A comparison of NTP servers connected to the same reference clock and the same network

Andrew N. Novick and Michael A. Lombardi
Time and Frequency Division, National Institute of Standards and Technology, Boulder, Colorado, USA

ABSTRACT

Network Time Protocol (NTP) servers maintained by national timing laboratories are ideally synchronized to a one pulse-per-second (pps) signal from the official national time scale. This method allows timing laboratories to distribute the official time of their respective nations via NTP, as opposed to distributing time obtained from Global Navigation System (GNSS) satellites or another reference. Distributing the official time is necessary because some sectors, including stock markets and financial exchanges, are often legally required to utilize time from a specific national time scale, such as UTC(NIST), the time scale maintained by the National Institute of Standards and Technology in the United States. To investigate how well a nation's official time can be distributed via NTP, this paper compares the accuracy and stability of four commercially-available NTP servers that are referenced to a 1 pps signal from UTC(NIST). Measurements were performed with all of the servers residing on the same subnet at NIST, where they were accessible to the general public via the Internet. The servers were continuously compared to a NTP client, located on a different NIST subnet, that was referenced to UTC(NIST) with 100 ns resolution. The results of these comparisons reveal the time differences between the servers and UTC(NIST), the relative time differences between the servers (obtained in common-view mode), and the time stability of each server.

1. INTRODUCTION

Network Time Protocol (NTP) servers are now sold commercially and deployed by numerous organizations. The manufacturers typically provide several options for automatic synchronization, with the Global Positioning System (GPS) being the most commonly selected option. When GPS is selected, the server is delivered with an integrated receiver and antenna and the NTP packets transmitted by the server are synchronized to Coordinated Universal Time (UTC) as obtained from GPS. Transmitting time obtained from GPS, however, is not acceptable for all applications. For example, national metrology institutes (NMIs) who are responsible for transmitting the official time for a given nation, such as the National Institute of Standards and Technology (NIST) in the United States, must transmit time that is obtained from their own time scale and is independent of all other sources [1]. Other applications require the server to be synchronized to a specific, non-GPS, source of time. Perhaps the best known example of this situation occurs in financial markets, where computers involved in stock market transactions are required to be synchronized to time from NIST in the United States [2, 3] or, in the case of European stock exchanges, to be synchronized to time from any NMI that contributes to UTC [4].

Synchronizing an NTP server to time from a specific NMI is ideally done by connecting a 1 pulse per second (pps) signal that originates from the NMI's time scale to the server. This situation is preferable because it ensures that the server's on-time marker (OTM) will be directly obtained from the time scale and the server clock will "tick" at the same frequency as the time scale clock. However, because the 1 pps signal only provides the "tick", a source of time-of-day also originating from the NMI must be available to name or label each second, but after the initial synchronization the time-of-day source will only be needed again when synchronization is lost, for example after a power outage or a 1 pps signal interruption.

To accommodate the needs of NMIs and others who must synchronize NTP servers with an autonomous source of time, several server manufacturers offer models that allow synchronization via 1 pps signals. This paper directly compares four different commercially-available NTP servers by connecting them to the same reference clock, a 1 pps signal from the UTC(NIST) time scale, and to the same network. Section 2 provides a description of the four servers. Section 3 describes the measurement configuration. Measurement results are provided in Sections 5 and 6, and Section 7 provides a short summary.

2. DESCRIPTION OF NTP SERVERS INVOLVED IN COMPARISON

The four NTP servers involved in the comparison were produced by three different manufacturers and are summarized in Table 1. Note that Server A had the fastest network interface. Also note that Servers C and D were produced by the same manufacturer and have the same model number, but Server D differs from Server C because it has an internal rubidium time base oscillator and does not accept an external time base. All external time bases and synchronization sources were obtained from the UTC(NIST) time scale.

Table 1. Summary of NTP servers involved in comparison.

