Internet Timekeeping Around the Globe

David L. Mills, A. Thyagarajan, B. C. Huffman
University of Delaware
http://www.eecis.udel.edu/~mills
mills@udel.edu

alautun, Maya glyph

Introduction

- Network Time Protocol (NTP) synchronizes clocks of hosts and routers in the Internet
- Provides submillisecond accuracy on LANs, low tens of milliseconds on WANs
- Primary (stratum 1) servers synchronize to UTC via radio, satellite and modem; secondary (stratum 2, ...) servers and clients synchronize via hierarchical subnet
- Reliability assured by redundant servers and diverse network paths
- Engineered algorithms used to reduce jitter, mitigate multiple sources and avoid improperly operating servers
- Unix NTP daemon ported to almost every workstation and server platform available today - from PCs to Crays
- Well over 100,000 NTP peers deployed in the Internet and its tributaries all over the world

NTP configurations

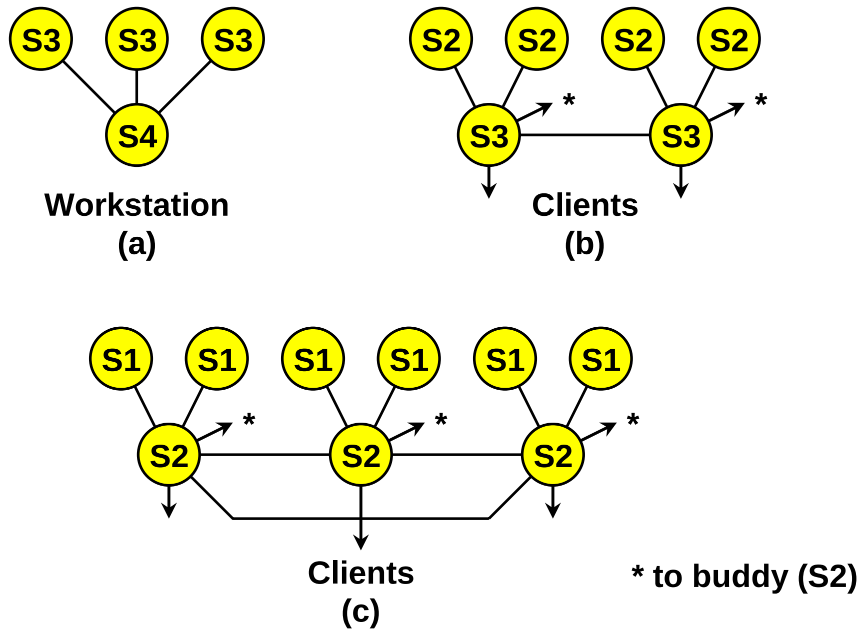

- (a) Workstations use multicast mode with multiple department servers
- (b) Department servers use client/server modes with multiple campus servers and symmetric modes with each other
- (c) Campus servers use client/server modes with up to six different external primary servers and symmetric modes with each other and external secondary (buddy) servers

NTP architecture

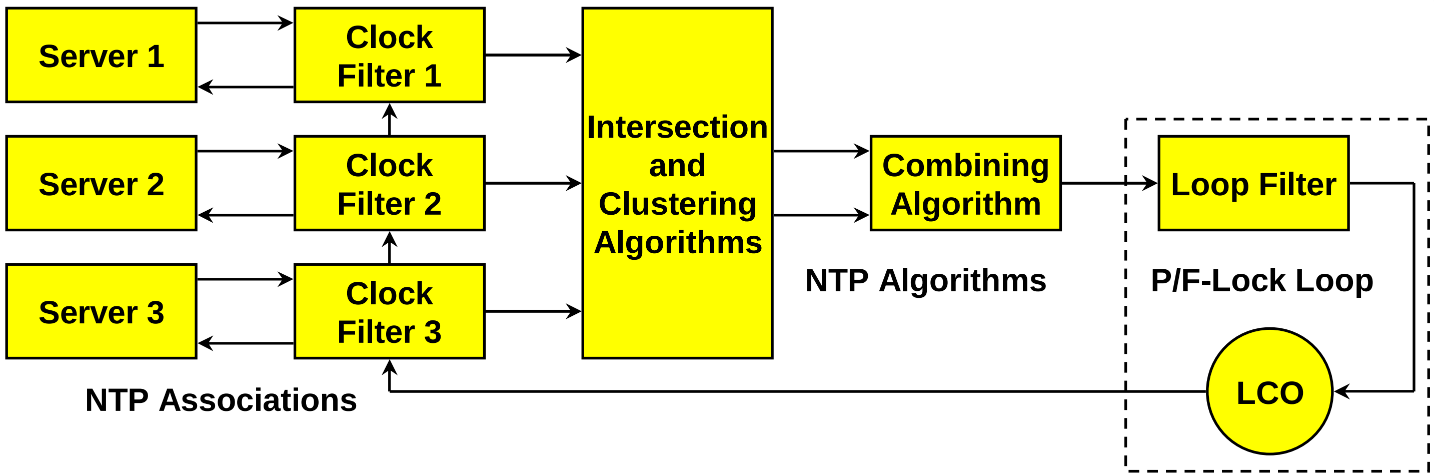

- Multiple synchronization peers provide redundancy and diversity
- Clock filters select best from a window of eight clock offset samples
- Intersection and clustering algorithms pick best subset of servers and discard outlyers
- Combining algorithm computes weighted average of offsets for best accuracy
- Loop filter and local clock oscillator (LCO) implement hybrid phase/frequency-lock feedback loop to minimize jitter and wander

