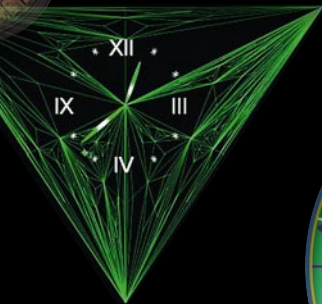


what time is it ?



The second is the standard international unit of time.



One second is the time that elapses during 9,192,631,770 cycles of the radiation produced by the transition between two levels of the cesium 133 atom..



Time Synchronization with IEEE 1588



The need for accurate time synchronization in SmartGrid applications is ubiquitous. There I said it, “SmartGrid”; I promise not to exploit this term again! The range of applications in the domain of electrical utility automation requiring time synchronization is as varied as the accuracies required. Ranging from one millisecond for sequence of event recording needed for post mortem fault analysis to sub microsecond for high precision fault location, the accuracy required from time synchronization is the critical factor that determines how to approach the problem.

THE PRECISION TIME PROTOCOL (PTP) which is embodied by the IEEE 1588 standard is poised to provide a solution that addresses the most demanding applications for time sync, yet is still cost effective to be used throughout the utility. This article will describe some of the technical aspects of PTP so the reader can make informed decisions about IEEE 1588 and how it might best serve particular needs for time synchronization.

The multitude of applications needing time synchronization was covered thoroughly in a PAC World article in the summer 2007 edition entitled “Substation Time Synchronization”. This excellent article gives a broad overview of the requirements, solutions and history and is worth reading for those new to the subject. Since the theme of this summer’s 2009 edition is IEC 61850, I will focus on two applications in this article and how PTP can provide a solution. Namely, sequence of event recording (SoE) and 61850-9-2 sampled values (SV) from a merging unit (MU).

Sequence of event recording is not unique to IEC 61850 and has been crucial for after the fact power

system analysis. Knowing exactly when the protection element picked up, when the trip was initiated and finally when the breaker opened and the fault was cleared are obviously useful for ensuring protection schemes are working as designed. ‘When’ is a key part of post analysis. Unfortunately too often only the sequence of events is known because time synchronization is lacking. One of the bigger challenges encountered after the 2003 blackout in North America was stitching together the timeline of unsynchronized SOE records.

The IEC 61850 international standard for “Communication Networks and Systems in Substations” is continuing to gain acceptance and deployments around the world. Of critical importance to merging units is high precision time synchronization in the realm of single digit microseconds. This accuracy requirement is similar for synchrophasors as defined by IEEE C37.118-2005 which demands a total vector error (TVE) less than 1% which equates to a maximum time error of 26us for a 60Hz system. Rule of thumb suggests having the time accuracy an order of magnitude

better, leading to the need for one microsecond.

Before going on to describe PTP in detail, it is important to realize that the advances in protection and control automation systems rely on modern communications networks and specifically Ethernet and TCP/IP. IEC 61850 specifies Ethernet as the communications data link and the powerful GOOSE and SV facilities depend on multicast traffic which Ethernet supports. In fact, all the major application layer protocols such as Modbus, DNP, and IEC 60870-5-104 have mappings over Ethernet/TCP/IP. PTP is no exception as it targeted Ethernet based networks from the outset.

IEEE 1588 was born to provide highly precise time sync over packet networks. To be specific, precise meant nanoseconds from the outset. One of the early highly demanding applications sought was to provide nanosecond level precision. It was targeted at the outset for such exotic applications as particle accelerators. The version 1 standard was ratified in 2002 and focused on PTP on an Ethernet LAN with a focus on automation systems. The version

Roger Moore is the Engineering Vice President of RuggedCom Inc. a leading manufacturer of industrially hardened communications technology for mission-critical applications in harsh environments. Prior to joining RuggedCom Roger was a project manager for General Electric’s Power Management division where he developed advanced protective relaying systems and substation automation technology.

Roger graduated from the University of Toronto in 1990 with a Bachelor of Applied Science degree majoring in computer science and physics.

He holds patents related to advances in communications and protective relaying technology.

PTP is Precision Time Protocol.

