REFLECTIONS ON TEN YEARS OF NETWORK TIME SERVICE

Richard Schmidt
Time Service Department
U.S. Naval Observatory
Washington, DC 20392-5420, USA
E-mail: *res@tuttle.usno.navy.mil*

Abstract

The year 2004 marks the 10th anniversary since the start of U.S. Naval Observatory (USNO) time dissemination on the Internet using the Network Time Protocol (NTP). In 1994, our service was inauspicious: two 50MHz/32MB servers on a 56kb WAN link handling one packet every 17 seconds. Today, three servers in Washington, D.C., process over five thousand packets per second from millions of clients across the U.S. and 63 other nations. At the USNO Alternate Master Clock at Schriever AFB, Colorado, two additional servers provide NTP. Seventeen USNO servers with embedded GPS provide U.S. regional coverage from Alaska to Hawaii, from Washington state to Florida, from southern California to Maine. For the past 6 years, USNO has provided time service on the SIPRNET from Washington, D.C., and Schriever AFB. SIPRNET timing will soon expand with remote SAASM GPS servers. Throughout its history USNO, with the assistance of cooperating agencies, has provided free public time dissemination, from time balls to telegraphic time, wireless broadcasts, telephone time, LORAN, GPS, and Internet time. NTP is lightweight, reliable, accurate, and robust. Freely distributed, it has been ported to numerous devices and operating systems. The architecture of NTP is permanently in evolutionary development and in remote system maintenance.

I. EARLY TIME DISSEMINATION AT USNO

The mission of the U. S. Naval Observatory is to:

1. Determine the positions and motions of celestial bodies, motions of the Earth, and precise time.

2. Provide astronomical and timing data required by the Navy and other components of the Department of Defense for navigation, precise positioning, and command, control, and communications.

3. Make these data available to other government agencies and to the general public.

1

4.   Conduct relevant research, and perform such other functions as may be directed by higher authority **[1].**

The Observatory's role as public provider of time began in December, 1844, in an order to Superintendent Lt. Matthew Fountaine Maury signed by Navy Secretary John Y. Mason:

> You will be pleased to devise some signal by which the mean time may be made known every day to the inhabitants of the city of Washington. When you are prepared to put your signal into operation, you will give notice of the kind you have adopted in the city newspapers, and at the same time inform the Department **[2]**.
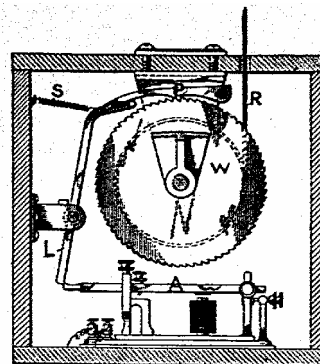
HOW THE BALL IS DROPPED
A rope, *R*, runs up through a flag-staff, out at top, down to ball at base; windlass, *W*, lifts the ball to top of pole; Observatory clock sends a current through magnet, *M*, by which armature, *A*, is instantly lowered; lever, *L*, held back by resting against end of armature, is now freed; spring, *S*, jerks it forward, knocks pawl, *P*, out of ratchet and ball comes down.

TELEGRAPHIC TIME

The signal that Maury chose was a time ball that was operated for 40 years from the observatory site in Foggy Bottom. In 1885, the time ball was moved to the roof of the State, War, and Navy Building, and that quaint service continued until 1936 **[3]**. The time ball was a line-of-sight signal of use to ships in the Potomac, but limited in scope. But in 1865, the Observatory director William Harkness connected his clocks to the city fire alarm telegraph system. This brought Observatory time to the Western Union Telegraph Company at the State Department, and within a few years this useful service was extended from coast to coast **[4]**. In 1877, the Naval Observatory and the Western Union

Company established a national time service that included remotely synchronized time balls and a vast number of telegraphically controlled clocks. So popular was the Western Union "self-winding clock" that by the

1920's over 120,000 subscribers paid monthly fees for Naval Observatory time **[5]**. As late as 1964, all Western Union stations maintained their Observatory time feeds, but a year later the company discontinued relaying Observatory time.

WIRELESS TIME
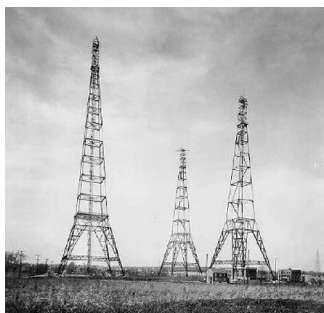
In 1904, the Naval Observatory first broadcast time by wireless from a naval station in Navesink, New Jersey **[6]**. By 1912, the Navy had erected Arlington Towers, station NAA, three 550-ft towers on the Potomac River in Virginia. These were the tallest free-standing towers in the country, and their broadcasts could be received at the Eiffel Tower in Paris. The Observatory connected its time signal via telegraph to the tower transmitter, and began a time broadcast that could be picked up with a $100 radio **[7]**. By 1915,

USNO time service operated daily from eight Navy transmitters, as shown in the following table [8]:

Station	Frequency	Controlled from	Schedule
Arlington, VA	120 kHz	USNO Washington	11:55 - noon, 21:55 – 22:00 EST
Key West, FL	300 kHz	" "	11:55 – noon EST
New Orleans, LA	300 kHz	" "	11:55 – noon EST
Mare Island, CA	120 kHz	USNO Mare Is.	11:55 – noon, 21:55 – 22:00 PST
Eureka, CA	214 kHz	" "	11:55 – noon PST
Point Arguella, CA	400 kHz	" "	11:55 – noon PST
San Diego, CA	150 kHz	" "	11:55 – noon PST
North Head, WA	150 kHz	" "	11:55 – noon PST

Time broadcasts were expanded to 24 times daily by 1938. A magazine advertisement of the period touted the "General Electric Clock, regulated by comparison with Naval Observatory Radio Time Signals. Time from the stars … Arlington time, reported by radio … right because the impulses of alternating current from your power station are right … kept constant by comparison with radio time signals from the U.S. Naval Observatory." Ironically, during World War II the Navy ordered the Observatory to cut back transmissions to just two per day. By the end of the war, the Observatory's radio time service was overtaken by the new time broadcasts of the National Bureau of Standards.

