Mobile Backhaul Synchronization Architecture

A Technical Paper prepared for SCTE/ISBE by

Jennifer Andreoli-Fang, PhD
Distinguished Technologist
CableLabs
Boulder, CO
303-661-3838
j.fang@cablelabs.com

John T. Chapman
CTO Cable Access and Cisco Fellow
Cisco Systems
San Jose, CA
408-526-7651
jchapman@cisco.com

Table of Contents

List of Figures

List of Tables

Introduction

The growth in mobile data consumption has been putting pressure on the mobile network operators (MNOs) to build out small cell networks. All this traffic needs to be backhauled to the mobile core. While traditional choices for backhaul focus on fiber and microwave, hybrid fiber coaxial (HFC) networks have been making advancements. HFC is now being considered as a backhaul contender by the MNOs thanks to its capacity growth, cost efficiency and speed of deployment.

Traditional mobile base stations need to be frequency synchronized to guarantee handover performance, and this service is provided by the backhaul. In the DOCSIS 3.1 specification, the DOCSIS Time Protocol (DTP) was designed into the DOCSIS 3.1 specification to support precision timing from the CMTS to the cable modem (CM). This would allow a CM to provide backhaul services to a mobile base station for backhauling via the DOCSIS link. However, DTP is just one piece of the puzzle, as it needs to work with other elements of the operator network to provide timing to the base stations. This synchronization framework has yet to be defined. Furthermore, each operator network has differing levels of timing support in their existing hardware. This complicates system level designs.

In addition to frequency synchronization, Long-Term Evolution Time-Division Duplex (LTE-TDD) and LTE-Advanced features such as coordinated multipoint (CoMP) and enhanced inter-cell interference coordination (eICIC) all require stringent time and phase synchronization. Supporting these features places additional requirements on the synchronization framework.

In this paper, we review the technologies that can support frequency, time, and phase sync. We propose several architecture options, discuss their corresponding deployment scenarios, and the implications of each option on operations, cost of ownership, and time to market. Finally, we make recommendations on the device requirements and identify optimal designs based on operator deployments.

Drivers for Modern Backhaul Synchronization Requirements

The LTE downlink air interface utilizes orthogonal frequency division multiple access (OFDMA), while the LTE uplink uses single-carrier frequency division multiple access (SC-FDMA). OFDM is attractive for high speed wireless communications mainly due to its ability to combat frequency selective fading in multipath environments without the need for complex equalization techniques. But OFDM also requires orthogonality between the OFDM subcarriers, i.e., that 2 consecutive subcarriers must be non-overlapping in spectrum. Errors in frequency synchronization lead to loss of frequency orthogonality that can cause inter-carrier interference (ICI). In LTE systems, evolved node B (eNB) and user equipment (UE) must be frequency synchronized to 50-250 parts per billion (ppb) to allow the UE to demodulate LTE signals correctly, and to be able to transmit on the uplink.

While traditional macrocell networks only require frequency synchronization, the expected proliferation of small cells poses new challenges on timing distribution both technically and financially. What drives these new challenges is the focus of this section.

Before we begin though, depending on the deployment scenario, it is possible that a small cell deployment does not place further constraints on the synchronization requirements. Table 1 shows a list of deployment scenarios. For example, in rural outdoor deployments where the cell sites will always be able to receive signals from the Global Navigation Satellite System (GNSS) due to having a clear view of the sky, GNSS can be deployed in each cell site. Another scenario is if the traditional LTE FDD (Frequency Division Duplex) mode is deployed, only network frequency synchronization is required using PTP (Precision Time Protocol) which is part of IEEE-1588v2.

Table 1 – Deployment scenarios

	LTE-A Interference Management	Sync Requirements	Sync Methods
Dense urban outdoor (hotspot)	Needed	Frequency, time, phase	GNSS or PTP
Dense urban indoor (venue, MDU, enterprise)	Needed	Frequency, time, phase	No GNSS visibility → PTP
Suburban indoor residential	Not needed	Deployment can be TDD, so frequency, time, phase	No GNSS visibility, need cheap sync method → PTP
Rural	Not needed	Frequency	GNSS

Apart from these particular scenarios, a small cell deployment will introduce additional, and rather stringent requirements on time and phase synchronization. Figure 1 shows the difference between the different kinds of synchronization. If the cable operators want to leverage their DOCSIS networks for wholesale backhaul business, or to backhaul their own small cell networks, they need to have a repertoire of tools to use to solve these technical problems. We will discuss the toolkit in the remainder of this paper after we discuss the drivers.

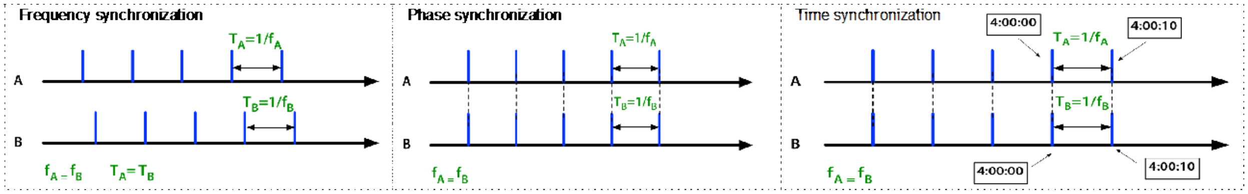

Figure 1 – Frequency, time, and phase synchronization

1. LTE TDD

With the expected deployment in the 3.5GHz spectrum in the US, the popularity of TDD has grown. LTE TDD is also prevalent in Europe. But TDD requires tight time synchronization.

In LTE TDD, uplink (UL) and downlink (DL) transmissions occur at the same frequency but are separated in time. The eNBs have to easily inform the UEs in the cell whether they should be listening or transmitting. The 3GPP defines 7 TDD subframe configurations so that the UEs know which subframe is for transmit or receive, although most small cells today support subframe configurations 1 and/or 2 only.

The TDD frame structure for subframe configuration 1 is shown in Figure 2. A special subframe denoted as the "S" subframe is defined to include a partial UL and a partial DL subframe, with a "guard period" sandwiched in the middle for switching between the UL and DL transmissions.

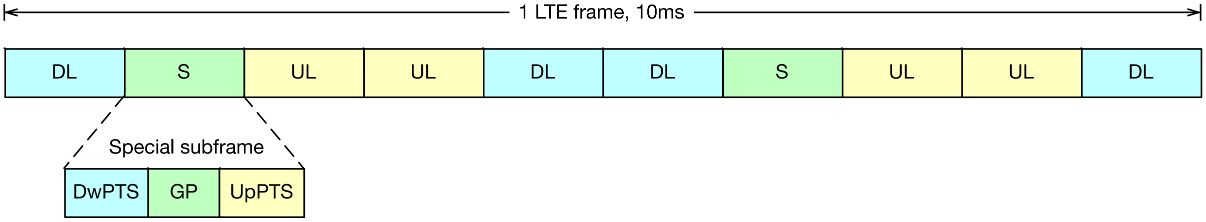

Figure 2 – Frame structure for TDD subframe configuration 1

To ensure maximum spectrum reuse, the eNBs operate in the same frequency. Additionally, to minimize interference, a cluster of eNBs are configured to use the same subframe configurations, so that they are either transmitting or receiving at the same time. Consequently, the adjacent eNBs must be synchronized almost perfectly to avoid an UL transmission interfering with a DL transmission in the neighboring cell. The 3GPP has specified in 0 that the neighboring eNBs must be phase aligned to within 3µs.

2. Heterogenous Networks and Dense Deployments

Small cell deployments are intended to address the ever-increasing mobile demands in both indoor and outdoor scenarios. For outdoor deployments, small cells are deployed in the same coverage area as the macrocells to fill in the capacity gaps. It is preferable to deploy small cells in different spectrum from the macrocells, but it is not always possible, due to limited spectrum availability. In co-channel or in-band deployments, small cells operate on the same frequency as the macros to maximize spectrum utilization. Such networks are called heterogeneous networks, or HetNets. The small cells in the HetNets experience inter-cell interference, because the macrocells transmit at significantly higher power levels.

As a large percentage of mobile traffic is consumed indoors, operators need to deploy ultra-dense small cells to fulfill the capacity needs. These co-channel eNBs situated in close proximity cause interference to one another, particularly at the overlapping cell edges.

