

White Paper
Precise Timing for Base Stations in the Evolution to LTE

August 2019



Contents

1	Revision History	1
1.1	Revision 2.0	1
1.2	Revision 1.0	1
2	Introduction	2
3	Using IEEE1588-2008 for Time and Phase Delivery	5
3.1	Basic PTP Operation	7
3.2	PTP Clock Types and On-Path Support for 1588 Packets	8
4	Reference Network	9
5	Time Error Budget	12
6	TC and BC Comparison	15
7	Test Results	17
8	Summary and Recommendations	23

1 Revision History

The revision history describes the changes that were implemented in the document. The changes are listed by revision, starting with the most current publication.

1.1 Revision 2.0

Revision 1.0 was published in August 2019. In this revision, the document was updated to the latest template.

1.2 Revision 1.0

Revision 1.0 was published in February 2013. It was the first publication of this document.

2 Introduction

As mobile networks are evolving to support new wireless technologies such as LTE and LTE-Advanced, and the associated higher data consumption bandwidth requirements of mobile devices and the users, the existing mobile access network is struggling to keep up. As a consequence, there are three primary technologies that need to evolve. Firstly, the base stations and their radio technology. Secondly, the ability to deliver dramatically more bandwidth at lower cost per bit to these new base stations. Finally, the distribution of the timing to the base stations that is so fundamental to the operation of the radios. This paper discusses the former two at a high level in the next few paragraphs before focusing on network timing in more detail.

With the ongoing rollout of LTE services worldwide, service providers are simultaneously upgrading their base stations with more efficient and cost-effective equipment incorporating Software Defined Radio (SDR) technology. This allows the same base station not only to deliver the new LTE services, but also continue to provide coverage for the existing 3G and sometimes even legacy 2G devices with the same equipment. These new multimode SDR-enabled base stations have initiated a significant transition in the backhaul networks. Rather than connecting to the radio access network using SONET/SDH or PDH E1/ T1 connections as in the past, they now connect to the network using much more efficient packet-based multi-Gigabit Ethernet connections. This means that the wireless backhaul network is upgraded to a more cost-effective broadband packet network, which can deliver vastly more bandwidth to base stations at significantly lower cost per bit. The transition to broadband IP/Ethernet backhaul networks is well underway, but will continue for years to come.

In addition to upgrading the macro base stations (often called macro cell sites) providing basic LTE coverage for the first wave of LTE-capable devices, more widespread adoption of LTE by mobile users requires a completely new base station technology called small cells. Small cells provides the needed increased capacity and coverage provided by the macro base station networks. These small cells have a smaller footprint, lower power, and lower cost than the macro cells, and will be mounted on top of lampposts, traffic signals, and are attached outside the buildings. To provide sufficient indoor coverage, small cells will also be deployed inside the shopping malls, stadiums, and large office buildings.

Small cells are a relevant topic for this white paper because the backhaul network must provide timing to the small cell radios, which creates additional challenges. After cell location and connectivity, delivery of precise timing is the most important element in radio access networks, as timing synchronization directly affects radio spectral efficiency and in turn throughput. While macro base stations previously could receive accurate frequency synchronization from legacy E1/T1 TDM connections and Time of Day (ToD) through GPS/GNSS satellite receivers, E1/T1 connections are decommissioned everywhere to reduce recurring access operating costs. Many of the current mobile base stations use GNSS/GPS as the source for the all-important signal known as 1 pps. This signal is used to manage the ToD calculation for next second rollover as well as to synthesize the fundamental source radio frequencies. But GPS is neither a cost-effective nor reliable timing technology in urban environments and indoors because there is often no unobstructed line of sight to the satellites. In the last few years there have also been increasing concerns about the ease of jamming and spoofing GPS signals using inexpensive and widely available technology, and timing has therefore become a major security concern in radio access networks. These concerns are amplified for small cells that are located at street level and are typically hard to secure.

Newer wireless technologies also increasingly rely on Time Division Duplex (TDD) technology to increase spectral efficiency and reduce the cost of spectrum licenses. TDD technology requires not only frequency synchronization between base stations, but also very precise and stable phase alignment. Use of advanced technologies for beam forming such as Multiple Input – Multiple Output (MIMO) and Coordinated Multipoint Transmission (CoMP) also rely heavily on precise phase alignment between radios in the base stations. The [Frequency and Phase Accuracy Requirements for Mobile Technologies \(see page 3\)](#) table shows the frequency and phase accuracy requirements at the air interface for different mobile technologies. The terms Frequency Division Duplex (FDD) and TDD, describe the structure of signaling used at the air interface. Frequency accuracy is self-defining. With TDD the phase accuracy is referencing the point in time in which the inter-channel guard band separating usage periods occurs. For the newer mobile technologies, timing is a parameter that severely affects the throughput. Although the following table shows the minimum requirements it should be noted that achieving higher time/phase and frequency accuracy lead to improved performance/throughput and to higher systems and network margins to accommodate residual link or systems impairments.

Table 1 • Frequency and Phase Accuracy Requirements for Mobile Technologies

Application	Frequency Transport/Air Interface	Phase Air Interphase
ITU-T G.8261 Network Target/ 3 GPP Max Error		
GSM/UMTS/W-CDMA	16 ppb/50 ppb	None
UMTS/W-CDMA Femtocells	n/a /250 ppb	
ITU-T G.8261 Network Target/ 3 GPP Max Error		
CDMA 2000	16 ppb/50 ppb	±3 μs to 10 μs
TD-SCDMA	16 ppb/50 ppb	±1.5 μs
ITU-T G.8261 Network Target/ 3 GPP Max Error		
LTE (FDD)	16 ppb/50 ppb	None
LTE (TDD)	16 ppb/50 ppb	±1.5 μs

Application	Frequency Transport/Air Interface	Phase Air Interphase
LTE-A MBSFN	16 ppb/50 ppb	$\pm 1 \mu\text{s}$ to $32 \mu\text{s}$
LTE-A Hetnet Co-ordination	16 ppb/50 ppb	$\pm 5 \mu\text{s}$
LTE-A CoMP (Network MIMO)	16 ppb/50 ppb	$\pm 5 \text{ ns}$ (within cluster)

This whitepaper shows how a backhaul networks with IEEE1588 Precision Time Protocol (PTP) with hardware timestamping can fulfil the requirements for phase synchronization at the air interface even for TD-LTE and LTE-Advanced wireless standards. It also discusses how backhaul technologies can differ in their ability to minimize time errors even in the presence of PTP support in hardware, and how time errors in a typical backhaul network add up to the overall time error budget at the air interface. It is concluded that 10 ns accurate 1588 Boundary Clocks (BCs) and Transparent Clocks (TCs) are needed in most fiber-connected backhaul gateways and routers and macro and small cells themselves. It also shows how widespread use of 1588 TCs can lower the cost of the timing recovery at the end-node, which will be particularly important for cost sensitive small cell deployments. Strategic placement of BCs is recommended at network interconnect points to manage PTP processing scale.

