AdvancedTCA™

Brough Turner, Senior Vice President, Chief Technology Officer
NMS Communications

The Advanced Telecommunications Computing Architecture (AdvancedTCA) is a relatively new open industry standard for building high-performance telecommunications and data communications systems. Developed by the PICMG consortium (http://www.picmg.org), AdvancedTCA (http://www.advancedtca.org) appears poised to make considerable inroads into a market that has traditionally been dominated by vertically integrated proprietary systems.

A comparison of the telecom and computer industries is striking. The computer industry moved to open standards decades ago, as the IBM PC and its competitors first created a desktop market and then evolved to dominate the rest of the computing industry. As a result, the computer industry has segmented horizontally with separate markets for chips, computers, operating systems, databases and applications. And each segment has many vendors yielding vibrant competition and continuous innovation.

In contrast, the bulk of the telecom equipment market remains vertically integrated. The telecom equipment industry leverages components from the computer industry, but the resulting systems are based on proprietary hardware in custom-designed chassis. There has been interest, particularly from the computer industry, in providing standard chassis and standard internal system interfaces useful to telecom equipment manufacturers. VME chassis were adopted for some telecom systems and the introduction of CompactPCI® with telecom-specific features brought an upsurge in standard chassis. But the total market share of open, industry-standard platforms in the telecom equipment market remains small. As standards originating in the computer industry, neither VME nor CompactPCI completely address the needs of the telephony market.

Today, several trends are leading to a change of approach in telecom equipment design.

- The much-hyped convergence of voice and data is actually happening — perhaps not as rapidly as was once predicted, but happening nonetheless.
- Data rates in core network systems are increasing to the point where performance is limited by I/O issues.
- Enhanced service platforms are employing an increasing number of standard computers (typically SPARC® Solaris™ or Linux® on Intel®).

- Major telecom equipment vendors have downsized and are seeking to outsource everything not core to their business.

AdvancedTCA appears to have arrived on the scene at just the right time.

WHERE DID AdvancedTCA COME FROM?

The AdvancedTCA effort began early in 2001 with a small working group outside of PICMG (the Santa Barbara group) who saw the need for "an open platform standard that meets the need of telecom infrastructure equipment for the next ten years." The initial group included a cross section of industry interests, including both telecom and computer companies, and it started from a clean sheet — not constrained by one company's desire to get a "rubber stamp" on an existing design. Most importantly, its problem statement was clear and focused; the work group was not attempting to specify a universal platform for multiple markets.

In November 2001, the Santa Barbara group turned its groundwork over to the PICMG organization, which kicked off development of the PICMG 3.x series of specifications. Interest was extremely high. Over 100 companies signed on to participate and more than 11,000 person-hours went into conference calls and face-to-face meetings over the next 15 months. As of mid-2003, there is a complete set of specifications and a number of real products that have gone through interoperability testing.

AdvancedTCA OVERVIEW

The focus of the AdvancedTCA work within PICMG was to define a telecom platform. While the AdvancedTCA platform may be useful in other contexts, all necessary tradeoffs were made in favor of telecom industry requirements.

Equally important for the work group was providing a structured growth path for the next decade or more. As a result, AdvancedTCA supports multiple switch fabrics, while maintaining standard mechanics, backplane, system management, power distribution and cooling.

Serial Links and Switch Fabrics

As silicon performance has improved, traditional I/O capacity has fallen behind. To address this, new data transport architectures have emerged. The first stage was a move from parallel buses to high-speed self-clocked serial data streams that eliminate the clock skew problems of parallel buses. Multiple high-speed self-clocked serial streams are more efficient than a parallel bus for a fixed set of pins, wires or traces on a circuit board.

Next came the move to switch fabrics. Serial links work best when connected point-to-point, avoiding the extra capacitance of multi-drops. Point-to-point links, however, are less likely to get data where it needs to go without switching. But as improving silicon performance has made switch chips affordable, high-performance data transport systems are adopting serial links interconnected by switches. The most widespread example of this has been the transition of Ethernet LANs from shared cables to hubs, and now to switches. Other notable switch fabrics include Fiber Channel, Infiniband and StarFabric. In each case, the result is a very high-performance system that scales easily, simply by adding more links and additional switches.

