Synchronous Ethernet – A RAD White Paper

Yaakov (J) Stein, Chief Scientist, RAD Data Communications, Ltd.
Alon Geva, Timing specialist, RAD Data Communications, Ltd.

Abstract

As more and more traffic is being delivered in Ethernet format, carriers are realizing the advantages to converging on a pure Ethernet infrastructure, but are not able to do so with Ethernet as presently defined. New ITU-T Recommendations, such as Y.1731 for Ethernet OAM and G.8031 for Ethernet protection switching, provide features essential for operational maintenance of Ethernet networks. Recently proposed mechanisms for configuring point-to-point connections in Ethernet networks close another gap between Ethernet and SONET/SDH networks.

The final difference between conventional TDM-based networks and Ethernet is that the former also transport frequency information, needed for some applications, while Ethernet does not. Numerous methods have been suggested for augmenting Ethernet to distribute frequency and/or timing information, such as IEEE 1588, IETF's NTP, and methods developed for use of TDM pseudowires.

In this article I will talk another method for transport of synchronization over Ethernet networks, for those applications that need it. This method locks the timing of the Ethernet physical layer, much the same way that SONET/SDH does. This enables simple recovery of the physical layer timing, which can be used directly by locking downstream networks and applications, or indirectly by using the physical layer timing as a common reference.

Worlds apart – SONET/SDH and Ethernet

Ethernet is by far the most widely deployed network technology, being used for transport over local and metro area networks, and the demand has been steadily increasing for it to become the standard handoff interface for network interconnection.

Wide area networks have conventionally been based on TDM technologies, from low rate T1 and E1 trunks, through PDH to SONET/SDH infrastructures. While originally designed for transport of multiplexed voice traffic, over time they have been adapted to carry data traffic in general, and Ethernet traffic in particular. However, TDM technologies are still optimized for constant rate signals, and cumbersome for bursty packet sources. So typically two different network infrastructures are employed, one optimized for constant rate inputs and one for packet sources.

As more and more traffic is being delivered in pure Ethernet format, carriers are realizing the advantages to converging on a pure Ethernet infrastructure, but are not able to do so with Ethernet networks as presently defined. The high reliability and short service restoration times built into TDM infrastructures and required by network operators are still largely absent in Ethernet networks. New ITU-T Recommendations, such as Y.1731 for Ethernet OAM and G.8031 for Ethernet protection switching, provide features essential for operational maintenance of Ethernet networks. Recently proposed mechanisms for configuring point-to-point connections in Ethernet networks close another gap between Ethernet and SONET/SDH networks.

Even once these upgrades are made there remains fundamental difference between TDM-based networks and Ethernet. TDM networks require, for their own proper functioning, highly accurate frequency information to be recoverable from their physical layers. Isolated physical clocks, such as piezoelectric crystals, can not be expected to agree with each other over time. Various effects, such as temperature changes and aging, cause even reliable frequency sources to wander. Were unrelated frequency sources to be used on opposite sides of TDM links, frequency discrepancies would lead to valid bits being lost and invalid bits inserted, phenomena known as "bit slips". In order to avoid bit slips, somewhere in every TDM network there is at an extremely accurate Primary Reference Clock (PRC) from which all other TDM clocks in the network directly or indirectly derive their timing. Clocks derived in this manner are said to be traceable to a PRC.

Applications Require Accurate and Stable Frequency

The unconditional requirement for distribution of highly accurate frequency information throughout the TDM network can be turned into a feature. Some communications and data communications applications require accurate frequency, and when these are connected to TDM networks they receive this information essentially gratis. One example of this requirement is cellular phone sites. Each cellular base-station requires an accurate frequency reference from which to derive its transmission frequencies. Were neighboring cellular sites to have differing frequency references there would be mutual interference between them. Cellular base-stations in GSM networks derive their frequencies based on the timing inherent in the TDM links that deliver the incoming traffic. Were these links replaced by Ethernet-based ones, this frequency reference would be lost.

The original CSMA/CD Ethernet only transmitted when there was information to be transferred, and the transmitted frames were preceded by a preamble designed to enable the receiver to lock onto the transmitting clock. Thus, rather than maintaining global synchronization, each receiver locked onto the transmitter's timing on a frame-by-frame basis. Today's fast and Gigabit Ethernet transmit continually, but still do not maintain network-wide synchronization. In any case, highly accurate synchronization is not as crucial as it is for TDM networks as received packets are buffered, and thus small timing discrepancies do not lead to bit slips. Thus even modern full-duplex Ethernet does not have the inherent capability of providing timing information for applications that require it.

In this paper we emphasize the distribution of frequency information, but acknowledge that there are also applications that require knowledge of absolute time (known as "wall-clock" time).

