# Scalable, High Speed, Internet Time Synchronization

Advanced Research Projects Agency
Contract DABT 63-95-C-0046

Quarterly Progress Report
1 December 1996 - 28 February 1997

David L. Mills
Electrical Engineering Department
University of Delaware

## 1. Introduction

This report covers the work done in support of the DARPA Information Technology Office program in computer networking. Contributors to this effort include Prof. David L. Mills and graduate students Ajit Thyagarajan and Bradley Cain. The project continues previous research in network time synchronization technology jointly funded by DARPA, NSF, US Navy and US Army. The technology makes use of the Network Time Protocol (NTP), widely used in the Internet, together with engineered modifications designed to improve accuracy in high speed networks, including SONET and ATM, expected to be widely deployed in the next several years. Specific applications benefiting from this research include multicast topologies, multimedia, real-time conferencing, cryptographic systems, and management of distributed, real-time experiments.

Recent projects reported in papers, technical reports, project reports and technical memoranda include advances in precision timekeeping technology, improved clock discipline algorithms, and engineered security models and protocols for scalable, distributed server systems. Past projects include a relatively inexpensive precision timing receiver using the LORAN-C radionavigation system, and an optimum matched-filter receiver/decoder using DSP technology. Software developed with joint funding includes the NTP Version 3 implementation for Unix and Windows and a set of precision-time kernel modifications for major Unix workstation manufacturers. Finally, the joint projects involve the conduct of experiments designed to evaluate the success of the research and assist technology transfer to computer manufacturers and network providers.

## 2. Present Status

Most of the effort during the quarter was concentrated on the refinement and documentation of the Network Time Protocol, Version 4, as documented in the reports listed later in this report. Work continued on the autoconfigure scheme proposed for NTP and the subject of a pending dissertation by Mr. Ajit Thyagarajan. The theoretical work on graph-theoretic algorithms and heuristics has been completed and now being documented.

Work continued on the refinement of the NTP clock discipline algorithms to achieve reliable submillisecond accuracy on computer systems without direct connection to stabilized time or frequency sources. This work extends previous work reported by Judah Levine at NIST Boulder, resulting in a definitive analysis and modelling of typical computer clock hardware and client/server paths in the global Internet.

## 2.1 NTP Version 4 Progress

Conversion to the GNU autoconfigure scheme, which greatly simplifies program maintenance and porting to new architectures, has been almost completed. Volunteer Harlan Stenn worked on his own time and installed many GNU programming tools in the process. Production of a new software distribution is now almost completely automated. While system-specific peculiarities of most Unix workstation kernels, such as Digital, Sun, HP, SGI and Intel, has been overcome, occasional porting problems probably remain to be discovered.

A major effort during the quarter was with the analysis and simulation of the new clock discipline algorithm, as reported in the technical report listed at the end of this report. The results of this effort provide a major improvement in timekeeping stability and accuracy, as well as allowing the poll intervals to be substantially increased without materially affecting accuracy. Much of the analytical development was conducted using the NTP simulator program ntpsim, which includes the operative algorithms of the NTP daemon supported by a suite of synthetic-noise generators and performance evaluation utilities.

In order to evaluate the performance of the new algorithms, a body of timestamp data were collected between various Internet primary time servers in the U.S. and other countries. Since the participating time servers were synchronized to external sources of time, the timestamp data represented actual one-way network path delays and allowed the errors due to these delays to be separated from those due to clock oscillator instabilities. The latter were determined from previous experiments measuring free-running clock offsets relative to directly or indirectly connected precision pulse-per-second (PPS) signals from a cesium oscillator. Further information on this project, including a status report and set of briefing slides, is on the web at http://www.eecis.udel.edu/~mills/precise.html.

Graduate student Ajit Thyagarajan is in the final stages of implementation for the autonomous configuration scheme planned for NTP Version 4. Much of the work, including the intricately crafted association model is documented in a technical memorandum to be included in a future specification document for NTP Version 4. Much of the code has already been implemented and tested for the manycast/anycast modes described in previous reports. The remaining code to be implemented and tested is mainly for the Span-Limited, Add-Drop, Greedy (SLAG) algorithm, which is used to control the population of clients allowed to participate with each server. Further information on this project, including a status report and set of briefing slides, is on the web at http://www.eecis.udel.edu/~mills/autonomous.html.

