Legacy telecom hits the 21st century: TDM circuits on AdvancedTCA switch fabrics

By R. Brough Turner

AdvancedTCA has gained substantial traction in the telecommunications industry because it's extremely scalable and it fully supports telecom approaches to power, cooling, packaging, reliability, management, and serviceability. The only thing it hasn't done well, at least so far, is support traditional Time Division Multiplexed (TDM) circuit switching. From the beginning, AdvancedTCA was focused on meeting the needs of the evolving telecommunication industry by, for example, supporting the move to packets and taking advantage of emerging high-speed serial interconnects and next generation silicon.

Substantial TDM legacy

But TDM has not gone away. In fact, for voice traffic, the conversion to Voice over IP (VoIP) has barely begun. IP-PBX sales began to take off in 2003 and consumer VoIP (Vonage, ATT CallAdvantage, Skype) has gotten considerable publicity over the past 12 months, but penetration rates are tiny in the context of more than 1 billion fixed and 1.7 billion mobile telephones in the world today. And globally, the mobile phone industry is adding several hundred million new subscribers each year. Some of these new users get packet-based access to data services, but they all get voice services, and virtually all these new mobile voice services use traditional circuit-switching technology. Yes, the conversion to Voice over Packet will happen, but it will take decades, during which the new VoIP infrastructure will have to interwork with traditional TDM networks and TDM devices.

Of course this need was not lost on AdvancedTCA working groups. The AdvancedTCA base specification includes a full set of redundant TDM clocks. These clocks are also available on Advanced Mezzanine Cards (AMCs) and are being incorporated into the evolving MicroTCA specification. In each case, the assumption has been that TDM interface modules convert incoming TDM signals to packet protocols so TDM can be transported over the native switch fabric (Ethernet, StarFabric, and other switch fabrics) in the AdvancedTCA chassis. But available packet protocols couldn't meet all the requirements for legacy TDM support.

VoIP protocols such as RTP/UDP/IP emulate individual phone calls over the Wide Area Network (WAN). As such they introduce substantial latencies, for example, 5 ms, 10 ms, 20 ms for basic packetization, and three times that or more, on an end-to-end basis. Furthermore, schemes for multiplexing multiple VoIP streams have always been an afterthought, not widely standardized.

There are also two competing WAN protocols called TDM over IP (TDMoIP) and Psuedo Wire Emulation Edge-to-Edge (PWE3). But both of these were designed for edge-to-edge over a WAN. As such they include sophisticated clock recovery schemes, but no switching. For TDM on an AdvancedTCA backplane, we don't need clock recovery as we already have robust clock signals, but we do need switching for many applications. So work on TDM transport has been ongoing within PICMG since late 2003, and the solution is now available.

Internal TDM

In March 2005, PICMG announced final ratification of the Internal TDM specification (PICMG SFP.1), a.k.a. I-TDM. This is a TDM-over-Packet standard optimized for AdvancedTCA and similar packet backplane applications. Indeed, I-TDM provides more capability than classic TDM buses like the H.110 bus used with CompactPCI. And while early focus has been on I-TDM over Ethernet, I-TDM is transportable on a wide variety of AdvancedTCA recognized switch fabrics.

The AdvancedTCA requirements for I-TDM are:

- Dense TDM interconnection between modules
- Low latency TDM switching (500 μsec desirable; < 1.6 ms mandatory)
- High capacity (much greater than H.110; OC-12's worth per module desirable)
- Efficient (support high call rate and high switching rates)
- Easy to interface to general purpose CPUs

The sidebar gives an I-TDM overview. The next two sections show system examples where I-TDM solves fundamental implementation issues and subsequent sections present I-TDM protocol details.

Wireless backhaul optimizer

In mobile networks, a mobile switching center handles traffic for hundreds or thousands of cell sites. The links between the radios at these cell sites and the central switching facility use TDM or ATM protocols carried on T1/E1 facilities, typically leased from the local incumbent

SIDEBAR

I-TDM Overview

- High capacity: Approximately 10K DS0s per module
  – 100Ks of DS0s per shelf
- Low latency: 1 ms and 125 msec packet rates
- Low overhead: Up to 512 DS0s per packet
- Software friendly: 64-bit aligned fields
- TDM switching via grooming
  – Easy grooming of DS0 channels within packet flows
- Options for multi-DS0 *bonded* channels and transport to CAS signaling

carrier at considerable expense. Backhaul optimizers reduce the required bandwidth using lossless compression. Figure 1 shows the blades in the central site equipment shelf that handles hundreds of backhaul links. A key issue is latency. Mobile telephony already involves substantial delay, which you can observe by calling mobile-to-mobile to a friend in the same room. Even a few milliseconds of additional delay can impact voice quality.