3 Using IEEE1588-2008 for Time and Phase Delivery

IEEE1588-2008, PTP is deployed into service provider mobile wireless backhaul networks. This standard supports both frequency and time/phase distribution directly through the backhaul network using timestamped packets and a protocol exchange to derive the actual time from the timestamped packets.

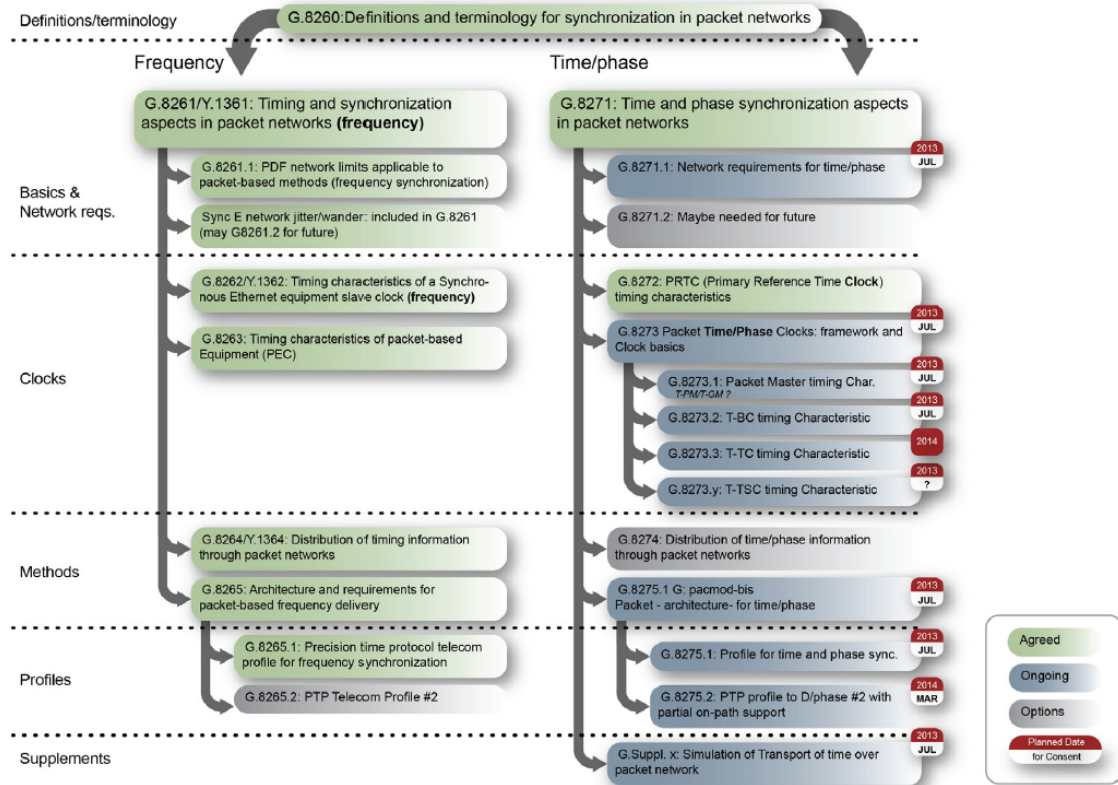
The IEEE1588-2008 standard was written with several annexes that allow various standard or industrial organizations to extend the basic operation of the PTP protocol to fit their specific needs.

For example, ITU-T uses these annexes to develop so-called PTP profiles for telecommunications networks. The ITU-T has so far developed a PTP profile for frequency-only delivery over PTP-unaware networks, meaning that the network elements between the PTP grandmaster and the PTP ordinary clock slaves are not modifying the PTP packets. The Packet Delay Variation (PDV) resulting from sending PTP traffic through PTP-unaware network elements is slowly filtered out by an intelligent low-pass filtering algorithm (also called "servo") in the PTP ordinary clock slave. This profile is specified in ITU-T G.8265.1. The servo algorithm used in the ordinary slave clocks is not a part of the 1588-2008 standard and so it is proprietary to the clock vendor. It should be noted that filtering out the PDV only allows the slave to recover the frequency and not the time/phase, and that the long time constants of the filtering algorithm result in very long frequency acquisition/settling times and also susceptibility to long-term phase wander. Most of the "secret sauce" in intelligent PTP frequency algorithms involves keeping this long-term wander in check.

While this can provide a frequency-only solution, it cannot provide a synchronization solution for phase and ToD.

The following figure shows the current progress the ITU-T is working on for a set of standards including a new PTP profile for time and phase delivery over networks with full (and possibly partial) PTP on-path support. This PTP profile, which will be released as G.8275.1. ITU-T, is also working on the architecture and performance specifications for each type of equipment needed to transfer time and phase, as specified in a number of documents in the G.827x series.

Figure 1 • Completed and Ongoing Work in the ITU-T for Frequency and Time Delivery



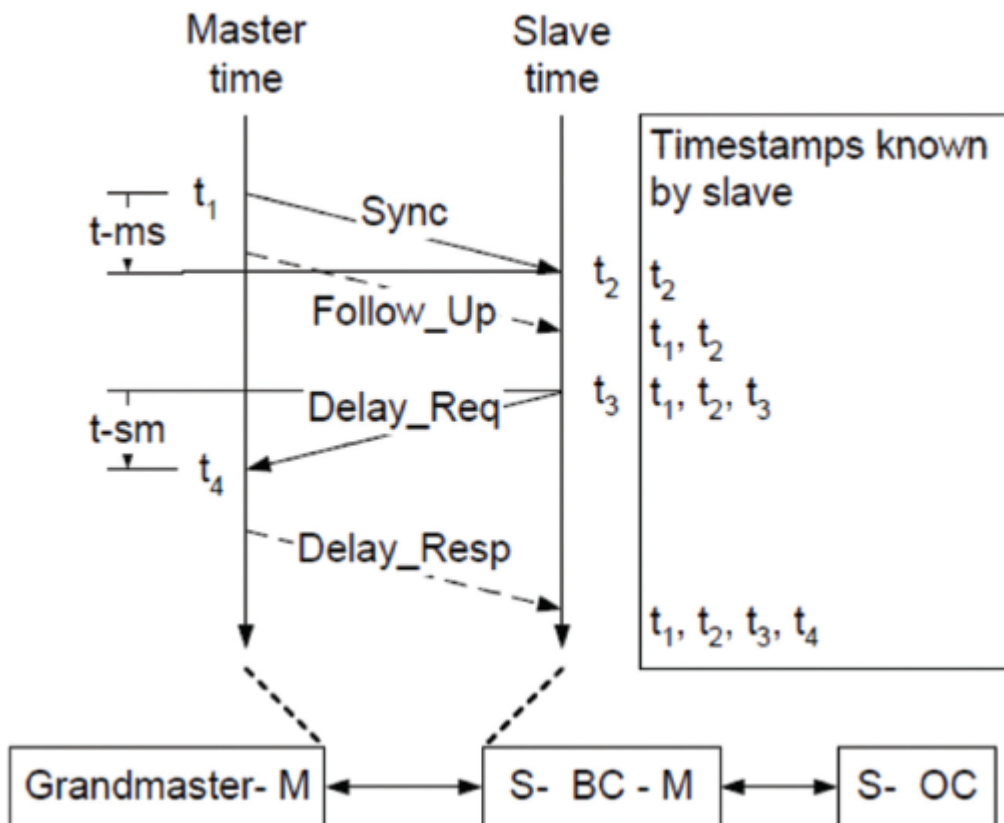
PTP-unaware networks are not designed to provide a symmetric delay, and traffic patterns as well as equipment architecture, typically causes significant and varying asymmetry in the upstream and downstream delay, making it impossible to recover the time and phase over PTP-unaware networks. It is therefore generally recognized in the industry that for time and phase delivery over the network, the network has to provide "PTP on-path support". This means all network elements between the PTP grandmaster and the PTP ordinary clock slaves operate as either PTP BCs or PTP TCs (as described) to eliminate or significantly reduce the impairments introduced by static asymmetries in network links, plus PDV caused by variations in packet traffic through these elements.

