

Reducing the Cost of Wireless Backhauling Through Circuit Emulation

Abstract

Data rate requirements of backhaul connections for wireless base transceiver stations (BTSs) continue to increase, while the cost of available Gigabit Ethernet connections decreases. As a result, IP/Ethernet backhauling has become a prime choice for new installments. However, for the hundreds of thousands of basestations deployed with time-division multiplexed (TDM) connections (E1/T1), carriers must pay much more for TDM lease lines, but can cut their costs by using circuit emulation service (CES) to transport the signals over a less expensive Ethernet network.


For any CES application, clock recovery at the far end of the link represents a major challenge, as the interworking function (IWF) blocks need to support differential and adaptive clock recovery. For more demanding applications, such as synchronization of wireless basestations, carriers can apply a hybrid timing generator (HTG). This paper discusses the protocol details of CES, issues around clock recovery, and the network, node, and device architectures of CES solutions. The paper will also cover the benefits of FPGA-based implementations.

Introduction

As an increasing number of applications demand higher bandwidth, and Internet usage continues to expand, telecom operators need to adapt their networks. Most of the networks used by telecom operators today are based on TDM systems. These are ideal for transporting voice traffic and supporting leased-line applications, but with data traffic having become more dominant over the last few years, operators are migrating their networks to packet-based systems. In fact, some operators are even planning all-packet or all-Internet protocol (IP) networks as their next-generation solutions.

For a smooth migration from a TDM-based to a packet-switched network (PSN), operators need to create a homogeneous network that uses packet technology in its core backbone, the metro, and the access networks. With reduced capital and operational expenses, a single network can be a substantial commercial advantage. However, even as operators plan all-IP networks, they still need to allow transport of TDM traffic over a PSN, since the TDM equipment used by business and government is not likely to be replaced at the same time. What's more, network providers generate much more revenue from leased-line services than they do from data networks (Figure 1). Heavy competition from new entrants in the data services market forces established companies to lower their rates to keep up. Within the existing laws of Europe and the United States, even when telecom operators migrate to an all-IP network, they are legally bound to provide TDM access points for third-party resellers and wholesalers.

Figure 1. Comparison of TDM Leased Line vs. Ethernet Services Revenue (Source: IDC, 2003)



For TDM and leased line services, pricing pressure is less of an issue. In fact, many established network providers are seeing an increase in leased-line revenues due to the growth of networks for wireless backhaul. With the revenue difference between data and leased-line services, CES has emerged as a popular method for established players to maintain their revenue stream while the core of the network switches from native TDM lines to a packet-based infrastructure. This approach only works, of course, if end users are guaranteed the quality and performance levels they expect.

Impact on the Wireless Network

Many countries are experiencing stability in the number of voice calls, but an exponential increase in data services usage. This trend is also showing up on wireless backhaul networks. As data usage increases, the technology for high-speed data access on mobile networks is rapidly evolving to support the demand.

However, the demand is causing increasingly visible strain on wireless backhaul networks. A typical radio basestation (RBS) can support a few hundred voice calls in parallel. On the backhaul link to a mobile switching center (MSC), a bandwidth of about four E1 or T1 lines is reserved per RBS. While a moderate mobile data connection today provides 384 Kbps, the next-generation networks already offering downlink speeds ranging from 1.8 to 14.4 Mbps (UMTS HSDPA phase 1). In the following planned phases (HSPA Evolved and Long Term Evolution [LTE]), the downlink data rate ultimately reaches 200 Mbps, while the uplink may provide up to 100 Mbps. To support this exponential growth in access link speeds, wireless backhaul networks need to be rebuilt using bigger pipes, such as those built with Ethernet technology. Continued bundling of E1 or T1 TDM pipes for transport simply will not be adequate to support these increases in data traffic.

