The NTP subnet in New Zealand

Paul Ashton
Department of Computer Science
University of Canterbury
paul@cosc.canterbury.ac.nz

TR-COSC 01/99, March 1999

Abstract

The Network Time Protocol (NTP) is used to synchronise the clocks of a significant percentage of the computers that make up the internet. Occasional large adjustments made to the local clock of our department's main NTP server (`kaka`) lead us to investigate the state of the NTP subnet in New Zealand. We examined the structure of the NTP subnet within New Zealand, recorded performance information for over 350 New Zealand NTP servers and investigated the reasons for the occasional large adjustments made to `kaka`'s clock.

The overall NTP performance figures for New Zealand NTP servers compare favourably with those reported in a world-wide study of NTP performance. However, the lack of any primary NTP servers in New Zealand, and the concentration of all secondary servers at a single site, were major factors in the occasional large adjustments made on `kaka`. Also, a feature of the NTP filtering algorithm was a factor in some of the large adjustments. At a more detailed level, several cases were found in which an NTP configuration could be improved.

Our overall conclusion is that the accuracy of NTP in New Zealand could be improved by installing two or three primary servers, backed up by several secondary servers.

1 Introduction

The Network Time Protocol (NTP) is used extensively throughout the Internet to keep computer time of day clocks close to Universal Time (UTC) [4, 2]. A host on which the NTP software runs is known as an NTP server. A small proportion of NTP servers are primary servers that are connected to external sources of UTC, such as GPS receivers or radio clocks. All other NTP servers synchronise, directly or indirectly, to one or more primary servers. The most significant source of error in NTP timekeeping is variability in network delays experienced by NTP packets that contain timing information.

Many New Zealand organisations operate NTP servers. Several reasons exist for believing that the accuracy achieved by NTP in New Zealand may be worse than that achieved in other parts of the Internet:

- There are no public primary servers in New Zealand at present. The New Zealand NTP subnet relies on overseas primary servers (most in the United States and Australia) that are accessed via long distance network links that are frequently congested.
- There are three public[1] secondary NTP servers in New Zealand—all operated by the University of Waikato. With all three secondary servers on the same site, the New Zealand NTP subnet is more vulnerable to congestion and failure of network links to Waikato than it would be if secondary servers were available at several sites.
- In earlier work on measuring the accuracy of NTP using a locally-developed test bed [1], it was noted that our department's main NTP server would sometimes step its clock by a large amount (an amount over 128 milliseconds). Such steps are signs that NTP is not working well.

Consequently, a study has been undertaken to investigate the health of the New Zealand NTP subnet. Also of interest was the extent to which NTP is being used in New Zealand, and the structure of the New Zealand NTP subnet.

In the next section, we give an overview of how NTP operates, and discuss previous NTP studies. This is followed by a description of our experimental procedures, details of the structure of the New Zealand NTP subnet, and the overall performance of NTP within New Zealand. Next, we detail some reasons behind the large steps observed on our local NTP servers. Finally we present our conclusions.

Note that in making observations about the New Zealand NTP subnet there is no intent to be critical of the system administrators involved. Their configuration may well meet their clock synchronisation needs very well.

[1] That is, listed in `http://www.eecis.udel.edu/~mills/ntp/servers.htm`

2 Background

The Network Time Protocol (NTP) has evolved over a period of nearly 20 years. Version 3 is very widely used—the number of computers whose clocks are synchronised by NTP worldwide is probably in the hundreds of thousands [3]—and development of version 4 of the protocol is well advanced. The main design goal of NTP is to tightly synchronise to UTC the clocks of a large number of computers distributed throughout the Internet. Important sub-goals are that NTP should be able to scale to huge numbers of hosts, and that NTP should operate well in the wide range of network architectures and traffic conditions present in today's Internet.

Each NTP server exchanges timing information with between 1 and a few hundred other NTP servers (commonly this number is less than 10). Communication links between NTP servers are arranged in a broadly hierarchical fashion. A relatively small group of NTP servers are attached to devices such as GPS receivers and radio clocks that are direct sources of UTC timing information. An NTP server with direct access to a source of external time is said to be at stratum 1 (level 1) of the NTP hierarchy. Every other NTP server synchronises its clock, directly or indirectly, to the clocks of one or more stratum 1 servers.

An NTP server at stratum N is configured (by hand) to exchange timing information with 1 or more *peers*, at least one of which is a stratum $N-1$ NTP server, with the others being stratum N servers (peers at the same stratum are primarily used to provide backup synchronisation paths to be used when lower stratum peers cannot be contacted or are faulty). For each peer, an NTP server regularly estimates the *offset* between its own clock and that of the peer based on time stamps present in a message sent to the peer and a reply message sent back by the peer. The main source of error in a clock offset estimate is asymmetry in the network delays experienced by the two packets. While the delay experienced by each message is unknown, the *round trip delay* (the sum of the two delays) is known, and the maximum error due to packet delay asymmetry is simply half the round trip delay.

An NTP server uses filtering, intersection and clustering, and combining algorithms to translate the streams of offset estimates produced for each peer to a single estimate of the offset between the NTP server's local clock and UTC. Each stream of peer offset estimates is filtered so that estimates calculated from packets that exhibited a large delay are ignored because of the large error that could be present in the offset estimates. Filtering produces one offset estimate per peer. These offsets are put through intersection and clustering algorithms which discard some offsets and retain those deemed to be of the highest accuracy. The final overall offset is a weighted average of the offsets that survive the clustering and intersection algorithms. The peer whose offset had the most influence on the overall offset is deemed the *system peer*, and the stratum of the NTP server is 1 more than the stratum of the system peer.

The structure of the NTP subnet is determined in a somewhat anarchic fashion by the system administrators of the various NTP servers. Advice in the NTP documentation suggests that for large sites, there should be three on-site NTP servers at stratum N, with remaining on-site NTP servers arranged in a synchronisation subnet beneath them. Consequently, the remaining on-site servers run at stratum $N+1$ or more. It is recommended that each of the on-site stratum N servers synchronises to at least two independent off-site servers at stratum $N-1$ (making a total of at least 6 stratum $N-1$ servers). Also, each of the stratum N servers should synchronise to each other to provide further redundancy. It is strongly recommended that the stratum $N-1$ off-site servers be accessed via different network links so as to ensure no single point of failure exists.

If highly accurate synchronisation is required at a particular site, it is likely that the main on-site servers would be equipped with external sources of time, and operate at stratum 1. Otherwise, the main on-site servers usually operate at stratum 2 or 3 by peering with off-site servers that operate at strata 1 or 2. Other on-site servers therefore operate at strata 3, 4 or 5 (few hosts in the global NTP subnet are at stratum 6 or more, even though the software can support NTP servers operating at stratum 15). Within a site, it is possible to distribute time (at the cost of some reduction in accuracy) via lightweight mechanisms based on broadcast and multicast communication.

To assist a system administrator in selecting lower stratum NTP hosts to peer with, a web page is maintained listing public primary (stratum 1) and secondary (stratum 2) time servers. A public NTP server is one that is willing to provide time information to NTP servers at other sites. At the time of writing there were 70 public primary servers and just over 100 public secondary servers. No formal mechanism exists for getting information on public servers at stratum 3, or on other primary or secondary servers.