2 standard was ratified in 2008 and introduced new concepts to accommodate applications from a much wider set of applications - from telecom networks to Audio Video Broadcasting (AVB). Version 2 introduces such concepts as transparent clock to deal with jitter from store and forward switches; peer to peer path delay measurement; higher sync message rates; UDP protocol mapping; security; profiles and more. The new flexibility found in version 2 introduces the potential for incompatibility between devices. Hence the creation of PTP profiles that specify a particular subset of PTP capabilities.

As it is usually with most new communication protocols, PTP adoption was slow at the outset; new protocols often take upwards of ten years before reaching critical mass. PTP is quickly approaching the critical mass phase nearly one decade after inception. Annual plug-fests at ISPCS reveal that PTP does work and delivers high precision, there are many committed vendors, and interoperability is for real.

The IEEE Power Systems Relaying Committee (PSRC) task force H7, started in 2008, is working on a profile to facilitate adoption of PTP in power system automation systems. It was determined that a profile was necessary to ensure the specific needs of electric utility automation are met, and to ensure device interoperability. The results of the task force will be embodied in the new standard C37.238 which should be complete in 2010. Progress on the PTP profile has been swift to date and an initial 'plug-fest' is anticipated in September 2009 which will allow equipment vendors to test device interoperability and accuracy. Once the standard is complete and equipment vendors

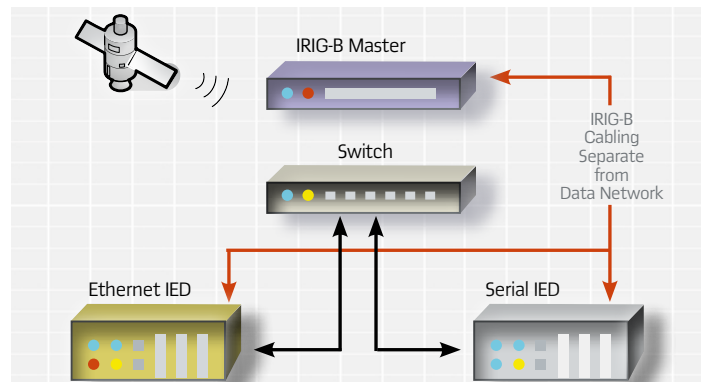
begin to release product to the marketplace, the pace of adoption should noticeably increase.

IRIG-B is today's de-facto standard for time synchronization in substations. IRIG-B is simple to understand and implement, can achieve microsecond precision, and is supported by many master clocks and IEDs. The main drawback of IRIG is the dedicated cabling infrastructure required to support it. Typically coaxial cable is daisy chained or star connected to all devices requiring time sync. The total input impedance of all IEDs must not exceed the drive capabilities of the master clock which may necessitate multiple IRIG-B circuits. IEDs and master clocks must internally have optical isolation to prevent ground loops. In contrast, PTP uses the same cabling infrastructure as the Ethernet LAN which is becoming ever more and more present in substations. Ethernet network topologies employing the Rapid Spanning Tree Protocol (RSTP) are varied and can be a star, ring, mesh, or any combination of those. There is no issue of whether impedance matching and optical isolation is already provided if fiber based Ethernet is in use. Finally, PTP should become as simple if not simpler to deploy than IRIG-B and provides additional benefits like the Best Master Clock (BMC) which

provides automated redundancy of the master clock. (Figure 1).

With the pervasiveness of GPS in society today, one must ask the question: "Why not simply use GPS in every device?" There is some merit to that statement: the cost of GPS receivers is very low and the small size allows them to be placed in even the smallest IED; no cabling infrastructure would be required and accuracy would be very high. Arguments from the past that fear the U.S. government could halt GPS service in times of emergency no longer apply; the world economy is too dependent on GPS for anyone to 'pull the plug'. The main drawback is antenna placement and cabling. For reliable timekeeping, a minimum of four GPS satellites must be in view which requires a reasonably uninterrupted sky-view. Indoor or underground applications end up requiring a cabling infrastructure that is more costly than IRIG-B! Long cable runs require high grade coaxial cable to minimize RF losses. If antennas are to be shared, expensive RF splitters are needed. As with the comparison to IRIG-B, the major benefit of PTP over GPS is the re-use of the Ethernet LAN for time distribution; there is no additional cabling required. And though the cost of GPS enabled devices keeps getting lower, the cost of PTP will ultimately be lower since the electronics are simpler.

1 Substation LAN with IRIG-B



The main drawback of IRIG-B is the need for a dedicated cabling infrastructure to support it.