3.1 Basic PTP Operation

The PTP protocol uses timestamped frames between the PTP master and one or more PTP slaves. The simple form of the PTP interface is listed as follows.

1. The PTP master sends out a PTP Sync frame that contains the ToD of the master at the time of transmission, T1.
2. When the PTP slave receives the Sync frame it registers the arrival time, which is measured using the local time in the PTP slave, T2.
3. The PTP slave sends a Delay_req frame to the PTP master and registers the transmission time of the Delay_req frame, measured using the local time in the slave, T3.
4. The PTP master registers the arrival time of the Delay_req frame, T4, using the PTP master time. It sends a Delay_resp frame back to the PTP slave containing the T4 time.
5. The PTP slave calculates the path delay between the master and the slave using the following formula (assuming that the network path delay is symmetrical): Path Delay: $[(T2-T1) - (T4-T3)]/2$
Offset: Slave Time – Master Time = $[(T2-T1) + (T4-T3)]/2$
6. The PTP slave corrects the local time based on the calculated offset.

Figure 2 • PTP Protocol Flow



To make the time as accurate as possible it is important that the timestamps are as accurate as possible and that the path delay is symmetrical. The protocol supports compensation of known and static asymmetry in the path delays, but this of course requires that the protocol knows the asymmetry, and that the asymmetry can be measured and is not variable so that it can be effectively compensated.

The influence of timestamping accuracy is discussed later, but it is important to emphasize that any uncompensated asymmetry in the network will directly translate to errors in the time/phase derived from the 1588 ordinary slave clock. There are multiple factors that can contribute to large asymmetries in packet networks listed as follows.

- Uplink/downlink path differences introduced by actual fiber path differences in fiber networks, or different modulation formats in microwave (MW) and millimeter-wave (MMW) links
- I/O serialization delays and speed mismatches
- Queuing and forwarding delays in network processors, traffic managers, and switch fabrics

While the first impairment is static and could potentially be physically measured, and therefore be compensated for in software, the second and third impairments have large variable components. Queuing delays in particular can easily amount to 100's of microseconds of delay variations and asymmetry, easily exceeding the time error specifications in [Frequency and Phase Accuracy Requirements \(see page 3\)](#) table by orders of magnitude even in a single network element. These types of variable delays must be corrected in hardware on a packet-by-packet basis.

PTP on-path support in networks can largely eliminate time errors introduced by the impairments by proper implementation of I/O-level timestamping and the use of (distributed) TCs everywhere with limited deployments of BCs for scaling purposes.

3.2 PTP Clock Types and On-Path Support for 1588 Packets

The PTP protocol specifies the following four different clock equipment types.

- **Grandmaster** - This device defines the time for the PTP domain and distributes this time information to other PTP clock units. In the architecture of the network this is the master timing source.
- **Ordinary clock slave** - This device receives the time from a grandmaster or a BC. This device will generate timing signals for the timing sink system.


Two additional clock types were identified in later versions of the standard – "Boundary Clocks" and "Transparent Clocks". Both of these clock types provide "1588 on-path support" to network elements in a packet-based network. Advantages and disadvantages of these two clock types and their proper use in networks is discussed in detail later in this paper.

- **Boundary clock (BC)** - This device is a PTP slave on its upstream facing port and PTP master on the downlink facing ports. This device reduces the load on the grandmaster clock thus providing PTP scale to the network. When present, PTP slaves communicate with the BC instead of directly with the grandmaster clock. The BC inserts its computed value of time in place of that generated by the grandmaster clock. The PTP slaves are unable to "see" the grandmaster clock and any time offset error between the BC slave port and the domain grandmaster is passed to the slave.
- **Transparent clocks (TC)** - These devices do not directly take part in the protocol between PTP masters/BCs and PTP slaves, they are therefore "transparent" to the PTP network. TCs are network elements, such as switches and routers that measure the delay that the processing/ forwarding of each PTP Sync and Delay_req frame takes inside the network element. This time is called the Residence Time (RT). The RT for each PTP Sync and Delay_req packet is added to a PTP field in the frame called the Correction Field (CF). Each TC in the PTP path adds its computed value of residence to the current contents of the CF. These devices do not modify the timestamp value placed in the PTP frame by the grandmaster.