Server population by stratum

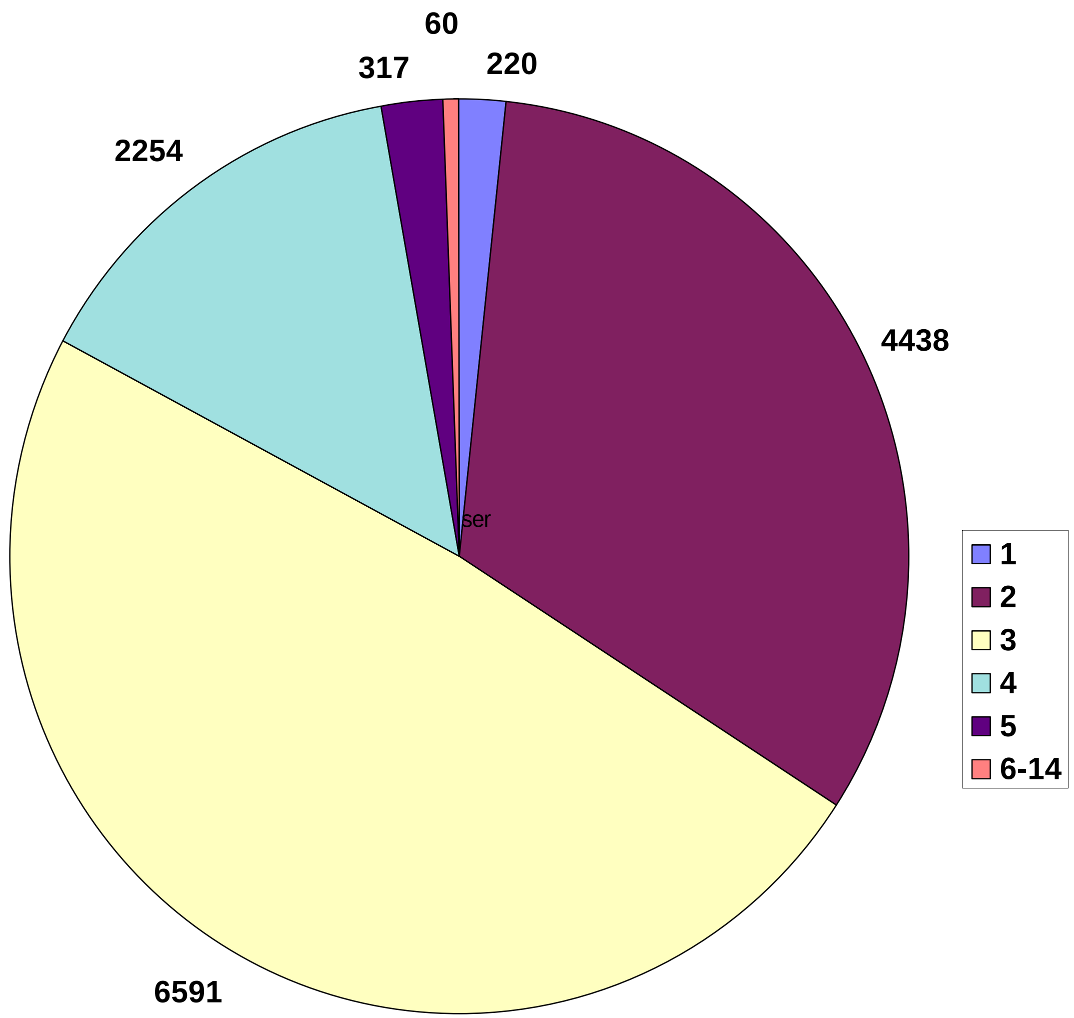

Association population by stratum

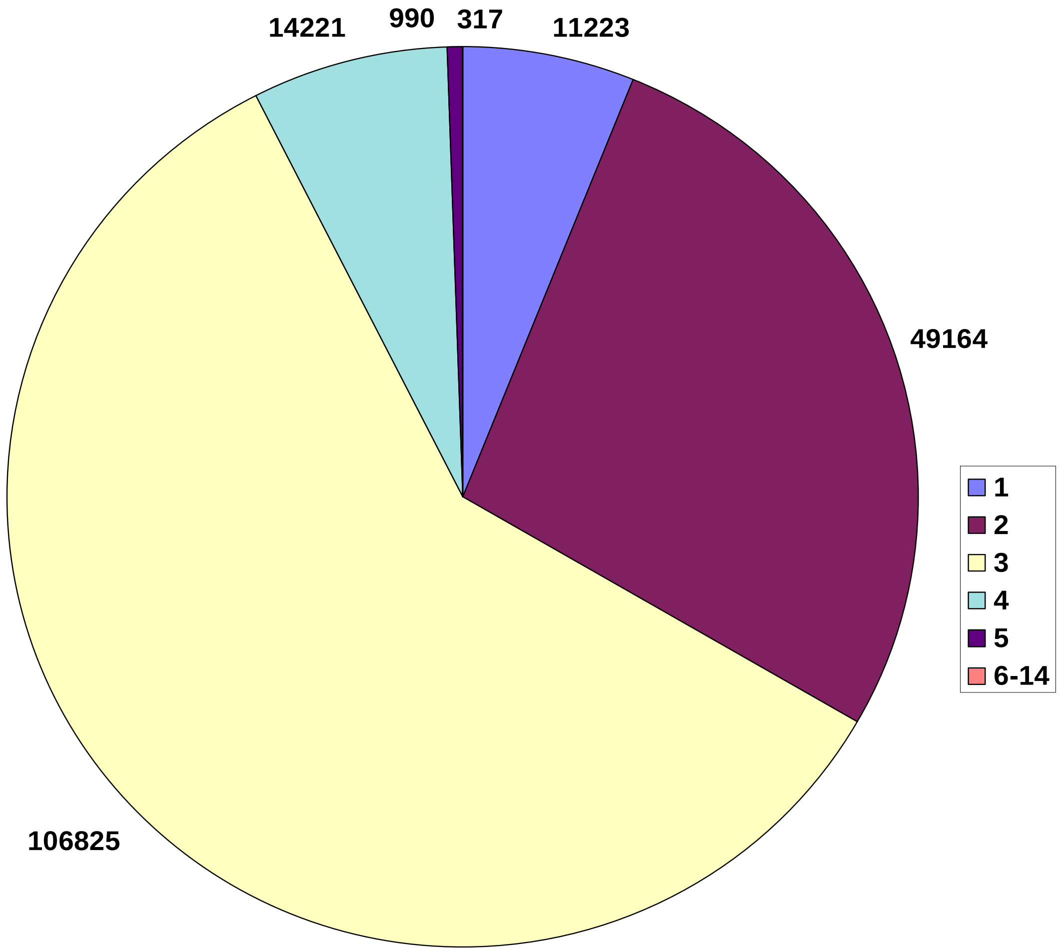

Associations per server population by stratum

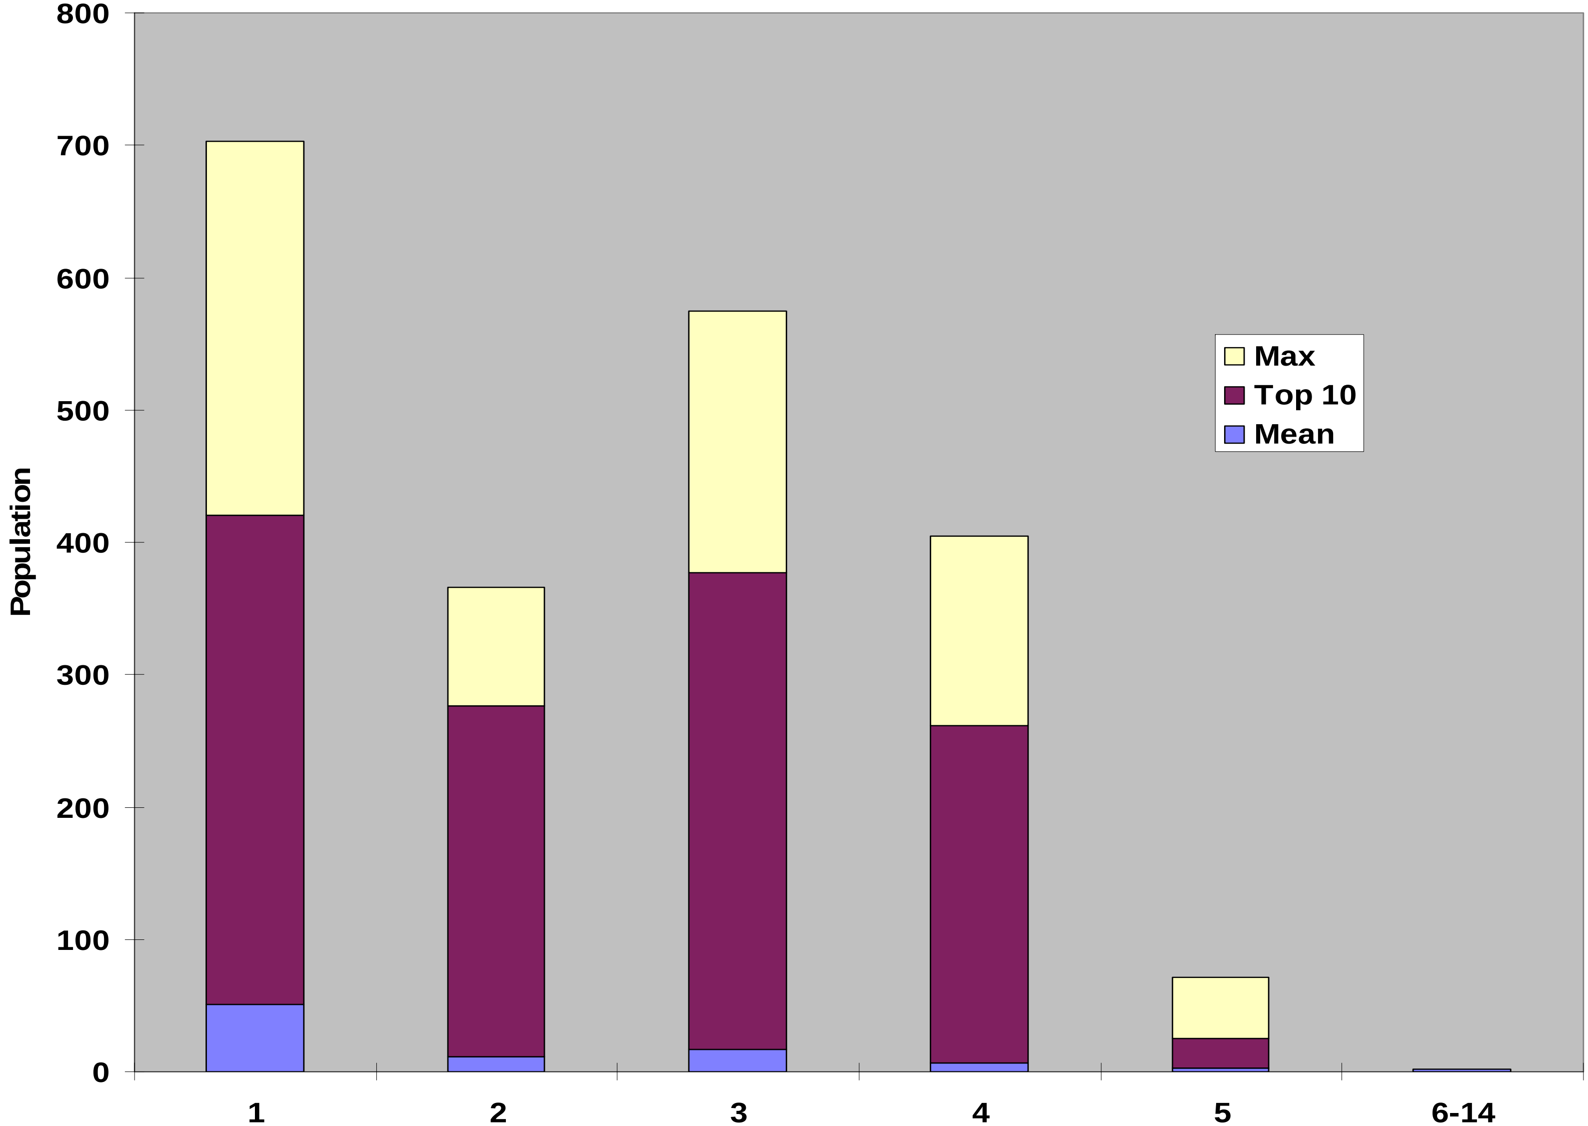

Time offsets

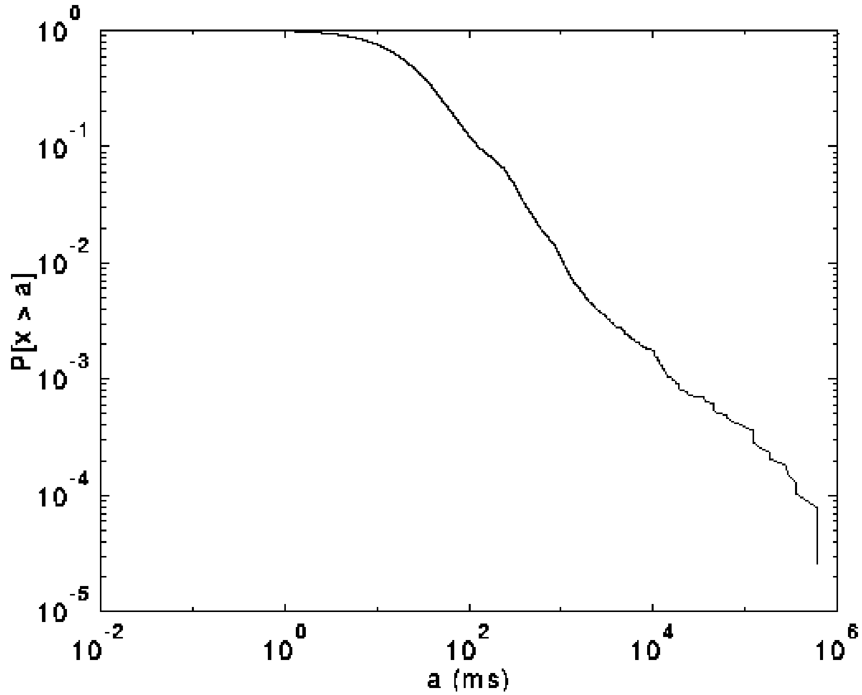

- Cumulative distribution function of absolute time offsets
  - 38,722 Internet servers surveyed running NTP Version 2 and 3