IEEE PES PSRC Working group H7 has the task to define a profile of PTP for adoption in protection and automation systems.

PTP promises high accuracy time synchronization and should deliver 'plug and play' operation over the Ethernet LAN. However simple PTP may appear to the end user, underneath PTP is a complex standard requiring over three hundred pages to describe in full.

Understanding the inner workings of PTP is not a simple task though I will attempt to provide a quick overview. In simplistic terms, PTP is very much like NTP (Network Time Protocol) in that the goal is to provide time synchronization over a packet based network. PTP differs from NTP in one important aspect, in that it defines a specific 'on-time' mark for several link layer protocols (Ethernet, ProfiNet, and ControlNet) that allows for hardware assisted time stamping. The ability to trigger in hardware for time-stamping in a manner not unlike IRIG-B is what sets PTP apart from NTP and allows the high degree of precision down to the nanosecond realm.

IEEE 1588 defines four types of clocks in a PTP system: ordinary, grandmaster, boundary, and transparent. An ordinary clock, as the name implies, is the most typical and applies to all slave clocks in the system; an ordinary clock also can become a grandmaster clock. A grandmaster is the reference clock for a PTP domain; in other words, a grandmaster is the time source to which all slaves synchronize. A PTP domain contains only one grandmaster at any given time; redundancy of the grandmaster can be accomplished via the best master clock (BMC) algorithm. A boundary clock is a multi port device that bridges between PTP domains meaning that it is a slave clock on one side, and a master clock on the other side. A boundary clock may bridge between different PTP

profiles and or between different link layer technologies. For example, the boundary might be between an Ethernet LAN using multicast PTP and an IP cloud using unicast over a TDM network. Finally, a transparent clock is a multiport device, typically a packet switch that manipulates PTP messages and corrects for non-deterministic time jitter introduced by store and forward queuing switch fabrics.

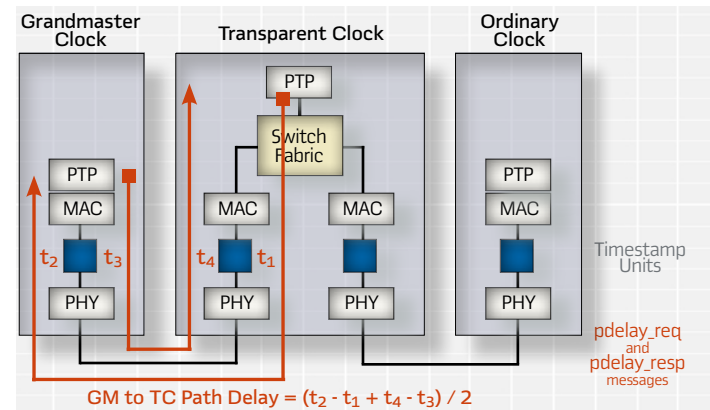
The sync message is the basic time synchronization packet for PTP and is what contains the actual time. Sync messages can be sent as either multicast or unicast messages and are a good example of why profiles are needed. A substation Ethernet LAN is more suited to multicast similar to GOOSE, whereas a telecom network is more suited to unicast sync messages. In a substation, the grandmaster will periodically multicast the sync message to all ordinary slave clocks. The time stamp in the sync message is very accurate because of the on-time mark and hardware assistance. A hardware assisted slave has the ability to extract the time stamp and update its internal clock with high precision as well. IEEE 1588 hardware assistance is placed between the PHY and MAC layer of an Ethernet device and there are several cost effective silicon solutions on the market today allowing equipment vendors to implement PTP. (Figure 2).

Hardware assistance takes one of two forms called either one-step or two-step. One step operation manipulates the sync message 'on the fly' as the packet is being transmitted; an accurate time stamp is inserted into the message and the Ethernet CRC is updated. Two step operation requires a follow up message to be sent that contains an

accurate time stamp for the previous sync message. While two step operation appears more complicated, it results in simpler silicon (i.e. less cost) and makes it possible to include security features like hashing. In order to achieve high precision time synchronization, hardware time stamped sync messages are needed, plus a measurement of how long it took the sync message to travel from master to slave.

