International Comparisons of Network Time Protocol Servers

Michael A. Lombardi, National Institute of Standards and Technology (NIST), United States, lombardi@nist.gov
Judah Levine, National Institute of Standards and Technology (NIST), United States
J. Mauricio Lopez, Centro Nacional de Metrología (CENAM), Mexico
Francisco Jiménez, Centro Nacional de Metrologia (CENAM), Mexico
John Bernard, National Research Council (NRC), Canada
Marina Gertsvolf, National Research Council (NRC), Canada
Harold Sanchez, Instituto Costarricense de Electricidad (ICE), Costa Rica
Oscar G. Fallas, Instituto Costarricense de Electricidad (ICE), Costa Rica
Liz Catherine Hernández Forero, Instituto Nacional de Metrología (INM), Colombia
Ricardo José de Carvalho, National Observatory (ONRJ), Brazil
Mario N. Fittipaldi, National Observatory (ONRJ), Brazil
Raul F. Solis, Centro Nacional de Metrologia de Panama (CENAMEP), Panama
Franklin Espejo, Instituto Boliviano de Metrología (IBMETRO), Bolivia

ABSTRACT

This paper describes a recently designed system that measures the time transmitted by network time protocol (NTP) servers located in North, Central, and South America. Direct measurements of the time transmitted by each server are obtained by comparing the received time stamps to the Cordinated Universal Time scale, UTC(NIST), located at the National Institute of Standards and Technology (NIST) in Boulder, Colorado. International comparisons of the time differences between servers are obtained by utilizing UTC(NIST) as a common-view observation signal. Results of both direct and common-view measurements are published in real time via the web site of the Sistema Interamericano de Metrologia (SIM) Time and Frequency Metrology Working Group, providing verification of a given server's accuracy. The paper describes the NTP measurement system, and presents results from direct comparison and common-view measurements. It also discusses factors that contribute uncertainty to NTP measurements; in particular network asymmetry caused by inconsistent routing and/or network congestion.

I. INTRODUCTION

The Network Time Protocol (NTP) [1] is designed to synchronize clocks via packet-switched, variable latency networks; in particular via the public Internet. Because NTP servers are relatively inexpensive and easy to maintain; and because accurate computer clocks are important to many industrial and financial applications, the timing laboratories at many national metrology institutes (NMIs) now utilize NTP as the dominant or sole method for distributing the national time. For example, at some NMIs, older time distribution systems that utilized telephone lines or radio broadcasts have been supplemented or replaced by NTP servers. In the case of recently established laboratories that are looking to start their first time service, providing the national time via NTP is the wisest and most cost effective choice. In all cases, NMIs must be able to verify the accuracy of the time sent via NTP, because the distribution of incorrect time by an official time provider is unacceptable.

II. THE USE OF NTP SERVERS IN THE SIM REGION

The Sistema Interamericano de Metrologia (SIM) is a regional metrology organization that includes 34 NMIs located in North, Central, and South America. As of November 2014, 22 of these NMIs either maintain, or have selected a designated institute (DI) to maintain, a national time standard that is regarded as an official source of time in their respective nations. Each time standard is continuously compared to other SIM time standards via common-view satellite observations through the SIM Time Network [2]. Once a national time standard has been established and validated through international comparisons, the obvious next step is to make the time available to the nation's citizens. As a result, eleven SIM NMIs now distribute their national time via one or more NTP servers, and eight of these 11 have at least one server that is currently monitored by the measurement system described in this paper.

The NTP servers being measured are synchronized with the national time standard in their respective countries. Table 1 lists the servers and their synchronization source.

Table 1. NTP servers operated by SIM NMIs/DIs.

NMI or DI	Country	Synchronization Source
CENAM	Mexico	1 pps from local time scale
CENAMEP	Panama	1 pps from local time scale
CMEE*	Ecuador	1 pps from local time scale
IBMETRO	Bolivia	1 pps from oscillator disciplined to SIM Time Scale [3]
ICE	Costa Rica	1 pps from local time scale
INDECOPI*	Peru	1 pps from local time scale
INM	Colombia	1 pps from local time scale
NIST	United States	1 pps from local time scale or NIST Automated Computer Time Service [4]
NRC	Canada	1 pps from local time scale
ONRJ	Brazil	1 pps from local time scale
SLBS*	St. Lucia	1 pps from oscillator disciplined to SIM Time Scale [3]

* Server not yet monitored due to restricted accessibility.

III. DESCRIPTION OF NTP MEASUREMENT SYSTEM

The SIM NTP measurement system became operational on May 31, 2014 (MJD 56808) at NIST in Boulder, Colorado. It currently has the capacity to measure the time from 20 servers (sufficient for monitoring one or two servers from each NMI/DI in Table 1) but can be expanded as necessary. The system includes an internal clock board with 100 ns resolution that is continuously synchronized by a 1 pps (pulse per second) signal from the UTC(NIST) time scale. The cable connecting the measurement system to UTC(NIST) has a calibrated (and compensated) delay of 480 ns. The time kept by the clock board is compared to the time stamps in the NTP packets transmitted by the servers to estimate the uncertainty of the server's clock. Figure 1 provides a block diagram.