Frequency Distribution in Ethernet Networks

Numerous methods have been suggested for augmenting Ethernet to distribute frequency and/or wall-clock information. The IEEE developed a protocol for precision clock synchronization in order to support networked measurement and control systems. IEEE-1588 enables submicrosecond synchronization of clocks by having a master clock send multicast synchronization message frames containing timestamps. All IEEE-1588 receivers correct their local time based on the received timestamp and an estimation of the one-way delay from transmitter to receiver.

The Network Time Protocol NTP is an IETF protocol used widely to distribute wall-clock time over IP networks and the Internet. In it too, time servers send timestamps in packet, but these are usually in response to request messages from a client, who may be in contact with several servers.

Certain communications applications that require timing have their own frequency distributions mechanisms. For example, TDMoIP® has a mechanism for recovery of accurate frequency without use of explicit timestamps under the assumption that the packets sent were created at a constant rate locked to the reference to be recovered. VoIP gateways recover approximate time based on embedded timestamps.

A few years ago there were discussions in IEEE regarding distribution of timing information via special timing frames for home entertainment applications, such as music rendition. Work on this topic is ongoing as part of the Audio/Video Bridging (AVB) Task Group. Some proponents called this method "synchronous Ethernet". The "synchronous Ethernet" we now introduce is quite different.

Synchronous Ethernet

There is yet another way to transport timing over Ethernet networks, for those applications that need accurate timing. Modern dedicated-media full-duplex Ethernet, in both copper (100BASE-TX, 1000BASE-T, 2BASE-TL, 10PASS-TS) and optical (100BASE-FX/LX/BX, 1000BASE-SX/LX/BX) physical layer variants, transmits continuously. In keeping with CSMA/CD traditions, the physical layer transmitter clock of these signals is derived from an inexpensive +/- 100ppm crystal, and the receiver locks onto it. There is no need for long-term frequency stability as the data is packetized and can be buffered. For the same reason there is no need for consistency between the frequencies of different links.

However, one *could* elect to derive the physical layer transmitter clock from a high quality frequency reference, by replacing the crystal with a frequency source traceable to a primary reference clock. This certainly would not effect the operation of any of the

Ethernet layers, which would be oblivious to the change. The receiver at the other end of the link would automatically lock onto the physical layer clock of the received signal, and thus itself gain access to a highly accurate and stable frequency reference. Then, in TDM fashion, this receiver could lock the transmission clock of its other ports to this frequency reference.

So, by feeding one network element in an Ethernet network with a Primary Reference Clock, and employing Ethernet PHY circuitry with well-engineered timing recovery circuitry of the type standard in SONET/SDH networks, we could set up a fully time synchronized network. Unlike TDM networks this timing accuracy is not required for the proper functioning of the data plane, which could function perfectly well with relatively inaccurate and inconsistent physical layer clocks. Rather it provides access to a highly accurate and stable frequency reference to applications that need it, such as cellular base-stations and TDMoIP gateways. An example application is depicted in the figure.

