# On the Accuracy and Stability of Clocks Synchronized by the Network Time Protocol in the Internet System[1]

David L. Mills
Electrical Engineering Department
University of Delaware

**Abstract**

This paper describes a series of experiments involving over 100,000 hosts of the Internet system and located in the U.S., Europe and the Pacific. The experiments are designed to evaluate the availability, accuracy and reliability of international standard time distribution using the Internet and the Network Time Protocol (NTP), which has been designated an Internet Standard protocol. NTP is designed specifically for use in a large, diverse internet system operating at speeds from mundane to lightwave. In NTP a distributed subnet of time servers operating in a self-organizing, hierarchical, master-slave configuration exchange precision timestamps in order to synchronize host clocks to each other and national time standards via wire or radio.

The experiments are designed to locate Internet hosts and gateways that provide time by one of three time distribution protocols and evaluate the accuracy of their indications. For those hosts that support NTP, the experiments determine the distribution of errors and other statistics over paths spanning major portions of the globe. Finally, the experiments evaluate the accuracy and reliability of precision timekeeping using NTP and typical Internet paths involving ARPANET, NSFNET and regional networks. The experiments demonstrate that timekeeping throughout most portions of the Internet can be maintained to an accuracy of a few tens of milliseconds and a stability of a few milliseconds per day, even in cases of failure or disruption of clocks, time servers or networks.

## 1. Introduction

How do hosts and gateways in a large, dispersed networking community know what time it is? How accurate are their clocks? In a 1988 survey involving 5,722 hosts and gateways of the Internet system [MIL88a], 1158 provided their local time via the network. Sixty percent of the replies had errors greater than one minute, while ten percent had errors greater than 13 minutes. A few had errors as much as two years. Most host clocks are set by eyeball-and-wristwatch to within a minute or two and rarely checked after that. Many of these are maintained by some sort of battery-backed clock/calender device using a room-temperature quartz oscillator that may drift seconds per day and can go for weeks between manual corrections. For many applications, especially those designed to operate in a distributed internet environment, much greater accuracy, stability and reliability are required.

The Network Time Protocol (NTP) is designed to distribute standard time using the hosts and gateways of the Internet system. The Internet consists of over 100,000 hosts on over 800 packet-switching networks interconnected by a comparable number of gateways. While the Internet backbone networks and gateways are engineered and managed for good service, operating speeds and service reliabilities vary considerably throughout the regional and campus networks of the system. This places severe demands on NTP, which must deliver accurate, stable and reliable standard time throughout the system, in spite of component failures, service disruptions and possibly mis-engineered implementations.

NTP and its forebears were developed and tested on PDP11 computers and the Fuzzball operating system, which was designed specifically for timekeeping precisions of a millisecond or better [MIL88b]. An implementation of NTP as a Unix 4.3bsd system daemon was built by Michael Petry and Louis Mamakos at the University of Maryland. A special-purpose hardware/software implementation of NTP was built be Dennis Ferguson at the University of Toronto. At least 16 NTP primary time servers are presently synchronized by radio or satellite to national time standards in the U.S., Canada and the U.K. About half of these are connected directly to backbone networks and are intended for ubiquitous access, while the remainder are connected to

1 Sponsored by: Defense Advanced Research Projects Agency contract number N00140-87-C-8901 and by National Science Foundation grant number NCR-89-13623.

regional and campus networks and intended for local distribution. It is estimated that there are well over 2000 secondary servers in North America, Europe and the Pacific synchronized by NTP directly or indirectly to these primary servers.

This paper describes several comprehensive experiments designed to evaluate the availability, accuracy, stability and reliability of standard time distribution using NTP and the hosts and gateways of the Internet. The first is designed to locate hosts that support at least one of three time protocols specified for use in the Internet, including NTP. Since Internet hosts are not centrally administered and network time is not a required service in the TCP/IP protocol suite, experimental determination is the only practical way to estimate the penetration of time service in the Internet. The remaining experiments use only NTP and are designed to assess the nominals and extremes of various errors that occur in regular system operation, including those due to the network paths between the servers and the radio propagation paths to the source of synchronization, as well as the intrinsic stabilities of the various radio clocks and local clocks in the system.

This paper does not describe in detail the architecture or protocols of NTP, nor does it present the rationale for the particular choice of synchronization method and statistical processing algorithms. Further information on the background, model and algorithms can be found in [MIL89c], while details of the latest NTP protocol specification can be found in [MIL89a]. This paper itself is an edited and expanded version of [MIL89b].

### 1.1. Standard Time and Frequency Dissemination

In order that precision time and frequency can be coordinated throughout the world, national administrations operate primary time and frequency standards and maintain Coordinated Universal Time (UTC) by observing various radio broadcasts and through occasional use of portable atomic clocks. 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 and are capable of maintaining fractional frequency stability to $10^{-13}$ and time to 100 ns when operated in multiple ensembles at various national standards laboratories.

The U.S. National Institute of Standards and Technology (NIST - formerly National Bureau of Standards) operates radio broadcast services for the dissemination of standard time [NBS79]. These include short-wave transmissions from stations WWV at Fort Collins, CO, and WWVH at Kauai, HI, long-wave transmissions from WWVB, also at Fort Collins, and satellite transmissions from the Geosynchronous Orbiting Environmental Satellite (GOES). These transmissions and those of some other countries, including Canada and the U.K., include a timecode modulation which can be decoded by special-purpose radio receivers and interfaced to an NTP time server.

Using high-frequency transmissions, reliable frequency comparisons can be made to the order of $10^{-7}$, but time accuracies are limited to the order of a millisecond [BLA74]. Using long-wave transmissions and 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. Using GOES the accuracy depends on an accurate ephemeris and correction factors, but is generally of the same order as WWVB. Other systems intended primarily for navigation, including LORAN-C [FRA82], Global Positioning System (GPS) [BES82], OMEGA [VAS78], and various very-low-frequency communication stations in principle can be used for very precise time and frequency transfer on a global scale; however, these systems do not provide timecodes including time-of-day or day-of-year information.

