Circuit Emulation Services-over-Packet and Cellular Backhaul

by Peter Meyer, Zarlink Semiconductor

Circuit Emulation Services (CES)-over-Packet is an emerging technology that allows circuit-switched services to be carried across a packet-switched network. In previous TechNotes, we have introduced this technology and outlined current CES-over-Packet standardization work. Here we'll discuss using CES-over-Packet in a cell site backhaul application, with specific focus on latency, bandwidth and timing synchronization issues.

Traditionally, 2G, 2.5G and 3G basestations are connected to basestation controllers through multiple T1/E1 lines which may be configured for Nx64 kbps DS0 traffic, ATM traffic using Inverse Multiplexing for ATM (IMA), or packet traffic using multi-link PPP (ML-PPP).

In an effort to reduce operating costs, carriers are increasingly replacing these T1/E1 access links with a packet-based connection, such as fixed wireless or Gigabit Ethernet (GE) fiber. CES-over-Packet technology offers a solution for seamlessly carrying TDM traffic over the new packet network.

Cell Site Backhaul Application

Cell site backhaul using CES-over-Packet is illustrated in Fig. 1. The traditional leased T1/E1 lines between the 2G, 2.5G and 3G basestations and their respective controllers have been replaced by a fixed wireless or GE fiber packet network. An external box emulates multiple T1/E1 lines from multiple basestations using CES-over-Packet over a GE fiber. The CES-over-Packet technology could likewise be built into a line card in the basestation and basestation controller.

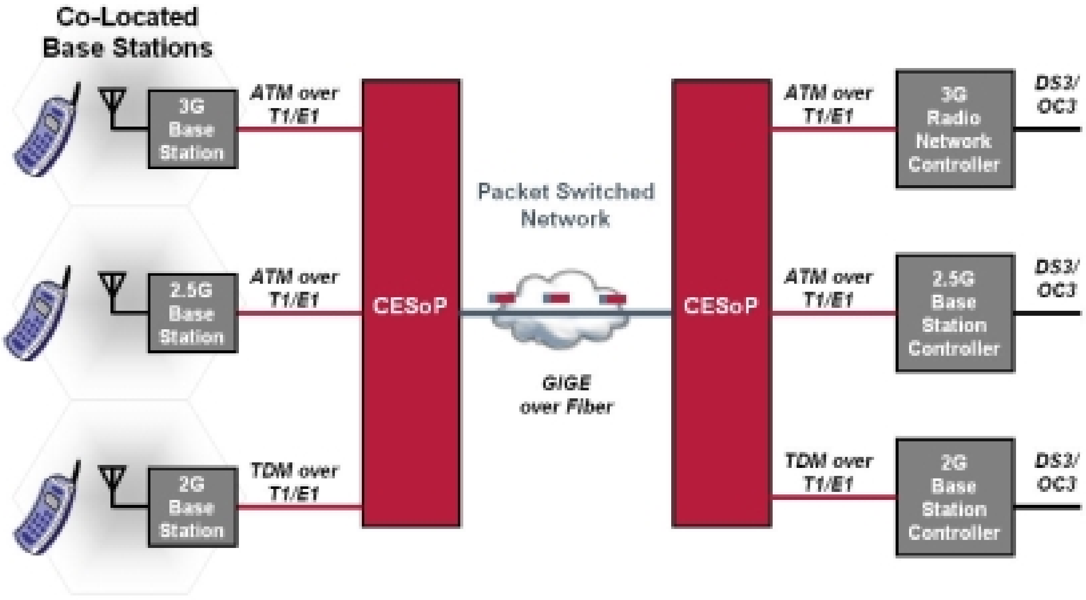

Fig. 1: Cell Site Backhaul Using CES-over-Packet

The use of CES-over-Packet technology would allow for lower cost transport between the two locations, as it replaces more expensive leased T1/E1 lines. The T1/E1 lines may all be synchronous, or each T1/E1 line may be asynchronous. The timing synchronization is carried end-to-end to ensure the T1/E1 lines at the cell site still meet the relevant T1/E1 timing standards.

Latency and Bandwidth

Latency and bandwidth are closely related, and in fact have an inversely proportional relationship when performing CES-over-Packet. A reduction in latency will result in an increase in bandwidth utilization, whereas an increase in latency will result in a decrease in bandwidth utilization. This trade-off requires that the CES-over-Packet inter-working function (IWF) be optimized to minimize latency while maximizing bandwidth utilization in a bandwidth-limited environment.

