Scalable, High Speed, Internet Time Synchronization

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

Quarterly Progress Report
1 March 1996 - 31 May 1996

David L. Mills
Electrical Engineering Department
University of Delaware

1. Introduction

This report covers the work done in support of the Information Technology Office of the DARPA on reliable, accurate network time synchronization. Contributors to this effort include Prof. David L. Mills and graduate students, Ajit Thyagarajan and Bradley Cain.

The Network Time Protocol (NTP) is widely used in the Internet to synchronize computer clocks. It is supported by over 100 primary time sources located in several countries, together with a hierarchical network of over 100,000 servers and clients scattered all over the globe. Management and configuration of this network has become almost unworkable. As the network is still growing rapidly, a means is required to automatically configure the hierarchy in response to changing topology and available server and network resources.The approach to this work involves distributed algorithms that collect information on the current timekeeping network topology using a combination of multicast and anycast messaging, directory services and network management resources. Existing and proposed service location protocols are expected to play a part as well. Distributed algorithms are used to process these data and construct a forest of spanning trees rooted on the primary servers (those synchronized to radio clocks or modem services. These algorithms attempt to maximize a metric according to the most accurate time, subject to constraints designed to protect the server and network resources.Recent projects reported in papers, technical reports, project reports and technical memoranda include precision timekeeping devices for Sun Microsystems workstations, improved clock discipline algorithms, and protocol upgrades for the Network Time Protocol (NTP). Other 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 and distributed to the research community includes the NTP Version 3 implementation for Unix and Windows and a set of precision-time kernel modifications for major Unix workstation manufacturers. Finally, ongoing 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

There are a number of ongoing research projects funded by DARPA, US Army, US Navy and NSF which mutually support each other. Current projects can be loosely grouped into four areas corresponding to the primary emphasis in work proposed to each of these agencies. The work reported below on autonomous configuration is the prime focus of DARPA funding; the work on cryptographic authentication is the prime focus of US Army funding; the work on precision syn-

chronization devices and kernel modifications is the prime focus of US Navy funding; the work on algorithms and service deployment is the prime focus of NSF funding. Since these projects interact with each other in important ways, they will be summarized as a whole.

2.1 Autonomous Configuration

Present progress in autonomous configuration includes refinements to client and server software to distribute time information via multicast and anycast messaging. An authentication scheme is under development to insure certifiable traceability to national standard time sources. A preliminary version of the new multicast/anycast and authentication schemes has been adapted for the Simple Network Time Protocol (SNTP), which is a subset of NTP. A revised specification document (RFC) has been completed and now in circulation for review and comment.

We have performed an in-depth analysis of the autonomous configuration problem from both a theoretical and practical standpoint. A literature survey reveals that this kind of problem has not been well-studied; however, “good” approximation algorithms have been developed for related problems. We have developed an algorithm with guaranteed worst-case running times within log(N) of the optimal with arbitrary network configurations, where N is the number of nodes. However, the algorithm performs much better with typical network configurations likely to be found in practice. Still, for large N in the order of hundreds or thousands, the performance is poor.

If additional constraints are imposed, such as assuming that the network nodes correspond to points in a geometric space with defined distance metric, the problem becomes easier to analyze. These assumptions are not unreasonable, since most real networks do have such characteristics. We have developed an algorithm which guarantees that the resulting the solution will be no greater than about six times the optimal solution time. However, there is still considerable room for improvement.

The ultimate objective in this effort is to develop a distributed algorithm which, operating with little or no topological information, automatically configures a network of hundreds or thousands of clients and servers. It is not likely that a solution can be found with optimum running time or even with polynomial running time. We have therefore focused our efforts on developing efficient heuristics that give near-optimal solutions most of the time and “poor” solutions rarely.

We have developed a candidate heuristic algorithm which optimizes the NTP synchronization hierarchy, subject to resource constraints. It uses a divide-and-conquer approach which enables an efficient implementation as well. The details of protocol are currently being worked out and we expect a complete description of the semantics shortly. To support this work, Mr. Thyagarajan and Mr. Cain have designed and implemented enhancements to current Internet multicasting support and distributed these enhancements throughout the Internet research community. These developments have been reported at a recent IETF meeting.