Server	Rack Height	Internal Time Base	External Time Base	Synchronization Source	**Network card**
A	1 U	Quartz	1 pps	1 pps	1 Gb
B	2 U	Quartz	10 MHz	1 pps	10/100 Mb
C	1 U	Quartz	10 MHz	1 pps	10/100 Mb
D	1 U	Rubidium	None	1 pps	10/100 Mb

3. MEASUREMENT CONFIGURATION

Prior to the measurements, the four servers were installed and connected to a 1 pulse per second (pps) signal from UTC(NIST) which served as their time reference. The servers were initially set on time manually to a UTC(NIST) source because the 1 pps signal does not designate the time of day. The cables for the time reference were of equal length from a distribution amplifier, and the delay of 502 ns from UTC(NIST) was entered as a reference delay into each server, so the servers all had an identical reference clock. Two of the servers (B and C) were also connected to a 10 MHz signal from UTC(NIST) which replaced their internal quartz time base oscillator. Server A locked its quartz oscillator to the 1 pps signal, and Server D utilized its internal rubidium oscillator as its time base.

The servers were connected to the same network switch with equal-length Category 5 (Cat 5) cables and the switch was connected to the network outside the network firewall. This was done to replicate the accessibility of a typical NTP server which can accept timing requests from computers anywhere on the Internet. When an NTP timing packet request was made, it propagated through a layer 2 switch, an internal router, and then through a firewall to the NTP servers and back. An access control list (ACL) was used to designate the path so it did not vary. The additional hardware layers between the server and the client potentially add network asymmetry and timing uncertainty.

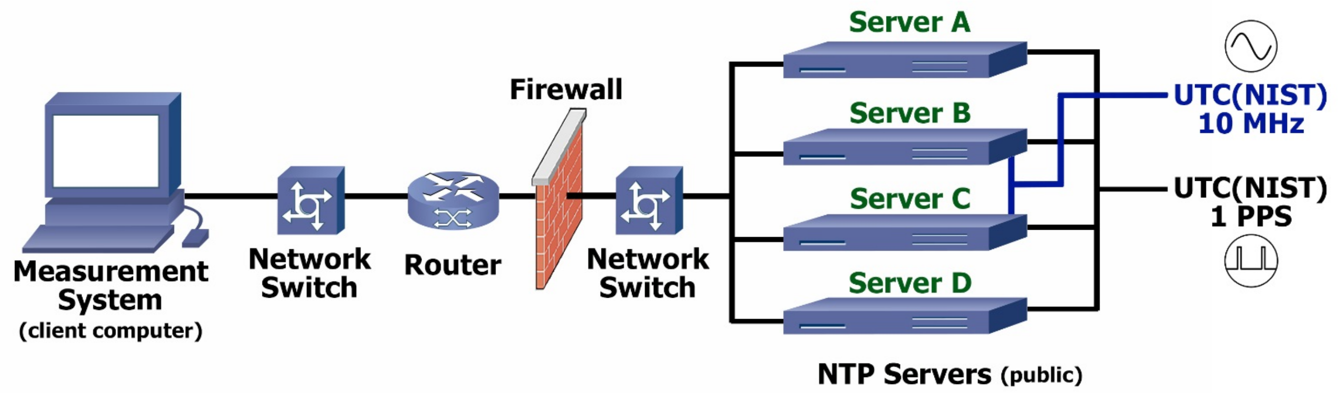

Figure 1. A diagram of the measurement configuration.

The measurement system was a computer operating as an NTP client, utilizing a hardware timing board as its system clock [1, 5] synchronized to UTC(NIST) with a 1 pps reference signal. A previous paper [5] described settings for the network interface card (NIC) and the computer's basic input output system (BIOS) in order to reduce client asymmetry to a minimum and the same settings were utilized in this experiment. The delay of the reference signal (438.0 ns) was compensated for in software, thus the NTP servers under test and the NTP measurement system had essentially the same reference clock (within ~1 ns of each other). The measurement system requested an NTP packet from each of the four servers every 10 s, compared the results to its system clock, and recorded the time difference (TD) and the round trip delay (RTD). At the end of each 10-minute segment, the 60 measurements collected during the segment were examined and the one with the smallest round trip delay was stored and used in the analysis provided in Sections 4 and 5. Ten-minute averages were also stored and were used in the stability analysis in Section 4 and the discussion of server anomalies in Section 6. A diagram of the measurement configuration is provided in Fig. 1.