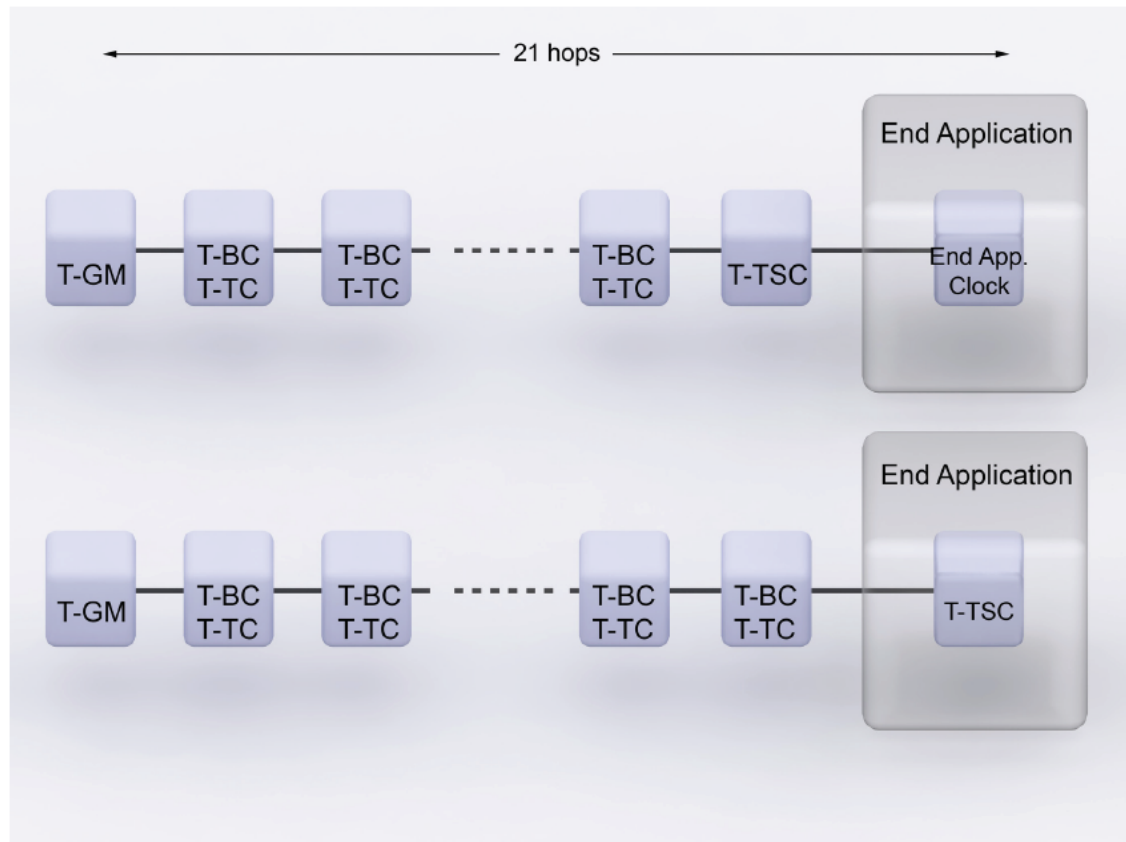
4 Reference Network

The following figure shows the typical Backhaul network. In this reference diagram, the Backhaul network provides Ethernet access to the mobile base stations and redundant PTP grandmasters are centrally located in the backhaul network. In the following figure, T-GM are Time Grandmasters receiving time from a Primary Reference Time Clock (PRTC). This is normally a GNSS/GPS receiver. T-TSC are the Time Ordinary Clock Slaves located at the base stations that derive timing from the T-GM using the 1588 PTP, and T-BC and T-TC are Boundary and Transparent Clocks providing 1588 on-path support within the network.


Figure 3 • Typical Backhaul Network



ITU-T SG15/Q13 is in the process of developing IEEE1588-2008 profiles for time and phase delivery over PTP-aware networks using the following network reference model. See the [Reference Model for Phase Delivery \(adapted and modified from ITU-T G.8271-1\) \(see page 10\)](#) figure, where there are 20 network elements and 21 hops between the PTP grandmaster and the ordinary clock slave service application clock located at the base stations. The initial network models and simulations are using BCs with and without Synchronous Ethernet as frequency reference, but reference models using TCs with and without SyncE as frequency reference will follow as part of the on-going standards developments.

Figure 4 • Reference Model for Phase Delivery (adapted and modified from ITU-T G.8271-1)


Recall the discussion on the emerging small cell rollouts providing the necessary capacity for LTE mobile devices, to provide additional perspective on realistic network deployments. Most of the macro base stations can and will be fiber-fed except in rural environments where the cost of laying fiber is too high and MW backhaul is still the primary access technology. But, as mentioned earlier, outdoor small cells will be primarily deployed at street level on top of lampposts and traffic signals. The lampposts are not connected by fiber, and only a few traffic signals will have DSL access where there are previous installations of security cameras at major intersections. Infonetics therefore estimates that nearly 90% of all outdoor small cells will be connected by MW and MMW technology. The following figure shows a typical deployment in an urban environment. While the macro cell on top of the office building may have fiber connectivity to the backhaul network, the small cells would not. Furthermore, the small cells have to connect back to a central pre-aggregation router which may be co-located with the macro cell. So that, a connection between macro cell on top of the office building and street level will have to be bridged by a MW or MMW link as well.

Figure 5 • Typical Small Cell Deployment and Backhaul Scenario


The network timing has to be delivered to these small cells. The GPS timing for small cells has various challenges that make network-based 1588 PTP timing a necessity. The problem with MW and MMW backhaul and DSL-based backhaul is that these access technologies have large PDVs because of the modems required. For MW and MMW modems, adaptive modulation schemes for various weather conditions will further lead to large changes in link delays, possibly even asymmetric. Small cells may also have to be connected in a daisy chain or partial mesh topology because of deployment constraints, and therefore multiple MW and MMW links may be present before a fiber-connected (pre)-aggregation router is reached. Even with Microsemi's proprietary distributed TC technology for such links, PDV will likely be higher for MW, MMW, and DSL connections than for fiber access networks with properly designed TC or even BC implementations within the IP Edge Routers and Switches.

It cannot be assumed that each of the 21 network hops in the ITU's reference model have the same time error budget allocation. It is not unreasonable to assume that even a well-designed MW or MMW link with distributed TC on-path support for the last link to a small cell will consume around 100 ns of the entire network budget, leaving only a few nanosecond per hop for the remaining links in the network. Time error budgets is discussed in the next section.

5 Time Error Budget

The 1588 on-path support in the form of BCs or TCs is needed in backhaul networks when phase and/or ToD synchronization is required to eliminate the influence of large asymmetric queuing delays common in switches and routers as well as asymmetric modem delay variations for MW and DSL links. But, there are significant differences in TC and BC implementation architectures, with large variations in the accuracy that can be achieved between different systems architectures, timestamping solutions, and filtering systems. Before discussing some of these differences and showing experimental test results on state-of-the-art equipment, the overall time error budget that is available in LTE networks must be looked into, and the precision a TC or BC on-path implementation needs to meet the overall time error budget at the air interface.

Static and dynamic time errors must be considered. Static time errors, can in principle be compensated for if they are known and measurable, but there will always be some residual error. For example, cable length asymmetries can only be controlled and measured to a certain degree. Dynamic time errors on the other hand result from timestamping accuracy, any dynamic delay variations in the I/O path outside of the timestamping engine(s), and the degree of phase alignment between 1 pps signals in distributed timestamping architectures; in particular when the 1 pps signal needs to be synchronized across a link as is the case for example for a distributed TC across a MW link. If the distribution of a dynamic time error is well behaved (example shows a narrow and symmetric Gaussian distribution), filtering can be applied to reduce the time error. However, if the shape of the distribution changes over time, filtering may be difficult or expensive since a highly stable oscillator is required to filter out the variations. Furthermore, it is necessary to ensure that the PTP slave or BC does not convert the dynamic time errors into a large static time error or static time error, often called "wander".