A timing packet is requested from each NTP server being monitored every 10 s. This request is made by sending a 48-byte packet via the user datagram protocol/Internet protocol (UDP/IP) to port 123 of each server. The last eight bytes of the packet include the time of the request, T_1, as obtained from UTC(NIST) via the clock board.

Each server responds to the timing request by returning a data packet. The entire packet is decoded by the measurement system, including three 64-bit time stamps. These time stamps utilize 32 bits to represent integer seconds, and an additional 32 bits to represent fractional seconds, providing resolution of 2^{-32} s (233 ps).

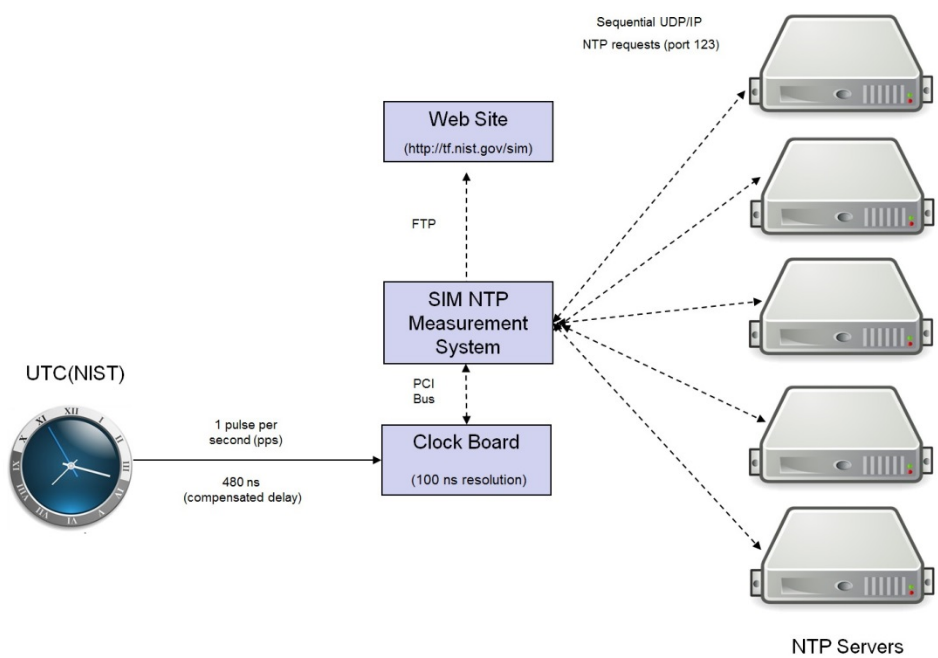

Figure 1. The SIM NTP Measurement System.

One of the time stamps returned by the server simply echoes back T_1, the time when the measurement system (client) made the request. Two other time stamps contain the time, T_2, when the request was received by the server; and the time, T_3, when the server transmitted its response. When the measurement system (client) receives the packet, it again queries the clock board and records T_4, the time of the packet's arrival. The difference between the time transmitted by the server and UTC(NIST) is obtained using the standard NTP equation for clock offset [1],

$$TD_{NIST} = \frac{((T_2 - T_1) + (T_3 - T_4))}{2}, \qquad (1)$$

where TD_{NIST} is the time difference with respect to UTC(NIST). Using these same four time stamps, the round trip delay between the client and server can be calculated as [1]

$$RT_{Delay} = (T_4 - T_1) - (T_3 - T_2). \qquad (2)$$

Note that the "divide by 2" in Eq. (1) assumes that the delay from the server to the client is equal to one half of the round trip delay. If this assumption were true, the network path would be symmetrical and dividing by two would fully compensate for the path delay. In practice, however, the network path always has some asymmetry. This asymmetry, and thus the uncertainty of the time received by the client, generally (but not always) increases as a function of the round trip delay.

The results of both the time difference and round trip measurements are updated every 10 s on the measurement system's display (Fig. 2) and recorded in a file. As noted previously, the system can currently measure up to 20 NTP servers, but only 12 measurements are displayed on the screen at one time.

As shown in Fig. 2, common-view time differences larger than 20 ms are displayed on a red background. The display also includes a row showing the last round trip delay for each measurement. Round trip delays that exceed 100 ms are displayed on a red background. Another row serves as a health indicator, displaying a "green light" if a server is healthy and a "red light" if it is not. The NTP packet does not include a health flag, but it has long been an accepted practice to set both leap second bits (in the packet's LI field) to 1 to indicate that the server's clock is not synchronized [5]. Thus, this method is utilized to obtain a health flag for the display. However, because the system is designed to detect time errors (and because NTP clients may knowingly or unknowingly ignore this health indicator), measurements obtained from an unhealthy server are still recorded.

For diagnostic purposes, the system allows the contents of the packet transmitted by any individual server to be viewed. This feature is activated by clicking the "Check Selected Server" button. The decoded NTP packet is displayed in the window with a white background in the lower left corner of the display screen (Fig. 2).

Figure 2. Display screen of the SIM NTP Measurement System.