TELEPHONE TIME

Since the 1970's, the Naval Observatory has provided a direct telephone time service using a Weatherchron system. The original Weatherchron featured the voice of announcer Fred Covington and was recorded on a magnetic drum. The current service is digitally encoded on EPROMs [9]. A digital time service for modems was also provided which supports remote loop back at 1200 baud, 8N1. Historical traffic logs for these services are shown above.
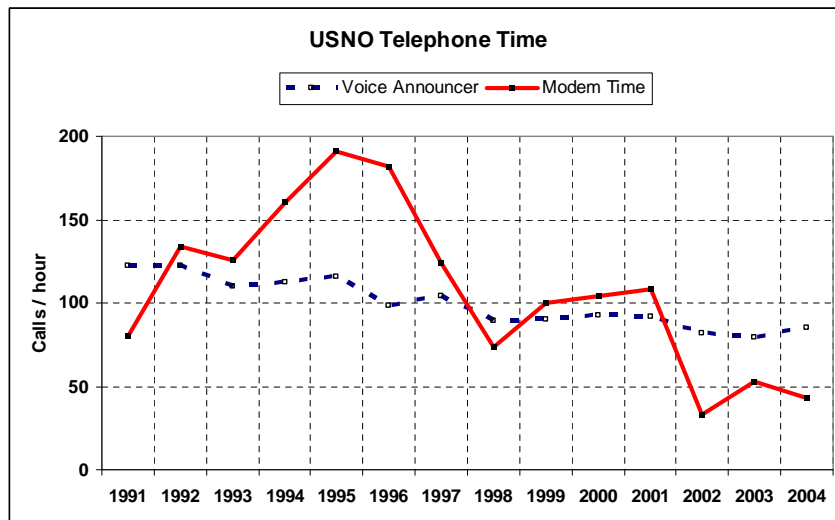
WEB TIME SERVICE

The Observatory's first public Web server was the Time Service Department server *http://tycho.usno.navy.mil*, which went live on 22 November 1994 [10]. Within a few months, traffic was doubling every two months. This Web server offered pages of information on the USNO Master Clocks, GPS operations, LORAN, and related Time Service products, but the most-accessed links have always been the time of day displays (over a thousand hits per day in July, 1995.) The original static time display *http://tycho.usno.navy.mil/cgi-bin/timer.pl* was updated on demand until the hit rate exceeded one hit per second. Then a cached page was created once per second on the second.

US Naval Observatory Master Clock Time

Nov. 30, 22:09:45 UTC
Nov. 30, 05:09:45 PM EST
Nov. 30, 04:09:45 PM CST
Nov. 30, 03:09:45 PM MST
Nov. 30, 02:09:45 PM PST
Nov. 30, 01:09:45 PM AKST
Nov. 30, 12:09:45 PM HAST

3

Currently this time display is hit 10 times per second.

In July, 1996, "real-time" animated clock displays were added.   The most utilized was the animated GIF clock, which produced a live stream of GIF animation displaying digital hours, minutes, and seconds. This program used server push technology with non-parsed headers ("Content-type: multipart/x-mixed-replace"). Even though it is not supported on the later Microsoft Internet Explorer browsers, it has many advantages. It can appear as a clock on a remote Web page, it has a low bit transmission rate, and it allows precise sequencing of times of transmission. For example, the server checks system clock time and delays the start of transmission until the precise on-time second mark. The main drawback of this clock is that it requires spawning a new process for each connection. The executing program has a MIME extension of "gif," displaying as an animated image. In July, 1998, the animated GIF program was called 600 thousand times per day (7 hits/sec) **[11]**.

II. NETWORK TIME PROTOCOL

The most accurate method for computer time synchronization is the Network Time Protocol (NTP), providing millisecond accuracy. The mechanics of this protocol were proposed by Dr. David Mills, in DARPA RFC-958 in 1985 **[12]**. The current standard is RFC-1305 **[13]**. NTP has been ported to virtually every computer platform, and is currently running on perhaps 100 million machines **[14]**. In the full NTP implementation, the ntpd daemon synchronizes a client in frequency and in phase to one or more peers or higher stratum servers. It is designed to operate over the packet-switched Internet, and it mitigates loss of connectivity while maintaining system time accuracy at the millisecond range.

Today, publicly advertised stratum 1 servers are in operation in 26 nations around the globe. But 10 years ago far fewer stratum 1's were available, and most of those relied on millisecond radio clocks.   The Naval Observatory's first NTP servers were the HP 9000/a747i machines *tick.usno.navy.mil* and *tock.usno.navy.mil*.   These basic industrial workstations (32 MB RAM, 50 MHz CPU's, 1 GB disks) synchronized to the USNO Master Clocks using IRIG-B- disciplined VME bus clocks (Datum-Bancomm 635vme, TrueTime VME-SG).   Time could be accessed without latency over the VME bus using a modification of an HP-UX 9.0 operating system VME driver originally developed for charge-coupled device photometry.
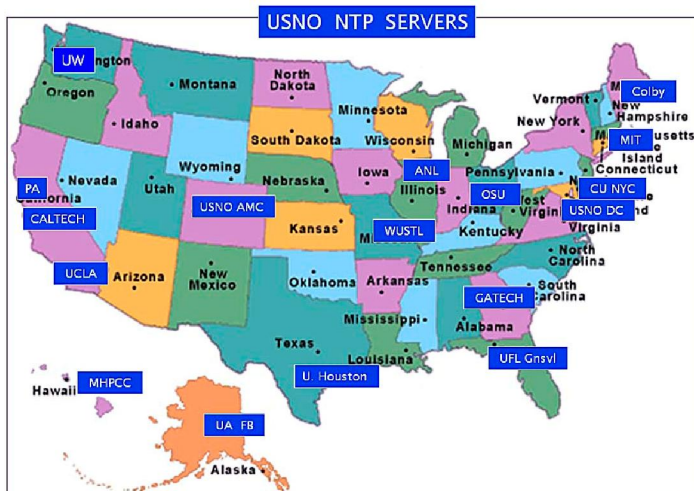
 In December, 1995, a remote USNO NTP server was set up at Washington University in St. Louis, Missouri, using a TrueTime GPS-VME card with a Magnavox GPS receiver. By 1997, 14 USNO NTP servers were in operation on the Internet in 11 US cities, and SIPRnet servers were established at USNO Washington and at the USNO Alternate Master Clock Facility at Schriever AFB.   In 2004, there are 23 USNO NTP servers in 18 U.S. cities.

Currently, these are Hewlett-Packard A500 and A180 business servers. Each has a PCI GPS receiver (Brandywine Communications Syncclock-32 PCI with Motorola UT+ GPS). The HP-UX servers have RAID

mirrored disks and operate for thousands of days without maintenance.

