

Examination of Time and Frequency Control Across Wide Area Networks Using IEEE-1588v2 Unicast Transmissions

A. Novick, M. Weiss, K. Lee, D. Sutton
of National Institute of Standards and Technology (NIST)
Boulder, CO, USA
Email: novick@nist.gov

Abstract - The IEEE 1588-2008 Precision Time Protocol (PTP) version 2 (IEEE 1588v2) can be used to synchronize a slave clock to a grandmaster clock over a wide area network (WAN). However, many of the algorithms the slaves use to steer to the master are optimized for a scenario where both devices are on the same subnet or local area network (LAN). This paper is a study of existing PTP hardware from a number of different manufacturers in unicast mode. We characterize the performance of the equipment, beginning with the timing outputs of the masters that are locked to their built-in Global Positioning System (GPS) receivers. Next, we compare the results of steering unicast clients to their masters through a LAN versus several wider-area network configurations such as virtual-LANs and the public Internet. Analysis of the results will show how clients of different manufacture handle the various network paths. It is our hope that these comparisons will instigate changes to clock steering and synchronization algorithms, which may help improve the overall capabilities of PTP for telecom and other networking environments. As network synchronization techniques improve, the quality of the PTP masters will become more significant. Therefore, the performance and calibration of PTP masters with respect to UTC(NIST) is also discussed.

I. INTRODUCTION

Network timing has become of great interest over the last decade, especially since the implementation of IEEE 1588-2008, where Precision Time Protocol (PTP) was defined to enable precise synchronization of clocks in measurement and control systems through packet based networks [1]. PTP uses a two-way timing method [2] to characterize the delays between two different points across a network. By repeatedly exchanging time-stamped messages between two devices (a master and a slave), the slave clock can be synchronized to, and will ‘follow’, the master as long as data are being exchanged. The client attempts to determine the delay of the packets through the network and adjusts the timing output to compensate. The use of 1588v2 in a multicast environment is limited primarily to local area networks (LANs) where the client devices (slaves) search the network for masters and begin to exchange time stamps. Network providers may or may not allow multicast on public networks, so the unicast method should be used outside of a LAN [3]. A unicast client can find a master outside the LAN if it has the IP address of the master to look for. We employ the unicast method because we are using the public network in some cases, and we also want to guarantee that we are using a designated PTP master.

The purpose of this paper is to compare the timing synchronization abilities of PTP hardware from several manufacturers in different real-world network settings to show the capabilities and weaknesses of these devices. We were not merely trying to determine which one is the ‘best’ or to advocate particular hardware. No brand names are used. Also, no boundary clocks or transparent clocks are implemented in the path, only *ordinary clocks* (master and slave) are used, and there is no use of simulations or network traffic generators.

In order to assess the output of a remotely steered client device, the master must first have a reference time source. Typically, the masters synchronize to signals from an internal Global Positioning System (GPS) receiver, a Network Time Protocol (NTP) server or a 1 pulse-per-second (pps) input. Ideally, we would synchronize the master to the NIST timescale, UTC(NIST), and measure the output of the slave compared to UTC(NIST). However, some of the master devices do not allow a 1 pps input as the reference and many facilities may not have an on-time 1 pps available. Therefore, we looked at comparisons using the GPS receiver built in to the master as the reference. Also, the internal receivers were allowed to survey their own positions, as opposed to entering known coordinates. Later in this paper, we will discuss the time offset and stability of the GPS reference inside each master unit and how we calibrated them with UTC(NIST).

II. EXPERIMENTS

We set up a test bed for multiple PTP masters and clients, including GPS antennas on the roof, time interval counters and multiple network connections, including one outside the NIST firewall on the public network. Figure 1 shows the first measurement setup to compare the 1 pps outputs of the PTP masters to UTC(NIST) by use of a commercial time interval counter (TIC) with an external time base frequency input from UTC(NIST). The results of the measurements were recorded directly into files every second on a desktop computer with a serial connection to the TIC. We calibrated the masters for twenty-four hours and entered the average delays into the units to compensate for their offset from UTC(NIST) due to position errors, cable delays and hardware delays.




Figure 1. Comparing the PTP master (M) to UTC(NIST)

The PTP masters have internal GPS disciplined oscillators (GPSDOs) that can vary significantly from device to device [4]. Figure 2 shows a comparison of the 1 pps outputs of the four PTP masters, with peak-to-peak variations that range from 30 ns to greater than 110 ns.