		NIST-B	NIST-G	NRC	CHU	ICE-1	ICE-2	INM-1	ONRJ-1	ONRJ-2	CENAM	CNM-2	IBMET	CNMEP
	United States		-0.9	7.6	2.8	-12.0	-8.4	-33.5	0.2	0.3	4.8	8.4	60.6	-8.7
	United States	0.9		8.5	3.8	-11.0	-7.5	-32.6	1.1	1.2	5.7	9.3	61.5	-7.8
	Canada	-7.6	-8.5		-4.7	-19.5	-16.0	-41.1	-7.4	-7.3	-2.8	0.8	53.0	-16.3
	Canada	-2.8	-3.8	4.7		-14.8	-11.3	-36.4	-2.6	-2.6	1.9	5.6	57.8	-11.6
	Costa Rica	12.0	11.0	19.5	14.8		3.5	-21.6	12.2	12.2	16.7	20.4	72.6	3.3
	Costa Rica	8.4	7.5	16.0	11.3	-3.5		-25.1	8.7	8.7	13.2	16.9	69.0	-0.3
	Colombia	33.5	32.6	41.1	36.4	21.6	25.1		33.7	33.8	38.3	41.9	94.1	24.8
	Brazil	-0.2	-1.1	7.4	2.6	-12.2	-8.7	-33.7		0.1	4.6	8.2	60.4	-8.9
	Brazil	-0.3	-1.2	7.3	2.6	-12.2	-8.7	-33.8	-0.1		4.5	8.1	60.3	-9.0
	Mexico	-4.8	-5.7	2.8	-1.9	-16.7	-13.2	-38.3	-4.6	-4.5		3.6	55.8	-13.5
	Mexico	-8.4	-9.3	-0.8	-5.6	-20.4	-16.9	-41.9	-8.2	-8.1	-3.6		52.2	-17.1
	Bolivia	-60.6	-61.5	-53.0	-57.8	-72.6	-69.0	-94.1	-60.4	-60.3	-55.8	-52.2		-69.3
	Panama	8.7	7.8	16.3	11.6	-3.3	0.3	-24.8	8.9	9.0	13.5	17.1	69.3	

This table was created at 10-24-2014 (MJD 56954) 17:24:52 UTC and will refresh every five minutes. Values are in milliseconds, time differences > 50 ms are marked in red.

Click on a server name or country name to graph the time difference between the server and SIMT(NIST). Click on a number to graph the time difference between two servers.

Figure 3. The NTP measurement grid on the SIM web site.

Every 10 minutes, the system records the average values (obtained from the last 60 measurements). These average values are processed with what we arbitrarily call the "AVG method." To attenuate the effects of asymmetry, the 10-minute data file records not only the average values, but also records the measurement result when the round trip delay had the shortest duration. This allows us to also process the data with what we arbitrarily call the "MIN method", where only one of 60 (1.67 %) of the measurements are included. The "MIN method" normally provides a lower uncertainty estimate of the difference between the server clock and UTC(NIST) than the "AVG method." However, it will not reduce the uncertainty if the source of asymmetry (for example, sustained network congestion) persists throughout the 10-minute segment

After the 10-minute data files are stored, they are transmitted via the file transfer protocol (FTP) to a server that processes the data. The results are graphically displayed on the web site of the SIM Time and Frequency Metrology Working Group (http://tf.nist.gov/sim). Here, in a similar fashion to the measurement system display (Fig. 2), the common-view time differences between the servers are displayed in a grid (Fig. 3). The column headings in the grid include the logo of the laboratory, and the row headings include an icon depicting the flag of the SIM nation. By glancing up and down a row or column of the grid, users can obtain the current time difference of a server with respect to all of the other servers. Common-view time differences larger than 50 ms are displayed in cells with a red background.

The server designations on the grid are the same as those listed in Table 2. Users can click on a server name or country name to graph the time difference between the server and UTC(NIST) as obtained through a direct comparison. Users can also click on a time difference number to graph the time difference between two servers, obtained by utilizing UTC(NIST) as a common-view reference. The graphs show the time difference and round trip delay data (10-minute samples) as processed with both the AVG and MIN methods. Up to 200 days of data can be plotted on one graph. In addition to the graphs, tabular data is displayed that includes the IP addresses and locations of the servers, the mean and range of the time differences, the frequency offset (calculated from the slope of the phase), and estimates of both frequency stability, $\sigma_y(\tau)$, and time stability, $\sigma_x(\tau)$, where τ_0 equals 10 minutes.

The following sections contain results from both direct comparison and common-view NTP measurements.

IV. RESULTS OF DIRECT COMPARISONS OF NTP SERVERS TO UTC(NIST)

The 13 servers listed in the grid (Fig. 3) were each compared to UTC(NIST) for the 100-day period ending on 10/27/2014 (MJD 56957). The time estimates for each server were obtained using both the AVG and MIN methods for intervals of one day. The average time differences did not exceed 40 ms for any of the servers, with a few servers returning average time differences of less than 1 ms (Fig. 4).