While an Ethernet-based PSN is ideal to transport data services, voice calls are TDM-based and need special handling on a PSN. Furthermore, a wireless basestation needs a high-quality clock reference to ensure its radio signal does not disturb neighboring basestations and to allow hitless handover of voice calls as mobile users move from station to station. Most basestations currently use a TDM network interface to transport voice and data traffic, as well as distribute network synchronization from a central source to all stations. Synchronization distribution is not supported by current Ethernet networks.

CES and Transport Protocols

A CES IWF allows a smooth migration from TDM traffic to packet traffic and vice versa. Instead of using fixed bandwidth multiplexed TDM circuits, the TDM traffic is packetized and aggregated with other TDM or IP traffic onto a single network. Traditional TDM networks have an implicit dual purpose: to transport data, and to maintain network wide signal synchronization. Today's PSNs, however, do not support network synchronization. To compensate for this, network operators have created overlay networks for synchronization, or use expensive, GPS-based reference clocks.

For any CES application, clock recovery at the far end of the link represents a major challenge. The IWF blocks support differential and adaptive clock recovery, but for more demanding applications, such as synchronization of wireless basestations, design engineers can apply an HTG. A CES IWF supports several independent E1/DS1 TDM-over-packet data engines. The TDM data from each input port is processed and forwarded as a packet stream. At the remote end, the packets are reassembled into a TDM data stream. Each of the TDM ports may be independently timed on ingress, with the reassembly function supporting independent clock recovery to retim the TDM data on egress. [Figure 2](#) shows one method for encapsulating the TDM data over the PSN, where the implementation uses the popular and standardized SAToP and CESoPSN variants. (Both standards are compatible with those described by the Metro Ethernet Forum (MEF-8)).

A key feature at the IWF decapsulation side of the link is the jitter buffer, which copes with delay variation in the PSN. Each packet switch on the PSN has internal buffers, queuing, and scheduling functions that introduce a certain amount of delay. As the CES packet travels through the PSN, this delay accumulates, often differing for each individual packet. This behavior is quantified as the packet delay variation (PDV) or packet jitter of the PSN. The decapsulation side must have a large enough jitter buffer to allow for a certain PDV.

Typical implementations of IWF have a jitter buffer of approximately 10 to 50 ms per channel. When data transport over a CES link begins, the receiver waits a specified amount of time before beginning playout of the buffer to the TDM circuit. A good starting condition begins when the buffer is approximately half full. Without this delay, a change in PDV would quickly lead to an underrun of the buffer. Optimally, the jitter buffer should be large enough to deal with the full range of PDV observed on a PSN. Conversely, there are upper bounds to how much delay or latency can be allowed on an end-to-end TDM circuit. To allow operation with various types of networks and to find an optimum balance between PDV and latency, the size of the jitter buffer should be a configurable parameter.

In addition to jitter buffer size, the TDM packetization length also determines latency. The SAToP encapsulation method allows the user to determine how many TDM bits are accumulated into a single packet. For an E1 TDM link, the default is 256 bytes per packet, or 1 ms of TDM data. The implementation should allow for configuration of the amount of packetization in steps of 125 μ s. To reduce latency, choose the 125 μ s configuration. It should be noted, however, that the amount of overhead per packet is a fixed amount. With a low latency encapsulation, the packet overhead becomes a major factor and can more than double the data rate from TDM to the packet layer. The SAToP encapsulation provides a straightforward and simple emulation service, and the installation and service turn-up of a SAToP link is easy because there are only a few general connection items to be provisioned (e.g., the destination system and the circuit identifier to indicate the TDM channel).

CESoPSN

The alternative encapsulation method is CESoPSN (defined in IETF RFC5086), which provides more bandwidth efficiency than SAToP. It allows the user to select which timeslots (64-kbps channels or DS0 channels) will be transported over the PSN. In bandwidth-limited access networks, such efficiency makes economic sense. For the IWF to identify these DS0s, a TDM framing function is needed to find the frame alignment signal and the individual timeslots. The packetized data is aligned with the TDM frame structure by one or more complete TDM frames sent in a single packet over the network. The far-end IWF requires a frame generator to recreate a complete E1 or T1 frame structure with a locally regenerated frame, and to fill in the DS0s as received from the packet stream. This IWF implementation of CESoPSN supports a software-configurable timeslot table that determines whether each DS0 is transported over the PSN. The CESoPSN standard describes various options that allow further optimization of bandwidth usage and how signaling bits are processed.