With the benefit of trunking, in many scenarios it is possible to focus solely on minimizing latency when network bandwidth is sufficient. Ideally, the target end-to-end latency for an emulated service should be as low as one would expect with a non-emulated service.

There are typically two main latency restrictions for a cell site backhaul application. First, there is an overall 150 ms latency end-to-end in the network. This latency value is divided amongst the various nodes in the network, including the cell phone and the basestation. After the division of the latency budget, the CES-over-Packet IWF may be budgeted for as low as 250 µs latency, as it would replace a TSI or framer function in the system. Alternatively, the T1/E1-to-T1/E1 latency across a CES-over-Packet connection between the basestation and basestation controller may be limited to under 1 ms, even as low as 500 µs, in order to have the same order of magnitude of delay as a traditional T1/E1 circuit.

Second, there is a latency value imposed by the maximum offset between T1/E1 lines running IMA or ML-PPP. An IMA or ML-PPP buffer, where the ATM cells or PPP packets are recombined from multiple T1/E1 lines, may support only 5 ms of latency offset between the T1/E1 lines. To ensure this requirement is met the connection between basestation and basestation controller may be limited to an absolute latency of 5 ms.

The bandwidth restrictions, or lack thereof, are dependent on the packet network between the basestation and basestation controller. Where the packet network is wireless there may be significant bandwidth restrictions, even as little as a 5% maximum packet overhead, a restriction that may require it to be as little as a sequence number and a connection identifier -- where the latency is also at a minimum. Where the packet network is a dedicated fiber connection, for example with Gigabit Ethernet transport, then there may be no restriction. This packet network would allow for all-out reduction of latency with a significant packet overhead penalty.

End-to-end latency may be estimated as the transmit latency + packet network latency + receive latency. The transmit latency is the sum of the transmit processing and the number of frames per packet * 125 µs. The receive latency is the sum of the receive processing and the delay through the receive jitter buffer which is programmed to compensate for packet network PDV. The end-to-end latency parameters controlled by the IWF are the transmit processing and the receive processing, both of which should be kept to a minimum by the IWF function.

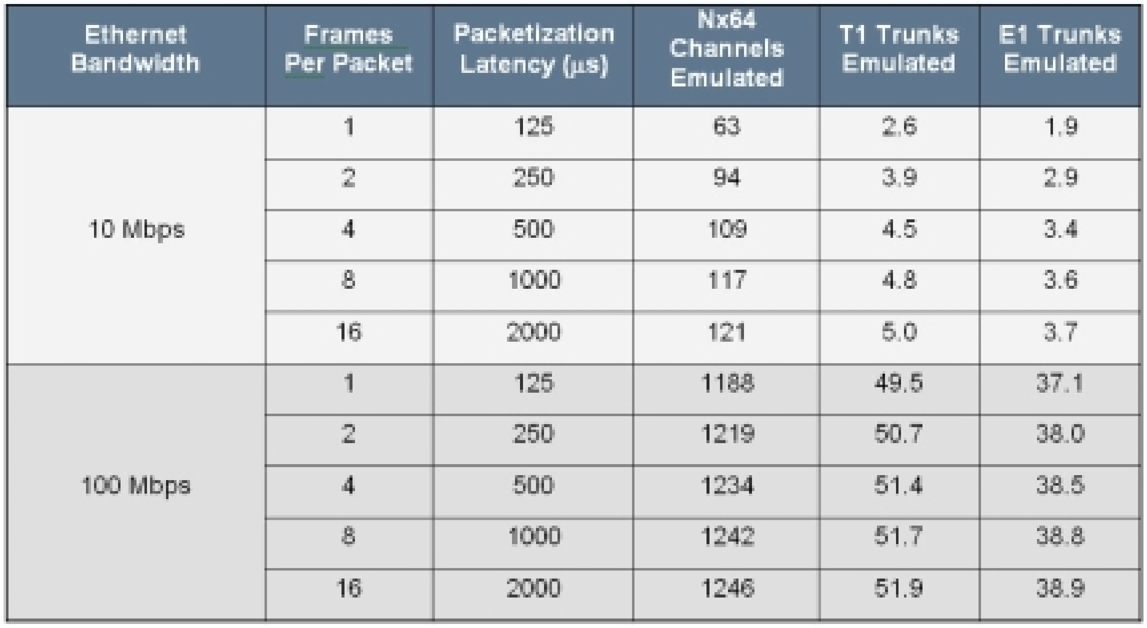

Ethernet Bandwidth	Frames Per Packet	Packetization Latency (µs)	Nx64 Channels Emulated	T1 Trunks Emulated	E1 Trunks Emulated
10 Mbps	1	125	63	2.6	1.9
	2	250	94	3.9	2.9
	4	500	109	4.5	3.4
	8	1000	117	4.8	3.6
	16	2000	121	5.0	3.7
100 Mbps	1	125	1188	49.5	37.1
	2	250	1219	50.7	38.0
	4	500	1234	51.4	38.5
	8	1000	1242	51.7	38.8
	16	2000	1246	51.9	38.9