Switch fabrics have an additional advantage over multi-drop connections, whether serial or parallel. With a multi-drop bus on the backplane, the failure of any one board can drag down the whole bus and potentially the whole system. With point-to-point links, only the links to the failed board are affected. And with dual switch fabrics, even the failure of a switch won't affect system interconnectivity.

Among switch fabrics, Ethernet has the largest market share, but other fabrics have performance advantages for specific applications and each fabric has an evolution path, so there is no single best choice. AdvancedTCA has been designed to support a variety of fabrics as the industry evolves over the next decade or so, and a range of topologies that can be optimized for specific applications.

Structure of the Specifications

AdvancedTCA is defined by a set of specifications: a base specification that includes a common backplane and separate fabric specifications that detail how specific switch fabrics can be implemented on the standard backplane. The specifications are:

- PICMG 3.0 — AdvancedTCA Base Specification, the master document that defines everything except the implementation of specific switch fabrics
- PICMG 3.1 — Ethernet Switch Fabric, which defines 1 and 10 Gigabit Ethernet fabrics and an option for Fiber Channel
- PICMG 3.2 — Infiniband, which defines how to build systems using Infiniband switch fabrics
- PICMG 3.3 — StarFabric, which defines how to build systems using StarFabric switch fabrics
- PICMG 3.4 — PCI Express, which defines the use of PCI Express Advanced Switching
- PICMG 3.5 — a recently chartered work group that is developing a specification for RapidIO on the AdvancedTCA backplane

AdvancedTCA Features

The basic features of AdvancedTCA include:

- Chassis based on 600 mm (~24") ETSI frame equipment practice, with provisions for 19" (~480 mm) and 23" (~580 mm) rack versions
- A full chassis, with active cooling, that fits in a 21" (~533 mm) (12U) vertical rack space; smaller chassis are permitted
- Front-to-back clearances that allow simultaneous front and rear I/O taking into account connector bodies and fiber bend radii
- NEBS and ETSI standards for shock, vibration and environment
- 2 to 16 boards per chassis
- 280 mm (~11") by 322 mm (~13") board form factor
- 1.2" (30.48 mm) board pitch, which permits large heatsinks and larger optical connectors
- Hot swap capability on all boards and active modules
- Optional rear transition module that supports rear I/O
- Support for industry standard mezzanine modules
- Dual redundant -48 VDC power feeds
- Up to 200 watts thermal dissipation per slot (3,200 watts per chassis) with forced air cooling
- Dual-redundant Ethernet control plane (independent of data transport switch fabric)
- Sophisticated system management, which includes electronic keying, module power control, health monitoring, active cooling control
- Single board failure domain (no parallel buses)
- Redundant star and full mesh data transport backplane topologies, able to support diverse switch fabrics
- Up to 2.4 terabits per second aggregate bandwidth per chassis when full mesh topology is used

Comparison with CompactPCI, cPSB and CompactTCA

AdvancedTCA provides many advances over VME and CompactPCI. It eliminates the parallel buses (such as the PCI bus), where a single failure can affect multiple slots. It dramatically simplifies power distribution while also providing redundant power. In addition, AdvancedTCA uses a substantially larger board than either VME or CompactPCI — a board form factor adequate to support modern CPUs with substantial heatsinks.

Of course CompactPCI is evolving as well. The CompactPCI Packet Switching Bus (cPSB defined by PICMG 2.16) provides a redundant star of 1 Gigabit Ethernet links in a CompactPCI chassis. This is adequate for some applications, but still modest when compared with AdvancedTCA's options of either a dual star or a full mesh with 10 Gbps Ethernet on each link. AdvancedTCA's comprehensive management scheme also goes far beyond what was proposed (but seldom implemented) in CompactPCI.

Because AdvancedTCA represents such a substantial advance over VME and CompactPCI, there is a proposal to take many of the new AdvancedTCA features and roll them back into the CompactPCI form factor under the name CompactTCA. Work on such a specification has begun.

MECHANICAL CONFIGURATION

At first glance, an AdvancedTCA chassis looks like many other telecom equipment shelves (see Figure 1). Boards slip into a chassis not unlike a VME or CompactPCI chassis. But where VME and CompactPCI rely on IEEE 1101 mechanical standards, AdvancedTCA defines a new and flexible mechanical configuration that is optimized for high-volume, low-cost sheet metal fabrication.

