Autonomous Configuration

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

Briefing roadmap on NTP technology and performance

- NTP project page http://www.eecis.udel.edu/~mills/ntp.html/.
  - Network Time Protocol (NTP) General Overview
  - NTP Architecture, Protocol and Algorithms
  - NTP Procedure Descriptions and Flow Diagrams
  - NTP Cryptographic Authentication (Autokey)
  - NTP Clock Discipline Principles
  - NTP Precision Synchronization
  - NTP Performance Analysis
  - NTP Algorithm Analysis
  - Long-range Dependency Effects in NTP Timekeeping

NTP architecture review

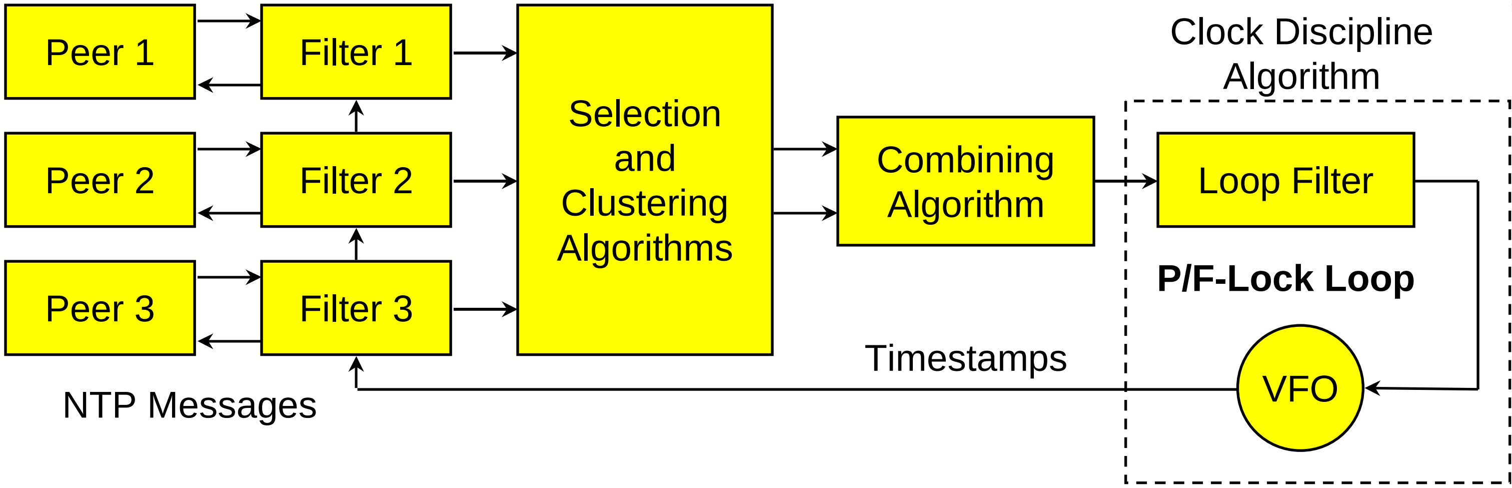

Peer 1

Peer 2

Peer 3

Filter 1

Filter 2

Filter 3

Selection and Clustering Algorithms

Combining Algorithm

Clock Discipline Algorithm

Loop Filter

P/F-Lock Loop

VFO

Timestamps

NTP Messages

- Multiple servers/peers provide redundancy and diversity.
- Clock filters select best from a window of eight time offset samples.
- Intersection and clustering algorithms pick best *truechimers* and discard *falsetickers.*
- Combining algorithm computes weighted average of time offsets.
- Loop filter and variable frequency oscillator (VFO) implement hybrid phase/frequency-lock (P/F) feedback loop to minimize jitter and wander.

The NTP subnet

department servers (stratum 3)

campus secondary servers (stratum 2)

Internet primary servers (stratum 1)

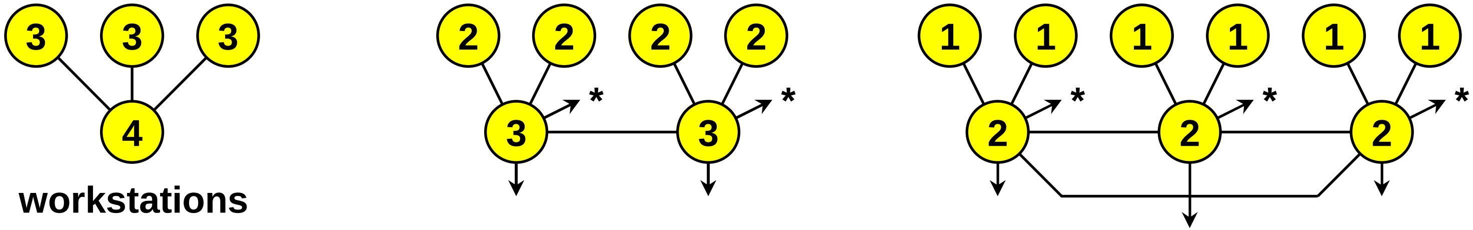

workstations (stratum 4)

