White Paper
Assisted Partial Timing Support (APTS)

November 2018

a Microchip company

Synchronization Drives Mobile Applications

To ensure basic handover between base stations and to provide continuous high quality mobile services, the frequency and phase of radio base station clocks must be carefully synchronized. This synchronization process is specific to the radio technology used. For LTE-FDD-based mobile networks, the inter-cell frequency alignment at the air interface between neighboring base stations must be within ±50 ppb of a common reference. To meet this requirement, the frequency signal into the base station must be within ±16 ppb allowable error. LTE-TDD phase-based networks are specified with a maximum of ±1.5 µs of time error between the radio interfaces and the maximum allowable end-to-end time error from UTC (the globally specified reference clock) to the eNB is ±1.1 µs.[1] This time error budget includes reference clock inaccuracies and random network delays due to transport node or link noise—all of which can cause network asymmetry. The transport network is allocated ±1 µs of the total time error allowable. Transport networks, however, are heterogenous and dynamic; they evolve according to changes in the technologies used, demographics, and usage patterns. This adds a further layer of complexity when designing the clocking architecture because a synchronization plan for a modern mobile network has to be both tightly engineered and flexible.

Synchronization Architectures

Frequency-based synchronization networks using physical layer time signals are traditionally architectured as a center-weighted hierarchical system. A centralized source clock generates a frequency that is propagated hop-by-hop over the transport network elements to the end application, such as FDD base stations.

Over the past decade, mobile networks have evolved from TDM to packet backhaul and replaced physical layer synchronization with systems carrying a timing signal using Precision Time Protocol (PTP) at the packet layer. The first wave of PTP deployments were for FDD applications and PTP has now been successfully implemented with PTP grandmaster clocks such as the Microsemi TimeProvider 5000 deployed in hundreds of mobile networks worldwide.

Increasingly, the adoption of advanced network services for LTE (LTE-A) and planning for 5G services is driving next generation mobile networks using phase-based applications deployed at the mobile edge. There is consequently a migration from grandmaster clocks engineered for frequency delivery to primary reference time clocks (PRTC, defined in ITU-T Rec.G.8272), that require a GNSS input and use phase-specific PTP profiles. The network architectures for these phase-based applications are subtly different from those developed for frequency. In particular, it is recommended (ITU-T G.8275) that PRTC clocks be distributed closer to the edge of the network. These distributed clocking systems are then backed up by high accuracy core PRTC/ePRTC that can generate and hold time for extended periods.

The following two figures show the evolution of mobile network architectures.

Typical FDD Macro Mobile Network Showing Physical and Packet Layer Services

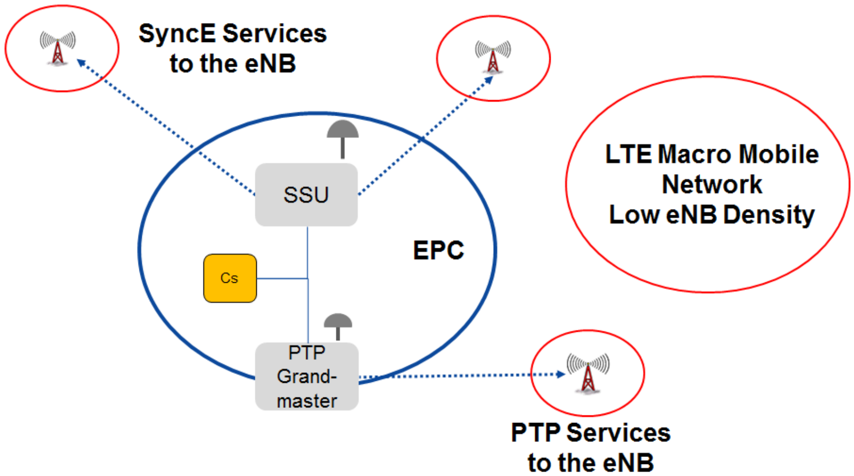

Typical LTE-A TDD Mobile Network Showing Distributed PRTC at Mobile Edge

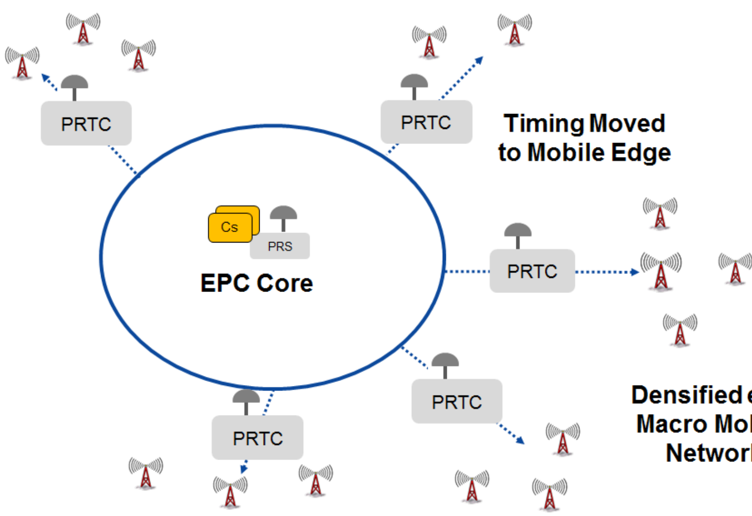

Synchronization Options for the Mobile Edge in Phase Networks

G.8265.1 grandmaster clocks, designed for frequency services using PTP, are often deployed at the RAN aggregation point, several hops from the eNB. Frequency transfer has some inherent elasticity that enables propagation over an asynchronous network with confidence as long as well-established engineering guidelines are followed.[2]

The delivery of phase services is engineered according to the time error budget limits imposed by the 3GPP (for radio interfaces) and ITU-T for the network interfaces and reference clocks. However, whereas the delivery of frequency using PTP is well understood, the same is not necessarily true of the transfer of phase timing using PTP. Sending a timecode across an asynchronous packet network with inherent noise and delay to deliver synchronization within ±1.1 µs time error relative to UTC can be a significant challenge.