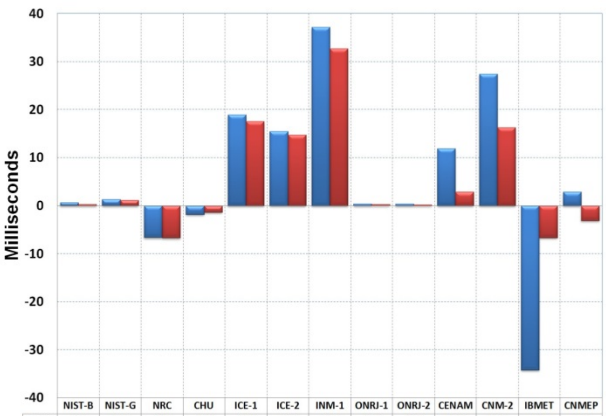

	NIST-B	NIST-G	NRC	CHU	ICE-1	ICE-2	INM-1	ONRJ-1	ONRJ-2	CENAM	CNM-2	IBMET	CNMEP
AVG	0.7	1.3	-6.6	-1.8	18.9	15.5	37.2	0.4	0.4	11.9	27.5	-34.2	2.9
MIN	0.4	1.1	-6.7	-1.4	17.5	14.7	32.8	0.3	0.3	2.9	16.3	-6.6	-3.1

Figure 4. Time differences of servers compared to UTC(NIST) using AVG (blue) and MIN (red) methods during the 100-day period ending 10/27/2014.

When operating properly, the server clocks are synchronized to UTC, at least to within ±1 ms (the synchronization sources are listed in Table 1). Therefore, most of the time error in the packet measurements is usually not attributable to the server clock, but instead to the effects of network asymmetry, which will be discussed in more detail in Section VI. The MIN method nearly always reduces the effects of asymmetry. This reduction is most apparent of the case of the server located at the Instituto Boliviano de Metrologia (IBMETRO) in Bolivia, where the average time difference with respect to UTC(NIST) was -34.2 ms when utilizing the average method, and only -6.6 ms when utilizing the MIN method.

The comparison data indicates that the effects of asymmetry are not necessarily a function of the distance travelled by the packets. Surprisingly, the two servers with the smallest time differences with respect to UTC(NIST) were located further away from Boulder than any of the others. These two servers, ONRJ-1 and ONRJ-2, are located at the National Observatory in Rio de Janeiro (ONRJ) in Brazil, about 9500 km (great circle distance) from Boulder. Figure 5 shows the one-day averages of the ONRJ-1 – UTC(NIST) time difference for the 100-day period ending on 10/27/2014 (MJD 56957). There was one sustained period on MJD 56918 where the time difference approached 10 ms. However, as indicated in Fig. 4, the average time difference for both servers was 0.4 ms when utilizing the AVG method, and 0.3 ms when utilizing the MIN method.

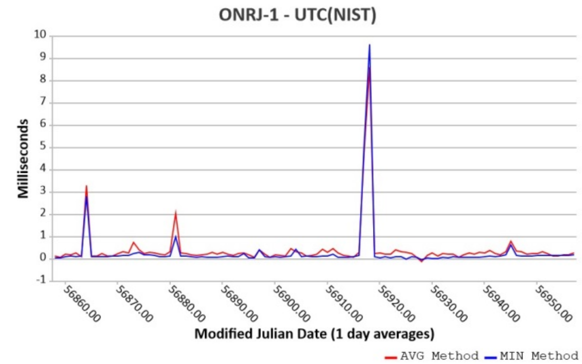

Figure 5. ONRJ-1 – UTC(NIST) over 100-day period.

V. RESULTS OF COMMON-VIEW NTP MEASUREMENTS

The common-view time transfer method allows two or more remote clocks to be compared to each other. A simple example would involve comparing two remote clocks to a common-view signal, *CVS*, that is receivable at both locations. This results in two measurements, *Clock 1 – CVS* and *Clock 2 – CVS*. By subtracting these two measurements, the time from *CVS* is cancelled, and what remains is an estimate of *Clock 1 – Clock 2*.

Normally, common-view measurements are implemented with electromagnetic signals that pass through the atmosphere [6], such as signals broadcast from a satellite. However, the SIM NTP system implements common-view by directly comparing two server clocks to UTC(NIST) and then subtracting the results of the two measurements. This technique (Fig. 6) is a convenient way to estimate the time difference between two server clocks.

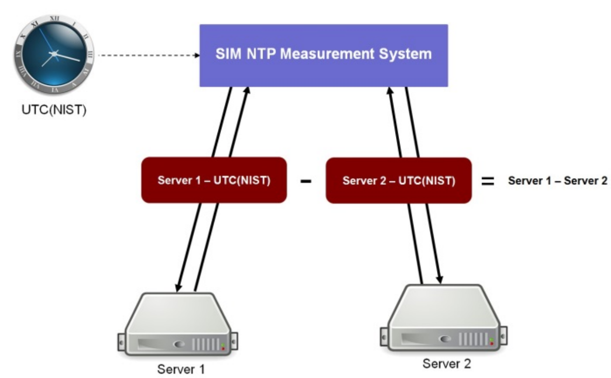

Figure 6. The common-view technique.