Our goal was to compare multiple PTP clients over several different network scenarios, so we began with the simplest setup: a master and slave on the same subnet in the same room (Figure 3). Each client was configured to follow the master of the same manufacture, although subsequent experiments showed cross-operability. The devices were not connected together with a crossover cable but, instead, connected to the LAN via separate jacks in the wall.

The clients were configured to be in unicast mode, with a synchronization rate of 64 packets per second. Also, we allowed the clients ample time to steer to the PTP data before recording any measurements. Our interest was in how well they can do in a steady-state, as opposed to how quickly they can lock to the master. We compared the 1 pps output of each client to UTC(NIST). On the same subnet, the synchronization varied between the different PTP clients, as shown in Figure 4.




Figure 2. Comparisons of PTP masters to UTC(NIST)




Figure 3. Measuring a slave (S) on the same subnet as the master (M)




Figure 4. Comparisons of PTP clients on the same subnet

The average time offset for each client was less than 60 ns compared to its master, but the peak-to-peak variations ranged from 45 ns to over 170 ns. It is interesting that the PTP-A master had the tightest range, but the PTP-A client had the largest range. Figure 5 shows how closely each client followed its master. Different steering methods by manufacturers cover a wide range of possibilities for how closely a PTP client follows the master. Following too closely (PTP-B) causes worse stability in the network scenarios we measured, especially when, in this case, the master is the least stable. The PTP-C slave seems to over-react to changes in the master with high gain adjustments.

Next, we put the PTP masters on the public network outside the firewall at the NIST location in Boulder, Colorado, while the clients were still on the internal LAN (see Figure 6). The masters and clients were in the same building, but there were several network elements in between. These elements (referred to as “hops”) can cause noise and asymmetries in the timing of the data exchange. In this case there were five hops between the devices.

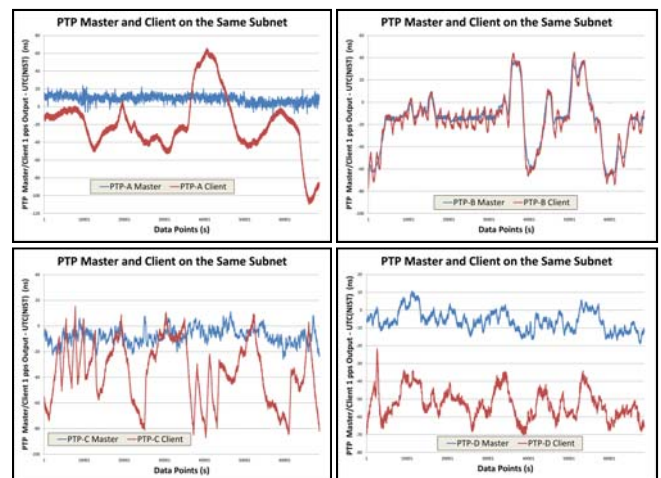


Figure 5. Graphs showing how well each client followed its master




Figure 6. The master is outside the NIST local area network

The differences between the client devices with masters outside the local network are shown in Figure 7. The average time offset of PTP-D was 2.3 μs and the peak-to-peak variation was 4.2 μs . PTP-B had a similar average, but the stability was worse, varying by 9.4 μs . PTP-C had a large average time offset of 22.9 μs , but a good stability, with a peak-to-peak range of 2.0 μs . The timing stabilities are shown in the time deviation (TDEV) plot in Figure 8. The timing requirement for some telecommunication networks is one microsecond [5,6]. None of the devices meet this criterion of accuracy in this case. However, some of the devices remain below one microsecond in stability, so if the network path were calibrated and the delay amount were entered into the client device so its output was advanced by the delay amount, meeting the requirement may be possible. However, much longer data runs would be necessary to see how the network delays change over time. Also, a reference at the client site (like GPS) would be necessary in order to calibrate the PTP client output. We measured an NTP client with a server outside the local network as well. It had an average time offset of 55.4 μs , and $\sigma_x(\tau)$ is shown to be worse than one microsecond for $\tau > 10,000$ s (Figure 8). It should be noted that NTP has a considerably slower synchronization rate and a much different intended use than PTP.




Figure 7. Performance of the clients with the master outside the LAN

The International Telecommunication Union (ITU) Option 2 (less stringent) requirement mask for a locked slave clock in telecommunication networks [7] is shown in Figure 8. The slaves did not stay under this mask for $\tau > 100$ s (best case).

