Timing gets smart

Hyperscale computing and LTE-Advanced are driving demand for high-performance timing solutions. By James Wilson.

Ethernet has come a long way since IEEE 802.3 was first published in 1980. First envisioned as a technology to connect PCs and workstations, it has gradually evolved to become the networking technology of choice for a broad range of applications across enterprise computing, data centre, wireless networks, telecommunications and industrial sectors.

Its ubiquity and the ever-decreasing cost of the hardware needed to support it, means Ethernet will continue to gain in popularity in these applications.

Some of the most interesting technology transformations are underway as 100G Ethernet is being adopted in data centres and wireless radio access networks. These migrations to high-speed optical Ethernet are driving the need for higher performance clock and frequency control products.
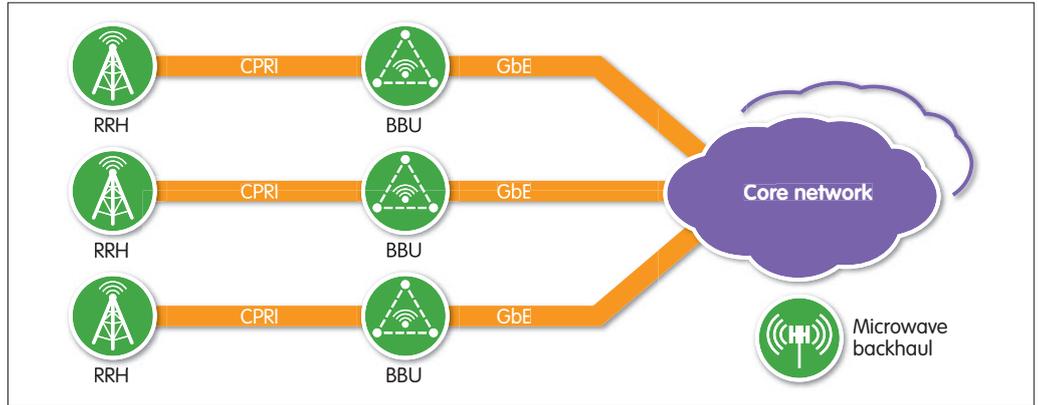
Traditional enterprise workloads are migrating quickly to public cloud infrastructures, driving an investment boom in data centres. In addition to increasing demands for lower latency, data centres share a challenge in that most data traffic stays within the data centre as workload processing is distributed across multiple compute nodes.

Modern data centres are optimising their network architecture to support distributed, virtualised computing by connecting every switch to each other, a trend known as hyperscale computing. One of the underlying technologies that makes hyperscale computing attractive commercially is high-speed Ethernet and data centre switches are migrating to 25G, 50G and 100G Ethernet to accelerate data transfer and network efficiency.

The migration from 10G to 25/50/100G Ethernet is driving data centre equipment manufacturers to upgrade switch and access ports to higher speeds. This, in turn, fuels the need for higher performance, lower jitter timing solutions. Ultra-low jitter clocks and oscillators are necessary in these applications because high clock noise can result in unacceptably high bit-error rates or lost traffic. Table 1 highlights the typical timing requirements for Ethernet PHYs, switches and switch fabrics. The safe and proven approach for high-speed Ethernet is to use an ultra-low jitter clock source that delivers excellent jitter margin to these specifications.

Wireless radio access networks
Wireless networks are poised to go through tremendous change as they migrate from 4G/LTE to LTE-Advanced and 5G. Next-generation wireless networks will be optimised for carrying
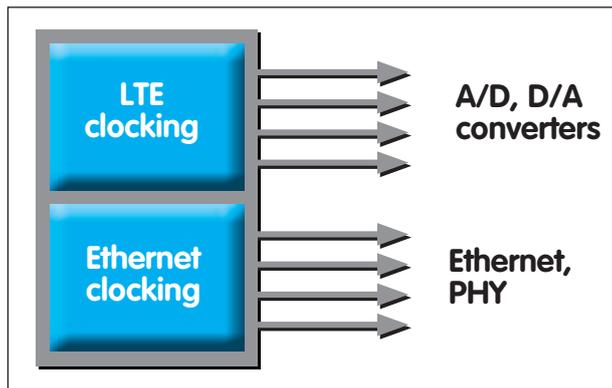
Figure 1: 4G/ LTE Radio Access Network (CPRI Fronthaul)

Figure 2: A single-chip clock IC generates LTE and Ethernet clocks in RRH, small cells and pico cells

mobile data which is expected to grow to 49Exabyte per month by 2021, a sevenfold increase over 2016. As a result, wireless networks are being re-architected and optimised for data transport. The wide-scale adoption of high-speed Ethernet in radio access networks (RAN) is expected to play a critical role.