There are three ways to solve this problem:

Solution A: GNSS. The operator can deploy GNSS at every eNB.
Caveats: Every eNB must be populated with GNSS, and the GNSS antenna must have continual line of sight to the satellite signal. LoS is not always possible because the view of the satellite can be impaired (both attenuated and delayed by reflections); for example, by vegetation, by shadows caused by high rise buildings (urban canyon), or because the eNB is deployed underground or indoors. Additionally, GNSS is susceptible to both intentional and unintentional jamming and spoofing. Ubiquitous GNSS can also be costly from an OPEX perspective.

Solution B: Embedded Time Boundary Clocks (T-BC). For this architecture, the transport network must be engineered with a hardware based de-jitter function known as time boundary clock (T-BC) embedded in every NE.
Caveats: The T-BC hardware and software must be deployed on every transport node on the clock chain, which often requires an onerous network investment cycle. Even when deployed on every NE, the BC does not necessarily guarantee that the timing signal will be within the required specification unless the network is carefully engineered to ensure that there is no hop-to-hop asymmetry on the links in addition to the permitted asymmetry from constant timing error allocated for each T-BC.

Solution C: Distributed PRTC. Lightweight PRTC can be moved to the edge of the network to reduce the hop count between the clock and the eNB such that phase based timing using PTP can reach the eNB within the recommended ±1.1 µs time error limits.
Caveats: Requires investment in lightweight clocks deployed around the edge of the network—a new distributed timing architecture.

Of these three solutions, locating the PRTC closer to the eNB may enable cost reduction compared to deploying T-BC hardware on every NE or installing GNSS at every cell site. Cost will be an increasingly important factor when planning for densification of eNB for LTE-A and 5G services.

With Recommendation G.8275, the ITU-T recognized that the stringent time error timing requirements at the eNB makes it difficult to deploy centralized PRTC clocks and simultaneously guarantee the viability of the phase signal to the end application. Moving the PRTC closer to the end application reduces the probability that noise and asymmetry from the network transport will negatively impact the PTP flow,[3] but also has an impact on the form-factor and capacity requirements of the PRTC.

At the core of the network where extremely accurate time and extensive holdover are required, the clocking infrastructure can include high performance, high capacity ePRTC with multiple Rubidium and ePRC Cesium devices that are not appropriate for deployment at the edge of the network. Distributed edge PRTC on the other hand can be much smaller and much lower cost.

ITU-T Recommendation G.8275—PRTC Deployed at the Network Edge

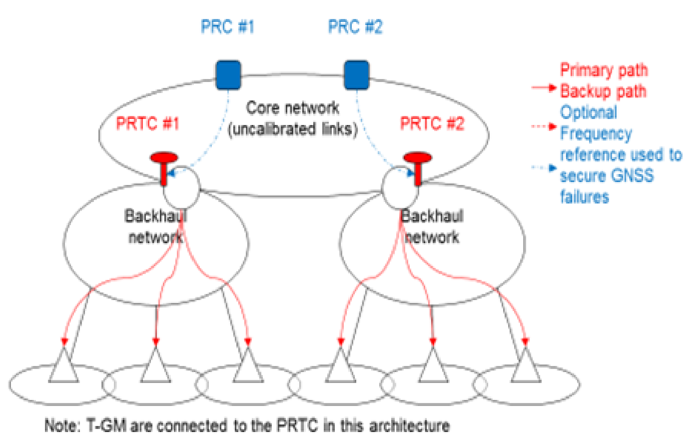

However, small PRTC distributed at the edge of the network as self-contained systems without a timing connection to the core, are isolated from the upstream centralized clocks. This can be a problem for continued operation if the device loses accurate GNSS connectivity as the lower cost oscillators used in such small PRTC will not be able to provide extensive holdover at ±100 ns level of accuracy. Holding ±100 ns for extended periods of time is the domain of high performance oscillators, not of the low cost OCXO or TCXO typically found in edge devices. Once a GNSS input is lost, then PRTC populated with such oscillators will quickly drift outside the ±100 ns specification. Also, since each unit is independent, there will be no coordination in timing between neighboring cell sites as discussed shortly. This is shown in the following two diagrams.

PTP Output Using GNSS—No APTS

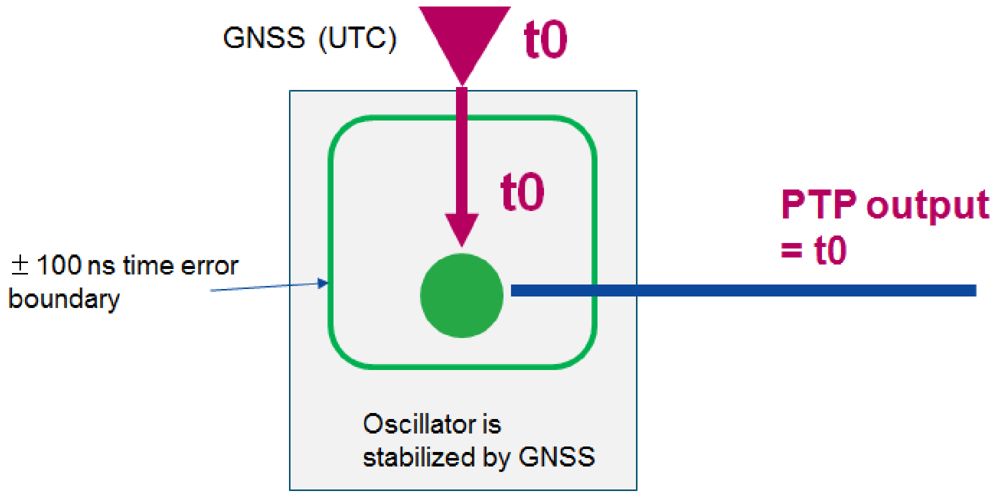

PTP output is driven by the GNSS input.

