Characterizing Quality of Time and Topology in a Time Synchronization Network

Cristina D. Murta, Pedro R. Torres-Jr.
Department of Computer Science
Federal University of Paraná
81531-990, Curitiba, Brazil

Prasant Mohapatra
Department of Computer Science
University of California Davis
95616, Davis, CA

***Abstract*— As Internet computing gains speed, complexity and becomes ubiquitous, the need for precise and accurate time synchronization increases. In this paper, we present a characterization of a clock synchronization network managed by Network Time Protocol (NTP), composed by thousands of nodes, including hundreds of Stratum 1 servers, based on data collected recently by a robot. NTP is the most common protocol for time synchronization in the Internet. Many aspects that define the quality of timekeeping are analyzed, as well as topological characteristics of the network. The results are compared to previous characterizations of the NTP network, showing the evolution of computer clock synchronization in the last fifteen years.**

I. Introduction

Clock synchronization in computer networks and distributed systems may be not only a desirable characteristic in those systems but also a request for the correct operation of some applications and services. As network computing gains speed, complexity and becomes ubiquitous, the need for precise and accurate time synchronization increases, as well as the challenges involved in providing accurate time to systems and applications.

There are two main purposes of time synchronization. The first purpose is to ensure that events occur on time, in the correct sequence. Therefore, synchronization is necessary to start scheduled events and to register their occurrence. Many activities in commerce, banking, financial, business, transport, medicine, services, as a few examples, may need to guarantee that tasks are timely scheduled, and concurrent and cooperating processes interoperate correctly.

The second purpose is tracing, that is, retrieving information concerning past events, whenever is necessary, regarding when the events occurred and in what sequence. This task is possible only if accurate timestamps of each event are available. Tasks related to authentication, digital signatures, fault diagnosis, crime investigation, and workload and network traffic characterization may require reliable timestamps so that the events can be listed in chronological order, composing a timeline.

Many computer programs and applications depend on reliable clock synchronization, for example, computer file systems, e-mail and database servers, back-up systems, tools for automatic program execution such as `cron` and `at`, and operating systems utilities to automatic generation of code based on the time of the last modification of the files, such as `make`. In fact, there are numerous services, protocols and systems that rely on one or more timekeeping characteristics, mainly precision, accuracy, and monotonicity.

The lack of synchronization can lead to many problems such as failure in the execution of programmed processes and tasks, loss of data, security holes and loss of credibility of systems and companies. As computers and networks become faster and execute more tasks per unit of time, the time scale of events becomes smaller and the need for accurate and precise timestamps rises as well.

Network Time Protocol (NTP) [1] is the standard protocol for computer clock synchronization throughout the Internet. The goal of NTP is to synchronize clocks of Internet hosts with UTC time servers. NTP builds and maintains a synchronization network for permanent and accurate timekeeping, being a very dynamic, flexible and reliable peer-to-peer network. NTP algorithms, running on Internet hosts, query available peers frequently to collect time information, based on which they compute the best source to synchronize with, and finally they update the own host clock.

This article presents a characterization of a NTP synchronization network, based on recent data collected by a robot that crawled the NTP network from an Internet host. Our focus is to characterize the quality of timekeeping in the discovered NTP network, as well as some aspects of the Internet topology.

Some characterizations were made earlier between 1989 and 1999 [2], [3], [4], [5]. However, in last the six years, the Internet has grown and evolved in many aspects, from the number of users and machines to infrastructure. Thus, we consider opportune to make a characterization of the current NTP network and to compare our results to the results of previous characterizations, discussing about the evolution of timekeeping in the Internet.

This article is organized as follows. In Section II, we describe the architecture of the NTP network and some aspects of the NTP protocol and terminology. In Section III, we present the procedures and events related to data collection and organization. In Section IV, we present the results of the analysis of the dataset and an extensive comparison to results from previous works. Finally, the conclusion finishes the work.

II. The Network Time Protocol

The following definitions and terminology are useful to a better understanding of this work. Accuracy is the degree of conformity of a measurement to its true value or standard reference. In the case of NTP network, the standard time reference is the Coordinated Universal Time (UTC) and the accuracy is, so, how well a clock timestamp compares with UTC. Precision or resolution is the degree to which the time measurement can be determined in a clock, whether second, millisecond, nanosecond or any multiple.