Figure 1: AdvancedTCA Shelf and Board

AdvancedTCA boards are larger. Their height is 8U (versus the typical 6U for VME or CompactPCI boards) or 332.25 mm (~13") and their depth is 280 mm (~11"), giving 2.5 times more board area than standard CompactPCI boards. Inter-board spacing of 30.48 mm (1.2") allows individual boards to hold large components including the heatsinks necessary for today’s state-of-the-art CPUs, optical modules and SIMMs in dense vertical arrays.

The default chassis configuration is designed to fit into the 600 mm by 600 mm racks (a little less than 24" by 24") that are standard in Europe and many other parts of the world. Minor variants then support the 19" and 23" racks that are widely used in the US and Canada.

The depth of the chassis and its position within a 600 mm (or 24") deep rack are such that an AdvancedTCA chassis can support both front I/O and rear I/O from the same chassis. Adequate space is available at the front and the rear for fiber optic connectors and the typical fiber optic cable minimum bend radius.

As expected for a telecom chassis, all field replaceable units (FRUs) — boards, fans and any optional devices added to a particular chassis configuration — are hot swappable.

Also visible in Figure 1 are two different types of backplane connectors — one optimized for power distribution and base functionality, and the other optimized for high-speed data. As discussed later in this paper, rear I/O completely bypasses the backplane, using part of the open space above the backplane visible in Figure 1.

SHELF AND SYSTEM MANAGEMENT

The high reliability of telecom systems requires a sophisticated management infrastructure. Historically, each equipment vendor has developed its own proprietary management system. In addition, earlier open standard platforms like VME and CompactPCI have skirted the issue of standardization, providing optional management systems that have not been widely adopted.

But the sophistication of the AdvancedTCA architecture demands a standard approach to management that ensures that boards, backplanes and chassis from different vendors can work together as a system. Accordingly, the AdvancedTCA specification includes the industry’s most comprehensive management architecture and makes shelf management functions mandatory for all compliant products.

Fortunately, the AdvancedTCA management system didn’t have to be invented from scratch. It leverages work done for CompactPCI System Management Bus (PICMG 2.9) and the Intelligent Platform Management Interface (IPMI) specification, originally developed by Intel, Hewlett-Packard, NEC and Dell to define interfaces for use in monitoring the physical health of servers, such as temperature, voltage, fans, power supplies and chassis.

Shelf Management Functions

The AdvancedTCA shelf management system can be used to retrieve part numbers, serial numbers, revision levels and local sensor data from any module, plus details of software loads and software revisions. In addition, under a scheme called electronic keying, each board is interrogated to determine its power and cooling requirements before power is enabled. The chassis is also interrogated to determine total available power and cooling. Individual boards are only enabled if the shelf configuration can meet their power and cooling requirements.

With electronic keying, boards also communicate with the shelf management infrastructure before they connect to the data transport fabric. Thus shelf management can ensure that boards designed to use an Ethernet fabric don't attempt to transmit Ethernet signaling to boards with Infiniband receivers and so on.

AdvancedTCA shelf management is able to monitor all active elements and report anomalies so higher level system management can determine corrective actions. Then AdvancedTCA shelf management is available to provide active control of the power, configuration and operation of system elements.

Finally, AdvancedTCA management architecture is flexible enough to support multi-tenant equipment shelves. These are situations where one owner (the landlord) provides the AdvancedTCA shelf and other owners (tenants) own components that are hosted within the landlord's shelf.

Shelf Management Components

There are several major components in the AdvancedTCA shelf management system. At the lowest level, each field replaceable unit (FRU) — the AdvancedTCA board, fans, fan tray, power-entry modules — includes an IPM controller, which monitors and manages that unit. The IPMI infrastructure provides logical and physical communications between the distributed IPM controllers. Dual-redundant IPM buses (IPMB) physically connect the distributed IPM controllers. IPMB is based on the Inter-Integrated Circuit (I^2C) bus, a simple bus that uses only two pins (data and clock).