The operators need to implement interference management techniques to address the interference issues unique to small cell deployments.

Traditional LTE includes simple physical (PHY) layer techniques such as heavy coding or OFDM's built-in cyclic prefix to combat interference. However, the techniques have all been designed for single cell operation. In case of HetNets and ultra-dense deployments, these methods are not enough.

To address this, a number of LTE Advanced (LTE-A) interference management techniques have been developed. We will now look at 2 techniques: eICIC and CoMP.

Ultimately, these techniques, while improving the small cell system capacity, pose stringent requirements on both synchronization and latency.

2.1. ICIC, eICIC

Suppose we have two neighboring eNBs operating on the same frequency. The UEs situated in the overlapping coverage area will experience high interference. This is because while the eNBs transmit to the UEs situated at the cell center with low power, they must transmit at higher power to the UEs at the cell edge in order to reach them with good enough SINR (signal-to-interference-plus-noise ratio). The situation is depicted in the left side of Figure 3.

Rather than transmitting blindly to the edge UEs with high power that would cause severe interference at the UEs, the two eNBs exchange information about what portion of the frequency spectrum it is planning to transmit with high power. In this way, the interference posed on each UE’s data channel (PHY downlink shared channel, or PDSCH) is reduced. The right side of Figure 3 shows an example situation: while eNBs A and B operate on the same frequency resource f1, A will transmit in resource f3 with high power, while B will transmit in resource f2 with high power. This is in essence a way for the eNBs to partition the spectrum, so that they would not be transmitting with high power in the same OFDM subcarriers. This technique, developed in LTE Rel-8, is called inter-cell interference coordination (ICIC).

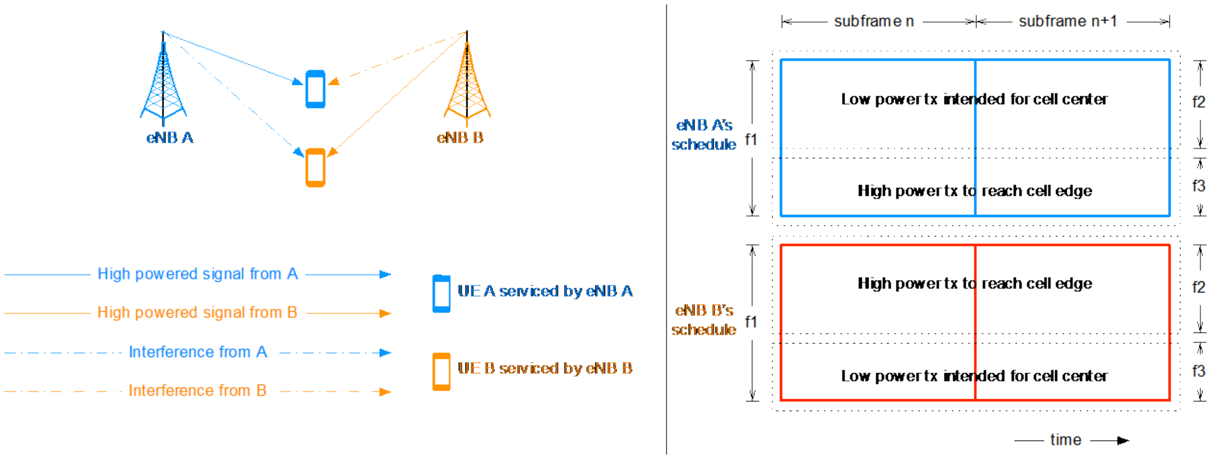

Figure 3 – Overlapping cells (left), and frequency domain inter-cell coordination (right)

While frequency partitioning works well for data channels, it does not solve the interference issue on the control channel. In each LTE subframe, the first 1-3 OFDM symbol(s) includes a broadcast control channel, i.e., the PHY downlink control channel (PDCCH), as shown in Figure 4, that includes subframe format indication and how the subframe is being scheduled to each UE. This channel must be received correctly in order for the UEs to decode the rest of the subframe.

To mitigate interference on the control channel, the LTE Rel-10 defines the eICIC with the concept of “almost blank subframe (ABS).” ABS is essentially subframe muting, and is shown in the left side of Figure 4. One of the eNBs provides information on which subframes in the near future it will mute, and sends this information to the other eNB. The negotiation involves message exchanges and takes place on the X2, which is a point-to-point logical interface between two eNBs.

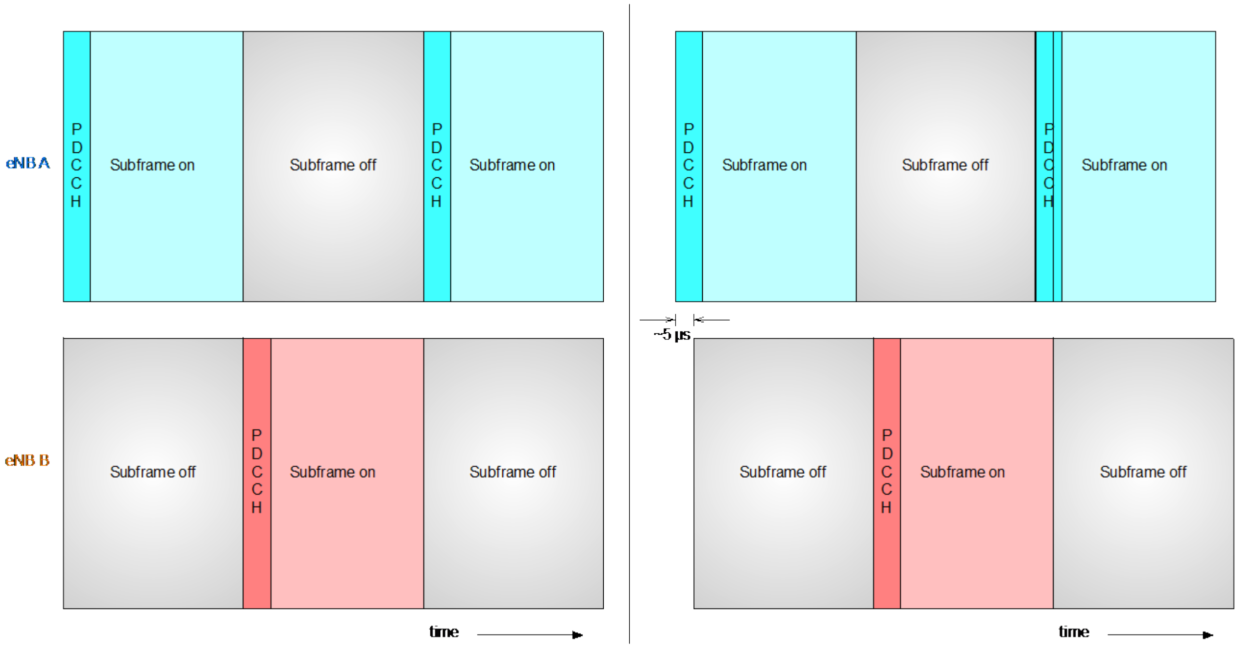

Figure 4 – ABS with perfect phase sync (left), and without (right)

ABS works when the two eNBs are in perfect phase synchronization. It is therefore critical for the clocks of the participating eNBs to be phase aligned, so that the subframes of the overlay eNBs do not overlap when one cell transitions its transmission to on while the other transitions to off. Otherwise, the PDCCH of one of the eNBs will experience severe interference as shown on the right side of Figure 4. The 3GPP does not formally define this phase sync limit. But various eNB vendors have quoted that the participating eNBs must generally be phase aligned within 5 µs, about the size of the cyclic prefix for the first LTE subframe symbol, in order for the technique to result in substantial performance gain.

2.2. CoMP

While eICIC improves the interference level experienced by the UEs at the cell edges, the UE’s throughput is limited to what can be achieved in a single cell due to the frequency and time partitioning of the spectrum and airtime. CoMP, featured in LTE Rel-11, enables multiple eNBs to simultaneously serve the UEs residing at the cell edge at the same time, analogous to a MIMO system, to increase the signal level and thereby achieve better edge UE throughput. Furthermore, while eICIC works on a semi-static time frame which is not suited for fast-changing channel conditions, CoMP allows eNBs in the coordinating set to negotiate resources dynamically.

