TIME SYNCHRONIZING USING IEEE® 1588 + SYNCHRONOUS ETHERNET IN A "TIME-SETTING" MODE OF OPERATION

P. Stephan Bedrosian
LSI Corporation
300 Brickstone Square, Andover, MA 01810, USA
stephan.bedrosian@lsi.com

Abstract

The IEEE 1588 precision time protocol (PTP) time-setting mode of operation can be used to dramatically improve the quality of time-based systems that rely on a physical layer syntonization source in addition to PTP. Based on the standardized PTP protocol and common network frequency distribution techniques, the PTP time-setting mode of operation is capable of distributing time over packet-based facilities with little or no impact by network packet delay variation (PDV). In essence, the timing performance at a slave clock is directly related to the traceability, accuracy, and stability of physical layer packet-based systems (e.g., Synchronous Ethernet) or existing syntonization-based networks (e.g., SONET/SDH, DS1, or E1).

INTRODUCTION

IEEE 1588-2008 **[1]** has been proposed as a method for both synchronizing and syntonizing wireless base stations over packet-based networks. In the case of time division duplex (TDD) node B systems, time synchronization has been proposed as a way to achieve an operational phase alignment tolerance of ±1.25 μs **[2]** relative to UTC, as required by the application. For this case, IEEE 1588 precision timing protocol (PTP) is optioned in the two-way mode of operation where the total propagation delay between the PTP grandmaster and the PTP slave can be measured and used to compensate received timestamps.

Although IEEE 1588 is a very well-documented protocol, performance aspects relating to its use or its ability to perform deterministically in real-world packet networks are not part of the standard. Because of the nature of best-effort packet networks, PDV is a significant factor that limits the performance of IEEE 1588-based systems. Changes in delay asymmetry, common in heavily loaded packet networks, also can cause time offsets that can be problematic to some end-user services requiring absolute time. In addition, because of the lack of metrics and masks to limit PDV at packet interfaces, it is not possible to specify or enforce packet-delay behavior that is favorable to adaptive clock recovery (ACR) systems.

Another factor common to IEEE 1588 systems is the long convergence time that is required by systems to achieve frequency and time lock. IEEE 1588 sync packets are typically sent at a rate between 1 Hz to tens of Hz to a synchronizing slave. Each sync packet represents a "significant instant" that a timing recovery process can use to achieve frequency and phase lock to an IEEE 1588 synchronizing grandmaster. Because of the relatively low rate of "significant instants," several minutes (or longer) might be required to achieve stable operation and phase lock to the synchronizing grandmaster. This situation is in contrast to common TDM synchronizing signals that exchange significant instants at a much faster rate (e.g., 1.544 MHz, 2.048 MHz, or N x 51.84 MHz).

IEEE 1588 + SYNCHRONOUS ETHERNET SLAVE CLOCK

One method that has been proposed to decrease the convergence time of IEEE1588-based timing recovery systems is to use the physical layer to transport traceable frequency-based information in addition to sending IEEE1588 sync packets. Physical layer syntonization methods, such as Synchronous Ethernet **[3]**, allow the IEEE 1588 slave clock to achieve syntonization in a matter of seconds because of the high rate of significant instants. Then, the syntonized clock that is synchronized to the time of the grandmaster can be used as a stable frequency reference for the slave's local time-base. For time applications, the synchronization process also requires that the one-way delay be computed to determine the exact time at the IEEE 1588 slave clock. The one-way delay can be computed in a number of ways, including the exchange of two-way PTP messages or by reading the correction field in PTP event messages.

Because of the maturity and standardization of Synchronous Ethernet **[4]**, it has been proposed as a syntonization reference for IEEE 1588-2008 in a mode called IEEE 1588 + Synchronous Ethernet. Specific information to support this mode, including a corresponding PTP profile, is currently under development in ITU-T Question 13/15. Other physical layer syntonization methods can be used to support IEEE 1588 to decrease the convergence time. These methods include SONET/SDH/ or PDH signals common to telecommunications networks. As long as these signals are traceable to a common reference (e.g., UTC) and meet the stability requirements of the IEEE 1588 slave, they can be used in place of Synchronous Ethernet. However, for the purposes of this document, Synchronous Ethernet is used as the method for providing a physical layer-based syntonized reference.

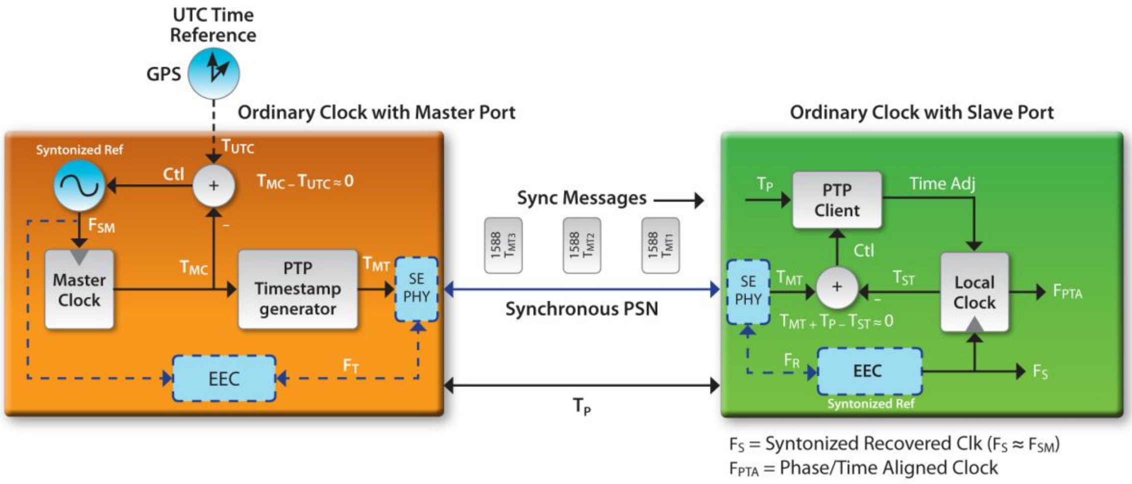

Figure 1. Symbolic example of IEEE 1588 + Synchronous Ethernet Timing System.