The network limits for static time error and dynamic time error have not been defined by the ITU yet, but they require a different metric than used for frequency delivery. The end applications have requirements on the absolute time error, or to the maximum time error compared to another application (mobile base stations) in the vicinity of the end application, so this puts a requirement on the maximum static time error delivered to the end application. The MTIE and TDEV masks defined for PTP frequency delivery are measurements of frequency error over a period of time and not a measure of "phase or time" against a known reference, and therefore of little value. The maximum time error that can be generated by the series of network elements and links cannot exceed the air interface time error specified by each specific wireless standard. Therefore, the maximum time error in a network is the sum of all time errors of all hops in the network.

In the following table, rules to the ITU reference model of the [Reference Model for Phase Delivery adapted and modified from ITUT G8271 1 \(see page 10\)](#) figure with 21 hops, with each hop contributing equally to the overall time error budget is applied. However, the PRTC accuracy of the eNodeB phase recovery, and uncompensated static errors from cable length asymmetries are treated separately.

Table 2 • Time Error Budget for LTE

	LTE Advanced	LTE (TDD)
Air interface accuracy required	500 ns	1.5 μ s
Budget for base station	100 ns	100 ns
PRTC TE budget	100 ns	100 ns
Fiber link asymmetry	100 ns	100 ns
Budget left for PTP equipment	200 ns	1.2 μ s
Max static TE per PTP unit (21 hops)	9 ns	57 ns

The PRTC accuracy is required to be able to deliver an accuracy of ± 100 ns (including both static time error and dynamic time errors) relative to UTC. This estimate is supported by a number of measurements done on commercially available telecom GPS receivers with good and calibrated antennas and highly stable local oscillators (the GPS antenna cable length affects the time accuracy if not compensated for). This requirement is specified in the ITU-T G.8272 specification released 9/20/12. It should be noted that if neighboring base stations within a LTE-A MIMO/CoMP network are receiving the time/phase information from the same PRTC, then the PRTC error can be removed from the budget.

The time accuracy requirement at the end of the PTP chain depends on the end-node and its radio interface specification. In [Table 1 \(see page 3\)](#), the accuracy of the air interface for LTE varies from 1.5 μ s to 500 ns, and a yet unspecified amount also needs to be reserved for the base station equipment. A budget of 100 ns for the timing end station (either eNodeB itself or cell site gateway/router) is assumed for time being.

Some budget is needed for cable asymmetries and unless each cable asymmetry is measured and compensated for this could take a huge part of the remaining budget. In this budgeting exercise, take out 5 ns per link (equal to approximately 2.5 m cable skew, as this requires a good cabling control), or 105 ns total.

If the remaining time error budget of 195 ns is equally divided among the 21 links of the ITU network model, this results in a time error budget of only 9 ns per hop for LTE-Advanced, and 57 ns per hop for TD-LTE.

Note: The recently released power distribution smart grid budget has per hop error limits specified to demonstrate that other users of PTP have taken this step as well.

The most stringent requirements are currently for LTE-Advanced, and the example budget above gives a maximum static time error per PTP BC/TC switch/router of only 9 ns. Some dynamic time error is allowed on top of this, assuming that the dynamic time error can be filtered off by the end-node and is not converted to a static time offset in any of the nodes in the network path.

Realistically not all hops in the network will be able to maintain 9 ns filtered dynamic error per hop, and some of the links in the chain may need a larger time error budget. This is the case for small cell backhaul where definitely the last, but likely multiple links, rely on MW or MMW links to connect the small cells back to a fiber-connected aggregation router. Significantly, less time error budget is left for the rest of the network links.

Though the wisdom of the ITU-T is not questioned in requiring 21 hops between PRTC and the base station for phase and ToD delivery—when the hop count for frequency delivery is only 10 hops, a realistic network model is proposed in the following table, taking into account different backhaul equipment that is still able to meet the LTE-advanced timing requirements at the air interface.

Table 3 • Modified Time Error Budget for Typical LTE Deployments

Equipment	Max Time Error	Hop Count
eNodeB/small cell	10 ns	3
Cell site gateway	10 ns	3
MW/MMW link	100 ns	2
Aggregation router	10 ns	15
Fiber length asymmetry	5 ns	15
Maximum total time error	485 ns	

In the modified budget in this table, 100 ns of time error budget is allocated each to two MW links at the expense of reducing the hop count from 21 to 15. It is also assumed that all eNodeB/cells that need to be phase synchronized are fed by the same PRTC, so that the budget previously allocated to differences in phase between different PRTCs can be eliminated.

Note: The time error budget is tightened for the cells themselves from 100 ns to 10 ns, and have maintained the tight accuracy of 10 ns per hop for fiber-fed routers and gateways.

The [Test Results \(see page 17\)](#) section shows that the example budgets in preceding table are quite achievable given proper design of TC and BC on-path support and choice of timestamping architecture.

6 TC and BC Comparison

There are major differences in the functions of a TC versus a BC, and both have their advantages and disadvantages. When designing a network, careful consideration of these positive and negative elements should be considered. The optimal solution may be a combination of BCs located at network interconnection points and TCs in the subsequent tributaries and rings. Generally, most network elements will be able to support both BC and TC, in particular if they could also serve as a time slave. However, network elements that are always located at transit points in the network can be upgraded to highly precise 1588 TC functionality just by upgrading the I/O-level PHYs with 1588 PHYs like the ones from Microsemi, without the cost of a full BC implementation. There are also advantages in having as many TCs as possible in the network to reduce the cost of filtering at the eNodeB.

A TC registers the arrival time of PTP Sync and Delay_req frames at ingress, calculates the residence time at egress, and updates the CF inside the frame. There are no master or slave ports, so the PTP traffic direction can change on the fly, without any need for reconfiguration or delay. This makes the TC useful in ring structures where the PTP traffic flow direction changes after a failure. Also, the TC operation is done without any CPU load for processing PTP frame updates.

The BC architecture is a PTP slave on the upstream port and PTP master on the downstream ports. A CPU in the system must be running the clock recovery algorithm adjusting the computed clock value based on the PTP frames received from the selected PTP master based on the best master selection algorithm. The BC must then generate PTP frames to the slaves connected to it. The BC needs to support switching between redundant masters, Best Master Clock Algorithm (BMCA), and perhaps even support some degree of time holdover in case of failure on the connected master. It will require a very stable reference clock, often in the form of an expensive OCXO. In reality, each BC is a proxy clock for the domain grandmaster. The BC operation is far more complicated and expensive than the TC operation because of the additional timing functionality and components required.

