

Synchronous Ethernet: A Method to Transport Synchronization

Jean-Loup Ferrant, Alcatel-Lucent

Mike Gilson, British Telecom

Sebastien Jobert, France Telecom

Michael Mayer and Michel Ouellette, Nortel

Laurent Montini, Cisco

Silvana Rodrigues, Zarlink

Stefano Ruffini, Ericsson

ABSTRACT

This article discusses the evolving transport architecture, covering some of the synchronization distribution problems to many endpoints where mobile backhaul and TDM emulation occur. It shows how Synchronous Ethernet fits into both the Ethernet and synchronization architectures, and discusses how this helped development in standardization bodies. Standardization allows the key building blocks of Ethernet silicon and specific timing devices to be developed, which allows a robust system implementation to be constructed while allowing interworking with and migration from existing SONET/SDH-based transport infrastructure. Finally, results are shown that indicate a very high level of performance is achievable with Synchronous Ethernet not subject to the normal packet delay variation and traffic load conditions that can occur in packet based networks.

INTRODUCTION

A number of key changes are occurring in current telecom networks that will have wide and far reaching consequences for the transport environment. These changes also provide some key challenges, especially in the area of frequency synchronization.

Ethernet has become the universal low-cost mature technology of choice in both the enterprise and residential markets. It is only a natural step for this to be applied in the wide area network (WAN) environment. In many cases it is the access network that is the first part of the architecture impacted by this evolution. However, Ethernet was not designed for the transport of synchronization, which is a key requirement for certain existing applications, especially concerning time-division multiplexing (TDM) emulation and mobile backhaul. Frequency synchronization is also a key aspect of future

applications primarily covering next-generation mobile, and may also provide support to the distribution of very accurate time of day at the physical layer.

Synchronous Ethernet (SyncE) is a key development of the evolution of Ethernet into a carrier grade technology suitable for the WAN environment where frequency synchronization is required.

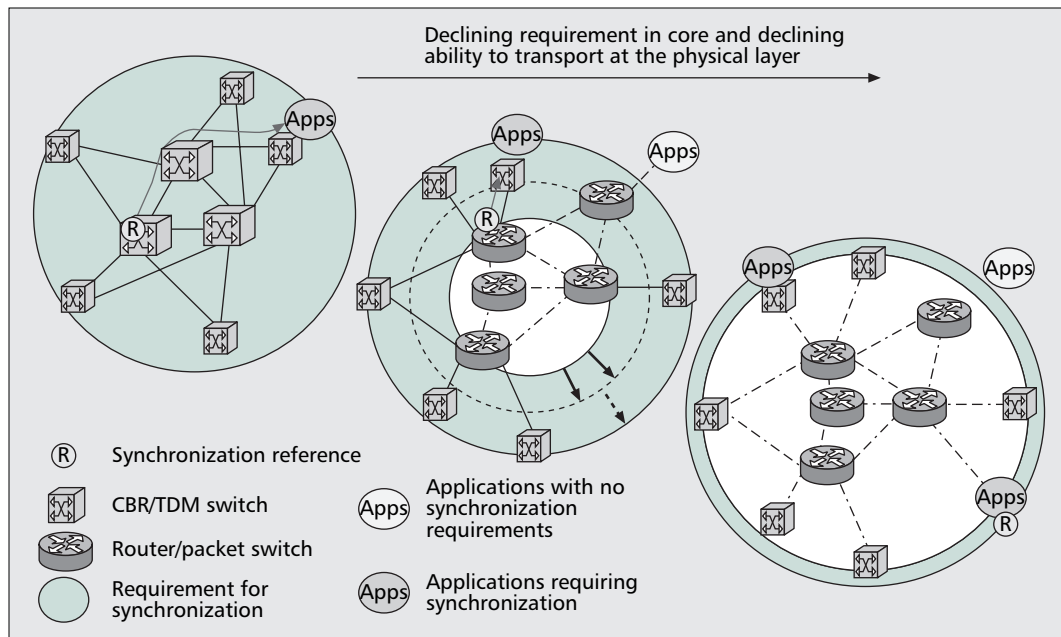
EVOLVING THE TRANSPORT ARCHITECTURE

KEY ISSUES

Migration from existing circuit-switched technology based on TDM networks to packet-switched technology based on Ethernet is seen as the future in the telecommunication networks. Communications providers (CPs) are under pressure to reduce costs but increase bandwidth for new service types. The need for flexibility to carry different types of services over the same network (convergence of services and networks) is also critical to create increasing efficiency. Synchronization and its evolution in these converged networks is a key issue that needs addressing.

SYNCHRONIZATION EVOLUTION

Switching from TDM networks to packet networks could at first be seen as an important change requiring careful study. TDM networks (e.g., synchronous optical network/digital hierarchy [SONET/SDH], plesiochronous digital hierarchy [PDH]) are technologies that natively have the capability to carry a timing reference at the physical layer. Packet technologies were initially designed to work in asynchronous mode where the oscillators in the equipment are free running. Although this allows the underlying infrastructure to operate, many applications exist that require frequency synchronization. For example, frequency synchronization and stability are



■ **Figure 1.** Frequency synchronization requirement in an evolving packet-based network.

Transport of synchronization has been traditionally performed for years with TDM networks by the line signals at the physical layer level, using well-known principles, engineering rules and experience. This approach can also be applied to Ethernet based packet networks using Synchronous Ethernet technology.

required at mobile base stations in order to make efficient use of the radio spectrum. In addition, TDM emulation may also require a synchronized and stable frequency to be available at the emulation points, which are at the edge of the network. Clearly the requirement for frequency synchronization is moving away from an infrastructure requirement of the core network toward a requirement of the edge application. Also, the ability to distribute synchronization from center to edge declines as infrastructure evolves toward a packet-based architecture. This can easily be represented as a diagram that increasingly looks like a doughnut (Fig. 1), where the hole describes the declining requirement and ability to transport frequency synchronization in the center.