Figure 1 shows a symbolic example of a timing system using IEEE 1588 + Synchronous Ethernet. The PTP syntonization function is performed by the Ethernet equipment clock (EEC) located in both the PTP grandmaster node and slave node. In the PTP grandmaster, a syntonized reference (F_{SM}) is used to control the counting rate of the master clock. The frequency of the syntonized reference is adjusted such that the grandmaster clock time (T_{MC}) is equal to the UTC time (T_{UTC}) of the external GPS time reference. These components form a time-locked loop (TLL) in the grandmaster. The EEC also uses the syntonized reference as a source clock to generate the transmit clock frequency (F_T) for the Synchronous Ethernet PHY. In addition to frequency, the EEC also generates Ethernet synchronization messaging channel (ESMC) packets used to denote physical-layer traceability of the transmit clock. The PTP timestamp generator transmits timestamped PTP packets as “snapshots” of the master clock time T_{MT} at specific intervals. These timestamps are used by various PTP event messages, including sync messages as shown in Figure 1.

In the PTP slave, the Synchronous Ethernet PHY recovers the receive clock frequency (F_R) from the PSN physical layer and sends it to the EEC. The EEC also receives the ESMC messages and uses them as a basis to select the appropriate input PHY frequency. The EEC performs a phase filtering of the (F_R) clock phase and generates the syntonized reference clock (F_S) for use by the PTP slave and other node functions requiring a traceable frequency reference. These components form a time-controlled loop (TCL) in the slave. The syntonized reference is also used to establish the counting rate of the slave’s local clock. The local clock time (T_{ST}) is compared with the received master timestamps (T_{MT}) plus the master to slave propagation time (T_P) and used by the PTP client to adjust the time of the local clock such that the difference between the master time and slave time is approximately zero. Because of a number of factors related to the PTP operating mode and packet transport over the PSN, the PTP client may also need to perform one-way delay time measurement, compensation, and PDV filtering.

PTP AND SYNCHRONOUS ETHERNET TRACEABILITY

One requirement for using a physical layer (PL)-based clock and IEEE 1588 is that these clocks must be source traceable to the same time reference. Figure 2 shows how a source-traceable relationship tends to keep the maximum time interval error (MTIE) generation between these two clocks bounded. For time-based systems, this situation means that the time error is bounded between the two clocks. For example, a 1 pulse per second (pps) timing signal that is generated by a source-traceable PL-based clock and a PTP-based clock always has a bounded time/phase error relationship.

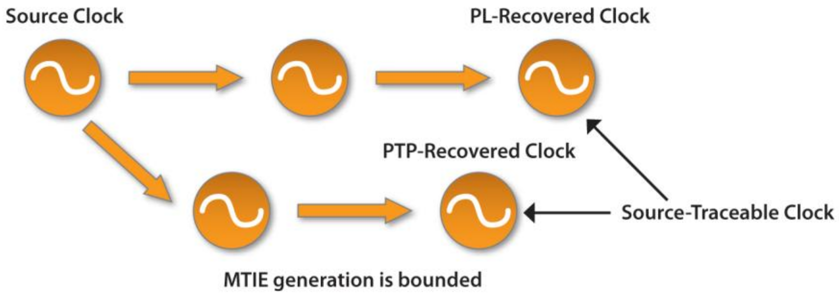

Figure 2. Source-traceable relationship between physical layer-based clocks and PTP-based clocks.

Figure 3 shows that the source-traceable relationship is in contrast to the plesiochronous relationship. In this case, the timing chains for the PL-based and PTP-based clocks originate from two different source clocks. For example, the physical-layer clock could be traceable to a PRS **[5]** reference, and PTP grandmaster could be traceable to UTC. In this case, MTIE could be generated at a maximum rate of .01 ns/sec (e.g., 0.01-ppb FFO). As long as a frequency offset exists between the PL-based source clock and the PTP grandmaster, MTIE or time error is generated and will be unbounded. Therefore, a plesiochronous timing relationship is not suitable for end-user applications requiring absolute time.

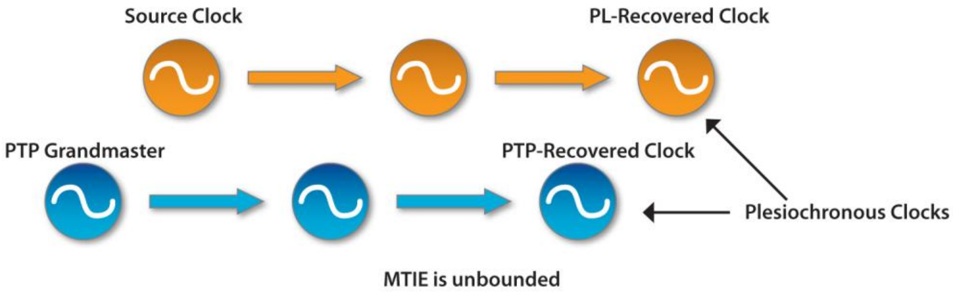

Figure 3. Plesiochronous relationship between physical layer-based and PTP-based clocks.