The apparent common-view time differences between two server clocks will still be dominated by the network asymmetries, which do not cancel when the measurements are subtracted. This is especially true in the case of servers located in different countries that are being measured via very different network paths. If we consider that the "worst case" uncertainty of a direct server clock comparison (Section IV) is equal to 1/2 the round trip delay (meaning that all of the delay is in one direction), and assume that the paths between NIST and each of the servers involved in the common-view comparison are uncorrelated (a reasonable assumption), then the "worst case" uncertainty of the common-view comparison, U_{cv}, can be estimated as

$$U_{cv} = \sqrt{\left(\frac{RT_1}{2}\right)^2 + \left(\frac{RT_2}{2}\right)^2}, \qquad (3)$$

where RT_1 is the round trip delay to server 1 and RT_2 is the round trip delay to server 2. Of course, the network asymmetry will seldom approach worst case conditions, thus the actual uncertainty will typically be much smaller.

The common-view method via UTC(NIST) is implemented in the on-line grid shown in Fig. 3. As noted previously, this allows users to quickly determine if a given server has a time offset with respect to each of the other servers in the grid. The row and column with a red background in Fig. 3 indicates that the time being received from the server at IBMETRO in Bolivia differs by more than 50 ms from each of the other servers, and could potentially be a problem.

If a server clock is not synchronized; for example if it is off by a full second due to the incorrect labelling of its 1 pps reference, the red cells that appear in the grid will immediately reveal the error. The red cells appear when the 50 ms threshold is exceeded. This threshold was chosen because it is sensitive enough to call attention to potential problems. However, while a 50 ms time difference certainly provides evidence of an asymmetric network path between NIST and a particular server, it still may not exceed the uncertainty of the common-view measurement. In other words, a red cell in the grid does not necessarily mean that a server clock is not synchronized. As a general rule, if the grid consistently reports time differences larger than about 200 ms for a given server, we can conclude that this error is too large to be caused by network asymmetry even in a region as large as the SIM region, and that it likely is due to an unsynchronized server clock. Time differences ranging from about 50 to 200 ms should be investigated to determine if the error is attributable to an unsynchronized server clock or to current network conditions.

When two servers at the same location are synchronized to the same source and connected to the same network, then most of the network asymmetry does cancel, and the common-view method can be used to accurately estimate the small time differences between servers. For example, Fig. 7 shows the results of a 100-day comparison for the period ending on 10/27/2014 (MJD 56957) between ONRJ-1 and ONRJ-2. The average time difference between the two servers was just tens of microseconds (estimated as -10 µs with the AVG method and 34 µs with the MIN method).

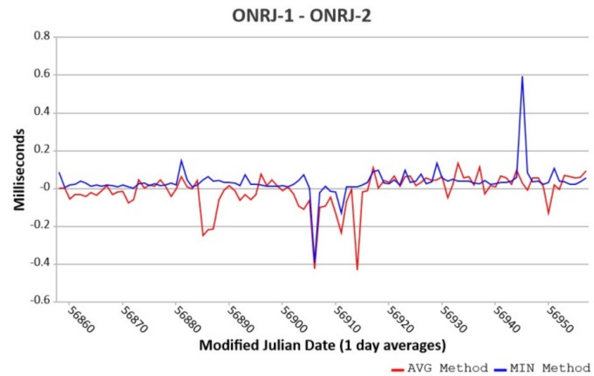

Figure 7. ONRJ-1 – ONRJ-2 over 100-day period, as obtained via common-view UTC(NIST) from Boulder, Colorado.

VI. CAUSES AND EFFECTS OF NETWORK ASYMMETRY

Network asymmetry can originate from several sources. A significant source of asymmetry is inconsistent routing of the NTP packets. As NTP message packets travel between servers and clients they traverse multiple communication routing nodes. Sophisticated routing algorithms control the lag of the packet at the router and the next leg of the route or hop. Based on traffic congestion, Ethernet links availability and routing tables, the NTP packets may be sent using different hubs at different times and when travelling in opposite directions. As a result, the delay from server-to-client can differ significantly from the delay from client-to-server.

The non-reciprocity of the paths of the NTP packets can result in significant time errors, because dividing the round trip delay by two, as in Eq. (1), does not fully compensate for the network path delay. For example, Fig. 8a shows the time differences between a server located at the National Research Council (NRC) in Ottawa, Canada and UTC(NIST), and Fig. 8b shows the corresponding round trip delay for the months of September and October 2014 (MJD 56901 to 56961).

Figures 8a and 8b indicate that the NRC – UTC(NIST) time difference is correctly indicated as near 0 ms when the round trip delay is near 40 ms. However, during periods when the round trip delay increases to ~55 ms, the time corresponding increases to about 7.5 ms, indicating that only about half of the 15 ms increase in the round trip delay is being removed.