For a system administrator in New Zealand, there are only two entries in the lists of public NTP servers that mention New Zealand. In the primary list, `ntp.cs.mu.OZ.AU` (located in Melbourne, Australia) includes New Zealand in its service area. Three New Zealand stratum 2 servers are listed in the secondary list—all located at Waikato University. These facts underscore some of the potential problems with the New Zealand NTP subnet raised in the introduction. With no primary NTP servers in New Zealand, the quality of all timing information is dependent on the delays experienced on our international Internet links. Also, with all three public secondary servers located at one site, New Zealand system administrators either have to synchronise only to the three Waikato servers (which creates a single point of failure, and makes accuracy dependent on low delays on paths to and from Waikato) or maintain links with overseas primary and secondary servers (which involve much longer delays, and which in some cases will incur traffic charges).

3 Experimental procedures

Two of our major goals were to uncover the current structure of the New Zealand part of the NTP subnet, and to compare v performance measures taken from New Zealand NTP hosts with those reported in studies of the global NTP subnet.

A data collection program, `ntpcrawl`, was written to gather data on both the structure and accuracy of the New Zealand NTP subnet. The program uses the standard `xntpdc` program to interrogate NTP servers. `ntpcrawl` takes as a command line argument the host name of the first NTP server *start* that is to be queried. Four `xntpdc` queries are made of that host:

- `sysinfo`, which returns the internet address of the system peer, the stratum of NTP server *start*, its root distance and root dispersion.
- `loopinfo`, which returns the current loop offset and the frequency difference between the local oscillator and UTC.
- `peers`, which returns for each peer its IP number, its stratum, the interval between polls initiated by *start*, the reachability mask (shows which of the last 8 messages sent to the peer were replied to), round trip delay to the peer, estimate of the offset between the peer's clock and that of *start*, and the dispersion measure associated with that offset.
- `monlist`, which returns a list of internet addresses of NTP servers that have communicated with *start*.

Once information from *start* has been gathered, `ntpcrawl` begins working its way through the NTP servers listed in the output of the `peers` and `monlist` requests. As new NTP servers are discovered, they are added to the list of servers to be visited in the future. `ntpcrawl` keeps working its way through NTP servers until all NTP servers discovered have been visited.

Some NTP servers listed by `peers` and `monlist` (for some NTP server S) are not visited, or have only a subset of the four requests sent to them. This happens in the following situations:

- If S is not in New Zealand, only the `sysinfo` and `loopinfo` requests are attempted (in other words, when a server outside New Zealand is reached no attempt is made to continue the search past S). A server is deemed to be in New Zealand if its DNS hostname ends in ".nz". This is not a fool-proof test of whether a host is in New Zealand or not (some `.nz` hosts we encountered were clearly in Los Angeles) but its accuracy is reasonably high.

- If either the `sysinfo` or `loopinfo` request fails, then the `peers` and `monlist` requests are not made.
- If the stratum of S is 16 (NTP sets the stratum to 16 if the local clock is not synchronised) then the `peers` and `monlist` requests are not made.
- A host H returned by `monlist` is visited only if:
  - H has been communicating with S on port 123 (port 123 is used by NTP daemons)
  - H has exchanged at least 25 packets with S (indicates regular contact between H and S which makes it likely that H is a client of S).
  - the version of NTP being used by H must be 3 (there was a concern that some of the data collected from earlier versions of NTP might have been calculated in a different way).
  - H and S must have communicated in the last 1024 seconds.

  Some of these tests are made to ensure that we narrow down the hosts returned by `monlist` to those that are likely to be currently active clients or peers of S.

Inevitably, `ntpcrawl` did not find all NTP servers in New Zealand. Many NTP servers are hidden behind firewalls, and were inaccessible to the computer on which `ntpcrawl` was run. At each site these inaccessible servers tend to be servers at higher strata. The key lower stratum servers need to be able to communicate with off-site peers, so are generally accessible through firewalls.

In addition `ntpcrawl` will not find NTP servers that have no connection with the three Waikato secondary servers. Also, some NTP servers refuse to respond to `xntpdc` queries. Nevertheless, a sizeable number of New Zealand NTP servers (355) were successfully visited. Data collected from these servers is presented in the next two sections.

One thing to bear in mind when looking at the figures is that because the data was collected on one of our machines, a group of 91 stratum 4 servers in the `cosc.canterbury.ac.nz` domain were included in the data collected. These machines are not accessible through the University firewall, and would have been excluded had the monitoring been done outside the University.

4 The structure of the New Zealand NTP subnet

To collect data over a range of network conditions, `ntpcrawl` was run four times a day (starting at 0305, 0905, 1505 and 2105) for seven successive days. With the exception of one run during which there was a network outage, each run took between 42 and 67 minutes to complete, with only 5 of the 28 runs taking over 50 minutes.

The number of hosts visited on each run was pretty consistent. With the exception of two low values that most likely resulted from network outages, the number of hosts visited ranged from 603 to 667, and on all but 6 of the 28 runs more than 630 hosts were visited. The results given below on the structure of the New Zealand NTP subnet are based on analysis of the largest log file, which was recorded between 2105 and 2152 on 1998-11-28.

Of the 667 "hosts" discovered in that run, three were actually local clocks, so 664 real NTP servers were discovered. These 664 servers can be split into 3 groups based on their location. Hostname lookup failed in 101 cases—we are unable to say whether these servers are in New Zealand or overseas (although the majority are probably in New Zealand). Of the remaining 563 hosts, 523 were located in New Zealand and 40 were overseas.

Table 1 shows for each group of NTP servers the number of servers at each stratum. Of the 17 overseas stratum 1 servers, 8 were in the United States, 2 were in Japan, 3 were in Australia, 3 were in Germany, and 1 was in the United Kingdom. The NTP servers whose stratum is unknown are those that would not respond to the `sysinfo` request (the server stratum is one of the fields returned by `sysinfo`).

Location	Servers	Servers at each stratum						Stratum unknown
		1	2	3	4	5	16	
Overseas	40	17	9	14	0	0	0	1
Unknown	98	0	3	4	1	1	0	89
New Zealand	526	0	12	183	153	8	1	169

Table 1: Breakdown of the number of NTP servers at each stratum for each of the three server groups.

Group	Servers	% servers at each stratum					
		1	2	3	4	5	6-14
Internet	13880	1.5%	32.0%	47.5%	16.2%	2.3%	0.4%
New Zealand	356	0.0%	3.4%	51.4%	43.0%	2.2%	0.0%

Table 2: Comparison of the proportions of NTP servers at each stratum as found by us within New Zealand, and Mills *et al.* in their Internet-wide study

It is interesting to compare the proportions of New Zealand NTP servers that belong to each stratum with those typical of the global Internet. Table 2 compares the percentages extracted from our data with percentages computed from a recent Internet-wide survey. The lack of public primary servers in New Zealand means that New Zealand servers tend to be concentrated at strata 3 and 4, whereas in the internet as a whole servers are concentrated more at strata 2 and 3. The average stratum for New Zealand servers is 3.4 as against 2.9 for the internet as a whole. With New Zealand having no stratum 1 servers, you might have thought the difference between the averages would be 1. The fact that the difference is 0.5 indicates that the New Zealand hierarchy is somewhat shallower than the internet-wide hierarchy, which is not surprising given that the New Zealand NTP subnet is much smaller.

