Symmetricom®

NGN SERIES

Services and Applications Perspective

IEEE 1588 Precision Time Protocol – Frequency Synchronization Over Packet Networks

Overview

Telecommunications networks are rapidly shifting from circuit switched to packet switched technologies to meet exploding demand for bandwidth in both core and access networks. Traditional circuit switched TDM networks were engineered to carry precise frequency synchronization throughout the network. Access platforms such as wireless base stations and MSANs (multi service access nodes) rely on synchronization delivered over the network backhaul connection to assure high QoS for end user applications. A key dependency in the evolution to Ethernet backhaul in telecom networks is the ability to deliver carrier grade synchronization over Ethernet to remote wireless base stations and access platforms.

Time Transfer Protocols

The initial telecom applications for time transfer protocols will be to drive servo control loops for reference oscillators in remote network elements such as street cabinet access platforms and wireless base stations. Reference oscillators in these remote network elements have traditionally recovered synchronization from T1/E1 TDM backhaul connections. As long as the TDM transmission network was traceable to a primary reference clock (PRC), remote elements could employ relatively simple servo controls to lock their oscillators to the PRC traceable backhaul feed. The problem occurs when the backhaul connection transitions to Ethernet—isolating the remote element from its source of synchronization. This paper examines the use of IEEE 1588 Precision Time Protocol (PTP) over Ethernet to deliver synchronization to remote network elements.

While Ethernet has proven to be a ubiquitous and inexpensive medium for connectivity, it has not been well-suited for applications requiring precise synchronization. By nature it is nondeterministic, which creates difficulty for real-time or time sensitive applications that require synchronization. PTP overcomes the Ethernet latency and jitter issues through hardware time stamping at the physical layer of the network. The result can be an unprecedented accuracy in the 100 nanosecond range that is achieved using an Ethernet network to carry the timing packets, allowing for remarkable cost savings.

Synchronization Options for Next Generation Networks

GPS based satellite receivers provide sub 100 nanosecond accuracies, and are often used where precision time and frequency synchronization is mission critical: in telecommunication, military, and aerospace applications. But improved accuracy comes at a cost. GPS based systems require outdoor antenna installations to assure a direct view of the sky to receive the low power satellite transmissions, which are not only an added expense but which create an extra burden on the physical infrastructure of the facility. For this reason, GPS is best suited to be used in a central location as the PRC for a telecom network with other technologies utilized to distribute synchronization and timing to remote locations.

Telecommunication carriers and equipment manufacturers are examining several new methods to deliver synchronization over Ethernet networks.

- Adaptive Clock Recovery (ACR): Many proprietary solutions based on circuit emulation services (CES) utilize ACR techniques to reproduce the network clock at the far-end downstream element. ACR based solutions are seeing some initial deployment. However, major carriers are generally reluctant to base wide scale deployment of new services on proprietary solutions, and ACR is not proven to meet telecom requirements under carrier load models.

- Synchronous Ethernet: The ITU has recently completed work on definition of Synchronous Ethernet (G.8261, G.8262, G.8263) to address the need for delivery of frequency synchronization over Ethernet transmission networks. The basic difference between native Ethernet and Sync-E is the transmit PHY clock. Today IEEE 802.3 requires the transmit clock to have a free-running clock of 100ppb. In Sync-E the transmit clock must be 4.6ppb and traceable to a Stratum 1 clock via an external SSU/BITS reference or the receive clock. By simply enabling the transmit and receive clocks of Ethernet to be linked together Sync-E can be used interchangeably with SONET/SDH.

IEEE 1588 PTP delivers synchronization over Ethernet to meet telecom requirements.

- Network Time Protocol (NTP): NTP is widely adopted as the most popular protocol for time synchronization over LANs and WANs. NTP is relatively inexpensive to implement, requiring little in the way of hardware. However, NTP in its current version and implementation practices does not meet the higher precision requirements for telecom network synchronization. Symmetricom is actively engaged in IETF studies to improve NTP for future application to meet telecom network synchronization requirements.