Traditional TDM technologies and applications require frequency synchronization in all parts of the infrastructure (shaded areas in Fig. 1) in order to work. When the infrastructure is no longer synchronized, data is lost, the application experiences impairment, and the service is degraded. Maximum levels of phase instability, described in terms of jitter and wander, are governed by the International Telecommunication Union — Telecommunication Standardization Sector (ITU-T) G.81x series of standards. With the shift toward packet-based technologies a hole opens up in the center of this architecture, as it is increasingly packet-based where synchronization is no longer required in the core; synchronization is only required at the edge by the application. This creates a challenge for the delivery of frequency synchronization.

Transport of synchronization has traditionally been performed for years with TDM networks by the line signals at the physical layer level, using well-known principles, engineering rules, and experience. This approach can also be applied to Ethernet-based packet networks using Synchronous Ethernet technology.

SCALABILITY

Another important consideration for frequency synchronization, especially for large networks with many endpoints, is the advantage of distributing the synchronization from some centralized location (i.e., the primary reference clock, PRC) toward the edge of the network. To get an idea of the scale of the number of endpoints that may have to be served, see Fig. 2. There are a number of methods by which frequency synchronization can be provided. Using a key advantage from TDM, one of the methods that can guarantee the best quality in delivery of a stable frequency is via the physical layer as part of the physical connection delivering the data.


With the convergence of fixed and mobile services and the evolution from TDM to Ethernet, Synchronous Ethernet provides an evolutionary solution to deliver synchronization. Synchronous Ethernet is specified to run over the Ethernet media types defined in IEEE 802.3 [5]. An important consideration, especially in the access portion of the network, is the ability of Synchronous Ethernet to interwork with synchronization networks operating over the full range of infrastructure, covering optical fiber, copper, and microwave. Although Synchronous Ethernet is currently defined only for media contained within IEEE 802.3, the concept of frequency transfer may be provided over any physical media, provided it operates in a constant bit rate fashion. Fiber access and microwave are possible examples.

SYNCHRONOUS ETHERNET ARCHITECTURE AND STANDARDS

KEY ARCHITECTURE ELEMENTS

Synchronous Ethernet provides a mechanism to transfer frequency over the Ethernet physical layer, which can then be made traceable to an

Specific functions, which are responsible for any formatting changes, are defined at the boundaries between network layers. "Trail termination functions" add any OAM information necessary to manage the layer network.



■ Figure 2. Typical scale of synchronization delivery in a large communications provider network.

external source such as a network clock. As such, the Ethernet link may be used and considered part of the synchronization network. Synchronous Ethernet must "fit" within the general architecture of an Ethernet network. To make use of its ability to transfer timing, Synchronous Ethernet must also fit within the general architecture of synchronization networks. The Ethernet and synchronization network architectures are based on the functional modeling approach described in ITU-T Recommendations G.805 and G.809.

Ethernet Architecture and Synchronization

— Architectures have been developed within the ITU-T for various technologies such as Ethernet, multiprotocol label switching (MPLS), asynchronous transfer mode (ATM), SONET/SDH, and optical transport network (OTN). Architectures describe the network in terms of its information transport capability and what is needed in terms of operations, administration, and maintenance (OAM) to manage a multidomain network.

ITU-T defines the network architecture in terms of layers. While these are somewhat similar to the way existing IEEE specifications are described, based on the open systems interconnection (OSI) reference model, the ITU-T methodology has some subtle differences. The ITU-T definition of a layer includes the concept of a monitored trail, which is necessary to achieve manageability in a multidomain network. Additionally, the ITU-T network layer models are not restricted to seven layers as in the OSI model. New layers are created where necessary. Each layer, however, must be fully defined, again in order to ensure manageability across

multidomain and possibly multitechnology networks. There are additional differences between the two approaches, but they do not need to be covered in order to understand aspects of the model related to synchronization. However, adherence to the strict conventions used to define architectures is mandatory to result in networks that are deployable in existing and new networks.

A layer network is defined in terms of the information to be transported across the network (i.e., the characteristic information, CI). Specific functions (i.e., adaptation and trail termination functions) responsible for any formatting changes are defined at the boundaries between network layers. Trail termination functions add any OAM information necessary to manage the layer network.

G.8010 Ethernet Architecture — Ethernet has been described in ITU-T Recommendation G.8010 as a two-layer network, the ETH and ETY layers. Simply, the ETY layer is the physical layer as defined in IEEE 802.3, while the ETH layer represents the pure "packet" layer. Ethernet medium access control (MAC) frames at the ETH layer are carried as a client of the ETY layer. Various protocols and functionality defined within the IEEE standards are mapped to specific functions within the layer network. In OSI terminology, ETY is layer 1, ETH layer 2.

Synchronization-related work within the ITU-T has been ongoing in parallel with the development of the Ethernet architecture. Experts within ITU-T Study Group 15 responsible for studying timing and synchronization have been investigating the need to provide explicit additions to general transport architectures in order

The proposal to specify the transport of a reference clock over Ethernet links was brought by operators to ITU-T Study Group 15 in September 2004. Analysis indicated that such a feature is consistent with the IEEE 802.3 standard, which specifies Ethernet clocks to be within ± 100 ppm.