Changes to the current NTP specification and implementation are required to implement this algorithm. Changes to several algorithms used to associate clients and servers have been designed and documented. The new anycast mode, similar to that used to discover nearby servers in other similar protocols, has been incorporated in the design. A new distributed mode of operation, with objectives similar to the matrix time synchronization method described in the literature, has been incorporated as well. A capsule summary of current activity in the autonomous configuration

area, as well as a set of briefing slides, is in the web page http://www.eecis.udel.edu/~mills/status.html. These pages are updated on a routine basis.

2.2 Cryptographic Authentication

NTP contains provisions to cryptographically authenticate individual servers as described in the protocol specification; however, the existing protocol model does not provide a scheme for the distribution of cryptographic keys, nor does it provide for the retrieval of cryptographic certificates that reliably bind the intended server identification data with the associated keys and related public values.

However, using conventional key agreement and digital signatures with time-critical applications and large client populations such as NTP can cause significant performance degradations. In addition, there are problems unique to NTP in the interaction between the authentication and synchronization functions, since reliable key distribution requires reliable lifetime control and good timekeeping, while secure timekeeping requires reliable key distribution.

Our current work is designed to produce a cryptographically sound and efficient methodology for use in NTP and similar distributed protocols. A literature review of cryptographic techniques and related mathematics has been completed. A detailed assessment of existing schemes proposed by various IETF task forces, specifically the IPSEC community, has been completed. Preliminary evaluation of SKIP, Photuris and ISAKMP schemes suggest that a scheme which is cryptographically sound, efficient and interoperable with existing and proposed directory and certificate services is a challenging task.

As a direct result of this research, a syllabus has been constructed for a graduate-level class in cryptographic theory and practice. This course was taught for the first time to a class of combined Electrical Engineering and Computer and Information Science students.

A report detailing analysis of the NTP security model and authentication scheme is nearing completion. A revised security model and strawman authentication scheme is in progress. A capsule summary of current activity in the cryptographic authentication area, as well as a set of briefing slides, is in the web page http://www.eecis.udel.edu/~mills/status.html. These pages are updated on a routine basis.

2.3 Software Upgrades

Much of the grunt activity during the recent quarter was major software upgrades to the operating systems for the workstations on our research network DCnet (128.4). These machines have grown an incredible degree of complexity with respect to connected radio clocks, PPS signals, network technology and various kernel modifications. The SunOS machines were upgraded to SunOS 4.1.3 and a DEC MIPS machine to Ultrix 4.4. In addition, minor upgrades were done to Solaris 2.4 systems and Digital OSF/1 systems. Some machines were reconfigured on Ethernet, FDDI and DS3 networks and the entire routing architecture was updated. As some of these systems had not been upgraded for several years, many things required intricate and time consuming repair. Maybe the worst thing was the PI had to move himself, office workstations and network umbilicals to a larger office.

We have installed one-time passwords (S/KEY) and secure shell, tightened RPC and TCP source verification, and activated Kerberos. While Kerberos support is running on most machines, we do not yet have specific services configured to use it. One of the reasons for this effort is the observed increase in attempted breakins, as evidence from source verification monitors. In fact, there are currently a number of domains from which such attacks have occured and are now denied access.

2.4 Hardware Upgrades

The local power utility, Delmarva Power, kindly donated a used Austron 2200 GPS receiver to serve as backup for our Austron 2201A receiver, which has timed the tick for most of our campus NTP primary servers for several years. The receiver was returned to the factory for refurbishment and upgrade to the 2200A version. Upon return to us, it was found defective and again returned for repair. The intent is to use it in a mutually redundant configuration where each receiver can back up the other. Backup is now handled by a 16-year old Spectracom 8170 WWVB receiver, which has been plagued by local EMI from undetermined source. However, it is our understanding that the WWVB transmitter is being upgraded to increase power, which may revive the 8170 as respectable backup.