GNSS Is Lost—PTP Output Wanders

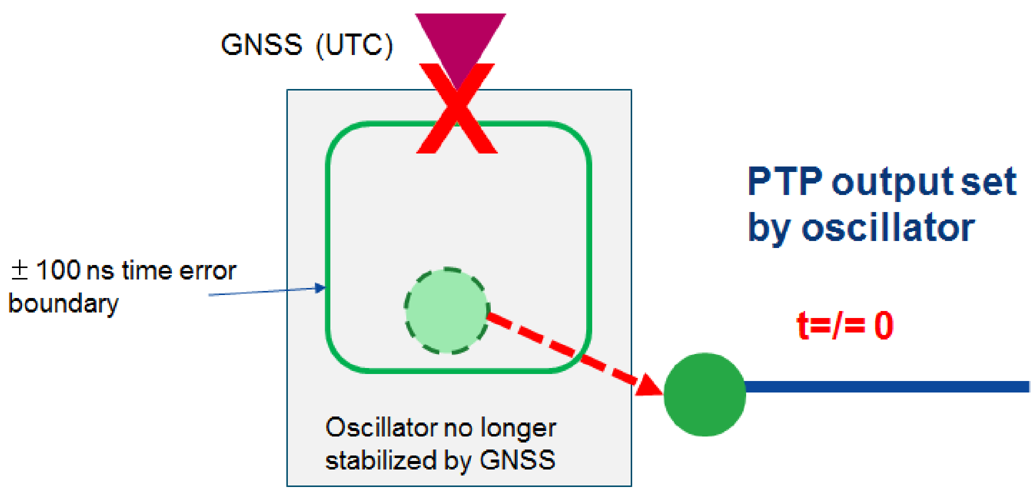

If the oscillator wanders, the PTP output quickly loses the time reference.

In normal circumstances, once the GNSS is lost (as shown previously), the PRTC immediately signals the loss of GNSS connectivity to the attached clients. This has ramifications for the eNB. In some client implementations, as soon as the PRTC signals GNSS connectivity is lost (by sending a clockClass7 flag, for example), the client will immediately disqualify the PTP input flow and go into holdover based on the internal oscillator in the radio device.

In this situation, if the oscillator in the eNB is populated with a low-cost oscillator, it will not be able to remain within ±1.1 µs of UTC for more than a few minutes. All eNB that disqualify the incoming PTP signal will drift independently. They will rapidly wander apart because the oscillators in each eNB will react differently to the individual environmental constraints and the speed, direction, and stability of the accumulating time error will be different for each eNB. Moreover, these radios will continue to generate RF and this will contribute to increasing and less controlled interference for other active eNB in the vicinity from the same or other operators.

Assisted Partial Timing Support

To avoid a situation where the edge PRTC is isolated or the event that a GNSS failure can no longer provide phase services, Microsemi developed the idea of connecting the edge PRTC to the centralized core clocks using a PTP flow. This idea was adopted by the ITU-T and consented as Recommendation G.8273.4—Assisted Partial Timing Support.

In this architecture, the incoming PTP flow is time stamped by the GNSS used by the core PRTC. The PTP flow from the core PRTC to the edge PRTC is configured as a unicast protocol, G.8265.1 or G.8275.2. The PTP input is calibrated for time error using the local edge PRTC GNSS. This GNSS has the same reference (UTC) as the upstream GNSS. The incoming PTP flow can be considered as effectively a proxy GNSS signal from the core with traceability to UTC.

If the edge system GNSS now goes out of service for any reason, the edge PRTC can fall back onto the incoming calibrated PTP flow as the timing reference and continue to generate outbound PTP time stamps that are aligned with GNSS.

This is shown more clearly in the following illustration.

PTP APTS Flows as a Backup for Edge PRTC

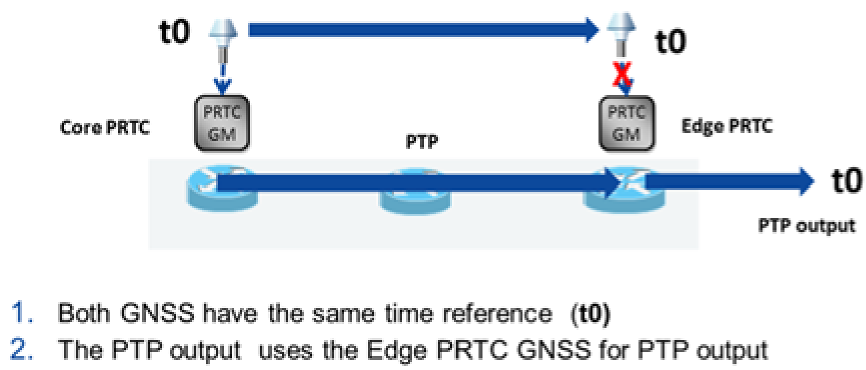

The ITU-T formal statement of the G.8273.4 architecture is shown in the following illustration.