4. MEASUREMENT RESULTS VIA DIRECT COMPARISON TO UTC(NIST)

Because the four servers were connected to the same clock and shared the same network, the goal of the measurement was to determine how closely each server was able to synchronize its clock to the 1 pps signal from the NIST time scale. In an attempt to accomplish this, data were collected for 40 days (10/04/2016 to 11/13/2016, MJD 57665 to 57705) from the measurement configuration shown in Fig. 1. The TD and RTD results for the 40-day period are shown in Fig. 2, with one data point recorded every 10 minutes.

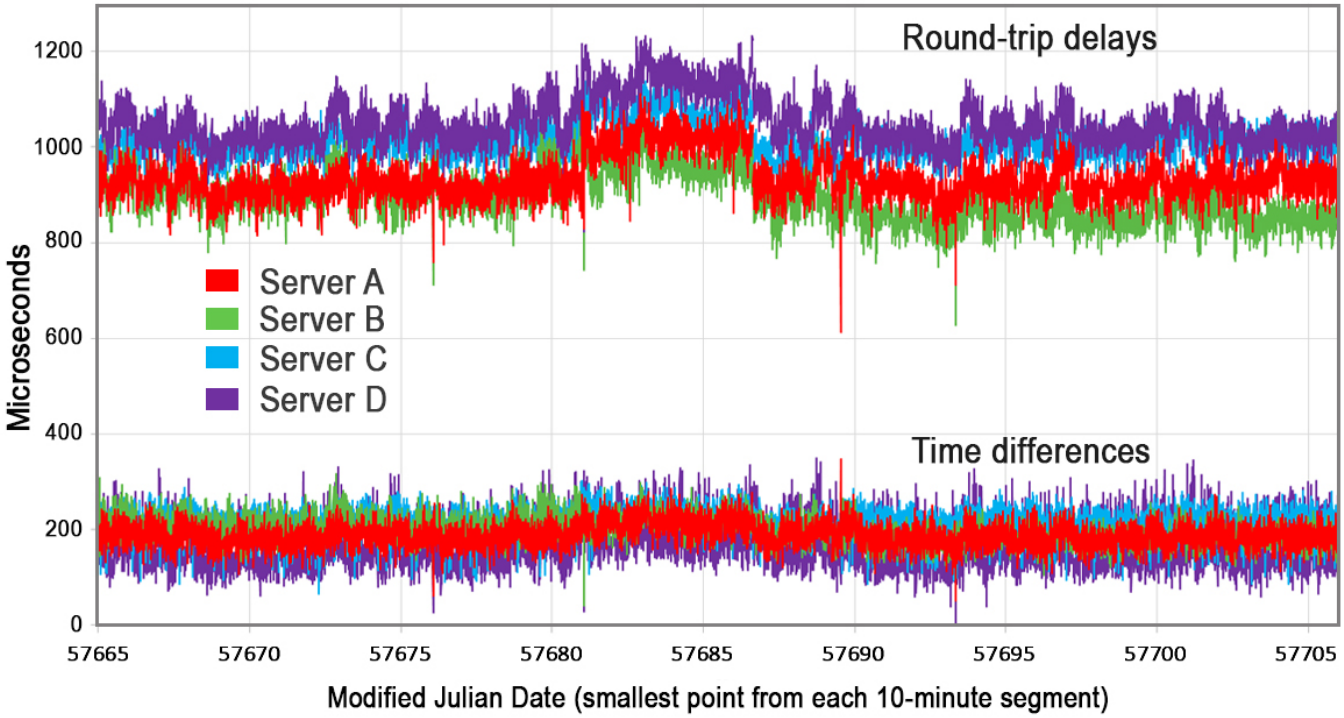

Figure 2. Common clock comparison of round trip delays and time differences between four NTP servers.

Figure 2 indicates that the average RTD was near 1000 µs (1 ms) and that the average TD was near 200 µs (0.2 ms), or about 20 % of the RTD. If all of the 200 µs TD was due to asymmetry, it would mean that the difference between the path delays to and from the server would be near 400 µs, or in this case about 700 µs in one direction and only about 300 µs in the other. It is well known that network asymmetry is the largest contributor to NTP time uncertainty, and that asymmetry cannot be ruled out as the sole source of uncertainty until the TD exceeds 50 % of the RTD (a situation that could only exist if all of the network delays were in one direction). Even so, a TD equivalent to 20 % of the RTD seems unusually large and it seems likely that at least some of the uncertainty in the TD values is due to server synchronization uncertainty and not to network asymmetry. This argument is supported by the fact that Server D had the smallest average TD (152.5 µs) and the largest RTD (1050.4 μs), as indicated in Table 2. The table summarizes the average values and stability of the TD and RTD for each server, with the green cells indicating the smallest value in each category.