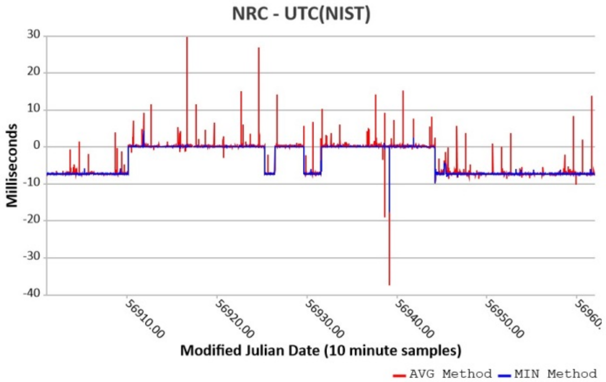

a)

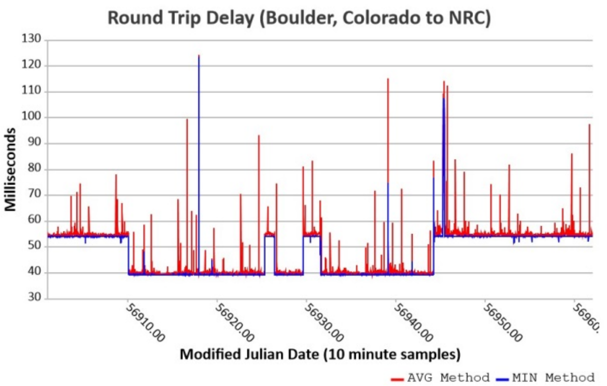

b)

Figure 8. a) Graph of time differences between NRC server and UTC(NIST). b) Graph of round trip delays between NRC server and NIST.

To investigate, we used the *tracert* utility in Microsoft Windows to analyze the link between two computers located at NRC and NIST. This utility traces the route between the origin and destination computers and reports the delay of each hop in the route. Table 2 displays the results (July 2014) of *tracert* from a NIST computer accessing a NRC computer and, inverted alongside it, the results of *tracert* from a NRC computer accessing a NIST computer.

We make several observations. First, the route is different in the two directions: the rows in the table were aligned to match the node IP (down to the subnet) used in different directions and the gaps indicate where the routes diverge. Second, the RT_{Delay} is greater for hop #10 than for hop #11 in the NRC to NIST route. This may indicate that the return routes from hops #10 and #11 are different or that there is a delay on hop #10 when it is returning the *tracert* packet, rather than forwarding it to hop #11. If the routes were static, we could evaluate the asymmetry of each link between the two computers and apply a correction to the final offset value. However, in reality network operators control the routing tables and algorithms, and end users are not allowed to request specific routes for applications such as NTP.

Table 2. The *tracert* output collected at similar times on NRC and NIST computers. The results were sorted and aligned to best demonstrate the asymmetries in route. The non-white background marks the different nodes used in the two directions.

NIST to NRC			NRC to NIST		
hop	Min RT_{Delay} (ms)	node IP	hop	Min RT_{Delay} (ms)	node IP
1	<1	132.163.136.254			
2	<1	132.163.6.1	16	*	
3	<1	132.163.3.251	15	*	
4	<1	140.172.2.33	14	*	Request timed out
5	<1	140.172.2.25	13	81	140.172.2.26
6	1	128.117.243.9	12	81	128.117.243.11
			11	**80**	137.164.26.2
			10	**92**	137.164.25.49
			9	88	207.231.245.129
			8	64	205.189.32.174
			7	53	205.189.32.182
7	23	192.43.217.222	6	44	205.189.32.178
8	76	205.189.32.98	5	36	205.189.32.180
9	81	206.130.255.11	4	5	206.130.255.13
	*	Request timed out	3	<1	132.246.0.25
	*		2	<1	132.246.0.53
	*		1	<1	132.246.52.1

We used *tracert* several times to monitor its stability and we observed changes in the number of hops and the round trip delay. For example, the results of NRC to NIST "trace routes" run in July 2014 and October 2014 are shown in Table 3. Both the number of hops and the round trip delay were significantly smaller in October than they were in July.

Table 3. The *tracert* output recorded at different times. The results were aligned to best demonstrate the differences in the routes. The non-white background marks the different nodes on the two occasions.

July 2014			October 2014		
hop	Min RT_{Delay} (ms)	node IP	hop	Min RT_{Delay} (ms)	node IP
1	<1	132.246.52.1	1	<1	132.246.52.1
2	<1	132.246.0.53	2	<1	132.246.0.53
3	<1	132.246.0.25	3	1	132.246.0.25
4	5	206.130.255.13	4	5	206.130.255.13
5	36	205.189.32.180	5	16	205.189.32.117
6	44	205.189.32.178			
7	53	205.189.32.182			
8	64	205.189.32.174			
9	88	207.231.245.129	6	34	198.71.45.2
10	92	137.164.25.49	7	31	198.71.45.8
11	80	137.164.26.2	8	54	192.43.217.223
12	81	128.117.243.11	9	54	128.117.243.11
13	81	140.172.2.26	10	53	140.172.2.26

Another common cause of asymmetry is network congestion, which occurs when the amount of traffic carried by the network exceeds the amount of available bandwidth. Periods when network congestion is worse than usual are usually identified by an increase in the round trip delay, because packets are buffered or rerouted.