The USNO NTP stratum 1 servers were advertised as "open access," and in 1994 the incoming NTP traffic was about one request per 17 seconds. As each NTP packet is only 90 bytes in size (48-byte NTP packet + 8 byte UDP header + 20 byte IPv4 header + 14 byte Ethernet frame header), it seemed unlikely that NTP traffic would ever significantly load even a T-1 network. But the phenomenal growth of the Internet brought exponential growth of the connected PC population. NTP traffic has more than doubled each year.
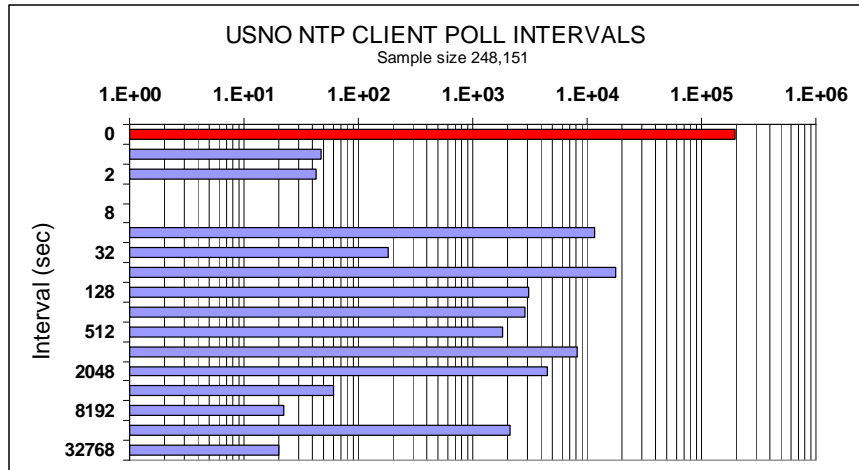
In August, 2004, with NTP requests coming in at 4,000-7,000 per second (implying millions of individual clients) in Washington, D.C., alone, a USNO NTP client sample was undertaken.   Over the course of a week, over 2,000,000 unique client IP's were logged (a rate of 15,000 new clients per hour).   Reflecting the general Internet population, the distribution of client domains (that could be resolved with reverse DNS lookups) was as shown in the following table:
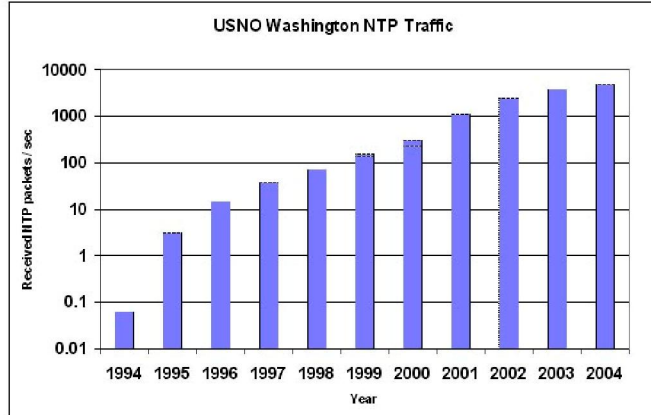
| Domain	.net	.com	.edu	DSL	.mil	.gov
	41%	41%	10%	4%	2%	2%

In addition, NTP clients from 188 foreign nations were identified. The majority (15%) were from Canada, followed by Germany and the Netherlands (9% each), Japan and France (6% each), and Belgium and Australia (4% each).

In November, 2004, a survey was made of 248,000 USNO NTP clients, looking at incoming NTP polling interval flags.   It was found that 79% of current clients have   poll-interval=0, indi- cating Simple Network Time Protocol (SNTP) **[15]** or SNTP-like   client   software. This indicates that the majority of USNO clients are not acting as stratum 2 relays to local workgroups, although some are known to relay to tens of thousands of their own NTP clients.   In the graph on the following page, the SNTP-like traffic at USNO Washington, D.C. shows up as hourly peaks in one day's traffic log. (Here, packets received and transmitted by the server farm are totaled. Transmitted packets are fewer due to discards of "rate-exceeded" clients.) The full NTP daemon does not poll on the hour, and is randomized across the day

by the volume of clients. SNTP, however, is typically launched at regular intervals by a UNIX cron-like scheduler. Humans tend to schedule events on the zero minute.