Further analysis is focused on the 526 New Zealand hosts. In fact, most of the analysis is done for 355 hosts in New Zealand whose stratum could be determined (in 169 cases the stratum could not be determined), which were operating at a valid stratum (1 host was operating at stratum 16) and that returned a valid list of peers (1 at stratum 3 did not). We will call these datasets the "allNZ" dataset and the "respondNZ" dataset (note that the respondNZ dataset is a subset of the allNZ dataset).

The first analysis we did was to determine how evenly NTP usage was spread across the New Zealand second and third level domains. Table 3 summarises for each level 2 domain the number of NTP servers in that level 2 domain, and the number of different level 3 domains these servers belonged to. Information is given for both datasets.

The right hand column is present to allow comparisons between the distribution of NTP servers

Domain	# of hosts		# of level 3 domains		Percent IP domains
	allNZ	respondNZ	allNZ	respondNZ	
ac.nz	324	222	9	9	0.7%
co.nz	99	48	28	19	85.6%
net.nz	62	56	12	11	3.8%
govt.nz	22	18	7	4	1.4%
gen.nz	18	11	6	4	1.5%
school.nz	1	0	1	0	1.4%

Table 3: For each second level domain, this table shows the number of NTP servers present in the domain, and the number of of level 3 domains the servers are spread across.

Domain	NTP servers
clear.net.nz	10
netgate.net.nz	11
manawatu.gen.nz	12
clix.net.nz	15
comnet.co.nz	16
otago.ac.nz	23
ihug.co.nz	26
massey.ac.nz	34
auckland.ac.nz	58
waikato.ac.nz	87
canterbury.ac.nz	94

Table 4: The level 3 domains in which at least 10 NTP servers were visited (allNZ dataset).

across the various level 2 domains with the sizes of the various level 2 domains. The percentages in the right hand column[2] show the percentage of the 17460 IP-connected organisations in New Zealand that belong to each level 2 domain. It would have been better to know the percentage of the 174201 New Zealand internet addresses that belong to each group, but this information was not available. The percentages add up to 99.8% (none of the NTP servers visited were classified as being in the `cri`, `iwi` or `mil` level 2 domains).

Our method for associating each host with a level 3 (and therefore a level 2) domain needs some explanation. It was assumed that all hosts in the same (class B or C) network were in the same level 3 domain. So for each class B and C network number, all hosts in that network were counted as being in the same level 3 domain. However, there were 6 networks that contained hosts that belonged to two or more level 3 domains. For example, 16 servers were found in network 131.203. Of these, 1 was in the `irl.cri.nz` domain, 3 were in `grace.cri.nz`, 1 was in `huttcity.govt.nz` and 11 were in `comnet.co.nz`. In the figures presented in Table 3, all 16 121.203 hosts were counted as belonging to `govt.nz`. If the second column was recalculated so that it was based on host names, the only significant changes would be to increase `.co.nz` to 108, reduce `govt.nz` to 7 and to add a count of 4 for `cri.nz`.

From the table it is clear that despite the fact the `ac` (academic) domain accounts for a very small proportion of IP connected organisations, it nevertheless contains a significant majority of the NTP servers visited. The average number of NTP servers at each site is high, because the 324 servers are spread across only 9 level 3 domains. It is also clear that given its size, relatively few NTP servers were visited in the `co` (commercial) domain. Finally, the `net` domain contains quite a few NTP servers in relation to its size, which indicates that ISPs and operators of backbone links have a considerable interest in synchronising the clocks of their computers.

Table 4 lists the top 10 individual level 3 domains in terms of the number of NTP servers found in the domain in the allNZ dataset.

Having examined the distribution of NTP servers over second and third level domains, we now describe the structure (that is, the links between servers) of the New Zealand NTP subnet. Table 5 gives a per-stratum summary of the number of peers each NTP server synchronises to. A general trend apparent from the table is that as the stratum level decreases, the average number of peers synchronised to increases. This is not surprising. Each low stratum NTP server generally provides time service for a number of higher stratum servers. Consequently, low stratum NTP servers generally synchronise to at least 3 peers so that good accuracy is maintained. High stratum NTP servers have few, if any, clients. Accuracy and reliability of the time service is a lower priority for such NTP servers, and so they generally synchronise to fewer peers.

Table 6 gives a per-stratum summary of the number of clients that synchronise to each NTP

[2]These percentages were computed from the 1998-12-01 figures presented in `http://webpages.netlink.co.nz/~mark/netsites-growth.html`.

Stratum (S)	In column N are the number of servers at each stratum synchronising to N peers															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
2			4	3	2	1	1									1
3	18	58	39	17	6	14	3		3	9	12					3
4	112	11	24	4						1		1				
5	8															

Table 5: Per-stratum counts of the number of peers each NTP server synchronises to (respondNZ dataset).

S	In column N are the number of NTP servers at each stratum with N clients.																			
	0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19
2	5	1	2								1									
3	97	12	16	9	7	6	12	6	1	1	2		1		1	3	3		1	2
4	145	3			1	1	2													

Table 6: Per-stratum counts of the number of clients that synchronise to each NTP server (respondNZ dataset). Omitted are three stratum 2 counts (62, 123, 171) and two stratum 3 counts (21, 91).

server. Taking into account the additional information in the caption, the general trend is that as the stratum level decreases, the average number of clients per NTP server increases. Again, this is as expected. There are relatively few low stratum NTP servers, and they commonly pass on timing information to many clients. At higher strata, many NTP servers are leaves of the hierarchy and have no clients.

The following sub-sections provide a more detailed analysis of the information presented in Tables 5 and 6, as well as providing new information.

4.1 Top level servers

This section discusses the highest level (lowest stratum) servers in the New Zealand NTP subnet. In the `ntpcrawl` run we are reporting, a total of 12 New Zealand stratum 2 servers were detected.

Even though it did not respond in this particular run of `ntpcrawl`, there is actually a private stratum 1 NTP server in New Zealand located at the University of Otago (this server has responded on other occasions). The presence of the stratum 1 server is reflected in our data. Three of the stratum 2 servers detected are located in Otago, and are clients of the Otago stratum 1 server. None of these three stratum 2 servers has any clients. This means that the stratum 1 Otago server and its small number of clients is a self-contained hierarchy that does not contribute any timing information to the wider New Zealand NTP subnet. Two other stratum 2 servers (`gateb.opus.co.nz` and `aitken.scitec.auckland.ac.nz`) also have no clients.

Of the remaining seven stratum 2 servers, four have a small number of clients. A stratum 2 server at Victoria University with an unknown host name (130.195.12.1) has a single client. `wiatemata.ait.ac.nz` has two other AIT machines as clients. `cisco2.atrix.gen.nz` has two clients at Saturn. `tui.grace.cri.nz` has 10 immediate clients, and 10 further clients below them. Because each NTP server can synchronises to multiple servers, many of the servers that synchronise (directly or indirectly) to one of the non-Waikato stratum 2 servers also synchronises to one or more of the three Waikato stratum 2 servers.