At the Central Office, multiple T1 or E1 trunks are delivered to the optimizer over redundant OC-3 optical facilities. One OC-3 pair faces the cell sites while the other OC-3 pair faces the central equipment. The 6+1 redundant AccessGate 1000 cards shown in Figure 1 are protocol processors capable of handling compression for several hundred bidirectional DS0 equivalents. Since the basic increment is either E1 or T1, I-TDM is used to carry multiple E1/T1 payloads between the optical interfaces and the media cards. The connections at the bottom of the figure show the active I-TDM flows. These packet flows are physically carried over the backplane Ethernet fabric (shown at the top of the figure). The entire system is TDM-centric. I-TDM provides the equivalent of a TDM bus and switch fabric but with less cost and substantially greater scale than traditional TDM bus implementations.

Media processing platform

Figure 2 shows an enhanced services platform that includes both the applications processor(s) and the media processing components to support multiple applications like voice mail, interactive voice response, conferencing, and delivery of other audio content to fixed and mobile telephones. For systems of a few hundred channels or less, all media services can be provided with software running on a general purpose processor. At higher densities, special purpose DSP blades provide better performance per watt and per slot. In either case, with VoIP protocols, each media stream has its own packet flow of 50, 100, or 200 packets per second, so several hundred media flows require tens of thousands of packets per second. Processing full RTP/UDP/IP headers at such packet rates can consume more than half the MIPS of a general purpose processor and add substantially to the requirements on a DSP blade.

When providing media services into a TDM network, it makes even less sense to convert the TDM signals to VoIP protocols within the chassis. A much more efficient approach is to use I-TDM within the AdvancedTCA chassis. With I-TDM at a 1 millisecond packet rate, a single packet flow of 1,000 packets per second handles hundreds of media channels and the processing overhead for the 64-bit aligned I-TDM headers is significantly less than that required for VoIP protocols. As a result, processing overhead at the media processor is reduced from 50 percent of CPU MIPS to perhaps 2-4 percent of CPU MIPS.

At the line interface, conversion from TDM (T1/E1/OCx) to I-TDM has roughly the same hardware complexity as an H.110 bus interface. Conversion from VoIP protocols to I-TDM protocols has similar complexity. So even in a purely VoIP system, it can be advantageous to convert VoIP protocols to I-TDM at the line interface so as to make more efficient use of media processor resources.

I-TDM on a range of AdvancedTCA fabrics

I-TDM is simple protocol that can be implemented on top of a variety of switch fabrics. The principal things I-TDM needs from a Layer 2 protocol are a flow ID and timestamps. When these are not available, a shim layer called System Fabric Plane (SFP) is available to fill the gap. In particular, Gigabit Ethernet is the most widely used AdvancedTCA switch fabric and the first target for I-TDM. As Ethernet headers do not include timestamps or a flow ID, SFP is used.

For PCI Express Advanced Switching Interface (ASI), an adaptation layer that supports flow IDs and timestamps is already available in the form of Peripheral Interface 2 (PI 2). The Serial RapidIO case is still to be determined. See Figure 3.

I-TDM protocol details

Figure 4 shows the relative arrangement of headers for I-TDM over the SFP shim over Ethernet. Notice that header fields are 64-bit aligned for easy manipulation in software on general purpose processors. SFP timestamps are in 5 μsec increments. The flow ID identifies the destination as well as the payload type (control or data), the payload format, and the packet rate.

Within I-TDM, there are two basic payload formats, one used to emulate low-latency TDM transport and switching at a nominal 125 μsec packet rate and one focused on media processing that uses a nominal 1 ms packet rate. The 1 ms packet format is required while the 125 μsec packet format is optional. Figure 5 shows the internal packet organization for 1 ms I-TDM packets. This is ideal for media processing and most other purposes, except TDM switching. In this format, the 16-bit channel ID identifies a specific destination, for example a specific DSP on a media processing board containing multiple DSPs.

Figure 6 shows the internal packet organization for 125 μsec I-TDM packets. In the 125 μsec payload format, individual channels are identified by their location within the packet. A 64-bit channel management structure, sent once per packet, manages up to 512 bytes (DS0s) per packet flow. Each channel management command must be acknowledged and activated before additional commands can be sent, so updating 512 DS0s can take as much as 512 ms. On the other hand, a bonded group of 24 or 30 channels can be established with one command, so systems like that of Figure 1 are able to establish a complete mapping of multiple E1s in just a few milliseconds.

Because channel identification is based on position, the endpoints must maintain state information. When there are no new commands, the channel management structure is used for a cyclic reaffirmation protocol so endpoints can confirm their synchronization.