ITU-T G.8273.4 Assisted Partial Timing Support Architecture

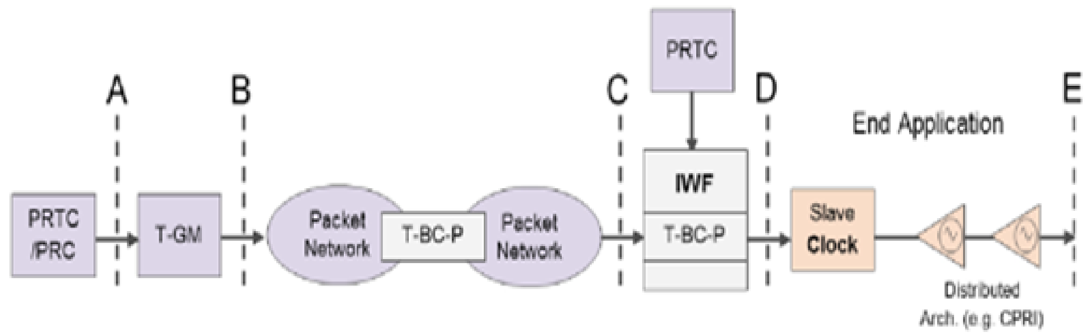

APTS Operation

APTS operation is quite simple in concept.

- Both the core PRTC and edge PRTC have a GNSS input referenced to UTC time.
- The core PRTC T-GM delivers PTP time stamps to the downstream edge PRTC/GM clock using a PTP profile.
- The PTP flow is configured as a G.8265.1 or G.8275.2 unicast profile.

Note: There is ongoing work in the ITU T to also use the G.8275.1 multicast profile.

- The edge PRTC compares the PTP time stamp to the local GNSS time.
- The edge PRTC accumulates information about the PTP flow from the PTP time stamps and from message exchanges with the Core PRTC. It thus understands the overall delay and time error on that specific PTP path.
- The edge PRTC adjusts the PTP flow information and calibrates the incoming PTP flow by compensating for the measured time error so that it is now equivalent to the local GNSS time.

This process is shown in the following illustration.

APTS G.8273.4: One PTP Input Is Calibrated for TE

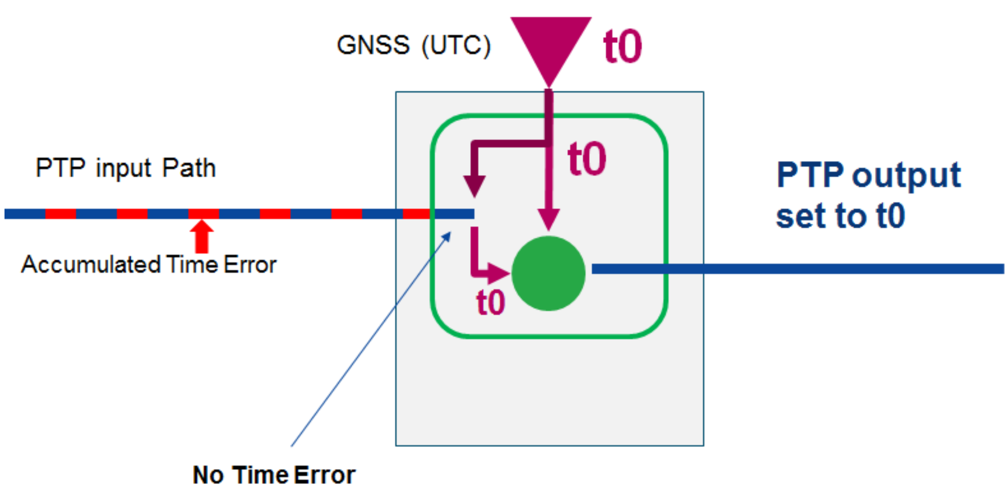

With APTS, the time error on a single PTP input path is removed.

Once the APTS algorithm is operating, the incoming PTP flow can be used as a proxy for the GNSS. If the GNSS on the local PRTC is lost, then the system will use the calibrated incoming APTS flow as the reference clock, as shown in the following illustration.

APTS G.8273.4. GNSS Is Lost: The Calibrated PTP Input Can Be Used to Maintain the Reference Time

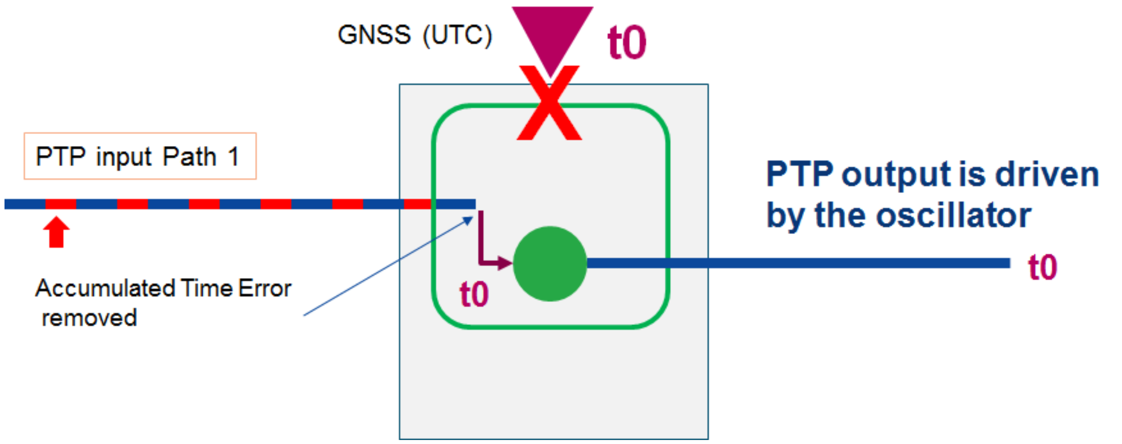

When GPS is lost, the PTP output stays locked to t0. Eventually, the oscillator will drift out away from t0.

