Survivable, Real Time Network Services

Defense Advanced Research Projects Agency
Contract F30602-98-1-0225, DARPA Order G409

Quarterly Progress Report
1 July 1999 - 30 September 1999

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 Qoing Li and Robert Redwinski. The project continues previous research in network time synchronization technology jointly funded by DARPA and US Navy. 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. Specific applications benefiting from this research include multicast topologies, multimedia, real-time conferencing, cryptographic systems, and management of distributed, real-time systems.

This quarterly report is submitted in traditional report form on paper. As the transition to web-based information dissemination of research results continues, almost all status information and progress reporting is now on the web, either on pages belonging to the principal investigator or to his students. Accordingly, this and future progress reports will contain primarily schedule and milestone data; current status and research results are reported on web pages at www.eecis.udel.edu/~mills in the form of papers, technical reports and specific briefings.

1.1 NTP Version 4

Work continued on the NTP Version 4 reference implementation and distribution for Unix, VMS and Windows. We have incorporated numerous patches required by the over two dozen ports of the code to various Unix architectures and operating systems. Following is a partial list of improvements.

1. The autoconfigure system has been considerably enhanced to handle literally dozens of operating systems, each with multiple versions and updates and some with bugs requiring contrived workarounds.

2. The development environment has been upgraded to support a revision control system. Repositories have been identified and a certain number of volunteer codecutters identified and authorized.

3. The clock discipline loop code has been enhanced to operate with the new nanokernel discipline, as well as interoperate with legacy versions of the NTP code and prior kernel modifications.

4. The clock filter and mitigation algorithms have been updated as the result of experience with infrequently connected ISDN and dialup links.

5. Several new device drivers have been contributed for a total of 34 drivers supporting over 40 receivers for every computer-readable time dissemination service in the world.

6. A shakedown review of the issues affecting year-2000 performance was completed by AT&T and several minor patches incorporated in the code.

7. Several changes were made to support ongoing autokey and autoconfigure facilities. One of the more crucial changes was to support multiple logical interfaces for a single physical interface.

8. The cryptographic routines other than MD5 have been moved to an optional directory and the software autoconfigure system modified to detect whether the routines are actually present and, if so, to compile and link them in the ntpd binary. This was done to simplify export procedures, with the expectation that the overseas user could incorporate the routines from local sources.

We have studied the characteristics of typical computer oscillators in detail and developed analytical models that predict their performance and heuristic algorithms that automatically optimize the clock discipline algorithm parameters such as time constant, selection of frequency or phase as the discipline method and so forth. We have implemented a special purpose discrete event simulator included most of the algorithms used in NTP and explored performance with different algorithms and tuning parameters.

The result of experiments with the simulator include optimized algorithms and tuning parameters for the NTP daemon and insight into how to implement frequency and phase steering using precision timing signals generated by some radios and laboratory devices such as cesium oscillators. A status report and briefing slides on this work are at www.eecis.udel.edu/~mills/status.htm.

2. Pulse-per-second Application Program Interface

Over the years a number of ad-hoc interfaces for the pulse-per-second signals generated by some radio clocks and laboratory instruments have been implemented. Perhaps the most widely used was the "ppsclock" STREAMS module for SunOS 4.1.x designed by Van Jacobson and Craig Leres at Lawrence Berkeley Laboratory. It uses the carrier detect (DCD) modem control signal lead and level converter such as the “gadget box” described at www.eecis.udel.edu/~mills/resource.htm. However, this module works only with the console device driver in SunOS 4.1.x and is not easily ported to other hardware and operating systems.

PPS interfaces have been implemented for other hardware and operating systems, including Ultrix, Alpha, FreeBSD and Linux; however, these all use different application program interfaces (API) and in general do not conform to the same model. In an effort to specify, design and implement a common PPS interfacing standard, a group including engineers from Sun, Digital (Compaq), FreeBSD and Linux collaborated on a proposed IETF standard API. The specification has evolved over five versions and reference implementations. It has been implemented and tested in SunOS 4.1.3 and in the final stages of implementation and test in Digital Unix, FreeBSD and

Linux. However, difficulties remain with Solaris due the extreme complexity of the kernel architecture and kernel build process.

