Survivable, Real Time Network Services

Defense Advanced Research Projects Agency
Contract F30602-98-1-0225, DARPA Order G409/J175

Quarterly Progress Report
1 July 2001 - 30 September 2001

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 student Harish Nair. MS student Tamal Basu has graduated and taken a position at Acorn Networks. Phd student Qiong Li has graduated and taken a position with Phillips Research.

The project continues previous research in network time synchronization technology funded by DARPA, US Army and NASA/JPL. 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 and sensor networks.

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.

2. Autonomous Authentication

The IETF STIME group has reviewed the latest Internet Draft and suggested a number of changes and additions. Some concern was expressed about the cookie used in client/server mode; in particular, the relative ease an intruder can intercept it and use it to disrupt protocol operations. As mentioned in the last report, the Autokey protocol client/server mode operations were modified to include a Diffie-Hellman agreement exchange where the agreed key is used as a stream cipher to encrypt and decrypt the cookie. The agreed key is discarded after use, so the server remains stateless. The scheme has been implemented and tested in the Autokey Version 2 protocol.

However, after testing with some older architectures, in particular a SPARC IPC running SunOS 4.1.3, the Diffie-Hellman agreement proved a very bad idea. The running time for a complete agreement computation could be as long as fifteen seconds, which certainly is a showstopper for a time synchronization protocol. Accordingly, the scheme was modified so that the initiator sent its public key, the responder rolled a random value and returned the encrypted value to the initiator. The original scheme was intended to conform to the strict interpretation of export controls that forbid encryption in any context except in connection with digital signatures, so technically the revised scheme does not conform to that model. As a practical matter, it probably makes little difference.

With the revised cookie scheme, there is no need for Diffie-Hellman parameters, which avoids another problem where formerly the most recent parameters would flood from servers to dependent clients, but not the other way. However, it does require some fancy footwork in symmetric modes when both peers attempt to initiate a cookie exchange at the same time. If both peers happen to send a cookie request before either of them has received the cookie response from the other peer, each will send a cookie response with its encrypted random value. The result is that each peer will have two cookies, one it generated and the other generated by the other peer. In this case, each peer computes the working cookie as the exclusive-OR of the two cookies.

A specific deficiency mentioned in the previous report is the lack of support for the Digital Signature Standard (DSS). This has been remedied and the standard distribution works for all digest/ signature schemes supported by OpenSSL.

The current status of the Autokey protocol and implementation is available via the web at www.eecis.udel.edu/~mills/autokey.htm.

3. Autonomous Configuration

Work continues on the development and test of the Manycast autonomous configuration scheme. In this scheme a multicast client sends an ordinary request packet to a multicast group address and potential multicast servers in TTL range have the opportunity to respond with an ordinary unicast message. Each received message causes the client to mobilize an ephemeral association, which then proceeds as if the server were originally explicitly configured. As mentioned in the last report, this scheme is working well in tests both with CAIRN routers and in our local environment. There is one caveat, however; some of the CAIRN routers are still running FreeBSD 3.4, in which the OpenSSL library does not link correctly. We have verified the problem does not exist in FreeBSD 4.3, but only a few of the CAIRN routers have this version. It is expected that most if not all CAIRN routers will be upgraded to the later version soon.

As mentioned in the last report, there was a problem in some configurations where the server selection code could frequently change the selections among servers of similar quality or “clock-hop”. This is particularly burdensome, since every change requires the client association to restart the Autokey protocol from the beginning. The solution was to incorporate a measure of hysteresis in the selection algorithm in the following way. When an update is received for a server already in the clique of three servers which discipline the clock, an association variable is initialized and subsequently allowed to decay exponentially. In the clustering algorithm which determines the clique membership this variable contributes to the distance metric used to define the members. The effect is to selectively favor the existing members unless one of them dies or another server not already a member shows much better performance. In such cases the newcomer can displace another member already in the clique.

The performance of the modified scheme is highly satisfying. Typically, a clique forms when a client first comes up and completes the Autokey dance, then holds its members unless something bad happens or until the association restarts for some reason, such as when a server rolls a new private value.

The current status of the Manycast mode and implementation is available via the web at www.eecis.udel.edu/~mills/autocfg.htm.