  - Offsets: median 23.3 ms, mean 234 ms, maximum 686 ms
  - Offsets < 128 ms: median 20.1 ms, mean 28.7 ms

Roundtrip delays

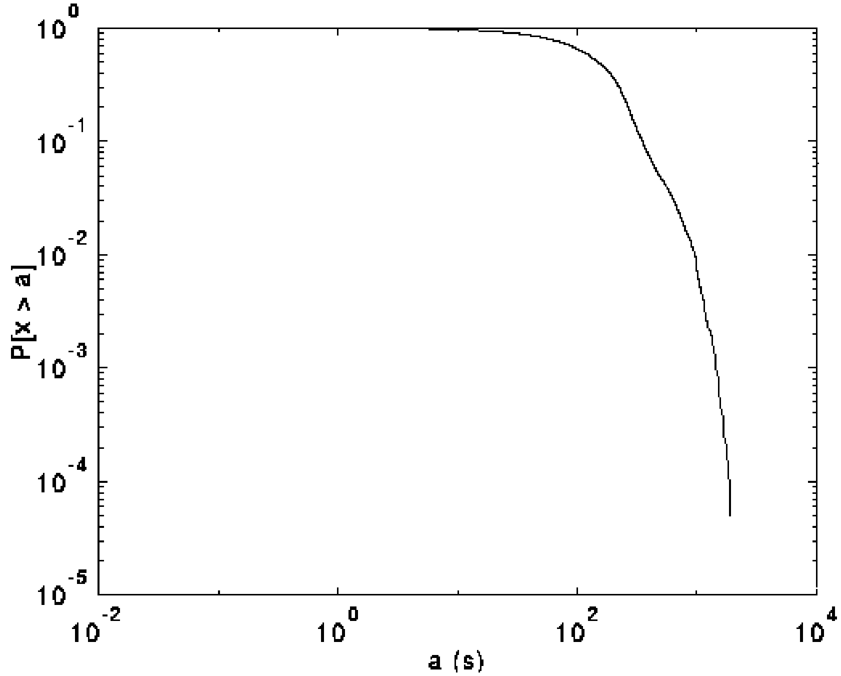

- Cumulative distribution function of absolute roundtrip delays
  – 38,722 Internet servers surveyed running NTP Version 2 and 3
  – Delays: median 118 ms, mean 186 ms, maximum 1.9 s(!)
  – Asymmetric delays can cause errors up to one-half the delay