Not surprisingly, the three public secondary servers at Waikato have much larger numbers of clients than any of the other stratum 2 servers. Table 7 summarises for each of Waikato's stratum 2 servers the number of clients at each stratum. The dominance of Waikato's stratum 2 servers is clear. The machine with the smallest number of clients (`johnwayne`) has more clients than the other 9 New Zealand stratum 2 servers put together.

Server	Clients at stratum		
	2	3	4
ankh	7	152	2
ramses	3	117	2
johnwayne	3	58	1

Table 7: Client counts for the three public secondary servers at Waikato.

Most clients of the Waikato secondary servers are at stratum 3 (they are synchronised to a stratum 2 server). Some clients are at stratum 2 (for example, each Waikato machine is synchronised to the other two). Links between servers at the same stratum enable backup synchronisation paths. In a couple of cases clients are at stratum 4. This is an unusual situation in which the client has decided that some other server at stratum 3 is a better synchronisation source. In one case, `ns1.manawatu.gen.nz` is a client of all three Waikato stratum 3 servers, but has decided that `mailhost.manawatu.gen.nz` is a better synchronisation source, even though it in turn has `ramses` as its system peer.

Another way of assessing the importance of the various NTP servers is to analyse the tree that results when you consider the links between each NTP server and its system peer. Although an NTP server often computes corrections to its local clocks by combining timing information from several peers, the system peer is the one deemed to be most influential (for example, the stratum of an NTP server is 1 greater than that of its system peer).

The system peer of `ankh` was `time-a.nist.gov`. A total of 239 NTP servers, at strata 3, 4 and 5, are directly below `ankh` in the tree. The system peer of `ramses` was `time-b.nist.gov`, and 23 stratum 3 and 4 NTP servers were below `ramses`. Finally, the system peer of `johnwayne` was `tictoc.tip.CSIRO.AU`, and 60 stratum 3 and 4 NTP servers were below `johnwayne`. It is interesting that even though `johnwayne` had many fewer clients than `ramses`, it had many more clients for which it was system peer. Overall then, 322 NTP servers had (directly or indirectly) one of the Waikato secondary servers as their system peer.

Only two other stratum 2 New Zealand servers had one or more clients for which they were system peers (the other seven either had no clients or had clients whose system peers were other NTP servers). The system peer of `tui.grace.cri.nz` was `tictoc.tip.CSIRO.AU`, and 14 stratum 3, 4 and 5 NTP servers were below `tui`. Also, `waitemata.ait.ac.nz` had one NTP server beneath it.

Finally, a small number of machines at stratum 3 had system peers (of stratum 2) outside New Zealand. Five `clix.net.nz` servers (which from their names were actually located in Los Angeles) had between them as their system peers two stratum 2 servers at the University of Delaware. `puka2.ait.ac.nz` had as its system peer `sun2.bnl.gov` (which is interesting, because `puka` was also a client of a much closer stratum 2 server—`waitemata.ac.nz`). `mx4.ix.net.nz` had as its system peer `eagle.tamu.edu`. Two stratum three `ix.net.nz` servers had `DNS.CIT.CORNELL.EDU` as their system peer, and three stratum four servers were in turn below them.

Adding up all NTP servers encountered that were not below one of the Waikato secondary servers gives a total of 38, compared to the 325 you get for the hierarchies of the Waikato secondary servers.

4.2 Servers that synchronise to large numbers of peers

Each of the three Waikato secondary servers peered with the other two. In addition, `ankh` and `ramses` were clients of three overseas servers and `johnwayne` was the client of five overseas servers. However, it is clear from Table 5 that some NTP servers had many more peers (between 9 and 16) than the three main New Zealand servers. This section looks at the 30 NTP servers that synchronised to 9 or more peers.

A stratum 2 server at Victoria University (somewhat heroically) synchronised to 16 peers. Ten

were overseas (nine of which were at stratum 1), four were stratum 3 servers at Victoria (which goes against advice not to synchronise to peers at a higher stratum), and the others were `ankh` and `ramses`. It is not clear why this machine had so many peers. It is not as if the machine had a large number of clients. In fact it had only one client (`sinbad.regy.vuw.ac.nz`). Ironically, `sinbad` is also a client of `ankh`, and `ankh` was the system peer of `sinbad` when the monitoring run occurred.

Many of the NTP servers with large numbers of peers were operated by Clear Communications, and belonged to the level 3 domains `clear.net.nz` and `clix.net.nz`. These included 3 servers with 16 peers, 6 with 11 peers and 8 with 10 peers. All 17 servers peered with `ankh` and with between 8 and 10 other clix/clear servers. In addition, three of these servers (all with "LosAngeles" in their host names) peered with the same 7 NTP servers in the United States. Evidently Clear were keen to maintain time services for their backbone machines in the face of many faults. Paths to stratum 1 servers were via `ankh` and via the three "LosAngeles" hosts with links to United States. Interestingly, a fourth "LosAngeles" host had only one peer (`ankh`).

Another group of backbone hosts with large numbers of peers was present in the `ix.net.nz` domain. Two had 10 peers and three had 9 peers. These NTP servers each synchronised to 3 or 4 other `ix.net.nz` servers, 4 overseas servers and `ankh` and `ramses`.

The other group of servers with many peers were to be found in the `manawatu.gen.nz` domain (six had 11 peers, one had 12). All synchronised to 7 or 8 other `manawatu.gen.nz` hosts and the three Waikato secondary servers.

4.3 More detailed observations on structure

The NTP documentation contains guidelines about how NTP should be established within an administrative domain. Some are worth paraphrasing here to allow comparison with the structures within New Zealand.

> In the case of a gateway or file server providing service to a significant number of workstations or file servers in an enterprise network it is even more important to provide multiple, redundant sources of synchronization and multiple, diversity-routed, network access paths. The preferred configuration is at least three administratively coordinated time servers providing service throughout the administrative domain including campus networks and subnetworks. Each of these should obtain service from at least two different outside sources of synchronization, preferably via different gateways and access paths. These sources should all operate at the same stratum level, which is one less than the stratum level to be used by the local time servers themselves. In addition, each of these time servers should peer with all of the other time servers in the local administrative domain at the stratum level used by the local time servers, as well as at least one (different) outside source at this level. This configuration results in the use of six outside sources at a lower stratum level (toward the primary source of synchronization, usually a radio clock), plus three outside sources at the same stratum level, for a total of nine outside sources of synchronization.
>
> When planning your network you might, beyond this, keep in mind a few generic don'ts, in particular:
>
> 1. Don't synchronize a local time server to another peer at the same stratum, unless the latter is receiving time from lower stratum sources the former doesn't talk to directly. This minimizes the occurrence of common points of failure, but does not eliminate them in cases where the usual chain of associations to the primary sources of synchronization are disrupted due to failures.
> 2. Don't synchronize more than one time server in a particular administrative domain to the same time server outside that domain. Such a practice invites common points of failure, as well as raises the possibility of massive abuse, should the configuration file be automatically distributed to a large number of clients.