Table 2. Summary of NTP server differences.

	Time Differences (μs)		Round Trip Delays (μs)	
Server	**Average**	**TDEV (τ = 10 min)**	**Average**	**STDEV**
A	188.1	22.9	935.3	46.7
B	200.5	21.5	895.0	51.1
C	220.2	25.7	1014.9	37.1
D	152.5	38.6	1050.4	51.1

Although Server D had the smallest average TD, it was least stable server, with the largest time deviation (TDEV), $\sigma_x(\tau)$, of 38.6 μs and tied with Server B for the least stable RTD (STDEV of 51.1 μs). Interestingly, Server B also had the smallest average RTD (895.0 μs).

Server C had opposite characteristics from Server D in some results; it had the largest average TD from UTC(NIST), 220.2 μs, and the lowest RTD STDEV (37.1 μs). This is especially interesting because the two servers (C and D) are produced by the same manufacturer, with the only specified difference being that Server D has to rely on its internal rubidium oscillator as its time base. Without an external time base input, it perhaps makes sense that the stability of Server D is lower than the other servers, which all lock their time base frequency to a signal from UTC(NIST). As previously indicated in Table 1, Server A had the fastest network interface, 1 Gb/s. However, this did not produce an obvious advantage in our results.

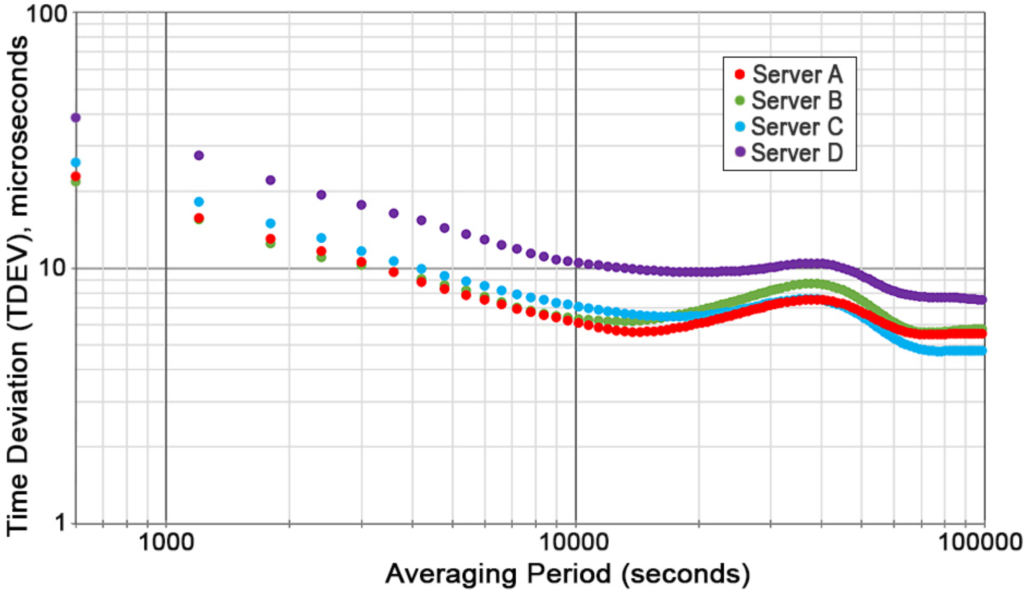

Figure 3. A comparison of the time stability of the four servers.

Figure 3 compares the time stability of the four servers by showing their time deviation (TDEV), $\sigma_x(\tau)$, for periods ranging from 600 s (10 minutes) out to more than one day. The stability of Servers A, B and C is below 10 μs after a period of slightly more than one hour, reaching an initial noise floor of about 6 μs in less than five hours. Server D, with its internal rubidium time base oscillator, was the least stable of the four servers at all averaging periods. The stability at one day reaches or drops below the initial noise floor for each of the four servers. All four servers show a bump in server stability at a period near 40 000 s, or about half a day. The ten-minute TD points with the smallest RTD shown in Fig. 2 do not clearly indicate any diurnal variations which would cause the bump in the stability graph in Fig. 3. However, Fig. 4 shows the ten-minute average of the time difference for Server A during a 14-day period near the end of the measurement run (10/26/2016 to 11/9/2016, MJD 57687 to 57701), revealing increased time delays during the daytime hours on the weekdays (for example, from MJD 57692 through 57696). During evenings and weekends, the network traffic onsite is at a minimum. The weekday traffic diurnal is most likely the cause of the half-day bump in the TDEV.