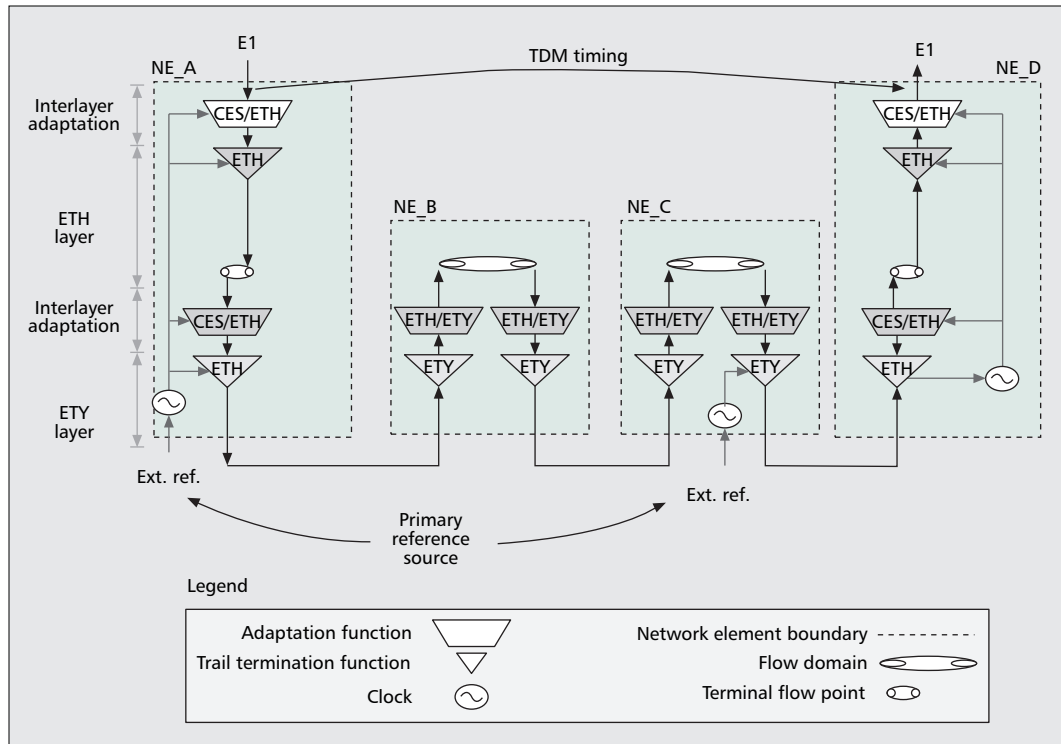


Figure 3. CES and synchronous Ethernet shown together with the ETY and ETH functions defined in G.8010.

to support synchronization. The initial work was related to circuit emulation services (CES) where a layer 1 service (e.g., T1/E1) is carried over a layer 2 network (e.g., Ethernet). The CES case is an important case architecturally, since timing is a distinct part of layer 1 service; however, it must be achievable over a packet network that does not inherently transfer timing. One of the key network issues for certain types of CES mappings is how a network timing reference is made available to the mapping points within a network.

Both CES and Synchronous Ethernet are additions to the existing ITU-T models. The intent of both is to not require changes to the existing IEEE defined protocols, but to extend and work within these protocol definitions to provide the functions on a network scale. Architecturally, CES is viewed as a higher layer, or client, to the Ethernet network, and as such is defined in the necessary adaptation functions. Synchronous Ethernet, on the other hand, is not defined as a function, but simply describes how a frequency reference is provided to currently defined functions (e.g., a trail termination function that implements one of the IEEE802.3 PHYs). Figure 3 shows the G.8010 architecture representing both the CES aspects and Synchronous Ethernet relative to the ETH and ETY layers defined in G.8010. This figure shows the two Ethernet architectural layers, ETH and ETY, together with higher-level adaptation functions representing mapping and demapping of an E1 circuit using circuit emulation with differential clock recovery. In this example the various functions are grouped into four network elements (NEs), as indicated with dotted boxes. Network timing is provided to NE_A and NE_C

via external interfaces (Ext.Ref.) on those network elements. For simplicity, the synchronization network that carries the timing between the primary reference source and the NEs is not shown. When it may be impractical to provide an external clock at NE_D (e.g., due to cost or where access is impractical), timing to NE_D is provided via Synchronous Ethernet. The link between NE_C and NE_D shows an example of Synchronous Ethernet. In this case the external clock on the third NE (NE_C) provides timing to the Ethernet physical layer (ETY). The fourth NE (NE_D) can then recover this clock from the physical layer (ETY). In this example the clock signal is used by the NE as a reference to reconstruct the E1 timing carried by circuit emulation.

While the figure is primarily intended to show frequency transfer, the model also allows greater understanding of other related functionality and provides a starting point to investigate the interactions of frequency synchronization (e.g., Synchronous Ethernet) with other forms of frequency and time transfer. The modeling also provides a common understanding between various groups within the ITU-T.

DEVELOPMENT IN STANDARDS BODIES

The proposal to specify the transport of a reference clock over Ethernet links was brought by operators to ITU-T Study Group 15 in September 2004. Analysis indicated that such a feature is consistent with the IEEE 802.3 standard, which specifies Ethernet clocks to be within ± 100 ppm (parts per million). Synchronous Ethernet clocks are within ± 4.6 ppm, which is within the frequency range of Ethernet interfaces. In addition, by externally timing the Ethernet clock, PRC traceability of the interface is achievable.

