Considerations for Edge Master Deployment

About this Document

This document gives some general guidelines to assist in planning a packet-based synchronization strategy utilizing edge grandmasters.

The information given here is for guidance only and shouldn't be taken to imply the performance or suitability of any particular network architecture. Instead, it is important that the synchronization performance of the network is measured during deployment and ideally monitored on an ongoing basis to ensure expected performance metrics are being met.

Introduction

Traditionally, cellular base stations required only frequency synchronization (for example a frequency accuracy of better than 16 ppb) and were interconnected using wireline technologies such as T1/E1 or SONET (SDH) that allowed for a frequency source with excellent accuracy and long-term stability to be derived directly from the transport physical layer. This meant that the synchronization function of the base station (or other end equipment) consisted of just a low bandwidth PLL, potentially with a stable oscillator to provide holdover in the event the link is lost, and there was little special consideration needed to design of the network.

However, in recent years two significant changes have occurred that have a major impact on synchronization. Firstly, starting with CDMA and accelerating pace with LTE-FDD and LTE Advanced, base stations have increasingly required not only frequency control but also phase alignment for correct operation. For example, LTE-TDD requires adjacent cells to maintain a concept of time between them of as close as 3 μs in some cases. Secondly, networks have rapidly shifted to packet-based architectures, typically using Ethernet but also including other standards in the access portion, such as xPON, that typically don't provide the frequency transport that T1/E1 and SONET (SDH) did. Synchronous Ethernet, in which the physical layer of the Ethernet link is frequency synchronized in a very similar way to SONET (SDH), partly addresses this shift by restoring the ability to provide frequency delivery over the network but none of the traditional physical-layer based mechanisms can provide time synchronization. Therefore, an alternative approach is needed.

Two complementary techniques have arisen to provide time synchronization to base stations interconnected by packet networks:

- GNSS (GPS) – Each base station has a co-located GNSS receiver that provides time synchronization directly. There are however many drawbacks with this approach:
  - Cost
  - Practicality – for example deploying a GNSS antenna for a small cell buried deep in a building
  - Reliability – although GNSS systems are well established they are easily jammed, either deliberately or accidentally
  - Accuracy – the standard accuracy of GNSS-recovered time might not be good enough for some emerging requirements (see the discussion of fronthaul later)
- PTP (IEEE 1588-2008) – This uses a packet-based protocol to transfer time between a grandmaster (itself typically receiving time from GNSS) to one or more slave devices (the base stations) by timestamping the sending and receiving of the packets and applying various filtering algorithms to compensate for network delays. The timing packets can flow directly between the grandmaster and slave or can pass through one or more Boundary Clocks, which are PTP-aware network devices that effectively act as clock repeaters. An edge master, such as Orolia's EdgeSync, is a boundary clock targeted for deployment towards the edge of a network (for example, in the access portion). PTP can be combined with Synchronous Ethernet to further improve performance in some cases.

An emerging technology is that of Assisted Partial Timing Support (APTS) where PTP and GNSS are used together to provide a robust synchronization solution. In this scenario, each base station typically relies on PTP alone, but the intermediate boundary clocks can combine both PTP and GNSS time for improved accuracy and reliability. Alternatively, the base stations themselves could accept both GNSS and PTP time. Orolia's EdgeSync supports APTS operation, allowing synchronization to be taken from GNSS, PTP or a combination of both.

Backhaul vs. Fronthaul Requirements

An important consideration when designing a synchronization strategy is the time alignment required across base stations or other equipment. This, in turn, depends on the nature of the cellular network design but in general can be split into so-called backhaul and fronthaul applications. The backhaul portion of a network is that part that is used to connect base stations, or other devices providing baseband processing, to the network core, whereas the fronthaul portion is used to connect the baseband processing function to the devices providing the actual air interface (the radios). In a traditional cellular architecture, such as that used prior to 4G LTE, the base stations and radios were either fully integrated, or connected via a dedicated link such as CPRI, meaning that only the backhaul network portion is required and therefore only this needs to be considered for synchronization. However, newer networks with a distributed radio access network (RAN) introduce the fronthaul portion to interconnect the baseband processing unit with multiple remote radio heads.

