


# NTP Performance Analysis

David L. Mills  
University of Delaware  
<http://www.eecis.udel.edu/~mills>  
<mailto:mills@udel.edu>




Sir John Tenniel; *Alice's Adventures in Wonderland*, Lewis Carroll

## Typical local clock phase offsets (1997 survey)



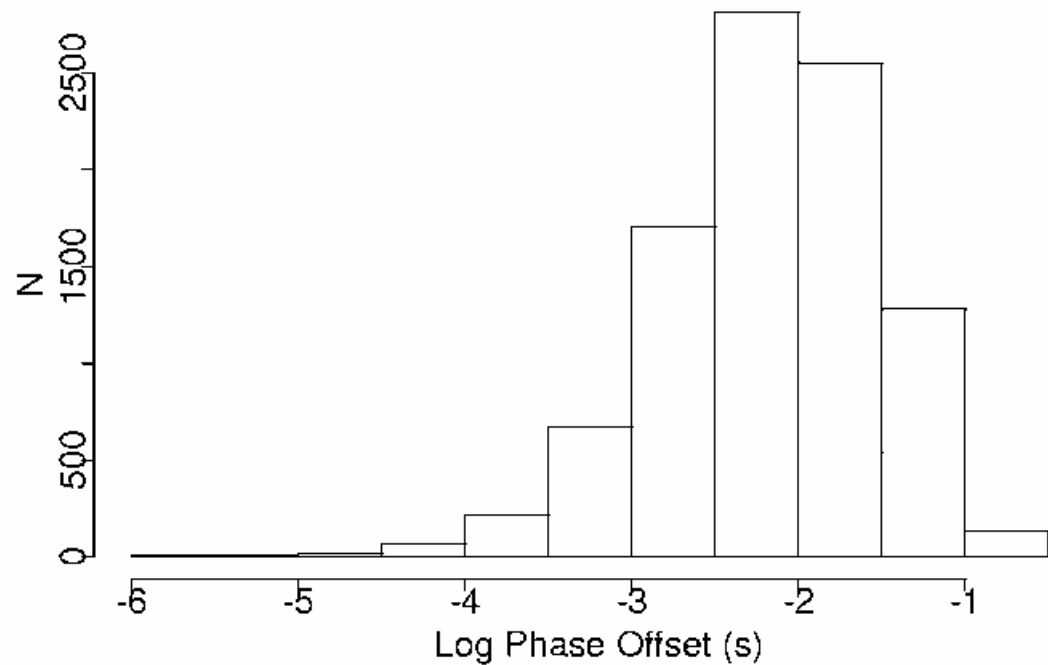
- Histogram of local clock absolute phase offsets
  - 19,873 Internet peers surveyed running NTP Version 2 and 3
  - 530 offsets equal to zero deleted as probably unsynchronized
  - 664 offsets greater than 128 ms deleted as probably unsynchronized
  - Remaining 18,679 offsets: median 7.45 ms, mean 15.87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

## Typical local clock frequency offsets (1997 survey)




- Histogram 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

## Typical local clock phase offsets (from survey)



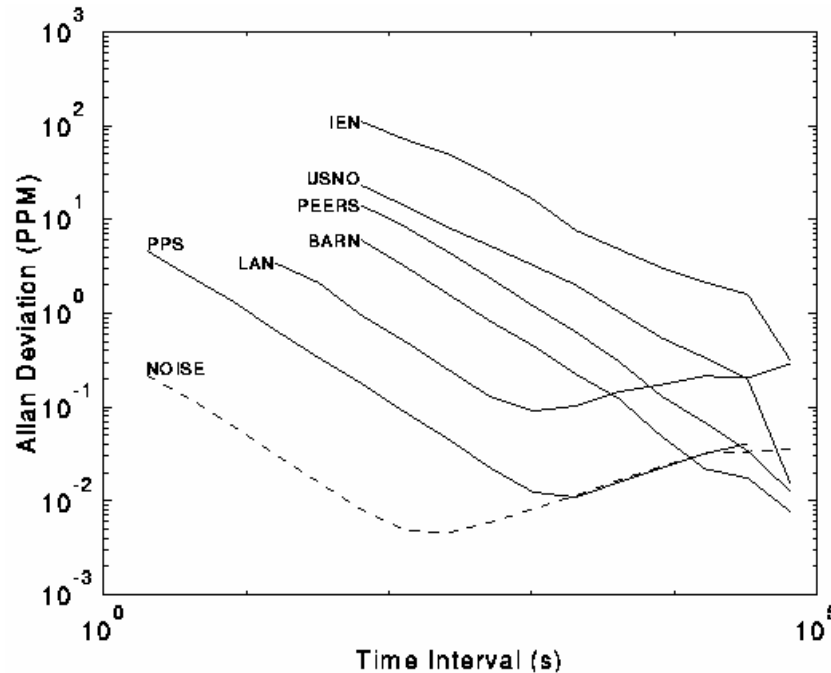
- Histogram of local clock absolute phase offsets
  - 19,873 Internet peers surveyed running NTP Version 2 and 3
  - 530 offsets equal to zero deleted as probably unsynchronized
  - 664 offsets greater than 128 ms deleted as probably unsynchronized
  - Remaining 18,679 offsets: median 7.45 ms, mean 15.87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