We were also able to compare some of the units at remote NIST sites. Radio stations WWV and WWVB broadcast time signals from a NIST site near Ft. Collins, Colorado, which is about 80 km (direct path) from the NIST location in Boulder. There is a dedicated T1 (1.544 Mbit/s) leased line between the two sites that is considered to be virtual LAN (VLAN) inside the NIST network. We set up the client devices at the remote site (Figure 9) and compared them to 1 pps signals from a backup UTC(NIST) generated at that site, which is typically within 30 ns of the official UTC(NIST) in Boulder. There were four network hops between the master and client in this setup.




Figure 8. TDEV plot of the clients with the master outside the LAN




Figure 9. Remote PTP synchronization setup

PTP-A and PTP-D performed the best in this case (Figure 10), but each had an average time offset > 30 μs . PTP-D had the best peak-to-peak variation (< 1.4 μs). As before, if the network path were calibrated, and the outputs of the clients were adjusted to be centered around zero, this still may meet the needs of some applications, assuming the characteristics of the network do not change. The TDEV plot in Figure 11 shows the time stability of PTP-D (and maybe PTP-A) remaining below a microsecond, but none of the clients were within the ITU requirements for a locked slave [7] for $\tau > 250$ s.




Figure 10. Performance of clients with remote master




Figure 11. TDEV of clients with remote master

Next, we compared the clients when they were synchronized to the masters by use of the public Internet. The masters were at the NIST Boulder site on the public network, and the clients were at the NIST radio broadcast site, also on the public network (Figure 12). There were nine network hops between the sites.

PTP-D was the “best” with an average time offset of just under 10 ms (Figure 13) and a peak-to-peak variation of 56 μ s. The NTP measurement of a remote server using the public Internet was actually better, remaining below 3 ms but varying by more than 2 ms peak to peak. The TDEV plots (not shown) for PTP and NTP clients with a master on the public Internet do not approach the quality of the telecommunication requirements. There may be client settings on some of these PTP devices that would work better in this environment, but from our observations, PTP clients do not work well on the public Internet and should be steered after much more averaging, at least to match the average time offset of NTP. However, not only are the network paths very asymmetric to begin with, but the packets can change course

from one synchronization to the next, so calibration of the client output may not improve the timing performance.




Figure 12. PTP master and slave across the public internet




Figure 13. Clients with remote masters the public Internet

The last remote master experiment was conducted between NIST in Boulder, Colorado and NIST in Gaithersburg, Maryland, which is over 2400 km away. There are dual T3 (44.7 Mbit/s) leased lines composing the NIST VLAN, with a significant amount of traffic between the two sites. In this case there were five network hops. Because there is no equivalent to UTC(NIST) generated at the Gaithersburg site, the master there used GPS as the reference. The client was in Boulder, and its output was compared to UTC(NIST). Figure 14 shows that the average time offset was 473.0 μ s, with peak-to-peak variations of 128.7 μ s. PTP over a leased line may not have large timing steps due to packets taking different paths, but the timing variation is large, possibly caused by variable (and abundant) network traffic.




Figure 14. Results with client at a site 2400 km away

III. CALIBRATION OF THE MASTERS

Each master's GPS receiver was allowed to survey its own antenna position, but the exact locations of the antennas were well known due to geodetic surveys performed recently. Also, the antenna cable delays were measured. Table I shows the errors in latitude/longitude and height of the surveyed coordinates compared to the geodetic solution. Also, the table shows the measured antenna cable delays and the delays entered into the PTP masters to adjust the output to match UTC(NIST). The difference between these two delays is the resulting master delay, which is affected by the height survey error and the delay through the master itself. It is not known why PTP-D has a negative receiver delay, but it is possible that the manufacturer advances the output of the device to compensate for *presumed* delays. Most important, without calibration of the master, the output can be in error, which will result in client error on the other end.

TABLE I. GPS ERRORS IN THE PTP MASTERS

	Lat/Lon Survey Error	Height Survey Error	Antenna Cable Delay	Output Adjustment (based on calibration)	Resulting Master Delay
PTP-A Master	0.4 m	18.2 m	122.1 ns	163 ns	40.9 ns
PTP-B Master	9.0 m	13.3 m	123.3 ns	147 ns	23.7 ns
PTP-C Master	2.7 m	-0.9 m	108.9 ns	112 ns	3.1 ns
PTP-D Master	1.9 m	19.1 m	101.6 ns	56.0 ns	-45.6 ns