### 1.2. The Network Time Protocol

An accurate, reliable time distribution protocol must provide the following:

1. The primary time reference source(s) must be synchronized to national standards by wire, radio or portable clock. The system of time servers and clients must deliver continuous local time based on UTC, even when leap seconds are inserted in the UTC timescale.
2. The time servers must provide accurate, stable and precise time, even with relatively large statistical delays on the transmission paths. This requires careful design of the data smoothing and deglitching algorithms, as well as an extremely stable local clock oscillator and synchronization mechanism.
3. The synchronization subnet must be reliable and survivable, even under unstable conditions and where connectivity may be lost for periods extending to days. This requires redundant time servers and diverse transmission paths, as well as a dynamically reconfigurable subnet architecture.
4. The synchronization protocol must operate continuously and provide update information at rates sufficient to compensate for the expected wander of the room-temperature quartz oscillators commonly used in ordinary computer systems. It must operate efficiently with large numbers of time servers and clients in continuous-polled and procedure-call

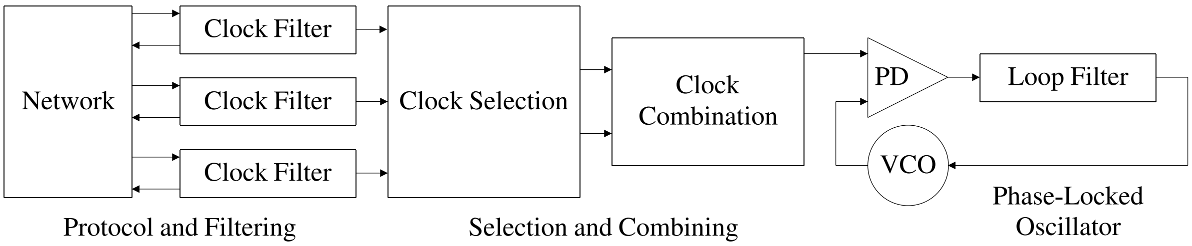


Figure 2. Network Time Protocol

modes and in multicast and point-to-point configurations.

5. The system must operate with a spectrum of systems ranging from personal workstations to supercomputers, but make minimal demands on the operating system and supporting services. Time server software and especially client software must be easily installed and configured.

In addition to the above, and in common with other promiscuously distributed services, the system must include generic protection against accidental or willful intrusion and provide a comprehensive interface for network management. In NTP address filtering is used for access control, while encrypted checksums are used for authentication [MIL89a]. Network management presently uses a proprietary protocol with provisions to migrate to standard protocols where available.

In NTP one or more primary time servers synchronize directly to external reference sources such as radio clocks. Secondary time servers synchronize to the primary servers and others in a configured subnet of NTP servers. Subnet servers calculate local clock offsets and delays between them using timestamps with 200 picosecond resolution exchanged at intervals up to about 17 minutes. As explained in [MIL89a], the protocol uses a distributed Bellman-Ford algorithm [BER87] to construct minimum-weight spanning trees within the subnet based on hierarchical level (stratum) and total synchronization path delay to the primary servers.

A typical NTP synchronization subnet is shown in Figure 1a, in which the nodes represent subnet servers and normal stratum number and the heavy lines the active synchronization paths. The light lines represent backup synchronization paths where timing information is exchanged, but not necessarily used to synchronize the local clock. Figure 1b shows the same subnet, but with the line marked *x* out of service. The subnet has reconfigured itself automatically to use backup paths, with the result that one of the servers has dropped from stratum 2 to stratum 3.

Besides NTP, there are several protocols designed to distribute time in local-area networks, including the

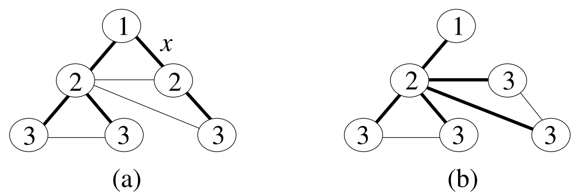


Figure 1. Subnet Synchronization

DAYTIME protocol [POS83a], TIME Protocol [POS83b], ICMP Timestamp message [DAR81] and IP Timestamp option [SU81]. The DCN routing protocol incorporates time synchronization directly into the routing protocol using algorithms similar to NTP [MIL83]. The Unix 4.3bsd time daemon *timed* uses a single master-time daemon to measure offsets of a number of slave hosts and send periodic corrections to them [GUS85]. However, these protocols do not include engineered algorithms to compensate for the effects of statistical delay variations encountered in wide-area networks and are unsuitable for precision time distribution throughout the Internet.

### 1.3. Determining Time and Frequency

In this paper to synchronize frequency means to adjust the clocks in the network to run at the same frequency, to synchronize time means to set the clocks so that all agree at a particular epoch with respect to UTC, as provided by national standards, and to synchronize clocks means to synchronize them in both frequency and time. A clock synchronization subnet operates by measuring clock offsets between the various servers in the subnet and so is vulnerable to statistical delay variations on the various transmission paths between them. In the Internet the paths involved can have wide variations in delay and reliability, while the routing algorithms can select landline or satellite paths, public network or dedicated links or even suspend service without prior notice.

In statistically noisy internets accurate time synchronization requires carefully engineered filtering and selection algorithms and the use of redundant resources and diverse transmission paths, while stable frequency synchronization requires finely tuned local clock tracking loops and multiple offset comparisons over relatively long periods of time. For instance, while only a few comparisons are usually adequate to resolve local time for an Internet host to within a few tens of milliseconds, dozens of measurements over many hours are required to achieve a frequency stability of a few tens of milliseconds per day and hundreds of measurements over many days to achieve the ultimate accuracy of a millisecond per day.

Figure 2 shows the overall organization of the NTP time server model. Timestamps exchanged with possibly many other servers are used to determine individual roundtrip delays and clock offsets relative to each server as follows. Number the times of sending and receiving NTP messages as shown below and let $i$ be an even integer.