A monotonic clock is a clock that provides non-decreasing timestamps. Although this is not a problem in most clocks, software clocks may be easily updated to a timestamp in the past. As a consequence, errors and misbehaviors may be generated by systems and applications that are not programmed to that situation. Time offset is the difference between two clocks. The stability of a clock is a measure of the frequency fluctuations of a clock, that is, how well it can maintain a constant frequency.

The NTP network is organized in a hierarchical client-server model. The hierarchy is driven by servers' stratum. Stratum 1 or primary servers are at the top (root) of the hierarchy. They synchronize their clocks to very accurate time sources such as GPS or radio clocks. Stratum 2 or secondary servers synchronize their clocks to primary servers. The hierarchy follows to lower level servers until stratum 16 servers, which are at the lowest level. Servers with the higher stratum numbers synchronize their time with the lower stratum number servers.

NTP is pretty well a peer-to-peer network, as it allows the nodes to connect to other peers, forming associations, to provide redundancy and reliability. The network is designed to be self-organizing and self-maintaining. Peers are free to connect symmetrically. Even primary servers may maintain associations with other servers for redundancy.

In NTP terminology, *host* refers to an instantiation of the protocol on a local processor, while *peer* refers to an instantiation of the protocol on a remote processor. NTP algorithms exchange messages with peers to collect time information. The clock selection and clock filter algorithms select the best time source, and the clock update procedure updates the local clock. NTP timestamps represent the main product of the protocol. To issue a timestamp, the NTP protocol need to compute the clock offset, the round-trip delay and the dispersion, all of them relative to a selected reference clock.

III. Data Collection and Organization

To collect data from the NTP network, we implemented a spider, similar to the one described in [5], but suitable for the current network environment. For example, the NTP spider recognizes both IPv4 and IPv6 addresses, and implements a filtering mechanism to prevent queries in servers that do not possess public IP address. New addresses can be added to the filter to prevent queries in specific networks.

The spider, or robot, queried the NTP daemon in each NTP host using the `ntpdc`, which is a NTP query program. Information related to the NTP state in each NTP network node was retrieved by the spider. To begin the first run, the spider was fed with a set of 263 stratum 1 and 2 public NTP time servers, whose addresses were available at [7]. The robot did run from August 30th to September 5th of 2005.

The first run resulted in a set of 1,278,834 NTP hosts (unique IP addresses). A second run was executed from September 26th to September 30th of 2005. This run queried all addresses found in the first run, and discovered new NTP hosts, totalizing 1,290,819 addresses. Only 148,307 responses to queries were received, and we excluded 1,056 due to errors and lost information, ending up with 147,251 complete responses, on which we based our analysis. Only 388 hosts implemented IPv6 addresses. The database has 870 Mbytes.

The spider was executed in a high-end computer running GNU/Linux operating system version sarge, with support to IPv6, connected to the PoP – PR/RNP gigabit network, which is connected to Internet backbones through high speed links.

The information retrieved by the spider includes system variables, that represent the state of the protocol in the local computer, including operating system environment and local clock mechanism, and peer variables, which represent the state of the protocol in the peers (servers) of the host. Each host also maintains a monitor list, which is a list of peers that have queried the host recently.

Security was a fundamental concern in our data collection. System administrators may be worried about those queries, although they are legitimate once NTP makes them available to everyone. To clarify the origin and purpose of the accesses, we set up a DNS record for the domain `ntpsurvey.arauc.br`, and a Web page (`www.ntpsurvey.arauc.br`) explaining the survey and giving contact information. Many system administrators contacted us, but only one asked to be out of the survey, and its IP address was added to the filter.

IV. Results

This section describes the results of the characterization of the discovered NTP-based synchronization network. Wherever is possible, the results are compared to results published in other studies. To our knowledge, this is the fifth published study about the NTP synchronization network. References to papers and the respective year of the data collection are shown in Table I. The year refers to when the data was collected, not the year of publication. Tables presented in this Section show the year of data collection as a reference to the respective publication, according to the relationship presented in Table I.

