

How Do Synchronous Ethernet And Precision Time Protocol Work?

Part 1: Frequency Distribution and Synchronous Ethernet

by Per Lembre, Director of Product Marketing, Xelerated

With new mobile devices and applications being introduced every day, service providers are scrambling to keep up with demand: but how? In deciding how to best prepare mobile networks to accommodate growing bandwidth and application use, service providers and network equipment vendors alike must take a number of factors into consideration: one important one being synchronization.

To illustrate the importance of synchronization in Ethernet design, this TechNote will review two different methods. First, we will examine synchronization at the physical layer (Synchronous Ethernet), followed by a review of the packet layer synchronization method (PTP).

Mobile Backhaul: The Need For Speed

Mobile Broadband is quickly becoming a reality across the globe. 3G dongles, the Apple iPhone and other smart phone devices have turned mobile broadband into an everyday experience. Now that consumers are accustomed to mobile access – e-mail, Twitter conversations, television services and Internet video on the move – they are hungry for more.

To accommodate high growth rates of mobile data, mobile operators are looking into upgrading backhaul networks to Carrier Ethernet based on new standards. The Long Term Evolution (LTE) technology framework, as the 3GPP version 8 standard is often referred to, offers more bandwidth for the individual user and more efficient use of capacity across every part of the network. In essence, LTE is about bringing heavyweight broadband to mobile users and fulfilling this “need for speed.” For service providers the goal is to seamlessly evolve the network, step-by-step, while prolonging the use of legacy networking equipment supporting mobile data services.

When upgrading Carrier Ethernet access systems designed for mobile backhaul, providers must also plan well to preserve synchronization during the transition. Why is synchronization important? Synchronization isn't just a design relic of transport technologies from the infancy of digitalized voice; rather, there are growing application demands for synchronization. For example, in cellular systems wireless frequencies need to be tightly managed for seamless handover and connection integrity. Basestations – whether legacy 2G, 3G, or future LTE/4G and WiMAX – will continue to rely on precision time to operate. To illustrate, Fig. 1, overleaf, compares 2G, 3G and LTE architectures in this context.

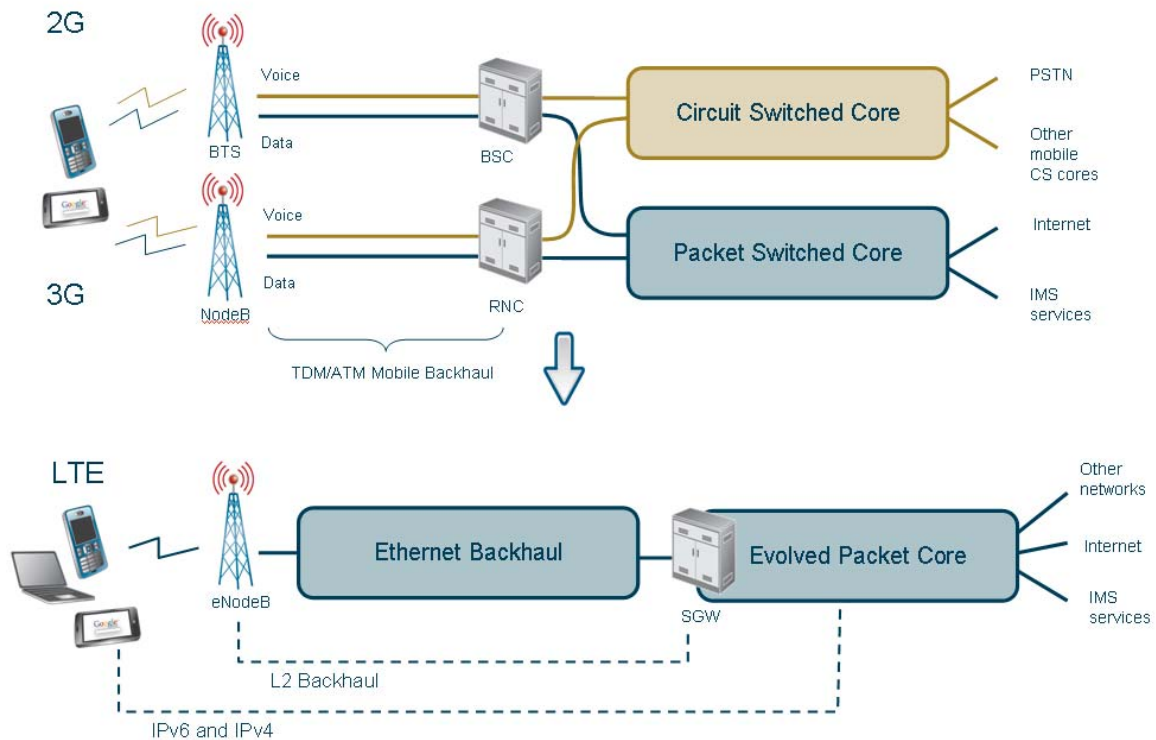


Fig. 1: Architecture Comparisons: 2G, 3G, and LTE

LTE differs from 2G and 3G architectures as it uses IP in all parts of the network, from the mobile terminal all the way through the basestation, backhaul, and the evolved packet core. (Note, Fig. 1 shows the 3GPP versions of 2G and 3G systems.)