Even with APTS, however, if the GNSS stays disconnected, then eventually the system oscillator will drift away from the ± 100 ns requirement. The time this takes will depend on the quality of the oscillator, environmental factors, the strength of the APTS algorithm, and especially on the stability of the incoming PTP path.

One of the weaknesses of the standard APTS implementation (G.8273.4) is that if the PTP path should be re-routed while the GNSS is inactive, the system will not have knowledge of the time error on the new path. In addition, to re-routing other events, such as change in traffic, load in the switches can induce additional asymmetry that mimics a classic re-route.

In other words, the standards based APTS is not resilient to a network re-arrangement that affects the incoming PTP flow. Modern OTN- or MPLS-based core networks can be very dynamic with network intermittent re-arrangement of the network paths. This can clearly be a problem for PTP flows that are optimized across a single static path.

Engineering Resiliency—Protection Against PTP Input Path Rearrangement

An end-to-end PTP system can be made more resilient by configuring more than one PTP path into the edge PRTC. However, the G.8273.4 recommendation only provides for frequency correction for any alternative PTP inputs that are used, any other PTP path is not corrected for time error.

While the frequency input can help stabilize the edge PRTC oscillator, it clearly cannot be considered as a true representation of the upstream PRTC clock which requires reference to UTC. Without implementing a time error correction on the second PTP input flow, the system is vulnerable to the dynamic network changes typical of a modern routed network. As the network re-arranges the PTP path, the system will lose the time error compensation ability and as a result the PRTC will move more quickly away from the ± 100 ns limit than it will with a PTP flow that is well calibrated for time error.

This is shown in the following two illustrations.

G.8273.4: The Second PTP Flow Is Frequency Only

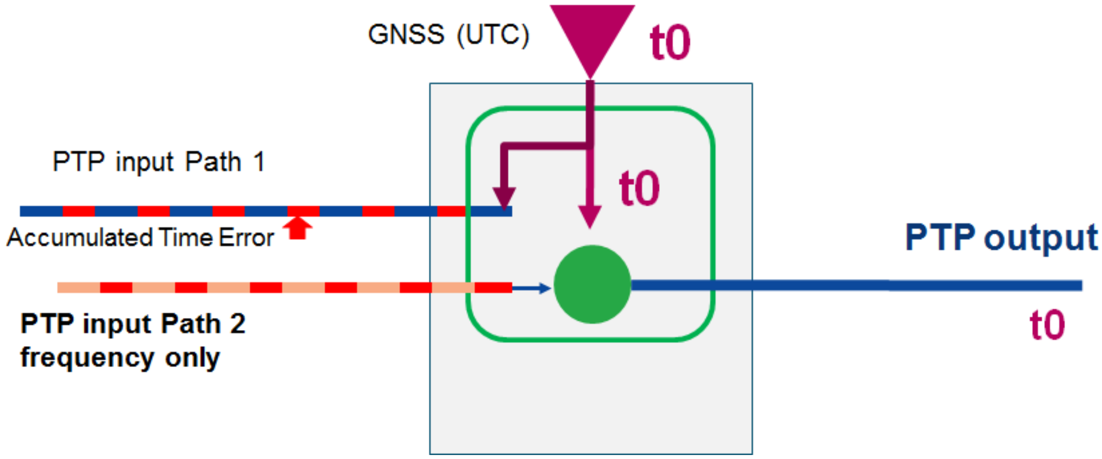

In G.8273.4, only one PTP inputs is calibrated for time error. All other PTP inputs are calibrated for frequency only.

If the PTP input path changes, then APTS cannot operate and the system falls back to frequency stability.

A Purely Frequency Disciplined Oscillator Will Drift Away Quickly from Accepted PRTC TE Limit

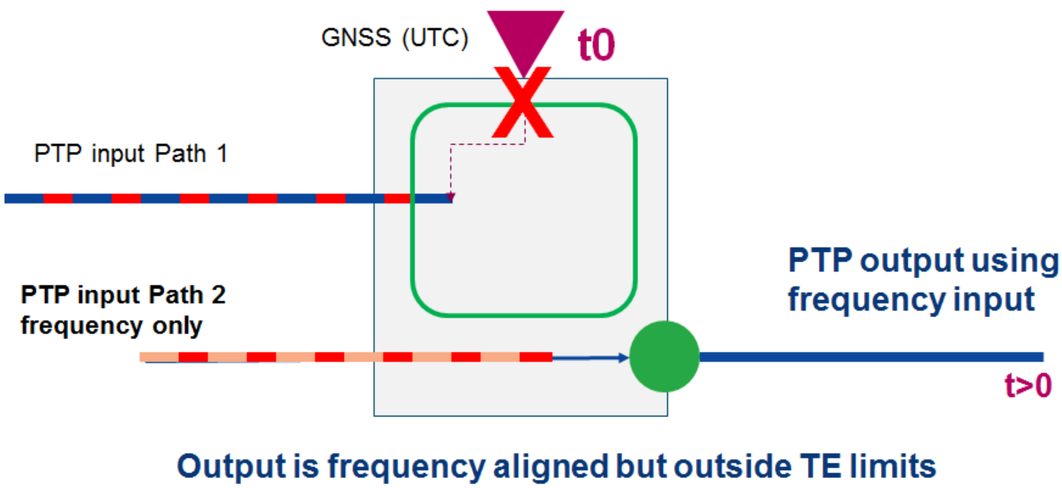

After path re-arrangement, the PTP input only stabilizes the frequency plane, rapidly the PTP output will wander with the frequency and drift outside the TE limits.