Waikato had followed these guidelines reasonably closely in configuring their three secondary servers. Each server peered with the other two. Also, each server peered with between 3 and 5 servers overseas. There was some overlap between the three groups of overseas servers, although between them the three Waikato secondary servers synchronised to 7 distinct overseas servers. All of the overseas servers that `ankh` and `ramses` synchronised to were in the United States. `johnwayne` synchronised to three primary servers in Australia as well as to two primary servers in the United States. It could be argued that all three should be synchronised to at least one Australian primary server and one American primary server so that all three secondary servers will remain in contact with a primary server even if all direct links to the United States or Australia are cut.

The other NTP servers at Waikato can be sub-divided into three groups. One group consisted of stratum 3 servers that synchronised to just one (one server), two (13 servers) or three (18 servers) of the Waikato secondary servers. A second group consisted of seven NTP servers in the Physics department. All were clients of the Waikato secondary servers (three were clients of two secondary servers, the other four were clients of all three secondary servers). In addition, each of the seven Physics servers was a client of between two and five of the other six Physics servers. If the NTP configuration advice was followed, all seven Physics servers would synchronise to the three Waikato secondary servers, with no links between the Physics servers.

A third group is the servers within the Computer Science department at Waikato. Four servers were at stratum 3. Three were clients of two of the secondary servers, the fourth was a client of all three. Also, each was a client of two or three of the other stratum 3 servers. All four were also clients of `lucy`, a stratum four server, and `lucy` was a client of all four. `lucy` had a single stratum 5 client (`mallowpuff`), which synchronised only to `lucy`. `xena`, one of the stratum 3 servers, had three stratum 4 clients (`sql`, `cdwriter` and `smilla`), with all three synchronising only to `xena`. `sql` had one stratum 5 client, and `smilla` had four, with each of the stratum 5 clients synchronising to a single server. This hierarchy is somewhat idiosyncratic. Nine servers synchronised to a single server, making the subnet fragile. Also, the hierarchy had greater depth than justified by the number of machines involved (six of the eight stratum 5 NTP servers discovered were in the Waikato Computer Science department). Errors accumulate at each level of the NTP hierarchy. Again, there were links between servers at the same stratum in cases where the two servers have the same set of upstream links. There were also cases of a stratum 3 server synchronising to a stratum 4 server.

Looking at NTP subnet organisation within each level 3 DNS domain, several patterns can be observed. Where the NTP subnet within a domain was small, it was common for all of the NTP servers within the domain to synchronise to between one and three of the Waikato secondary servers. This was the case for 18 domains, in which a total of 7 servers synchronised to all three secondary servers, 19 synchronised to two of the three, and 14 synchronised to one. The level 3 domain with the most hosts contained 10 machines, and in 12 cases the domain contained a single machine.

Where a domain contained more than a few machines, the NTP subnet was usually organised in a hierarchical fashion. One pattern observed was where there was a single stratum 3 NTP server that synchronised to a sub-set of the Waikato secondary servers, with all other servers in the domain at stratum 4 and synchronising only to the local stratum 3 server. This was the case for `cosc.canterbury.ac.nz` (stratum 3 server synchronised to two secondary servers, with 91 clients), `massey.ac.nz` (synchronised to two secondary servers; had 3 clients) and `citylink.co.nz` (synchronised to three secondary servers; had 3 clients). These subnets are vulnerable to a failure of the single stratum 3 server.

A more robust version of this pattern is where a domain has a two or more stratum 3 servers (that synchronise to some sub-set of the Waikato stratum 3 servers), and a greater number of stratum 4 servers that synchronise to a sub-set of the local stratum 3 servers. This was the case for:

- `topnz.ac.nz`—two stratum 3 servers; two stratum 4 servers synchronised to both stratum 3 servers.

- `acneilsen.co.nz`—two stratum 3 servers; four stratum 4 servers synchronised to both stratum 3 servers.

- `otago.ac.nz`—three stratum 3 servers; 13 stratum 4 servers synchronised to all three stratum 3 servers In addition, three stratum 2 servers were encountered that synchronised to Otago's stratum 1 server as well as to the three local stratum 3 servers.

- `auckland.ac.nz`—six main stratum 3 servers (two had 9 clients, three had 7, one had 4) and two minor stratum 3 servers (one had 1 client, one had 0); 17 stratum 4 servers (12 synchronised to three stratum 3 servers, one synchronised to 2 servers, 4 synchronised to 1 server).

  Auckland University had some NTP servers outside this main hierarchy. One was at stratum 2. It synchronised to a stratum 1 server in Australia, `ankh` and a local stratum 3 server (having a server synchronise to a higher stratum server is strange). This stratum 2 server was an orphan—it had no clients.

  There was a small subnet within the Computer Science department. Two stratum three servers (`kakapo` and `cs20`) synchronised to all three Waikato secondary servers and to each other. A stratum 4 server (`data`) synchronised to `kakapo`, and a stratum 5 server (`mail`) synchronised to `data`. Again, a hierarchy this deep is not needed.

Of the 15 stratum 3 servers listed above, 2 synchronised to one secondary server, 5 to two secondary servers and 8 to three secondary servers.

There was an interesting variation on this theme for `plain.co.nz`, which had three stratum three servers. Each server synchronised to two Waikato secondary servers, with each synchronising to a different pair. Each server also synchronised to the other two.

The CRI and comnet NTP servers were arranged in something close to a hierarchy. The main servers were `tui.grace.cri.nz` and `selket.grace.irl.cri.nz`. `tui` is a stratum 2 server that synchronised to `tictoc.tip.CSIRO.AU` (stratum 1), `ankh` and a stratum 3 comnet server (another case of a higher stratum peer). `selket` was a stratum 3 server that synchronised only to `tui`. The remaining servers synchronised to various groups of peers. Seven synchronised to both `tui` and `selket`. All chose `tui` as their system peer (and were therefore at stratum 3)—synchronising to `selket` is redundant as the only server `selket` synchronises to is `tui`. Five servers synchronised to `ramses` and `selket`. All chose `ramses` as their system peer and so are at stratum 2. Three servers synchronised only to `selket`, so were at stratum 4. One of these three (`brahms`) had a stratum 5 client (`invermay`). One server synchronised to `tui` and a stratum 3 comnet server, and one server synchronised to `tui`, `selket` and `ramses`.

Having one main server synchronise exclusively to another of the main servers does not make for a robust design. This particular group of hosts would be better restructured as a two level hierarchy, with the top level consisting of two or three hosts at stratum 3 (or 2) and the bottom level hosts at stratum 4 (or 3). Another benefit of this approach is to eliminate the stratum 5 server (again, the number of servers involved does not justify a server operating at stratum 5).

Another pattern is for most or all hosts in a domain to be at stratum 3, and for there to be many links between them (even though these links often do not give access to any new upstream NTP servers). Examples of this pattern include:

- NTP servers on 24 hosts were encountered in the `clix.net.nz` and `clear.net.nz` domains. One of these was simply a client of `ankh`, and will not be discussed further. The remaining servers can be divided into two groups. Four of the 23 had "LosAngeles" as part of their host name (a fifth failed to respond to `xntpdc` queries). Three of the four LosAngeles machines were synchronised to 7 stratum 2 and 3 servers in the USA, as well as `ankh` and 8 `clear` and `clix` servers. The fourth LosAngeles server, `b1-fa0-0-0.losangeles.clix.net.nz` was synchronised to just one NTP server—`ankh`.