This Section is divided into three subsections. The first one describes and compares the size and structure of the network. The second subsection characterizes quality of timekeeping and time-related metrics. The last subsection presents data about the topological aspects of the NTP network.

A. Network Growth and Architecture

The growth in the number of network hosts over the years can be seen in Table I. In this Table, "network size" means the number of hosts discovered in the searches, and "replies"

Year	Reference	Network Size	Replies
1989	[2]	8,455	946
1993	[3]	15,000	7,251
1995	[4]	—	38,722
1999	[5]	647,401	175,527
2005	This paper	1,290,819	147,251

TABLE I

NTP NETWORK EVOLUTION: YEAR OF STUDIES AND REFERENCES; NUMBER OF HOSTS AND REPLIES TO QUERIES.

means the number of hosts that effectively answered the NTP queries. In the period, both the number of network hosts and the size of the NTP network (replies) grew about three orders of magnitude. Comparing to the last study, in 1999, the network size almost doubled. However, in this same period of time, the fraction of hosts that answered to queries dropped from 27% to 11% of the hosts discovered.

About 89% of the discovered hosts could not be reached or refused to answer the queries, presumably for many reasons. Some hosts may be blocked by a firewall. Some hosts may have queried a NTP server but they may not have a NTP daemon running. In this case, its IP address was identified but it is not able to get the query. The query itself, composed of UDP packets, might have been lost due to network failures or congestion. Another explanation suggests that some servers are not able to answer the queries because of differences of protocol versions or implementations.

The main characteristic of the NTP architecture is the hierarchy of servers, driven by the stratum number. Table II shows, in the last column, the frequency of servers per stratum in our study. Most of the servers act as stratum 2, 3 or 4, being stratum 3 servers the most common. In absolute numbers, we have 533 stratum 1 servers, 28,687 stratum 2 and 82,145 stratum 3. Similar result was observed in [5], which is reproduced in the fourth column of Table II. This composition is a result of the hierarchical nature of NTP. Primary (stratum 1) servers are the primary reference of time and synchronize their clocks to international time standards using a reliable source of time. Secondary servers (stratum 2) are in the network core and are consulted by many stratum 3 servers, which, in turn, distribute information of time for local networks.

Table II also presents the distribution of the servers per stratum over the years. The percentage of stratum 1 servers is small in all surveys, but we can also observe a considerable drop in this percentage in the last years. This trend may have two explanations. First, we can mention the high cost of acquisition and maintenance of a primary server that allows public access, including security issues. The second explanation is due to the free-scale model [8], which is a model of growth and evolution of the Internet, and its extensions [10], [11].

One of the main principles of those models is the preferential attachment, that is, new nodes make connections preferentially to nodes with an already large connectivity, a phenomenon called “rich get richer”. This model relates the lifetime of a site to its connectivity. Nodes highly connected (hubs) do not die and rarely lose their connections. According to this model, the sprouting of new primary servers would be discouraged because they would have few clients. Moreover, the small number of stratum 1 servers strengthens their huge importance for the quality of timekeeping in the synchronization network.

We can also observe in Table II that the number of stratum 3 servers, that was about 47% in the previous studies, almost reached 59% in our study. This result may be an evidence of the increasing need for synchronization. A larger number of autonomous systems is looking for synchronizing its machines to NTP servers closer to primary servers.

Stratum	1993	1995	1999	2005
1	0,91	1,58	0,55	0,38
2	20,36	31,84	15,29	20,60
3	46,53	47,28	48,61	58,98
4	27,60	16,17	21,84	17,73
5	0,52	2,27	4,06	1,28
6	0,08	0,43	0,94	0,23
7 a 15	0,50	0,43	0,55	0,80
16	3,50	0	2,77	0

TABLE II

EVOLUTION OF THE DISTRIBUTION OF SERVERS PER STRATUM (%).