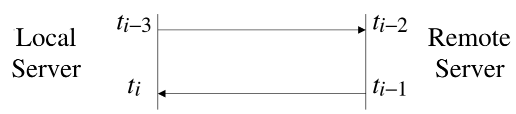


Then $t_{i-3}$, $t_{i-2}$, $t_{i-1}$, $t$ are the values of the four most recent timestamps as shown. The roundtrip delay $d_i$ and clock offset $c_i$ of the receiving server relative to the sending server is:

$$d_i = (t_i - t_{i-3}) - (t_{i-1} - t_{i-2}) \,,$$
$$c_i = \frac{(t_{i-2} - t_{i-3}) + (t_{i-1} - t_i)}{2} \,.$$

This method amounts to a continuously sampled, returnable-time system, which is used in some digital telephone networks [MIT80]. Among the advantages are that the transmitted time and received order of the messages are unimportant and that reliable delivery is not required. Obviously, the accuracies achievable depend upon the statistical properties of the outbound and inbound data paths. Further analysis and experimental results bearing on this issue can be found in [COL88], [MIL85a] and [MIL85b].

As shown in Figure 2, the computed offsets are first filtered to reduce incidental noise and then evaluated to select the most accurate and reliable subset among all available servers. The filtered offsets from this subset are first combined using a weighted average and then processed by a phase-locked loop (PLL). In the PLL the phase detector (PD) produces a correction term, which is processed by the loop filter to control the local clock, which functions as a voltage-controlled oscillator (VCO). Further discussion on these components is given in subsequent sections.

## 2. Discovering Internet Timetellers

An experiment designed to discover Internet time server hosts and evaluate the quality of their indications was conducted over a nine-day interval in August 1989. This experiment is an update of previous experiments conducted in 1985 [MIL85b] and early 1988 [MIL88a]. It involved sending time-request messages in each of three

| Protocol | Valid | Timeout | Error | Unknown |
|---|---|---|---|---|
| ICMP | 11533 | 61343 | 265 | 532 |
| TIME | 8441 | 1400 | 2293 | na |
| NTP | 784 | 713 | 6956 | na |
| Totals | 20758 | 63456 | 9514 | 532 |

Table 1. Time Responses by Protocol

time distribution protocols: ICMP Timestamp, TIME and NTP, to every Internet address that could reasonably be associated with a working host. Previously, lists of such addresses were derived from the Internet host table maintained by the Network Information Center (NIC), which contained 6382 distinct host and gateway addresses as of August 1989.

With the proliferation of the Internet domain-name system used to resolve host addresses from host names [MOC87], the NIC host table has become increasingly inadequate as a discovery vehicle for working host addresses. In a comprehensive survey of the domain-name system, Mark Lotter of SRI International recently compiled a revised host table of 137,484 entries. Each entry includes two lists, one containing the Internet addresses of a single host or gateway and the other containing its associated domain names. For the experiment this 9.4-megabyte table was sorted by address and extraneous information deleted, such as entries containing missing or invalid addresses, to produce a control file of 112,370 entries.

The experiment itself was conducted with the aid of the control file and a specially constructed experiment program written for the Fuzzball operating system [MIL88b]. The data were collected using experiment hosts located at the University of Delaware and connected to the University of Delaware campus network and SURA regional network. The experiment program reads each entry from the control file in turn and sends time-request messages to the first Internet address found. If no reply is received after one second, the program tries again. If no reply is received after an additional second, the program abandons the attempt and moves to the next entry in the control file. The program accumulates error messages and sample data for up to eight samples in each of the three time protocols. It abandons a host upon receipt of an ICMP error message [DAR81] and abandons further hosts on the same network upon receipt of an ICMP net-unreachable message. Using this procedure, attempts were made to read the clock for 107,799 distinct host addresses.

In the experiment the clock offsets were measured for each of the three time protocols relative to the local clock used on the experiment host, which is synchronized via radio to NBS standards to within a few milliseconds. The maximum, minimum and mean offset for up to eight replies for each protocol was computed and written to a statistics file, which contains valid responses, ICMP error messages of various kinds, timeout messages and other error indications. In the tabulation shown in Table 1 the timeout column shows the number of occasions when no reply was received, while the error column shows the error messages received, including ICMP time-exceeded, ICMP host-unreachable and ICMP port-unreachable messages. The unknown column tabulates occurrences of a specially marked ICMP Timestamp reply that indicates the host supports the protocol, but does not have a synchronized time-of-day clock.

In summary, of the 107,799 host addresses surveyed, 94,260 resulted in some kind of entry in the statistics file. Of these, 20,758 hosts (22%) were successful in returning an apparently valid indication. Note that there may be more than one attempt to read a host clock and that some clocks were read using more than one protocol. The valid entries were then processed to delete all except the first entry received for each address and protocol. In addition, if a host replied to an NTP request, all other entries for that host were deleted, while, if a host did not reply to an NTP request, but did for a TIME request, all other entries for that host were deleted. This results in a list of 8455 hosts which provided an apparently valid time indication, including 3694 for ICMP Timestamp, 7666 for TIME and 789 for NTP.

In order to discover as many NTP hosts as possible, the NTP synchronization subnet operating in the Internet was explored starting from the known primary servers using special monitoring programs designed for this purpose. This search, together with those discovered using the domain-name system and additional information gathered by other means, resulted in a total of about 990 NTP hosts. These hosts were then surveyed again, while keeping track of ancillary information to determine whether they were synchronized and operating correctly. This resulted in a list of 946 hosts apparently synchronized to the NTP subnet and operating correctly.

The methodology used here can miss a sizeable number of NTP hosts, such as personal computers, hosts not listed in the NIC or domain-name database and implementations that do not respond to the monitoring programs. In fact, extrapolating from data assembled from personal communications, the grand search described here discovered much less than half of the NTP-speaking hosts.