The Arbiter company donated a 1088A GPS receiver, which has been installed at the Backroom test site. This site, already equipped with WWV, CHU and WWVB receivers, is in a somewhat quieter EMI environment than campus and is used for software development, in particular, development of radio clock drivers. A driver for the Arbiter clock has been completed and integrated in the NTP software distribution.

2.5 Precision Time Kernel Upgrades

The precision time kernel modifications for Sun and Digital workstations have been upgraded in response to recent upgrades in the Unix kernels supplied with these machines. New versions were completed, tested and supplied to those manufacturers. HP seems to have lost some enthusiasm and has not yet volunteered system and sources upgrade to current HP-UX 10.0. We have very good rapport with Sun and Digital and usually get upgrades in a timely manner.

We have recently received system and sources upgrade to Digital Unix 4.0 for our two Alphas. This system has incorporated all of the precision time kernel modifications and NTP Version 3 in the standard product, except for PPS signal support; however, the features must be enabled during kernel configuration. Proposals for adding PPS support are pending at Digital. We would like to see a common standard interface for PPS signals in all major workstations. We regard this as a major coup, at least for Digital.

In response to some concerns about the stability of the kernel modifications in extreme cases of disruption due to misconfigured kernels, improper operation of the synchronization daemon, or electronic warfare attack, a detailed simulation of the local clock discipline algorithm was implemented in Matlab with the Control Systems Toolbox. This is an extension of a prior simulation implemented in Mathmatica. The Matlab simulation was configured in both continuous and discrete modes and with maximum worst-case parameters and loop delays. The results show clearly that the algorithm is stable for delays up to eight times the update interval and loop compensation parameter mismatches from half the optimum value to twice this value.

A massive overhaul and upgrade of the NTP distribution to support GNU autoconfigure is nearing completion. This will provide a much more reliable software configure and build environment, as well as much more convenient porting to other architectures and operating systems. As part of this upgrade, the documentation was completely rewritten and recast as HTML pages suitable for browsing. The document pages are included in the NTP distribution for Unix and Windows, as well as in the NTP web site http://www.eecis.udel.edu/~ntp.

There has been increasing interest in using the NTP service model and architecture in other protocol stacks, in particular, IPv6 and OSI. In general, NTP can work in an IPv6 stack as it does now with only a minor change in header format. The change involves increasing the size of the reference identifier field, which holds the server address currently selected for synchronization. This field must be increased from four octets for IPv4 to 16 octets for IPv6. In the case of OSI, NTP can work in the stack just above the connectionless transport layer with a reference identifier field size of 24 octets to hold a full NSAP and length indicator . Detailed specifications on how to do this are in the RFC draft “Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI,” which is now in preparation.

2.6 NTP Version 4 Progress

We have completed analysis and design of association matching procedures necessary for NTP Version 4. The NTP Version 3 procedures require matching only on the source address and destination address. This caused problems and confusion in cases where the same two peers can operate in two different modes at the same time, as when calibrating the propagation delay in multicast mode. The new design requires a match on the mode as well.

A new mode called anycast has been designed for local area server discovery. In this mode a client sends an ordinary client/server request to the NTP multicast group address. Cooperating servers in scope reply to this request using their ordinary unicast address as the source address. The client binds to the first reply received and ignores subsequent replies from the same or other servers and continues operation using only unicast addresses.

The literature review during preparation for a new course in cryptography theory and practice suggested several new angles of attack on the problem of cryptographic server authentication in NTP multicast mode. In particular, a set of extensions to the present NTP authentication scheme has been designed. This and a detailed security analysis of NTP is in preparation as a technical report for early publication.

2.7 Pentium PC Hardware and Software Upgrades

We have purchased a Pentium PC to serve as a development platform and general resource machine. It normally runs a Windows 95 disk partition with the usual complement of word and image processing applications and presentation managers. It is equipped with a color scanner, which we use to scan in pictures and artwork for the web pages. We also replaced an ageing PostScript printer with a new model and installed a low-end color printer for presentation hardcopy and overhead slides.