The 3GPP defines several types of CoMP: coordinated scheduling (CS) including beamforming, and joint processing, including dynamic point selection and joint transmission.

CS is in essence a dynamic version of ICIC, but with frequency resource partitioning occurring dynamically at every subframe. The left side of Figure 5 shows a snapshot of signal vs. interference in subframe n after 2 eNBs have coordinated their scheduling. When eNBs A and B operate on the same frequency resource f1, it is possible through CoMP signaling to optimize the bandwidth usage. In this example, eNB A will transmit in resource f3 to reach its edge UE A, while eNB B will transmit in resource f2 to reach its edge UE B at subframe n indicated in the figure. The right side of Figure 5 shows that the spectrum resources are partitioned and that the scheduling of frequency resources can adapt dynamically on a subframe-by-subframe time scale.

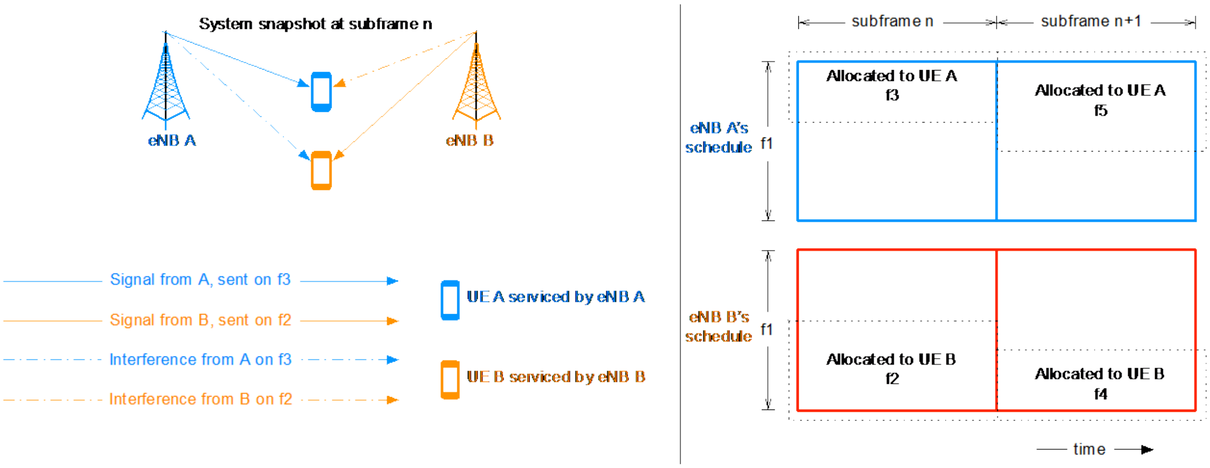

Figure 5 – Coordinated scheduling: snapshot of signal vs. interference for subframe n (left), and subframe scheduling after negotiation (right)

For CS, the data is only available at the UE's serving cell. In our case, eNB A serves as the master eNB for UE A, and eNB B serves as the master eNB for UE B. Scheduling and beamforming decisions are made by using the channel state information (CSI) shared between the eNBs in the coordinating set.

With joint processing, user data is available at multiple eNBs. Dynamic point selection is one type of joint processing. It is similar to CS in that only a single eNB transmits to an edge UE at a given time. The difference is that any eNB can serve an edge UE, compared to just the master eNB in the CS case. Figure 6 shows an example.

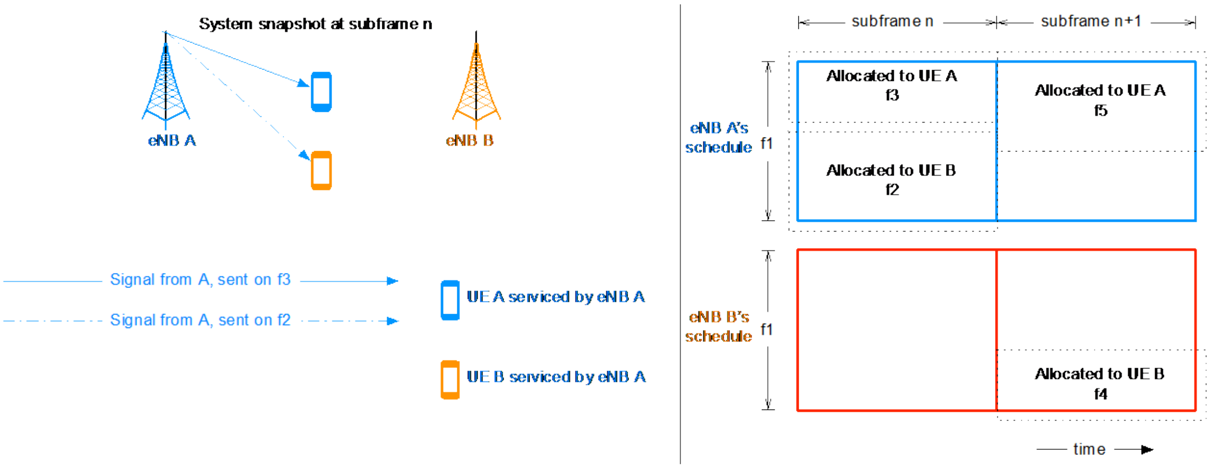

Figure 6 – Dynamic point selection: snapshot of signal vs. interference for subframe n (left), and subframe scheduling after negotiation (right)

With joint transmission, multiple eNBs can send the same data simultaneously to the edge UE at the same time and frequency resource. This improves the received power level at the UE and therefore improves the throughput.

A CoMP resource coordinator coordinates the schedules between multiple eNBs. It can reside in the eNBs in a distributed fashion, or be close to the evolved packet core (EPC) in a centralized fashion. Figure 7 gives a high level view of how CoMP works when a resource coordinator (RC) is located centrally. Referring to the steps in Figure 7: the UE in CoMP mode measures the CSI from all eNBs it can hear and sends CSI feedback to its master eNB (step 1). The eNBs forwards the CSI from the CoMP UEs to the Resource Coordinator (RC, step 2). The RC performs scheduling functions (step 3), and the scheduling information is then conveyed back to each eNB (step 4). If the CSI is delayed on the X2 interface, then the performance gain for the CoMP UEs will degrade.

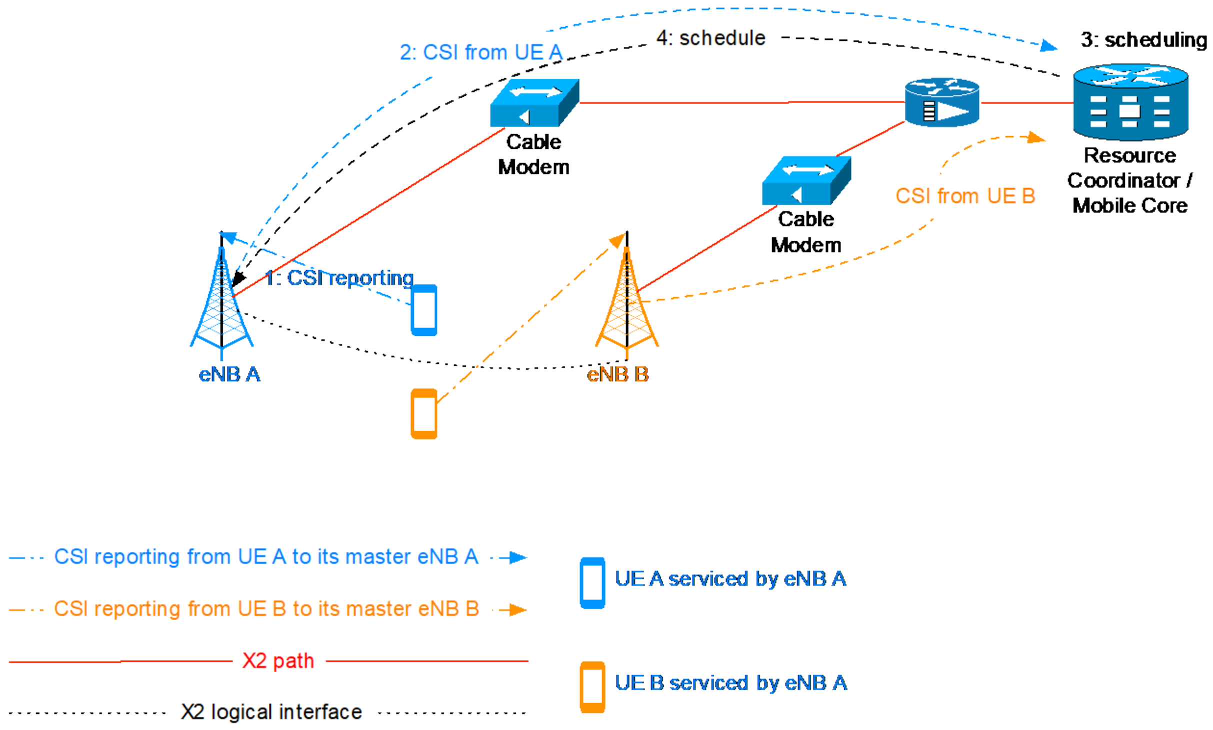

Figure 7 – CoMP operations

So, to support inter-eNB CoMP, the clocks of the neighboring eNBs must be time and phase synchronized to align the radio frames transmitted from different eNBs to the UE. The UE performance degrades with less accurate time and phase synchronization; and the amount of degradation depends on the CoMP technique.