Clock Recovery and Network Synchronization

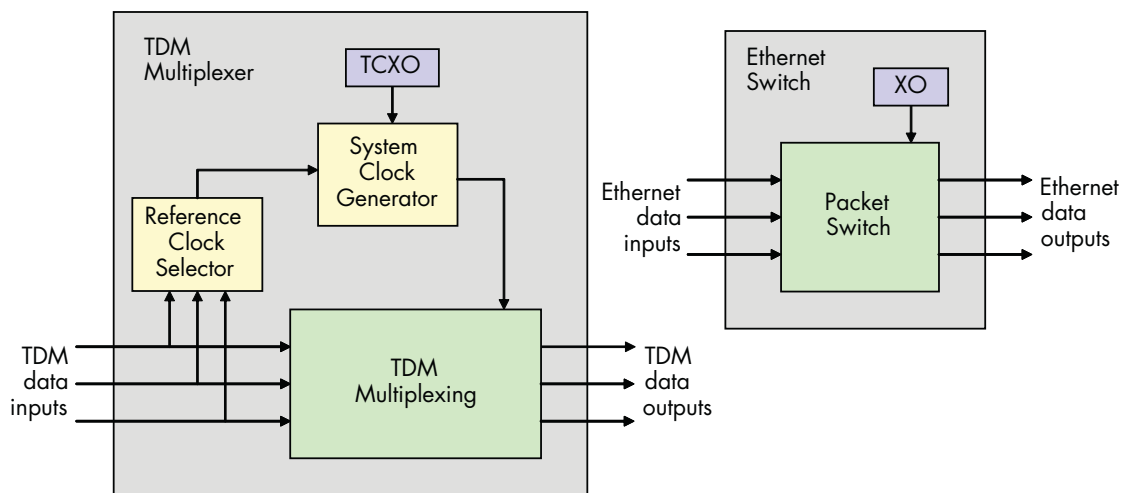
The clock recovery function ensures that the clock rate of the TDM transmitter at the decapsulation side aligns closely with that of the TDM receiver at the encapsulation side. Without accurate clock recovery, TDM bit errors result, or there is excessive jitter and wander. Bit errors are obviously unacceptable, but wander is also something that must be tightly controlled. ITU-T recommendations G.823 and G.824 define the maximum wander for TDM interfaces. As an example, the wander at the output of an E1 interface must be kept below 18 μ s. In addition, a new standard (ITU-T G.8261) specifically targets the performance for TDM interfaces carried via CES over a PSN. According to this standard, the wander must be kept below 4.3 μ s in specific scenarios. Requirements that are even more stringent apply when the TDM circuit is used as a reference clock input for a wireless basestation.

A key feature of any CES transport system is the clock recovery at the far end. In most applications using TDM circuits, such as E1, T1, or SONET/SDH, the link is used to transport the information bits and to distribute clock information through the network. While basic CES ensures reliable transport of the TDM information bits, there is no guarantee that the emulated circuit can be used to distribute the clock as well. Before examining the clock recovery aspects of a CES, normal TDM circuits and the network clock distribution must be considered.

A TDM or SONET/SDH system typically has a central clock generator (defined in ITU-T G.813) and is known as the SDH equipment clock (SEC), which is used by all output interfaces as a reference clock (Figure 4, left). A set of synthesizers, phase-locked loops (PLLs), and filters is used to generate the various internally required clock frequencies and to reduce jitter and wander accumulation in the network. A SEC can operate in free-running mode, but typically is locked to one of the input interfaces. By provisioning or via a protocol, the system software determines which interface will serve as the reference clock input. Often, redundant inputs are assigned in such a manner that a link failure also results in the reference clock switchover. With this basic setup, a complete clock distribution network can be constructed, where one node is linked to a primary reference clock (PRC) and all other nodes form a tree-shaped network for distributing the synchronization. Each end system has a clock that is traceable to the same PRC.