  The second group (of 19 servers) synchronised to at least 6 peers (five synchronised to 11 peers, seven synchronised to 10 peers). All 19 synchronised to `ankh` and two LosAngeles

servers (`b1-fa0-0-0` and `b2-fa-0-0-0`). The other three to eight peers synchronised to are drawn from the group of 19. Each of the 19 had between 0 and 18 of the other members of the group as clients.

This was one of the few large New Zealand subnets that had some synchronisation paths that did not pass through the Waikato secondary servers. Three hosts in the subnet each synchronised to seven NTP servers in the USA. Also, all hosts except `b1-fa0-0-0.losangeles` synchronised to `b2-fa-0-0-0`, one of the three servers that synchronised to the seven USA servers.

Even so, there are some unusual aspects to the subnet. The large number of links between stratum 3 servers runs counter to the NTP configuration advice. Also, the fact that 21 servers synchronised to `b1-fa0-0-0` is strange, as it relies entirely on `ankh` for its timing information. This is odd for two reasons. First, `ankh` is a stratum 2 server that was getting its timing information via long-distance links to the United States. The timing errors that occur on such links are then compounded by `b1-fa0-0-0` using such a link to get timing information off ankh, and compounded again for each New Zealand server that synchronises to `b1-fa0-0-0`. Second, there were two `clix` machines in Los Angeles that had 7 US peers that could have been chosen instead of `b1-fa0-0-0`.

- Eight NTP servers in the `manawatu.gen.nz` and `manawatu.net.nz` domains communicate extensively with each other. All synchronised to the three Waikato secondary servers, and to eight or nine of the others (except one that synchronised to only one of the others). All had 6 or 7 of the others as clients. In this case all eight servers fall back on their local clocks as a synchronisation source if connections to Waikato fail, which might explain why there is so much synchronisation between stratum 3 servers.

- The `ix.net.nz` subnet followed this pattern to some extent, and the hierarchical pattern to some extent. Five servers at stratum 3 all synchronised to `ramses`, `ankh` and four overseas NTP servers (three in the United States, one in Australia). Each of the five stratum 3 servers synchronised to three of the others. There are two stratum 4 servers that synchronised to the stratum 3 servers (one synchronised to three stratum 3 servers, the other to four).

4.4 Summary

At a micro level, several opportunities for improving the subnet structure have been identified. Many are related to ensuring that each client synchronises to more than one peer, and avoiding synchronising to excessive numbers of peers at the same stratum.

At a macro level, the main structural shortcomings of the New Zealand NTP subnet are the lack of any public primary servers, and the concentration of all public secondary servers at a single site.

The fact that the three Waikato secondary servers had many more clients than other servers within New Zealand (`ankh` has 161 clients) is not of itself a problem. Mills *et al.* encountered many servers with more than 200 clients, including a stratum 1 server with 652 clients [3]. This indicates that `ankh` was not dangerously overloaded, even though it had the most clients of any New Zealand NTP server.

5 Accuracy measures for the New Zealand NTP subnet

In addition to collecting structural information, `ntpcrawl` also collected various pieces of data for each NTP server that indicate how well clocks are being synchronised. By analysing this data we can determine whether the structural shortcomings of the New Zealand NTP subnet (such as lack of a stratum 1 server and reliance on three secondary servers located on the same site) have an obvious adverse effect on the accuracy of clock synchronisation achieved by NTP.

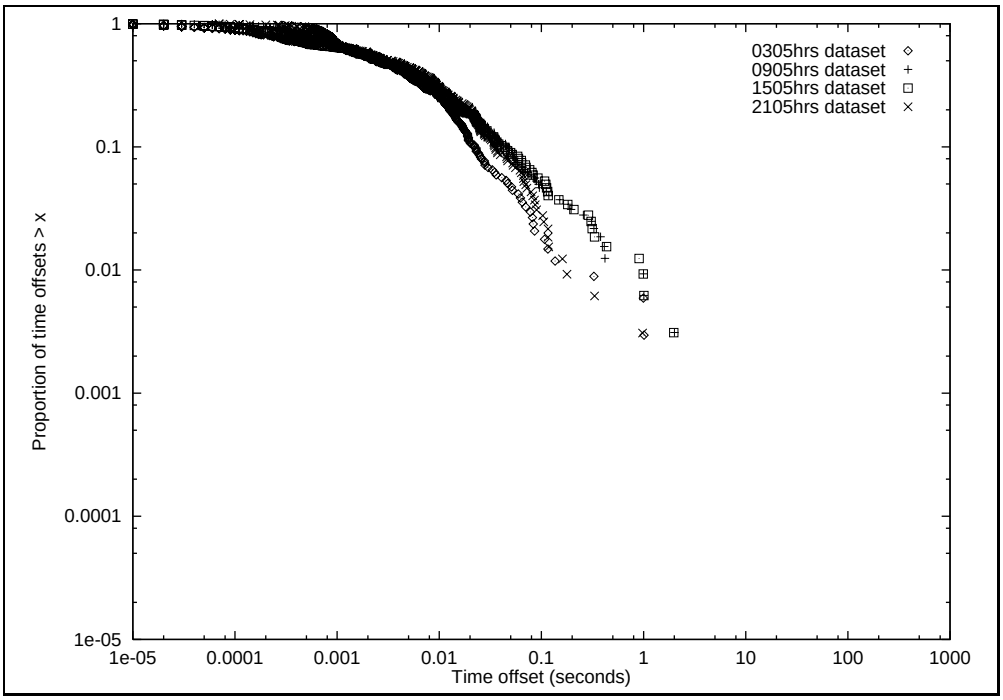

Figure 1: Cumulative distribution function of the estimated time offset with respect to the system peer for the four data sets collected on 1998-11-24.

Results of a recent large-scale study of NTP performance around the globe are described by Mills *et al.*[3]. In their study, data was gathered from 13,880 NTP servers. The main performance indicators presented is time offset, which is NTP's estimate of the current offset between the local clock and that of the system peer. Three other performance indicators are presented as well, including round trip delay with respect to the system peer, and the estimated frequency offset between the local clock and UTC.

Mills *et al.* present their results as cumulative density functions. Figure 1 shows time offsets we measured presented as CDFs. For a particular time offset t, the corresponding Y value is the proportion of the time offsets recorded whose value is greater than t. Where t is very small, nearly all time offsets recorded are greater than t, so the Y value plotted is close to 1. As t grows, so the proportion of the values recorded that are greater than t shrinks.

For each of our seven day's worth of data, we produced four cumulative density function plots to compare with the four presented in [3]. In all 28 cases (four plots for each of seven delays), our curve was below that of Mills *et al.* This means that the time offsets (and other indicators) for something over 300 New Zealand NTP servers were (overall) lower than those reported by Mills *et al.* for nearly 14,000 NTP servers. Lower time offsets, indicate better synchronisation.

In this paper, we have included three cumulative density function plots as examples. All show results from the 4 datasets recorded on 1998-11-24. In all three figures results are shown for hosts that meet the criteria given above for the "respondNZ" dataset.