## 2.2 Collaboration Projects

There has been a continuous collaboration for many years between U Delaware, US Naval Observatory, US Coast Guard and NIST. This includes the exchange of new ideas on computer network synchronization, advice on configuring time server resources and, in the case of USCG and USNO, equipment grants and loans. Both NIST and USNO have installed additional Internet time servers at various locations in the U.S. The U Delaware contribution has been monitoring and measurement projects, as reported in a paper listed in the previous quarterly report. The available resources are deployed as follows:

There are presently six NIST servers located at four sites: NIST Boulder, CO, NIST Gaithersburg, MD, National Center for Atmospheric Research (NCAR) Boulder, CO, and Microsoft Corp., Redmond, WA. There are two redundant servers at the NIST Boulder and NIST Gaithersburg sites. All servers are based on Digital Alpha workstations and all except the NIST Boulder servers are synchronized by telephone modem and the Automated Computer Time Service (ACTS) service operated by NIST Boulder. The NIST Boulder servers are connected directly to the NIST master clock cesium ensemble.

There are presently eleven USNO servers located at eight sites: Massachusetts Institute of Technology, Cambridge, MA, Falcon AFB, CO, Washington University at St. Louis, MO, USNO Washington, DC, Georgia Institute of Technology, Atlanta, GA, UCLA Los Angeles, CA, University of Houston, TX, and Digital Equipment Corp., Palo Alto, CA. There are three redundant servers at the USNO Washington site and two at the Falcon AFB site. All are based on Hewlett Packard workstations and all except two of the USNO Washington servers are synchronized by GPS. The remaining two are connected directly to the USNO master clock cesium ensemble.

## 2.3 Infrastructure

The Austron 2200 GPS receiver donated by Delmarva Power has been installed and connected to a department file server. A Spectracom 8170 WWVB receiver and Hewlett Packard 5061A cesium clock is to be installed later. The equipment stack will duplicate the current stack used by the rackety.udel.edu public primary time server. The 2200 receiver has consistently misbehaved due to one problem or another. After a massive swap of circuit boards with its near-twin 2201 GPS receiver, the problem was traced to a bad oscillator module, which was replaced by the manufacturer. An additional problem was discovered in the antenna radome, which over time had become flooded with water. The water was drained, the interior circuitry cleaned and a weep hole drilled to prevent future flooding. Both of these expensive receivers have a long and dreary service history and both have on occasion spent months at the factory awaiting repair.

The TrueTime Corporation generously donated a NTS-100-GPS NTP time server, which was installed on the UDELnet campus network and made available for public access. The campus now has three public GPS primary servers, one on UDELnet and two on DCnet and three private GPS primary time servers, one each for the campus and backroom subnets and one for the UDel router between DCnet and DARTnet.

The massive overhaul and upgrade of personal web pages and various project web pages reported in previous quarters continues. At this time, all papers, reports and memoranda produced for the last eleven years is now in PostScript form on this investigator's home page www.eecis.udel.edu/~mills. Project descriptions and status reports have been updated for all project activity areas. Briefing slides for all meetings and most paper presentations are provided in HTML format produced from Windows Power Point presentations. Desktop publishing functions have been migrated from Corel Ventura Publisher to Frame Products FrameMaker, which has been installed on Unix and Intel workstations. Production of most publication-quality graphics has been migrated from AT&T S language to Math Works Matlab and Corel PhotoPaint.

## 2.4 Intruder Watch

As reported in the previous quarter, there have been three incidents involving apparently accidental intrusion on resources of the DCnet (128.4) research network. All three of these involved misconfigured hosts on other networks attempting to access nonworking addresses on the 128.4 network. All were reported to the CERT for archiving, but none required direct assistance or intervention by the CERT.

A closer watch with the tcpdump program at the backroom test site indicate the problem is much more serious than previously reported. There has been a continuous stream of misdirected IP datagrams to the 128.4.2 subnet, all resulting in spasms of ARP packets to nonexistent IP host addresses. In general, these pass unnoticed on networks operating at Ethernet speeds of 10 Mb/s, which is the case for most campus networks. However, the backroom subnet operates over a ISDN line, so that intrusions are much more likely to be noticed. On occasion, this has resulted in moderate deterioration in service to that subnet. As reported previously, identification of the source of the misdirected datagrams and correcting the configuration problems at the source requires a good deal of detective work.

## 2.5 Long Range Dependency

As reported for the last quarter, graduate student Qiong Lin has been collecting reams of data representing one-way delays in the global Internet. Intrepid NTP primary time server pogo.udel.edu has been configured to watch 23 other primary time servers on all continents except Antarctica (coming soon?). The peer-peer paths have very widely varying characteristics, some by domestic wire or fiber, some by undersea cable and some by satellite, but all are individually synchronized to UTC by radio, satellite or modem to within a millisecond or two.