PTP, on the other hand, offers the cost effectiveness of NTP by using existing Ethernet distribution networks, and it exceeds the NTP accuracy through the use of hardware based time stamping. PTP can coexist with normal network traffic on a standard Ethernet network using high speed switches, and yet provides synchronization accuracy to the microsecond level. With the addition of IEEE 1588 boundary clocks or transparent switches, sub microsecond synchronization accuracy is achievable. The key to this caliber of performance is hardware assisted time stamping.

How PTP Works: Hardware – Assisted Time Stamping

The two primary problems that must be overcome in network timekeeping are oscillator drift and time transfer latency. Regardless of the protocol used, oscillator drift can be mitigated by using higher quality oscillators and by deriving time from a more accurate source, such as GPS. The time transfer latency problem is more difficult and is two-fold in nature: there is latency associated with processing of time packets by the operating system, and network latency created by the routers, switches, cables and other hardware that exist between clocks. It is in the area of reducing operating system latency and jitter that PTP is most successful.

PTP combines time stamping units (TSU) with an innovative method for exchanging time stamp detail between master and slave clocks. A TSU placed between the Ethernet Media Access Control (MAC) and the Ethernet PHY transceiver sniffs both inbound and outbound traffic and issues a time stamp when the leading bits of an IEEE 1588 PTP packet are identified, precisely marking the arrival or departure of PTP time packets (See Figure 1).
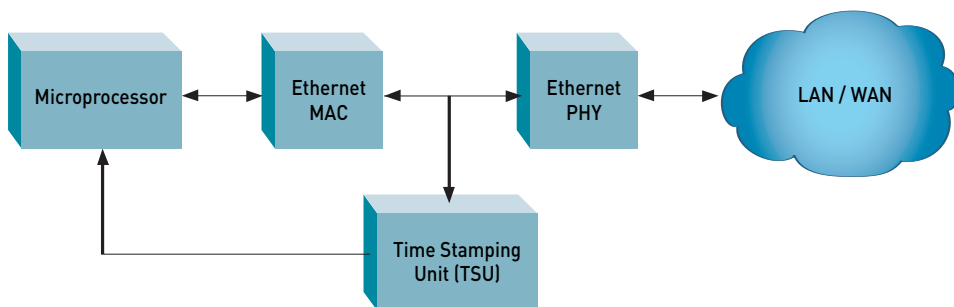
FIG. 1 Hardware set up for a system deploying IEEE 1588. The time stamp unit monitors packets at the MII layer, between the Ethernet MAC and PHY, and issues a precise time stamp when a 1588 timing packet passes through it.

In order to estimate and mitigate operating system latency, the master clock periodically sends a Sync message based on its local clock to a slave clock on the network. The TSU marks the exact time the Sync message is sent, and a Follow_Up message with the exact time information is immediately sent to the slave clock. The slave clock time stamps the arrival of the Sync message, compares the arrival time to the departure time provided in the Follow_Up and is then able to identify the amount of latency in the operating system and adjust its clock accordingly.

Network related latency is reduced by measuring the roundtrip delay between master and slave clock. The slave periodically sends a delay request message (Delay_Req) to the master clock which issues a delay response message (Delay_Resp). Since both messages are precisely time-stamped, the slave clock can combine this information with the detail from the Sync and Follow_Up messages to gauge and adjust for network induced latency. The protocol for exchanging precise time stamps is detailed in Figure 2.

Hardware assisted time stamping is the key to PTP accuracy.

Determining Target Accuracy

As detailed above, IEEE 1588 PTP is similar in concept to NTP. Both protocols distribute time packets inband with the payload traffic. NTP is ubiquitous and operates at the upper layers. PTP is a specialized layer two protocol with hardware time stamping TSUs to provide sub-microsecond accuracy. Performance over a telecom WAN is determined by three main factors:

• The resolution and accuracy of the time stamping engines in the master and slave (the accuracy you begin with).