A Shelf Manager, preferably redundant, watches over all managed devices and detects hot swap events. The shelf manager connects to the IPMBs via a variant of an IPM controller called a Shelf Management Controller. The Shelf Manager plays the role, in AdvancedTCA, of IPMI's Baseboard Management Controller.

While AdvancedTCA boards are intelligent FRUs with their own IPM controllers, the specification also provides for managed FRUs that must be represented to the IPMI infrastructure by another intelligent FRU that can provide the required management and control services. Examples of managed FRUs include shelf power entry modules or mezzanine modules that plug into an AdvancedTCA board.

A higher-level service based on TCP/IP performs remote booting, SNMP management, remote disk services and other IP-based services. This system manager, which may be external, controls one or more shelf managers and possibly one or more systems. Connections between the system manager and the shelf manager(s) are IP-based. The AdvancedTCA specification focuses on shelf management, leaving the system manager as a logical entity outside the scope of the specification (see Figure 2).

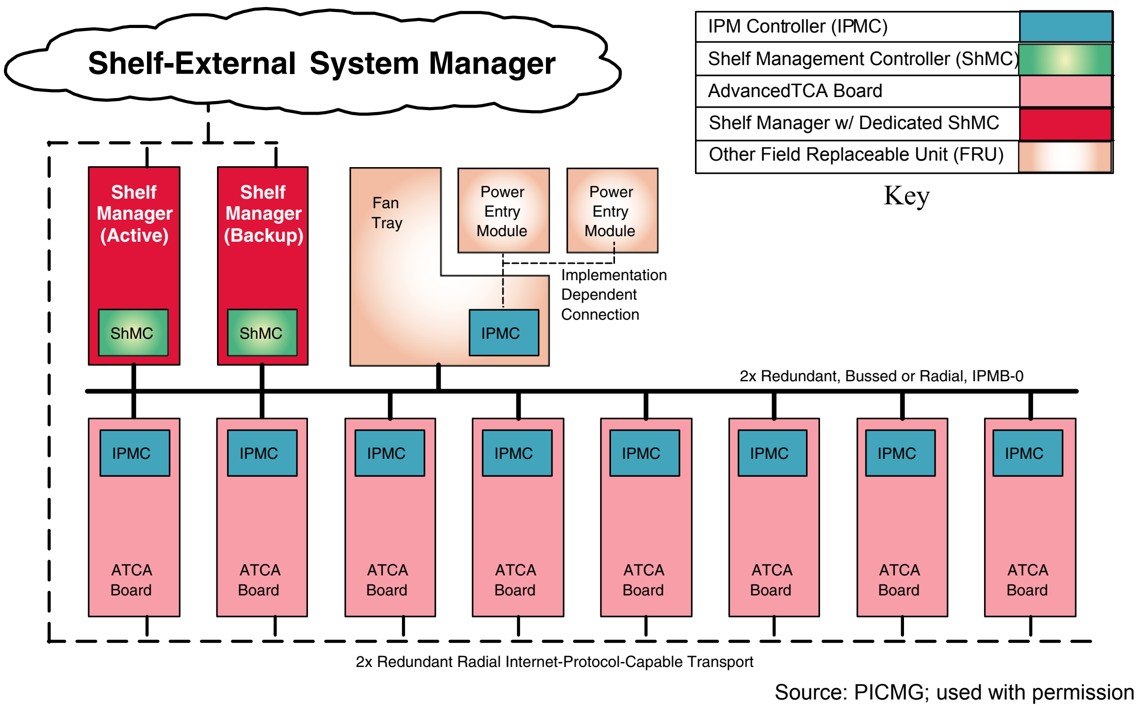

Figure 2: Example Shelf Management Components

POWER AND COOLING

The high power requirements of today's telecommunications systems requires that special attention be paid to power distribution and cooling. AdvancedTCA's large boards can be crammed with powerful components, such as CPU chips and other advanced integrated circuits that frequently dissipate tens of watts. AdvancedTCA provides for up to 200 watts of power per slot (3,200 watts per chassis) with forced air cooling.