**NTP Traffic December 22, 2004**
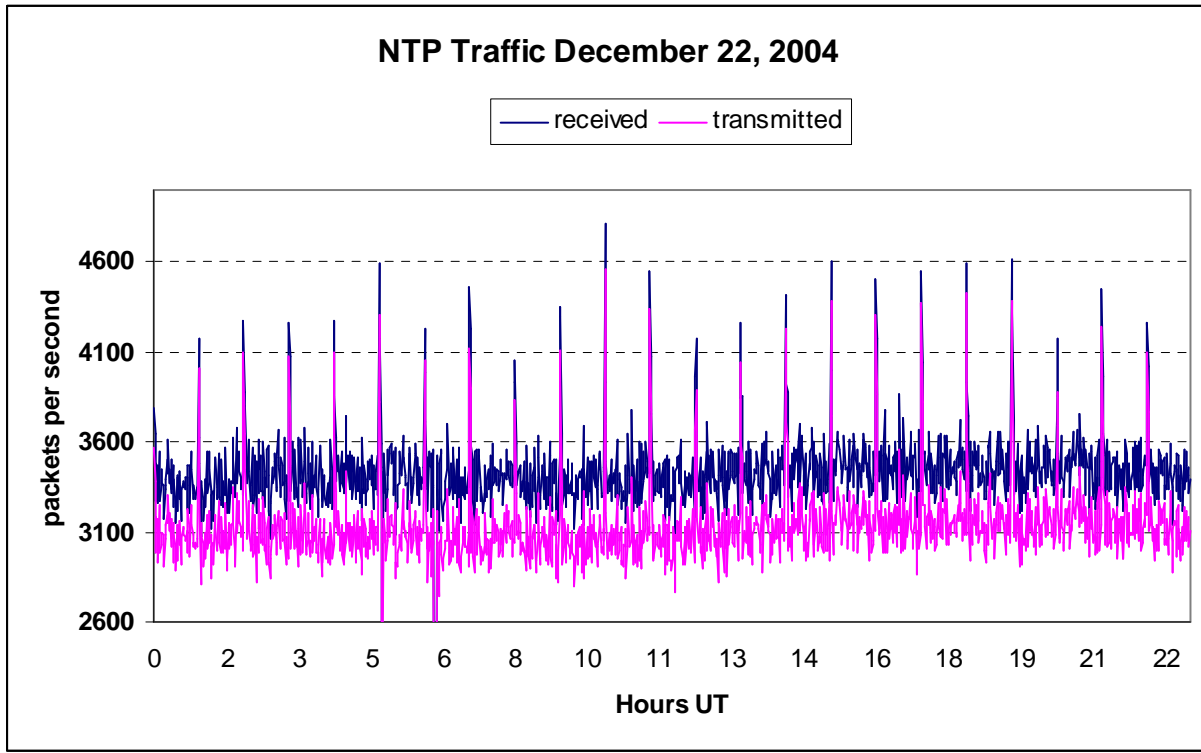
NTP SERVERS AT USNO WASHINGTON

In order to provide automatic failover in the event of loss of individual NTP servers, the well-known NTP hosts *tick, tock,* and *ntp2.usno.navy.mil* are now load-balanced behind a CAI Networks Webmux LE IP load balancer. This arrangement enables incoming traffic to be routed to any or all of the NTP severs. Routing can be based on weighted or unweighted round-robin or based on the number of current connections. Persistence can be enabled at the expense of load balancer memory. The NTP server farm can be upgraded without creating additional server host names.

The NTP servers are each synchronized to both Master Clock 1 and 2 via IRIG-B, as shown in the figure at right. Each server has two IRIG-B- synchronized bus clocks, providing redundancy in the event of loss of a single input. A third input, GPS for example, would be desirable in order to filter out a discordant IRIG input signal.
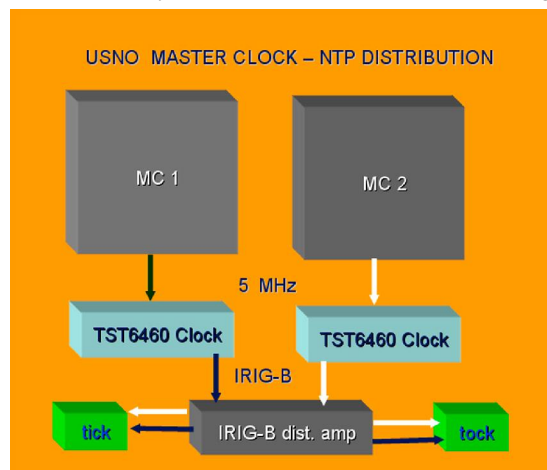
III. CURRENT USNO NTP DEVELOPMENT

In October, 2004, USNO integrated a new Hewlett-Packard rx1600 (Itanium processor) server with a Symmetricom bc635pci-U bus clock under HP-UX 11.23. This high-performance server can process 24,000 NTP requests per CPU-second with an NTP process load of 4.5% of the total CPU. The role of the PCI bus clock is illustrated in the figure below. The "ntpd" daemon forms a transfer vector with functions to start, poll, and shutdown the PCI device. These operate through kernel driver routines that attach the PCI device to the process, open it as a UNIX device file, and perform ioctl's to obtain memory through the HP-UX WSIO PCI I/O services. These driver routines were developed by David Johns (USNO) based on his earlier PCI driver for HP-UX 11.11. The Symmetricom bus clock's dual-port RAM appears as virtual addresses in user space. One of these is a pair of 32-bit registers with major time in the UNIX "struct tm" format with microsecond to 100-nanosecond time in the following word. This register can be accessed in 3.24
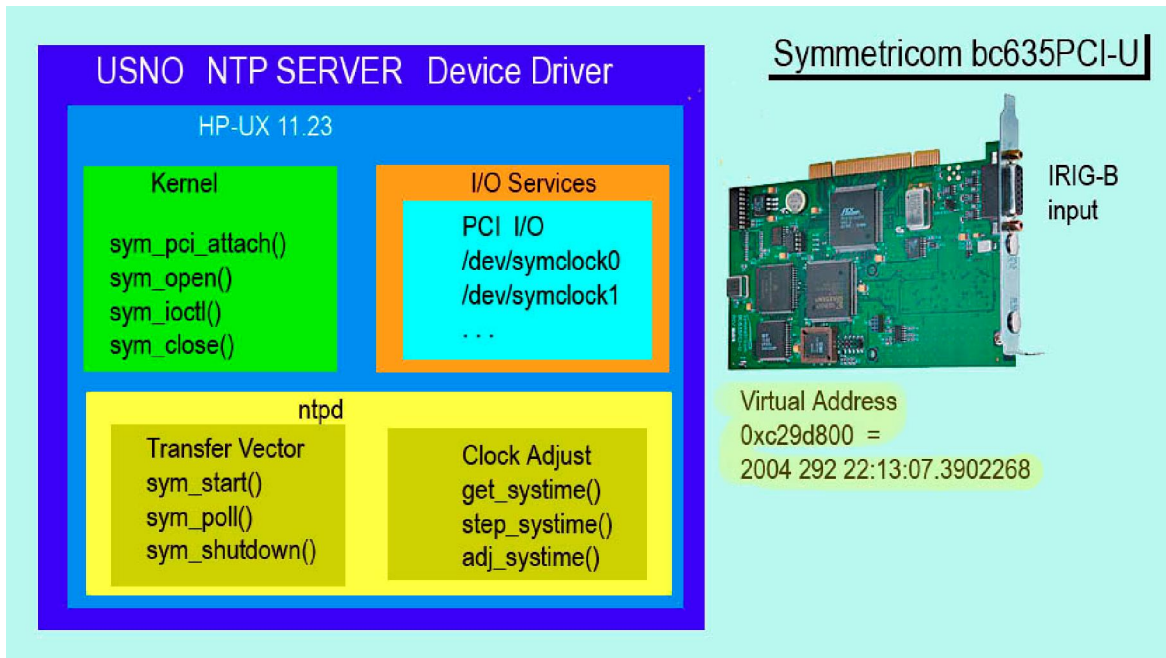

microseconds (μs) of real time. It feeds into the NTP clock adjust algorithm, which makes frequency changes to the system clock via the UNIX adjtime system call. The figure right shows from the top an FEI-Zyfer Gsync PPS GPS receiver outputting IRIG-B to a Hewlett-Packard rx1600 server, and a Hewlett-Packard


A500 server with a Brandywine PCI GPS bus clock.

The NTP daemon can be instructed to monitor the phase error of the UNIX system clock with respect to the PCI bus clock reference, the so-called "loopstats" values. The following figures show the time deviation and


overlapping Allan deviation of the current (rx1600 and A500) NTP servers and of the previous HP9000/748i servers [16].