In addition to the phase synchronization requirement, clearly, CSI must be sent expeditiously for the information to stay relevant and allow the cells to coordinate scheduling according to the dynamically varying channel conditions. This leads to a set of latency requirements (and possible solutions), which is covered by the authors in 0.

3. eMBMS / MBSFN

Another LTE feature is the enhanced multimedia broadcast multicast services (eMBMS), with a common use case being mobile broadcasting of live sporting events. It is supported in LTE over the multicast broadcast single frequency network (MBSFN). MBSFN allows multiple eNBs to transmit identical waveforms at the same time and frequency resources. The UE combines the multiple waveforms as multipath components of a single eNB. So the synchronization requirement is driven by the OFDM cyclic prefix in order to avoid inter-symbol interference.

4. Summary

Table 2 summarizes synchronization requirements for LTE TDD, and LTE-Advanced features such as CoMP, eICIC, eMBMS described in this section.

Table 2 – LTE and LTE-Advanced synchronization requirements

	Frequency	Phase	Notes
LTE FDD	± 50 ppb	None	3GPP TS 36.104 0 §6.5.1
LTE TDD	±50 ppb (wide area) ±100 ppb (local area) ±250 ppb (home)	10 µs (wide: cell radius >3km) 3 µs (local: cell radius <3km) 1.33 µs + Tprop (home eNB radius >500m) 3 µs (home eNB radius <500m)	3GPP TS 36.133 0 §7.4.2
CoMP	None	±1.5 µs	
eICIC	None	±1.5 – 5 µs	
eMBMS / MBSFN	None	±10 µs	

Synchronization Toolkit

There are 3 types of technologies to provide synchronization: physical layer mechanism such as synchronous Ethernet (SyncE), packet-based method such as precision time protocol (PTP) as defined in IEEE-1588v2, and GPS-based method. The latter two can also support time and phase synchronization.

Mobile operators today have various solutions in their repertory to distribute frequency, time, and phase synchronization. One straightforward method shown in Figure 8 uses a distributed architecture where the reference signal is distributed through the satellite signals. The GNSS receiver extracts the signal and is co-located with the end applications, in this case, the cellular radio sites. As such, this method does not require timing support from the network elements. This is an important point, because attaching a GNSS receiver to each cell site versus upgrading all or part of the network elements to support timing distribution poses financial tradeoffs for an operator, a point of consideration with dense small cell deployments. However, a GNSS-based method is not useable in most indoor deployment scenarios, and is not always reliable in the outdoor scenarios either due to atmospheric effects. So, operators must consider alternate technologies, which will be the focus of the remainder of this section.

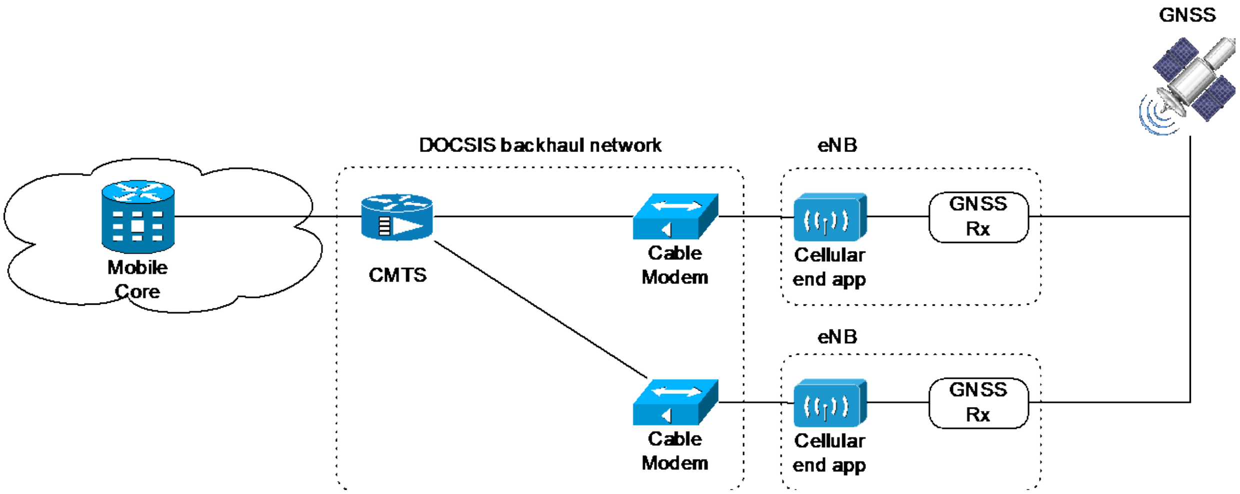

Figure 8 – Distributed GNSS-based synchronization without network support

5. Synchronous Ethernet

Synchronous Ethernet (SyncE) is a physical layer based frequency distribution algorithm. It is based on the Ethernet standards but additionally transmits a PHY transmit clock. In traditional Ethernet, a free running PHY clock is transmitted over the network medium. In SyncE, the Ethernet transmit clock is locked to a reference clock. Because the clock is continuously transmitted, SyncE is not subject to impact of the noise sources as PTP and Network Time Protocol (NTP). As such, SyncE is used as a mechanism to stabilize the frequency clock in case of failure events, i.e., when the timing can no longer be traced to a Primary Reference Time Clock (PRTC).

6. IEEE 1588 Precision Time Protocol

The IEEE 1588 precision time protocol 0, also known as PTP, is a packet-based synchronization technology that provides time, frequency, and phase synchronization. Since numerous tutorials exist on PTP (e.g., 0), this section will only provide a brief overview.

The basic principle is to distribute time sync reference by means of a 2-way timestamp exchange. As the mean path delay is half of round trip delay, the basic assumption is that a symmetric path between the packet master and the packet slave is required.

In contrast to SyncE that defines the timing content of the signals based on the significant edges of a data signal, PTP relies on the transmissions of timing messages. The series of time messages allows a PTP slave to recover the clock by estimating its timing offset from a PTP master.

The level of precision of the recovered clocks is contingent upon the PTP client’s ability to filter out the noise sources that will affect the accuracy of the recovered clock. The operator also has the ability to build out the network to minimize the effect of the noise sources, which include but are not limited to:

- Reference clock drift
- Timestamp error at the PTP grandmaster and slave
- Packet delay variation (PDV)
- Network asymmetry

To reduce the PDV that results from the queuing delay of the event messages, IEEE 1588 defines boundary clocks (BC) and transparent clocks (TC). Both BC and TC are switches or routers that participate in the timing protocol but in different ways. The BC terminates the PTP protocol, i.e., all PTP messages, on its slave port, uses the timing message to set its clock, and regenerates the PTP messages on its master port(s). In contrast, the TC does not set its clock based on the event messages, but instead, adjusts the event message timestamp to reflect the propagation time for the message to traverse through the equipment.

7. ITU-T G series recommendations

There has been confusion about the relationship of the PTP and the telecom profiles. The protocol as defined by the IEEE 1588 committee is the protocol that defines a set of message exchanges between two nodes. PTP alone does not guarantee meeting the end application’s performance requirements. Equipment that implement the PTP may not interoperate with each other, and may not satisfy any end application performance requirements. So, the International Telecommunication Union (ITU) worked and agreed on a set of architectures, telecom profiles, and performance specifications, all aimed towards meeting the performance requirements for telecom applications. In this section, we will outline the set of ITU-T Recommendations, what they are, what the relationship between them is, and as an operator, which Recommendations should be the focus depending on the deployment scenarios.