### 2.1. Evaluation of Timekeeping Accuracy by Protocol

In evaluating the quality of standard time distribution it is important to understand the effects of errors on the applications using the service. For many applications the maximum error under all conditions is more important than the mean error under controlled conditions. In these applications conventional statistics such as mean and variance are inappropriate. A useful statistic has been found to be the error distribution plotted on log-log axes and showing the probability $P(x>a)$ that a sample $x$ from the population exceeds the value $a$ on the $x$ axis. Figure 3 shows the error distributions for each of the three time protocols included in the survey. The top line in Figure 3 is for ICMP Timestamp, the next down is for TIME and the bottom is for NTP.

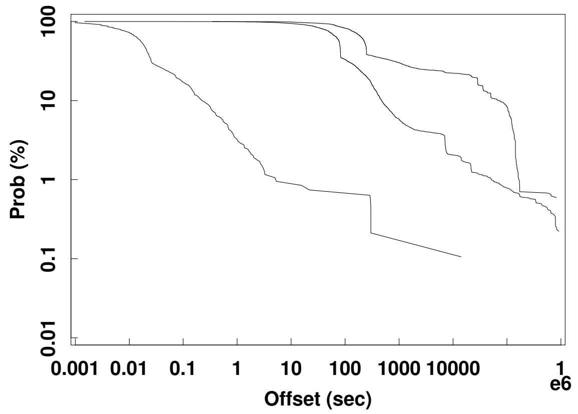


Figure 3. Offset Error Distributions by Protocol

The graphs shown in Figure 3 suggest several conclusions. First, the time accuracy of the various hosts varies dramatically over at least nine decades from milliseconds to over 11 days. To be sure, not many hosts showed very large errors and there is cause to believe these hosts either were never synchronized or were operating improperly. In the case of NTP, for example, which is designed expressly for time synchronization, eight hosts showed errors above ten seconds, a value considered barely credible for a host correctly synchronized by NTP in the Internet. It is very likely that some or all of these hosts, representing about one percent of the total NTP population, were using an old NTP implementation with known bugs. On the other hand, one percent of the ICMP Timestamp hosts show errors greater than a day, while one percent of TIME hosts show errors greater than a few hours. Clearly, at least on some machines running the latter two protocols, time is not considered a cherished service.

At the other end of the scale, Figure 3 suggests that at least 30 percent of the hosts in all three protocols make some attempt to maintain accurate time to about 30 ms with NTP, a minute with TIME and a couple of minutes with ICMP Timestamp. Between this regime and the one-percent regime the accuracies deteriorate; however, in general, NTP hosts maintain time about a thousand

| Host | Source | Hops | Samples | Offset | Delay |
|---|---|---|---|---|---|
| FORD | GOES | 10 | 8097 | 2 | 190 |
| ISI | WWVB | 12 | 2214 | -12 | 269.5 |
| MIT | WWV | 11 | 991 | -8 | 178 |
| NCAR | WWVB | 8 | 1563 | 6 | 231 |
| UIUC | WWVB | 7 | 3986 | -9 | 198 |
| UMD | WWVB | 5 | 16105 | 1 | 60 |

Table 2. Offset and Delay of Primary Servers (ms)

times more accurate than either of the other two protocols.

## 3. NTP Performance Analysis and Measurement

The above experiments were designed to assess the performance of all time servers that could be found in the Internet, regardless of protocol, system management discipline or protocol conformance. The remaining experiments described in this paper involve only the NTP protocol and the algorithms used in NTP implementations to synchronize the local clock.

### 3.1. Accuracy and Stability of NTP Primary Time Servers

In this experiment a number of NTP primary time servers was surveyed for overall accuracy and stability. Primary servers are synchronized by radio or satellite to national standards and located at or near points of entry to national and international backbone networks. Since they are monitored and maintained on a regular basis, their performance can be taken as representative of a managed system.

The experiment operated over a two-week period in August 1989 using paths between six primary servers on the east coast, west coast and midwest. All measurements were made from an experiment host located at the University of Delaware. Most of the paths involve links operating at 1.5 Mbps or higher, although there are over a dozen links on some paths and some lower speed links are in use. Samples of roundtrip delay and clock offset were collected at intervals from one to 17 minutes on all six paths and the data recorded in files for later analysis.

Table 2 shows the results of the survey, which involved about 33,000 samples. For each server the name, synchronization source, number of gateway/router hops and number of samples are shown. The offset and delay columns show the sample medians for these quantities in milliseconds. Note that the number of samples collected depends on whether the server is selected for clock synchronization, as determined by the NTP clock-selection algorithm described in [MIL89a].

As in previous surveys of this type, statistics based on the sample median yield more reliable results than those based on the sample mean. However, statistics based on the trimmed mean (also called Fault-Tolerant Average [KOP87]) with 25 percent of the samples removed are within a millisecond of the values shown in Table 2.

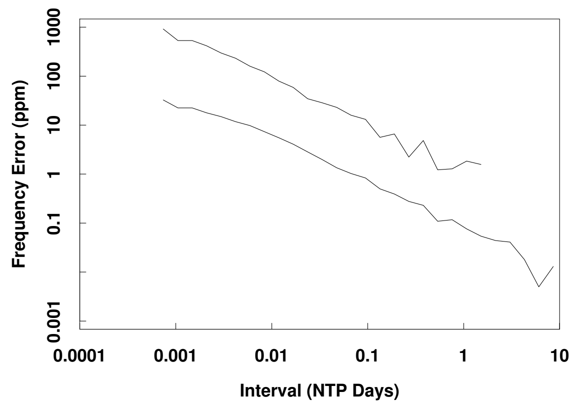


Figure 4. Frequency Stability of Primary Servers

The residual offset errors apparent in Table 2 can be traced to subtle asymmetries in path routing and network/gateway configurations. If these can be calibrated, perhaps using a portable atomic clock, reliable time transfer over the Internet should be possible within a millisecond or two if measurements are made over periods consistent with the two-week experiment. Assuming successive offset measurements can be made with confidence to this order, frequency transfer over the Internet could in principle be determined to the order of $10^{-9}$ in two weeks.