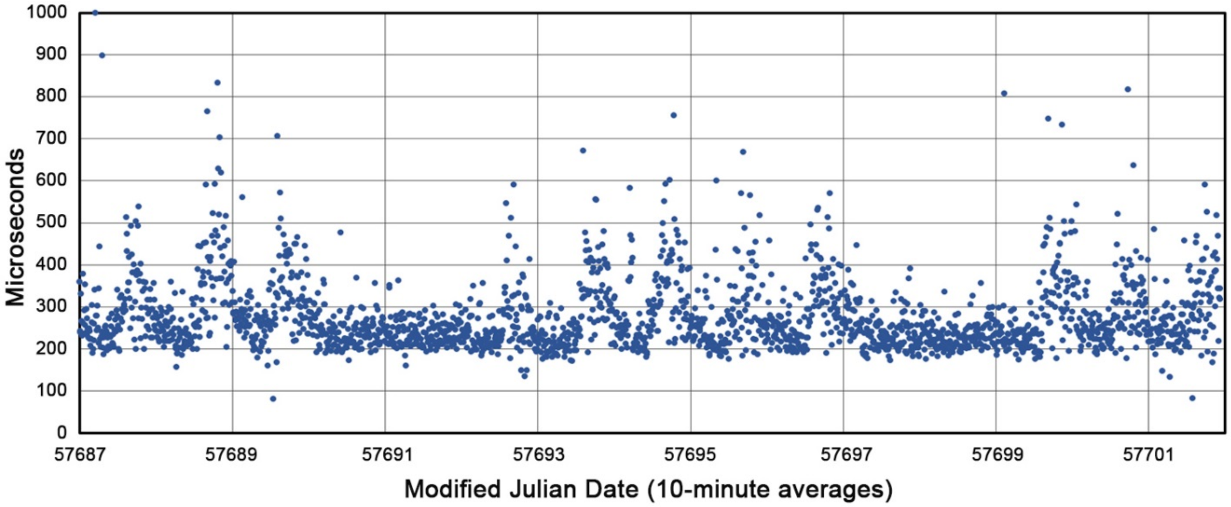

Figure 4. Ten-minute average time differences for Server A for a 14-day period.

5. MEAUREMENT RESULTS VIA COMMON-VIEW

Conclusively separating the uncertainty due to network asymmetry from the uncertainty in server clock synchronization is difficult, but it is easy to determine that there are repeatable differences in the time transmitted by the four servers by utilizing the common-view method with UTC(NIST) serving as the common-view signal source. For example, if two nearly simultaneous measurements are made, the first of Server A – UTC(NIST) and the second of Server B – UTC(NIST), we can subtract the first measurement from the second to obtain the relative difference between Server A – Server B. As noted, common-view measurements show relative differences between servers, and not absolute differences with respect to UTC(NIST), which was the initial goal of our experiment. However, common-view measurements do show that the uncertainty of synchronization with respect to UTC(NIST) varies from server to server, which is useful information. If the servers being compared via common-view are widely separated and connected to different networks, common-view measurements might not be useful for comparing server clocks because the path delays are not equivalent. Here, however, the server clocks being compared are on the same network as each other and only a few hops (pieces of network hardware) away from the common source, so the path delays are relatively equivalent allowing the time differences between server clocks to be revealed.

For the same measurement period utilized in Fig. 2 (10/04/2016 to 11/13/2016, MJD 57665 to 57705), common-view comparisons were made, showing the time differences between all possible server combinations for the entire 40-day period. Figure 5 shows one-day averages of the comparisons and the differences between all servers exceed 10 μs.

