Scalable, High Speed, Internet Time Synchronization

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

Quarterly Progress Report
1 June 1998 - 30 September 1998

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, graduate students Qoing Li and Robert Redwinski, and undergraduate student Douglas Miller. 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. Specific applications benefiting from this research include multicast topologies, multimedia, real-time conferencing, cryptographic systems, and management of distributed, real-time systems.

Recent quarterly reports have been 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. Clock Discipline Algorithm Improvements

Comments from the referees on a paper previously submitted suggested there may be a flaw in the clock discipline algorithm proposed in the paper. The design of the algorithm was based on the results of simulation using data collected during operation of a previous version of the algorithm. At the time the paper was written, the implementation of the proposed algorithm was not yet complete, so the claimed performance was based on simulation. This fact was clearly stated in the paper. However, when the implementation was completed, the performance predicted by the paper analysis did not agree with the observed behavior of the algorithm. At first, this appeared to be a major setback and gave credence to the referee comments.

After a good deal of analysis and measurement, a bug was discovered in the implementation. When it was fixed, the performance predicted by simulation agreed within nominal expectations. The flaw cited by the referee turned out to be a misleading statement in the description of the new algorithm. The statement was reworded and the paper prepared for publication. It is to appear in IEEE/ACM Trans. Networking at an early date.

3. Network Time Protocol Version 4

Work continued on the implementation and refinement of the NTP Version 4 distribution for Unix, Windows and VMS. The GNU *autoconfigure* system (not to be confused with the NTP autoconfigure technology), which is the basis supporting the over two dozen ports of the distribution to various architectures and operating systems, continued to evolve as each new bug in each new operating system became exposed. We received a donated Alphastation 433au from Digital Equipment and used it to test the net NTP Version 4 algorithms with its 64-bit little-endian architecture. In fact, there were several problems which had not been detected earlier.

A good deal of work has been expended over the Summer reworking the NTP daemon to remove excess baggage and in general polish it up. The infrastructure code, like I/O, timer management, sockets and so forth date from 1988 and has been seriously abused over the years since then by more than two dozen ports to different architectures and operating systems. Finding all the little broken things that spill out when some part or another of the code is reworked to be simpler, smaller and esthetically more pleasing to follow turned out to be a major operation. Unfortunately, the volunteer corps that love to build in new features run for the hills when presented with this challenge.

Progress on the autonomous configuration implementation continues to be slow since the departure of its implementor Ajit Thyagarajan. Mr. Thyagarajan has taken an interim job with Torrentnet, an upstart router manufacturer. He has not finished his dissertation, which would involve in part fixing the bugs remaining in the implementation. So far, a sufficiently attractive carrot has not been found to persuade him to do that, although hope remains. This work may have to be delayed while other more pressing issues are resolved. Investigation of this issue will continue into the next quarter.

4. Infrastructure

We have once again expanded and reconfigured our research network DCnet. There are now several subnets for the Internet research machines and several more for related research activities in the Department. Two hubs serving two routers and three workstations at the Backroom site were installed and interconnected with fiber optic cables. Another hub was installed to service the growing number of workstations and related gear in this investigator’s office. Spectracom donated a new GPS receiver, which was installed in the machine room and connected to the file/web/NFS server for the Internet research machines.

The network configuration for the Backroom test site has become a testing issue. The site is connected to campus via a T1 line that homes to a UDELnet router and also by an ISDN Ethernet bridge as backup. While the DCnet subnets are administered by this investigator and designates, the UDELnet routers are not. The original intent was to use tunnelling so that, in effect, the UDELnet facilities would become transparent. However, this would require the UDELnet router at the Backroom end, as well as the router at the campus end would have to be configured to support the tunnel. So far, the campus Information Technology Office has not been able to do that.

We did manage to put up a multicast tunnel between campus and the Backroom site and it works as expected. We also brought up gated in workstations on campus and at the Backroom site and configured it to accept and provide routes to our CAIRN/DARTnet router. The result is a fairly

astonishing volume of routing updates that stumble over the ISDN link. Since routing updates transit the ISDN but do not transit the T1 line. The result is an interesting combination of unicast and multicast routes. To fix the problems one and all, The Torrentnet folks donated one of their prototype routers. It is to be installed on the campus subnet and used as a tunnel portal.

5. Plans for the Next Quarter

Our plans for the next quarter include continued testing and refinement of the NTP Version 4 protocol model, specification and implementation. Specifically, we plan to resolve the problems with the Unix socket interface mentioned in the previous report, so that the NTP autoconfigure feature is really useful. In addition, we plan to continue the collaboration with Coastek InfoSystems in the design and implementation of the cryptographic certification algorithm. 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.

6. 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 Proponent. The project descriptions are cross-indexed so that the various interrelationships are clearly evident. 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.