.
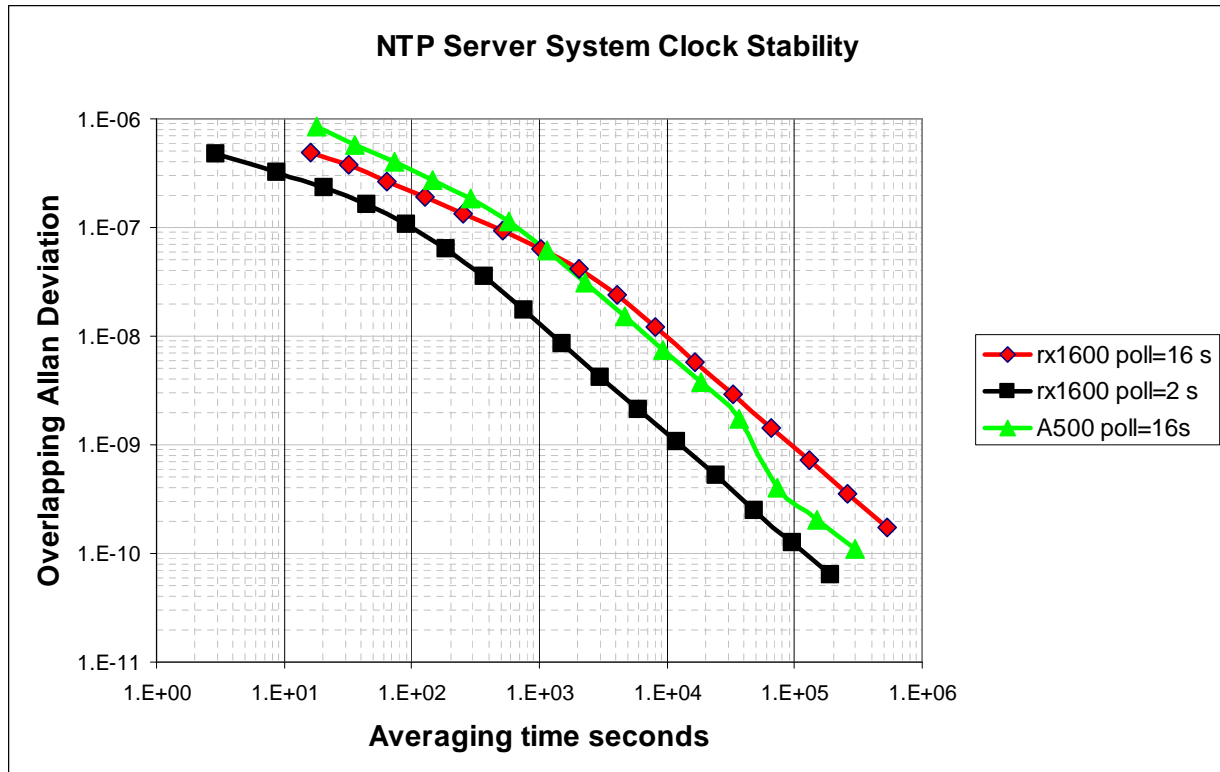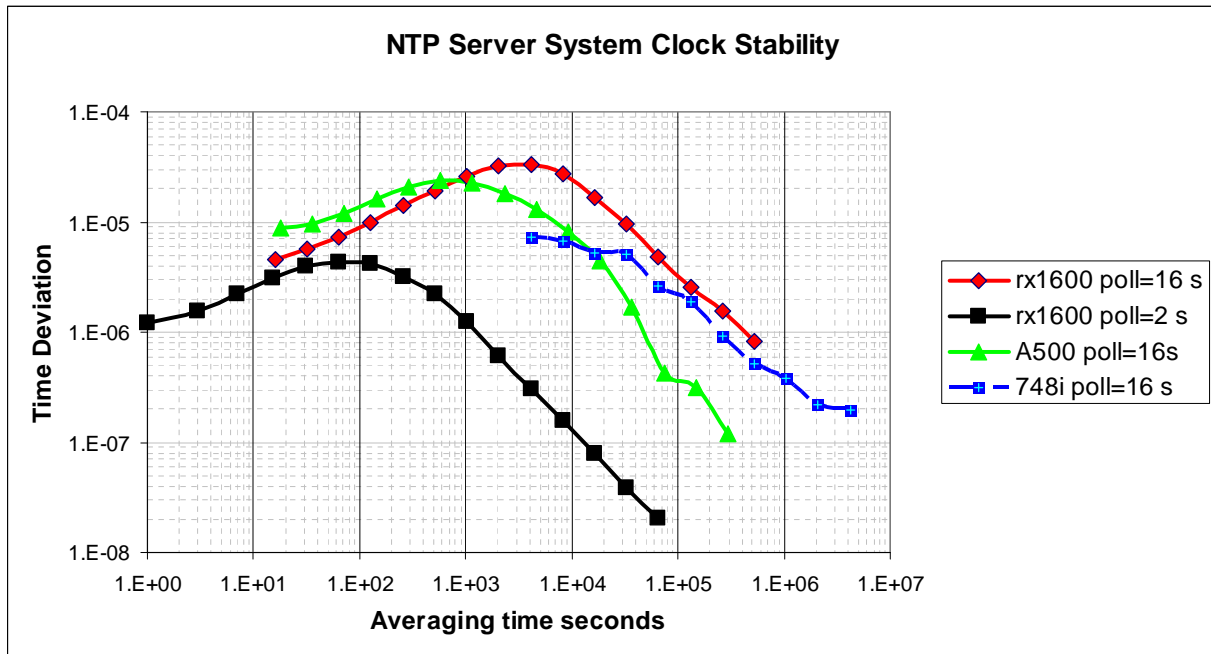
Note the improved stability resulting from decreasing the NTP refclock poll interval from the nominal 16 seconds to 2 seconds, where the standard deviation of the NTP loopstats (system clock offset from PCI bus clock) decreased from 54 μs to 7 μs. A loopstats histogram is shown below.

