Network Working Group
Request for Comments 1128

D. L. Mills
University of Delaware

October 1989

Measured Performance of the Network Time Protocol in the Internet System

Status of this Memo

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 DARPA/NSF Internet and the Network Time Protocol (NTP), which is specified as an Internet Standard in RFC-1119. 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 subnet 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 DARPA, NSFNET and other agency networks. The experiments demonstrate that timekeeping accuracy throughout most portions of the Internet can be ordinarily maintained to within a few tens of milliseconds, even in cases of failure or disruption of clocks, time servers or networks.

This memo does not specify a standard. Distribution of this memo is unlimited.

Keywords: internet clock synchronization, standard time distribution, Internet protocol, timekeeping experiments.

Table of Contents

List of Figures

List of Tables

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 as much as a second per day and can go for weeks between manual corrections. For many applications, especially distributed internet applications, much greater accuracy and reliability is 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 and reliable standard time 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 called *ntpd* 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. Over a dozen NTP primary time servers are synchronized by radio or satellite to national time standards in the U.S. and Canada. About half of these are connected directly to international backbone networks and are intended for ubiquitous access, while the remainder are connected to regional networks and intended for regional and local access. 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 large scale experiments designed to evaluate the availability, accuracy 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 in various types of errors that occur in regular system operation, including those due to the network paths between the servers, the radio propagation path to the source of synchronization and the radio clock itself.

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 [MIL89a], while details of the latest NTP protocol specification can be found in [MIL89b].

1.1. Standard 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 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 UTC frequency 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 three radio services for the dissemination of standard time and frequency information [NBS79]. 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 composition throughout the day and season and result in unpredictable delay variations at the receiver (see Section 3.2). While these services and those operated by the National Research Council of Canada (CHU) and other countries can be received over large areas of the world, 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].

A second service operated by NIST is the low-frequency (LF or CCIR band 5) transmissions from Boulder, CO (WWVB), which 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]. However, there is only one station and it operates at modest power levels.

The third service operated by NIST uses ultra-high frequency (UHF or CCIR band 9) transmissions from the Geosynchronous Orbiting Environmental Satellite (GOES). There is some speculation on the continued operation of GOES, especially if the LORAN-C [FRA82] and Global Positioning System (GPS) [BES82] radiopositioning systems operated by other U.S. agencies continue to evolve as expected. While the OMEGA [VAS78] radionavigation system operated by the U.S. Navy and other countries can in principle provide worldwide frequency and time distribution, this system is unlikely to long survive the operational deployment of GPS.

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 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 up 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 crystal oscillators used in ordinary computer systems. It must operate efficiently with large numbers of time servers and clients in continuous-polled and procedure-call 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 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 the configured subnet using NTP. Subnet peers calculate clock offset and delay between them using timestamps with 200 picosecond resolution exchanged at intervals of 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 and total synchronization path delay to the root.

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

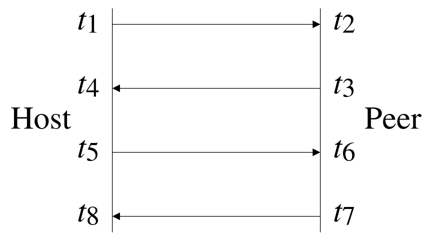

Figure 1. Calculating Delay and Offset

them in both frequency and time. A clock synchronization subnet operates by measuring clock offsets between the various peers 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 subnets accurate time synchronization requires carefully engineered filtering and selection algorithms and the use of redundant resources and diverse transmission paths, while accurate frequency synchronization requires finely tuned oscillator tracking loops and multiple phase 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 resolve frequency to within a few tens of milliseconds per day.

In NTP the roundtrip delay and clock offset relative to a peer time server are calculated as follows. Number the times of sending and receiving NTP messages as shown in Figure 1 and let i be an even integer. Then t_{i-3}, t_{i-2}, t_{i-1} and t_i are the contents of the four most recent timestamps as above. The roundtrip delay d_i and clock offset c_i of the receiving host relative to the sending peer 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 order and timing 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].

The computed offsets are first filtered to reduce stochastic noise and then evaluated to determine the most accurate and reliable selection among usually several peers. The filtered offsets from the selected peer are used to adjust the phase and frequency of the host clock, which is implemented by an adaptive-parameter, first-order, phase-locked oscillator. The loop parameters are carefully engineered to provide low frequency errors, so that the host clock will retain good accuracy during subnet outages of a day or more.