All-IP Long-Term Evolution (LTE) Pushes All-Ethernet Backhaul

The LTE standard, finalized in early 2009, defines major architectural changes to meet the goals of LTE (listed overleaf in Table 1). Important to note is the focus on an all-IP approach. Previous mobile standards used a mix of circuit-switched, asynchronous transfer mode (ATM) and IP technologies to enable support of legacy services. In the LTE standard, both the data and control planes are based on IP protocols: with no exceptions. This allows for a flatter network approach and a simplified operational model.

Goal	Main Drivers
Higher download bit rates, 100 Mbit/s (for 20 MHz spectrum)	True end-user broadband experience
Higher upload bit rates, 50 Mbit/s (for 20 MHz spectrum)	True end-user broadband experience
Reduce packet latency to <5 ms in RAN	VoIP to cell phones requires low latency
Reduce control plane latency to <100 ms	Increase user experience by reducing transition time from idle to active
All-IP control plane	Simplified operations and decreased opex
Increased control plane capacity >200 users per cell (for 5 MHz spectrum)	Increased capacity to support more users
Flat architecture	Scale bandwidth to cope with increased load
Lower power consumption	Long battery time to meet always online behavior; Green networking
More efficient use of spectrum	New and enhanced use of radio spectrum allows for great coverage and efficient utilization of scarce resources
Support of other open technologies	Interworking with CDMA2000, EV-DO, GPRS, UMTS and WiMAX networks; Smooth migration to LTE from various starting points

Table 1: LTE Main Goals

(Source: 3GPP TR 25.913 v8, December 2008, Xelerated)

To prepare for LTE and accommodate 3G build-outs, service providers are re-engineering the mobile backhaul infrastructure using Carrier Ethernet technologies.

Synchronization Methods

There are two basic methods for distribution of precision time and frequency over a packet infrastructure: by the physical layer (Synchronous Ethernet) or by exchanging time messages through a packet protocol (PTP). Table 2 compares and contrasts the two methods at a glance.

Function	PTP IEEE 1588	Synchronous Ethernet ITU G.8261
Architecture	Packet-based transport, independent of physical layer	Uses physical layer of Ethernet (ETY, IEEE 802.3)
Forwarding principle	Ethernet, UDP over IPv4/IPv6	Ethernet frames used for synchronization status messages only
Distribution of frequency	Yes	Yes
Distribution of time-of-day	Yes	No
Dependency of packet processing hardware	Uses packet processing hardware to measure link and resident time. Impairments on packet layer affect synchronization performance.	Independent of delays and delay variations at the packet processing layers.

Table 2: Precision Time Protocol (PTP) And Synchronous Ethernet Compared

Distribution of Frequency

Network service performance based on Synchronous Ethernet is defined in ITU G.8261, which translates to component requirements defined in G.8262. Performance targets for PTP can be defined in application specific profiles. Both Synchronous Ethernet and PTP may be used in combination to achieve a high level of frequency synchronization with a common defined time. Realistic studies in large scale deployment scenarios indicate an accuracy of 50 ns.

Synchronous Ethernet distributes precise frequency via tight control of Ethernet link bit rates and is dependent on an unbroken chain of nodes with additional hardware support to function. In theory this is not the case for PTP; however, traversing non-supporting PTP nodes significantly impairs the synchronization quality.

Distribution of time-of-day

Synchronous Ethernet delivers a high level of frequency accuracy, but can't deliver time-of-day information, measured by GMT (Zulu) as provided by GPS signals. Conversely,

PTP supports time-of-day information, used by billing and service level agreement reporting functions in mobile networks.

Dependency of access link types

Both PTP and Synchronous Ethernet can be supported in point-to-point Ethernet aggregation switches. E-PON systems, on the other hand, cannot use the PTP protocol. Link asymmetry impairs the timing accuracy while other means of distributing precise clock signals transparently across E-PON environments are being developed.

We will review each method in detail, starting with the physical layer.