The current method for using the IEEE 1588 + Synchronous Ethernet is to keep both of these protocols active indefinitely. This method means that a frequency is continuously recovered from the Synchronous Ethernet physical layer and is used to establish the frequency reference for the slave's local clock. In addition, IEEE 1588 packets are continuously sent at a rate defined by the PTP profile and required by the PTP slave to support synchronization of its local clock. For time applications, PTP messages supporting a two-way mode of operation must also be sent at defined intervals to measure the one-way delay between the PTP grandmaster and slave. Although this arrangement works in some networks, the overall accuracy and stability of the clock recovered at the PTP slave is dependent on the operation of the interconnecting packet switches and background traffic conditions.

Packet delay variation adversely affects the performance of ACR systems, which includes IEEE 1588. Because the magnitude and frequency aspects of PDV are unspecified and unlimited, the IEEE 1588 slave clock must rely on a statistical analysis of a sufficient sample of received packet delays to determine the phase and time alignment of the grandmaster clock. As more PTP slaves are added to the grandmaster's domain, more PTP-related traffic is generated, which adds to the PDV problem. This problem is also compounded as packet networks grow and the background traffic increases. Any background traffic that competes with PTP-related traffic for resources in an intermediate packet switch creates PDV that can adversely affect the timing performance of PTP systems or other applications that rely on ACR methods.

Though transparent clocks (TCs) can be used to reduce the effects of PDV at the PTP slave, they must be placed at each packet switch between the grandmaster and slave. The issue of OSI layer violation involved with the modification of the correction field without updating the MAC source address exists as well as other issues related to security **[6]** of the PTP packets.

TIME-SETTING CONCEPT

The fundamental function of a PTP slave is to use the same epoch, epoch offset, and counting rate as the grandmaster. If these conditions are met, then the PTP slave is synchronized to the grandmaster. In other words, the time at the slave is the same and progresses at the same rate as time at the grandmaster. In an IEEE 1588 + Synchronous Ethernet system, the counting rate is established by the timing source carried by the Synchronous Ethernet physical layer. For our purposes, the Synchronous Ethernet timing source must be source traceable to the source of the grandmaster. In this way, the MTIE between the Synchronous Ethernet and the grandmaster is bounded. The grandmaster epoch and epoch offset are sent by the exchange of PTP event messages to the PTP slave. The accuracy and convergence time of establishing the epoch and epoch offset are influenced by the rate of exchanged event messages and network PDV characteristics.

To minimize the effects of network PDV and to decrease the convergence time of a PTP slave, you can use a newly defined operational mode called "time-setting." The time-setting mode requires that Synchronous Ethernet or other UTC-traceable physical-layer source is used as a reference by the slave's local clock. The slave is then allowed to establish a session with the grandmaster to achieve a time/phase lock with the grandmaster. After synchronization has been established, PTP event messages cease to be exchanged for an extended interval, and the slave's local clock continues to keep time using the traceable syntonized reference provided by the slave's EEC.

Figure 4 shows the series of events that occurs during startup of the time-setting mode. At time t_0, the slave's local clock becomes syntonized by using a physical-layer frequency reference that is traceable to the grandmaster reference. Note also that between time t_0 to t_1, the slave's local clock is in an "unsynchronized mode," where the time value of the local clock is not traceable to the grandmaster. At time t_1, the slave clock starts an active PTP session with the grandmaster, where event messages are exchanged. At time t_2, the slave's local clock achieves synchronized operation with the grandmaster, where it is determined that the slave time is equal to the grandmaster time. The determination of synchronous operation is accomplished by a number of methods, but generally involves a point when significant changes to the local clock time are not necessary. Synchronized operation requires that the one-way delay be computed and used to correct the time of the local clock in the slave. This computation can be accomplished through the use of either a two-way mode of operation (e.g., delay request-response or P-delay mechanism) or by using the time value of the correction field if a network of transparent clocks is used. At time t_3, the active PTP session with the grandmaster is terminated. By terminating the PTP session, event messages cease to be exchanged between the grandmaster and the slave. From time t_3 onward, the slave's local clock remains synchronized to the grandmaster by using the Synchronous Ethernet frequency reference to maintain the counting rate used to update the time offset and match the progression of time at the grandmaster.

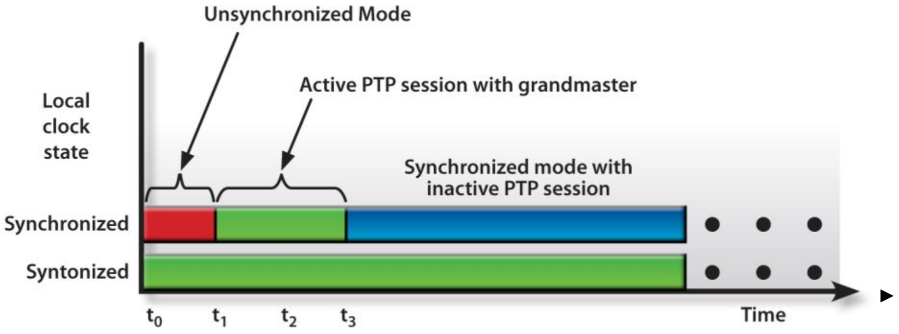

Figure 4. Startup of time-setting mode at the slave clock.