I-TDM grooming

With multiple DS0s per packet flow and active media processing that includes new calls arriving and old calls ending, it is inevitable that individual DS0 channels within a packet flow will get fragmented and the data structure will develop gaps where active calls have terminated. To handle this, I-TDM includes a grooming function that allows easy reallocation of TDM channels within one or more I-TDM packet flows. The source node is responsible for bandwidth optimization. As channels are reassigned, the destination node reacts accordingly.

I-TDM call control

Nodes within an AdvancedTCA system are typically autonomous, that is they are loosely coupled with other nodes using well defined protocols over a packet-based control fabric or a control VLAN running on top of a shared control and data fabric. Establishing and clearing I-TDM connections between autonomous nodes is done


Figure 1

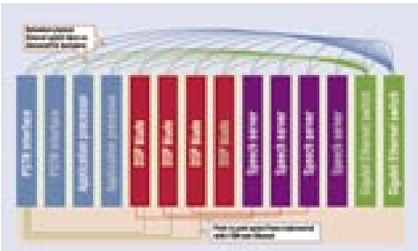
Figure 2

using standard call control stacks, preferably Session Initiation Protocol (SIP). This may be by direct interaction of the two endpoint nodes or it may be done with SIP 3rd party call control if a separate application processor needs to directly control connections between modules.

In a practical implementation, there may be a lower level software interface, for example, a device driver for an I-TDM chip or a device driver for the operating system on a general purpose processor. Due to the variety of potential operating environments, the interface to such a driver has not been standardized. On the other hand, the multiplexing of channels into a single packet is completely specified and is handled by the chip, firmware or device driver implementing I-TDM. In other words, at the SIP layer, connections look like point-to-point TDM links, not multiplexed TDM packets. This allows SIP to simply connect channels (DS0s or bonded groups of DS0s). The lower layers completely manage the multiplexing of many channels into a single packet, hidden from the SIP layer.

I-TDM status

The I-TDM specifications grew out of work that had been going on within specific companies between 2000 and 2003. A PICMG work group was formed in December 2003 to pursue standardization of both the I-TDM and the SFP protocols. The work was substantially completed in late 2004 and formally adopted by PICMG in March 2005. At this time there are pre-standard implementations that are commercially available and proprietary implementations of the standard but, as yet, no off-the-shelf I-TDM components or interoperability tests.

Figure 3

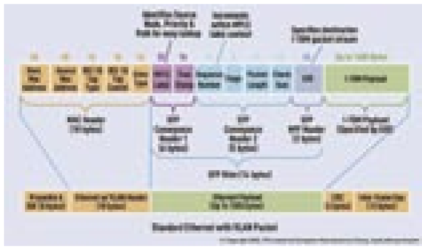
Figure 4

I-TDM provides a standard solution for a critical need, so it appears likely to gain widespread adoption – and not just within the AdvancedTCA community! With legacy TDM systems likely to persist for decades, I-TDM provides a convenient path to leverage modern switch fabrics and modern processors in any context that needs to support legacy telecom. We expect to see I-TDM demos in the latter part of 2005 and widespread deployments in 2006.

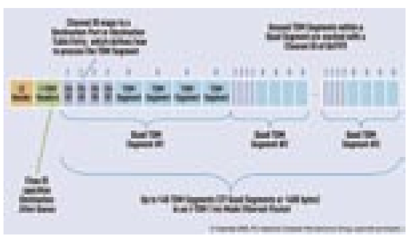
Figure 5

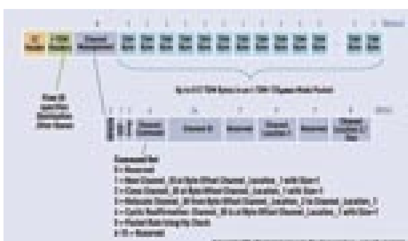
Figure 6

R. Brough Turner *is SVP, CTO, and co-founder of NMS Communications, where he oversees the evolution of NMS's technology and product architectures and works on business strategy and new market development. Brough has broad business experience, but focuses on engineering, technology, and products. Besides AdvancedTCA, his current interests include mobile video, multimodal applications, and wireless infrastructure.*

Brough writes and is quoted widely on telecommunications topics in both trade and general business publications, and he is a frequent speaker at telecom industry events around the world. He holds a BSEE from the Massachusetts Institute of Technology.

For more information, contact Brough at:

NMS Communications
100 Crossing Blvd.
Framingham, MA 01702
Tel: 800-533-6120 • Fax: 508-620-9313
E-mail: rbt@nmss.com
Website: www.nmscommunications.com