The ITU published a comprehensive set of Recommendations for distributing frequency, time, and phase synchronization, particularly geared towards telecom applications, since they have been the ones driving the tightness of the clock requirements. Within the set, the following are for time and phase synchronization:

- G.8260: general definitions and metrics
- G.8271 0: methods for distributing phase and time
- G.8271.1 0: maximum network limits on time errors and requirements on network elements
- G.8272 0: performance requirements for PRTC and T-GM
- G.8273: packet based phase / time clocks
- G.8273.1: performance requirements for T-GM
- G.8273.2 0: performance requirements for T-BC and T-TSC
- G.8273.3: performance metrics for T-TC
- G.8275 0: general architecture for distributing time and phase sync using PTPv2
- G.8275.1 0: PTP profile assuming full timing support from the network
- G.8275.2 0: PTP profile assuming partial timing support from the network

In particular, the operator should first focus on G.8275, where general architecture, along with protection mechanisms (holdover, which we will discuss shortly) are defined. It also specifies telecom-specific clock types, which are more rigorously defined than in the IEEE specs.

7.1. Types of telecom clocks

A primary reference timing clock (PRTC) is capable of providing frequency, time, and phase synchronization for other clocks in a network, by providing reference signals to a telecom grandmaster (T-GM). It is typically traceable to a universal time standard such as the UTC obtained from GNSS. G.8272 specifies the accuracy requirements for the PRTC and the T-GM.

A telecom boundary clock (T-BC) is an IEEE 1588 boundary clock with additional performance requirements defined in G.8273.2.

A telecom transparent clock (T-TC) is an IEEE 1588 transparent clock with additional performance requirements yet to be defined.

A telecom time slave clock (T-TSC) is an IEEE 1588 ordinary clock with only a slave port (i.e., cannot be a grandmaster) with additional performance requirements defined in G.8273.2.

7.2. Telecom profiles

The IEEE 1588-2008 introduced the concept of “profile” which includes a specific set of modes of operations, messages, message rates and attributes designed to satisfy an end application’s requirements.

The G.8275.1 telecom profile requires the network to provide full timing support. That is, boundary clocks must be implemented at every network node on the timing distribution path between the PTP grandmaster and the client. The profile defines a set of PTP parameters used to guarantee interworking between implementations. It specifies aspects such as the PTP messages to be used in the profile, 1-step vs. 2-step masters, message rates, protections, etc.

In contrast, the G.8275.2 telecom profile only requires the network to provide partial timing support. This really means that not every node in the timing distribution chain has to fully participate in the PTP, or to satisfy the performance requirement for T-BC. It is designed for operators who have no full timing support capability and cannot upgrade every switch and router in the timing distribution chain immediately.

The G.8275.2 telecom profile introduces additional clock types such as T-BC-P, T-TC-P, T-TSC-P, and T-TSC-A, where “-P” indicates “partial,” and “-A” indicates “assisted.” But performance characteristics of these clock types have yet to be defined.

Table 3 contains a summary of all the synchronization technologies discussed in this section.

Table 3 – Summary of synchronization technologies

	Pros	Cons
GPS/GNSS	• Global coverage with great precision	• Penetration is poor in indoor and dense urban with high rise • Upgrading every client to GNSS capability is expensive • Susceptible to jamming • Can be expensive if every cell site requires a receiver

	Pros	Cons
Packet-based, e.g., PTP	• Capable of providing frequency, time, and phase sync • Can be implemented for any deployment locations • Not every cell site must be upgraded (but still needs equipment upgrade on timing distribution chain)	• Performance accuracy subject to noise sources • Operators must do careful testing and measurements of the entire timing distribution chain in order to ensure end application performance requirements can be met
SyncE	• PHY layer technology means it is not subject to the noise sources from packet-based distribution mechanisms • Can be used in conjunction with other protocols to increase holdover performance	• Does not support time and phase sync • Point-to-point protocol means if a node in the chain is broken, synchronization for client cannot be achieved • P2P protocol means must upgrade every switch in the chain

Time and Phase Distribution over the DOCSIS Network

As discussed in Section 4, timing and synchronization requirements for small cells are stringent compared to the macrocell, due to interference management techniques and the use of TDD that may be required for small cell deployment. On top of this, DOCSIS is a packet based network. As such, it has the issue of network asymmetry. This makes distributing timing synchronization even more challenging. Luckily, DOCSIS Time Protocol that has been defined as part of the DOCSIS 3.1 specifications 0 several years ago. We will not be discussing the DTP in this paper. But for an excellent tutorial on DTP, see 0.

8. General architecture for phase and ToD distribution over DOCSIS backhaul

The general reference architecture for distributing time and phase using the DOCSIS backhaul network is shown in Figure 9. The PRTC provides timing reference for the timing distribution chain to the end application or the client. The PRTC can get the reference from a GNSS signal. An additional physical layer frequency synchronization signal can be included in the form of a primary reference clock (PRC, used for frequency synchronization only). We will discuss how SyncE can help improve the stability and accuracy during failure events shortly.

The PRTC is attached to a packet master clock known as the T-GM that implements packet-based distribution protocol such as IEEE 1588-2008. From there, the timing distribution chain can include a series of T-BCs and/or T-TCs. Networks with full timing support will only implement T-BCs compliant with the ITU-T G.8273.2 0 spec. Networks with partial timing support can include T-TCs.

The distribution chain includes the backhaul network serviced by a DOCSIS network. In the DOCSIS portion of the chain, the CMTS and the CM can participate in the IEEE 1588-2008 timing protocol. If so, the CMTS-CM pair can form an IEEE 1588 Boundary Clock that terminates the PTP domain by the CMTS and regenerates the PTP timestamp by the CM.

The timing reference signal will eventually reach packet slave clock(s) in the form of Telecom Time Slave Clock(s) (T-TSCs). The T-TSC may be integrated with the end application, in this case, the eNB.

The general architecture provides requirement flexibilities depending on an operator’s use case, performance requirement, and total cost of ownership (TCO) and time-to-market needs.

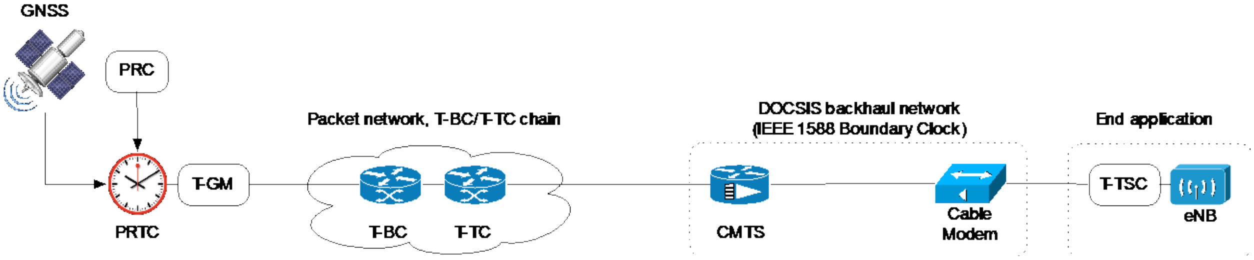

Figure 9 – General architecture to support phase and time distribution

The remainder of this section will cover various technology options for providing synchronization service using DOCSIS as backhaul. As we will see, each option has its merits and deficiencies.

9. Options for networks with full timing support

In the ideal case, each element in the operator network implements ITU-T G.8275.1 0 telecom profile, and is additionally compliant with certain performance requirements as specified in G.8272, G.8273.2.

9.1. G.8275.1 + DTP

This option uses the clocks implementing the G.8275.1 telecom profile to deliver time and phase synchronization throughout the distribution chain. Frequency synchronization is derived through time. The CMTS recovers time and frequency from the PTP messages, and translates the PTP timestamp into the DOCSIS timestamp. The CM regenerates the PTP timestamp based on the DOCSIS timestamp and passes it along to its downstream PTP slave which is the T-TSC. In this way, the CMTS-CM pair acts as a IEEE 1588 BC that terminates the 1588 domain at the CMTS, while the modem acts as a PTP master for the T-TSC in the end application client. Because the network is built with elements that are compliant with G.8275.1 and their corresponding ITU telecom standards (except the DTP domain elements), with time error budgeting, this option guarantees the delivery of frequency, time, and phase synchronization to LTE small cells that implement LTE-Advanced features such as eICIC, CoMP, and for LTE TDD deployments.