With a carefully designed clock distribution network based on E1/T1 or SONET/SDH, the reference clocks available at the edge of the network have such a high quality level that they can serve as the reference clock for other external networks. For example, a typical wireless basestation (RBS or BTS) can also use the E1/T1 network interface as a clock reference. This clock is then used again as a reference clock for the radio or air interface. When all neighboring radio basestations derive their clock reference from the same synchronization network, the amount of interference or bit errors on the radio network is limited.

Figure 4. Typical System Clock Architectures in a TDM Multiplexer and an Ethernet Switch



Data networks, on the other hand, are not built to distribute network synchronization among nodes as described for TDM networks. The typical architecture of an Internet protocol/multi-protocol label switching (IP/MPLS) router or Ethernet switch is that each interface card has a local clock oscillator that is used as a clock reference for local on-board transmit line interfaces (Figure 4, right). This local clock oscillator operates in free-running mode. For Ethernet interfaces, the only requirement is that the clock is within ± 100 ppm of the nominal value. Buffer overruns or underruns due to clock differences between the transmit side and receive side of a data link are handled through increasing or decreasing the idle time or interpacket gap (IPG) between packets. This mechanism is well suited to prevent data buffer overruns, but cannot distribute synchronization among nodes in a PSN. Each link has a free running clock that is not linked to other links, prohibiting any attempt at creating a clock distribution network.

In summary, two aspects need to be considered when a CES are deployed over a PSN. First, the PSN does not provide a reliable clock to use as reference for the outgoing TDM stream. The PDV of a PSN makes it very difficult to recreate the clock rate of the receiving end correctly. Secondly, without a synchronization network, the TDM circuits carried via CES cannot be used as a reference to nodes downstream in the synchronization distribution network. Very specific measures are needed to drive a wireless basestation clock from a CES.

Clock Recovery Options for Emulated Circuits

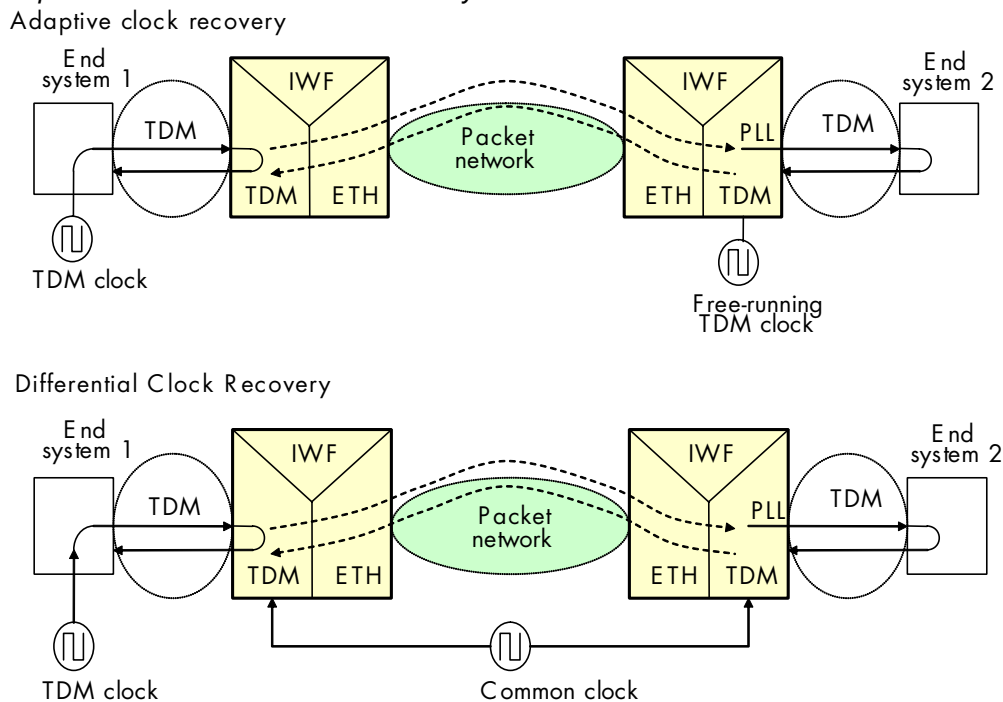
There are two main architectural network scenarios for the clock recovery of a CES: adaptive and differential mode. In both scenarios, the receiver needs to implement a frequency-locked loop (FLL), PLL or some combination of the two to precisely control the frequency and phase of the outgoing TDM stream.