• Latency/Packet Delay Variation (PDV) through the WAN (hop count, loading, and switch/router configuration)

• Servo processing gain and oscillator implementation at the slave side (how effectively PDV uncertainty can be filtered out)

PTP packets flow in both directions to compensate for path delay through the network.
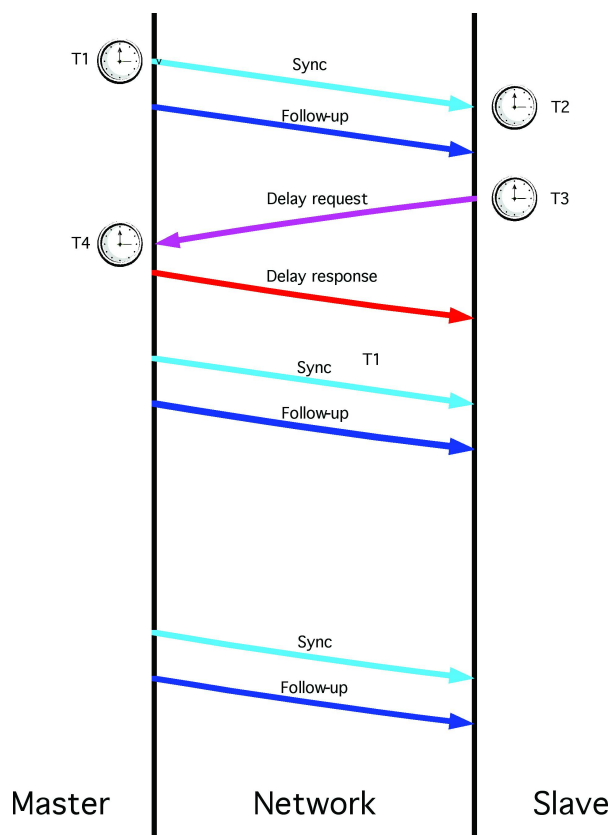
FIG. 2 The sequence of packets used to transfer time from the PTP master to a PTP slave. Sync packets are stamped as they leave the master (T1) and arrive at the slave (T2). Follow-up packets communicate the sync packet departure time to the slave. Delay response packets are stamped as they leave the slave (T3) and when they arrive at the master (T4). Sync and follow-up packet pairs are sent by the master on a periodic basis, as are the delay request and delay response packets. The formula for slave clock correction = 0.5 (T1-T2-T3+T4).

Given a high degree of starting precision, Packet Delay Variation (PDV) over telecom networks rapidly becomes the dominant source of error for packet based timing solutions. Layer two switched networks with attention to QoS provisioning and loading variations will provide the best PDV performance. This aligns well with IEEE 1588 PTP as it is optimized for layer two switching environments. Performance over layer three routed networks will show little time stamping stability difference between NTP and PTP due to the high PDV uncertainty through layer three software based routers. PDV will dominate for layer three software routed networks. Choice of oscillator stability, and servo design at the slave side will be a key performance factor to assure compliance to telecommunication network synchronization requirements.

Choosing broadcast intervals and oscillator types
In PTP, the desired accuracy of timing determines how often sync messages are broadcast and what kind of oscillator is used. More frequent broadcasts result in more accurate sync, but also in more network traffic, although the bandwidth used is extremely small. Higher quality oscillators also result in more accurate sync. It may

be tempting to try to achieve target accuracy more economically by using a lower quality oscillator while increasing broadcast frequency, but this is unadvisable. Low quality oscillators lack the stability needed to achieve high precision with PTP for telecom applications, so shortening the broadcast interval offers diminishing returns.