The Physical Layer

Synchronous Ethernet, as defined in ITU G.8261, uses the physical layer to synchronize all links to the same frequency. Because Carrier Ethernet deployments are commonly high-speed point-to-point connections, the physical layer (ETX) works well in asynchronous mode as packets are buffered in each node.

In Synchronous Ethernet, the link frequency is linked to a traceable primary reference source. All participating nodes – see Fig. 2 – (as in an unbroken chain) recover the clock at the PHY of the incoming interface and lock it to the phase-locked loop (PLL) oscillator on the board. This is similar to SDH/Sonet and allows for highly accurate frequency synchronization, known as synthonization. This method, however, doesn't allow for distribution of time of day-information, which limits reporting functions in mobile networks. One additional benefit is that the oscillator in switching nodes can be moderately advanced and are therefore inexpensive.

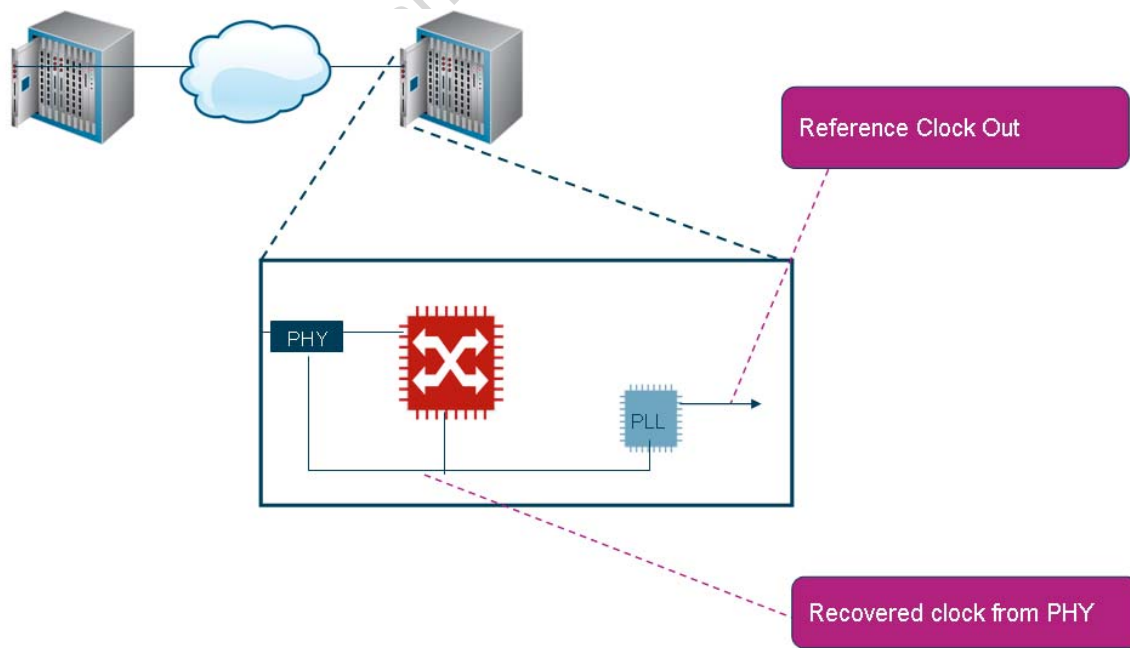


Fig. 2: Frequency Synchronization In Synchronous Ethernet

Timing Requirements

In reviewing Synchronous Ethernet functionality, it is critical to study the range of timing requirements across various networks. As described in Table 3 mobile networks have a variety of tight timing requirements on precision of reference frequency, phase and time.

WIRELESS TECHNOLOGY	INTERFACE	MAX FREQUENCY OFFSET	MAX PHASE OFFSET	TIME OF DAY MAX OFFSET
GSM	Radio	50 - 250 ppb (parts per billion)		
	Transport	0.01 ppb	4.3 μ s	
CDMA/CDMA2000	Radio	50 - 250 ppb		1 μ s
WCDMA/UMTS	Radio	50 - 250 ppb		2.5 μ s (TDD)
	Transport	15 ppb		
	RNC, net sync	0.01 ppb	4.3 μ s	
TD-SCDMA	Radio	50 - 250 ppb		
WiMAX Fixed	Radio	15 ppm		
	Transport	0.01 ppb	4.3 μ s	1 μ s (TDD)
WiMAX Mobile	Radio	15 ppm		1 μ s (TDD)
	Transport	16 ppb		
LTE	Radio	50 - 250 ppb	10 - 50 μ s (TDD)	

Table 3: Wireless Networks Timing Requirements
(Source: IETF internet draft TICTOC requirements, G.8261)

Synchronization requirements in wireless differ between technologies. Note: transport requirements may put higher demands than the air interface.