Theoretically, as long as the Synchronous Ethernet clock remains source-traceable to the grandmaster clock, no long-term time-error exists between the grandmaster and slave. Practically, however, it might be necessary or desirable for the slave to periodically re-establish an active PTP session with the grandmaster to verify its time accuracy. For this case, the slave renews the time-setting procedure with the grandmaster. This process requires the slave to compute the one-way delay time, add this value to the received sync message timestamps, and then compare this value with the slave's local time. If it is determined that the time difference between the grandmaster and slave exceeds a time accuracy threshold, a time correction of the slave's local clock could be made. This time-setting renewal procedure could be a periodic or scheduled event on the order of several hours or days, or during ideal periods of low background traffic loading and low network PDV.

Figure 5 shows the series of events that occurs during the time-setting renewal mode. At time prior to t_0, the slave clock is in a synchronized mode of operation with an inactive PTP session. For this case, it is assumed that the startup time-setting process shown in Figure 4 has already occurred. During this period, the slave's local clock remains synchronized to the grandmaster by using the traceable Synchronous Ethernet frequency reference to maintain the counting rate and match the progression of time at the grandmaster. At time t_1, the slave clock begins the time-setting renewal process by establishing an active PTP session with the grandmaster. This process involves the use of PTP general messages to begin the exchange of PTP event messages. At time t_2, the state of the network traffic load is assessed by performing an optional series of one-way delay measurements. If the analysis of these measurements exceeds a user-defined threshold, then the time-setting renewal process ceases and the slave continues to use the syntonized physical layer to maintain time. Otherwise, the synchronization of the slave's local clock is verified and adjusted, if necessary, to align with the grandmaster. This process can be accomplished, for example, by comparing the time value of received PTP sync messages with the time value of the local clock. If it is determined that the time difference exceeds the defined clock accuracy, the time of the local clock can be adjusted. Otherwise, the time at the slave's local clock remains unchanged. At time t_3, the active PTP session with the grandmaster is terminated. By terminating the PTP session, event messages cease to be exchanged between the grandmaster and the slave. From time t_3 onward, the slave's local clock remains synchronized to the grandmaster by using Synchronous Ethernet frequency reference to maintain the counting rate and to match the progression of time at the grandmaster.

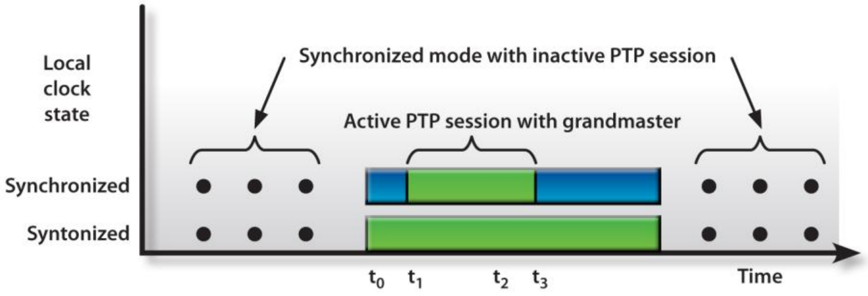

Figure 5. Time-setting renewal at the slave clock.

Between periods of time-setting renewal, slaves may exchange general messages with the grandmaster clock or boundary clock. In this way, slaves can monitor the health of a grandmaster or schedule a time-setting renewal event. In addition, PTP protection switching using the BMCA or other alternate methods of grandmaster selection can be performed by the slave without the need for exchanging PTP event messages between time-setting renewal events.

Figure 6 shows the series of events that occurs for the time-setting mode during a loss and recovery of the physical-layer syntonizing reference. At time t_0, the local clock is in a synchronized state and relies on the syntonized source to maintain time accuracy with the grandmaster. At time t_1, the slave experiences a loss of traceability of the syntonized reference. At this time, the local clock enters a “time-holdover” **[7]** mode of operation where the counting rate of the local clock is based on an atomic frequency reference (e.g., PLL in a holdover state). If Synchronous Ethernet is used, the EEC maintains a frequency-based holdover state until an acceptable input reference is made available or the failed reference is restored.

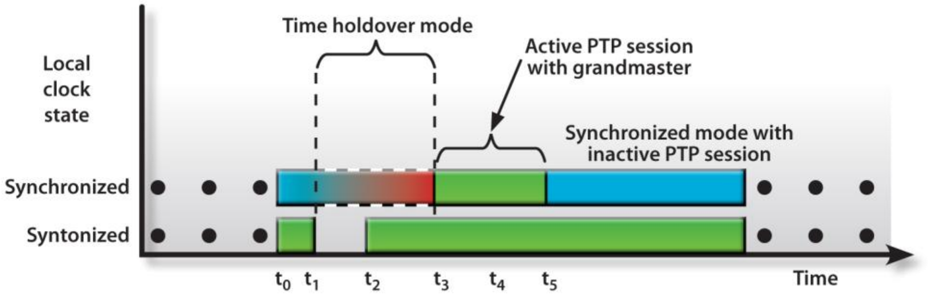

Figure 6. Loss and recovery of syntonized reference at the slave clock.