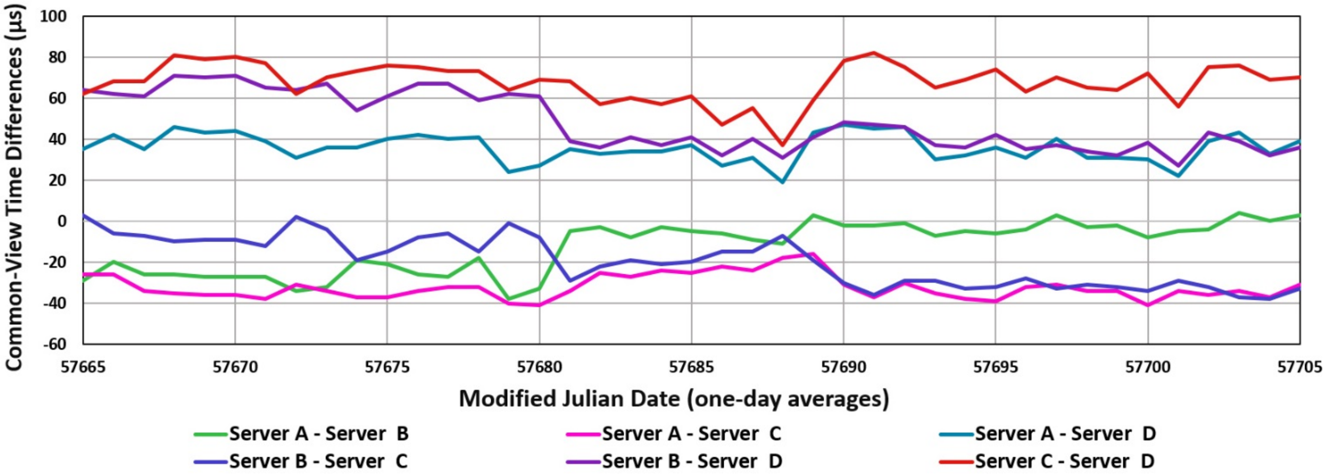

Figure 5. Results of intercomparisons of the NTP servers via the common-view method.

Table 3 summarizes the results, indicating that the largest difference is 67.8 µs between Servers C and D, which are made by the same manufacturer but utilize different time bases. The smallest time difference is Server A – Server B, which was 12.4 µs. These two servers had the lowest TDEVs for τ = 10 min and the smallest average RTDs.

Table 3. Summary of common-view time differences (one day averages) between NTP servers.

Server	A	B	C	D
A	---	12.4 µs	32.1 µs	-35.6 µs
B	-12.4 µs	---	19.8 µs	48.0 µs
C	-32.1 µs	-19.8 µs	---	-67.8 µs
D	35.6 µs	-48.0 µs	67.8 µs	---

6. SERVER ANOMALIES

We have examined the outputs of several NTP servers under the same conditions and how they vary, and the differences have been somewhat consistent. However, consider a case where the measured output of a server changes and the RTD does not. It could be explained by the delays in each direction of the packet changing in equal but opposite ways [6], thus flipping the asymmetry, but an anomaly in the server timing could also be the cause. For instance, Fig. 6 shows a period where the RTD did not change significantly and Servers A and B also did not have a significant change, but Servers C and D changed levels for more than one day and then returned to similar values. Server C moved by an average of -64 µs and Server D moved by an average of +134 µs. The temperature and humidity in the laboratory was stable during this period. This appears to be caused by anomalies in the servers themselves and it is very curious that two servers reacted at the same time. Again, these are the two servers of the same manufacturer and some characteristic of the internal architecture could be the cause. Consistent time differences can be removed by calibrating an NTP server, but anomalous changes in server outputs are unpredictable.