Octet number	Size	Field
1–6	6 octets	Destination address = 01-80-C2-00-00-02 (hex)
7–12	6 octets	Source address
13–14	2 octets	Slow protocol Ethertype = 88-09 (hex)
15	1 octet	Slow protocol subtype = 0A (hex)
16–18	3 octets	ITU-OUI = 00-19-A7 (hex)
19–20	2 octets	ITU Subtype
21	4 bits	Version
	1 bit	Event flag
	3 bits	Reserved
22–24	3 octets	Reserved
25–28	4 octets	QL TLV (type, length, reserve, SSM)
29–1532	32-1486 octets	Future enhancement TLVs and padding
Last 4	4 octets	FCS

Table 1. *Ethernet Synchronization Messaging Channel (ESMC) protocol data unit.*

Another key input that was also brought by operators was the need for interworking between SONET/SDH and Synchronous Ethernet equipment in order for them to have a single synchronization network to manage. This proposal was in line with the current extensions of SONET/SDH equipment with Ethernet interfaces. This resulted in three important decisions:

- To extend the scope of the G.803 [7] reference synchronization chain to Synchronous Ethernet equipment. This has been done in G.8261 [1], which defines the architectural aspects of Synchronous Ethernet.
- To specify Synchronous Ethernet clocks compatible with SONET/SDH clocks as defined in G.813 [8] and G.812 [9]. This has been done in G.8262 [2], which specifies the clocks for Synchronous Ethernet equipment.
- To use the synchronization status message (SSM), as defined in G.707 [10]. Synchronization status messages contain an indication of the quality level of the clock that is driving the synchronization chain, and is used to control, maintain, and restore the synchronization chains of Synchronous Ethernet equipment. The IEEE proposed a solution based on Slow Protocol to ITU-T. ITU-T further developed the protocol, which is included in Recommendation G.8264 [3].

The details of the synchronization functions that needed to be implemented in Synchronous Ethernet equipment have been added in a new

release of G.781 [4], which originally specified these functions only for SDH equipment.

ITU-T Recommendation G.8261 [1], released in 2006, was the first document to introduce the Synchronous Ethernet concept as part of the studies related to the synchronization aspects in packet networks. This technology has been standardized by the ITU-T as a direct result of the experience gained with the standardization of timing distribution over SONET/SDH networks.

Since the very beginning, G.8261 was recognized as the main reference for this technology, although the first version of the standard presented several aspects not fully defined. Due to that, the initial concepts included in the 2006 version of G.8261 have been expanded and complemented with a revised G.8261 (recently published). This revised version provides more detailed requirements and architectural guidelines in a similar manner as was done for the standardization of synchronization networks based on SONET/SDH. The standardization of Synchronous Ethernet was finally completed during the same period by two other fundamental Recommendations, G.8262 and G.8264.

ITU-T Recommendation G.8262 [2] defines requirements for clock accuracy, noise transfer, holdover performance, noise tolerance, and noise generation. It defines two options for clocks for Synchronous Ethernet; these clocks are called Synchronous Ethernet equipment slave clocks (EECs). The first option, referred to as EEC-Option 1, applies to Synchronous Ethernet equipment designed to interwork with networks optimized for the 2048 kb/s hierarchy. Requirements for this option are based on those found in G.813 Option 1 used in SDH networks. The second option, referred to as EEC-Option 2, applies to Synchronous Ethernet equipment designed to interwork with networks optimized for the 1544 kb/s hierarchy. Requirements for this option are compatible with the requirements for Stratum 3 clocks deployed in SONET network elements. These requirements are based on a combination of requirements from G.813 Option 2 and G.812 Type IV.

To allow a Synchronous Ethernet link to convey the SSM quality level (QL) as defined in G.707 and G.781, a specific channel has been defined based on IEEE 802.3, Organization Specific Slow Protocol (OSSP), currently specified in IEEE 802.3ay [6], a revision to IEEE 802.3-2005 PAR. The Ethernet Synchronization Messaging Channel (ESMC, Table 1) protocol is composed of the standard Ethernet header for a slow protocol, an ITU-T specific header, a flag field, and a type length value (TLV) structure. The use of flags and TLVs aims to optionally improve the management of the Synchronous Ethernet link and associated timing chain. ITU-T Recommendation G.8264 [3] currently defines two messages, Event and Information, both supporting the same mandatory QL TLV for SSM transmission. The two message types are necessary to meet the strict delay requirements in ITU-T Recommendation G.781 [4], while still meeting the message rate requirements placed on slow protocols. The protocol allows for future enhancements through the definition of new TLVs as appropriate.

ENGINEERING THE SOLUTION

SILICON AND TIMING DEVICE COMPONENTS

Several semiconductor companies are in the process of implementing Synchronous Ethernet functionalities. This includes both Ethernet transceiver manufacturers as well as timing device manufacturers.

As an example today, a typical Gigabit Ethernet Quad-PHY transceiver operating at 100BASE-TX, 100BASE-FX, and 1000BASE-X rates already uses a free-running input reference clock (TX clock), and has a clock data recovery (CDR) block and in most cases a recovered output clock (RX clock) which is extracted from the line. These already available and important functions form the basis of Synchronous Ethernet. Today, the free-running input reference clock and extracted recovered output clock from transceivers are essential functions for proper Ethernet data transmission and recovery. For 100BASE-TX transceivers, the input free-running reference clock as specified by IEEE 802.3-2005 has an accuracy of $25 \text{ MHz} \pm 0.01$ percent (or 100 parts-per-million) as well as the recovered output clock frequency tolerance. For Gigabit Ethernet the accuracy is $125 \text{ MHz} \pm 0.01$ percent.