Time-holdover is a “back-up” mode of operation and results in a degraded level of time performance of the slave’s local clock. Generally, the longer a slave is in time-holdover, the greater the time uncertainty or error of the slave’s local clock time. For TDD wireless applications, time uncertainty may cause the slave’s time to exceed the ±1.25-μs phase alignment requirement relative to UTC. Therefore, the suitability of the time at the slave during time-holdover mode may be in question for some wireless applications.

At time t_2, the syntonized reference is restored such that the slave’s syntonized reference is again traceable to the grandmaster. At this point, the counting rate of the slave’s local clock is equal to the grandmaster, and time-holdover mode is exited. However, because of the time uncertainty or drift during the time holdover interval, a time-offset error may exist between the slave and grandmaster. At time t_3, the slave initiates an active PTP session with the grandmaster by exchanging general messages and event messages and starts the process of holdover recovery. During this time, the time-offset error is measured and the process of “adjusting” the time at the slave’s local clock to be synchronized to the grandmaster begins. This process can be accomplished in a number of ways, including gradually adjusting the local clock time until the local clock is again synchronized with the grandmaster. At time t_4, the local clock achieves synchronized operation with the grandmaster and exits the time holdover recovery process. At time t_5, the active PTP session with the grandmaster is terminated, and event messages cease to be exchanged between the grandmaster and slave. From time t_5 onward, the slave’s local clock remains synchronized to the grandmaster by using the syntonized reference to maintain the counting rate, thereby matching the progression of time at the grandmaster.

TIME-SETTING – USE CASES

The following section covers a variety of use case scenarios that demonstrate how the PTP time-setting mode can be used to transport time over packet-switched networks. In addition, various methods and techniques are presented for using the PTP time setting to achieve deterministic time synchronization over PTP aware and non-PTP aware packet networks.

Scheduled Time-Setting Renewal

Unlike the current method of PTP synchronization where a PTP slave must continually recover timing over a variety of changing network delay conditions, PTP time setting allows PTP slaves to be resynchronized during periods of low network PDV when the delay floor is optimum. Network PDV levels are similar to diurnal wander. This situation is typically attributed to lower traffic loads during certain times of the day or night and results in lower levels of PDV with well-defined delay floors. Thus, it is possible to schedule time-setting renewal events for PTP slaves during periods when network PDV is low. Therefore, a network operator could take advantage of these diurnal traffic patterns and schedule PTP time-setting renewal events during periods of very low network traffic.

Manual Initiation of Time-Setting Renewal

For this method, a network operator could initiate a PTP time-setting event at will. For example, given a unique path between a grandmaster and a PTP slave, it is possible for a network operator to purposely reduce or offload all background traffic to a different path during the PTP time-setting renewal event. During this time, the PTP traffic between the grandmaster and slave experience very low PDV in both the forward PTP path and the reverse PTP path. This alternate routing condition lasts for the duration of the time-setting process and allows the background traffic to revert to the original routing at its conclusion. A further benefit of this scheme is that the path between the grandmaster and slave could be made "secure" by limiting other background traffic, thus reducing the likelihood of a cyber attack.

Active Load Measurment Initiation of Time-Setting Renewal

One of the most significant challenges facing time transfer over ordinary packet networks is the ability to accurately compensate for the propagation delay between the grandmaster and the slave clock in a network of non PTP-aware packet switches. Because of relationship between background traffic and the resulting traffic delay, it is not possible to accurately predict the delay that a PTP sync message experiences. This problem is further compounded by the delays caused by reverse packet traffic at intermediate switches.

The plot in Figure 7 shows how the packet delay variation changes as a function of background traffic load. This graduated series of delay curves is based on a simulated G.8261 ten-switch network using TM-2 background traffic loads. Each plot represents the range of packet delays experienced by a periodic series of test packets. This plot shows clearly how the PDV changes from a narrow Gaussian to a long-tailed distribution as the background traffic load changes from 0% to 80%.

In many adaptive timing recovery systems, it is desirable to use packets with the lowest propagation delay as the basis for establishing the frequency and phase of the recovered slave clock. The packets with the lowest propagation delay for a specific path define the delay floor. Statistical filters are typically used to find those packet delays at or near the delay floor and then uses those delays in the slave clock timing-recovery process. The probability of receiving a packet delay at or near the delay floor is inversely proportional to background traffic load over the same path. In order to determine the accurate location of the delay floor, a statistically significant number of packet delays must be received. Thus, the convergence time of the slave clock's recovery process tends to increase with increasing background traffic load.