B. *Characterization of Timekeeping*

The second characterization is about the quality of timekeeping. The NTP algorithms compute values of the metrics clock offset, network distance and dispersion, and rely on these measurements to select the best clock to synchronize. All measurements are relative to a selected reference clock. Clock offset represents the amount of adjustment necessary to the local clock to place it in phase with the reference clock, that is, it is the difference between these clocks. Network distance is measured by the round-trip delay between two hosts. Dispersion represents the maximum error of the local clock relative to the reference clock. Since most NTP hosts will synchronize to UTC via another peer time server, there are two measurements for each of these three metrics, one is measured by all local hosts relative to their synchronization peers, called host–peer, and the other is determined by the peer relative to the primary reference source of standard time, which we call peer–root. In this Section, we analyze these variables.

1) Offset to Synchronization Peer: Offset is the most important indicator of timekeeping performance. It represents the difference between the clocks of a host and its synchronization peer in the time synchronization network. The smaller the offset, the better is the synchronization. Fig. 1 shows the complement of the cumulative frequency of the host–peer offsets. The offset range spreads over many magnitudes, showing a huge variability.

We observe a strong fall in the plot at about 100 ms. NTP defines a threshold for a clock update, usually 128 ms.

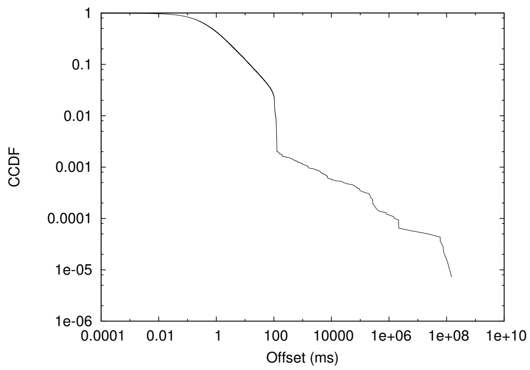

Fig. 1. CCDF of host–peer offsets

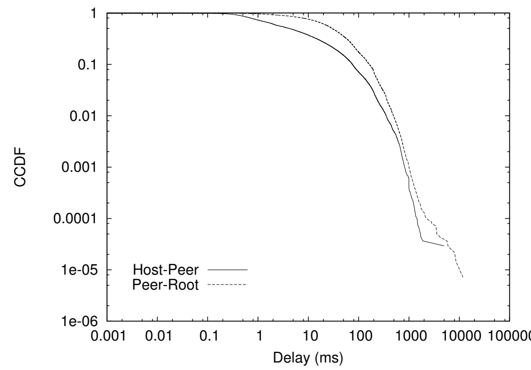

Fig. 2. CCDF of RTT to Synchronization Peer and to Root Server.

Year	Median	Mean	COV
1995	118	186	—
1999	32	33	3,48
2005	4,2	29	2,69

TABLE III

RTT (MS) TO SYNCHRONIZATION PEER IN THREE STUDIES.

Once the host clock was synchronized with the reference clock, the offset will hardly exceed the threshold, even in congestion conditions. Offset samples that exceeds this value are discarded, meaning that they are not considered by local clock algorithms for adjustment. However, if the offset remains above of the threshold for a long time, in general 900 s, the local clock is updated to the time indicated in the reference clock. In these cases, a step adjustment is necessary, otherwise a gradual adjustment is performed.

The long tail to the right indicates that a small portion of hosts presents a large offset, that can reach hours or even days. A comparison to previous studies shows considerable improvement in the quality of timekeeping. In the 1995 study [4], 9.9% of the hosts presented offset larger than 128 ms, while 3% of the hosts showed the same result in 1999 [5]. In our study, only 0.2% of the hosts had offset larger than 128 ms. Excluding values larger than 128 ms in all datasets, the offset median and mean in 1995, 1999, and 2005 are, respectively, 20.1 and 28.7, 1.8 and 8.2, and 0.7 and 7.0, all measurements in ms. These values show that we have substantially better timekeeping in the NTP network today.

2) Network Distance: The metric of network distance is the round-trip time (RTT) or delay between peers. RTT is an important measurement for NTP synchronization algorithms, since the quality of timestamps of a server depends on a small RTT to its synchronization peer and to the root. NTP measures the RTT between a host and its synchronization peer, and also computes the distance to the root, which is the accumulated RTT from a peer to the primary reference of time in the synchronization subnet.