One of the fundamental problems of time synchronization addressed by PTP is path delay - how long it takes a message to propagate from the grandmaster to the slave. Many delays are introduced on the path including: cable delay; encoding and decoding of symbols on the wire; switch fabric latency; store and forward of switches; and most troublesome, queuing from switches. Cable delay is proportional to the speed of light and is approximated as five nanoseconds per meter. Encoding and decoding symbols on the wire, often referred to as the PHY which is short for physical layer, is technology dependent; for Ethernet PHYs, an approximation of fifty nanoseconds is realistic for each side of the link. Switch fabric latency varies, though ten microseconds is a realistic estimate. Store and forward delay is proportional to the packet size and

2 Transparent clock residence correct and P2P Path Delay Measurement



symbol rate on the wire; for a one thousand byte packet on 100Mbps Ethernet, a delay of roughly eighty microseconds can be expected. Finally and most troublesome is the queuing delay from switches which can range from zero to many milliseconds. Queuing delay is problematic because unlike the other delays it is non-deterministic and can vary from packet to packet. Adding them all up yields a total delay from master that can easily exceed one millisecond for a multi hop Ethernet LAN. Obviously, correction for path delay must be done in order to achieve microsecond accuracy.

PTP corrects for path delay in one of two ways: end to end (E2E) or peer to peer (P2P). E2E path delay correction involves measuring the delay across the entire path from master to slave. The slave sends a delay_req message to the master which replies with a delay_resp message. Both messages contain hardware assisted time stamps and using simple arithmetic an estimate of the path delay can be obtained. E2E has two very critical assumptions about the underlying network: i) the delay is the same in both directions and ii) the delay is consistent from message to message. P2P path delay correction involves measuring the multiple delays between each hop between network nodes (i.e. switches and devices). In a fashion similar to E2E, each network node sends a pdelay_req message to its neighbor and receives a pdelay_resp from which the peer to peer delay can be computed. E2E path delay includes all delays from master to slave and is computed by each ordinary clock that is attempting to synchronize. P2P path delay in contrast, is computed for each 'hop' through the network and only includes the cable and PHY delays.

P2P has one key benefit over E2E delay in an Ethernet LAN: the measured path delay cannot change from message to message due to network topology change. A redundant network topology under

the control of RSTP may change the delivery path of sync messages after a network event such as a failed link. P2P delay measurement will instantly reflect the new path and so not experience any step change nor impact on slave phase locked loops; however, E2E will experience a step change in path delay that could affect accuracy while slave servo loops stabilize. For systems demanding one microsecond precision, the effect of a new path delay could result in accuracy straying out of spec during network failover events. How the path delay gets used exactly is best described in conjunction with the concept of transparent clock.

A communications network supporting PTP will have switches that support either boundary or transparent clock operation and possibly a combination of both. Boundary clocks serve as a grandmaster clock for a domain of downstream devices and may have a single port that synchronizes to an upstream master. Typically, boundary clocks will be utilized to bridge between different communications technologies or network zones. A boundary clock is the originator of sync messages for all the slaves in that time domain.

The transparent clock methodology is most suitable to an Ethernet LAN. In this approach, sync messages are forwarded by the switch from the master on towards the slave as with any Ethernet frame. However, the PTP sync is modified in transit to include a correction factor that provides the residence time of the sync message within the switch. This correction factor is vital to the accuracy of PTP in a switch Ethernet network, otherwise significant jitter in the sync message timing would result from switch queuing which is unpredictable. A transparent clock captures the arrival and departure times of a sync message, computes the difference and modifies the correction field accordingly. The ordinary slave clock then adds the

Loss of time synchronization is unacceptable in utility systems.

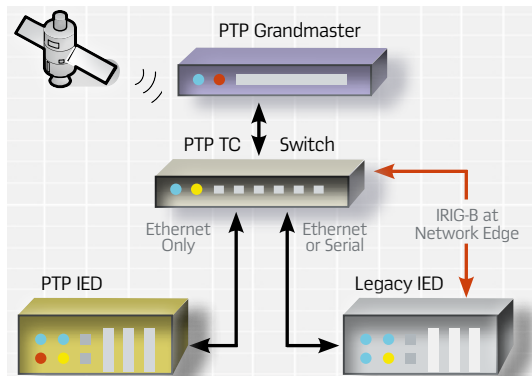
correction field to the time value in the sync message to update its clock. Further, when P2P path delay is in use, the transparent clock also adds the peer to peer path delay to the correction field. The end result is the slave receives a correction field that represents the total time it took that particular sync message to traverse the network. The combination of transparent clock and P2P path delay provides a simple and powerful way of achieving high accuracy on a switched Ethernet LAN. (Figure 3).