BCs offer one main advantage over TCs. They reduce the number of ordinary slave clocks accessing each domain grandmaster therefore improving the scalability of each grandmaster. In high PTP packet rate networks this becomes significant. An example would be if the Sync and Del_req packet rate was 64 pps, equaling 128 pps per slave, the grandmaster could easily become a packet bottleneck with 1000 slaves. If only End-to-End (E2E) TCs were used between a high number of slaves and a grandmaster, the grandmaster would need to respond to all the delay request frames from the slaves. In a fully PTP-aware network the packet rates would typically be much lower, 1 pps–16 pps, easing the load on the PTP master.

Peer-to-Peer (P2P) TCs are solving the problem by loading the grandmaster with Delay_req processing as they use the peer delay measurement mechanism specified by IEEE1588-2008. Instead of measuring the full path delay from slave to master, the path delay for each connection is measured by the TC and ordinary clock slaves. The TC includes the ingress path delay in the CF update of the Sync frames and replies to peer delay requests from any other devices. This requires CPU processing of these Pdelay_req frames, but the actual correction of the Sync frames can be done in hardware without CPU processing.

From a management point of view the use of TCs also has the advantage that the PTP configuration only needs to be handled at the grandmaster and at the slaves.

Table 4 • Comparison Between BC and TC

Configuration	BC	E2E TC (non-synchronized)	P2P TC
Noise accumulations	Medium (depending on the total number of nodes in the distribution chain as well as the PLL bandwidth and local reference stability of each BC)	Very low (depending only on the total number of nodes in the distribution chain)	Very low (depending only on the total number of nodes in the distribution chain)
Clock selection (BMCA) on every NE	Yes	No	No
Supporting multiple clock domains	Yes, but computationally problematic	Yes	Yes
Management requirements (OPEX)	Requires management on every node	Not required (plug and play)	Not required (plug and play)
Design complexity (HW/SW)	High as it requires additional SW to manage states, packets generation and termination and run BMCA for example	Very low. All the required support (HW mostly) is already built into the PHY	Very low. All the required support (HW mostly) is already built into the PHY
Local reference	Relatively expensive. Stratum 3 or better than Sync-E (G.8262 TCXO) for filtering and holdover	Very cheap (no filtering or holdover are needed)	Very cheap (no filtering or holdover are needed)
Temperature sensitivity	Medium–Even OCXOs are affected by temperature ramps and this causes time/phase drift	Very low	Very low
External recovered time (for example, 1 pps) monitoring capabilities	Yes	No	No

Most of the routers and switches will support both BC and TC functionality even though a BC is more expensive and complex to implement. The TCs will become more prevalent and can reduce the cost of 1588 on-path support as it gets closer to the cost sensitive edge of the network. Furthermore, TC results in a time error distribution that can be filtered out by low-cost oscillators in the time slaves, so there are significant benefits in enabling TC functionality even if the network elements can support both BC and TC. This is particularly relevant for timing distribution to cost sensitive small cells.

7 Test Results

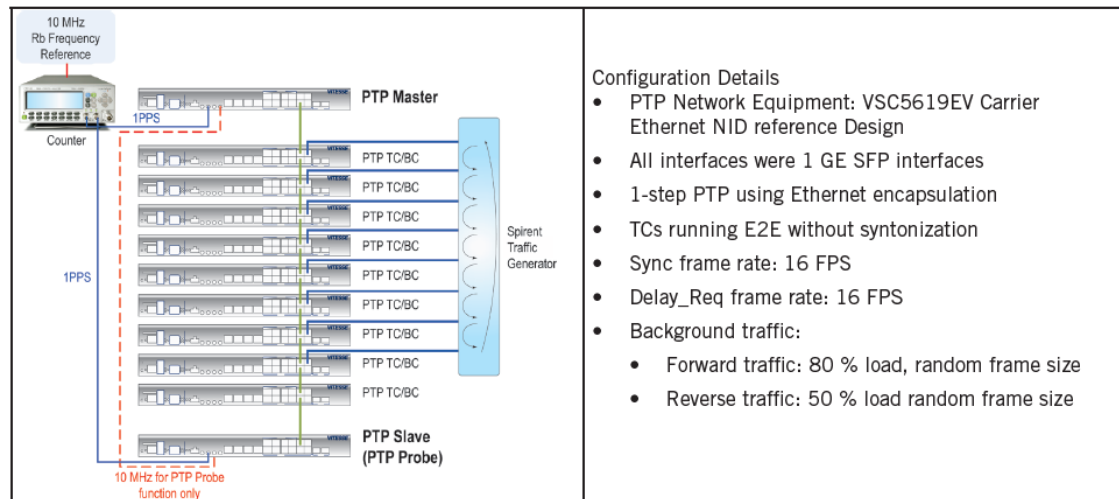
Microsemi has conducted a series of tests using the Carrier Ethernet NID reference design with the VSC7418 Serval™ mobile backhaul optimized Carrier Ethernet (CE) Switch. This device is designed to support carrier grade PTP with the accuracy needed for mobile backhaul. It can operate in PTP master, slave, BC, and TC modes.

The Serval switch used in this example is a unique CE switch that supports Provider Bridging and MPLS-TP. It is designed specifically for backhaul and network edge applications by handling a full suite of Operations Administration and Maintenance (OAM) functions, including 802.ag Maintenance Entity End Points (MEPs) for both MEP-Up and MEP-Down. The switch supports full hierarchical Quality of Service (QoS), including policing and traffic shaping.

All timestamping functionality (as in Microsemi 1588 PHYs) is implemented in hardware on a packet-by-packet basis, including asymmetry corrections for known asymmetries, and frame CRC updates. As a result, there is no CPU loading from the frame modification and forwarding functions, and no performance degradation as a function of network loading.

The following figure shows the test configurations similar to the standard G.8261.1 test setup.

Figure 6 • Test Setup



The following table lists the test results.