Peer roundtrip delays

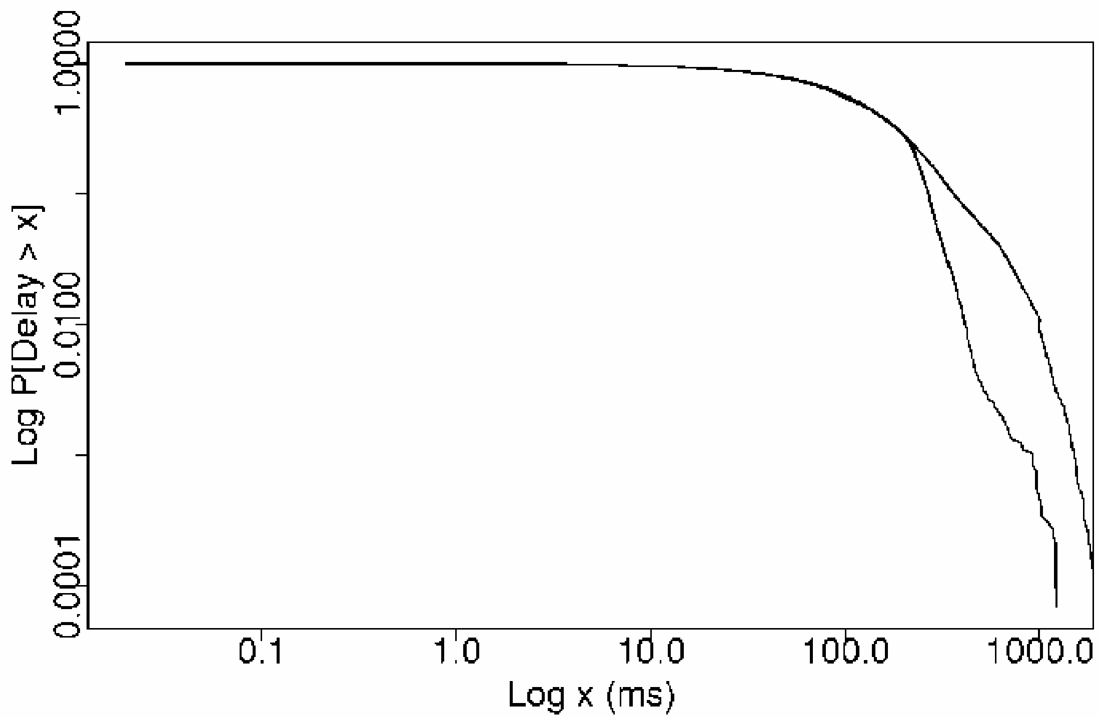

- Cumulative distribution of peer-peer absolute roundtrip delays
  - 182,538 samples excludes measurements where synchronization distance exceeds 1 s. since by specification these cannot synchronize the local clock
  - Upper curve: different subnets (median 118 ms, mean 173 ms, max 1.91 s)
  - Lower curve: same subnet (median 113 ms, mean 137 ms, max 1.40 s)