3. Miscellany

We are planning for an Abilene connection in the near future. This will be implemented using a 100-Mbps fiber strand to the Cisco router connected to Abilene. An IP-in-IP tunnel will be used to a CAIRN/ABONE PoP. To this end we reconfigured all DCnet (128.4) subnets to the Torrent router provided on indefinite loan by the manufacturer. This is a high performance unit with numerous 10-Mbps and 100-Mbps copper and fiber interfaces that should handle any conceivable expansion of the network, which now supports four research projects funded by NSF and NASA for other principal investigators in the Electrical and Computer Engineering Department and Computer and Information Sciences Department.

Both of our Austron GPS receivers failed their master oscillators, which use a temperature stabilized quartz crystal. As this component is prohibitively expensive to repair, both receivers were junked. To replace them, we have purchased a Spectracom GPS receiver and re-hosted an identical receiver previously donated by Spectracom. We plan to move all timing equipment, now scattered over three rooms, into the machine room in a common rack. There will be two identical stacks consisting of a SPARC time/file server, Spectracom GPS receiver as primary source, Spectracom WWVB receiver as backup source and various serial line breakout boxes. One stack will be used for public access, the other for DCnet and affiliated networks access and each stack will back up the other. This project also involves reconfigured and simplified connections to several serial line console ports for other machines and CSU/DSU units.

A cesium oscillator has failed the beam tube, which is at the heart of the instrument. This particular instrument was donated by the U.S. Naval Observatory and has already had one tube replacement. In order to reduce the cost of replacement, a used tube was provided. A new tube has a lifetime of about ten years; apparently, the used tube was near end life and failed after only about one year of operation. As used tubes cost about $1000 and new tubes about $15,000 selection of a good used tube with significant lifetime remaining is of prime importance.

4. Personnel

Tamal Basu joined the project in September 1999. He is currently funded half time by the DARPA project and half time by another NSF project. He is to continue the work initiated by Mr. Redwinski; in particular, extending the analytical model and routing protocol test suite. One of the things we are looking for in the work with very large networks is whether the behavior in the face of hostile attack has a precipitous threshold where, under conditions of escalating attack and destruction of network nodes or links, the routing algorithm experiences catastrophic failure. We are particularly interested in the behavior near the vicinity of this threshold and what mechanisms can be designed to mitigate this behavior. Once upon a time, Vinton Cerf described the situation in the then-ARPAnet as “waves of congestion washing from coast to coast”. This situation was also suspected to occur in the NSFnet backbone network deployed in the 1986-1988 time frame.

Ajit Thyagarajan continues work to complete his dissertation. His topic is the analytical and experimental study of autoconfigure algorithms suitable for deployment in a survivable internet.

He is finishing up extensions of the current centralized algorithms to operate in a distributed context. He expects to complete all requirements by the end of the Fall semester 1999.

David Mills was elected to IEEE Senior Member status in August 1999.

5. Meetings

Following is a list of presentations and meetings made during the current period of performance. Note that a full complement of briefing slides is archived at www.eecis.udel.edu/~mills/colloq.htm. In most cases the slides are organized as an extended briefing designed to be read via the web, but not all the slides were used at the presentation. Thus, a considerable amount of extra information beyond the oral presentation is available.

The Mars Internet Workshop was held at the University on 15 July 1999. Among the topics discussed was a transport mode called by this investigator the Pony Express Transport Service, in which messages are aggregated and launched at the origin and travel to the destination without explicit acknowledgement. Also, a routing paradigm called Hail Mary was proposed in which destination name-address translation are done only upon arrival at a planetary gateway. This investigator announced his intention of becoming a founding member of an Interplanetary Internet Chapter of the Internet Society.

A particularly noteworthy event was the SIGCOMM Tutorial on the Technical History of the Internet on 31 August. This investigator presented two briefings, one on the early history of routing protocols leading to the establishment of the Exterior Gateway Protocol described in RFC-904. The other briefing was on the NSFnet backbone network experience with routing and congestion instabilities.

6. Plans for Next Quarter

High on the list of activities for the next quarter is finishing up the autokey code, in particular updating the NTP Extension Field parsing code to support new field types. In addition, cryptographic interfaces and code needs to be built to generate cryptographic key pairs and provide the public values to autokey clients.