In 4G/LTE radio access networks, the RF and baseband processing functions are split into separate remote radio heads (RRH) and base band units (BBU). Each RRH is connected to a BBU over a dedicated fibre connection based on the Common Public Radio Interface (CPRI) protocol (see fig 1). This architecture enables the replacement of dedicated copper and coax cable connections between the radio transceiver and the base station and enables the BBU to be placed in a more convenient location to simplify deployment and maintenance.

While more efficient than legacy 3G wireless networks, this network architecture is limited because bandwidth is constrained by the speed of the CPRI link – typically up to 10Gbit/s. In addition, the CPRI connection is a point-point link and

10GbE	25GbE/40GbE	100GbE
1ps RMS (12kHz to 20MHz)	0.3ps RMS (12kHz to 20MHz)	0.15ps RMS (12kHz to 20MHz)

RRH and BBUs are typically deployed in close proximity – from 2km to 20 km – constraining network flexibility.

As part of the evolution to 5G, the wireless industry is having to rethink base station architectures and the connection between baseband and radio elements, known as the front haul network, is a key area for optimisation. Higher bandwidth front haul networks are required to support new LTE features that support high-speed mobile data, including Carrier Aggregation and Massive MIMO. In addition, network densification and the adoption of small cells, pico cells and micro cells will put additional bandwidth requirements on fronthaul networks.

To minimise investment and operating costs, 5G will use a Cloud-RAN (C-RAN) architecture that centralises baseband processing (C-BBU) for multiple remote radio heads.

New standards are being developed for fronthaul to support the evolution to C-RAN. IEEE 1904 Access Networks Working Group (ANWG) is developing a new Radio over Ethernet (RoE) standard for supporting CPRI encapsulation over Ethernet. This new standard will make it possible to aggregate CPRI traffic from multiple RRHs and small cells over a single RoE link, improving fronthaul network utilisation. Another working group, IEEE 1914.1 Next Generation Fronthaul Interface (NGFI), is revisiting the Layer-1 partitioning between RF and baseband to support more Layer-1 processing at the RRH. NGFI enables the fronthaul interface to move from a point-to-point connection to a multipoint-to-multipoint topology, improving network flexibility and enabling better coordination between cell sites. A new CPRI standard for 5G front-haul (eCPRI) will be released

in August 2017 that details the new functional partitioning of base station functions and includes support for CPRI over Ethernet.

These new fronthaul standards create the need for frequency flexible timing solutions that can support both LTE and Ethernet clocking in radio heads, small cells, and pico cells (see fig 2). These solutions will provide the opportunity for hardware designs to unify all clock synthesis into a single, small-form factor IC.

Another key challenge is accurate timing and synchronisation. Historically 3G and LTE-FDD mobile networks relied on frequency synchronisation to synchronise all network elements to a very precise and accurate primary reference, typically sourced from a signal transmitted by GNSS satellite systems (GPS, BeiDou). These systems require frequency accuracy within 50 ppb at the radio interface and 16 ppb at the base station interface to the backhaul network. LTE-TDD and LTE-Advanced retain these frequency accuracy requirements but add very stringent phase synchronisation requirements (±1.5µs). This key requirement enables new features, such as

Table 1: Typical phase jitter requirements for high-speed Ethernet

Author profile:
James Wilson is senior marketing director, timing products, with Silicon Labs.

Figure 3: An LTE-Advanced Radio Access Network

enhanced inter-cell interference coordination (eCIC) and coordinated multipoint (CoMP), which maximise signal quality and spectral efficiency. These phase synchronisation requirements are expected to be further tightened in the upcoming 5G standards.

Figure 3 shows the LTE-Advanced network architecture, in which multiple remote radio heads connected to a centralised base band unit over packet-based eCPRI networks and phase/frequency synchronisation is provided by IEEE 1588v2/SyncE. Time and phase synchronisation is supported by implementing IEEE1588/SyncE at the remote radio head and centralized base band unit. Higher bandwidth 100 GbE networks are used to backhaul traffic from each base band unit to the core network. Higher performance, more flexible timing solutions are now available that simplify clock generation, distribution and synchronization in LTE-Advanced applications.

Ethernet is being broadly adopted in data centre and wireless networks to enable higher network utilisation and lower cost data transmission and to enable new service provider features and services. The transition to packet-based Ethernet networks in these infrastructure applications is driving the need for more flexible, lower jitter timing solutions based on innovative architectures that enables frequency flexibility and ultra-low jitter.