In order to test this conjecture an experiment was designed to determine the stability of the apparent timescale constructed from the first-order offset differences produced in an experiment similar to that which produced Table 2. This is similar to the approach described in [ALL74a] to analyze the intrinsic characteristics of a precision oscillator. In the month-long experiment, measured offsets were filtered by the algorithm described in the next section. The resulting samples were averaged at given intervals from about a minute to about ten days. The difference in offsets at the beginning and end of the interval divided by the duration of the interval represents the frequency during that interval. The standard deviation $\sigma(\tau)$ calculated from the sample population for each given interval $\tau$ is shown in Figure 4. Among the primary servers listed in Table 2, the lower curve represents the "best" one (UMD) and the upper curve the "worst" one (ISI).

The results show that, even for the best server and using carefully filtered data averaged over periods in the order of days, reliable stabilities approaching .01 parts per million (ppm) - about a millisecond per day - are difficult to achieve without further processing. Techniques which

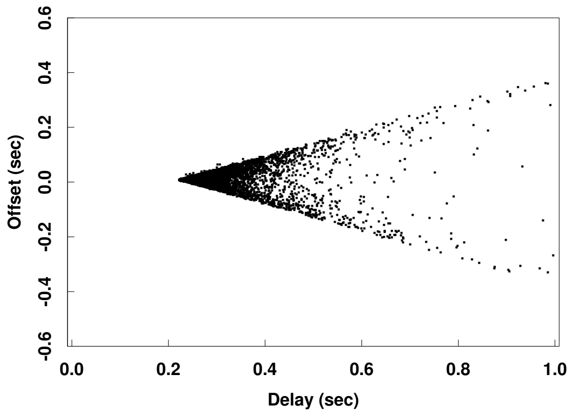


Figure 5. Offset-Delay Scatter Diagram

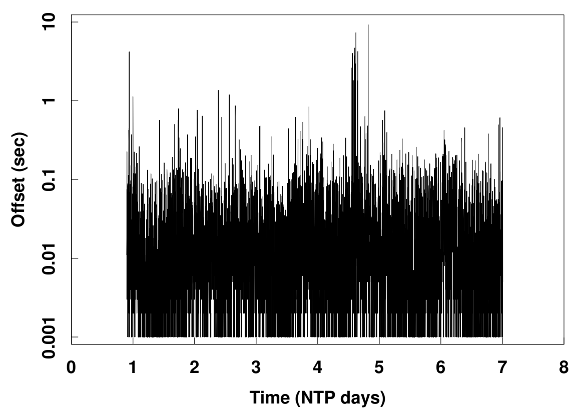


Figure 6. Raw Offsets

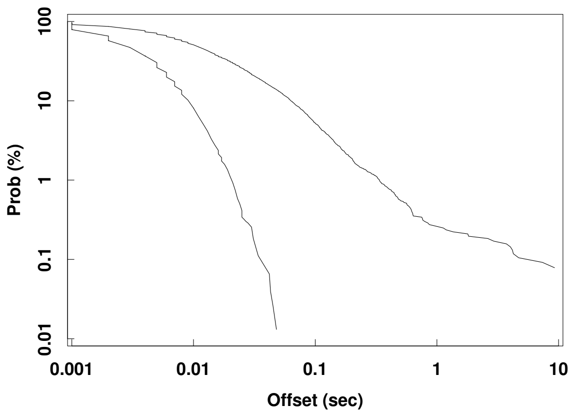


Figure 8. NTP Offset Error Distribution

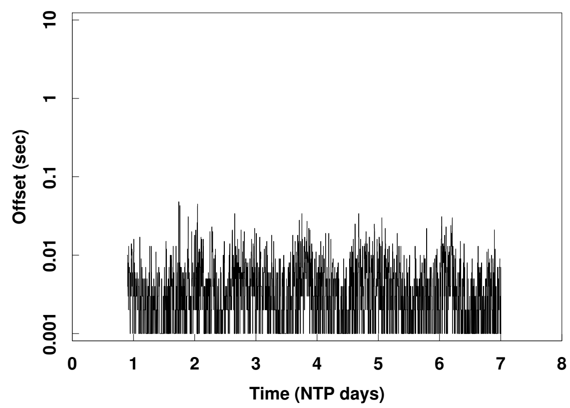


Figure 7. Filtered Offsets

can approach this goal will be presented later in this paper.

### 3.2. Effects Due to Filtering Algorithm

In order to more completely assess the accuracy and reliability that clocks can be synchronized using NTP and the Internet, the paths illustrated in Table 2 were carefully measured in several surveys conducted over a period of 18 months. Each survey used up to six time servers and lasted up to two weeks. A typical survey involves the path between experiment hosts at the University of Delaware and USC Information Sciences Institute, located near Los Angeles, over a complex path of up to twelve network hops involving NSFNET, ARPANET and several other regional and campus nets. This path was purposely selected as among the statistically noisiest in order to determine how well clocks can be synchronized under adverse conditions.

A number of algorithms for deglitching and filtering time-offset data are summarized in [MIL85a] and [MIL89c]. Experiments during the development of NTP Version 2 have produced an algorithm which provides high accuracy together with a low computational burden. The key to the new algorithm becomes evident through an examination of scatter diagrams plotting clock offset versus roundtrip delay. Without making any assumptions about the distributions of queueing and transmission delays on either direction along the path between two servers, but assuming the intrinsic frequency errors of the two clocks are relatively small, let $d_0$ and $c_0$ represent the delay and offset when no other traffic is present on the path and so represents the best estimates of the true values. The problem is to accurately estimate $d_0$ and $c_0$ from a sample population of $d_i$ and $c_i$ collected under typical conditions and varying levels of network load.