* to buddy in another subnet

- NTP synchronizes the clocks of hosts and routers in the Internet
- Time synchronization flows from primary servers synchronized via radio and satellite over hierarchical subnet to other servers and clients
- NTP provides submillisecond accuracy on LANs, low tens of milliseconds on typical WANs spanning the country
- NTP software daemon has been ported to almost every workstation and server platform available today, including Unix, Windows and VMS
- Well over 100,000 NTP clients and servers are now deployed in the Internet and its tributaries all over the world

NTP autonomous system model

- Fire-and-forget software
  - Single software distribution can be compiled and installed automatically on most host architectures and operating systems
  - Run-time configuration can be automatically determined and maintained in response to changing network topology and server availability
- Autonomous configuration (autoconfigure)
  - Survey nearby network environment to construct a list of suitable servers
  - Select best servers from among the list using a defined metric
  - Reconfigure the NTP subnet for best accuracy with overhead constraints
  - Periodically refresh the list in order to adapt to changing topology
- Autonomous authentication (autokey)
  - For each new server found, fetch its cryptographic credentials from public databases
  - Authenticate each NTP message received as sent by that server and no other
  - Regenerate keys in a timely manner to avoid compromise

Goals and non-goals

- Goals
  - Robustness to many and varied kinds of failures, including Byzantine, fail-stop, malicious attacks and implementation bugs
  - Maximum utilization of Internet multicast services and protocols
  - Depend only on public values and certificates stored in secure directory services
  - Fast operation using a combination of public-key and private-key cryptography
- Non-goals
  - Administrative restrictions (multicast group membership control)
  - Access control - this is provided by firewalls and address filtering
  - Privacy - all protocol values, including time values, are public
  - Protection against out of order or duplicated messages - this is provided by the NTP protocol
  - Non-repudiation - this can be provided by a layered protocol if necessary

Autonomous configuration and authentication - issues

- Configuration and authentication and synchronization are inseparable
- Autonomous configuration (autoconfigure)
  - Centralized configuration management does not scale to large networks
  - Finding optimal topologies in large subnet graphs under degree and distance constraints is NP-hard
  - Greedy heuristics may not produce good topologies in acceptable time
  - Solution may involve span-limited, hierarchical multicast groups and add/drop heuristics
- Autonomous authentication (autokey)
  - Centralized key management does not scale to large networks
  - Symmetric-key cryptosystems require pairwise key agreement and persistent state in clients and servers
  - Servers cannot maintain persistent state for possibly thousands of clients
  - Public-key cryptosystems are too slow for good timekeeping
  - Solution may involve a combination of public and private key cryptosystems

Autonomous configuration - approach

- Dynamic peer discovery schemes
  - Primary discovery vehicle using NTP multicast and anycast modes
  - Augmented by DNS, web and service location protocols
  - Augmented by NTP subnet search using standard monitoring facilities
- Automatic optimal configuration
  - Distance metric designed to maximize accuracy and reliability
  - Constraints due to resource limitations and maximum distance
  - Complexity issues require intelligent heuristic
- Candidate optimization algorithms
  - Multicast with or without initial propagation delay calibration
  - Anycast mode with administratively and/or TTL delimited scope
  - Distributed, hierarchical, greedy add/drop heuristic
- Proof of concept based on simulation and implementation with NTP Version 4

NTP configuration scheme

- Multicast scheme (moderate accuracy)
  - Servers flood local area with periodic multicast response messages
  - Clients use client/server unicast mode on initial contact to measure propagation delay, then continue in listen-only mode
- Manycast scheme (highest accuracy)
  - Initially, clients flood local area with a multicast request message
  - Servers respond with multicast response messages
  - Clients continue with servers as if in ordinary configured unicast client/server mode
- Both schemes require effective implosion/explosion controls
  - Expanding-ring search used with TTL and administrative scope
  - Excess network traffic avoided using multicast responses and rumor diffusion
  - Excess client/server population controlled using NTP clustering algorithm and timeout garbage collection

Discovery mechanisms

- The emphasis here is on autonomous configuration and repair; discovery schemes in themselves are secondary
- NTP multicast and/or anycast modes used to discover servers within the same hierarchical group; groups may be tiled over Internet
- Ancestors of hierarchical group discovered from NTP peer data, augmented by NTP monitoring data
- Authentication verified by DNS lookup and MD5 message digest
- Database is synthesized from all these data and distributed to "interested" servers and clients
- Interested servers and clients run a heuristic algorithm to construct hierarchical subnet topology

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