Fig. 2 presents the distributions of host–peer and peer-root RTT values. We observe that the network distance to root is larger than the distance to the synchronization peer, showing that most hosts synchronize with geographically nearby peers. There are a small percentage of NTP peers that synchronize throughout long or congested paths.

Statistical estimates of host–peer RTT in diverse studies are presented in the Table III, while the estimates of peer-root RTT are presented in Table IV. A comparison indicates a significant improvement in the latency of the paths used in the NTP network and, therefore, of the Internet. This improvement in the infrastructure of the Internet can explain improvements in the accuracy of the host clocks. The NTP protocol design assumes that network paths are symmetric, and errors due to asymmetric paths are bounded by half the RTT. As a consequence, an overall improvement in RTT values (lower RTTs) over the network means lower error values and better timekeeping.

Year	Median	Mean	COV
1993	∼ 100	∼160	1,07
1999	47	84	1,86
2005	30	64	1,81

TABLE IV

RTT TO ROOT SERVER (MS) IN THREE STUDIES.

3) Network Distance to Synchronization Peer per Stratum: Table V presents host–peer RTT values as a function of the stratum number of synchronization peer. The largest RTT values are to primary (Stratum 1) servers, in all studies, and those values are more than the double of the RTT to secondary servers. In recent years, we observe that the greater the stratum number, the smaller the RTT.

This result demonstrates the hierarchical arrangement of the NTP in a real environment, with stratum 1 servers being the farthest, and lower level servers being progressively closer to their clients. We also observe, in the same Table, the improvement in RTT to stratum 1 and stratum 2 servers over the years. This results is also an evidence of the improvement in the Internet infrastructure. The higher coefficients of variation in the last two studies may be an evidence of the heterogeneity of the systems and networks in the Internet. The networks are faster but more diverse.

	1993		1999		2005	
Stratum	Mean	COV	Mean	COV	Mean	COV
1	105	1,06	80	2,34	55	1,74
2	42	1,76	29	3,52	24	2,79
3	36	1,72	15	4,47	18	4,83
4	42	0,45	12	3,17	17	3,47
5	50	0,38	3	5,33	9	8,89

TABLE V

RTT (MS) TO THE SYNCHRONIZATION PEER PER STRATUM IN THREE STUDIES.

4) Dispersion to Root Server and Synchronization Peer: Dispersion is an NTP metric for the upper bound error of the local clock relative to the reference clock. NTP hosts compute dispersion to their peers and NTP peers compute dispersion to their root server. The dispersion to the root represents the accumulated error in a synchronization path from a peer to the root of the NTP subnet.

Fig. 3 presents the CCDF distribution of host–peer and peer–root dispersions. The dispersion to root is larger than the dispersion to peers, and it also presents a longer tail. The tail is composed by 3% of the data, and largest dispersion is about 16 hours. Comparing to previous studies, there is large evidence of improvement in synchronization accuracy. In the 1999 study, the largest dispersion was over one year. The tail was also composed by about 3% of the data. Taking out the tail of both datasets, the values for median and mean dispersion, in 1999 and 2005 are, respectively, 39 and 88, and 40 and 74 ms. We observe a decrease in mean dispersion but not in median dispersion. It may indicate that, while larger errors can get smaller, there is a lower bound for dispersion, which may be related to the quality of computer clocks.

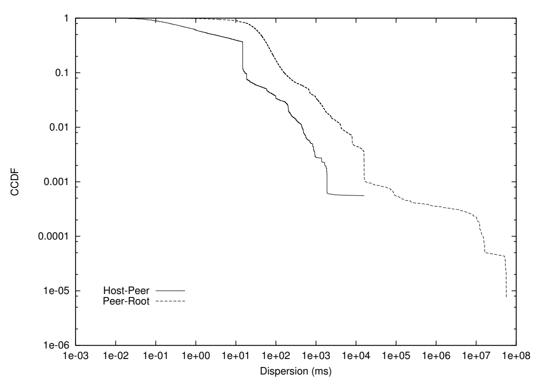

Fig. 3. Dispersion to synchronization peer and to root server.