The primary idea behind Synchronous Ethernet is that the input reference clock is not free-running with an accuracy of ± 100 ppm, but rather locked and traceable to a primary reference clock as defined in ITU G.811 (achieving long-term accuracy of ± 10 parts-per-trillion). If the input to a transceiver is very accurate and stable, the recovered clock of a transceiver should, in theory, be locked and traceable to that input. In addition, Synchronous Ethernet transceivers will offer the possibility to select from which port to extract the output clock (e.g., with multi-PHY transceivers) as well as, for example, a fast link failure mechanism and squelching functions to prevent degraded output references being used as a clock reference. These are important features for proper network synchronization distribution.

In addition to the transceiver manufacturers, several timing device manufacturers are also introducing specialized Synchronous Ethernet components that will be deployable within line cards and centralized timing cards. These components are introduced to provide functions the Ethernet transceivers themselves might or might not provide beyond those listed above due to time to market. They are introduced so that system vendors can make use of current transceivers on existing line cards, but offload some of the synchronization functions to those specialized components.

These specialized components use, for example, integrated digital phase locked loop (DPLL) that includes different functions depending on whether the components are deployed in a line card or a central timing card. Functions that are important for line cards include frequency conversion, configurable PLL bandwidth for noise filtering, and jitter attenuation. Functions that are important for central timing cards include clock accuracy, noise transfer, holdover perfor-

mance, noise tolerance, and noise generation, which are specified in G.8262 [2]. The functions found on central timing cards are primarily the same as those currently used within SONET/SDH.

NETWORK, SYSTEM AND EQUIPMENT IMPLEMENTATION

This section shows how some of the functions listed earlier can be applied to implement Synchronous Ethernet on existing or new switches and routers. For simplicity, Fig. 4 presents a diagram of two network elements connected via an Ethernet interface (in practice there will be a cascade of switches forming a synchronization network and timing chain). From a synchronization perspective we only show the distribution of synchronization from left to right, where the Master NE distributes synchronization and the Slave NE recovers synchronization. Two types of line cards are shown: Synchronous Ethernet capable line cards and conventional Ethernet line cards. It can be seen from the figure that the Synchronous Ethernet line cards can interface to a timing backplane to source and terminate timing. For simplicity, the data paths are not shown in the figure. Both card types are interchangeable from the perspective of the data transport capabilities.

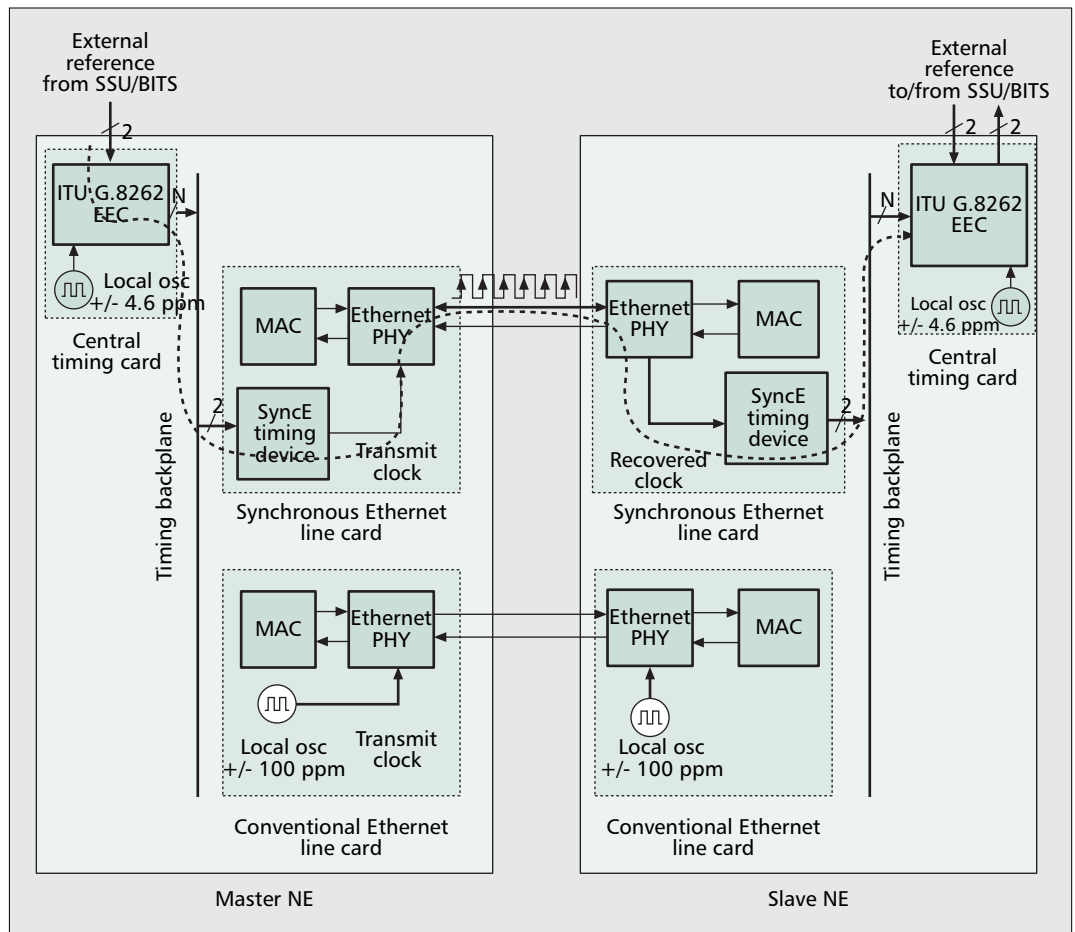
The Master NE takes external input timing references coming from the network clock (SSU or BITS). These interfaces are specified in ITU-T G.703 (e.g., 2048 kHz synchronization interface). In the future an external input timing reference might be Synchronous Ethernet based. These references are then used as input to the ITU G.8262 EEC clock, typically located on the central timing card of the NE. The figure also shows a free-running input crystal operating at ± 4.6 ppm, as specified in ITU G.8262 (same requirement as in G.813 option 1 and G.812 clock type IV). The EEC, for instance, specifies the level of jitter and wander at the input and output of the clock, as well as the specification under short-term and long-term transients such as loss of timing reference or link failure. It is also responsible for functions listed earlier such as timing reference selection and switching. The EEC output timing reference is then distributed via the NE backplane to reach the Ethernet line cards.