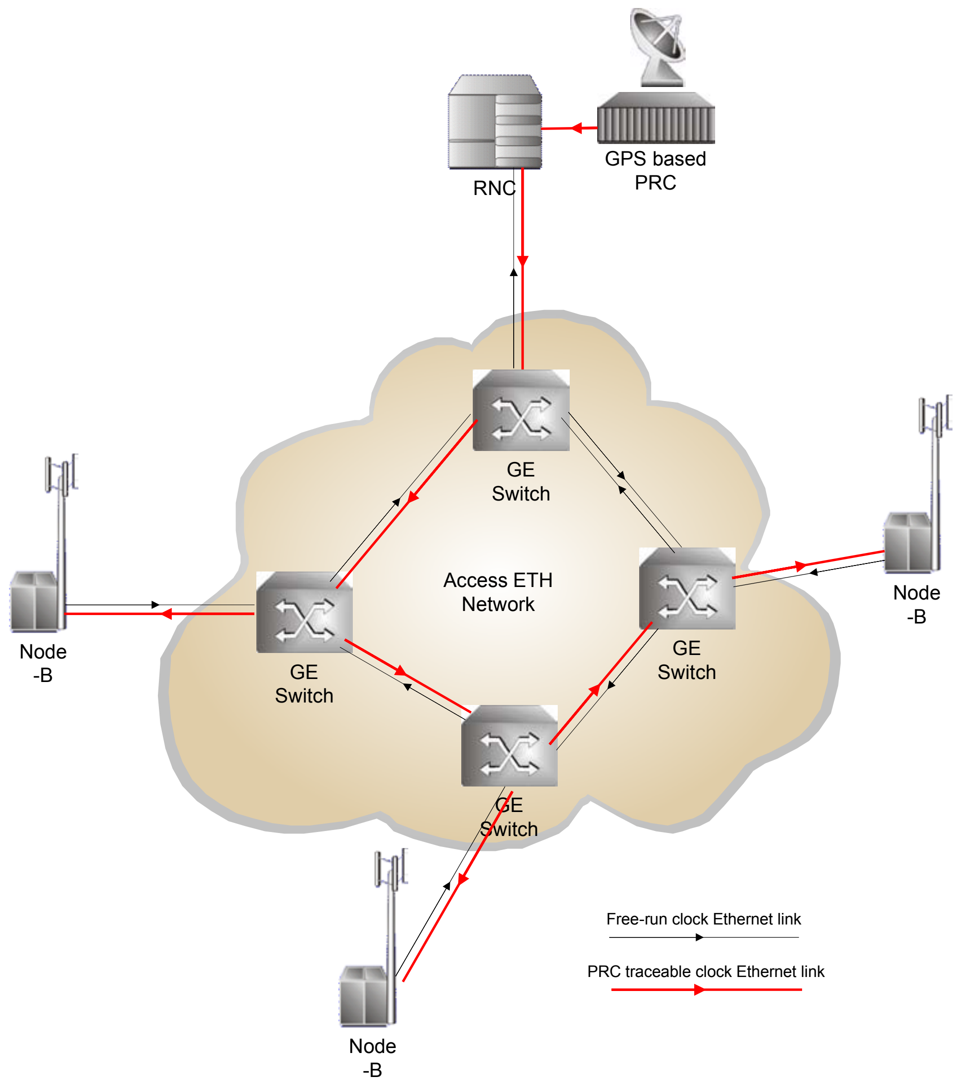

While such locking of the physical layer of Ethernet has doubtless been imagined before, the first attempt at its standardization as a network-wide solution is to be found in Clause 8.1.1 of the new ITU-T Recommendation G.8261. This standard mainly deals with timing aspects relating to transport of TDM over Ethernet networks, but while doing so explains how asynchronous Ethernet services may be mapped into a synchronous Ethernet network.

Operation of Synchronous Ethernet Networks

The advantage of using Synchronous Ethernet, as compared to methods relying on sending timing information in packets over an unlocked physical layer, is that it is not influenced by impairments introduced by the higher levels of the networking technology (packet loss, packet delay variation). Hence, the frequency accuracy and stability may be expected exceed those of networks with unsynchronized physical layers.

IEEE-1588 defines mechanisms to combat network degradations (such as ‘boundary clocks’ and ‘transparent clocks’) for difficult cases, but these mechanisms require hardware support from network elements. If new switches and routers are already required, it is difficult to see why one should choose not to implement synchronization of the physical layer, if only frequency accuracy is required.

Synchronization of the physical layer of a link means that the receiver can lock onto the frequency of the transmitter. Of course, the question remains as to the quality of the transmitter clock! Is it locked to a stratum 1 clock, or only a stratum 3? Perhaps the clock source failed and it is in holdover mode, deriving frequency from its own local clock? In TDM networks, network elements determine the traceability of the frequency being supplied using Synchronisation Status Messaging (SSM). This messaging enables signaling of stratum level, notification of failures, and taking of appropriate actions.
For SDH SSM is defined in G.707, and for lower rate TDM signals in G.704.

Appendix VIII of G.8261 describes a Synchronisation Status Messaging for synchronous Ethernet that enables downstream Ethernet switches to determine the status of those upstream. It exploits EFM OAMPDUs, and reuses G.707 status messages. When an Ethernet switch is informed of an upstream synchronization failure condition, the switch can take appropriate action, such as selecting an alternate synchronization source.

Other Applications

We have seen that synchronous Ethernet may be used to lock up all clocks in a network. However, at times an input source has its own independent clock and can not be locked to the synchronous Ethernet clock. Can such traffic be transported over the synchronous Ethernet network and recreated with its own clock upon egress?

Such scenarios are not new. ATM CES services provided a Synchronous Residual Time Stamp (SRTS) option with which the difference between an arbitrary input clocks and the physical layer clock could be encoded and transported across the network. At egress, the encoded difference could be used along with the physical layer clock to accurately recover the source clock. For MPLS networks a similar mechanism was defined in ITU-T Recommendation Y.1413, and this mechanism is directly applicable to the Ethenet case as well.

We have mentioned before that there are applications that require wall-clock, and not (just) accurate frequency. Synchronous Ethernet can be used to lock frequency throughout an Ethernet network, but does not deliver wall-clock time. On the other hand protocols such as IEEE-1588 and NTP are designed to distribute wall-clock, but as we have seen due to network degradation can not equal the performance levels of synchronous Ethernet in providing accurate frequency references.

The solution is to use the two techniques together. The delivery of wall-clock by NTP or 1588 is made difficult by the requirement that these protocols simultaneously lock on to frequency and time. When the network is synchronous the wall-clock distribution protocol can assume accurate frequency, and its job becomes much easier. Indeed, it can be shown that such a hybrid approach significantly reduces the time estimation error.