Systematic oscillator frequency offsets

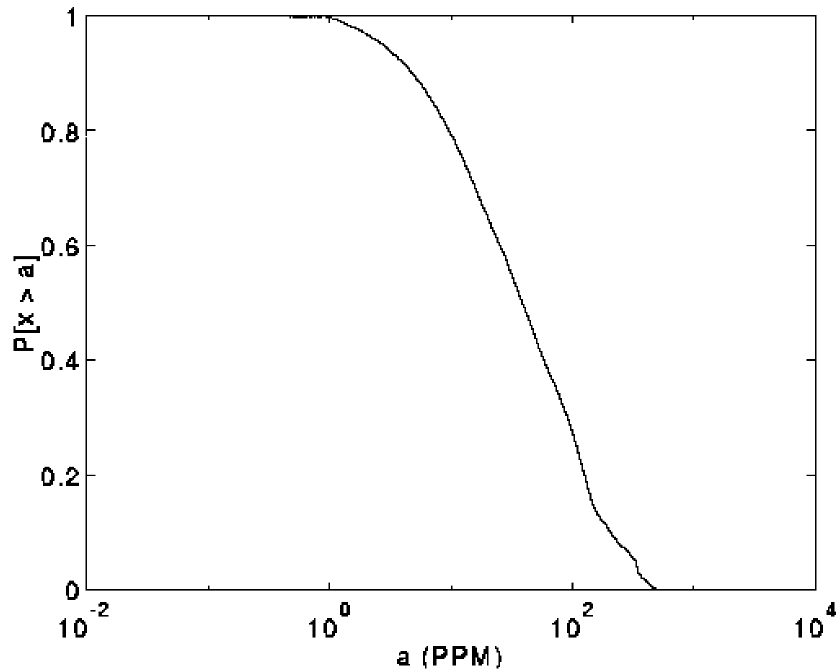

- Cumulative distribution function of absolute frequency offsets
  - 19,873 Internet servers surveyed running NTP Version 2 and 3
  - 593 outlyers greater than 500 PPM discarded as unsynchronized
  - Remaining offsets: median 38.6 PPM, mean 78.1 PPM

Local clock frequency offsets

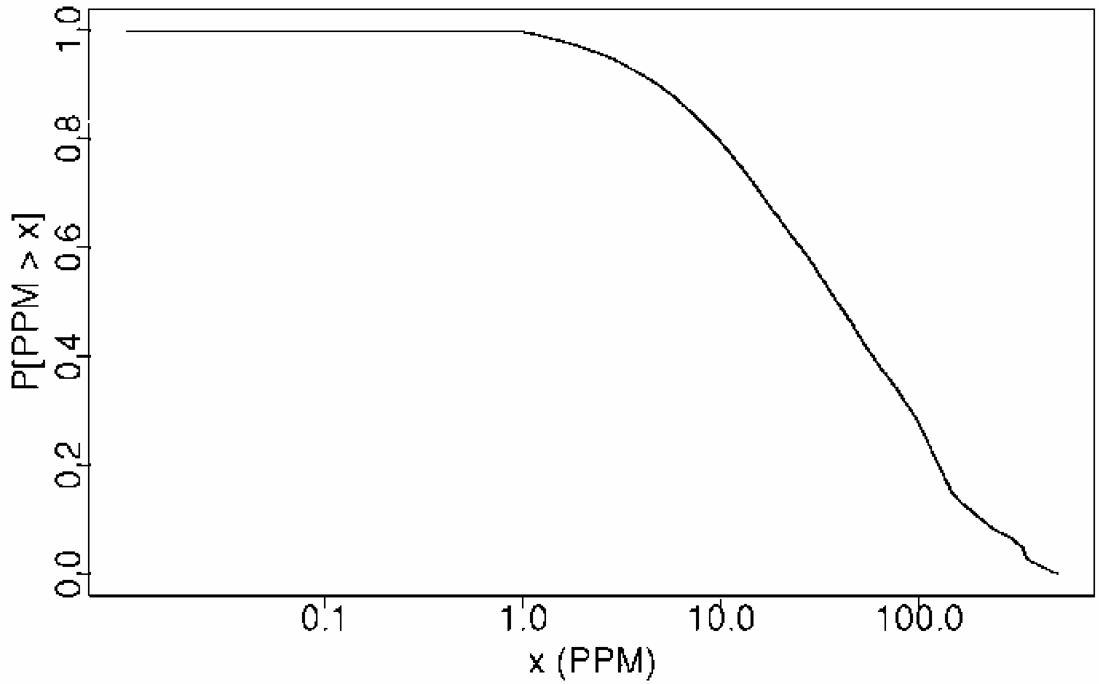

- Cumulative distribution of local clock absolute frequency offsets
  - 19,873 Internet peers surveyed running NTP Version 2 and 3
  - 396 offsets equal to zero deleted as probably spurious (self synchronized)
  - 593 offsets greater than 500 PPM deleted as probably unsynchronized
  - Remaining 18,884 offsets: median 38.6 PPM, mean 78.1 PPM