To deliver this kind of power to each slot, low-voltage buses at 5.0 or 3.3 volts would require hundreds of amperes. Then this power would have to be further stepped down to one volt or less to power the logic cores of modern CPUs, DSPs and network processors. Because on-board DC-to-DC converters will be required anyway, the choice is to provide standard telecom power, -48 VDC, to each slot and use on-board power converters to derive the voltages actually needed on each specific board. -48 VDC has been a telecom standard for most of the history of telecommunications, so reliable, cost-effective power supplies and DC-to-DC converters are commercially available.

To ensure reliability, the AdvancedTCA chassis receives, filters and distributes dual -48 VDC feeds to each board. The AdvancedTCA power connector (visible at the bottom of each slot in Figure 1) provides the power and handles power sequencing when boards are inserted or removed. This AdvancedTCA-specific connector also handles the redundant system management ports and provides a distinct geographic address for each slot. It also has pins to support telecom ringing voltage power supplies and telecom test ports.

With up to 200 watts per slot, cooling is a critical issue. While not all systems require this level of power or justify the corresponding expense, 3,200-watt systems are possible and can continue to function despite failures of individual fans. Early in the development of the AdvancedTCA specification, a large effort was put into thermal modeling. Figure 3 shows one of the computational fluid dynamics models for AdvancedTCA airflow. Today, AdvancedTCA chassis are available with forced air cooling that support 200-watt boards while guaranteeing that no board sees more than a 10°C temperature rise from air ingress to air egress. This is possible even with 55°C intake air and one failed fan.

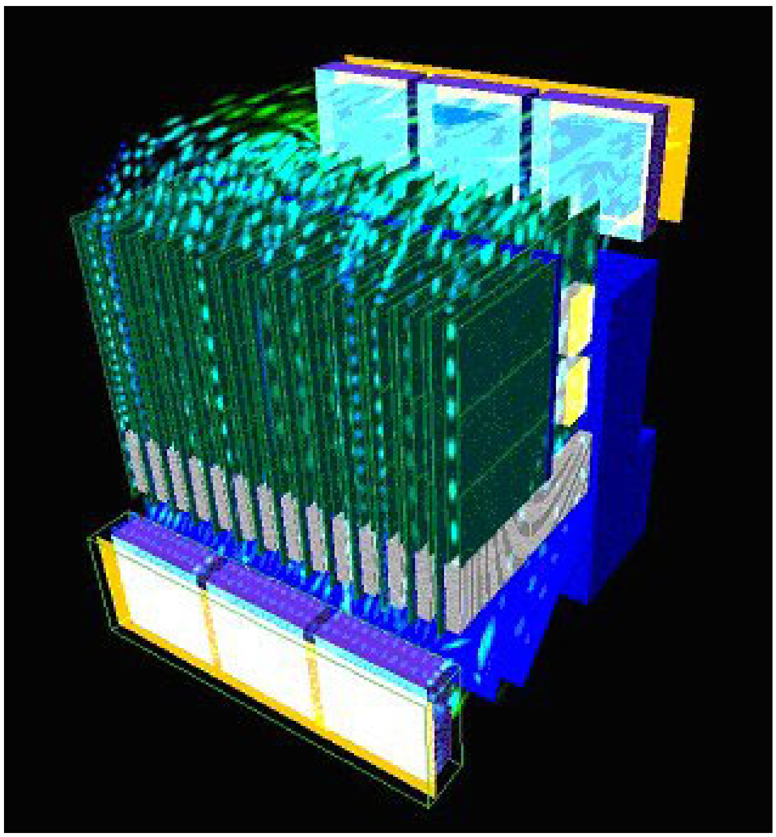

Source: PICMG; used with permission

Figure 3: Computational Fluid Dynamics Model of Airflow

DATA TRANSPORT

By focusing on switch fabric technology, AdvancedTCA supports inter-board data exchange at performance levels far beyond those of VME, CompactPCI and most prior proprietary designs. And AdvancedTCA does this in a fashion that can evolve with technology and market requirements.

The base specification (PICMG 3.0) defines the physical connectors used to mate boards with the backplane, the mapping of signals on these connectors, and the routing of signals across the backplane, i.e. between boards. AdvancedTCA data transport uses ZD connectors, which support high-speed differential signaling at up to 5 GHz, providing enough headroom to ensure that future technologies with higher signaling rates will be usable within the AdvancedTCA framework. In addition, the AdvancedTCA architecture provides a way to build an extremely flexible backplane, based on a full mesh of inter-slot connectivity.