## Typical local clock frequency offsets (from survey)



- Histogram 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

# Allan deviation - combined data





## Legend

IEN	IEN Torino, Italy
USNO	US Naval Observatory, Wash DC
PEERS	19 nonlocal time servers in Europe, Japan, Australia, North and South America
BARN	local time server on DCnet
LAN	free-running clock via Ethernet
PPS	free-running clock via PPS signal
NOISE	free-running clock (synthesized)

All servers are synchronized to GPS  
All NTP algorithms are operative


- Vee-shaped curves show local servers with free-running local clocks; other curves show remote servers synchronized to GPS
  - Lines with slope  $-1$  represent white phase noise due to network jitter
  - Lines with slope  $+0.5$  represent random-walk frequency noise due to clock oscillator wander
  - Intersection of phase and frequency noise lines is called the Allan intercept
  - In general, PLL is best when  $T_c$  is below Allan intercept; FLL is best above it

## Compare FLL, PLL and hybrid modes for USNO peer



- Left graph shows FLL and PLL weights, right graph shows standard error as a function of poll interval for FLL, PLL and hybrid modes
  - Note FLL is best above 200 s, PLL is best below this
  - Hybrid mode is best between most important range, 100 s to 1,000 s, and not much worse than FLL or PLL outside this range.
  - PLL comes unstable above 4,000 s due to loss of lock.

# Standard error for various network paths




- Legend
- IEN IEN Torino, Italy
  - USNO US Naval Observatory, Wash DC
  - PEERS 19 nonlocal time servers in Europe, Japan, Australia, North and South America
  - BARN local time server on DCnet

All servers are synchronized to GPS  
All NTP algorithms are operative

- Solid lines show hybrid mode performance, dashed lines PLL mode, both over a ten-day period
  - Hybrid mode better than PLL mode by a factor of ten over important range
  - Local time server better than 200  $\mu$ s standard error at poll = 64 s
  - All nonlocal time servers better than 2 ms at poll  $\leq$  1,024 s
  - Standard error of all nonlocal time servers (including best USNO) is better than any server separately



# NTP performance compared



## Legend

IEN	IEN Torino, Italy
USNO	US Naval Observatory, Wash DC
BARN	local time server on Dcnet
PPS	PPS signal (64-s poll clamp)

## Label Format

1	server name
2	mean poll interval (s)
3	mean error (ms)
4	RMS error (ms)
5	max error (ms)

All servers are synchronized to GPS  
All NTP algorithms are operative

- Typical performance of stratum-2 servers synchronized to remote primary servers
  - Except for PPS, which uses simulated phase noise, all use actual network noise measured in real time
  - Frequency noise is simulated with curve fit to PPS data