Accuracy is also a function of the IEEE 1588 master clock, called the grandmaster, that is the ultimate source of time on the network. Grandmasters are typically referenced to GPS so that they are both very stable and very accurate. Accuracy to UTC is typically 30 nanoseconds RMS or better. By starting with such an accurate clock with an absolute time reference, time on a PTP enabled network can be very well synchronized. A quality grandmaster also provides other measurement features to characterize the latency and jitter characteristics of network elements and to measure slave accuracy relative to the grandmaster.
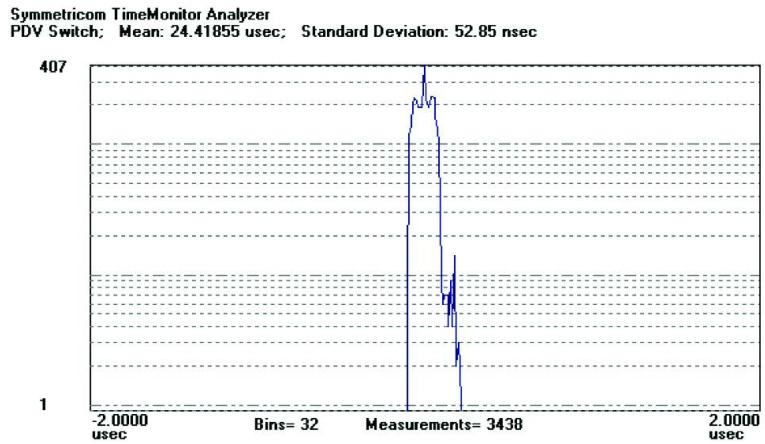
Selecting other hardware

PTP has been readily implemented to work on Ethernet networks where router buffer delay and switch latencies undermine the accuracy of time transfer. Figure 3 shows a comparison of latency and PDV measurements made on typical Ethernet switches, wire-speed routers, and software based routers. Advanced wire-speed routers can provide fast switching to rival traditional layer 2 switching in terms of latency and PDV making them very suitable for distribution of PTP synchronization. On the other hand, the high latency and PDV associated with software based routers can be a limitation as noted above.

PTP also introduces specialized elements known as boundary clocks and transparent clocks with added functionality to preserve accuracy. The boundary clock is a multi-port switch containing one port that is a PTP slave to a master clock, while the other ports are masters to downstream slave clocks. Boundary clocks provide a decent method for regulating synchronization to a number of subnets. But using cascading boundary clocks accumulates non-linear time offsets in their servo loops, resulting in unacceptable degradation of accuracy.
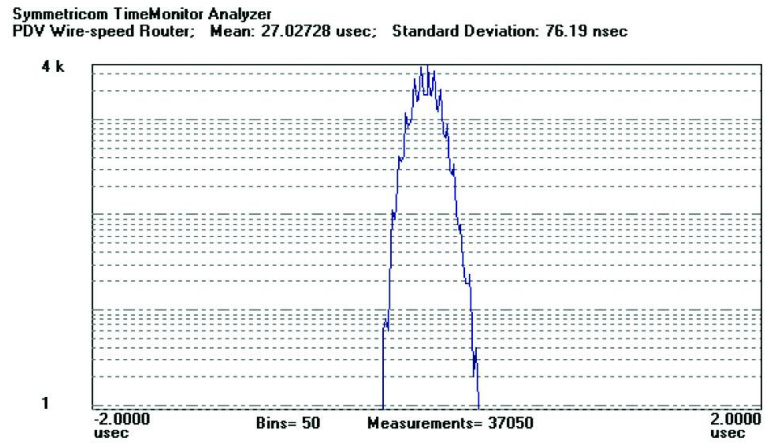
The transparent clock is another potential hardware option for the PTP-based network. This is a PTP enhanced switch which modifies the precise time stamps in the Delay_Resp and Follow-Up messages to account for receive and transmit delays within the switch itself. The result is improved sync between slave and master clocks. But the transparent clock can also create security issues when the original packet crypto checksum doesn't match the final packet arriving at the slave.