Clock oscillator phase errors

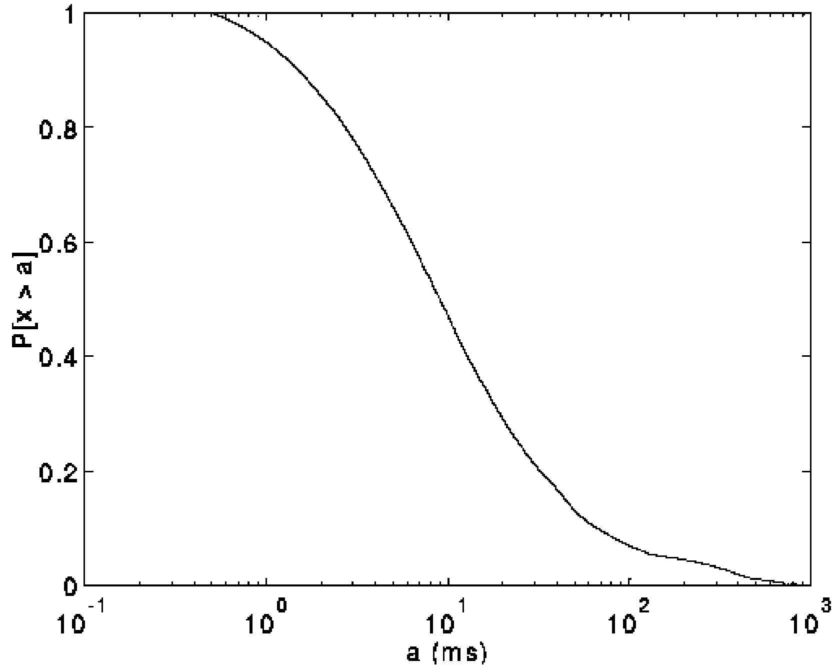

- Cumulative distribution function of absolute phase errors
  - 19,873 Internet servers surveyed running NTP Version 2 and 3
  - 131 outlyers with errors over 1 s discarded as unsynchronized
  - Remaining errors: median 9.1 ms, mean 37.0 ms

Local clock phase offsets

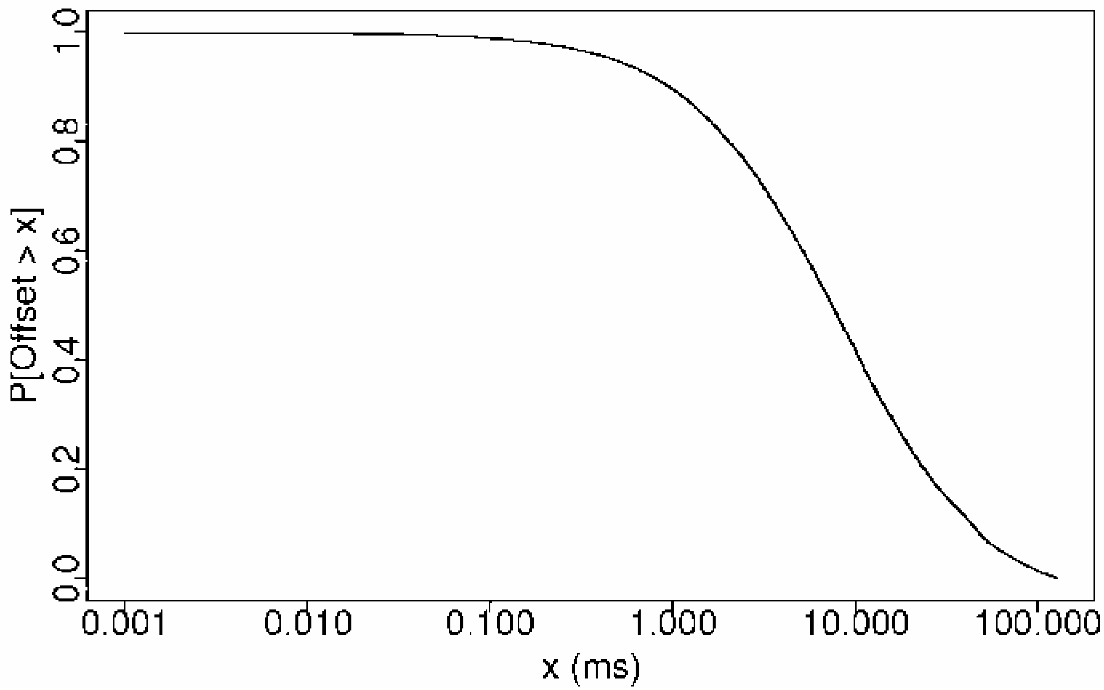

- Cumulative distribution of local clock absolute phase offsets
  - 19,873 Internet peers surveyed running NTP Version 2 and 3
  - 530 offsets equal to zero deleted as probably unsynchronized
  - 664 offsets greater than 128 ms deleted as probably unsynchronized
  - Remaining 18,679 offsets: median 7.45 ms, mean 15.87 ms

Peer clock offsets -same/different subnets

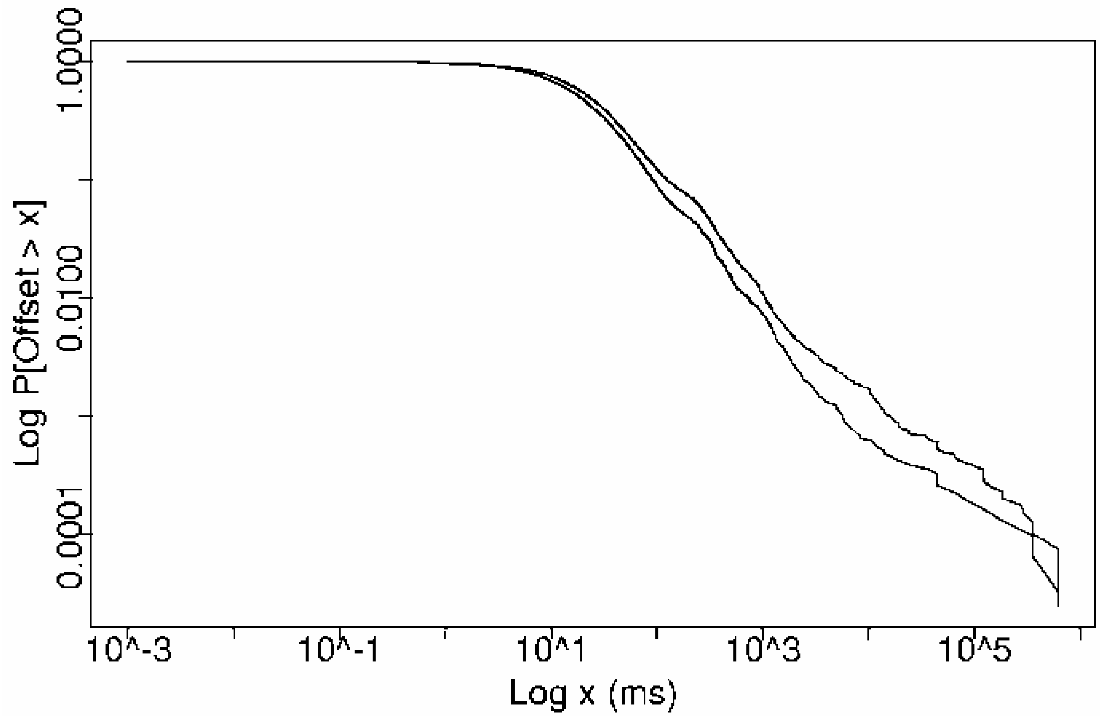

- Cumulative distribution function of peer-peer absolute clock offsets
  - 182,538 peers used by 34,679 clients, 85,730 on the same subnet, 96,808 on a different subnet.
  - Upper curve: different subnet (median 19 ms, mean 161 ms, max 621 s)
  - Lower curve: same subnet (median 13 ms, mean 188 ms, max 686 s)