Figure 5 shows a typical scatter diagram for the path under study, in which the points $(d_i, c_i)$ are concentrated near the apex of a wedge defined by lines extending from the apex with slopes $\pm 0.5$, corresponding to the locus of points as the delay in one direction increases while the delay in the other direction does not. From these data it is obvious that good estimators for $(d_0, c_0)$ are points near the apex and that the best offset samples occur at the lower delays. Therefore, an appropriate technique is simply to select from the $n$ most recent samples the sample with lowest delay and use its associated offset as the estimate. This is the basis of the clock filter shown in Figure 2 and the NTP Version 2 algorithm described in detail in [MIL89a].

Figure 6 shows the raw time-offset series for the path under study over a six-day interval, in which occasional errors up to several seconds are apparent. Figure 7 shows the time-offset series produced by the filtering algorithm, in which the large errors have been dramatically reduced. Finally, the overall performance of the path is apparent from the error distributions shown in Figure 8. The upper line shows the distribution for the raw data, while the lower line shows the filtered data. The significant facts apparent from the latter line are that the median error over all samples was only a few milliseconds, while the maximum error was no more than 50 ms.

### 3.3. Effects due to Other Processing Algorithms

Precision timekeeping requires an exceptionally stable local oscillator reference in order to deliver accurate time when the synchronization path to a primary server has failed. Furthermore, the oscillator and control loop must maintain accurate time and stable frequency over wide variations in synchronization path delays. For instance, in order to maintain time to within a millisecond per day without outside reference, the local oscillator frequency must maintain stability to within .01 ppm or better.

Stabilities of this order usually require a relatively expensive oven-compensated quartz oscillator, which is not a common component in everyday computer systems. The NTP local clock model uses an adaptive-parameter, type-II, phase-locked loop (PLL), which continuously corrects local oscillator phase and frequency variations relative to updates received from the network or radio clock. The (open-loop) transfer function is

$$F(s) = \frac{\omega_c^2}{s^2\tau^2}\left(1 + \frac{s\tau}{\omega_z}\right),$$

where $\omega_c$ is the gain (crossover frequency), $\omega_z$ the corner frequency of the lead network (necessary for PLL stability), and $\tau$ is a parameter used for bandwidth control.

Bandwidth control is necessary to match the PLL dynamics to varying levels of timing noise due to the intrinsic stability of the local oscillator and the prevailing path delays in the network. On one hand, the loop must track uncompensated board-mounted crystals found in common computing equipment, where the frequency tolerance may be only .01 percent and can vary several ppm as the result of normal room temperature changes. On the other hand, after the frequency errors have been tracked for several days, and assuming the local oscillator can be stabilized accordingly, the loop must maintain stabilities to the order of .01 ppm. The NTP PLL is designed to adapt automatically to these regimes by measuring the sample variance and adjusting $\tau$ over a 16-fold range.

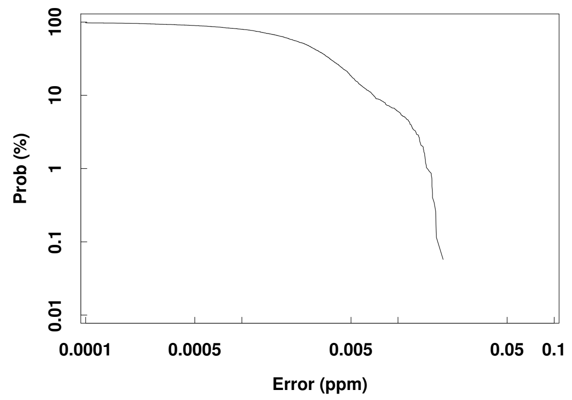


Figure 9. PLL Frequency Error Distribution

In order to assess how closely the NTP PLL meets these objectives, the experiment described in Section 3.1 above was repeated, but with the local clock of the experiment host derived from a precision quartz oscillator. The offsets measured between each of the six primary servers and the experiment host were collected and processed by a simulator that duplicates the NTP processing algorithms. However, in addition to the algorithms described in [MIL89a], which select a subset of quality clocks and from them a single clock as the synchronization source, an experimental clock-combining method involving a weighted average of offsets from all selected clocks was used. In principle, such methods can reduce the effect of systematic offsets shown in Table 2 [ALL74b]. However, these methods can also significantly increase the sample variance presented to the PLL and thus reduce the local-clock stability below acceptable levels. Thus, the experiment represents a worst-case scenario.

Figure 9 shows the frequency error distribution produced by the simulator using offset samples collected from all six primary servers over a four-week period. The results show that the maximum frequency error over the entire period from all causes is less than .02 ppm, or a couple of milliseconds per day. During this period there were several instances where other servers failed and where severe congestion on some network paths caused weighting factors to change in dramatic ways and $\tau$ to be adjusted accordingly. Figure 9 may thus represents the bottom line on system performance at the present level of NTP technology refinement.

| Time (UTC) | MUF (MHz) | 2.5 (MHz) | 5 (MHz) | 10 (MHz) | 15 (MHz) | 20 (MHz) |
|---|---|---|---|---|---|---|
| 0 | 19.8 | 9.5x2 | 9.5x2 | 9.5x2 | 8.6n1 | |
| 2 | 16.0 | 9.0n2 | 9.0n2 | 8.6n1 | 8.6n1 | |
| 4 | 14.7 | 9.0n2 | 9.0n2 | 8.6n1 | | |
| 6 | 13.9 | 9.0n2 | 9.0n2 | 8.6n1 | | |
| 8 | 13.4 | 9.0n2 | 9.0n2 | 8.6n1 | | |
| 10 | 13.0 | 9.0n2 | 9.0n2 | 8.6n1 | | |
| 12 | 12.8 | 9.5x2 | 9.5x2 | 8.6n1 | | |
| 14 | 28.8 | | 9.7j2 | 8.8j1 | 8.8j1 | 8.8j1 |
| 16 | 32.7 | | 9.7j2 | 8.8j1 | 8.8j1 | 8.8j1 |
| 18 | 33.9 | | 9.7j2 | 8.8j1 | 8.8j1 | 8.8j1 |
| 20 | 33.1 | | 9.7j2 | 8.8j1 | 8.8j1 | 8.8j1 |
| 22 | 29.9 | | 9.7j2 | 8.8j1 | 8.8j1 | 8.8j1 |

Table 3. Radio Propagation Delays (ms)

## 4. Accuracy and Stability of Radio Synchronization

In order to assess the overall system synchronization accuracy relative to UTC, it is necessary to consider the inherent accuracy, stability and precision of the radio propagation paths and radio clocks themselves. All of the radio clocks used in the surveys have a design precision within one millisecond and are potentially accurate to within a millisecond or two relative to the propagation medium. However, the absolute accuracy depends on knowledge of the radio propagation path to the source of standard time and frequency. In addition, the radio clocks themselves can be a source of random and systematic errors.