Servo control implementation and oscillator selection are key performance factors for telecom applications.
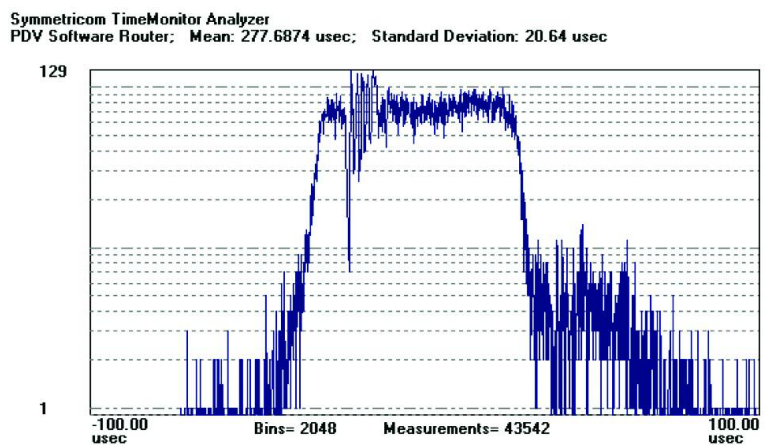
Ethernet Switch

Symmetricom TimeMonitor Analyzer
PDV Switch;   Mean: 24.41855 usec;   Standard Deviation: 52.85 nsec

Wire-speed Router

Symmetricom TimeMonitor Analyzer
PDV Wire-speed Router;   Mean: 27.02728 usec;   Standard Deviation: 76.19 nsec

PTP is optimized for layer two switching and wire-speed router based networks.

Software Based Router

Symmetricom TimeMonitor Analyzer
PDV Software Router;   Mean: 277.6874 usec;   Standard Deviation: 20.64 usec

FIG. 3 Histogram showing latency and PDV for Ethernet Switch (top), wire-speed router (middle), and software router (bottom). Wire-speed routers show comparable performance to layer 2 switches, while software routers show two orders of magnitude increased PDV.

PTP in Telecom Applications

Many telecom network equipment suppliers are targeting IEEE 1588 PTP as the most cost effective method to meet synchronization requirements for next generation wireless and access platforms. For example, all GSM and UMTS base stations must be frequency synchronized to +/- 50 ppb (parts per billion) to support handover as mobiles transition from one base station to another. Failure to meet the 50 ppb synchronization requirement will result in dropped calls. Base stations have traditionally met this requirement by locking their internal oscillators to a recovered clock from the T1/E1 TDM backhaul connection. When the backhaul transitions to Ethernet, the base station becomes isolated from its traditional network sync feed. Figure 4 shows a typical deployment scenario for a wireless network using PTP to deliver synchronization to remote base stations. The base stations will all incorporate a PTP slave device to recover timing packets used to discipline the base station's internal oscillators to meet the 50 ppb requirement. PTP slaves in the base stations rely on access to a carrier class PTP grandmaster clock deployed in the mobile switching center (MSC). Key deployment considerations include:

- Integration of PTP grandmaster functionality into existing MSC synchronization platforms (i.e. BITS—building integrated timing supplies, and SSU—synchronization supply units)

- Provisioning of Ethernet transmission elements—fast switching

- Oscillator selection and PTP slave/servo control

PTP is targeted for delivery of synchronization to next generation wireless base stations and multi-service access platforms.
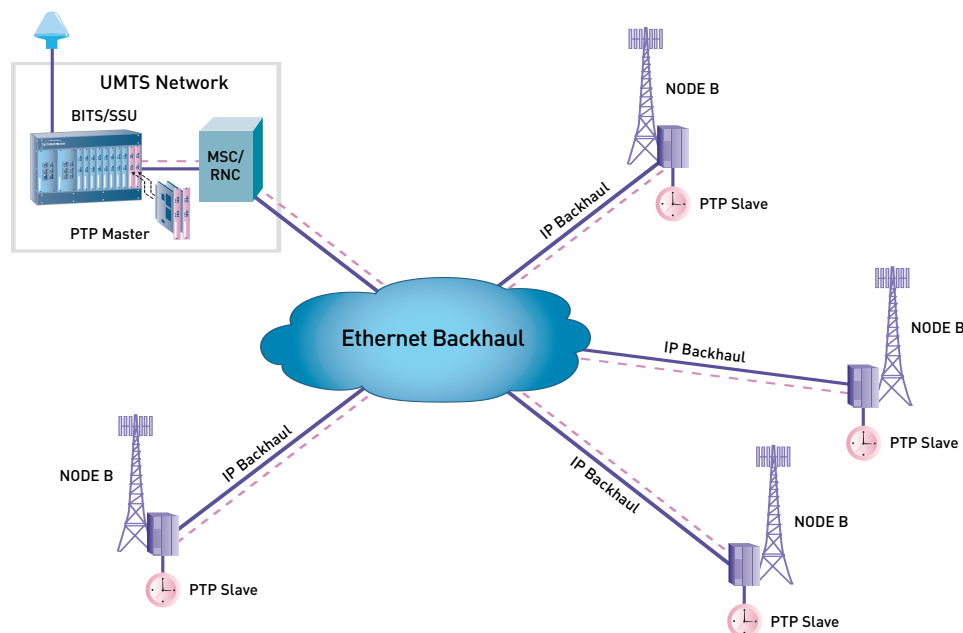
FIG. 4 Delivery of synchronization to next generation UMTS base stations will rely on PTP grandmaster clocks deployed in the MSC/RNC. Sync packets flow from the grandmaster clock to the slave clocks in the base stations.

