds). This not only facilitates conversion between NTP and conventional civil time, but makes the insertion of leap seconds much easier. All that is required is to change the modulus of the seconds counter, which on overflow increments the day counter. This design insures that continuity of the timescale is assured, even if outside synchronization is lost before, during or after leap-second insertion. Since timestamp data are unaffected, synchronization is assured, even if timestamp data are in flight at the instant and originated before or at that instant.

6. Summary

By international agreement the primary frequency reference for our civilization is the atomic oscillator. The standard second is determined as a specified number of atomic cycles, the standard day as 86,400 standard seconds and the standard (Julian) year as 365.25 standard days. The Julian clock is synchronized to atomic frequency and its timescale calibrated in standard days (JDN), where JDN 0.0 corresponds to 12h on 1 January 4713 BC. While the Julian clock is elegant and everlasting, it is not useful in precision chronometric calculation

of conventional civil time of day and day of year, since our solar day and solar year oscillators are slightly unstable and incommensurate.

In order to maintain the nominal solar year, the Gregorian calendar mandates the insertion of leap days, which can be determined in advance. In order to maintain the nominal solar day, leap seconds must be inserted at times which cannot be reliably determined in advance. The basis of civil time is the UTC clock, which was cloned from the Julian clock at 0h on 1 January 1972. Without knowledge of prior leap seconds, an event determined on the UTC timescale can appear many seconds late on the Julian timescale.

The NTP timescale is based on the UTC timescale and runs at atomic frequency; however, it is calibrated in standard seconds after 0h on 1 January 1900. The NTP timescale is, in effect, redefined at each leap second; that is, since the NTP oscillator is phase-continuous spanning a leap second, the number of apparent NTP seconds since UTC came into being does not change. However, the number of actual UTC seconds has in fact changed. Reconciliation of the NTP and UTC timescales requires global institutional memory in the form of a table of occurrences such as shown in Table 3.

7. References

- [ALL74] Allan, D.W., J.H. Shoaf and D. Halford. Statistics of time and frequency data analysis. In: Blair, B.E. (Ed.). *Time and Frequency Theory and Fundamentals*. National Bureau of Standards Monograph 140, U.S. Department of Commerce, 1974, 151-204.
- [BEL86] Bell Communications Research. Digital Synchronization Network Plan. Technical Advisory TA-NPL-000436, 1 November 1986.
- [CAL86] “Calendar.” *The Encyclopaedia Britannica Macropaedia*, 15th ed., vol. 15, pp. 460-477. Encyclopaedia Britannica Co., New York, NY, 1986.
- [DER90] Dershowitz, N., and E.M. Reingold. Calendrical Calculations. *Software Practice and Experience* 20, 9 (September 1990), 899-928.
- [DOD81a] Defense Advanced Research Projects Agency. Internet Control Message Protocol. DARPA Network Working Group Report RFC-792, USC Information Sciences Institute, September 1981.
- [DOD81b] Defense Advanced Research Projects Agency. Internet Protocol. DARPA Network Working Group Report RFC-791, USC Information Sciences Institute, September 1981.
- [JON83] Jones, R.H., and P.V. Tryon. Estimating time from atomic clocks. *J. Research of the National*

- Bureau of Standards* 88, 1 (January-February 1983), 17-24.
- [LIN80] Lindsay, W.C., and A.V. Kantak. Network synchronization of random signals. *IEEE Trans. Communications COM-28*, 8 (August 1980), 1260-1266.
- [MAR85] Marzullo, K., and S. Owicki. Maintaining the time in a distributed system. *ACM Operating Systems Review* 19, 3 (July 1985), 44-54.
- [MIL89] Mills, D.L. Network Time Protocol (Version 2) specification and implementation. DARPA Network Working Group Report RFC-1119, University of Delaware, September 1989.
- [MIL90a] Mills, D.L. On the accuracy and stability of clocks synchronized by the Network Time Protocol in the Internet system. *ACM Computer Communication Review* 20, 1 (January 1990), 65-75.
- [MIL90b] Mills, D.L. Network Time Protocol (Version 3) specification, implementation and analysis. Electrical Engineering Department Report 90-6-1, University of Delaware, June 1990.
- [MIL91] Mills, D.L. Internet time synchronization: the Network Time Protocol. *IEEE Trans. Communications* 39, 10 (October 1991) (scheduled to appear).
- [MOR83] Morley, S.G., G.W. Brainerd and R.J. Sharer. *The Ancient Maya*, 4th ed., pp. 598-600. Stanford University Press, Stanford, CA, 1983.
- [MOY82] Moyer, G. The Gregorian Calendar. *Scientific American* 246, 5 (May 1982), 144-152.
- [NBS88] *Automated Computer Time Service (ACTS)*. NBS Research Material 8101, U.S. Department of Commerce, 1988.
- [POS80] Postel, J. User Datagram Protocol. DARPA Network Working Group Report RFC-768, USC Information Sciences Institute, August 1980.
- [POS83] Postel, J. Time protocol. DARPA Network Working Group Report RFC-868, USC Information Sciences Institute, May 1983.
- [RAM90] Ramanathan, P., K.G. Shin and R.W. Butler. Fault-tolerant clock synchronization in distributed systems. *IEEE Computer* 23, 10 (October 1990), 33-42.
- [RAW87] Rawley, L.A., J.H. Taylor, M.M. Davis and D.W. Allan. Millisecond pulsar PSR 1937+21: a highly stable clock. *Science* 238 (6 November 1987), 761-765.
- [SCI91] ScienceScope. Sounding out the threat of global warming. *Science* 251 (8 February 1991), 615.
- [SHI87] Shin, K.G., and P. Ramanathan. Clock synchronization of a large multiprocessor system in the presence of malicious faults. *IEEE Trans. Computers C-36*, 1 (January 1987), 2-12.
- [TIM86] "Time." *The Encyclopaedia Britannica Macropaedia*, 15th ed., vol. 28, pp. 652-664. Encyclopaedia Britannica Co., New York, NY, 1986.
- [VAS88] Vasanthavada, N., and P.N. Marinos. Synchronization of fault-tolerant clocks in the presence of malicious failures. *IEEE Trans. Computers C-37*, 4 (April 1988), 440-448.