The Synchronous Ethernet line card includes the Ethernet MAC and PHY transceiver, as well as any timing device. The timing device provides functions listed earlier. Frequency conversion adapts backplane frequencies to frequencies that are required as input reference clocks to the transceiver (25 or 125 MHz). The output of the timing device serves as the TX clock reference into the transceiver, replacing the free-running crystals with ± 100 ppm accuracy in the conventional line card. The reference is then used to synchronize the physical line coding (e.g., 4B/5B for FE, 8B/10B for GE) of the interface toward the Slave NE.

At the Slave NE, the clock is recovered within the transceiver CDR in the Ethernet PHY block of the Synchronous Ethernet line card (current operation of transceivers). In some cases where the RX clock is not available at the

In addition to the transceiver manufacturers, several timing device manufacturers are also introducing specialized Synchronous Ethernet components that will be deployable within line cards and centralized timing cards.

It is important to note that Synchronous Ethernet can be applied to both existing networks and new networks. Synchronous Ethernet deployment is not required ubiquitously over an entire network, but only in cases where frequency distribution is needed.



■ Figure 4. Synchronous Ethernet network element and clock distribution.

transceiver, the use of an external CDR might be required to recover the clock. The recovered output clock is then fed into a timing device for functions (e.g., those found in the Master NE) such as jitter attenuation and transceiver frequency conversion to backplane frequency. The clock is then sent through the backplane to reach the Slave's central timing card. This timing reference then becomes a reference to the EEC (also known as a line-timing reference). As shown in the Slave NE, an EEC can accept line and external references, as well as the input of a ± 4.6 ppm local oscillator (used in situations where there are no line or external references available). The EEC can also provide timing outputs in the form of 2048 kHz synchronization interface to external SSU/BITS network elements (typically Stratum 2 NE).

From this point on, the Slave NE then becomes the Master NE for the next downstream NE, and synchronization is transported on a node-to-node basis, where each node participates in recovery and distribution. It is by this process that synchronization traceable to a primary reference source can be recovered and distributed within each NE. It is also important to note that Synchronous Ethernet does not disrupt the IEEE 802.3 architecture.

In a real-world implementation, however, synchronization distribution would be bidirectional, and provide redundancy through multiple central timing cards and line cards. The EEC as

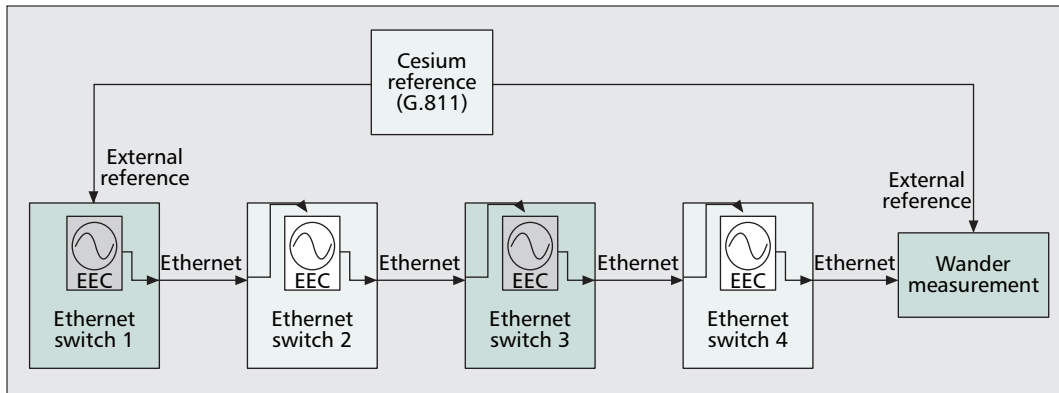
defined in ITU G.8262 was specified to interwork with SONET/SDH, and by doing so, Synchronous Ethernet can inherit most of the SONET/SDH synchronization design principles.

It is important to note that Synchronous Ethernet can be applied to both existing networks and new networks (sometimes referred to as green-field networks). Synchronous Ethernet deployment is not required ubiquitously over an entire network, only where frequency distribution is needed. In the green-field case it is expected that new Synchronous Ethernet equipment will be the preferred approach; however, in existing networks, as an example, it may be possible to simply replace asynchronous line cards with new synchronous line cards to provide this capability. But this is dependent on the architecture of the existing equipment.

In addition to the distribution of physical signals explained above, a simple communication protocol is needed between NEs. This protocol is specified in ITU G.8264 as the ESMC and is described above. ESMC is used to transmit, from NE to NE, the clock quality level value. The ESMC will typically be implemented outside the MAC and PHY, similar to other Ethernet OAM protocols.