rx1600 Server Clock Offset, poll=2s
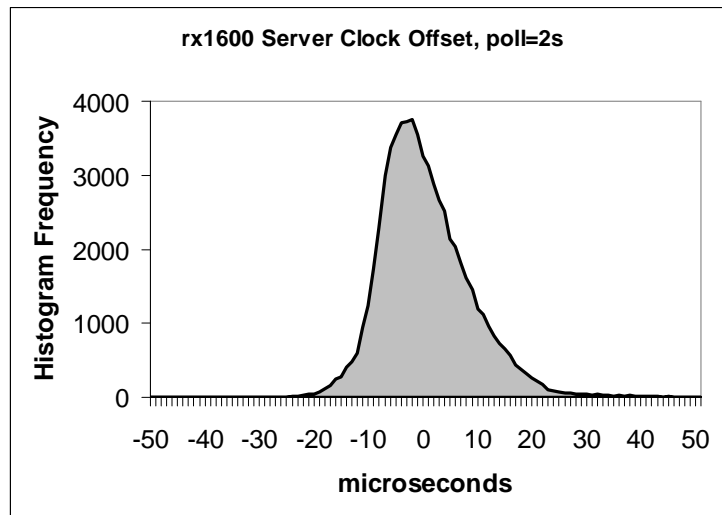
The figure below shows typical NTP clock frequency error in parts-per-million for the rx1600 Itanium server. With a 16-second bus clock-polling interval, frequency excursions are large and hint at environmental temperature effects.

rx1600 System Clock Frequency Error poll=16s
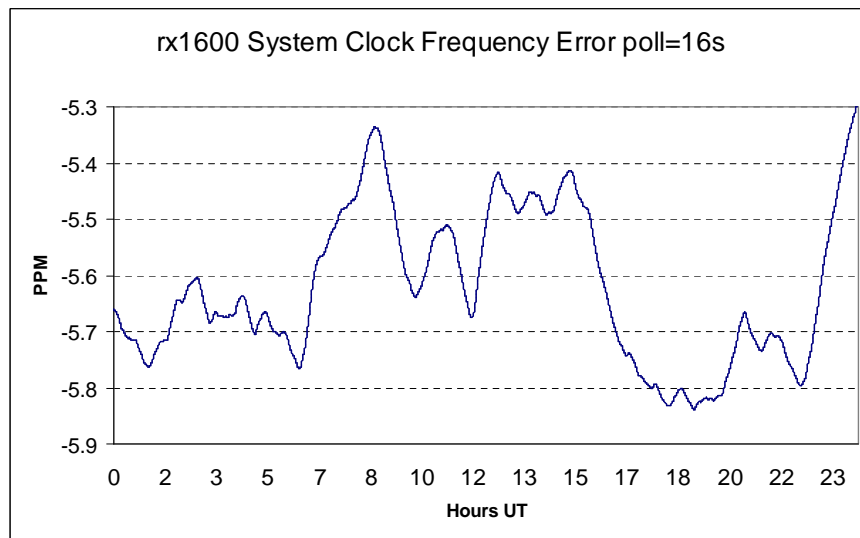
The next figure shows the corresponding frequency error with a polling interval of 2 seconds. The multi-hour effects are no longer dominant. Frequent polling can result in large wander in the event of loss of the reference clock, as the system clock error estimate is based on a shorter integration. Note that NTP prevents network polling at intervals of less than 16 seconds (NTP MIN_POLL = 2^4).

rx1600 System Clock Frequency Error poll=2 s
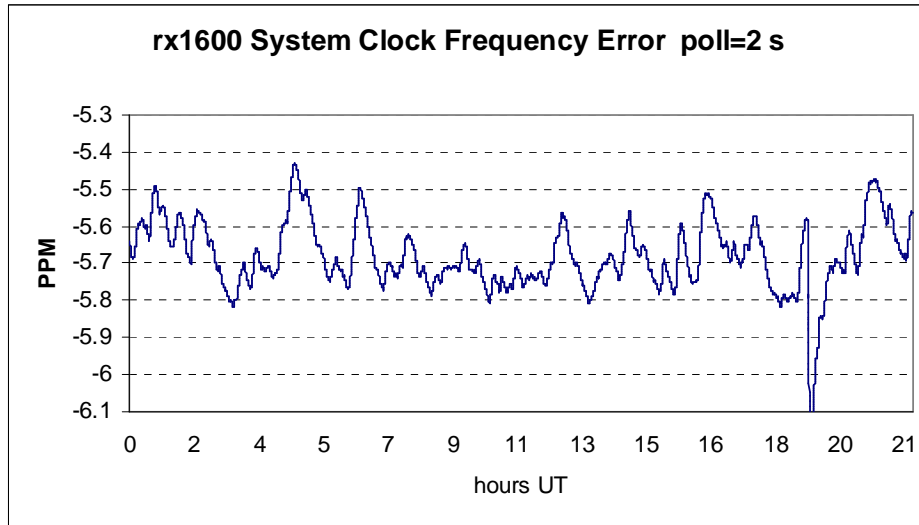
IV. GPS RECEIVER MONITORING

The USNO design for remote NTP servers provides basic GPS receiver status monitoring. In the case of the Zyfer Gsync with PPS GPS, a TCP/IP application on the NTP server (developed in ANSI C) opens a telnet session to the receiver network port and queries status, as shown below.

```
Zyfer: 015   Model: 391-2XX:OCXO     S/N: 3910230 IP: 10.1.4.51
Version: V1.22.00,Sep 29 2003,14:45:31, 385-3021,2.11 3/12/2003
Year: 2004 DOY: 358        18:29:54 UT             TFOM: 4
Latitude: 38 55.2202 N   Longitude: 77 03.9793 W        Alt: 57.97
Time Mode: UTC Time       Operating Mode: Time Locked
Anti-spoof mode: C/A mode         PPS Mode: PPS off      Receiver mode: Navigate
CV Keyload Status:        No key loaded
 Time valid  Almanac valid  Position valid


 PRN      SN       Status
 27       34       C/A tracking
 08       48       C/A tracking
 07       47       C/A tracking
 31       47       C/A tracking
 19       28       C/A tracking
 28       45       C/A tracking
 20       29       C/A tracking
 11       42       C/A tracking
```

V. NTP SECURITY

Our choice 10 years ago of the HP-UX operation system for NTP servers has become a major factor in insuring their network security. The NTP servers generally operate outside company firewalls and so must be locked down. It is simple to disable all unnecessary services in the UNIX environment. No browsers, mail programs, graphics systems, or other network holes are integral to the core UNIX operating system. One measure of operating system vulnerability is found in the Open Vulnerability Assessment Language (OVAL) vulnerability definitions database **[17]**. The statistics are based on number of different known vulnerabilities,

not on number of deployed systems. The current distribution of known vulnerabilities by operating system breaks down as follows:

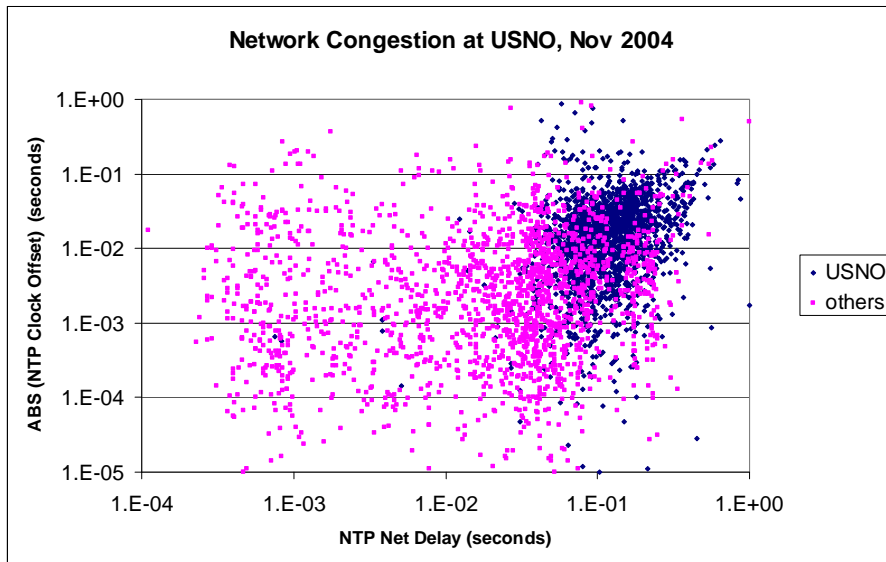| O/S	Microsoft Windows	Red Hat Linux	Sun Solaris	Hewlett-Packard HP-UX
	73%	20%	6%	1%

For information on choosing a secure operating system, see the Carnegie Mellon CERT Coordination Center site: *http://www.cert.org/tech_tips/choose_operating_sys.html*. One feature of network security of interest to clients is the NTP authentication scheme, which uses Public Key Infrastructure (PKI). Unlike traditional client-server roles, NTP PKI is designed to authenticate the server to the client, as protection against so-called "man-in-the-middle" security threats. The current NTP distribution supports a number of identity schemes, including private certificate, trusted certificate (X509 v3), and a strong challenge-response exchange. NTP authentication is enabled when the NTP daemon is compiled with support for Secure Sockets Layer. The "ntp-keygen" utility is used to create PKI certificates.

VI. NETWORK CONGESTION

In the spring of 2004, the growth of NTP traffic reached 8 Mbps and finally overloaded the USNO's network connection, which was a half-duplex 10 Mbps line shared by other agencies. By the summer of 2004, several attempts were made to shape USNO's network traffic attempts. Most involve dumping some incoming packets, which can stimulate the NTP clients to poll more frequently. This is because the NTP client-server poll interval, which ranges from 2^4 to 2^{10} seconds, is an adaptive parameter of NTP's type-II PLL, computed in part from the Allan variance of the clock update samples obtained from a server. When the sample noise is low, the client increases its poll interval to the maximum 17 minutes, in order to achieve a precision of about a millisecond per day. But when the server timestamps are noisy due to latencies or lost responses (for which NTP's MAX_DISPERSION is assumed), the client's loop bandwidth increases and its poll interval decreases. The effects of network overload on USNO's NTP quality of service are illustrated in a survey of 2,100 USNO clients that reported NTP estimates of network delay and clock offset with respect to USNO Washington. Plotting client NTP network delay vs. clock offset on an unsaturated network shows the familiar "NTP wedge plot," as seen in the first figure on the following page, taken from a 1997 study of USNO NTP servers [16].

NTP CLOCK OFFSET vs. NET DELAY

(Atlanta, Boston, Falcon, Houston, St.Louis, NYC } - Washington
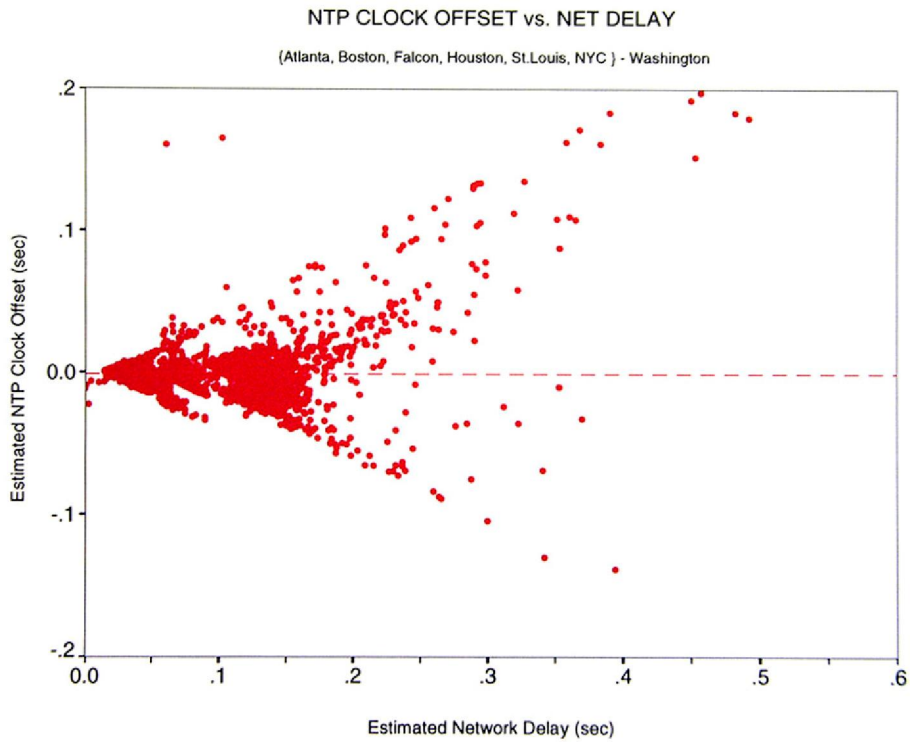
But the results on our congested network look quite different, as seen in the next figure.