A GPS disciplined oscillator (GPSDO) is sometimes considered a self-calibrating device. However, we have shown that even when the antenna cable delay is entered, position error and receiver delays have caused average time offsets over 40 ns in some cases. Also, even though the mid-

term frequency stabilities of the masters (Figure 15) are very good because they are steered to GPS, we noted earlier (Figure 2) that there were short-term peak-to-peak fluctuations of greater than 120 ns in the worst case. Figure 16 shows all of the masters within the ITU specification for time stability of primary reference clocks [8]. At present, there is no ITU requirement for time accuracy. In order to know how well a particular PTP master works with a GPSDO as its reference, it is advisable to have it calibrated by NIST or another National Metrology Institute (NMI) or a calibration laboratory [4]. Ideally, a GPSDO calibration should be performed with the antenna and cable that will be used in the field, leaving only the coordinate error. In order to achieve the best performance with a calibrated GPSDO, precisely surveyed coordinates should be entered into the receiver.

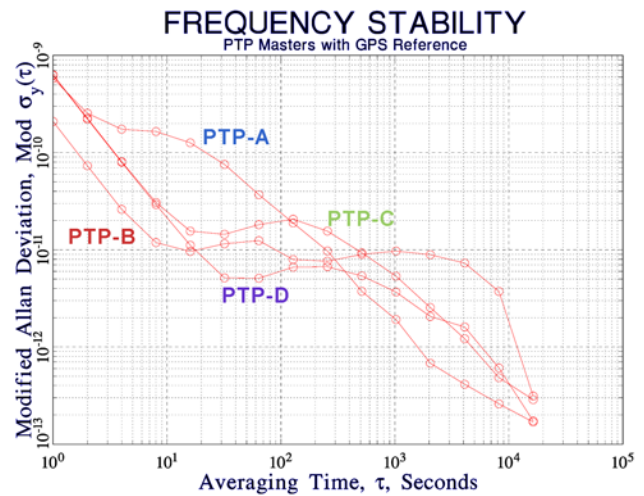


Figure 15. Frequency stability of the PTP masters




Figure 16. Time stability of the PTP masters

IV. CONCLUSION

These results show several real-world network scenarios and how different commercial PTP hardware reacts in each situation. We see that there is usually one that is worse than the other three, but it is not always the same one. As more distance and network elements are added to the network, the asymmetry of the paths can increase, causing larger time offsets. As PTP continues to grow as a tool for synchronization across networks, and the types of network situations continues to vary, we hope that providers of PTP hardware continue to develop their devices and steering methods to account for this. Also, increasing network infrastructure and opportunities for different levels of service (symmetric paths) from network providers will improve long-distance PTP synchronization. Another possibility of operation is GPS-assisted PTP, where the client would also have an embedded GPS receiver, and the PTP timing would be calibrated continually, so it could be used in a holdover mode if GPS becomes unavailable.

ACKNOWLEDGMENT

We would like to thank Robert Sorensen and Dr. Eugene Song of NIST for their help with these measurements.

REFERENCES

- [1] IEEE-1588, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems," July 2008.
- [2] J. Edison, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, A Tutorial," 2005 Conference on IEEE-1588, 10-12 October, 2005, Winterthur, Switzerland.
- [3] L. Cosart, "Precision Packet Delay Measurements Using IEEE 1588v2," on Proceedings of the 2007 International IEEE Symposium on Precision Clock Synchronization (ISPCS) for Measurement, Control and Communication, 1-3, October 2007, Vienna, Austria, pp. 85-91.
- [4] M. A. Lombardi, A. N. Novick and V. S. Zhang, "Characterizing the performance of GPS Disciplined Oscillators with Respect to UTC(NIST)," in *Proceedings of the 2005 Joint Meeting of the IEEE International Frequency Control Symposium and Precise Time and Time Interval (PTTI)*, 29-31 August 2005, Vancouver, Canada, pp. 677-684.
- [5] J. Ferrant, M. Gilson, S. Jorbert, M. Mayer, L. Montini, M. Ouellette, S. Rodrigues, S. Ruffini, "Development of the First IEEE 1588 Telecom Profile to Address Mobile Backhaul Needs," IEEE Communications Magazine, vol. 48, no. 10, October 2010, pp. 118-126.
- [6] 3rd Generation Partnership Project 2 (3GPP2), "Recommended Minimum Performance Standards for cdma2000 Spread Spectrum Base Stations: Release B," *C.S0010-B, Version 2.0*, February 2004.
- [7] ITU-T G.813, "Timing Characteristics of SDH Equipment Slave Clocks," March 2003.
- [8] ITU-T G.811, "Timing Characteristics of Primary Reference Clocks," September 1997.