A design for a fully distributed autoconfigure scheme is taking shape. It involves revisions in the automatic load leveler scheme and in the expanding ring server search. Briefly, these involve controlling load using aggregate packet rates rather than explicit client identification, and restricting Li, Qiong, and D.L. Mills. Investigating the scaling behavior, crossover and anti-persistence of internet packet delay dynamics. Proc. 1999 IEEE GLOBECOM 99 Symposium, (Rio de Janerio, Brazil, December 1999) (to appear - please do note cite or redistribute before publication).

7. Publications

All publications, including journal articles, symposium papers, technical reports and memoranda are now on the web at www.eecis.udel.edu/~mills. Links to the several publication lists are available on that page, as well as links to all project descriptions, status reports and briefings. All publications are available in PostScript and PDF formats. Briefings are available in HTML, PostScript, PDF and PowerPoint. The project descriptions are cross-indexed so that the various interrelationships are clearly evident. Also included are the documentation pages for various pub-

lic software distributions. Links to other related projects at Delaware and elsewhere are also included on the various pages. Hopefully, the organization of these pages, which amount to a total of about 300 megabytes of information pages and reference documents, will allow quick access to the latest results and project status in a timely way.

Following is a retrospective list of papers and reports supported wholly or in part on this project and the immediately preceeding project "Scalable, High Speed, Internet Time Synchronization," DARPA Order D012. The complete text of all papers and reports, as well as project briefings, status reports and supporting materials is at www.eecis.udel.edu/~mills.

7.1 Papers

1. Mills, D.L. Cryptographic authentication for real-time network protocols. In: *AMS DIMACS Series in Discrete Mathematics and Theoretical Computer Science, Vol. 45* (1999), 135-144.

2. Mills, D.L. Adaptive hybrid clock discipline algorithm for the Network Time Protocol. *IEEE/ACM Trans. Networking 6, 5* (October 1998), 505-514.

3. Li, Qiong, and D.L. Mills. On the long-range dependence of packet round-trip delays in Internet. *Proc. IEEE International Conference on Communications* (Atlanta GA, June 1998), 1185-1191.

4. Mills, D.L., A. Thyagarajan and B.C. Huffman. Internet timekeeping around the globe. *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Long Beach CA, December 1997), 365-371.

5. 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.

6. Mills, D.L. The network computer as precision timekeeper. *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Reston VA, December 1996), 96-108.

7. Mills, D.L. Improved algorithms for synchronizing computer network clocks. *IEEE/ACM Trans. Networks 3, 3* (June 1995), 245-254.

7.2 Technical Reports

8. Sethi, A.S., H. Gao, and D.L. Mills. Management of the Network Time Protocol (NTP) with SNMP. Computer and Information Sciences Report 98-09, University of Delaware, November 1997, 32 pp.

9. Mills, D.L. A precision radio clock for WWV transmissions. Electrical Engineering Report 97-8-1, University of Delaware, August 1997, 25 pp.

10. 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.

11. Mills, D.L. Proposed authentication enhancements for the Network Time Protocol version 4. Electrical Engineering Report 96-10-3, University of Delaware, October 1996, 36 pp.

12. Mills, D.L. Simple network time protocol (SNTP) version 4 for IPv4, IPv6 and OSI. Network Working Group Report RFC-2030, University of Delaware, October 1996, 18 pp.

13. Mills, D.L. Simple Network Time Protocol (SNTP). Network Working Group Report RFC-1769, University of Delaware, March 1995, 14 pp.

14. Mills, D.L. Simple Network Time Protocol (SNTP). Network Working Group Report RFC-1769, University of Delaware, March 1995, 14 pp.

7.3 Internet Drafts

15. Mills, D. L., T.S. Glassey and M.E. McNeill. Authentication Scheme Extensions to NTP. Internet Draft draft-mills-ntp-auth-coexist-01.txt, IETF, September, 1998.

16. Mogul, J., D.L. Mills, J. Brittenson, J. Stone and U. Windl. Pulse-per-Second API for UNIX-like Operating Systems, Version 1.0. Internet Draft draft-mogul-pps-api-05.txt, IETF, August 1999, 25 pp.