Reference clock sources

- In a survey of 38,722 peers, found 1,733 primary and backup external reference sources
- 231 radio/satellite/modem primary sources
  - 47 GPS satellite (worldwide), GOES satellite (western hemisphere)
  - 57 WWVB radio (US)
  - 17 WWV radio (US)
  - 63 DCF77 radio (Europe)
  - 6 MSF radio (UK)
  - 5 CHU radio (Canada)
  - 7 modem (NIST and USNO (US), PTB (Germany), NPL (UK))
  - 25 other (cesium clock, precision PPS sources, etc.)
- 1,502 local clock backup sources (used only if all other sources fail)
- For some reason or other, 88 of the 1,733 sources appeared down at the time of the survey

Timekeeping facilities at UDel - December 1997

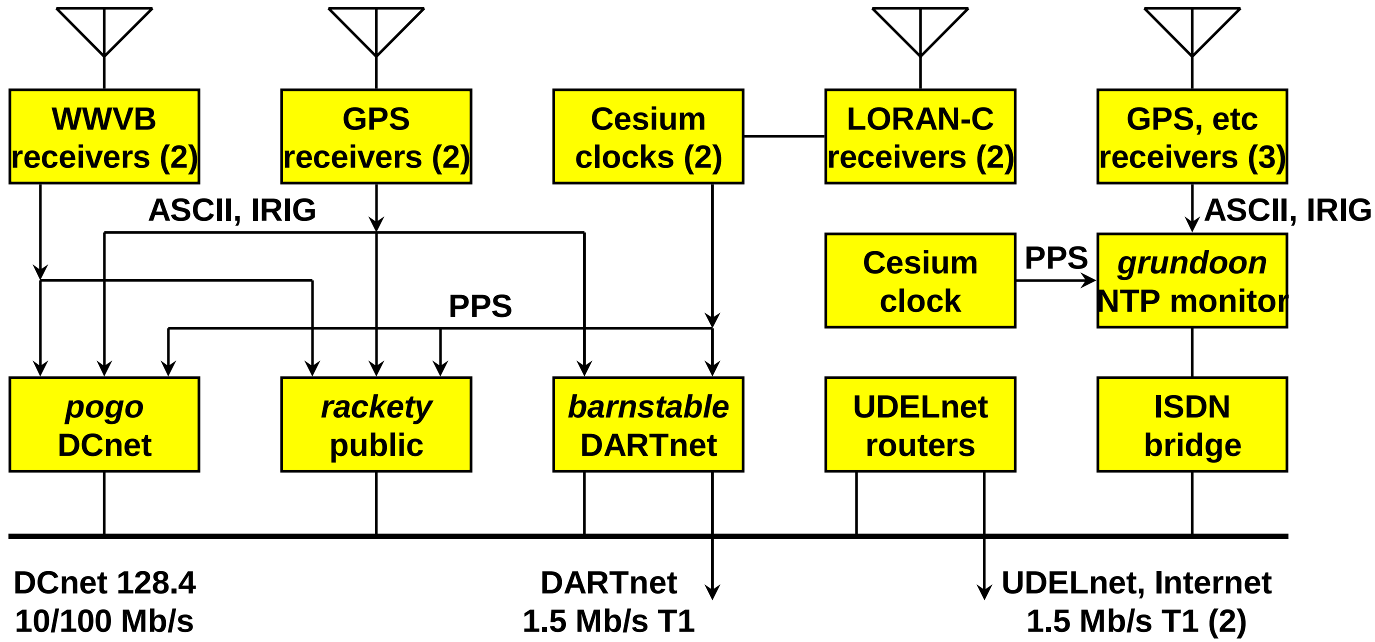

- Cesium oscillators are calibrated by U.S. Naval Observatory and checked continuously by Northeast US LORAN-C chain and GPS
- NTP primary time servers synchronize to ASCII, PPS and IRIG-B, all with kernel modifications for precision timekeeping
- NTP secondary servers (not shown) include SunOS 4/5, Ultrix 4, OSF/1, HP-UX, Cisco, Bancomm and Fuzzball (semi-retired)