Implementing G.8275.1 on every element in the timing chain except the DTP domain provides guarantees that the time error for each link will be within the maximum absolute time error (max|TE|). We will discuss more in the time error budgeting subsection.

The operator has the flexibility of deploying one or more T-BCs between the T-GM and the CMTS. However, this option requires all network elements between the CMTS and the T-GM to be upgraded to be compliant with the G.8275.1 and G.8273.2. This could increase TCO. Alternatively, a T-GM can be collocated with the CMTS to avoid the upgrade.

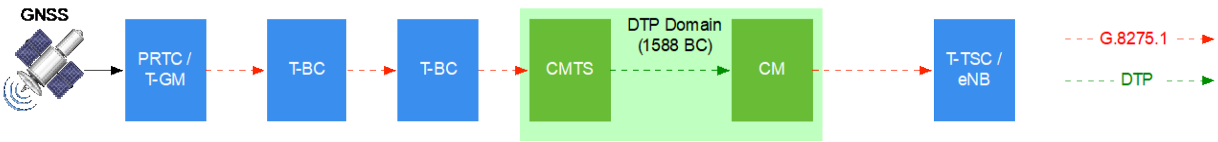

Figure 10 – G.8275.1 + DTP

Additional requirements on the CMTS and the CM include:

- The CMTS must support the IEEE 1588-2008 slave
- The CM must support the IEEE 1588-2008 master
- The CMTS and the CM must implement DTP

9.2. Time error budget analysis

In Section 4 Table 2, we described a list of LTE and LTE-Advanced features with their corresponding time and phase synchronization requirements. In order to deploy these features in their small cell networks, operators need to carefully design their network to distribute overall maximum absolute time error budget over each network element. In this section, we perform a sample time error budget analysis for the timing distribution chain to show what an operator would need to do to deploy LTE TDD in their networks.

LTE TDD operating mode requires 3 µs of phase synchronization between the adjacent home eNBs with radius of ≤ 500 meters. The following maximum absolute time error (|TE|) has been specified for each clock type by the ITU:

- |TE| ≤ 100 ns for PRTC 0. This allocation also works for combined PRTC and T-GM function
- A constant time error |cTE| ≤ 50 ns for Class A T-BC 0
- |cTE| ≤ 50 ns for Class A T-TSC 0

Additionally, a time error budget of 250 ns is assumed for holdover for the entire distribution chain, and 200 ns is assumed for dynamic time error budget (see 0 Appendix V Note 2). The total time errors incurred in the distribution chain in Figure 11 is:

$$|TE_{PRTC}| + N_{T-BC} \cdot |TE_{T-BC}| + |TE_{DOCSIS}| + |TE_{T-TSC}| + dTE' + TE_{HO} += 1500\ ns,$$

where dTE' denotes filtered dynamic time error (see 0 Appendix IV), and TE_{HO} denotes holdover error. Interested readers are directed to 0 for detailed discussion on constant, dynamic time error, and time error filtering.

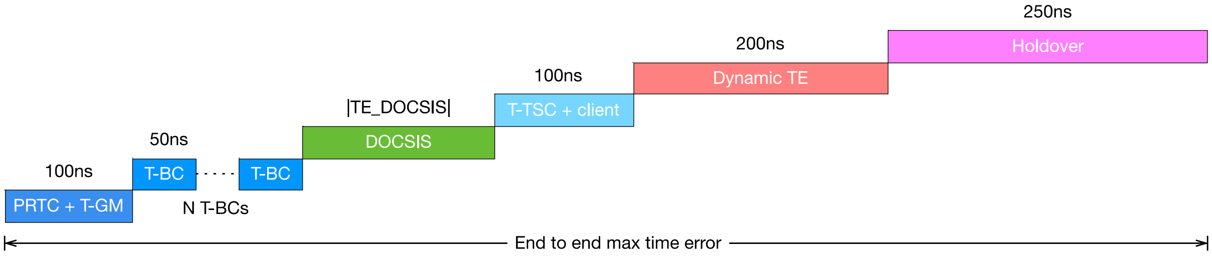

Figure 11 – Time error budgeting for networks with full timing support

DOCSIS 3.1 0 proposes 5 DTP system levels with varying degrees of CM-to-CM skew. Let us suppose an operator’s DOCSIS backhaul network satisfies the performance requirement for a Level IV DTP System as defined in Table 10-9 of 0. This means that the operator can allocate a total of only two T-BC in between the T-GM and the DOCSIS network. In other words, the PRTC must be located within 2 hops away from the CMTS.

Alternatively, if the DOCSIS equipment can instead satisfy a Level III DTP System, the operator can allocate 8 T-BCs, or 8 hops between the T-GM and the CMTS. This provides the operator the flexibility in architecting their timing distribution network and reduces the number of grandmaster clocks the operator must deploy.

While the 5 DTP system levels have been defined, further work may be needed to continue to refine the time error budgeting for each of the HFC network elements.

9.3. Protection mechanisms during link failure

As shown in the reference architecture in Figure 9, a series of master-slave clock pairs forms the timing distribution chain that extends from the grandmaster clock to the packet slave clock, or the eventual client which is the eNB. Since packet-based synchronization protocols rely on constant exchange of sync messages between the master and the slave, a disruption on a particular link means timing distribution is interrupted at the eventual client. When this occurs, the T-BC whose upstream link is disrupted will inform the client that the reference signal is no longer traceable to a PRTC.

Two protection mechanisms can allow time and phase to be continuously delivered to the client: redundancy and holdover. A network operator can deploy multiple PRTCs at different locations to provide redundancy. For the DOCSIS backhaul network, this means a CMTS is configured with communication paths to backup PRTCs. The T-BC involved in the failure event will run its PTP best master clock algorithm (BMCA) to look for a new PTP path, and thereby help the client to find a new PRTC and lock to a new traceable reference signal. During this period of network rearrangement, the client’s holdover mechanism can kick in to continue to generate clock from the last known traceable timing reference. Since holdover relies on the client’s free-running local oscillator, time error can accumulate during this period of network rearrangement. A better way of maintaining time accuracy is to have the client lock to a physical layer frequency reference, i.e., SyncE.

In the absence of a backup reference source, the synchronization stack on the client will enter the holdover state. Since there is no backup plan, the client needs to maintain accurate timing for a longer period compared to the period of network rearrangement until the reference source can be recovered. The 3GPP or the ITU does not define the holdover time and accuracy limits. Instead, typically, an operator specifies the time period and the accuracy limit for the equipment. Generally, a longer holdover period could be achieved with a higher quality oscillator. However, higher quality oscillators could be costly. While this does not cause an issue with macrocell deployments where cell radius is easily in the 10km range, the dense deployment in the small cell case can cause the cost of small cell equipment to escalate. Once again, having a PHY layer frequency sync support, i.e., SyncE, will improve the holdover performance.

9.4. G.8275.1 + DTP + SyncE

In addition to the requirements discussed in the “G.8275.1 + DTP” option in Section 9.1, each network element implements a Synchronous Ethernet Equipment Clock (EEC), as shown in Figure 12. While SyncE can improve system performance, it is worth noting that since SyncE requires specialized Ethernet hardware support, relying on SyncE will require the operator to replace all of its existing Ethernet equipment in the entire timing chain.

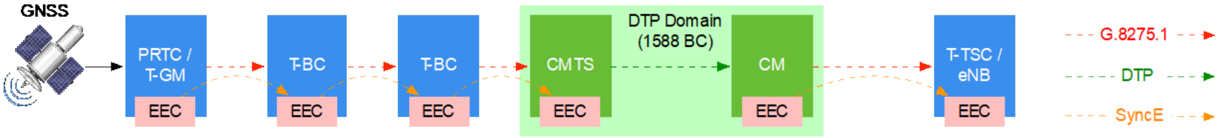

Figure 12 – G.8275.1 + DTP + SyncE

10. What if the network provides partial or no timing support?

In the ideal scenario that requires full timing support from the network, all network equipment must participate in the timing protocol and implement the G.8275.1 telecom profile. The timing distribution chain except the DOCSIS portion includes only T-BCs that are G.8273.2-compliant.