Figure 4 is a representation of the AdvancedTCA backplane. Zone 1 carries the power and system management buses discussed previously. Zone 3 is available for rear I/O, to be discussed later, and Zone 2 supports inter-board data transport.

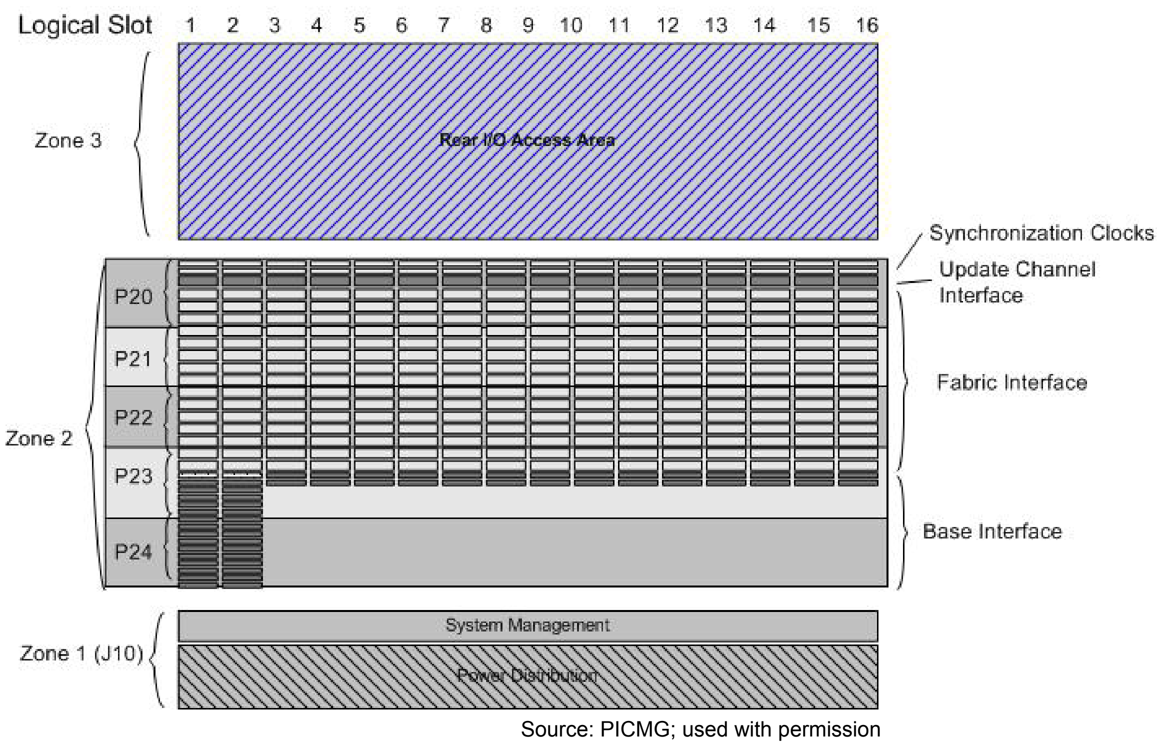

Figure 4: AdvancedTCA Backplane

In Zone 2, backplane signal mapping provides four separate interfaces between boards:

- **Fabric Interface**, which provides a full mesh of inter-board connections; the PICMG 3.0 specification defines the physical framework and subsidiary specifications (3.1, 3.2, …) define specific implementations
- **Base Interface**, which supports dual star 10/100/1000 Base-T Ethernet interconnections and is fully specified in the PICMG 3.0 specification
- **Update Channel Interface**, which connects pairs of slots so redundant devices in adjacent slots can communication with each other for private purposes
- **Synchronization Clock Interface**, which provides a set of redundant clock buses to support applications that require the exchange of synchronous timing information

Fabric Interface

Switch fabric systems utilize several different topologies. In a star topology, each board is connected to one board that contains a central switch. In a dual star, each board has separate connections to each of two separate switch boards. A dual-dual star supports two distinct and redundant switch fabrics across the backplane. Finally, in a mesh topology, each slot has a direct connection to every other slot. If a board includes a switch, it can directly route traffic to any other board. More importantly, all boards can intercommunicate simultaneously. This means a 16-slot AdvancedTCA shelf, with a full mesh of 10 Gbps interconnections, has an aggregate available bandwidth of 2.4 terabits per second.