To address the differing synchronization requirements of backhaul and fronthaul networks, a number of categories of synchronization have been defined as part of the IEEE 802.1CM standard, as shown in Table 1.

Category	Network Portion	Application	Required Accuracy
C	Backhaul	Baseband interconnect	1.36 µs
B	Fronthaul	RAN with non-contiguous carrier aggregation	110 ns
A	Fronthaul	RAN with contiguous carrier aggregation	45 ns
A	Fronthaul	RAN with distributed MIMO or TX diversity	10 ns

The required time accuracy defined for each category is derived by taking the alignment required for correct operation between adjacent nodes, dividing by two to allow for synchronization back to a common timing source (such as GNSS or a PTP grandmaster), since one device could be offset in one direction and the other in the opposite direction, and subtracting an allowance for time errors within the end equipment itself. For example, as stated earlier, the air interface of adjacent LTE-TDD small cells must be aligned to 3 µs or better. Dividing this by two gives 1.5 µs, from which 140 ns is subtracted as a local error allowance to give the 1.36 µs figure for category C.

As can be seen, the requirements for fronthaul networks are considerably more stringent than those for backhaul, with the A and A+ requirements even challenging the abilities of GNSS (which in general provides 100 ns alignment). In reality, 4G and first generation 5G networks in use or being planned now will need to meet category C for backhaul and category B for the fronthaul portion. However, it is expected that category A will be expected for second generation 5G networks, though there is currently no clear picture as to when, or if, class A+ will become a requirement.

In general, meeting any of the fronthaul synchronization categories (A, A+ or B) will require very well architected networks which will most likely use full on-path support (see the following section). Such networks are not the primary target of edge masters, such as the Orolia's EdgeSync, and therefore aren't covered further in this discussion.

PTP Architectures for Backhaul

The IEEE 1588-2008 standard that defines PTP deliberately allows for considerable flexibility to allow the protocol to be applied to multiple applications. It is then up to other organizations to define a subset of PTP for specific applications in the form of a PTP Profile. There are three profiles that are relevant to cellular applications, all defined by ITU-T:

- G.8265.1 – For frequency-only delivery over general networks
- G.8275.1 – For time delivery over networks in which all the switching and routing equipment operates as a boundary clock (so-called full on-path support), and in which Synchronous Ethernet, or other physical frequency source, is required
- G.8275.2 – For time delivery over networks in which only some (or potentially none) of the nodes operate as a boundary clock (partial on-path support)

APTS, mentioned earlier, is an extension of partial on-path support as defined by G.8275.2.