A major use for the PC is as development platform for DARTnet router software. The SPARC-based DARTnet routers are to be replaced by PCs running FreeBSD, so the PC system is configured with a FreeBSD disk partition and program development utilities, including a port of NTP. In addition, an older PC has been upgraded to Windows 95 and two new Pentium PCs running Windows 95 have been installed at the Backroom test site at no cost to the Government. All of these machines are connected to the DCnet research network via Ethernet and ISDN media and all share file systems with other machines on the network.

One or another of the four PCs provides networked access for a full suite of applications, including desktop publishing (Corel Ventura and FrameMaker), program development (Microsoft Visual C++, Visual Basic and Fortan), relational database (Borland Visual dBASE), office systems (Microsoft Office and Lotus Smart Suite), schematic capture and PC board layout (Accel Technologies Tango-PRO), image manipulation and editing (Corel PhotoPaint), and symbolic mathematics (Wolfram Mathematica and MathSoft MathCAD). Most of this software was provided at no cost to the Government.

In addition to the PC-based applications, we have secured Unix workstation licenses for schematic capture and PC board layout (Cadence), circuit simulation), and continuous and discrete system simulation (MathWorks Matlab). We have campus licenses for most Sun and Digital workstation software as well.

2.8 NTP Survey

In an effort to determine the extent of NTP service and its performance in the Internet, we have designed and are in process of a survey of all reachable NTP sites. This task is conducted as a background process so as not to affect traffic statistics in any measurable way. The intent is to determine reachability, time and frequency error distributions, stratum-level histograms and other related data. To date, we have found some 20,000 clients and servers running NTP on the Internet and are daily discovering more. However, this is only a fraction of the suspect population, as many sites are beyond firewalls, some use the RPC program ntpdate and some synchronize only among their own group. None of these sites are visible using the available monitoring tools and purpose-built detective kit.

We are still in process of collecting data, but expect to complete the program in the next month. Some of the preliminary results learned, including histograms of time and frequency residuals, have already been posted to briefing slides on the web http://www.eecis.udel.edu/~mills/status. When complete, an analysis and evaluation will be prepared as a report.

2.9 Web Pages Overhaul

A massive overhaul and upgrade of personal web pages and various project web pages has been completed. This includes research status reports (not administrative status reports like this one), briefings and all papers, technical reports, project reports and technical memoranda published by our laboratory in the last ten years. Maintenance of the huge amount of this stuff is a continuing burden; however, we believe this to be an effective technical disseminations channel, as well as an opportunity for us to advertise the university, department and our laboratory to prospective students and staff. Additional information can be found in the http://www.eecis.udel.edu/~mills and http://www.eecis.udel.edu/~millslab web pages.

2.10 Graduate Course in Cryptography Theory and Practice

A major project for the Spring semester was a new graduate level course in cryptography theory and practice. This turned out to be more popular than expected, with about twenty students enrolled from the Electrical Engineering and Computer and Information Sciences departments. Topics included the mathematical foundations of cryptography and practical techniques using private key and public key cryptosystems, one-way hash functions and random number generators. A syllabus can be found in the http://www.eecis.udel.edu/~mills/teaching.html web page.

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. We plan to analyze the autonomous configuration algorithm using empirical as well as probabilistic techniques. These methods have been successfully applied to problems such as the travelling salesman problem. It remains to be seen whether this analysis can be extended to a broader class of problems. Specifically, we plan to code the algorithm and begin integration with the current NTP Version 3 implementation for Unix and Windows. As the extensions are backwards compatible, the new features can be activated and tested in regular operation without impacting current users.

As mentioned previously, we plan to complete the SNTP Version 4 specification as an RFC, as well as the NTP security model, by the end of the next quarter. In addition, the NTP survey is to be completed and documented.

We also plan to complete the design and implementation of the new distributed mode feature, in which a mob of servers exchange timestamps, in order to construct a complete model of time offsets, in order to provide ensemble statistics and global clock steering, as well as mutual backup and robustness in the face of electronic warfare attack.