To illustrate this, Table 4 shows the average round trip delay (in milliseconds) between NIST in Boulder, Colorado and each server being measured for the months of September and October 2014 (61 days). It also shows the round trip delay during that same interval with the "MIN method", and the difference between the two methods. The difference is less than 2 ms for several of the servers, indicating consistent routing of the NTP packets and sufficient network bandwidth. This was expected, for example, in the case of NIST-B, which is located in the same building as the measurement system; but again was somewhat surprising in the case of ONRJ-1 and ONRJ-2, which are located some 9500 km away. Despite the long distance, time transfer across the network between Boulder and Rio de Janeiro has remained stable. This again suggests that the number and type of network hops matters more than sheer distance.

Table 4. Round trip delays between servers under test and NIST during September-October 2014.

Server	Location	RT_{Delay} (ms)		
		Average	Minimum	Avg – Min
NIST-B	Colorado, USA	3.2	2.1	1.1
NIST-G	Maryland, USA	52.1	51.1	1.0
NRC	Canada	47.9	46.9	1.0
CHU	Canada	77.2	63.9	13.3
ICE-1	Costa Rica	133.5	126.1	7.4
ICE-2	Costa Rica	132.0	122.4	9.6
INM-1	Colombia	162.5	150.2	12.3
ONRJ-1	Brazil	188.7	187.4	1.3
ONRJ-2	Brazil	189.8	188.3	1.5
CENAM	Mexico	106.0	72.7	33.3
CNM-2	Mexico	107.3	74.7	32.6
IBMET	Bolivia	410.0	228.5	181.5
CNMEP	Panama	114.8	98.5	16.3

The most extreme example of network congestion recorded by the SIM system was seen when measuring the server at IBMETRO in Bolivia. This network congestion is indicated in Table 4, which shows that the difference between the round trip delay, as estimated with the AVG and MIN methods, was 181.5 ms during September and October 2014, much larger than the differences recorded for the other servers. During this measurement, the IBMETRO server was connected to an asymmetric digital subscriber line (ADSL) provided by a local telecommunications provider. As its name indicates, ADSL is inherently asymmetric because its downstream rate is faster than its upstream rate, thus it is not well suited for NTP time transfer. In addition, the ADSL data transfer rate is not particularly fast, limited to 2.5 Mb/s downstream and 1 Mb/s upstream, so the bandwidth can easily be saturated. During periods of bandwidth saturation, both the round trip delay and the time error will substantially increase.

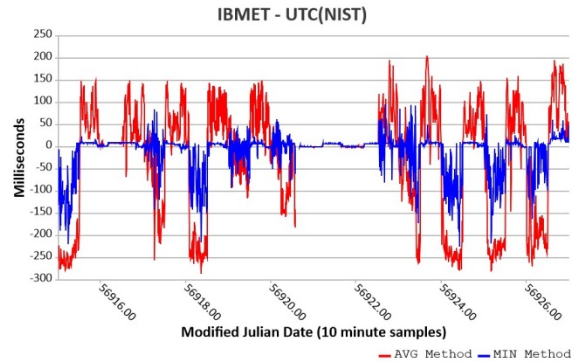

a)

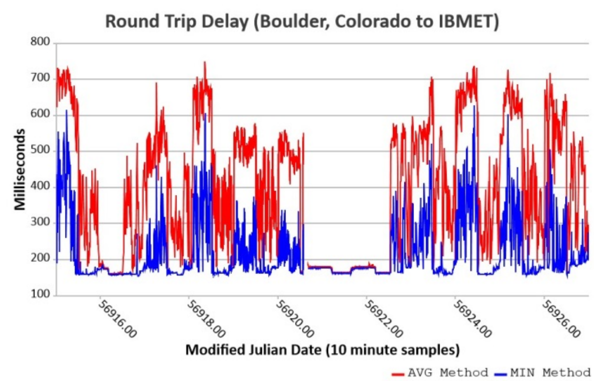

b)

Figure 9. a) Graph of time differences between IBMETRO server and UTC(NIST). b) Graph of round trip delays between IBMETRO server and NIST.

Figure 9a shows the time differences of IBMETRO's server with respect to UTC(NIST) and Fig 9b shows the corresponding variations in the round trip delay. Each graph covers the 12-day period ending on 09/26/2014 (MJD 56926). Both graphs show 10 peaks that represent two work weeks (Monday through Friday) when the ADSL connection was saturated. The quiet period between the two five-day work weeks was the two-day weekend (Saturday and Sunday). The time difference between the IBMETRO server and UTC(NIST) was less than 10 ms during the entire "quiet period" and less than 1 ms for brief intervals, but had a peak-to-peak variation during the noisy periods of nearly 500 ms, with the maximum time errors approaching -300 ms. The round trip delay using the AVG method ranged from less than 200 ms during the "quiet" periods, to more than 700 ms during periods of network congestion. Despite the noisy network, the average time difference of the IBMETRO server with respect to UTC(NIST) was just -6.6 ms for the 100-day period ending on 10/27/2014 (Fig. 4), indicating that the server clock was synchronized.

A similar, but less severe, effect of network congestion occurs in Mexico, where Internet access on the campus of the Centro Nacional de Metrología (CENAM) is obtained by a wireless microwave link from the nearby city of Querétaro. Network traffic exceeds the available bandwidth during the normal working hours at CENAM. This leads to excessive buffering of the NTP packets, causing both the network delays and the uncertainty of the time received by clients to substantially increase [7].