Figure 1 shows the cumulative density functions for the estimated offset between the NTP server's local clock and that of its system peer. The 50th percentiles for the four datasets range from 2.6 milliseconds to 3.1ms and the 95th percentiles from 46.8ms to 110.35ms. For probabilities less then 0.1 the points for the 0305 dataset are below the points for the other three datasets. This occurs because network loadings tend to be lower in the early hours of the morning, resulting in lower message delays and better clock synchronisation. This effect was observed in most of our cumulative density function plots.

Figure 2 shows the cumulative density functions for the round trip message delay between the

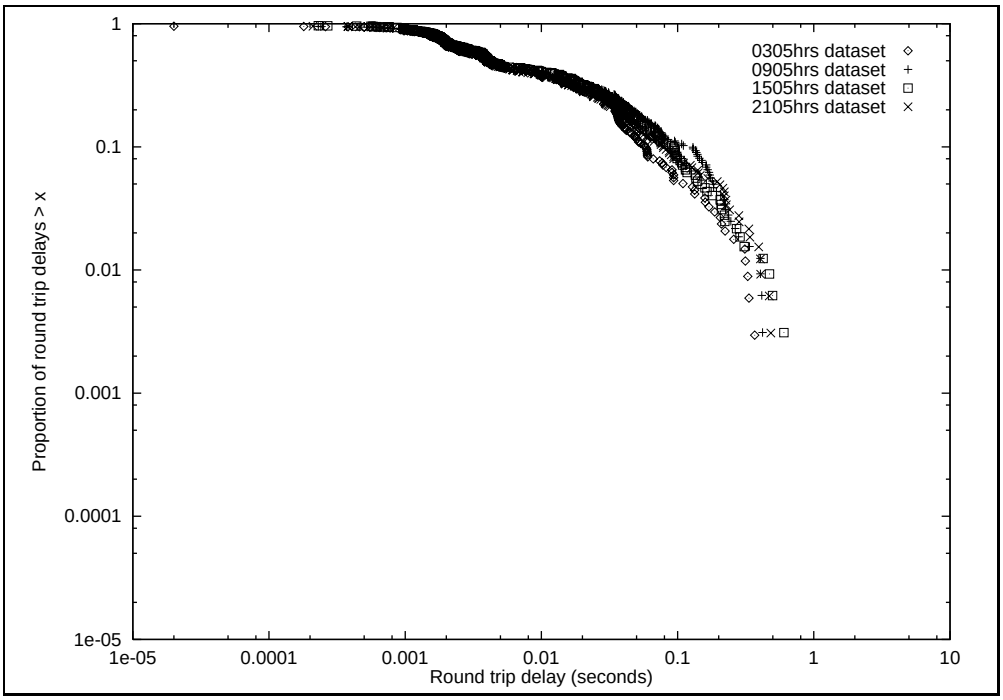

Figure 2: Cumulative distribution function of the round trip delay measured with respect to the system peer for the four data sets collected on 1998-11-24.

NTP server and its system peer. The 50th percentiles for the four datasets are 4.0ms in all cases and the 95th percentiles range from 109.5ms to 204.5ms.

Figure 3 shows the cumulative density functions for root dispersion. Dispersion is an estimate of the expected maximum error between an NTP server's clock and some other clock. Root dispersion is an estimate of the expected maximum error between the NTP server's clock and the stratum 0 clock it is ultimately synchronised to. In other words, the expected maximum error with respect to UTC. The 50th percentiles for the four datasets range from 63.2ms to 78.5ms and the 95th percentiles from 167.4ms to 683.6ms. Maximum dispersions in the early morning are clearly lower than at other times of the day.

Another observation based on the data collected relates to NTP's estimate of systemic frequency offset. A clock's systematic frequency offset is the difference between the rate of the local clock and that of UTC. Frequency offset is often expressed in units of microseconds per second. The frequency offset of the clock in a typical workstation usually lies between -100 and 100 microseconds per second. For all seven days the frequency offset cumulative density functions of all four datasets were very close to each other. This indicates that frequency offset measures are less susceptible to long message delays than clock offset estimates.

Overall, there was nothing in the data collected to indicate major persistent problems with the accuracy of NTP within New Zealand. The performance indicators reported in [3] are primarily aimed at assessing how well each NTP server synchronises to its direct peers (time offset with respect to the system peer is the main performance indicator used). No results are given that indicate how closely NTP servers are synchronised to UTC. Given that New Zealand has no public primary servers, it would have been interesting to compare our root dispersion measurements with those of Mills *et al.*, but system dispersion is not reported in [3].

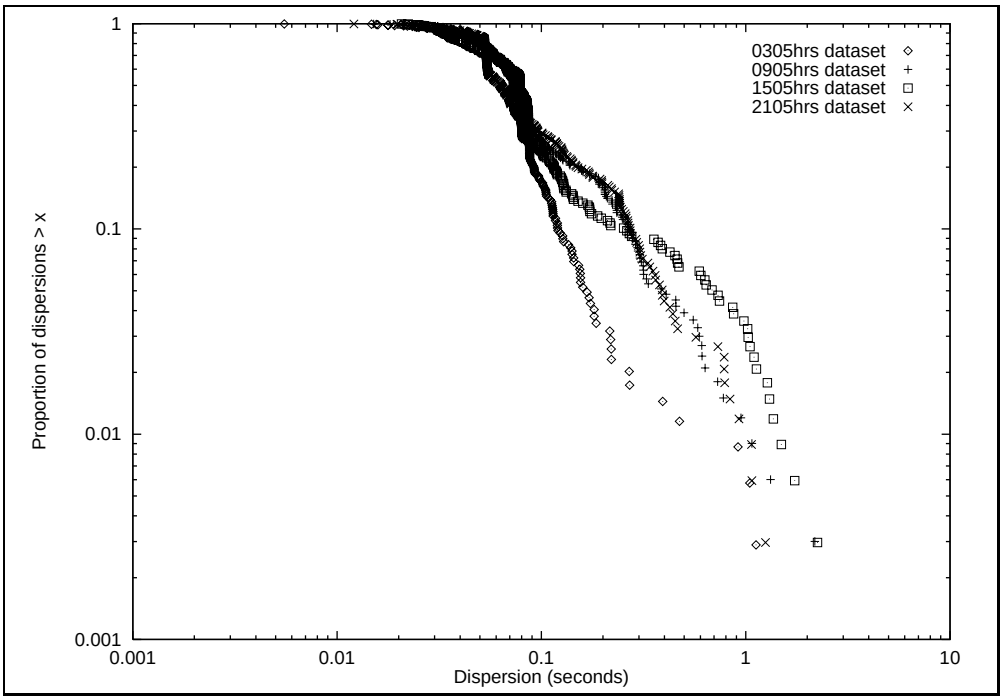

Figure 3: Cumulative distribution function of the system dispersion for the four data sets collected on 1998-11-24.

6 Reasons for the clock steps

Because no obvious major problems were found with the overall accuracy of NTP within New Zealand, an investigation of NTP logs was carried out to determine the unusual conditions that caused the clock steps recorded on `kaka`. Mills comments that "Corrections of this magnitude [over 128ms] are exceedingly rare, usually as the result of reboot, broken hardware or missed leap-second events." [5].