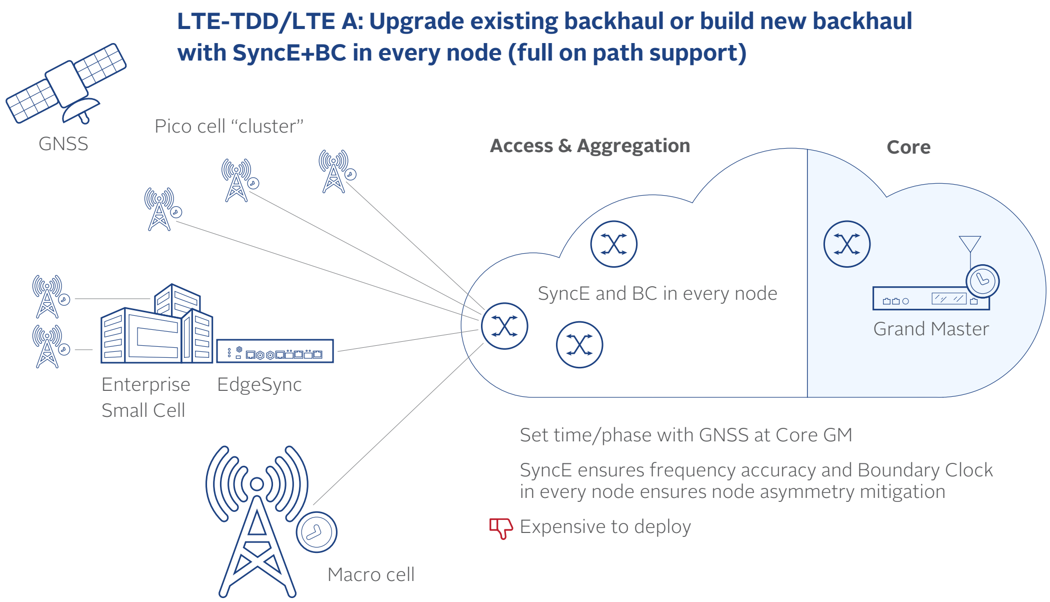

Figure 1 – Example G.8275.1 Network

Each profile defines many aspects of PTP operation, but a key one in terms of network design is that G.8265.1 and G.8275.2 typically operate using unicast messaging whereas G.8275.1 uses link-local multicast messages. Such messages are blocked by any compliant switch or router, effectively enforcing the requirement for each node to function as a boundary clock.

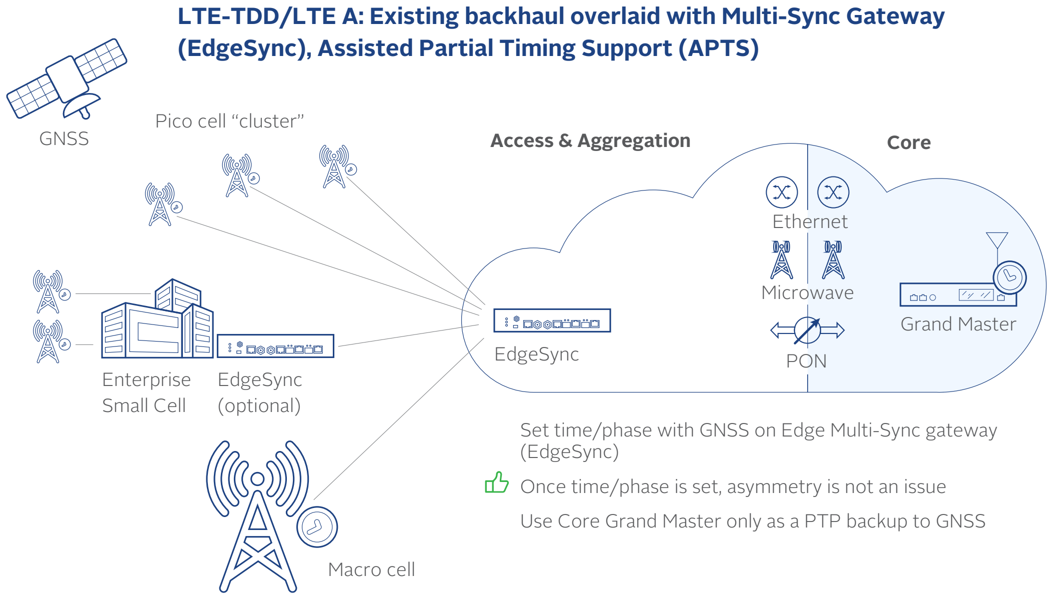

Figure 2 – Example G.8275.2 Network

In general, any given network will only utilize a single PTP profile.

Planning a PTP Deployment

The two primary considerations when planning a PTP deployment are:

- Which PTP profile to use
- What PTP equipment to deploy and where

With regard to the first of these, the profile to use may be dictated by an existing network if the new deployment is to be integrated into a network that already utilizes PTP, or it may be application-driven for a green-field deployment. In this latter case, it is likely that one would want to deploy a G8275.1 or G.8275.2 network to provide time alignment for future use, even if it's not immediately needed (for example in the case of an LTE FDD design) and the choice between the two will come down to required performance and budget.