MSANs and IP-DSLAMs also have the requirement to support legacy TDM applications such as T1/E1 drops to business parks. Equipment manufacturers are targeting PTP as a method to distribute synchronization to remote terminal based access platforms. The ITU has recently released G.8261 (formerly G.pactiming) as a standard to establish requirements for synchronization of packet based networks. The same deployment considerations apply as for wireless platforms as noted above. Key again is the integration of PTP grandmaster capability into telco central office based BITS and SSU platforms (Figure 5).
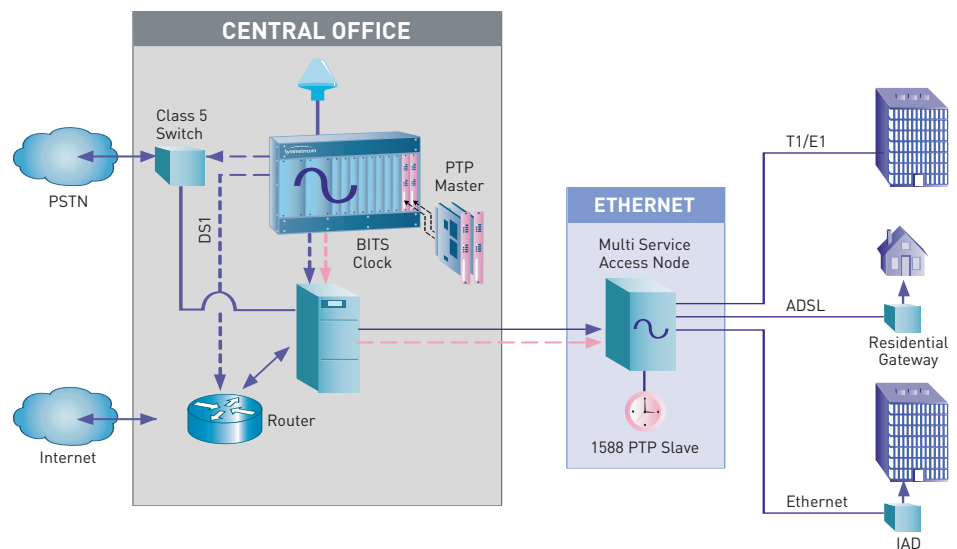
FIG. 5 MSANs supporting legacy TDM services require synchronization distribution over Ethernet backhaul from the central office. PTP slaves in the MSAN derive synchronization from the PTP grandmaster clock in the central office BITS/SSU.

The Future of PTP

PTP has justifiably received considerable attention since its introduction in 2002, and its influence is growing. A variety of silicon vendors are producing hardware that supports PTP with embedded hardware time stamping units. The next version of the 1588 protocol is currently being defined and is expected to increase accuracy even more. And improved fault tolerance and increased update rates are expected to enhance PTP performance for telecom applications. With its nanosecond accuracy, ease-of-deployment, and cost-effectiveness, PTP is poised to transform the landscape of synchronization applications in any number of fields.

To learn more about IEEE 1588, visit http://ieee1588.nist.gov/