## Performance of typical NTP servers in the global Internet




NTP Server	Location	Days	Mean	RMS	Max	>1	>5	>10	>50
Austron GPS	DCnet	91	0.0	0.012	1.000	0	0	0	0
rackety	DCnet	95	0.066	0.053	2.054	11	0	0	0
mizbeaver	DCnet	17	0.150	0.171	1.141	2	0	0	0
churchy	DCnet	42	0.185	0.227	3.150	15	0	0	0
pogo	DCNet	88	0.091	0.057	1.588	8	0	0	0
beauregard	DCnet	187	0.016	0.108	2.688	30	0	0	0
umd1	U Maryland	78	4.266	2.669	35.89	29	29	28	0
swifty	Australia	84	2.364	56.70	3944	27	27	27	13
ntps1	Germany	70	0.810	10.86	490.9	12	12	12	6
time_a	NIST Boulder	85	1.511	1.686	80.56	28	19	11	2
fuzz	San Diego	77	3.889	2.632	47.59	27	27	23	0
la	Los Angeles	83	0.650	0.771	17.84	28	8	3	0
enss136	NSFnet WashDC	88	0.657	1.203	32.65	38	23	10	0

- Table shows number days surveyed, mean absolute offsets (ms), RMS and maximum absolute error (ms) and number of days on which the maximum error exceeded 1, 5, 10 and 50 ms at least once
- Servers represent LANs, domestic WANs and worldwide Internet
- Results show all causes, including software upgrades and reboots

# Measured PPS time error for Alpha 433




## Performance with a secondary server via Ethernet




- Clock offsets for Sun SPARC 1+ and SunOS 4.1.1 over four days
  - Primary server synchronized to GPS with PPS
  - Spikes are due to Ethernet jitter and collisions
  - Wander is due to client clock oscillator instability

## Performance with a secondary server via T1 line




- Clock offsets measured for a NSFnet secondary server running NTP
  - Measurements use NSF server synchronized to a primary server via Ethernets and T1 tail circuit
  - This is typical behavior for lightly loaded T1 circuit

## Closed-loop characteristics of primary servers




(a) Clock Offset Relative to GPS



(b) Clock Offset between Two Primary Servers


- Clock offsets for Sun SPARC 1+ and SunOS 4.1.1 over one day
  - Two primary servers, both synchronized to the same GPS receiver (no PPS)
  - (a) Measured GPS receiver relative to the local clock of either server
  - (b) Measured one server across the Ethernet relative to the local clock of the other server
  - Note 300- $\mu$ s spike of unknown cause is visible in both (a) and (b)

## Performance with a modem and ACTS service



- Measurements use 2300-bps telephone modem and NIST Automated Computer Time Service (ACTS)
- Calls are placed via PSTN at 16,384-s intervals


## Time offsets with an Australian primary server




- Transmission path is one way via satellite, the other way via undersea cable
- This surely is an extreme case of network jitter and congestion



## Typical frequency variations with temperature




(a) Frequency Offset Measured by PPS




(b) Frequency Offset Measured by NTP

- Measured frequency offsets for free-running local clock oscillator
- (a) Measured directly using PPS signal and `ppsclock` clock discipline
  - Typical room temperature thermostatically controlled in winter
- (b) Measured indirectly using NTP and host synchronized to PPS signal
  - Room temperature follows the ambient in first nice days in spring

# Errors due to kernel latencies



(a) Latency for `gettimeofday()` Call



(b) Latency Distribution for (a)

- These graphs were constructed using a Digital Alpha and OSF/1 V3.2 with precision time kernel modifications (now standard)
- (a) Measured latency for `gettimeofday()` call
  - spikes are due to timer interrupt routine
- (b) Probability distribution for (a) measured over about ten minutes
  - Note peaks near 1 ms due timer interrupt routine, others may be due to cache reloads, context switches and time slicing
  - Biggest surprise is very long tail to large fractions of a second

## Further information

---



- NTP home page <http://www.ntp.org>
  - Current NTP Version 3 and 4 software and documentation
  - FAQ and links to other sources and interesting places
- David L. Mills home page <http://www.eecis.udel.edu/~mills>
  - Papers, reports and memoranda in PostScript and PDF formats
  - Briefings in HTML, PostScript, PowerPoint and PDF formats
  - Collaboration resources hardware, software and documentation
  - Songs, photo galleries and after-dinner speech scripts
- Udel FTP server: <ftp://ftp.udel.edu/pub/ntp>
  - Current NTP Version software, documentation and support
  - Collaboration resources and junkbox
- Related projects <http://www.eecis.udel.edu/~mills/status.htm>
  - Current research project descriptions and briefings