Mobile Radio Access Network (RAN) technologies require precise synchronization of all basestations to operate. Used to avoid interference and crosstalk between channels in the air interface, synchronization ensures seamless handover.

Fig. 3 outlines the architecture supporting basestation frequency synchronization using Synchronous Ethernet. The Primary Reference Source is fed to a synchronous network, based on Sonet/SDH or Synchronous Ethernet. The synchronization path follows a chain of nodes, through the core to mobile backhaul network before eventually reaching the basestations.

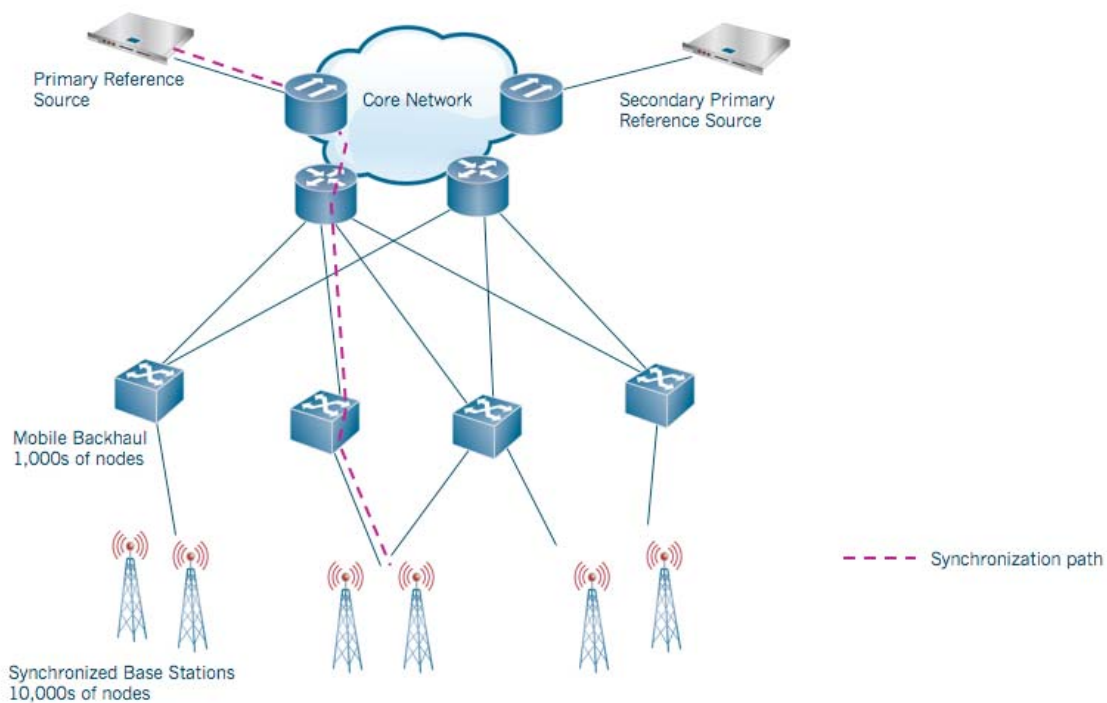


Fig. 3: Basestation Frequency Synchronization

Frequency Distribution

In today's mobile backhaul networks, precision time may be distributed by GPS or through plesiochronous digital hierarchy (PDH) connections. However, when moving to a packet infrastructure, the PDH connections are no longer available for free. In addition, deployment of GPS receivers at every cell site has a number of coverage, security and maintenance challenges.

Synchronous Ethernet requires hardware support, as the packet processor should be designed to flexibly select an incoming interface for clock recovery and then use this to synchronize all output interfaces to the same frequency. The hardware must also tightly monitor the frequency and phase of the data rate on incoming links.

Evaluation and Alternatives

As mobile bandwidth demands continue to rise, service providers must also respond to the *need for speed*. In preparing for LTE, service providers and network equipment vendors need to consider synchronization implications when upgrading legacy equipment or deploying new solutions.

There are two methods of synchronization: through the physical layer (Synchronous Ethernet) and through packet protocol (PTP). While the Synchronous Ethernet method avoids the packet processing delays in PTP, it cannot distribute time-of-day distribution, which limits billing and service level agreement reporting functions in mobile networks.

Because of these clock distribution challenges in the Synchronous Ethernet method, packet-based synchronization methods have emerged as a viable alternative. Part 2 of this TechNote will review this alternative – Precision Time Protocol – in detail, emphasizing the measurement of link delay and distribution of time information.



As Published on EN-Genius.NET