As can be seen in the previous illustrations, the standard implementation assumes that the network is static and that the PRTC will be able to rely on the incoming PTP flow to deliver a reference clock. However, modern asynchronous packet networks are dynamic; network rearrangements are quite common and PTP paths can and do change. One of the primary value adds of an MPLS or OTN network for example is that re-routes occur seamlessly without having to reserve alternative paths or provision extra bandwidth in the network. For frequency applications this may not be a major problem, depending on the number of hops the PTP packets have to traverse, but for a phase application that relies on well-engineered time error, a PTP flow network path change can be problematic since the new path will almost certainly have a different time error.

Microsemi has solved this problem by enhancing the standard with automatic asymmetry compensation, a patented method that allows time error compensation on up to 64 PTP paths per system.

Automatic Asymmetry Compensation (AAC)—A Microsemi Proprietary, Patented Feature

Automatic asymmetry compensation implemented by Microsemi significantly enhances the standardized APTS algorithm. The following illustration shows a simple representation of AAC.

APTS + AAC

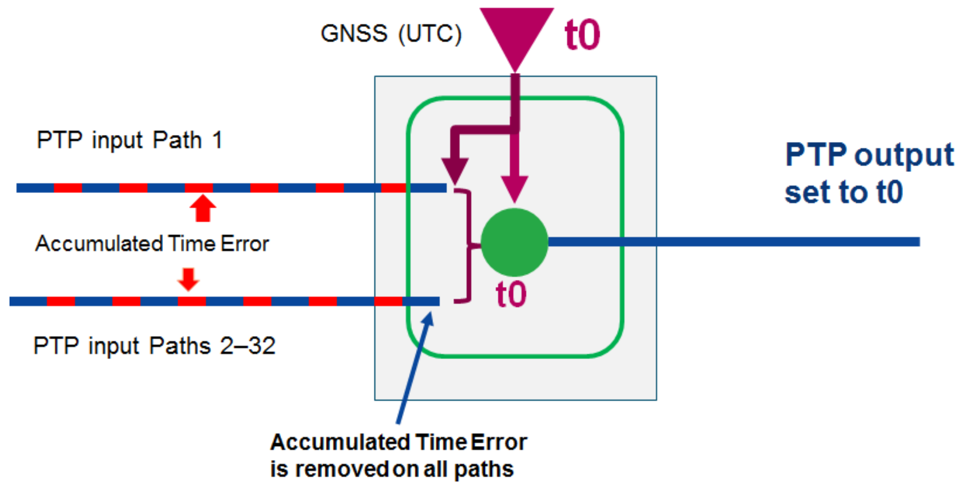

With AAC, up to 32 paths are corrected for time error.

As discussed previously, with G.8273.4 the system calibrates only one PTP path. Under these circumstances, a time error calibration is only viable as long as the calibrated path is viable. If the path between the core and edge PRTC should change under re-arrangement, then the inherent time error will change and the path compensation or calibration is no longer viable.

With automatic asymmetry compensation from Microsemi, a PTP input path time error table is maintained by the edge PRTC system for up to 32 input PTP timing paths. Each path is associated with the PTP master that provides the active flow. Moreover, in the case of Microsemi edge PRTC and gateway clocks, multiple clients can operate on the same system, each with the potential to calibrate up to 32 input paths for time error, for a total of 64 timing paths for a system.

Recall that APTS calibrates the PTP input using the local edge PRTC GNSS. The Microsemi AAC includes resiliency protection to ensure that this calibration is performed when the GNSS input is free of detectable jamming and spoofing. This is particularly useful in an urban environment where a stationary timing receiver can experience periods of degrade tracking every day as the GNSS satellite constellation interacts with the local blocking and reflection environment.

Asymmetry Correction: Detail

Just because the PTP flow is calibrated, this doesn’t mean it is actually providing correction to the PTP output. If GNSS is driving the phase/time outputs, then the incoming PTP flow is not involved in driving the output. An important point here is that the ability to generate asymmetry table entries and have a calibrated path is completely unrelated to whether or not the current PTP path is actually driving output or not. In other words, the APTS + AAC is active at all times, whatever the state of the local system, including the GNSS.[4]

The AAC feature dynamically builds up a history that enables the system to recall what has previously been seen. The table entries for asymmetry correction constitute a database that stores information about the PTP paths associated with the unique clock ID of the source PRTC. Each entry has a signature that is used for that path when GNSS is not available. Once identified, the stored asymmetry offset associated with that path is applied if it is seen again.

Network rearrangement can of course affect the PTP input as it can cause significant change in PTP flow characteristics, such as a complete loss of flow, change in noise characteristics (changes in floor packet percentage), or a change of round-trip time. When such a significant change occurs in the PTP flow, it needs to be re-evaluated, and then, if the right criteria are met, it can become a calibrated path. Of course, new asymmetry path entries cannot be created without GNSS availability (which provides the calibration reference).