Adaptive Clock Recovery

In the case of adaptive clock recovery, as shown in Figure 5, the control loop can only use the jitter buffer fill level as an input variable. When the average buffer fill level increases, the TDM clock rate must be increased to prevent overflow. Similarly, when the buffer drains too much, the TDM clock rate is decreased. This control loop must be able to handle various impairments that occur when data passes through a PSN. The most challenging aspects are PDV and packet loss. ITU-T recommendation G.8261 provides a full description of such network impairments, along with mandatory test cases.

An accurate and stable local clock is essential since the PSN does not provide a clock distribution network for TDM networks. For E1/T1 transport networks, a clock accuracy better than +/-50 ppm deviation from the nominal clock frequency is required. The main factor that affects clock stability, which determines the deviation of clock frequency over time, is temperature variation. A temperature-controlled crystal oscillator (TCXO) or an ovenized oscillator (OCXO) can reduce temperature dependency. Long-term clock stability is the key factor for adaptive clock recovery. As packets are carried over the PSN, the delay may vary over the course of the day. This 24-hour pattern on PSNs corresponds with human activity during the day: high during business hours and in the evening, due to busy network links, and progressively lower through the evening and night. This variation must be handled in the IWF, requiring a clock with a stable frequency. Failure to maintain a stable clock may result in wander performance that does not conform to ITU-T G.8261.

Figure 5. Adaptive and Differential Clock Recovery



Differential Clock Recovery

The alternative clock recovery method is differential clocking, in which both the encapsulation and decapsulation IWFs have access to a common reference clock. The encapsulation side inserts a timestamp into each transmitted packet and the decapsulation side retrieves the timestamp, using it as the main input parameter to the FLL or PLL to

control the frequency and phase of the outgoing TDM stream. The differential scenario handles the network impairments on the PSN much more easily, as does not rely directly on the arrival time of each packet. Instead, the timestamps are used to control the outgoing bitrate.

Several network architecture solutions have been proposed for network-level synchronization, with some already in use in provider networks. One popular scheme is to use a PRC at the core side of the CES link and a GPS receiver at the access side of the CES link. Although the GPS receiver provides a good quality reference clock for the CES IWF, many equipment sites face signal reception issues from GPS satellites. Cost is another concern because back-ups are based on high-quality clock oscillators. In addition, the reliance on technology over which telecom network providers have no control, such as if a GPS satellite network is switched off or jammed temporarily, can influence the performance of the network or even cause it to fail.