Precision Timekeeping Equipment

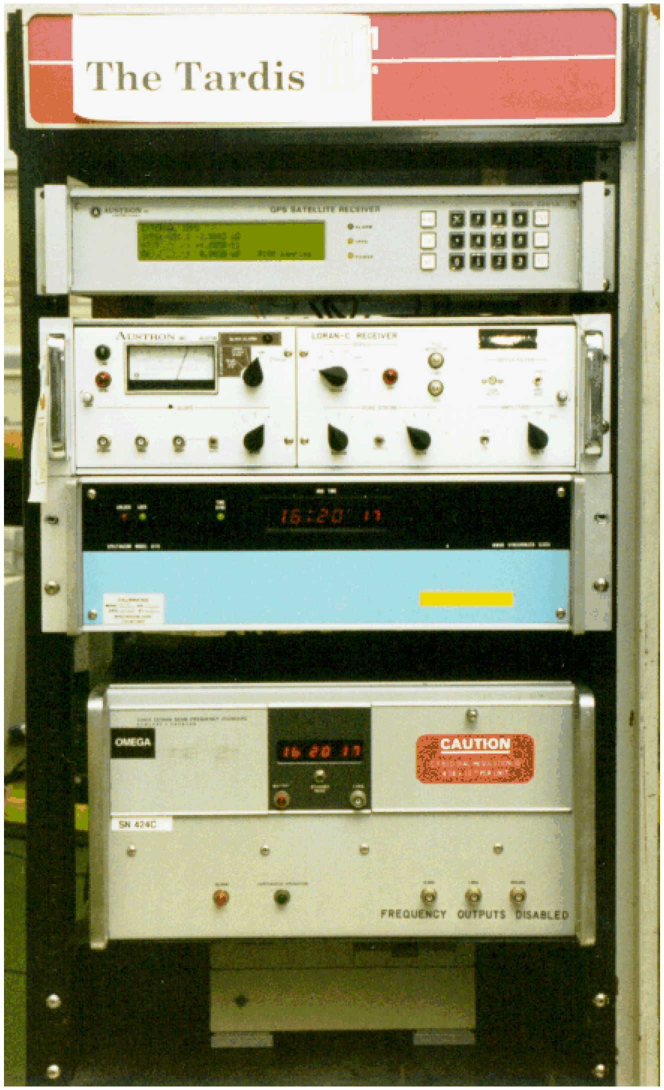

Austron 2200A GPS Receiver

Austron 2000 LORAN-C Receiver

Spectracom 8170 WWVB Reciver

Hewlett Packard 5061A Cesium Beam Frequency Standard

Squeezing the nanoseconds

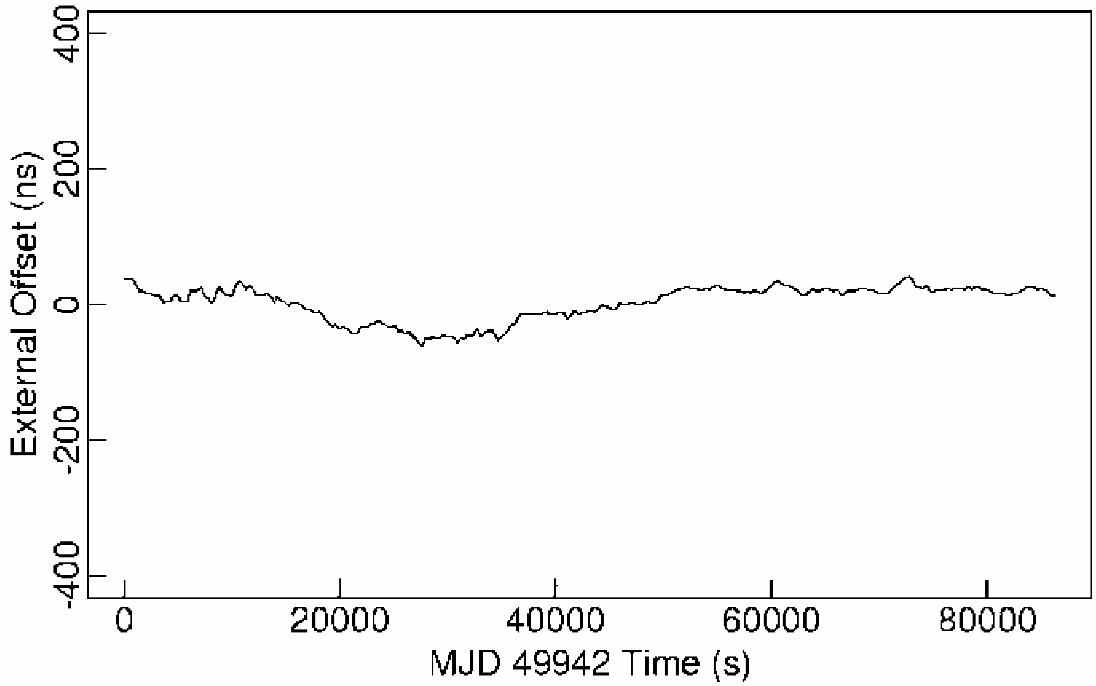

- This shows the residual error measured between the Austron 2201 GPS receiver and the HP 5061A cesium clock
- The GPS receiver is stabilized using the LORAN-C receiver, which improves its accuracy to about 50 ns, in spite of the intentional degradation introduced in the GPS signal available to the public

A day in the life of a busy NTP server

- NTP primary (stratum 1) server *rackety* is a Sun IPC running SunOS 4.1.3 and supporting 734 clients scattered all over the world
- This machine supports NFS, NTP, RIP, IGMP and a mess of printers, radio clocks and an 8-port serial multiplexor
- The mean input packat rate is 6.4 packets/second, which corresponds to a mean poll interval of 157 seconds for each client
- Each input packet generates an average of 0.64 output packets and requires a total of 2.4 ms of CPU time for the input/output transaction
- In total, the NTP service requires 1.54% of the available CPU time and generates 10.5, 608-bit packets per second, or 0.41% of a T1 line
- The conclusion drawn is that even a slow machine can support substantial numbers of clients with no significant degradation on other network services

The Sun never sets on NTP

- NTP is arguably the longest running, continuously operating, ubiquitously available protocol in the Internet
- USNO and NIST, as well as equivalents in other countries, provide multiple NTP primary servers directly synchronized to national standard cesium clock ensembles and GPS
- Over 230 Internet primary servers in Australia, Canada, Chile, France, Germany, Isreal, Italy, Holland, Japan, Norway, Spain, Sweden, Switzerland, UK, and US
- Over 100,000 Internet secondary servers and clients all over the world
- National and regional service providers BBN, MCI, Sprint, Alternet, etc.
- Agencies and organizations: US Weather Service, US Treasury Service, IRS, PBS, Merrill Lynch, Citicorp, GTE, Sun, DEC, HP, etc.
- Several private networks are reported to have over 10,000 NTP servers and clients; one (GTE) reports in the order of 30,000 NTP-equipped workstations and PCs

NTP online resources

- Internet (Draft) Standard RFC-1305 Version 3
  - Simple NTP (SNTP) RFC-2030
  - Designated SAFEnet standard (Navy)
  - Under consideration in ANSI, ITU, POSIX
- NTP web page http://www.eecis.udel.edu/~ntp
  - NTP Version 3 release notes and HTML documentation
  - List of public NTP time servers (primary and secondary)
  - NTP newsgroup and FAQ compendium
  - Tutorials, hints and bibliography
- NTP Version 3 implementation and documentation for Unix, VMS and Windows
  - Ported to over two dozen architectures and operating systems
  - Utility programs for remote monitoring, control and performance evaluation
  - Latest version on ftp.udel.edu in pub/ntp directory