Table 5 • Test Results

Configuration	Frequency Support	PDV at PTP Slave	1 pps MTIE (5000 s)	Average Phase Offset	1 pps Max TE
Master-slave	PTP		16 ns	2 ns	8 ns
Master -9 x TC - slave	PTP	33 ns	24 ns	10 ns	25 ns
	SyncE	26 ns	17 ns	2.9 ns	11 ns
Master -9 x T-BC - slave	PTP	48 ns	75 ns	17 ns	51 ns
	SyncE	25 ns	30 ns	2.7 ns	17 ns

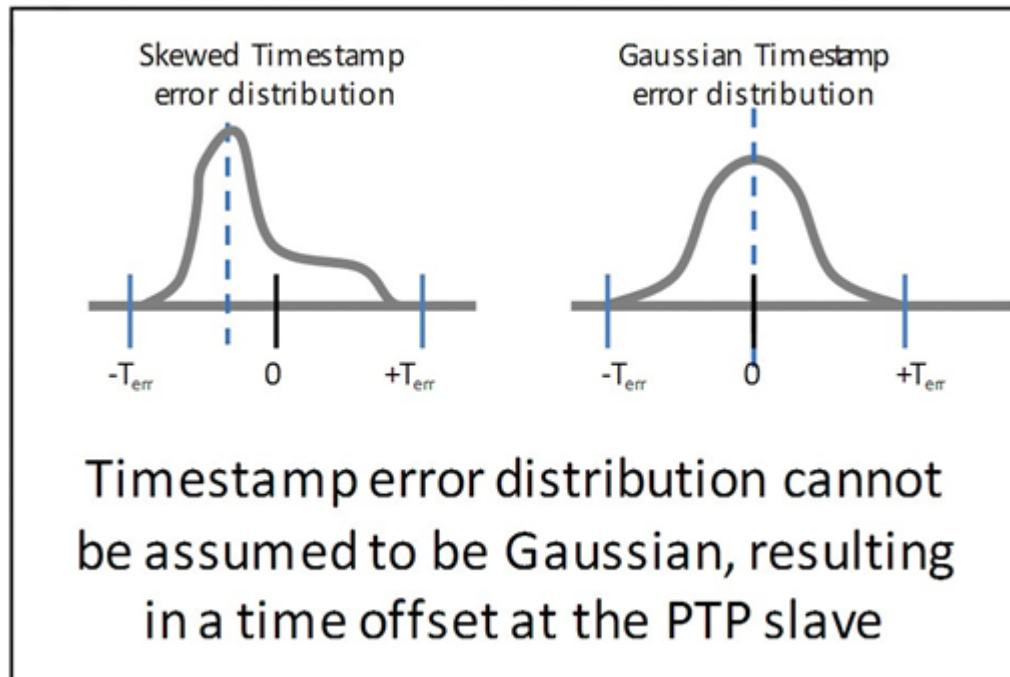
The results show that well designed TC and BC implementations can easily meet the maximum time error per network hop of fiber connected switches and routers in a backhaul network. There are performance differences between BC and TC network chains because of the high-accuracy 1588 on-path support of the Microsemi architecture, specifically, as the following.

- The packet delay variation measured at the end of a series of devices starting with a PTP master followed by 9 TCs or 9 T-BCs before reaching a PTP probe are comparable, with the TCs performing better than the T-BC.
- When measuring the performance of a PTP slave added to the end of the 9xTC/BC chain, the performance measured at the 1 pps interface is also better when the chain is configured in TC mode than when configured in BC mode. This is most likely a result of the noise added by the filtering functions in the BCs.
- Adding SyncE in parallel with PTP improves the performance of the BC chain significantly, but never reaches the accuracy level of the TC chain, even without SyncE support. The SyncE source was a Rubidium reference. SyncE wander effects were not included in the tests.
- Both the TCs and the T-BCs were unaffected by intervening traffic, as should be expected for network elements with proper 1588 on-path support.

The performance of the TC and BC on-path support in the Microsemi solution results from careful control and minimization of both static and dynamic time errors in each network element.


The static time errors are typically caused by uncompensated asymmetry in the physical layer devices and logic placed before the timestamp point of the system, timing distribution offsets inside the system and PDV filtering algorithms that are not optimized for full PTP support in the network.

Major contributors to dynamic time errors are the timestamping resolution and oscillator noise from the equipment. The timestamping of the PTP frames is done with a certain timestamp resolution and the lower the resolution the higher the amplitude of the time error noise. If no filtering is applied to the recovery of the time using PTP, then the worst case timestamp noise between a PTP grandmaster and a PTP ordinary clock slave or BC (excluding any other noise sources) will be equal to the timestamp resolution at each timestamp location. The frequency spectrum of the timestamp noise depends on many parameters, but it cannot always be assumed to have a Gaussian distribution around zero time error, especially not if SyncE is used in parallel with PTP. To ensure a minimum time error it is important also to keep the dynamic time error low, for example, having a high timestamp resolution.

Figure 7 • Timestamp Error Distribution

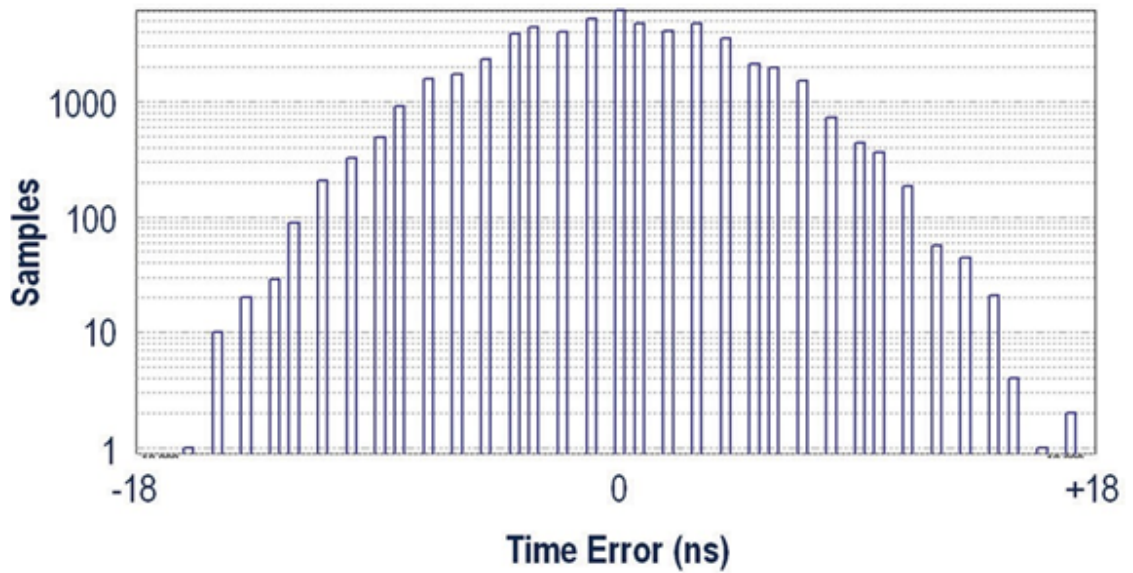
It is important to understand that the TCs operate with full hardware PTP function, so the PTP frame modifications and forwarding functions are not loading the CPU in any way. This also means that there is no filtering function or clock servo algorithm that can suppress any noise from the upstream nodes. This can be seen in the following figure, where the packet delay variation is shown taken after 9 TCs (the PTP slave acting as a PTP probe). This shows that the PTP master, 9xTC, and the PTP probe only generates a variable time error noise (timestamp noise) of 33 ns.