4. NTP and IPv6

The folks at Viagenie in Canada have completed an initial port of the NTPv4 daemon code for IPv4 to include IPv6 functionality. The project is ongoing and expected to take some time to fix up the utilities and, in particular, Autokey. However, in its present form an IPv4 configured client works with a IPv4 server and a IPv6 configured client works with a IPv6 server. The goal is to provide intermixed IPv4 and IPv6 service from a single server. We expect to continue with this project in a collaboration with Viagenie and UCL.

5. Timekeeping in the Interplanetary Internet

With support from Jet Propulsion Laboratories (JPL) and NASA, we have begun a study of timekeeping issues in the Interplanetary Internet, specifically for Mars missions and supporting infrastructure. Currently, the clocks for all space vehicles are coordinated on Earth using a software library called SPICE. With new technology, timekeeping in space would be distributed so that timed experiments could be accurately controlled autonomously and without intervention from Earth. There are three issues which distance the new technology from current practice: bodies are moving, connectivity between them is necessarily not continuous and transmission delays can be very long.

The Mars internet includes three segments. The Earth segment consists of the current Internet, including experiment hosts and NASA Deep Space Network (DSN) earth station gateways in California, Spain and Australia. The Mars segment consists of orbiters, base stations and rovers near and on Mars. The DSN segment consists of the space between DSN gateways, spacecraft transport buses and Mars orbiters. The orbiters function as gateways, experiment data buffers and supporting infrastructure, including time synchronization.

The current NTP technology has no provisions for mobile servers and clients, where range and range rates can vary with time, and only minimal provisions for intermittent connectivity. In the Mars internet, orbiters and surface stations may have only intermittent connectivity, while in the DSN segment real-time connectivity is possible only at scheduled opportunities and then only with very long delays. These considerations are mitigated by the fact that ranges and range rates can be predicted with some accuracy from the known positions of the spacecraft bus, orbiters and surface stations using ephemerides maintained by astronomical means.

The mission of this project is to develop an understanding of the timekeeping issues in the Interplanetary Internet, including Mars missions and beyond. The technology must be able to

1. Extract ephemerides data from mission housekeeping files both on Earth, the spacecraft bus and Mars orbiters to compute range and range rate for useful transmission opportunities.

2. Enhance the NTP protocol to encode and transmit ephemerides and/or orbit element data along with NTP timestamp data.

3. Develop new or modified algorithms using these data to determine accurate time and frequency offsets between moving objects.
4. Demonstrate a proof-of-concept in the form of a space simulator using suitably modified components of the existing NTP software distribution embedded in a special purpose discrete time simulator.

The current status of the Interplanetary Internet timekeeping project is available via the web at www.eecis.udel.edu/~mills/ipin.htm.

6. Infrastructure

There have been a continuous stream of contributions to the NTP software distribution from the developer's group. These include further refinements to the clock filter and clock discipline algorithm, new reference clock drivers and the huff-n'-puff filter mentioned in the last report. The entire distribution has been re-hosted in Bitkeeper technology as recommended by Sun and included in at least one mirror site at bitkeeper.org. There are now three public distributions: the current stable version, a development version incrementing from that version and a IPv6 development version incrementing from the development version.

Hacker interference continues to be a serious problem. In the most recent attack, one of our FreeBSD machines was subverted with code to generate IP packets as fast as the machine would go, resulting in a 100-Mb spray that consumed all bandwidth in the department net, campus net and most of the available bandwidth beyond our campus Abilene gateway. While it was quickly isolated and shut down, the experience prompted yet another internal review of defense mechanisms.

The attacker exploited a known bug in FreeBSD telnet daemon, so telnet service was immediately disabled in all DCnet servers and clients. After a good deal of sweat and pain, the ssh daemon was installed on all servers and clients that did not already have this capability, including some very old test hosts such as SunOS, Ultrix and HP-UX. Then, all internet services, including ftp, telnet and utilities serviced by the Unix inetd wrappers were disabled. The only things left now are the ssh services, routing, teleconferencing, NTP and mail. A similar thing was done on the CAIRN router. Perhaps the next attack might exploit the camera control program - we are watching.