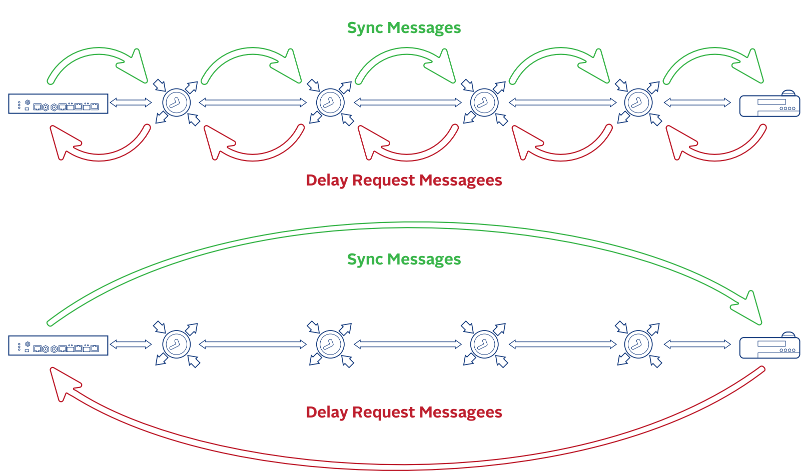

Figure 3 – On-path Support Network (top) vs. Network without PTP Support (bottom)

For the question of PTP equipment to be deployed, it has to be decided if it is desired to use PTP across the whole network or to deploy PTP in smaller local clusters towards the edge of the network. In the first case, at least one grandmaster must be deployed that can provide time to the entire network, with ideally one or more additional grandmaster to avoid the risk of a single point of failure. In the second case, lower-capacity edge masters, such as Orolia's EdgeSync can be deployed at multiple points around the edge of the network, each taking time from GNSS an providing synchronization to the local cluster of end devices downstream from the edge master. While a unified, PTP-capable, network might be considered the ultimate aim, deploying such a network is often far from trivial, especially if having to piece-meal update a legacy network. Therefore it is expected that many deployments will, at least initially, use the edge master-based distributed cluster model. For reasons outlined later, it may still be desirable to deploy edge masters in a PTP-capable network, either acting purely as boundary clocks or combining PTP with GNSS in an APTS application.

It must also be remembered that for a G.8275.1 network, every single switching or routing node in the network must also function as a PTP boundary clock, and Synchronous Ethernet must typically be deployed across the whole network. Again, this is likely to present a challenge when upgrading legacy networks, so it is expected that G.8275.2 will be the preferred profile supported by networks in the near future.

Considerations for a G.8275.2 Network Design

Based on the previous assumption that G.8275.2 will be the preferred architecture for backhaul networks, either with edge masters using GNSS only for isolated local clusters, or using PTP from higher up in the network and acting as a boundary clock, one then has to consider the optimal location of the edge masters.

There are a number of factors to consider in making this decision. These are listed below in order of importance:

- The effect of the network on time accuracy
- The network technology and asymmetry
- The ease of location for the edge master
- The number of slave devices per cluster
- Availability considerations

Each of these is considered in turn:

The effect of the network on time accuracy

This is by far the most significant consideration. The performance of PTP directly depends on the delays introduced in passing timing packets, in both directions, between the master and the slave. The absolute time for a packet to get from source to destination is not critical, but the variation in this packet-to-packet delay is. This is referred to as packet delay variation (PDV) and results from the fact that the PTP timing packets are only a small percentage of the overall network traffic, and which are therefore competing with other packets to pass through the network.

To understand this, consider packets flowing from master to slave through a single Ethernet switch without any other network traffic. The time taken by each packet will likely vary a little bit due to timing alignment within the switch etc., but such variation is likely to be insignificant compared to the actual transport time. Such a situation is described as having low PDV. Now, however, imagine that other packets, for example data to the base station, are arriving on a different switch port but being routed out on the same port as the PTP packets. Because Ethernet doesn't allow for preemption of frames, any given PTP packet may have to wait in an outgoing queue within the switch for an already in-progress packet from the other port to be completed. Such queuing, called head-of-line blocking can introduce considerable PDV. Most switches allow individual traffic types of be given different priorities, and it should always be the case that PTP traffic is treated with the highest priority. However, this only allows PTP packets to jump ahead of other packets waiting to be sent and cannot mitigate for the effect of packets already in progress. Therefore, competing traffic will always introduce PDV through switches or routers that do not have a boundary clock function.