Table 1: Latency-Bandwidth Relationship

Consider a scenario where an Ethernet connection is provided for network access. The connection may be either 10 Mbit/s or 100 Mbit/s full duplex. Table 1 shows the number of Nx64 kbit/s channels that may be trunked together and carried across a 10 Mbit/s or 100 Mbit/s connection. Table 1 also shows the equivalent number of T1 or E1 trunks, where a T1 trunk is considered to be 24 x 64 kbit/s channels or 1.5 Mbit/s of bandwidth ([24 x 64] ÷ 1024), and an E1 trunk is considered to be 32 x 64 kbit/s channels or 2.0 Mbit/s of bandwidth ([32 x 64] ÷ 1024). The numbers factor into account the bandwidth overhead required for 38 byte packet header and a network throughput of 80% of rated bandwidth. A 10 Mbit/s connection has sufficient bandwidth to carry over 3 E1 trunks or 5 T1 trunks with a 500 µs and 2 ms packetization delay. A 100 Mbit/s connection has sufficient bandwidth to carry well over 50 T1s or 38 E1s with a 250 µs packetization delay.

Timing and Synchronization

A primary concern when providing synchronous TDM circuit emulation service across an asynchronous packet network is clock synchronization. In a traditional TDM circuit-switched network the clock is "built-in" to the physical layer service, whether that be T1, E1, T3, E3, SONET or SDH.

In a packet-switched network, such as a typical Ethernet network, no reliable means of recovering or synchronizing the network clocks is provided. Here, clock synchronization functionality falls upon IWF. This is not a trivial challenge. An Ethernet network was designed to carry asynchronous data packets, and the quality of packet networks may vary greatly.

There is not a single timing synchronization standard that is applicable to all basestation applications. There are a few different timing standards that are requested depending on the service provided by the service provider. The most common timing standard is ITU-T G.823 traffic interface for E1 circuits or ANSI T1.403 for T1 circuits (which is equivalent to ITU-T G.824 traffic interface), which covers clock distribution. This standard is the basic timing requirement that should be met for most basestation applications.