Network Congestion at USNO, Nov 2004

Two populations are seen: the clients' view of USNO NTP servers in dark, and their view of other non-USNO servers seen in pink. The range of NTP dispersion as a measure of network stability is shown in the

following graph, where the USNO points are clustered in the range (128ms – 1 s) at which the client rejects the server as unstable.
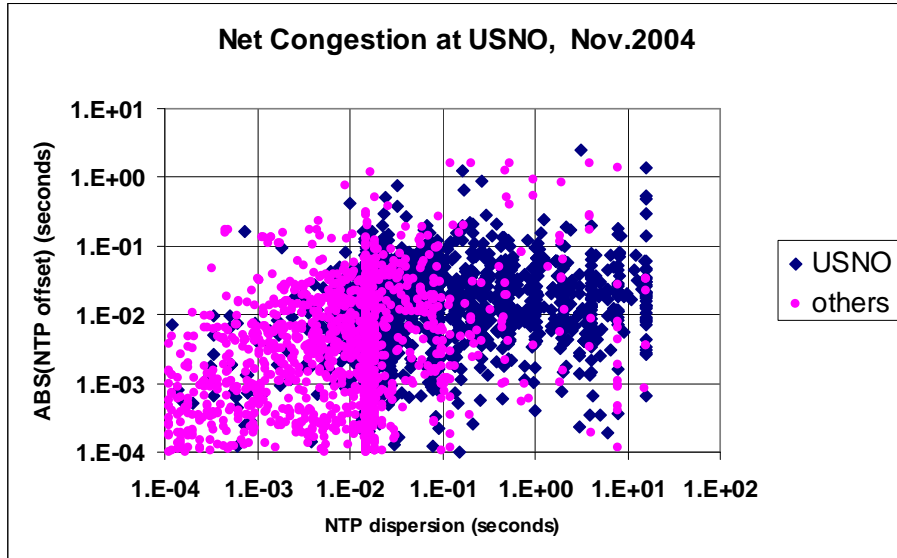
The current USNO network congestion is expected to be alleviated, at least for a few years, early in 2005. But the exponential growth of NTP client traffic appears to continue unabated, and a number of steps are being taken to survive the onslaught. The USNO stratum 1 servers have historically been listed as "open access" on the various public lists **[18]**. This was recently changed to "open access to .gov, .mil domains, and other stratum 2's by arrangement." By concentrating on providing reliable service to organizational stratum 2's which relay to multiple clients, rather than providing time to millions of individuals, we assure quality at the top of the NTP pyramid. Now a more active program of monitoring NTP traffic is targeting the small percentage of users who abuse the servers with unreasonable traffic. We use traffic-sampling software to identify clients that hit the servers many times per second. These are almost always misconfigured client networks.

VII. CONCLUSION AND ACKNOWLEDGMENTS

After 10 years of continuous operation, the USNO NTP service has proven its ability to meet its original goal, that of providing accurate and highly reliable network time service to a large client base. The use of standard protocols and interfaces, and of commercial off-the-shelf components, has enabled USNO to leverage past driver development over three generations of server upgrades. The successful integration of latest-technology servers, bus clocks, and PPS GPS enhances USNO's ability to meet future requirements in network time service.

Thanks to these staff members of the Time Service Department, US Naval Observatory: David Johns, for current PCI driver development in the HP-UX environment; Dr. Demetrios N. Matsakis, for support in all aspects of the NTP program; Dr. Edward Powers, for technical support and liaison; Francine Vannicola, for continued support of the SIPRNET NTP service; Paul Wheeler, for time distribution services; and Dr. Gernot Winker (retired) for insisting in 1994 that there must be a *tock* to accompany *tick*. Thanks also to Dr. Judah Levine, Time and Frequency Division, National Institute of Standards and Technology, and Dr. David Mills, University of Delaware, for many helpful discussions.

VIII. DISCLAIMER

Although some manufacturers are identified for the purpose of scientific clarity, USNO does not endorse any commercial product nor does USNO permit any use of this document for marketing or advertising. We further caution the reader that the equipment quality described here may not be characteristic of similar equipment maintained at other laboratories, nor of equipment currently marketed by any commercial vendor.

REFERENCES

[1] Department of the Navy OPNAV Instruction 5450.90C, 23 May 1984. This instruction directs the Naval Observatory to "manage the distribution of precise time by means of timed navigation and communications transmissions…"

[2] Communication from J. Y. Mason to Lt. M. F. Maury, 10 December 1844. Copy in Observatory archives.

[3] S. J. Dick, 2003, **Sky and Ocean Joined, The U.S. Naval Observatory 1830-2000** (Cambridge University Press), pp. 186 and 468.

[4] *ibid*, p. 182.

[5] A. W. Grumbine, 1985, "The Era of Morse Telegraphy," **Dots and Dashes** (Morse Telegraph Club). See *http://www.faradic.net/~gsraven/telegraph_tales/grumbine/grumbine_2.html*

[6] Dick*, op. cit*, p. 468.

[7] A. H. Orme, 1913, *"Regulating 10,000 Clocks by Wireless,"* **Technical World Magazine,** October, 1913, pp. 232-233.

[8] *"Transmission of Time Signals by Naval Radio Stations,"* **Radio Stations of the United States** (Bureau of Navigation, Dept. of Commerce), 1 July 1915 edition, pp. 169-172.

[9] Paul Wheeler, U.S. Naval Observatory, personal communication. The voice is that of the late Fred Covington.

[10] Monthly Report, November, 1994, Time Service Department, U.S. Naval Observatory, Washington, D.C.

[11] *Ibid*, July, 1998.

[12] D. L. Mills, 1985, *"Network time protocol,"* DARPA Network Working Group Report RFC-958, September 1985.

[13] D. L. Mills, 1996, Internet Engineering Task Force Request for Comments 2030, University of Delaware, October 1996.

[14]   D. L. Mills, 1997, "The *network computer as precision timekeeper*," in Proceedings of the 28th Annual Precise Time and Tim Interval (PTTI) Applications and Planning Meeting, 3-5 December 1996, Reston, Virginia, USA (U.S. Naval Observatory, Washington, D.C.), pp. 97-107.

[15] D. L. Mills, 1996, "*Simple network time protocol (SNTP), version 4 for IPV4, IPV6, and OSI,*" Network Working Group Report RFC-2030, University of Delaware, October 1996.

[16] R. E. Schmidt, 1998, *"SIPRNET network time service,"* in Proceedings of the 29th Annual Precise Time and Time Interval (PTTI) Systems and Applications Meeting, 2-4 December 1997, Long Beach, California, USA (U.S. Naval Observatory, Washington, D.C.), pp. 373-380.

[17] United States Computer Emergency Readiness Team, *"Open Vulnerability Analysis Language,"* *http://www.us-cert.gov/oval/about.html*

[18] See *http://ntp.isc.org/bin/view/Servers/WebHome* and *http://tycho.usno.navy.mil/ntp.html*