NETWORK PLANNING CONSIDERATIONS

The deployment of synchronization networks utilizing layer 1 Synchronous Ethernet provides a useful tool to network operators to distribute frequency that is not subject to packet delay varia-



■ **Figure 5.** *Synchronous Ethernet testbed.*

tion. Due to the characteristics of the clocks, in general, existing network planning and engineering rules can be employed, greatly simplifying the integration of Synchronous Ethernet technology into the SONET/SDH synchronization network.

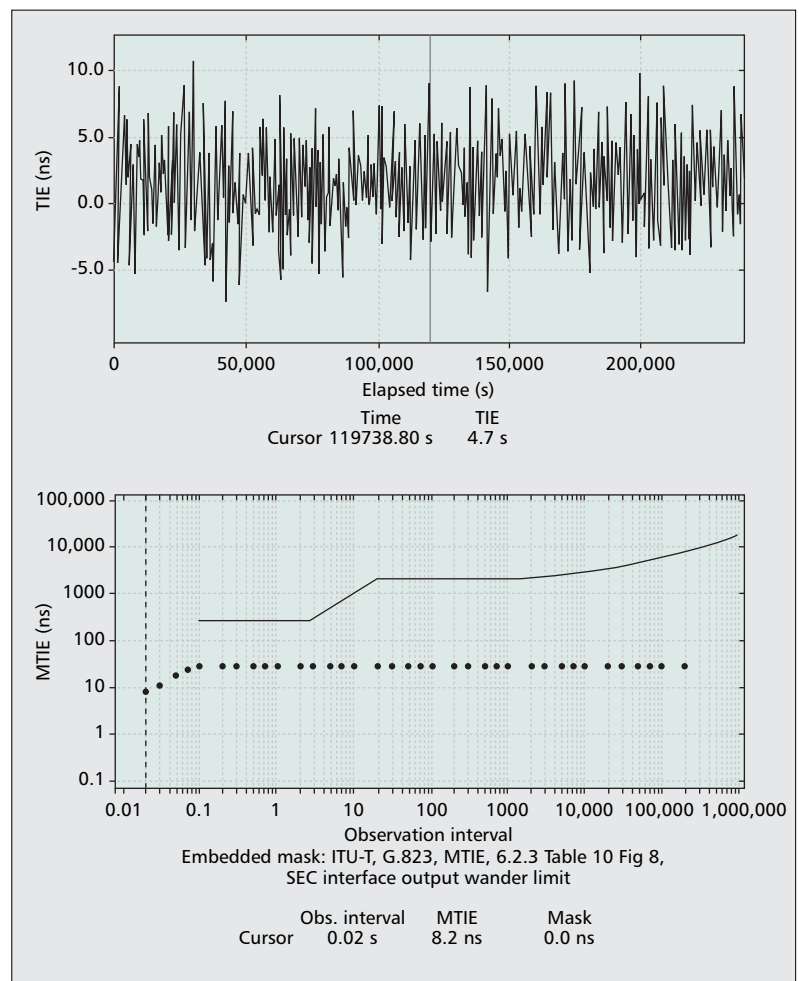
However, as networks evolve, it is expected that operators will be dealing with a mix of Ethernet equipment that may or may not provide Synchronous Ethernet capabilities. An extra step may be required in network planning simply to ensure that Ethernet links on equipment with appropriate Synchronous Ethernet functionality are selected as timing links. This extra step contrasts the SDH/SONET case, where all NEs by definition are capable of being integrated into the synchronization network.

ACHIEVABLE RESULTS

This section discusses the performance that can be achieved by Synchronous Ethernet based on a subset of results that have been presented within ITU-T for the development of G.8262.

The test and measurements were done according to Fig. 1a of ITU-T G.810 [11]. The testbed, shown in Fig. 5, consisted of a timing chain of four cascaded Ethernet switches, where the clock was transported across the network through 1000BASE-X and 100BASE-TX interfaces. In addition, Ethernet traffic, not shown in the figure, was sent through the Ethernet switches, and the last Synchronous Ethernet port was congested with a subscription ratio of 10:1.

Figure 6 shows the time interval error (TIE) and maximum TIE (MTIE) for a three-day period. TIE is a metric used to represent the phase movement of the recovered signal (here the Ethernet recovered clock at the end of the timing chain) against an ideal reference. The filtering bandwidth of the EEC clocks was set to 10 Hz for testing purposes, but could have been set as low as 0.1 Hz, which would have produced even lower TIE/MTIE performance. The TIE is bounded within approximately ± 10 ns, and the MTIE is bounded to about 25 ns; therefore, the measured MTIE meets the required mask. These results, combined with the fact that EEC Options 1 and 2 meet the same noise transfer and noise generation requirements as for SDH and SONET, respectively, clearly indicate that the clock chain requirements in G.803 are met. Based on this, it is expected that current network syn-



■ **Figure 6.** *Achievable results: TIE and MTIE for three days.*

chronization guidelines and practices can be retained as operators migrate from SONET/SDH toward Ethernet. In addition, the results are not impacted by packet network impairments such as traffic load, latency, or delay variation (10:1 oversubscription in this test).

CONCLUSION

Measured results show that very high levels of frequency transfer performance are achievable with Synchronous Ethernet, delivering phase

A key aspect to success is that it does not change fundamentally the Ethernet technology or the standards that describes Ethernet. Architecturally the well defined layers within Ethernet allow this evolution in Ethernet to occur.

performance in the region of a few tens of nanoseconds. Such performance replicates that achievable from SONET/SDH transport networks, providing security of synchronization supply as well as a migration path from TDM to Ethernet for service delivery. Furthermore, as such a solution uses the physical layer, it has been shown to be immune to traffic load and packet delay variation.