However, replacing and upgrading every network element to be compliant with G.8275.1 and T-BCs specs can become expensive, and the availability of equipment can be an issue. So, for operators who cannot upgrade every clock in their network immediately, an alternative option exists to still enable the delivery of frequency, ToD, and the phase synchronization needed to support LTE small cell deployments that implement LTE-A features and LTE TDD.

10.1. G.8275.2 + DTP (with or without SyncE)

In a network with partial or no timing support, non-participant or non-PTP-aware nodes, as well as T-TC(s) are allowed. In order to enable the delivery of frequency, time, and phase synchronization needed to service LTE-A techniques and LTE TDD, both the CMTS and the CM need to implement the G.8275.2 telecom profile and act as a T-BC-P. Note that the performance requirements for the T-BC-P node have not been formalized by the ITU.

As with the G.8275.1 option, the CMTS recovers time and frequency from the PTP messages, and translates the PTP timestamp into the DOCSIS timestamp. The CM regenerates the PTP timestamp based on the DOCSIS timestamp and passes it along to its downstream PTP slave which is the T-TSC-P. In this way, the CMTS-CM pair acts as a IEEE 1588 BC that terminates the 1588 domain at the CMTS, while the modem acts as a PTP master for the T-TSC-P in the end application client.

Since the performance of non-participant nodes are unknown, and most likely is worse compared to T-BC, the number of non-PTP-aware nodes, especially in a cascade, must be limited in a timing chain. One or more T-BC-P nodes can be placed strategically in the chain to reduce the effect of time error. To ensure required timing accuracy is achieved, the operator needs to perform proper testing, especially when the number of non-PTP-aware hops increases.

As discussed earlier, implementing SyncE at every node will improve frequency stability and holdover performance. However, requiring SyncE will require the operator to replace all of its existing Ethernet equipment with EEC-capable equipment in the entire timing chain – something the operator may have wanted to avoid in the first place by using G.8275.2 rather than G.8275.1. So, the options shown in Figure 13 and Figure 14 are well suited for installing a new Ethernet infrastructure for a new deployment region.

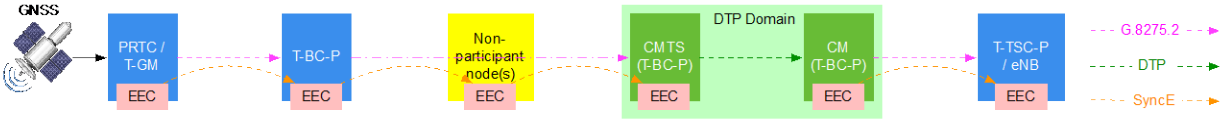

Figure 13 – G.8275.2 + SyncE + DTP

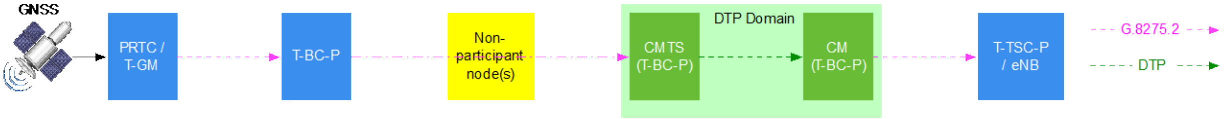

Figure 14 – G.8275.2 + DTP

11. What if DTP is not available

Since new modem silicon may be needed to implement DTP and PTP, none of the options discussed so far is deployable if an operator wishes to deploy the synchronization solution before the next modem silicon cycle. But even with the modem being a non-participant node in the timing distribution chain, depending on the availability of DTP on the CMTS, 2 options exist for distributing phase and ToD reference signals.

11.1. With CMTS participation

If the CMTS implements G.8275.2, it can act as a T-BC-P node to terminate the PTP timing messages from one of its upstream PTP-aware nodes such as a T-BC-P. The CMTS can then regenerate the PTP messages, sending them over the top through the CM directly to the T-TSC-P which could be part of the end application. In this case, the DOCSIS link is PTP and DTP-unaware.

Carrying the timing messages without the support of the network, and in case of DOCSIS, as regular data, will introduce a host of time errors due to PDV and network asymmetry, among other things. At least on the DOCSIS link, the PTP messages should be carried with unsolicited grant service (UGS), real-time polling service (RTPS), or high priority best effort upstream service flows to reduce the time error. Additionally, without DTP, MAC layer asymmetry cannot be determined.

Despite increased time error due to non-participant CMs, this option, shown in Figure 15, reduces the time error accumulated in the chain compared to the next option where the CMTS does not participate in the timing protocol.

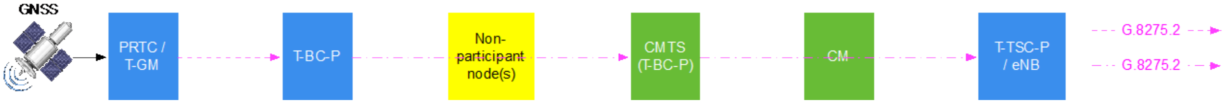

Figure 15 – G.8275.2 with CMTS participation

11.2. Without CMTS participation

If the CMTS does not participate in the timing protocol, neither PTP nor DTP, the PTP messages must be sent while both the CMTS and the CM are timing-unaware. Due to the accumulation of time errors from non-participant nodes, this option reduces the number of hops the end application can be placed away from the T-GM.

This option as shown in Figure 16, does have some advantages over other options discussed so far. Since no additional requirement is placed on the CM and the CMTS, the option is deployable today.

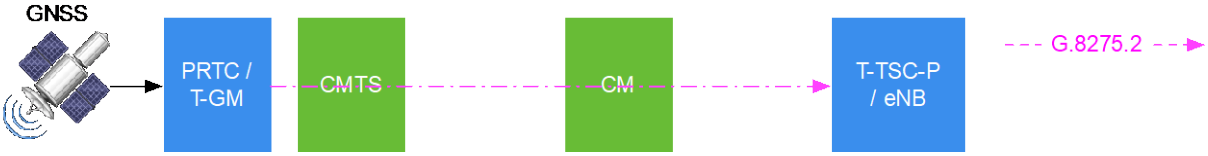

Figure 16 – G.8275.2 without CMTS participation

While the options illustrated in this subsection are deployable today, neither provides compensation for the time error introduced by the underlying HFC network. Thorough testing is required to understand the time error budget for the DOCSIS link. Therefore, while it may be possible to provide a coarse level of frequency synchronization with these options, it is highly improbable to provide phase and time synchronization with these options.

12. Summary

Table 4 compares the options discussed in Sections 9 – 11.

Table 4 – Comparison of DOCSIS-based options

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Use Case	Target solution for frequency/time/phase sync to small cells that implement LTE-A features and TDD. Allows install of 1 or more indoor and outdoor small cells on a new Ethernet infrastructure which supports SyncE, when improved holdover performance is required.	For delivery of frequency/time/phase sync to small cells that implement LTE-A features and TDD. Allows install of 1 or more indoor and outdoor small cells on a new Ethernet infrastructure which supports SyncE, when improved holdover performance is required. Allows cheaper T-GM install, if T-GM service is not directly provided by CMTS, or if T-GM is not collocated with CMTS.	Allows install of 1 or more indoor and outdoor small cells on an existing Ethernet infrastructure that does not support SyncE. If time error budget is large, CMTS does not need to be collocated with T-GM/PRTC. Allows cheaper T-GM install, if T-GM service is not directly provided by CMTS, or if T-GM is not collocated with CMTS. 8275.2 could be used as a workaround until DTP and 8275.1 are fully supported by entire network.	Allows install of 1 or more indoor and outdoor small cells on an existing Ethernet infrastructure that does not support SyncE. If time error budget is large, CMTS does not need to be collocated with T-GM/PRTC.