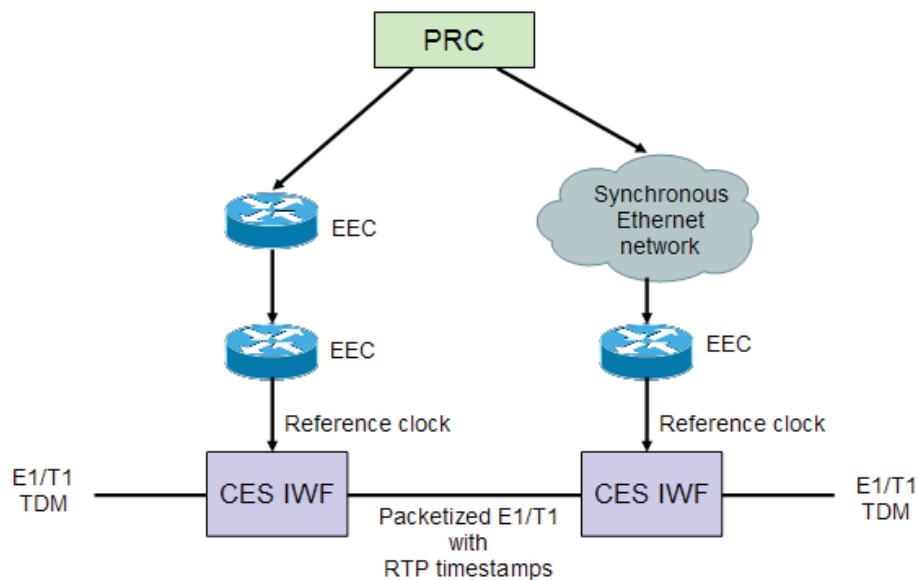
New Network Synchronization Architectures

Two new network synchronization architectures for the PSN have been proposed and are currently in process of standardization. The first is Synchronous Ethernet and is defined by ITU-T G.8262, the other is commonly referred to as Timing over Packet and is based on IETF NTP or IEEE1588 standards.

Synchronous Ethernet

With Synchronous Ethernet, expected to deploy in the second half of 2008, the physical layer of the Ethernet link will distribute the clock among nodes in the network (see Figure 6). Very similar to SONET/SDH, each node has a local clock, or Ethernet equipment clock (EEC, under definition in ITU-T G.8262) that determines the outgoing clock rate of each interface and is derived from the incoming clock of one of the input interfaces or a local free-running oscillator. The Synchronous Ethernet approach calls for careful construction of a tree-shaped network that includes all nodes to distribute a common clock reference derived from the PRC to each CES IWF, which uses differential clock recovery based on RTP to generate the clock rate precisely for the outgoing TDM stream.

Figure 6. CES With RTP and Differential Clocking Based on Synchronous Ethernet



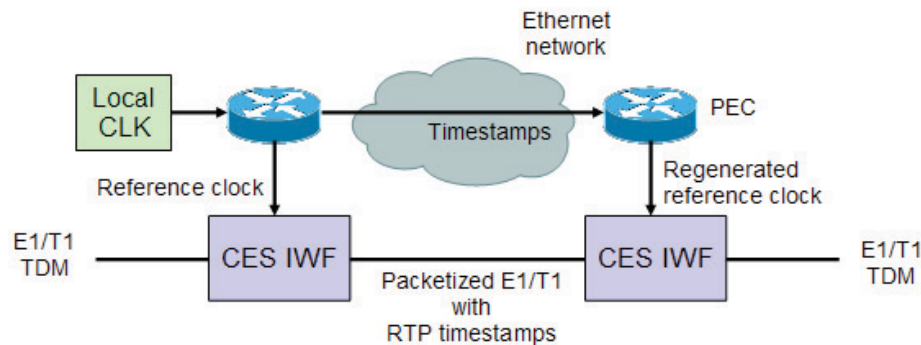
Based on the same architectural structure as SONET/SDH, Synchronous Ethernet is expected to yield a performance level comparable to that being achieved on SONET/SDH networks today. However, until Synchronous Ethernet is available for widespread use, this approach can only be used for end-to-end network synchronization when all intermediate switching nodes in the network have hardware and software support. Any Ethernet switch or router currently being manufactured will break the synchronization network, because each outgoing port is clocked from a

free-running oscillator (see bottom of [Figure 4](#)). Because of this, Synchronous Ethernet is currently only usable in green-field or new networks, such as small network domains, where long-term provider can cover the whole network through equipment upgrades.