Unfortunately, it is very hard to quantify the effect of PDV since different PTP slave designs will have different techniques for handling it – some better than others – and the performance will normally be dependent not only on the absolute value of the PDV (the difference in transit time between the fasted and slowest packets), but also the distribution of the individual delays within that range. However, it is always the case that the fewer the number of switches or routers between the edge master and the slaves, the better the performance. As a general rule, to meet the category C sync requirement there should be at most five switches or routers between the edge master and each base station.

The network technology and asymmetry

A fundamental assumption of PTP timing is that the path between the master and slave must be symmetrical. In other words, the average time taken for a timing packet to travel from master to slave, ignoring the effects of PDV, should be the same as the time taken from slave to master. If this isn't the case then the time recovered by the slave will have a permanent offset of exactly one half the asymmetry. For example, a 500 ns asymmetry would result in a static error of 250 ns. Although this error is within the category C limit, it erodes the margin for additional error resulting from PDV.

A big source of asymmetry arises from long fiber runs, either because of physical length differences in separate RX and TX fibers, or effective optical path length differences when both directions operate over a single fiber using different wavelengths. For example, a difference in length of 200 m, which could easily occur in a 100+ km fiber run, would result in an asymmetry of about 1.3 μs and a corresponding time error of 650 ns – nearly half the category C budget. For this reason, it makes sense to restrict a cluster of base stations served by an edge master to a geographically local network, for example within a single building, rather than across multiple sites separated by large distances.

Another consideration arises when a portion of the network uses a technology other than Ethernet. A typical example is the use of a PON technology, such as BPON or GPON, or an ATM technology such as VDSL for providing a last mile network connection. The nature of such technologies often makes them totally unsuitable for transporting PTP timing since they introduce extreme asymmetry or PDV (through the use of timeslot-based mechanisms for upstream data). Therefore, unless the provider of that network segment has taken special measures to support PTP, the edge master should always be located downstream of such a link. For example, in a building where a number of small cells are connected internally by Ethernet, but then connected to the outside world by GPON, the edge master should be located within the building itself on the Ethernet network.

Ease of location

In the case where the edge master will use GNSS, it most likely will require connection by coaxial cable to a roof-mounted, or at least exterior, antenna. Often the effort and cost associated with this will influence the optimal location of an edge master, particularly for in-building applications.

Number of slave devices per cluster

Each slave connected to a master will require a finite amount of processing resource within the master and therefore there is a limit on the number of base stations that a single edge master can serve. While this may be an important consideration in some instances, Orolia's EdgeSync is available with support for up to 128 separate slaves, which is likely to be more than enough for most applications.

Availability considerations

Consideration should be given to what overall effect will happen if an edge master is unable to provide a normal level of synchronization for some reason. One possible scenario is if the edge master itself fails, in which case the attached base stations will not have any source of synchronization and will rely on their internal oscillators to maintain timing, which they may only be able to do for a matter of minutes. For this reason, it may make sense to provide some level of redundancy. For example, two edge masters could be deployed within a building so that if one fails the base stations can automatically switch over to using the other.

Another availability consideration is what happens if the overall synchronization source fails. For example, in the case of an edge master using GNSS, what happens if GNSS is jammed by an external source. In this case, the edge master will enter a holdover mode in which it will use its internal oscillator to maintain synchronization. Orolia's EdgeSync is available with options to maintain category C synchronization limits for either four or eight hours when in holdover. This should be sufficient to ride over most synchronization outages.

Summary

The deployment of edge masters, such as Orolia's EdgeSync, as part of a G.8275.2 PTP network can be used to provide synchronization to clusters of base stations or baseband units, such as small cells in a high-rise building, for LTE-TDD and 5G type applications. However, consideration needs to be given to the topology, size and locations of these clusters to ensure the synchronization requirements are met.