After a preliminary round of processing of the one-way delay data, it was observed that some hosts showed clear patterns of long range dependency, while others did not. This was considered very curious, since most of the hosts involved are at considerable distance as the packet flies and many of the network paths overlapped. However, it has been observed in many similar experiments that the displayed data, in the form of scatter plots, are very sensitive to hidden resonances, missing data points and host failures. Accordingly, the delay data were processed again and intervals representing obvious host malfunctions were removed. Once this was done and the data reprocessed as before, all of the hosts showed the expected long range effects.

There remain a number of undetermined questions, such as whether the dependency originates in the hosts or in the network paths; we suspect the former. Plans are in progress to monitor the UDel web server and search its traffic characteristics looking for these effects.

## 2.6 Teaching

The course on cryptography techniques with computer network applications offered in Spring 1996 was well received. A seminar course in computer security is being offered in Spring 1997. The intent is to eventually develop a coordinated two-semester course in computer security, including both applied mathematics background and application to current Internet security models, architectures and protocols.

## 2.7 Meetings

A meeting was held with Prof. Scott Corson and representatives of the Army Research Laboratories in February 1997 to discuss possible collaboration in the area of routing algorithms.

## 3. Plans for the Next Quarter

Our plans for the next quarter include continued development of the NTP Version 4 protocol model, specification and implementation. Specifically, we plan to begin implementation of the new authentication scheme and integration with the current NTP Version 3 daemon for Unix and Windows. The daemon is to be tested first in the research net, then the DARTnet/CAIRN community. As the extensions are backwards compatible, the new features can be activated and tested in regular operation without impacting current users.

We plan to complete the design and implementation of the new autonomous configuration scheme, as well as the theoretical analysis, simulation and experimental justification of the scheme, as represented by Mr. Thyagajan's dissertation.

The work on long-range dependency is expected to grow as our understanding of the interesting phenomena grows.

## 4. Publications

Mills, D.L. The network computer as precision timekeeper. Proc. Precision Time and Time Interval (PTTI) Meeting (Reston VA, December 1996), to appear in print. URL: see http://www.eecis.udel.edu/~mills/papers.html.

Abstract

This paper describes algorithms to discipline a computer clock to a source of standard time, such as a GPS receiver or another computer synchronized to such a source. The algorithms are designed for use in the Network Time Protocol (NTP), which is used to synchronize computer clocks in the global Internet. They have been incorporated in the NTP software for Unix and Windows and, for the highest accuracy, in the operating system kernels for Sun, DEC and HP workstations. RMS errors on LANs are usually less than 10 ms and on global Internet paths usually less than 5 ms. However, rare disruptions of one kind or another can cause error spikes up to 100 ms on LANs and 100 ms on Internet paths.

Mills, D.L. Authentication scheme for distributed, ubiquitous, real-time protocols. Proc. Advanced Telecommunications/Information Distribution Research Program (ATIRP) Conference (College Park MD, January 1997), 293-298. URL: see http://www.eecis.udel.edu/~mills/papers.html.

Abstract

Cryptographic authentication methodology proposed for use in the Internet require substantial resources when very large client populations are involved. Resource provisioning becomes especially important when time-critical services are involved. In the cast of time-synchronization services, a special case exists, since cryptographic keys must enforce valid

lifetimes, but validating key lifetimes requires cryptographic keys. This paper proposes a scheme which minimizes server resources while resolving the apparent circularity.

Mills, D.L. Clock discipline algorithms for the Network Time Protocol Version 4. Electrical Engineering Report 97-3-3, University of Delaware, March 1997, 35 pp. URL: see http://www.eecis.udel.edu/~mills/reports.html.

Abstract

This report describes the analysis, implementation and performance of engineered algorithms to discipline a computer clock to a source of standard time, such as a GPS receiver or another computer synchronized to such a source. The algorithms are designed for the Network Time Protocol (NTP) Version 4, the successor to NTP Version 3, which is in widespread use to synchronize computer clocks in the global Internet. The report includes an overview of the new NTP architecture and process decomposition as related to the clock discipline algorithm, which is implemented as a hybrid phase/frequency-lock feedback loop. An extensive engineering analysis of this algorithm is presented along with the results of a detailed simulation to evaluate and validate its performance using both synthetic data and real measurements with Internet time servers in Europe, Asia and the Americas.