Timing Over Packet

Based on the deployment limitations of Synchronous Ethernet, the industry is also showing interest in Timing over Packet, an alternative that is less dependent on the intermediate network. Timing over Packet provides a short-term solution. Instead of using the physical-layer clock, it relies on timestamp messages to be distributed among the nodes in a network. Dedicated time servers in a network send out regular timestamps to client systems. The clients use the timestamp to measure client usage in order to synchronize their local clock frequency and phase. In addition, the concept of a packet equipment clock (PEC, under definition in ITU-T G.8263) has been proposed that will use packet-based timestamps to generate a reference signal for its local system clock, as shown in [Figure 7](#). A high-quality reference clock is used as a reference on one CES IWF, with the timestamp messages processed by the PEC to regenerate a reference clock for the second CES IWF. A dedicated HTG is under development to provide both Synchronous Ethernet and Timing over Packet protocol support for a CES IWF.

Figure 7. CES With RTP and Differential Clocking Based on Timing Over Packet



However, the existing Timing over Packet protocols need to be enhanced to create a telecom carrier-grade network synchronization architecture. Proposed enhancements include clock performance parameters that allow basestation synchronization in wireless networks, and redundancy mechanisms that allow timing clients to make controlled and hitless switches between multiple redundant timing servers if the node or part of the network fails.

Flexible FPGAs a Perfect Match for CES Design Challenges

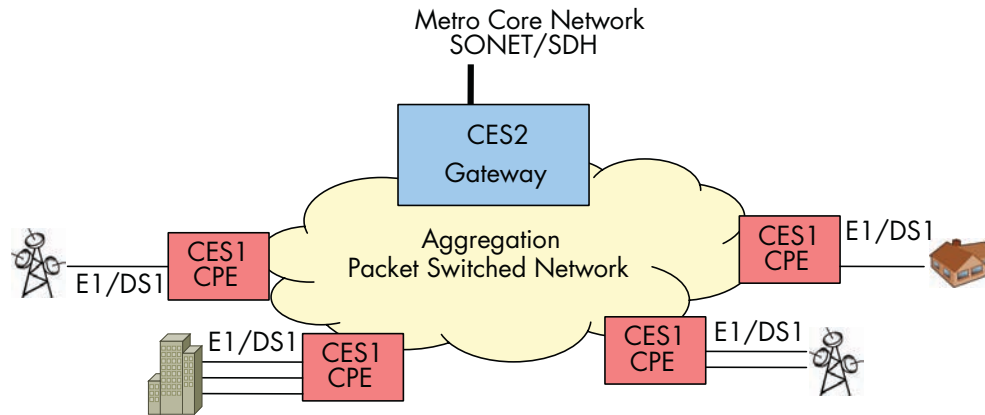
Through collaboration between FPGA manufacturers, such as Altera, and intellectual property (IP) providers, turnkey IP cores are now available that can implement all of the key aspects of a CES system on an FPGA-based platform. With such designs, system developers can easily customize the number and type of TDM interfaces needed, add custom logic or protocols to differentiate functionality, or adapt CES packet header fields. Given that CES network architectures continue to evolve, the reprogrammability of FPGAs is highly advantageous.

The SAToP and CESoPSN encapsulation modes are implemented with IP cores, using standard signal interfaces to integrate the cores easily into a production system. Of the two solution variants, CES1 targets cost-sensitive and compact end systems, such as customer premises equipment (CPE) or wireless network terminals where up to 32 TDM channels can be supported on low-cost FPGAs. The other, CES2, is ideally suited to high channel-count gateways or aggregation nodes. This gateway IP supports CES for hundreds of E1/T1 TDM channels, and can support channelized OC-12 or STM-4 capacity. Alternatively, this IP can be used as the core for a high-density E1/DS1 circuit-emulation card in a router or switch.

With CES1 and CES2 based on the same core architecture and implementation, users can choose the best fit for their application. [Figure 8](#) depicts a combined network design using CES1 at the edge nodes and CES2 in an aggregation site. Optimal interoperability is guaranteed where the specific enhancements built into the IP are available end-to-end

on the emulated circuits. In addition, the CES IWF cores can be extended with an optional on-chip high-density E1/T1 framer to provide TDM monitoring and framing or alternatively support a SONET/SDH mapper/demapper for E1/T1 signals.