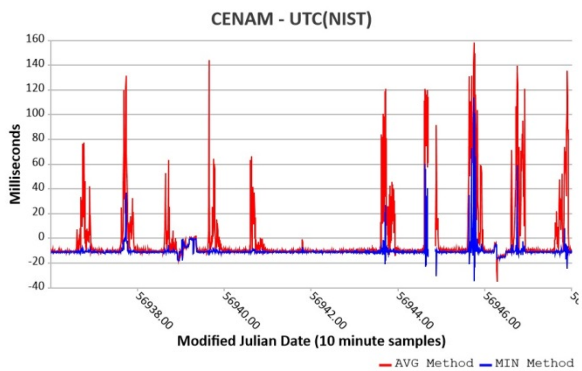

a)

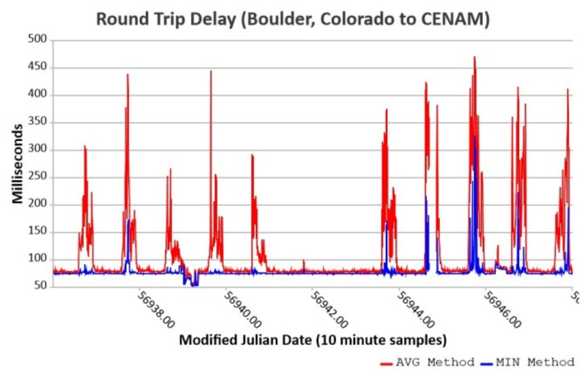

b)

Figure 10. a) Graph of time differences between CENAM server and UTC(NIST). b) Graph of round trip delays between CENAM server and NIST.

Figure 10a shows the time differences of one of CENAM's server with respect to UTC(NIST) and Fig. 10b shows the corresponding variations in the round trip delay. Each graph covers the 12-day period ending on 10/17/2014 (MJD 56947). Both graphs show 10 peaks that represent two work weeks (Monday through Friday) when the network on the CENAM campus experienced periods of congestion due to high traffic. The quiet period between the two five-day work weeks was the two-day weekend (Saturday and Sunday). The time difference between the CENAM server and UTC(NIST) was typically about 10 ms during the nighttime hours and the weekend, but exceeded 100 ms during periods of high network usage. The round trip delay (Fig. 10b) was typically 75 to 80 ms, but sometimes exceeded 400 ms during these same high traffic periods.

Notice that during Wednesday night of the first week, the time difference in Fig. 10a was near 0, during the same interval when the round trip delay (Fig. 10b) reached what might be its "true" minimum of 54 ms. This seems to imply that the network path was efficiently routed and symmetrical for a brief period, allowing the time offset of the server clock to be more accurately determined.

VII. SUMMARY

NTP is a low-cost and practical time transfer technique that allows laboratories to distribute their time to a large number of clients. In many nations, it has become the dominant or sole method for distributing the national time. Therefore, each laboratory must verify that the time they transmit via NTP is accurate, because the distribution of incorrect time is unacceptable. The SIM NTP measurement system was designed to efficiently verify the accuracy of servers that distribute the official time of North, Central, and South American nations, and to make the measurement results instantly accessible. In addition to serving as a time verification tool, the system allows us to analyze the delay asymmetry of the public Internet.

Future work could involve the expansion of this system to include other servers operated by NMIs and DIs both within and outside of the SIM region. In addition, NTP client measurement system could potentially be installed at multiple SIM laboratories. With a periodic data exchange, this would allow the unambiguous measurement of server clocks by fully compensating for network asymmetries.

ACKNOWLEDEMENTS

The authors thank Jeff Sherman and John Lowe of NIST, and Steven Sommars of Alcatel-Lucent for their review of this manuscript.

REFERENCES

[1] D. Mills, 2006, *Computer Network Time Synchronization: The Network Time Protocol*, Taylor and Francis (Boca Raton, Florida).

[2] M. Lombardi, et al., 2011, “The SIM Time Network,” *Journal of Research of the National Institute of Standards and Technology*, 116(2), 557–572.

[3] M. Lombardi, J. Lopez-Romero, and N. Diaz-Munoz, 2013, “Automated Control of National Time Standards via the SIM Time Scale,” *Proceedings of 10th International Conference on Electrical Metrology (SEMETRO)*, September 2013.

[4] J. Levine, 2011, “Timing in telecommunication networks,” *Metrologia*, vol. 48, pp. S203-S212.

[5] D. Mills, 1989, “Network Time Protocol (Version 2): Specification and Implementation,” *Network Working Group RFC-1119*, p. 51.

[6] J. Levine, 2008, “A review of time and frequency transfer methods,” *Metrologia*, vol. 45, pp. S162-S174.

[7] F. Jimenez-Tapia and J. López-Romero, 2012, "Sincronization de la Hora Oficial con Servidor NTP CRONOS del CENAM," in Spanish, *Proceedings of the 2012 Simposio de Metrologia*, Querétaro, Mexico, 6 pages, October 2012.