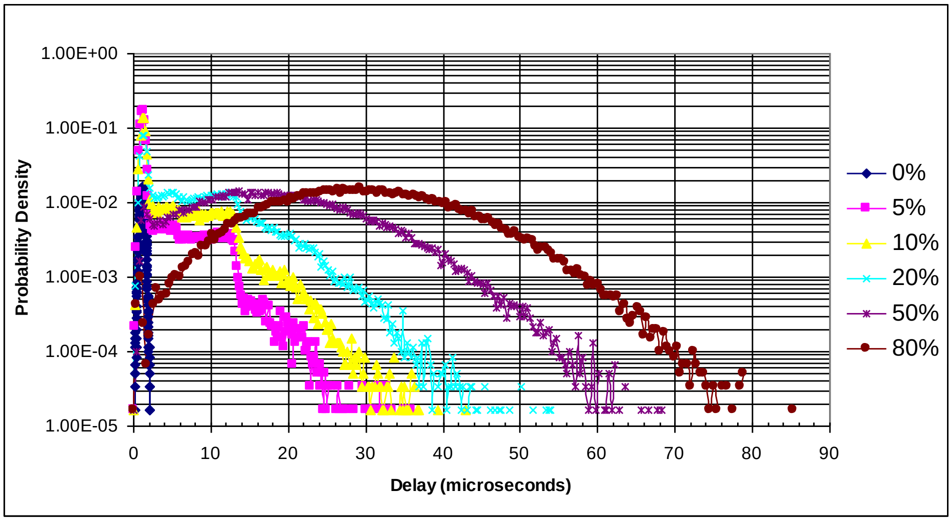

Figure 7. Packet delays as a function of background traffic load.

The minimum requirement of many time-based applications is to achieve a time accuracy of ~1 µs or less relative to UTC. The plot in Figure 8, based on the previous G.8261 simulation model, shows how the probability of receiving a packet with a delay variation of less than 1 µs lowers significantly as background traffic load increases. For ACR mechanisms, this low probability means that the windowing functions used to find the low-delay packets must increase greatly to ensure that a suitable packet within 1 µs of the delay floor is received. These delays are further compounded by the relatively low rate of PTP sync messages (typically fewer than 20 per second), as specified by the application- specific PTP profiles.

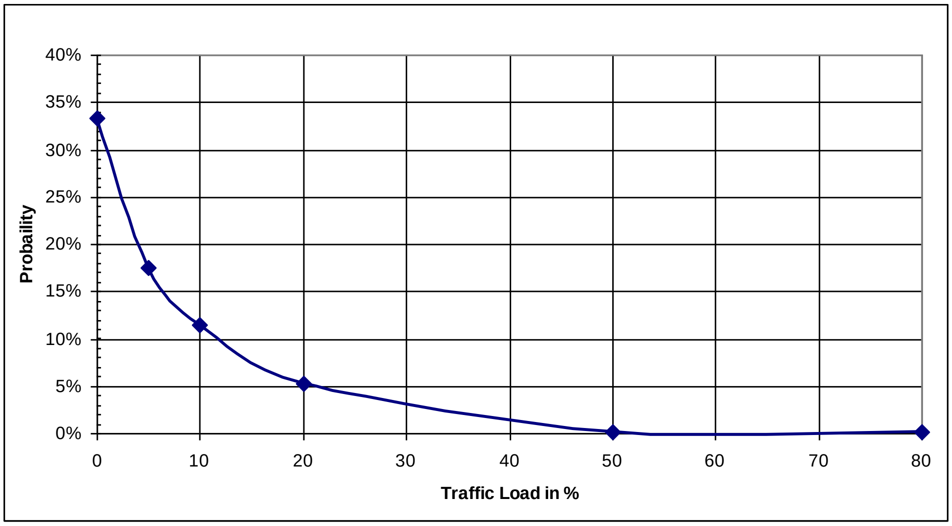

Figure 8. Probability of receiving a packet within 1 µs of the delay floor.

The need for large statistical windows used by ACR mechanisms places a phase-stability requirement on oscillators used in PTP slaves. Because of the large time intervals between updates, these oscillators must essentially hold a stable phase reference for intervals from tens of seconds to several minutes. In this state, these oscillators must employ phase-stabilizing techniques involving the control of temperature and voltage fluctuations. For example, the need for temperature stability can require the use of costly ovenized oscillators to deliver the needed phase stability at the PTP slave.

One common technique for the control or measurement of PTP is the use a network of peer-to-peer transparent clocks (P2P-TCs) or boundary clocks (BCs) to carry PTP messages. By relying on P2P-TCs to update the correction field of PTP sync messages, the propagation delay of each PTP event message can be measured and used to effectively compensate for any delay or delay variations that these sync packets experience in the PSN. A network of BCs can also be used where the delay request/response mechanism can be used effectively between BCs. Although the use of a network of P2P-TCs or BCs is an effective method of compensating for delay and delay variation of PTP event messages, this usage involves a complete build-out of PTP-aware devices in the PSN.

Because of the unknown variation in network delay, the delay request/response mechanism cannot be used reliably in a non-PTP aware network. The high variation in packet delays due to increasing traffic reduces the correlation between the actual propagation delay experienced by the sync message compared with the sampled one-way delay obtained by the delay request/response mechanism. This error is further increased by the asymmetric delay caused by reverse background traffic flows.

On the other hand, the accuracy of the delay request/response mechanism can be dramatically improved if the background traffic load approaches 0%. Because PTP is based on a system of absolute time measurement, it is possible to use various PTP delay measurement mechanisms to compute the average one-way delay and determine when low traffic load conditions exist. One characteristic that can be used to describe the PDV is the difference between the mean one-way delay values of a consecutive series of delay measurements compared with the minimum one-way delay of that series. That is to say, for a series of N received consecutive one-way delays D_1 to D_N, the mean delay value D_{MEAN} is defined as:

$$D_{MEAN} = \frac{1}{N}\sum_{X=1}^{N} D_X$$

Likewise, the minimum delay D_{MIN} of the same consecutive set is defined as:

$$D_{MIN} = \underset{X=1}{\overset{N}{MIN}}\, D_X$$

By taking the difference of these two values, the mean delay offset of the one-way delays can be computed. The mean delay offset Do_M is defined as:

$$Do_M = D_{MEAN} - D_{MIN}$$

Thus, the value of Do_M could be used as a metric to determine the level of relative measure of traffic load. To demonstrate how this metric can be used by the PTP time-setting process, consider the computed values of Do_M based on the previous G.8261 simulation model as shown in Table 1. Note the linear and proportional relationship between traffic load and Do_M. Although the accuracy of this measurement depends on gathering a significantly large sample size of consecutive packet delays, the data can be gathered while maintaining the PTP slave’s local clock with the stable physical-layer syntonizing reference.