A key aspect to success is that it does not fundamentally change the Ethernet technology or the standards that describe Ethernet. Architecturally, the well defined layers within Ethernet allow this evolution in Ethernet to occur. Lengthy analysis of this architecture and its description within standards, together with appropriate reuse of existing standards and components, has enabled Synchronous Ethernet to be developed relatively easily. It is also compatible with the current synchronization network.

Such a solution will allow communications providers considerable flexibility in serving many nodes and applications with defined levels of performance, and provides a well understood evolution path from existing transport networks to Ethernet-based transport networks.

REFERENCES

- [1] ITU-T Rec. G.8261, "Timing and Synchronization Aspects in Packet Networks," Feb. 2008.
- [2] ITU-T Rec. G.8262, "Timing Characteristics of Synchronous Ethernet Equipment Slave Clock (EEC)," Aug. 2007.
- [3] ITU-T Rec. G.8264, "Distribution of Timing Through Packet Networks," Feb. 2008.
- [4] ITU-T Rec. G.781, "Synchronization Layer Functions," Feb. 2008.
- [5] IEEE 802.3-2005, "CSMA/CD Access Method and Physical Layer Specifications," 2005.
- [6] IEEE 802.3ay (IEEE P802.3Rev) Maintenance #9 (Revision) Task Force, expected 2008.
- [7] ITU-T Rec. G.803, "Architecture of Transport Networks Based on the Synchronous Digital Hierarchy (SDH)."
- [8] ITU-T Rec. G.813, "Timing Requirements of SDH Equipment Slave Clocks (SEC)."
- [9] ITU-T Rec. G.812, "Timing Requirements of Slave Clocks Suitable for Use as Node Clocks in Synchronization Networks."
- [10] ITU-T Rec. G.707, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)."
- [11] ITU-T G.810, "Definitions and Terminology for Synchronization Networks."

BIOGRAPHIES

JEAN-LOUP FERRANT (Jean-Loup.Ferrant@alcatel-lucent.fr), graduated from INPG Grenoble, France, joined Alcatel in 1975, and worked on analog systems, PCM, and digital crossconnects. He has been working on SDH synchronization since 1990 and on SDH and OTN standardization for more than 15 years in ETSI TM1 and TM3, and ITU-T SG13 and SG15. He has been rapporteur of SG15 Q13 on network synchronization since 2001. He is one of Alcatel-Lucent's experts on synchronization in transport networks.

MIKE GILSON (mike.gilson@bt.com) is a technologist with BT Design at Adastral Park. He joined BT in 1983, and has played a major role in developing and implementing BT's Time and Timing strategy from 1988 to the present. He is currently actively contributing to the ITU and IETF standards bodies, and the ITSF and NIST synchronization forums. He has a B.A. (Hons) degree in business studies and is a member of the Institute of Engineering and Technology.

SEBASTIEN JOBERT (sebastien.jobert@orange-ftgroup.com) is an R&D engineer, Network Synchronization, France Telecom. He graduated in computer science and telecommunications from Paris 6 University, and joined France Telecom R&D in 2005 after one year working on VoIP aspects. He worked on data synchronization aspects for a year, participating in the Open Mobile Alliance (OMA) DS WG. He is currently working on network synchronization aspects, and follows Q13/SG15 at the ITU-T.

MICHAEL MAYER (mgm@nortel.com) graduated from Queens University (B.Sc.EE). In 1982 he joined Bell-Northern Research (now Nortel). For the past 12 years, his area of focus has been systems design and standardization, and he has actively contributed to standards in the areas of synchronization, OTN, SONET/SDH, and ASON control plane architecture. Within the ITU, he has been a technical editor of several ITU-T Recommendations covering the ASON optical control plane and synchronization related topics.

LAURENT MONTINI (lmontini@cisco.com) is a technical leader in the office of the CTO at Cisco. He has focused his work on frequency and time distribution over next-generation networks for the last two years, participating in the standardization efforts since 2005. Prior to this function he spent six years as a consulting system engineer for mobile operators dealing with service and traffic convergence in backbone and radio access networks where the emergence of pseudo-wire techniques raised the timing issue.

MICHEL OUELLETTE (ouellett@nortel.com) joined Nortel in 1997 and is a member of scientific staff in the CTO group. Current responsibilities include the design, performance, and standardization of network synchronization architecture, frequency and time distribution for carrier-Ethernet and 2G/3G/4G wireless backhaul, as well as the development of algorithms for Mobile IP. He is involved in ITU-T SG15/Q13, has published 15 journal/conference papers, and has eight patents granted.

SILVANA RODRIGUES (silvana.rodrigues@zarlink.com) is a system architect, Timing and Synchronization, Zarlink Semiconductor. She graduated in electrical engineering from Campinas University, Brazil. She started her carrier at the Telecommunication Research Center, Brazil. She has been working on synchronization since 2002. She participates in several standards groups; she is the editor of the G.paclock.bis and G.8262 Recommendations in ITU-T, and secretary for IEEE-1588. She is also an active participant in the OPTXS-SYNC standard committee and TICTOC at IETF.

STEFANO RUFFINI (stefano.ruffini@ericsson.com) is an expert in R&D, and manager of Ericsson's Network Synchronization Center. He joined Ericsson in 1993 and has been working on synchronization aspects for about 15 years. He has represented Ericsson in various standardization organizations (including ETSI, ITU, 3GPP, and IETF) and is currently actively contributing to ITU-T SG15 Q13 (editor of ITU-T Recommendation G.8261, which introduces Synchronous Ethernet technology) and IETF TICTOC. He is one of Ericsson's experts involved in the definition of equipment and network synchronization solutions.