Microsemi APTS + AAC: All PTP Paths Are Calibrated

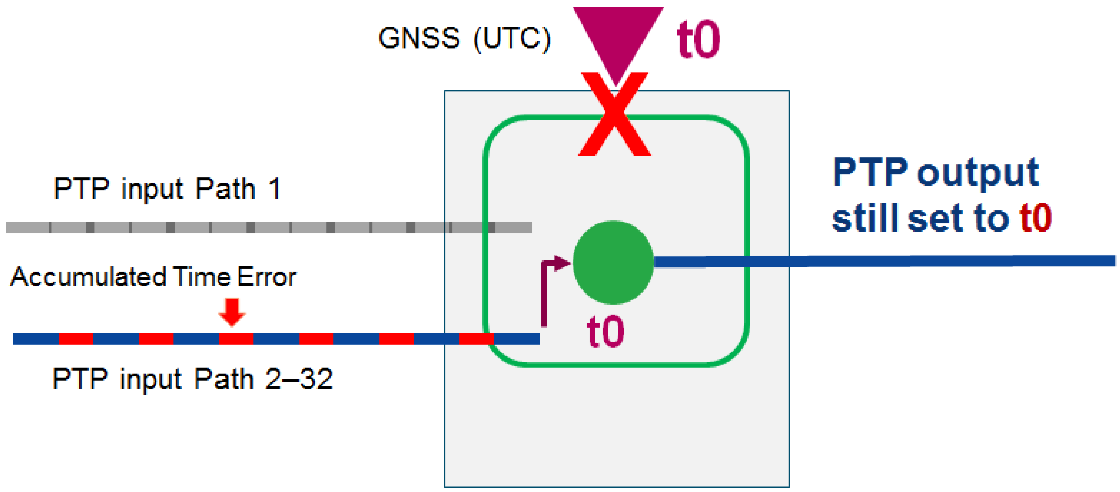

If Path 2–32 is active, the input PTP flow is still corrected for time error and is locked to t0.

With AAC all oscillators remain within ±100 ns for extended periods of time.

Behavior When Path Is Not Calibrated for Time Error

If the PTP input is driving the PTP phase/time output, phase adjustment will occur if the input is a calibrated path. If the PTP path has not been calibrated for time error, then only frequency adjustments will be applied. This behavior protects phase/time outputs from being impacted by unknown PTP asymmetry that would occur if phase/time adjustments relied on a PTP path that had not been calibrated for time error.

Example APTS AAC Operation

Consider the following scenario.

- The system is initially running with GNSS and PTP (with Microsemi AAC, the asymmetry feature is automatically enabled).
- GNSS is driving the outputs.
- All outputs are at t0.

Assume the current PTP path has an offset correction (time error due to asymmetry) of 3 µs. This becomes the calibrated path. The path is calibrated because the asymmetry adjustment is automatically applied while the GNSS is active.

GNSS is lost, so the PTP input takes over driving the outputs.

There are 2 scenarios:

- The PTP path with an offset correction of 3 µs becomes the calibrated path. The PTP input takes over and drives the phase output.
- Let us assume there is now a change in the PTP input path caused by some network phenomenon such as a fiber cut or re-route to improve load balance, a new PTP completely different signature appears (for example, a change in round-trip time). Since GNSS is not available to establish the asymmetry associated with this path, it cannot be calibrated (unless we have already seen this path before, which would mean an existing calibrating match exists with the new signature). In this case, the system will remain in GNSS Not Available and, since the new PTP path is not calibrated, there will be no phase adjustments. However, note that as per the G.8273.4 standard, there will be frequency corrections (but only frequency corrections). This will stabilize the phase outputs but there will be a slow divergence of the time error over time (wander) because of real world noise in the frequency corrections.

There are two main possibilities from here:

- The original PTP path returns. This will cause further system rearrangement. Detection of the known signature will result in the use of the already calibrated PTP input. Active phase control resumes.
- GNSS returns, which allows the new PTP signature to be calibrated, producing a new table entry. From this point forward, this signature can be identified and used for phase control if it is available and GNSS is not.

As we have already stated, for AAC to be functional, the local GNSS must be qualified and operational because the GNSS input is used as the calibration value; PTP input paths are compared and validated against this value. However, once at least one table entry has occurred, the asymmetry feature can function without GNSS.

AAC implemented by Microsemi also enables user adjustment of phase-aligned outputs (for example, PPS, PTP master output) when PTP is the selected input reference. This allows for user compensation of known, static asymmetry in the PTP input path. A use case is a scenario where the path between the source PRTC and the edge PRTC is known to have a fixed rate conversion in the path for example from 1GE to 100BASE-T. This conversion creates known asymmetry of about 6 µs, which would result in 3 µs of phase error (error due to asymmetry is always half the difference in the path lengths). However, the user must know the asymmetry on the path, and this will require measurement. Thus, this configuration option is only viable when the asymmetry in the PTP path is known and constant. If there is some dynamically changing asymmetry in the path, this capability is not helpful because it cannot adapt. The strength of AAC on the other hand, as the name implies, is that it automatically detects and compensates for asymmetry without having to implement a separate measurement and inject a value manually.

Conclusion

As mobile networks evolve from frequency-based networks to dense, highly distributed radio heads that require phase alignment to provide advanced 4G/5G services, it will be increasingly necessary to deploy PRTC clocks around the edge of the network. These PRTC can be protected by implementing assisted partial timing support, G.8273.4, an engineering tool that can be used to back up the PRTC at the edge from a core PRTC clock.