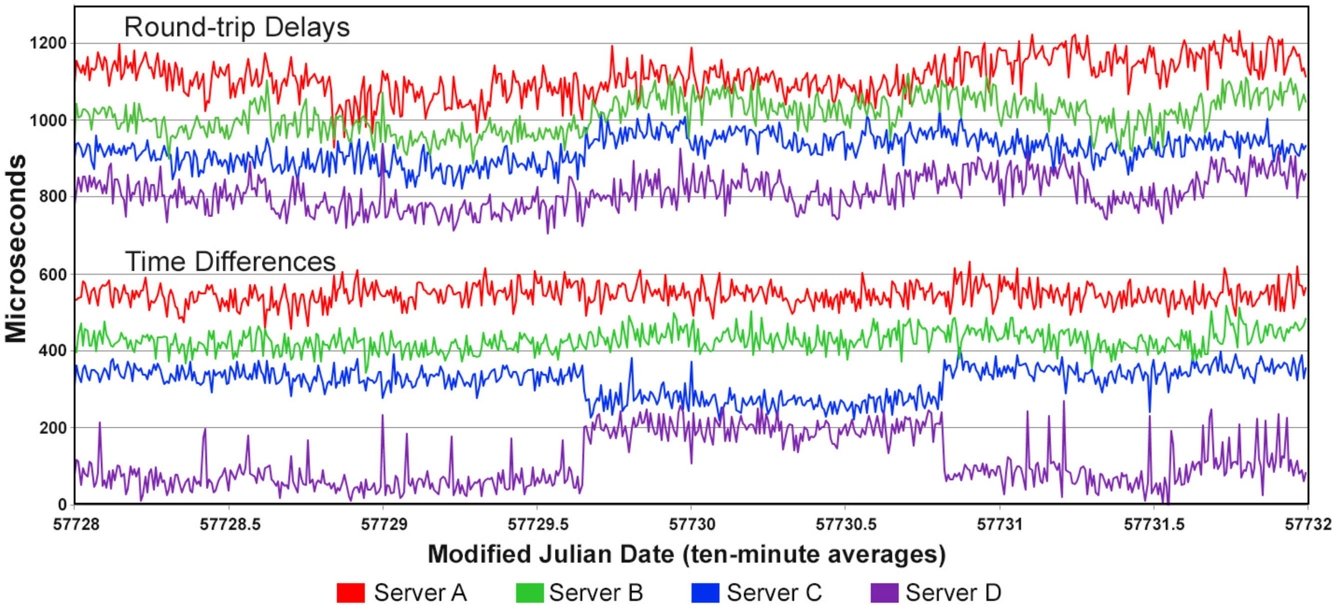

Figure 6. TDs and RTDs of four NTP servers (tracks are separated for clarity) during an anomalous occurrence.

7. SUMMARY

Timing laboratories can distribute their nation’s official time via NTP to a very large number of users and more attention should be paid to the uncertainty of the received time. By comparing commercial NTP servers with the same reference and on the same network, we have shown differences in server clocks at levels of tens of microseconds that do not appear to be related to network asymmetries. Due to uncertainties both in server clock synchronization and due to network asymmetries, it seems clear that all timing laboratories that distribute time via NTP should monitor the output of their servers, not just internally but preferably from the public Internet, to simulate the experience of customers who are requesting time from the public Internet. Future work could involve the development of better techniques for estimating the contribution of server clock uncertainty to the overall time measurement uncertainty, and measuring the holdover capabilities of NTP servers that have temporarily lost their reference clock.

8. REFERENCES

[1] M. Lombardi, et al. "International Comparisons of Network Time Protocol Servers," *Proceedings of the 46th Annual Precise Time and Time Interval Systems and Applications Meeting*, 1-4 December 2014, Boston, Massachusetts, 57-66.

[2] Financial Industry Regulatory Authority, “Clock Synchronization,” (FINRA, Rockville, MD), *Regulatory Notice 16-23*, 2016.

[3] United States Securities and Exchange Commission, “National Market System Plan Governing the Consolidated Audit Trail Pursuant to Rule 613 of Regulation NMS under the Securities Exchange Act of 1934,” (SEC, Washington, D.C.), 2016.

[4] European Securities and Markets Authority, “Regulatory technical and implementing standards – Annex I: MiFID II / MiFIR,” (ESMA, Paris, France), *ESMA/2015 /1464*, 2015.

[5] A. Novick and M. Lombardi, 2015, “Practical Limitations of NTP Time Transfer,” *Proceedings of the 2015 Joint Conference of the IEEE International Frequency Control Symposium and the European Frequency and Time Forum*, 12-16 April 2015, Denver, Colorado, 570-574.

[6] K. Vijayalayan and D. Veitch, 2016, "Rot at the roots? Examining public timing infrastructure," *Proceedings of the 35th Annual IEEE International Conference on Computer Communications*, 10-14 April, 2016, San Francisco, California, 1-9.