5) Clock Stability: The stability of a clock is a measure of the frequency fluctuations of a clock, that is, the variation of frequency. It is an inherent characteristic of an oscillator that determines how well it can produce the same frequency over a given time interval. Those variations can be random or systematic. The stability depends on many aspects that have influence over the quality of a clock oscillator such as temperature, humidity, pressure, vibration, magnetic and gravitational fields. NTP algorithms try to compensate the error in oscillators. The stability is measured as an exponential average of square root of the mean (RMS) frequency differences, given in parts per million (PPM).

Stability is also a system variable in the set of server NTP variables. Fig. 4 shows the complement of the CDF of stability of servers' clocks. The median is 0.015 PPM and the top quartile is 0.239 PPM. The long tail rises the mean to 228 PPM but when we get out of 3% of the larger values, the mean goes to 2.32 PPM and standard deviation of 9.6 PPM. It means a difference of approximately 200 ms a day.

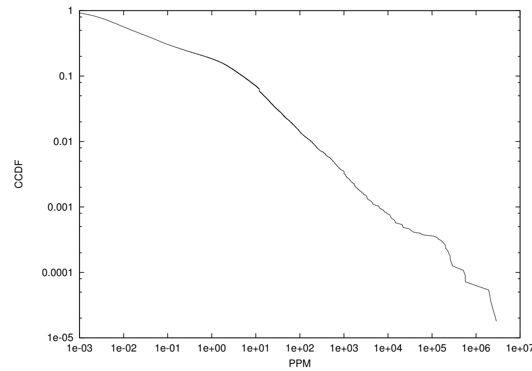

Fig. 4. Clock Stability (PPM).

C. Network Topology Characterization

There are three main aspects in the characterization of NTP network topology. The first aspect refers to the number of peer associations in a server. The NTP protocol collects and analyzes timing data from all peer associations of a host to select the best clock to be used for synchronization from possible several servers. Multiple time servers allows NTP to use statistical techniques to choose a reliable source of synchronization, as well to improve system redundancy.

The left plot in Figure 5 presents the CDF of the number of associations per server, with the servers classified per stratum. Stratum 1 and stratum 2 servers present similar distributions of the number of associations. In both cases, about 20% of the servers has only one association, and about 50% has up to three associations. Stratum 3 servers have even less associations. About 31% of stratum 3 servers has only one association and 75% has up to three associations.

There are servers that maintain a large number of associations, in the order of hundreds. The tail reveals the large range in connectivity in the NTP network. It is in accordance to the scale free model, and the preferential attachment phenomenon [8], [10], [11]. The large variability in the number of node links, or number of associations, as we see in Fig. 5 is a consequence of this phenomenon.

A second characterization of the NTP network topology is the number of peers from a server point of view, that is, the number of clients that has seek time information from a given server. This number is given by counting on the dataset, for

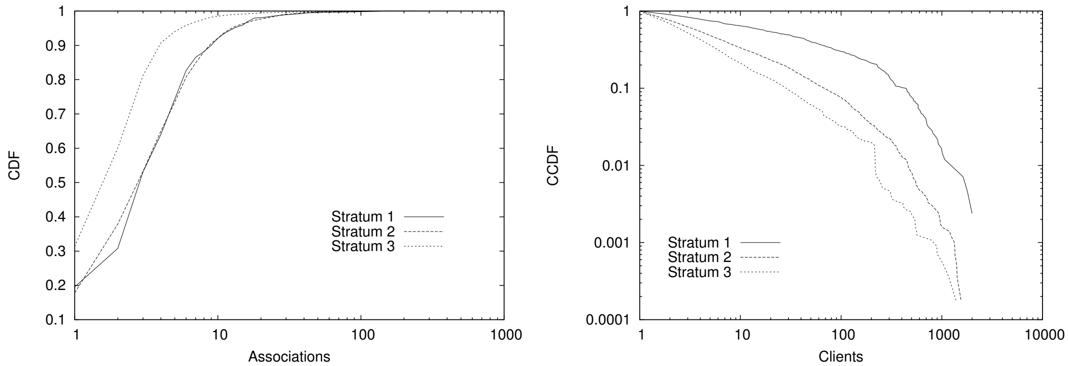

(a) Associations per server.