Loss of time synchronization is unacceptable in a utility automation system. PTP provides a mechanism called Best Master Clock (BMC) which provides an automated mechanism for choosing the best clock to be grandmaster in a domain. BMC uses a simple decision tree based on several parameters which include clock class, clock accuracy and user defined priority. If the current master fails, a second master is chosen automatically. Any device can potentially be the master which allows an islanded group of devices to continue using the best time source at their disposal.

PTP is new technology and requires long term planning and a migration strategy. Electrical utilities are known for having a cautious

PTP provides mechanisms for choosing the best clock to be grandmaster in a domain.

3 Substation LAN with IRIG-B



and conservative approach to new technology so even green-field installations may not be candidates for piloting PTP. However, the benefit of eliminating the IRIG-B cabling may be too beneficial to ignore; or the need for high accuracy synchronization where GPS poses difficulties may also be a driving force. However, early deployments will face challenges because of today's limited choice of equipment - IEDs, Ethernet switches or master clocks. Other challenges will include lack of experience and vendor interoperability.

For the near future, a hybrid substation LAN may make the most sense where the core Ethernet network supports PTP while the IEDs support traditional IRIG-B. The rationale is threefold: i) a utility will have a considerable investment in capital and expertise for a given vendor and model of IED; changing to a new vendor or model is a major decision ii) the Ethernet LAN is often a new component of substation automation so it is a new decision point iii) master clocks and Ethernet switches with PTP support seem to be more readily available at the moment than similarly equipped IEDs. Figure 4 gives an example of what such a network might look like. Time is distributed through the LAN using PTP and converted at the edge of the network to IRIG-B thus eliminating the majority of the IRIG-B cabling infrastructure.

Ethernet switches are an essential part of the migration to PTP. The majority of today's switch models do not support hardware assisted PTP. However existing Ethernet networks can still use IEEE 1588 albeit with lesser expectations for accuracy. With hardware assisted master and slave PTP devices, sub millisecond accuracy is attainable. PTP can even be used across the telecom cloud though results will vary depending on the packet delivery jitter introduced. If PTP is used across a typical IP network, the results are hard to predict and may

only marginally improve on NTP; however, PTP used over a dedicated TDM circuit with minimal jitter might show excellent results.

As part of any migration to PTP technology, a test strategy needs to be implemented. Though PTP has great promise, early adopters will inevitably experience growing pains. This is especially important when mixing PTP devices from different vendors. Any test strategy will depend on the timing accuracies being sought. One microsecond precision is usually associated with more demanding applications and it is also more difficult to verify. For example, one microsecond precision demands pulse per second (PPS) outputs and an oscilloscope to verify with confidence. Verification of one millisecond accuracy might be accomplished by simply paralleling contact inputs and examining SOE records. Factory acceptance testing (FAT) might want to stress worst case scenarios including rapid temperature fluctuations, heavy network traffic loads, network topology changes, and grand master clock failover. Field commissioning tests might want to include a simplified version of FAT testing at least until confidence and experience in PTP is gained.

Determining whether PTP is a good technology fit is also an

There is no doubt that accurate time synchronization is an essential component of grid modernization.

important part of any migration strategy. If the move to Ethernet and IP based communications has not been tackled then certainly considering PTP is probably premature. There also needs to be a certain willingness to live near the leading edge of technology. However new PTP may be, it is technology that is here to stay. Many different industries and applications are adopting it as their time synchronization system for the future. IEEE 1588 profiles are being developed for telecom networks with cellular backhaul being a prime application. The LXI consortium has already released a profile for test and measurement. Considerable activity is happening through IEEE 802.1 as AVB for broadcast quality distribution of audio and video streams. Chipsets with PTP hardware assist has been available from major silicon vendors for some time. And finally, ready to deploy products are available and more are steadily coming onto the market. (Figure 4). ■

1588 standard issues.

The migration to 1588 will take time as utilities upgrade their infrastructure with substations Ethernet networks being the starting point.

4 BMC allows local and/or remote masters