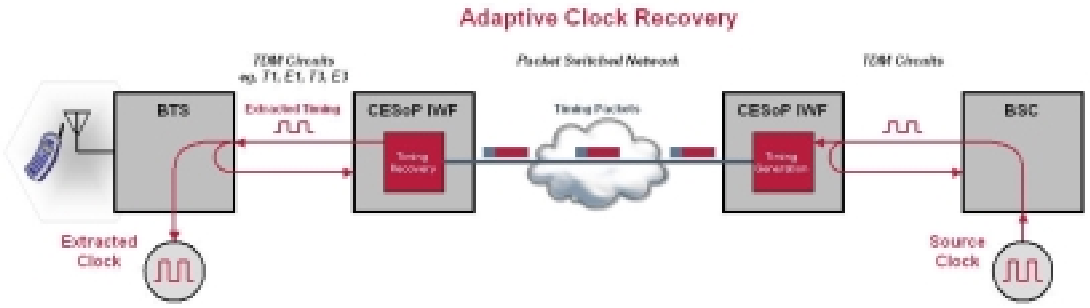
Adaptive Clock Recovery

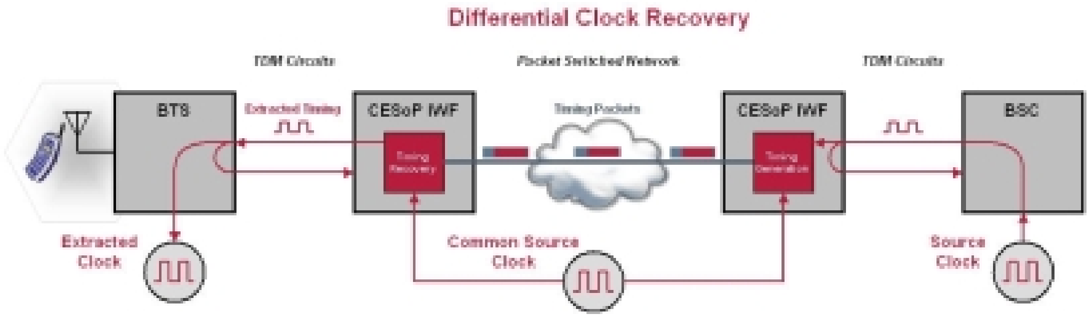
Differential Clock Recovery

Fig. 2: Clock Synchronization Across A Packet-Switched Network

A more stringent timing synchronization standard that is sometimes required is ITU-T G.823 synchronization interface for E1 circuits or ANSI T1.101 for T1 circuits (which is equivalent to ITU-T G.824 synchronization interface). This timing standard would permit the basestation to re-time additional portions of the network. This would be useful where one basestation would be connected in a daisy chain fashion to a second basestation by T1/E1 lines.

A third timing synchronization requirement is applicable for UMTS or GSM basestations. The 3GPP specification requires ±50 ppb absolute frequency accuracy. The ±50 ppb budget would be divided between multiple components inside the basestation. Thus, the target for timing synchronization from the CES-over-Packet connection should be approximately ±25 ppb.

A fourth timing synchronization requirement, and the least common, is applicable for CDMA basestations. This synchronization specification requires 3 μs absolute time error on every T1/E1 circuit.

The choice of mechanism to transmit and recover timing between a basestation controller and a basestation is dependant on the availability of a common reference source at both ends of the emulated circuit. In the absence of a common reference source, then adaptive clock recovery, typically using sequence numbers or timestamps in the header of the data packets, is used.

Differential clock recovery, which provides superior performance to adaptive clock recovery, may be used where a common reference source is available at both ends of the CES-over-Packet connection. The common reference source may be obtained using a GPS receiver at the basestation, and extracted from the physical wireless link for wireless backhaul applications. Alternatively, the common reference source may be extracted from a point-to-point fiber connection between the basestation and basestation controller that uses a 125-MHz Gigabit Ethernet clock. Should the packet network be run over a SONET/SDH physical connection, then the 19.44-MHz SONET/SDH clock may be used as the common reference source.

One unique aspect of the cell site backhaul application is that it is not uncommon for each T1/E1 circuit to be independently timed. A service provider may emulate multiple T1/E1 circuits from multiple carriers that have basestations co-located at a cell site. This scenario requires that each carrier, and perhaps each T1/E1 circuit, has independent timing transferred between the cell-site and the basestation controller locations. Alternatively, if the CES-over-Packet function is built into a basestation, then a single clock may be transferred between the basestation and basestation controller. That clock is then used to time all the T1/E1 circuits.

It should be noted that timing does not need to always use adaptive or differential techniques. Some applications, even where a packet network is used to carry CES-over-Packet traffic, will still use a T1/E1 line to carry a BITS clock to the basestation. Another approach is to use the GPS receiver to carry the synchronization clock, rather than as a differential clock source.

Measurements

A lab experiment measured the characteristics of a CES-over-Packet-based backhaul connection, and compared these results against the traditional backhaul connection.

Four T1 lines were setup between a basestation and a switch. The lines were configured for ML-PPP operation. Two of the T1 lines were connected directly between the basestation and switch. The other two T1 lines were emulated over a CES-over-Packet connection. The packet network was a simple Ethernet switch.

The CES-over-Packet connections (one per T1) were unstructured, packetizing 24 bytes per packet with a 1 ms jitter buffer. The measured round-trip latency was 4 ms (resolution of ±0.5 ms), or approximately 2 ms of one-way latency. Without using CES-over-Packet, the round-trip latency was 2 ms (resolution of ±0.5 ms), or approximately 1 ms of one-way latency. Thus, the addition of the CES-over-Packet IWF and packet network added 2 ms of round-trip latency, or 1 ms of latency in each direction. The additional measured latency was in line with the 1 ms of latency added by the jitter buffer, the 125 μs packetization delay and minimal delay through the packet network.

The timing for two of the T1 lines was embedded in the packet header of the CES-over-Packet connections. The timing synchronization performance met the requirements of the basestation. When the packet network Ethernet cable was disconnected, the switch and basestation reported AIS and LOF. Once the Ethernet cable was re-connected, the end-to-end connection again operated error free.

About The Author

Peter Meyer is an applications engineer with Zarlink Semiconductor's Packet Processing group.

as published in...
analog ZONE