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Accuracy	DTP solves DOCSIS asymmetry issue. Frequency sync through SyncE assists and improves time sync. SyncE reduces sync time after first contact, which means quick sync / fast re-sync. Implementing 8275.1 provides guarantees to achieve max time error per hop.	To ensure required accuracy is achieved, proper testing is needed. Time for testing increases when non-PTP-aware hops increase. But, with proper testing and tuning, it is possible to achieve same accuracy as 8275.1. DTP solves DOCSIS asymmetry issue. Frequency sync through SyncE assists and improves time sync. SyncE reduces sync time after first contact, which means quick sync / fast re-sync.	Requires short chain between T-GM and CMTS. Proper testing and optimization is required to keep time error low. Time for testing increases when non-PTP-aware hops increases. DTP solves DOCSIS asymmetry issue.	Requires short chain between T-GM and CMTS. Proper testing and optimization is required to keep time error low. Time for testing increases when non-PTP-aware hops increases. Since only small cells and T-GM implement 8275.2, DOCSIS MAC layer can introduce time error and link asymmetry that cannot be corrected. Use of RTPS or high priority BE services to carry 8275.2 traffic is a must.
Operation	SyncE extends holdover in case PTP fails. By using frequency and time, lower requirements on local oscillators, or can extend holdover time by using same oscillators.	SyncE extends holdover in case PTP fails. By using frequency and time, lower requirements on local oscillators, or can extend holdover time by using same oscillators. 8275.2 provides failover and/or better holdover solution for DTP while transport is recovering from service disruption, e.g., after reboot of a CM or CMTS.	8275.2 provides failover and/or better holdover solution for DTP while transport is recovering from service disruption, e.g., after reboot of a CM or CMTS.	In case of link failure, does not need improved holdover time on local oscillators, since sync messages can be transported over IP, as soon as data link is re-established. Reduces the service disruption time after CMTS reboot, since lengthy sync time is not required when carried over IP without DTP.

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Total Cost of Ownership	Either T-GM needs to be collocated with each CMTS, or all elements between T-GM and CMTS must be upgraded, as 8275.1 requires compliance for every hop.	Depending on PDV, asymmetry, and # of hops, same hardware deployed today can be used for 8275.2. Allows for partial upgrade of select nodes, by providing a T-BC in strategic locations to lower time errors. Since SyncE is required, all elements in the timing chain must be upgraded. Also Ethernet switches in the premises must be upgraded to support SyncE.	Backward compatible with existing indoor infrastructure. Support use of very low cost equipment on sites. Works without upgrading any hardware or software in backhaul chain. Depending on PDV, asymmetry, and # of hops, same hardware deployed today can be used for 8275.2. Allows for partial upgrade of select nodes, by providing a T-BC in strategic locations to lower time errors.	Backward compatible with existing indoor infrastructure. Support use of very low cost equipment on sites. Works without upgrading any hardware or software in backhaul chain.

	G.8275.1 + SyncE + DTP	G.8275.2 + SyncE + DTP	G.8275.2 + DTP	G.8275.2
Time To Market	Radio vendors support 8275.1 today. At least 1 modem silicon vendor supports PTP and DTP in hardware today. May requires new modem silicon to support SyncE master and PTP master.	Depending on radio vendor support of 8275.2, CM support of DTP and PTP, and when Ethernet equipment at the premise and in the timing chain can be replaced with SyncE. At least 1 radio vendor is testing 8275.2, and will support it by EOY 2017. May require new modem silicon to support SyncE master and PTP master.	At least 1 modem silicon vendor supports PTP and DTP in hardware today. CM needs to implement IP stack on CPE facing interface. This is likely a new board design, not a new silicon. But will require additional time.	At least 1 radio vendor is testing 8275.2, and will support it by EOY 2017. Since this solution does not require DTP or PTP support on CM, allows for deployment by EOY 2017.

Conclusion

In this paper, we discussed the drivers for backhaul synchronization requirements needed for modern LTE and LTE-A networks, which are significantly more stringent compared to the traditional macrocell deployments. Although frequency, time, and phase synchronization can already be supported by today's technologies, to guarantee accuracy and holdover performance, operators must architect their network carefully with the right set of equipment, testing, and optimization. With the correct options, DOCSIS-based backhaul networks can support LTE-A features and LTE TDD deployments.

Abbreviations

3GPP	3rd Generation Partnership Project
ABS	almost blank subframe
BC	boundary clock
BMCA	best master clock algorithm
CM	cable modem
CMTS	cable modem termination system
CoMP	coordinated multipoint

CS	coordinated scheduling
CSI	channel state information
DL	downlink
DTP	DOCSIS time protocol
EEC	Ethernet equipment clock
eICIC	enhanced inter-cell interference coordination
eMBMS	enhanced multimedia broadcast multicast services
eNB	evolved node B
EPC	evolved packet core
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HetNet	heterogeneous network
HFC	hybrid fiber-coax
ICI	inter-carrier interference
ICIC	inter-cell interference coordination
ITU	International Telecommunication Union
ITU-T	International Telecommunication Union-Telecommunication
LTE	long-term evolution
LTE-A	LTE advanced
LTE-FDD	long-term evolution frequency-division duplex
LTE-TDD	long-term evolution time-division duplex
MBSFN	multicast broadcast single frequency network
MNO	mobile network operator
NTP	network time protocol
OFDM	orthogonal frequency division multiplexing
OFDMA	orthogonal frequency division multiple access
P2P	peer-to-peer
PDCCH	PHY downlink control channel
PDSCH	PHY downlink shared channel
PDV	packet delay variation
PHY	physical
ppb	parts per billion
PRC	primary reference clock
PRTC	primary reference time clock
PTP	precision time protocol
RC	resource coordinator
RTPS	real time polling service
SC-FDMA	single-carrier frequency division multiple access
SINR	signal-to-interference-plus-noise ratio
SyncE	synchronous Ethernet
T-BC	telecom boundary clock
T-GM	telecom grand master, master clock only
T-TC	telecom transparent clock
T-TSC	telecom time slave clock
T-BC-P	telecom boundary clock-partial
T-TC-P	telecom transparent clock-partial

T-TSC-A	telecom time slave clock-assisted
T-TSC-P	telecom time slave clock-partial
TC	transparent clock
TDD	time division duplex
TCO	total cost of ownership
ToD	time of day
UE	user equipment
UGS	unsolicited grant service
UL	uplink
UTC	universal time coordinated

Acknowledgement

The authors would like to thank Yi Tang and Zheng Lu, both of Cisco Systems, for reviewing the white paper and providing valuable comments. The lead author would like to thank the following people involved in the discussions of the DOCSIS-based mobile synchronization architecture: Alvaro Simon of Vodafone Spain, Tilco van Dijk of Liberty Global, Pedro Antão of NOS Portugal, Robert Grimm of Vodafone DE, and Bruno Cornaglia of Vodafone Italy. Additionally, the lead author acknowledges Peter Percosan of Digital Strategy, Volker Leisse and Thomas Nogues of CableLabs for facilitating the work.

Bibliography & References

John T. Chapman, et. al., “The DOCSIS Timing Protocol (DTP), Generating precision timing services from a DOCSIS system,” Proceedings of INTX/SCTE Spring Technical Forum, 2011.

John T. Chapman, Jennifer Andreoli-Fang, “Low Latency Techniques for Mobile Backhaul over DOCSIS,” to appear in *Proc. of SCTE Fall Technical Forum*, October 2017, Denver.

3GPP TS 36.104, “Evolved Universal Terrestrial Radio Access (E-UTRA); Base Station (BS) radio transmission and reception.”

3GPP TS 36.133, “Evolved Universal Terrestrial Radio Access (E-UTRA); Requirements for support of radio resource management.”

IEEE Std 1588-2008, “IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems.”

ITU-T G.8271, “Time and phase synchronization aspects of packet networks.”

ITU-T G.8271.1, “Network limits for time synchronization in packet networks.”

ITU-T G.8272, “Timing characteristics of primary reference time clocks.”

ITU-T G.8273.2, “Timing characteristics of telecom boundary clocks and telecom time slave clocks.”

ITU-T G.8275, “Architecture and requirements for packet-based time and phase distribution.”

ITU-T G.8275.1, “Precision time protocol telecom profile for phase/time synchronization with full timing support from the network.”

ITU-T G.8275.2, “Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network.”

“Data-over-cable service interface specifications, DOCSIS® 3.1, MAC and upper layer protocols interface specification,” I-11, issued 2017-05-10, CableLabs.