During the time-setting renewal process, a series of delay request/response measurements could be taken without updating the PTP slave's local clock. The value of Do_M can then be computed using a predetermined sample size and compared with a predefined threshold value. This threshold value can be based on either a user-defined or default value that corresponds with a maximum allowed traffic load, as shown in Table 1. If the value of Do_M is equal to or lower than this preset threshold value, then the PTP slave's local clock is allowed to resynchronize as a part of the time-setting renewal process. If the value of Do_M is higher than the pre-set threshold value, the PTP slave's clock remains unchanged and continues to be updated using the syntonized physical layer.

Table 1. Mean-delay offset as a function of background traffic load.

Load %	**Do_M in μs**
0	1.118665487
5	2.832925888
10	4.570940347
20	8.016836889
50	18.67312531
80	30.1693796

This technique of using the Do_M is one of many possible metrics that could be used to evaluate background traffic load. Regardless of the actual metric and technique used, the overall objective is to determine when low traffic-loads exist and only allow the PTP slave's local clock to be resynchronized during those periods. The actual trigger event for a time-setting renewal process can be initiated either manually (by an operator), scheduled (to occur at a specific time), or based on periodic measurement algorithm.

DEPLOYMENT SCENARIOS

PTP systems using the time-setting mode of operation can be deployed in a variety of packet networks with various levels of PTP or physical-layer syntonization support. The following list outlines the various possible network deployment scenarios.

PTP aware and Synchronous Ethernet aware – These networks support various combinations of PTP devices (BCs or TCs) and maintain a synchronous physical layer by embedding the EEC functionality in each of these devices. The EEC source frequency must be traceable to UTC.

PTP aware with alternate syntonized physical layer support – These networks support various combinations of PTP devices (BCs or TCs) and maintain a synchronous physical by using traditional physical layer synchronized methods (SONET/SDH or PDH). The alternate synchronized source frequency must be traceable to UTC, which is beyond the plesiochronous requirements of G.811. Therefore, appropriate equipment or network modifications may be necessary to achieve the required phase stability of the end application.

Ordinary packet network and Synchronous Ethernet aware – These networks are not required to be PTP aware and can consist of ordinary packet switches. For best performance, these switches should have interconnection links at a 1-Gb/s rate or higher. The synchronized physical layer frequency is distributed by a series of EEC clocks and may or may not follow the same timing chain as the PTP path. The EEC source frequency must be traceable to UTC.

Ordinary packet network with alternate syntonized physical layer support – These networks are not required to be PTP aware and can consist of ordinary packet switches. For best performance, these switches should have interconnection links at a 1-Gb/s rate or higher. The synchronized physical layer frequency is distributed using traditional physical layer synchronized methods (SONET/SDH or PDH). The alternate synchronized source frequency must be traceable to UTC, which is beyond the plesiochronous requirements of G.811. Therefore, appropriate equipment or network modifications may be necessary to achieve the required phase stability of the end-user application.

The actual performance of each of these deployment scenarios depends on both the delay characteristics of the packet transport network, PTP transport awareness, and the ability of physical layer networks to meet the phase and time requirements of the end-user application. For the initial time-setting event, the convergence time depends on the actual background traffic load and the resulting PDV on the path between the grandmaster clock and the slave clock. Various standards organizations are studying the maximum number of packet-switching elements between the grandmaster and slave or the length of the Synchronous Ethernet timing chain to ensure compliant operation.

PTP TIME-SETTING – SUMMARY OF EXPECTED PERFORMANCE AND ADVANTAGES

Many performance advantages can be realized by using the time-setting method in contrast with the traditional method in which both IEEE 1588 + Synchronous Ethernet are fully active. A summary of the various performance characteristics and advantages follows.

Reduced cumulative packet traffic is sent from the grandmaster to a PTP slave – The overall cumulative PTP traffic sent between a grandmaster clock or a boundary clock to a slave is significantly reduced. Because any background traffic contributes to the network PDV, lower amounts of background traffic improve overall network PDV for other types of packet traffic, including circuit emulation-type traffic using ACR methods.

Increased number of PTP slaves synchronized to a single grandmaster domain – If the PTP grandmaster does not need to have a continuous session with each PTP slave, more PTP slaves can be associated with a specific grandmaster domain. In essence, the number of slaves to a grandmaster can be "oversubscribed," thus lowering the number of grandmasters needed to support a population of slaves in a domain in contrast with the dedicated PTP sessions currently required.

Increased number of available grandmasters for each PTP slave – For telecommunications applications, high-availability PTP distribution is required. Various techniques for alternate master selection, in addition to the BMCA, have been specified in IEEE 1588. Grandmasters in a time-setting mode of operation have more resources available at any given time than if they were engaged in continuous PTP sessions with dedicated slave clocks. This situation translates into better network coverage when a PTP grandmaster for a domain fails. Slaves can then switch to a different grandmaster on the same domain or different domain without transient effects on the local clock.