### 4.1.Estimation of Propagation Delays

An evaluation of the timekeeping accuracy of the NTP primary servers relative to national standards in principle requires calibration by a portable atomic clock; however, in the absence of a portable clock, the propagation delay can be estimated for the great-circle path between the known geographic coordinates of the transmitter and receiver. However, this can result in errors as large as two milliseconds when compared to the actual oblique ray path. Additional errors can be introduced by unpredictable latencies in the radio clocks, operating system, hardware and in the protocol software (e.g., encryption delays) for NTP itself.

It is possible to estimate the timekeeping accuracy by means of a detailed analysis of the radio propagation path itself. In the case of the WWVB and MSF services on 60 kHz, the variations in path delay are relatively well understood and limited to the order of 50 microseconds [BLA74]. In the case of the GOES service the accuracy is limited by the ability to accurately estimate the distance along the line-of-sight path to the satellite and the ability to maintain accurate stationkeeping in geosynchronous orbit. In principle, the estimation errors for either of these services is small compared to the accuracy usually expected of Internet timestamps generated with NTP.

However, in the case of the WWV/H and CHU services, which operate on HF frequencies from 2.5 through 20 MHz, radio propagation is determined by the upper ionospheric layers, which vary in height throughout the day and night, and by the geometric ray path determined by the maximum usable frequency (MUF) and other factors, which also vary throughout the day, season and phase of the 11-year sunspot cycle.

In an effort to calibrate how these effects affect the limiting accuracy of the NTP primary servers using WWV/H and CHU services, existing computer programs were used to determine the maximum usable frequency (MUF) and propagation geometry for typical ionospheric conditions forecast for January 1990 on the 2476-km path between Newark, DE, and Fort Collins, CO, by two-hour intervals. The results, shown in Table 3, assume a smoothed sunspot number (SSN) of 194 and include the time interval (UTC hour), MUF (MHz) and delay (ms) for frequencies from 2.5 through 20 MHz. In case no propagation path is likely, the delay entry is left blank. The delay itself is followed by a code indicating whether the path is entirely in sunlight (j), in darkness (n) or mixed (x) and the number of hops. A symbol (m) indicates two or more geometric paths are likely with similar amplitudes, which may result in multipath fading and unstable indications.

From Table 3 it can be seen that the delay decreases as the controlling ionospheric layer (F2) falls during the night (to about 250 km) and rises during the day (to about 350 km). The delay also changes when the number of hops and thus the oblique ray geometry changes. The maximum delay variation for this particular path is from 8.6 to 9.7 ms, a variation of 1.1 ms. While this variation represents a typical scenario, other scenarios have been

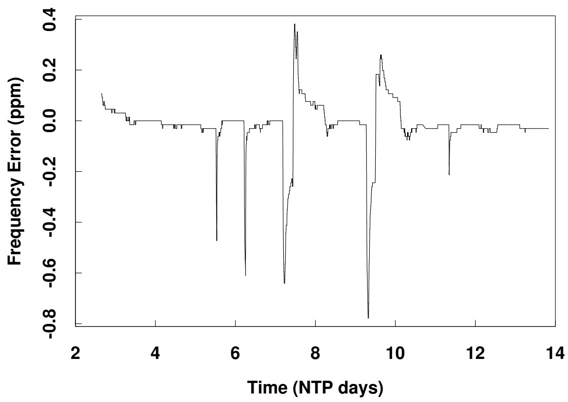


Figure 10. Radio Clock Frequency Errors

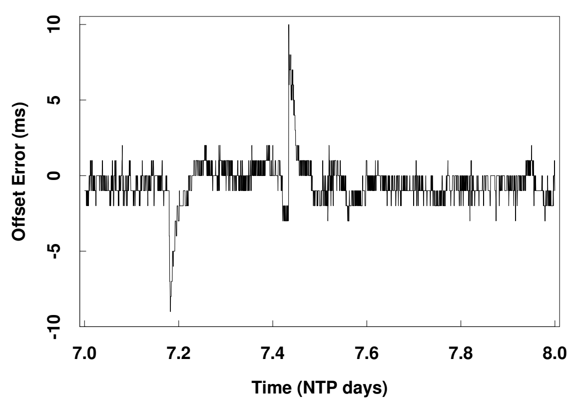


Figure 11. Radio Clock Offset Errors

found where the variations exceed two milliseconds. These results demonstrate that the ultimate accuracy of HF-radio derived NTP time may depend on the ability to accurately estimate the propagation path variations or to confine observations to the same time each day.

### 4.2. Accuracy and Stability of Radio Clocks

The final experiment reported in this paper involves an assessment of the accuracy and stability of a commercial WWV/H radio clock under typical propagation conditions. In order to separate these effects from those due to the measurement host, the local clock was derived from a precision oven-compensated quartz oscillator with rated stability of $\pm 5 \times 10^{-9}$ per day and aging rate of $1 \times 10^{-9}$ per day. The oscillator was set to within about $\pm 1 \times 10^{-8}$ relative to the 20-MHz WWV transmission under good propagation conditions near midday at the midpoint of the propagation path. The offsets of the radio clock relative to the local clock were filtered and processed by the NTP algorithms (open loop) and then recorded at 30-second intervals for a period of about two weeks.