(b) Clients per server.

Fig. 5. Distribution of the associations per stratum (left) and per clients (right).

each server, how many times its IP address is listed as a time source in the set of peers.

The right plot in Fig. 5 shows that primary servers (Stratum 1) have a larger number of clients than stratum 2 servers, and stratum 2 servers have a larger number of clients than stratum 3 servers. About 50% of the primary servers has up to 27 clients. About 10% of stratum 1 servers, 1.3% of stratum 2 servers and 0.3% of stratum 3 servers have more than 400 clients. This distribution also confirms the scale free model for Internet topology.

A last aspect in the characterization of NTP network topology is the load on NTP servers, measured by the number of stratum $n+1$ servers over the number of stratum n servers. Table VI shows the load values. It shows a considerable rise in the stratum 1 load index in our study, compared to previous studies. In all studies, the load index of primary servers is large, compared to stratum 2, 3 and 4 servers. This result may be explained by the drop in the number of primary servers, despite the growth in number of hosts in the NTP network. It also indicates the role of primary servers as hubs of the NTP network.

Stratum	1993	1995	1999	2005
1	22,36	20,17	28,03	53,45
2	2,29	1,48	3,18	2,86
3	0,59	0,34	0,45	0,30
4	0,02	0,14	0,19	0,07

TABLE VI

TOPOLOGICAL CHARACTERIZATION: LOAD INDEX IN FOUR STUDIES.

V. CONCLUSION

In this paper, we presented a characterization of a NTP synchronization network composed by thousands of nodes, including hundreds of primary servers. We compared our results to the results of five other studies made in the last fifteen years. The data analysis reveals that the NTP network experienced considerable evolution in quality of timekeeping in this period. The analysis shows evidence of improvement in Internet infrastructure, which helps to improve timekeeping performance. There is evidence that more and more systems are trying to keep synchronization to the NTP network, and they search for more reliable time sources.

The analysis of topological aspects confirms recent theories and models of Internet topology. As a main aspect, the drop in the number of primary servers and the corresponding increase in their index load shows that they act as hubs in the core of the NTP network, which is in accordance to the principles of those theories. Our dataset can be explored further in a number of ways and we are working on this.

REFERENCES

[1] Mills, D. L.: Network Time Protocol (Version 3) Specification, Implementation and Analysis. Request for Comments 1305 (1992).
[2] Mills, D. L.: On the Accuracy and Stability of Clocks Synchronized by the Network Time Protocol in the Internet System. ACM Computer Communication Review 20(1) (1990)
[3] Guyton, J. D., Schwartz, M. F.: Experiences with a Survey Tool for Discovering Network Time Protocol Servers. Proc. Usenix Summer Tech. Conf. (1994)
[4] Mills, D. L., Thyagarajan, A. Huffmann, B. C.: Internet Timekeeping around the Globe. Proc. Precision Time and Time Interval Applications and Planning Meeting (1997). http://www.eecis.udel.edu/m̃ills/database/papers/survey5.ps
[5] Minar, N.: A Survey of the NTP Network. http://www.media.mit.edu/ñelson/research/ntp-survey99/ (1999)
[6] Mills, D. L.: On the Chronometry and Metrology of Computer Network Timescales and their Application to the Network Time Protocol. SIGCOMM Comput. Commun. Rev. 21(5) (1991)
[7] NTP Project: Network Time Protocol Public Servers Project. http://ntp.isc.org (2005)
[8] Barabasi, A. L.: Linked: The New Science of Networks. Perseus Publishing (2002)
[9] Faloutsos, M., Faloutsos, P.,Faloutsos, C.: On Power-law Relationships of the Internet Topology. Proc. ACM SIGCOMM (1999)
[10] Dorogovtsev, S. N., Mendes, J. F. F.: Evolution of Networks: From Biological Nets to the Internet and WWW. Oxford University Press (2002)
[11] Pastor-Satorras, R., Vespignani, A.: Evolution and Structure of the Internet: A Statistical Physics Approach. Cambridge University Press (2004)