Achieve 5-nines (99.999%) **performance for time distribution over packet-based facilities** – It is possible to achieve a higher level of performance through the use of the PTP time-setting mode because the influence of PDV can be substantially minimized or even eliminated. Between time-setting renewal events, the PTP slave clock uses the out-of-band, but traceable, physical layer-based syntonization source that is not influenced by PDV. Therefore, the performance and reliability of the slave is directly related to the reliability of the physical-layer syntonization reference, which is typically maintained to a 5-nines availability.

Lower reliance on PDV metrics – At this time, no metrics are available to describe or limit the PDV at a packet interface. As such, there is currently no way to define, qualify, or determine if a PTP session will yield a stable time result at the slave. Because the PTP time-setting mode commonly uses a physical-layer based transport method to control the stability of the slave's local clock, it is not affected by "high" levels of PDV, unstable delay-floors, or other PDV anomalies. At such time when PDV metrics are available, they can be used to determine the best time to perform a time-setting renewal event based on the measurement of network PDV between the grandmaster and slave.

Supports a secure time transport capability – PTP security issues are based on a variety of issues surrounding the "spoofing" of PTP timestamp-bearing packets by a "man in the middle attack" to a denial of service scenario. For traditional PTP systems that rely on constant stream of PTP packets, the likelihood always exists that the PTP packets received by a PTP slave could be compromised. However, by relying on the syntonized physical layer to maintain slave timing and selectively using the PTP protocol to time-set the PTP slave, the susceptibility of these common cyber attacks can be greatly reduced if not eliminated.

CONCLUSIONS AND FUTURE WORK

The PTP time-setting mode of operation relies on two key concepts that result in the deterministic ability to transport time over packet networks. First, a syntonized physical layer that carries UTC traceable timing is used to maintain the counting rate of the PTP slave after the local clock is synchronized. By using well-established methods for controlling jitter and wander in the network, common physical layer transport technologies, including Synchronous Ethernet, can be used as enabling methods. Second, the ability to deterministically select when a PTP slave node is synchronized is very different from the common way that the protocol is used today. By only allowing the synchronization process to occur during optimum periods of low background traffic, the effects of PDV can be greatly minimized, thus achieving a faster convergence time. In addition, low background traffic conditions greatly increase the accuracy of the of the delay request/response mechanism for the measurement of one-way delay between a grandmaster and slave and the compensation of the time offset. These low PDV conditions allow PTP to accurately transport time over common-packet networks without the need for on-pass support. Though TCs or BCs can be used with the time-setting method, their use may not be mandatory in all cases.

There are many advantages of not sending PTP event messages during periods where the physical layer is used to maintain synchronization. One significant aspect of the orthogonal PTP application layer and the physical layer is the ability to tolerate cyber attacks. Relying on the secure aspects of the physical layer frequency distribution, the PTP time-setting mode has the ability to maintain time distribution during periods when packet traffic may be compromised.

Many aspects of this work must be explored and verified. The concept of using Synchronous physical layer frequency to maintain time clocks has been shown to achieve a very high level of performance under controlled laboratory conditions **[8]**. However, it is always valuable to conduct field trials with real network equipment and a variety of normal and fault conditions. Pending the outcome of this work, this PTP mode of operation should be considered for standardization by a recognized standards body with a defined PTP profile.

Another aspect that must be addressed in standards is the definition of a UTC-based frequency and time reference for use as a source of syntonizing and synchronizing signals. This device ensures that the output frequency reference that does not deviate in time or phase relative to UTC. This synchronous relationship also ensures that the long-term phase error between multiple reference devices will not drift in time or phase with respect to each other. The performance of this device contrasts with the performance of a G.811 clock, which is allowed to have up to a maximum phase error of .01 ns/sec. Therefore, new or additional specifications must be incorporated in G.8262, G.8264, or other physical layer-based syntonizing methods to support the transfer of time in a PSN.

REFERENCES

[1] IEEE Std 1588™-2008, *"IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems."*

[2] 3GPP TS 125.402 V8.1.0 (2009-07), *"3rd Generation Partnership Project; Technical Specification Group Radio Access Network; Synchronisation in UTRAN Stage 2 (Release 7)."*

[3] ITU-T G.8262, *"Timing characteristics of Synchronous Ethernet equipment slave clock (EEC)."*

[4] ITU-T G.8264, *"Packet over Transport aspects – Quality and availability targets."*

[5] ITU-T G.811, *"Timing characteristics of primary reference clocks."*

[6] A. Treytl and B. Hirschler, 2009, *"Security Flaws and Workarounds for IEEE 1588 (Transparent) Clocks,"* in Proceedings of the IEEE International Symposium for Measurement, Control, and Communication (SPCS), 12-16 October 2009, Brescia, Italy (IEEE).

[7] P. S. Bedrosian, 2007 ITU-T WD70, *"Holdover mode for Time Distribution Systems."*

[8] P. Moreira, J. Serrano, T. Wlostowski, P. Loschmidt, and G. Gaderer, 2009, *"White Rabbit: Sub-Nanosecond Timing Distribution over Ethernet,"* in Proceedings of the IEEE International Symposium for Measurement, Control, and Communication (SPCS), 12-16 October 2009, Brescia, Italy (IEEE).