Figure 8. CES1 Access and CES2 Gateway Cores Combined in a Network Design




The IP cores for CES provide a complete turnkey solution that uses a combination of hardware state machines and filtering techniques that relieves software and CPU performance bottlenecks. For example, the CES cores do not require an external CPU for continued operation. Only the configuration provisioning must be provided and the alarm conditions or performance counters regularly read out. All protocol processing, control loops for the FLL/PLL, and clock recovery are handled internally by the IP cores.

Summary

CES IP cores are built to comply with the strict requirements for carrier class networking, and users can be confident that extensive and thorough testing has been carried out to ensure conformance with all applicable standards for CES, TDM interfaces, and PSNs as well as network synchronization. Protocol interoperability of CES or synchronization with other off-the-shelf products or devices is part of the evaluation tests. IP cores should be extensively tested in real network scenarios to reduce the possibility of project risks associated when introducing a new technology in a system or network. This will allow system designers to focus on integrating system components, instead of worrying about the network performance of the emulated circuits.

The performance evaluation reports and measurement results of the various network test scenarios should be available to customers interested in using these IP cores. Reference configurations for each IP core, such as the one shown in Figure 9, are one way that system designers or network operators can evaluate the IP or qualify it for their unique CES configuration.

Figure 9. CES2 Reference Design Configuration



Further Information

- AimValley BV:
www.aimvalley.com
- Pseudowire and Circuit Emulation Services:
www.altera.com/end-markets/wireline/applications/pwe3/wil-pwe3-ces.html
- Overview of Pseudowire and Circuit Emulation Technology:
www.altera.com/end-markets/wireline/applications/pwe3/overview/wil-pwe3-ces-overview.html
- SAToP (IETF RFC4553):
www.ietf.org/rfc/rfc4553.txt?number=4553
- CESoPSN (IETF RFC5086):
www.ietf.org/rfc/rfc5086.txt?number=5086
- Metro Ethernet Forum (MEF-8):
<http://metroethernetforum.org/PDFs/Standards/MEF8.pdf>
- ITU-T recommendation G.813:
www.itu.int/rec/T-REC-G.813/en
- ITU-T recommendation G.823:
www.itu.int/rec/T-REC-G.823/en
- ITU-T recommendation G.824:
www.itu.int/rec/T-REC-G.824/en
- ITU-T recommendation G.8261:
www.itu.int/rec/T-REC-G.8261/en
- Synchronous Ethernet:
www.itu.int/rec/T-REC-G.8262/en
- Timing over Packet:
www.ietf.org/html.charters/ntp-charter.html
http://ieeexplore.ieee.org/xpl/freeabs_all.jsp?tp=&isnumber=4278948&arnumber=4278949&punumber=4278947

Acknowledgements

- Jan Venema, Chief Technical Officer, AimValley BV
- Sri Purisai, Strategic Marketing Manager, Communications Business Unit, Altera Corporation



101 Innovation Drive
San Jose, CA 95134
www.altera.com

Copyright © 2008 Altera Corporation. All rights reserved. Altera, The Programmable Solutions Company, the stylized Altera logo, specific device designations, and all other words and logos that are identified as trademarks and/or service marks are, unless noted otherwise, the trademarks and service marks of Altera Corporation in the U.S. and other countries. All other product or service names are the property of their respective holders. Altera products are protected under numerous U.S. and foreign patents and pending applications, maskwork rights, and copyrights. Altera warrants performance of its semiconductor products to current specifications in accordance with Altera's standard warranty, but reserves the right to make changes to any products and services at any time without notice. Altera assumes no responsibility or liability arising out of the application or use of any information, product, or service described herein except as expressly agreed to in writing by Altera Corporation. Altera customers are advised to obtain the latest version of device specifications before relying on any published information and before placing orders for products or services.