Figure 8 • 9xTC Sync Timestamp PDV – No SyncE




The following figure proves that the time error timestamp noise can be filtered down easily at the slave, as it has a normal distribution.

Figure 9 • 9xTC Sync Timestamp PDV Distribution – No SyncE



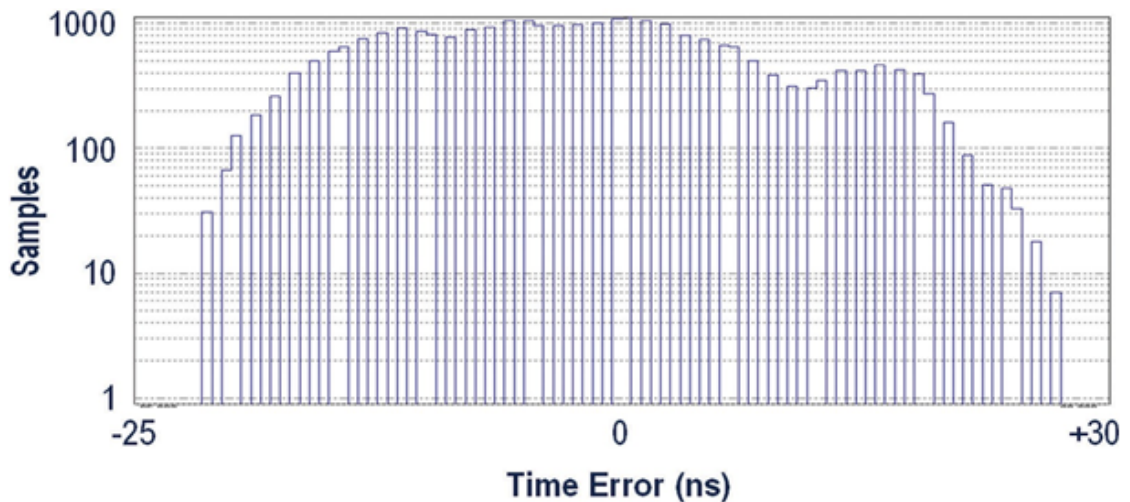
The PDV measured after a chain of 9 BCs show that the PDV is varying at a much slower pace, making it much harder for the end-node to filter off the PDV noise.

Figure 10 • 9xBC Sync Timestamp PDV Distribution –With SyncE



The following figure shows that the distribution of the time errors of the chain of Boundary Clocks. Unlike the Gaussian distribution of the TC chain of [9xTC Sync Timestamp PDV Distribution No SyncE](#) (see [page 20](#)) figure, the distribution is asymmetric and not centered around a specific value, making it much more difficult to filter and determine the exact time.

Figure 11 • 9xBC Sync Timestamp PDV Distribution – With SyncE



When compared, the time error distribution for the [TC chain in the 9xTC Sync Timestamp PDV No SyncE](#) (see [page 20](#)) figure, with its (relatively) high frequency distribution, can easily be filtered out by a low-cost oscillator at the slave. The filtering at every node for the BC chain means that a much more stable and expensive oscillator is required to filter out the residual time error at the end application.

Filtering the dynamic time errors in BCs using very low bandwidth filters can remove much of the dynamic time error, but causes the oscillator noise of the BC to become significant and contribute to the overall time error. The PTP ordinary clock slave at the end application will then need to have an even lower bandwidth filter to filter out the oscillator noise from the BCs, and this requires that the noise of the local oscillator at the end-node is stable to perform this filtering. This is similar to the normal layer 1 frequency transfer in telecom networks like SyncE, where the clock distribution is made using a chain of EEC devices with a required filter bandwidth of 1 Hz–10 Hz followed by an SSU with a lower filtering bandwidth.

The filtering of the dynamic timestamp error in a T-BC or T-TSC can also result in the generation of a static time offset, or an offset that is changing very slowly. This noise then transfers directly through the rest of the BCs in the chain and reach the connected slaves.

To reduce the cost of the slave oscillator at the Small Cell, it is therefore advantageous to deploy TCs as much as possible in the backhaul network.

8 Summary and Recommendations

The key points and recommendations condensed from the discussion, and important for the proper design of the backhaul network and base stations are listed as follows.

- The overall time error budget for TDD-LTE and LTE-Advanced is very tight and that every node and link in the backhaul network should be designed for minimal time error contributed to the overall time error allowed at the air interface.
- Not all network nodes and links are equal from a PDV and therefore timing precision perspective. The industry and standards bodies therefore need to look at different time error budgets for different classes of equipment and links rather than a uniform time error budget per hop.
- A typical backhaul network may consist of a series of fiber-connected switches and routers which could achieve maximum time errors as low as 20 ns–30 ns with properly designed TC implementations, while MW and MMW links even with properly designed Distributed TC architectures may have to be allowed to have as much as a 100 ns time error. Base stations /eNodeBs themselves also could have time errors as little as 20 ns–30 ns when using high-precision timestamping PHYs.
- The maximum time error of the entire backhaul network simply becomes the sum of all time errors allocated to each network element or hop.
- TCs are preferred in the backhaul network even if a network element can support both TC and BC. Although BCs can achieve sub-30 ns time error budgets per link, particularly with SyncE assist, the low-pass filtering at every BC node leads to low-frequency dynamic time error that can only be filtered out at the base station using very high stability oscillators, which would significantly contribute to the cost at the base station, in particular for small cells. The high-frequency noise from TC nodes can be filtered out much easier using low-cost oscillators. Also note, BCs are cascaded PLL's and will require significant time for the chain to relock in case of a perturbation or BMCA switch.
- BCs should be used strategically throughout the network to reduce the traffic load on the grandmaster clock.
- Timestamping should be performed at the PHY ingress/egress point to eliminate serialization and deserialization delay asymmetries, and measure dynamic packet delay variations across each network element as accurately as possible.
- High timestamping accuracy and a Gaussian timestamp error distribution are important to maximize performance in both BC and TC designs. No filtering algorithm can correct for inaccurate, and in particular non-Gaussian timestamp error distribution.
- Microwave backhaul equipment needs to implement distributed TCs across each link to minimize time errors from the modem and adaptive modulation schemes that can easily exceed the overall time error budget of the network. This is particularly important for small cell backhaul where multiple MW or MMW links may be daisy chained in an urban environment.
- Each base station macro and small cell alike should not contribute in a significant way to the time errors. This implies accurate timestamping must be done directly at the PHY I/O interface. In particular, in combination with TCs implemented in the backhaul network, this can lead to simpler servo algorithms that translate to cost-effective PTP clock slave implementations. Reducing BOM cost in highly cost sensitive small cells is particularly important.

**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

© 2019 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.

VPPD-03314