Our stratum 3 server, `kaka`, synchronises to two of the Waikato secondary servers (`ankh` and `ramses`) and so is configured in a manner common to many other New Zealand NTP servers. In a 31 day period covering late July and most of August, 1998, a total of 36 clock steps occurred on `kaka`. The availability of NTP logs covering that period allowed investigation into the causes of the steps. In some cases the information in the logs was inconclusive, but in others the causes of the steps were reasonably clear.

In 18 cases (nine pairs of steps) the steps resulted from large errors in offset estimates caused by long round trip delays. In each case a message exchange with one server with a round trip delay of between 300ms and 900ms (the minimum round trip delay is around 34ms) was immediately followed by a similar message exchange with the other server. If both message exchanges produce highly inaccurate clock offset estimates (which is possible if the round trip delay is high, but is not necessarily the case) then the NTP server falsely concludes that its own clock is in error by more than 128ms and steps the clock.

In nine cases we observed clock steps caused by the long round trip delays in the most recent message exchange with each server. In all nine cases, the round trip delays of the following few message exchanges were much lower. The server quickly realised that it had made a mistake in stepping the clock in the first place, and after about five minutes steps the local clock in the opposite direction by about the same amount.

One thing about this behaviour puzzled us. For each peer, NTP buffers details of the 8 most recent message exchanges so that it if an occasional message exchange with a long round trip

delay occurs NTP can choose to use one of the other 7 recent message exchanges in calculating the current time offset. In the cases we observed, there were always message exchanges within the last 8 that had low message delays. Clearly, the NTP server was ignoring these for some reason, otherwise the step would not have occurred.

To find out why NTP seemed to be overlooking recent message exchanges, the source code (version 3) was consulted. When looking through its buffer of 8 recent message exchanges, NTP ignores those that occurred over 800 seconds ago. When the 9 steps occurred, NTP was polling each server at intervals of 1024 seconds, which meant that all but the most recent message exchange were ignored. NTP version 4 has a different way of determining how many of the buffered message exchanges to consider, which may turn out to be better in practise than the version 3 method.

Other steps seem to have been caused by the clocks of `ankh` and `ramses` differing by a large amount (over 200ms), and `kaka` deciding to change its system peer from one to the other. This happened on a number of occasions on August 6th. During the period described, round trip delays were low (all were under 100ms) so the estimates of time offset with respect to each server were accurate to within 30ms or less. Initially `kaka` was synchronised to `ramses`. `kaka` then came up with time offsets of -29ms for `ramses` and 219ms for `ankh`, and decided to switch to `ankh` as its system peer (resulting in a 219ms clock step). Five minutes later, `kaka` came up with time offsets of 38ms for `ankh` and -256ms for `ramses`, and decided to switch to `ramses` as its system peer (the clock was stepped by -256ms). Five minutes later again, `kaka` came up with time offsets of 6ms for `ramses` and 273ms for `ankh`, and decided to switch to `ankh` as its system peer (the clock was stepped by 273ms).

The behaviour just described indicates that the clocks of two of the Waikato secondary servers differed by around a quarter of a second—a very large difference indeed. Just before this sequence of "clock hops" occurred, there was evidence in the logs that `ramses` clock changed rapidly (quite possibly it was stepped). In three successive message exchanges (at intervals of 1024 seconds), all with round trip delays of less than 60ms, the following offsets with respect to `ramses` were estimated: 12ms, -87ms, -232ms.

Later, `xntpd` log records from `ankh` and `ramses` covering a three day period in February 1999 were compared with log records from `kaka`. Several steps occurred on all three machines. In five cases, a step on `ankh` or `ramses` was followed within five to ten minutes by a step on `kaka` of about the same size. This confirms that some steps made on `kaka` are caused by steps in the clocks of the Waikato secondary servers (our July/August data indicated that such steps were occurring, but we were unable to access the Waikato logs to confirm this). One particularly large pair of steps was observed for `ramses`. Its clock was stepped by 588ms, then five minutes later stepped by -659ms.

7 Conclusions

We have described experiments performed to investigate the structure of the NTP subnet in New Zealand, and to record various performance indicators for the NTP servers encountered. We have also investigated the reasons behind some large time steps that occurred on our department's main NTP server.

Investigations into the structure of the New Zealand NTP subnet found that NTP was in use by many organisations. The structure of the NTP subnet is largely determined by the independent decisions of a large group of system administrators, without a great deal of coordination occurring. This is reflected in the many connection patterns we encountered. The major structural shortcomings of the New Zealand NTP subnet is that it lacks public primary servers, and that there are no public secondary servers outside the three at Waikato. We came across a number of cases in which it appeared that organisations could improve their NTP configurations (in terms of the their internal NTP server hierarchy and external peers). Scope exists for writing a tool to analyse an organisation's NTP subnet and reporting on potential configuration problems.

Comparing various performance indicators taken from the New Zealand NTP subnet with those reported in an Internet-wide NTP study does not indicate the presence of any major systemic

performance problems within the New Zealand NTP subnet. Investigation of clock steps that occurred on `kaka` shows that the structural problems identified above were a factor in at least some of them. Some clock steps were caused by long round trip delays for packet exchanges with all stratum 2 servers. Having all of New Zealand's public secondary servers at a single site increases the chances that long delays will be experienced with respect to all secondary servers synchronised to. This was not helped by the fact that the NTP version 3 filtering algorithm only considers results from the most recent message exchange when the poll interval is 1024 seconds (such a poll interval is common).

On other occasions a step of the clock of a Waikato secondary server lead to a step of kaka's clock. Clock steps on Waikato secondary servers most likely are a result of the fact that they have to rely on long and frequently congested international links to get access to timing information from primary servers. Having primary servers within New Zealand would help to avoid this.

In summary, the lack of stratum 1 servers in New Zealand and the concentration of all stratum 2 servers at a single site can be identified as potential problems from a knowledge guidelines on NTP configuration. By investigating reasons for steps in the clock of a local stratum 3 servers, we have confirmed that the New Zealand NTP subnet configuration does indeed lead to relatively poor synchronisation from time to time. Improvements to New Zealand's NTP subnet will be of benefit to anyone interested in sub-second clock synchronisation.

A configuration that should lead to much more accurate clock synchronisation within New Zealand would involve establishing two or three public primary NTP servers at different sites in New Zealand, supported by six or so public secondary servers distributed throughout the country.

References

[1] Paul Ashton and Jonathan Mackenzie. Initial experiences with a clock synchronisation test bed. In *Computer Science '99—Proceedings of the 22nd Australasian Computer Science Conference, ACSC '99*, pages 98–109, Auckland, January 1999.

[2] D. L. Mills. Improved algorithms for synchronizing network clocks. *IEEE/ACM Transactions on Networks*, pages 245–254, June 1995.

[3] D. L. Mills, A. Thyagarajan, and B. C. Huffman. Internet timekeeping around the globe. In *Proc. Precision Time and Time Interval (PTTI) Applications and Planning Meeting*, pages 365–371. Long Beach, Ca, December 1997.

[4] David L. Mills. Precision synchronization of computer network clocks. *ACM Computer Communications Review*, 24(2):28–43, April 1994.

[5] David L. Mills. Clock discipline algorithms for the network time protocol version 4. Technical Report 97-3-3, Electrical Engineering Department, University of Delaware, March, 1997.