The results of the experiment are shown in Figure 10 and Figure 11. Figure 10 shows the estimated frequency error by intervals for the entire period and reveals a frequency stability generally within .05 ppm, except for occasional periods where apparent phase hits cause the indications to surge. The times of these surges are near times when the path MUF between the transmitter and receiver is changing rapidly (see Table 3) and the receiver must change operating frequency to match. An explanation for the surges is evident in Figure 11, which shows the measured offsets during an interval including a typical surge. The figure shows a negative phase excursion of about 10 ms near the time the MUF would ordinarily fall in the evening and a similar positive excursion near the time the MUF would ordinarily rise in the morning.

Since the phase excursions are far beyond those expected due to ionospheric effects alone, the most likely explanation is that the increased noise in received WWV/H signals near the time of MUF-related frequency changes destabilizes the signal processing algorithms resulting in incorrect signal tracking. This particular problem has not been observed with WWVB or GOES radio clocks.

## 5. Conclusions

Over the years it has become something of a challenge to discover and implement architectures, algorithms and protocols which deliver precision time in a statistically rambunctious Internet. In perspective, for the ultimate accuracy in frequency and time transfer, navigation systems such as LORAN-C, OMEGA and GPS, augmented by portable atomic clocks, are the preferred method. On the other hand, it is of some interest to identify the limitations and estimate the magnitude of timekeeping errors using NTP and typical Internet hosts and network paths. This paper has identified some of what are believed to be the major limitations in accuracy and measured their effects in large-scale experiments involving major portions of the Internet.

The results demonstrated in this paper suggest several improvements that can be made in subsequent versions of the protocol and hardware/software implementations, such as improved radio clock designs, improved timebase hardware, at least at the primary servers, improved frequency-estimation algorithms and more diligent monitoring of the synchronization subnet. When a sufficient number of these improvements mature, NTP Version 3 may appear.

## 6. References


[ALL74a] 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.

[ALL74b] Allan, D.W., J.E. Gray and H.E. Machlan. The National Bureau of Standards atomic time scale: generation, stability, accuracy and accessibility. In: Blair, B.E. (Ed.). *Time and Frequency Theory and Fundamentals*. National Bureau of Standards Monograph 140, U.S. Department of Commerce, 1974, 205-231.

[BER87] Bertsekas, D., and R. Gallager. *Data Networks*. Prentice-Hall, Englewood Cliffs, NJ, 1987.

[BES82] Beser, J., and B.W. Parkinson. The application of NAVSTAR differential GPS in the civilian community. *Navigation 29, 2* (Summer 1982).

[BLA74] Blair, B.E. Time and frequency dissemination: an overview of principles and techniques. In: Blair, B.E. (Ed.). *Time and Frequency Theory and Fundamentals*. National Bureau of Standards Monograph 140, U.S. Department of Commerce, 1974, 233-313.

[COL88] Cole, R., and C. Foxcroft. An experiment in clock synchronisation. *The Computer Journal 31*, 6 (1988), 496-502.

[DAR81] Defense Advanced Research Projects Agency. Internet Control Message Protocol. DARPA Network Working Group Report RFC-792, USC Information Sciences Institute, September 1981.

[FRA82] Frank, R.L. History of LORAN-C. *Navigation 29, 1* (Spring 1982).

[GUS85] Gusella, R., and S. Zatti. The Berkeley UNIX 4.3BSD time synchronization protocol: protocol specification. Technical Report UCB/CSD 85/250, University of California, Berkeley, June 1985.

[KOP87] Kopetz, H., and W. Ochsenreiter. Clock synchronization in distributed real-time systems. *IEEE Trans. Computers C-36,* 8 (August 1987), 933-939.

[MIL83] Mills, D.L. DCN local-network protocols. DARPA Network Working Group Report RFC-891, M/A-COM Linkabit, December 1983.

[MIL85a] Mills, D.L. Algorithms for synchronizing network clocks. DARPA Network Working Group Report RFC-956, M/A-COM Linkabit, September 1985.

[MIL85b] Mills, D.L. Experiments in network clock synchronization. DARPA Network Working Group Report RFC-957, M/A-COM Linkabit, September 1985.

[MIL88a] Mills, D.L. Network Time Protocol (version 1) specification and implementation. DARPA Network Working Group Report RFC-1059, University of Delaware, July 1988.

[MIL88b] Mills, D.L. The fuzzball. *Proc. ACM SIGCOMM 88 Symposium* (Palo Alto, CA, August 1988), 115-122.

[MIL89a] Mills, D.L. Network Time Protocol (version 2) specification and implementation. DARPA Network Working Group Report RFC-1119, University of Delaware, September 1989.

[MIL89b] Mills, D.L. Measured performance of the Network Time Protocol in the Internet system. DARPA Network Working Group Report RFC-1128, University of Delaware, October 1989.

[MIL89c] Mills, D.L. Internet time synchronization: the Network Time Protocol. DARPA Network Working Group Report RFC-1129, University of Delaware, October 1989.

[MIT80] Mitra, D. Network synchronization: analysis of a hybrid of master-slave and mutual synchronization. *IEEE Trans. Communications COM-28,* 8 (August 1980), 1245-1259.

[MOC87] Mockapetris, P. Domain names - concepts and facilities. DARPA Network Working Group Report RFC-1034, USC Information Sciences Institute, November 1987.

[NBS79] *Time and Frequency Dissemination Services.* NBS Special Publication 432, U.S. Department of Commerce, 1979.

[POS83a] Postel, J. Daytime protocol. DARPA Network Working Group Report RFC-867, USC Information Sciences Institute, May 1983.

[POS83b] Postel, J. Time protocol. DARPA Network Working Group Report RFC-868, USC Information Sciences Institute, May 1983.

[SU81] Su, Z. A specification of the Internet protocol (IP) timestamp option. DARPA Network Working Group Report RFC-781. SRI International, May 1981.

[VAS78] Vass, E.R. OMEGA navigation system: present status and plans 1977-1980. *Navigation 25, 1* (Spring 1978).