Summary of APTS ACC Operation

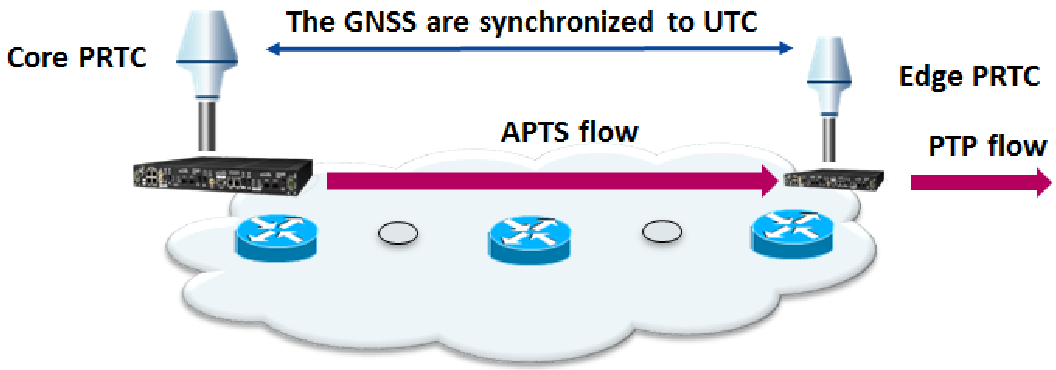

APTS synchronizes the core PRTC with the edge PRTC using PTP.

The PTP flow into the edge PRTC is calibrated using the local GNSS, which is the same time reference (UTC) as the core GNSS.

AAC enables the PTP input path to change and remain calibrated for time error—up to 32 different paths can be sustained.

APTS + AAC is a powerful engineering tool for 5G synchronization architectures.

APTS enables backup of phase-based clocking systems from a robust core clock such as an enhanced PRTC system leaning on Cesium operation to provide non-stop timing flows to the edge devices.

The following table compares APTS to APTS + AAC.

APTS to APTS + AAC

Parameter	APTS	APTS + AAC
Time error correction	Yes	Yes
Frequency correction	Yes	Yes
Multiple client inputs	No	Yes
Multiple PTP input paths	No	32
Total PTP inputs	1 or 2	64
Automatic operation	?	Yes
Works seamlessly and automatically in dynamic network environment	No	Yes

However, the standard APTS algorithm is limited to providing time error correction for one PTP input flow, and therefore lacks a fundamental resiliency—that is, the ability to calibrate and use more than one PTP input path that has been corrected for time error. Microsemi's patented Automatic Asymmetry Compensation is a powerful enhancement and imprvoement to the standard enabling the edge PRTC to calibrate up to 64 different PTP input paths and therefore remain in operation even with significant and frequent changes in the transport network.

Microsemi provides world-leading clock and synchronization systems, with an unmatched pedigree in development of technologies and tools that have enabled seamless operation of fixed and mobile networks for over 25 years.

As can be seen from this short paper, APTS + AAC is another significant value add in this long record of innovation.

Footnotes

1. LTE-A and 5G NR macro network services require adjacent radio-to-radio alignment of within a 3-µs window, **±**1.5 µs relative to UTC between the PRTC and the eNB radio.
2. Based on ITU-T Recommendation G.8265 (derived from Recommendation G.803).
3. In the evolving eCPRI environment, the backhaul network will evolve to fronthaul, midhaul, and backhaul.
4. Having paths entered in the TE table does not necessarily guarantee that the edge PRTC is currently (“at this moment”) able to provide asymmetry compensation. The ability to provide asymmetry compensation is simply stated as, “If (and only if) the current PTP flow has been signature matched with a table-entry, then (and only then) we are currently able to compensate for asymmetry.”

Microsemi Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA
Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996
Email: sales.support@microsemi.com
www.microsemi.com

© 2018 Microsemi. All rights reserved. Microsemi and the Microsemi logo are trademarks of Microsemi Corporation. All other trademarks and service marks are the property of their respective owners.

Microsemi makes no warranty, representation, or guarantee regarding the information contained herein or the suitability of its products and services for any particular purpose, nor does Microsemi assume any liability whatsoever arising out of the application or use of any product or circuit. The products sold hereunder and any other products sold by Microsemi have been subject to limited testing and should not be used in conjunction with mission-critical equipment or applications. Any performance specifications are believed to be reliable but are not verified, and Buyer must conduct and complete all performance and other testing of the products, alone and together with, or installed in, any end-products. Buyer shall not rely on any data and performance specifications or parameters provided by Microsemi. It is the Buyer's responsibility to independently determine suitability of any products and to test and verify the same. The information provided by Microsemi hereunder is provided "as is, where is" and with all faults, and the entire risk associated with such information is entirely with the Buyer. Microsemi does not grant, explicitly or implicitly, to any party any patent rights, licenses, or any other IP rights, whether with regard to such information itself or anything described by such information. Information provided in this document is proprietary to Microsemi, and Microsemi reserves the right to make any changes to the information in this document or to any products and services at any time without notice.

Microsemi, a wholly owned subsidiary of Microchip Technology Inc. (Nasdaq: MCHP), offers a comprehensive portfolio of semiconductor and system solutions for aerospace & defense, communications, data center and industrial markets. Products include high-performance and radiation-hardened analog mixed-signal integrated circuits, FPGAs, SoCs and ASICs; power management products; timing and synchronization devices and precise time solutions, setting the world's standard for time; voice processing devices; RF solutions; discrete components; enterprise storage and communication solutions; security technologies and scalable anti-tamper products; Ethernet solutions; Power-over-Ethernet ICs and midspans; as well as custom design capabilities and services. Microsemi is headquartered in Aliso Viejo, California, and has approximately 4,800 employees globally. Learn more at www. microsemi.com.

MSCC-0104-WP-01009-1.00-1118