As commentary, one should reflect on the results of these attacks, which have been an almost continuous threat over the last year or so. They consume a great deal of time to detect, deflect, document and defend against future attacks of the same kind. All staff and graduate students have been told never to unpack attachments of any kind from any source, including family and friends. Our wrappers deflecting known bad guys have upwards of 6000 host domains from which some evil attack or other has been detected. The mail filter for this investigator has several hundred host domains which have originated serious spam, including aol.com, att.net and all of China. Our department routers have every possible defense mechanism enabled.

Still, evidence from tcpdump wiretaps shows continuos barrage from somewhere to one or another nonworking local subnet addresses at rates sometimes over five packets per second. When it gets really bad, our department staff have to track down where those bogons are coming from and educate some system ad mist rat or on the rules of Internet etiquette. Perhaps this might be the legacy of being one of the first kids on the block to get a low IP network number like 128.4.

7. Future Plans

Our plans for the next quarter include further work on the NTP book mentioned in the previous report. Specifically, the chapter on time metrology needs to be updated due to recent changes recommended by the International Astronomical Union. We plan to extend the Autokey protocol to include dynamic trust dependencies involving host groups similar to PGP. We also plan to upgrade to version 3 of the X.509 certificate and include provisions for automatically verifying the certificate trail. Eventual plans are to integrate a certificate discovery function in the manycast scheme, which is necessary for a field-deployable sensor network.

We plan to continue work on the Autoconfigure scheme, in particular a rewrite of the configuration code and name resolution code. This would allow certificates to be retrieved from other than the manycast servers found, in particular from public repositories should this be useful. An interesting problem yet to be resolved is how to develop a security metric which could be used in conjunction with the current synchronization to refine the server selection algorithm.

We plan to continue investigation of timekeeping issues for the Interplanetary Internet, including synchronization paradigms where connectivity between moving objects is not continuous.

We plan to continue the refining of the Autokey protocol, as defined in the current Internet Draft, eventually to Internet Standard status. Eventual plans are to update the NTP Version 3 specification to reflect the current NTP Version 4 architecture, protocol and algorithms.

8. 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 public 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 preceding 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.

8.1 Papers

1. Li, Q., and D.L. Mills. Jitter-based delay boundary prediction of wide-area networks. *IEEE/ACM Trans. Networking 9, 5* (October 2001), 578-590

2. Levine, J., and D. Mills. Using the Network Time Protocol to transmit International Atomic Time (TAI). *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Reston VA, November 2000), 431-439.

3. Mills, D.L., and P.-H. Kamp. The nanokernel. *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting* (Reston VA, November 2000), 423-430.

4. Mills, D.L. A brief history of NTP time: confessions of an Internet timekeeper. Submitted for publication; please do not cite or redistribute.

5. Li, Q., 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).

6. 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.

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

8. Mills, D.L., T.S. Glassey, and M.E. McNeil. Authentication scheme extensions to NTP. Internet Draft draft-mills-ntp-auth-coexist-01.txt, University of Delaware and Glassey-McNeil Technologies, September 1998, 9 pp.

9. 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.

10. 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.

11. 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.

12. 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.

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

8.2 Technical Reports

14. Mills, D.L. Public key cryptography for the Network Time Protocol. Electrical Engineering Report 00-5-1, University of Delaware, May 2000. 23 pp.

15. Mogul, J., D. Mills, J. Brittenson, J. Stone and U. Windl. Pulse-per-second API for Unix-like operating systems, version 1. Request for Comments RFC-2783, Internet Engineering Task Force, March 2000, 31 pp.

16. 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.

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

18. 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.

19. 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.

20. 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.

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

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

8.3 Internet Drafts

23. Mills, D.L. Public-Key Cryptography for the Network Time Protocol. Internet Draft draft-ietf-stime-ntpauth-01.txt, University of Delaware, April 2001, 45 pp.

24. Mogul, J., D. 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, Compaq Western Research Laboratory, August 1999, 30 pp. (obsoleted by RFC-2783)

25. Mills, D.L., T.S. Glassey, and M.E. McNeil. Authentication Scheme Extensions to NTP. Internet Draft draft-mills-ntp-auth-coexist-01.txt, University of Delaware and Glassey-McNeil Technologies, September 1998, 9 pp. (expired)