Significantly, the full mesh configuration for the Fabric Interface of AdvancedTCA is a superset of the other useful topologies including star, dual star, dual-dual star and the full mesh, making it the most flexible backplane alternative. The electronic keying capability within AdvancedTCA shelf management is used to ensure that boards using the fabric interface are communicating with other compatible boards.

Of course, the functions on some boards may not justify the expense of an on-board switch and full mesh connection. AdvancedTCA boards designed for use in a dual star or even a single star will function in full mesh systems — they merely use a subset of the available capacity.

Fabric Interface Board Variants. Boards can be classified by how they support the Fabric Interface. There are four variants:

- **Mesh-enabled boards** support mesh connections to between 2 to 15 other boards; the number of channels supported determines the maximum number of boards that can be interconnected on a shelf
- **Node boards** are intended for use in a star or dual star topology where a Hub board provides the switching functions
- **Hub boards** are intended for use with star or dual star topologies and typically provide switching for all 15 of the other slots in a chassis; Hub boards may support

either or both of the base interface and one of the fabric interfaces defined in PICMG 3.x subsidiary specifications
- **Fabric Interface not supported** boards, which are suitable for applications that require only the capacity of the Ethernet base interface and do not need to connect to the fabric interface

Fabric Interface Technology Alternatives. The AdvancedTCA base specification (PICMG 3.0) defines a flexible backplane that includes 120 differential signaling pairs for a full mesh fabric interconnection. But it doesn't specify the technology that runs over these connections. This is left to subsidiary specifications, some available today and others that may be developed in the future. While there are fabric interface variants, however, there is general agreement on using Ethernet for control traffic, and complete agreement on using the PICMG 3.0-mandated infrastructure for system management.

At present, two subsidiary specifications have been ratified and three more are in varying stages of development:

- PCIMG 3.1 Ethernet/ Fiber Channel, ratified March 2003
- PICMG 3.2 Infiniband, ratified March 2003
- PICMG 3.3 StarFabric, ratified May 2003
- PICMG 3.4 PCI Express, ratified May 2003
- PICMG 3.5 RapidIO, working group chartered in June 2003

Base Interface

The base interface is configured for dual star operation with switches in logical slots 1 and 2. In other words, there are two base interface channels provided on every slot except slots 1 and 2. Each base interface channel consists of four differential pairs that carry 10/100/1000 Base-T Ethernet signaling between boards. Base interface capabilities are controlled by the shelf management system, however, the base interface drivers are allowed to begin auto-negotiation without waiting for the host FRU to be enabled.

Update Channel Interface

The update channel interface provides ten differential signal pairs for a point-to-point connection between two logically adjacent slots. A typical use of the update channel interface is to allow coordination between an active and a standby processor. The technologies and the protocols used over the update channel interface are not defined as it is expected that any two boards utilizing the update channel will be from the same vendor. Electronic keying is used to ensure that the two boards wishing to communicate are compatible.

Synchronization Clock Interface

While newer telecommunications systems utilize voice-over-packet or voice-over-cell technologies, many telecommunications applications still need to interface with traditional time division networks via T1/E1 trunks or SONET/SDH optical facilities. These applications typically require strict timing relationships between multiple interfaces and the external network. The synchronization clock interface provides a redundant set of clock buses that allow boards to exchange the necessary synchronization.

The synchronization clock interface has two redundant sets of three buses. The three buses are:

- CLK1 — an 8 kHz signal defined with the same clock tolerances as those used in the H.110 bus in CompactPCI
- CLK2 — a 19.44 MHz signal with 50/50 duty cycle, which is 1/8 of the line rate of a SONET/SDH OC3 interface
- CLK3 — a user-defined clock signal

As each bus is redundant, each signal appears in both an A and B form, resulting in, for example, CLK1A and CLK1B.

The synchronization architecture for AdvancedTCA can support either a centralized pair of redundant master clocks or a distributed system in which each board is capable of driving clock signals onto the backplane buses and synchronizing its own facilities either to external inputs or to backplane signals.

REAR I/O

AdvancedTCA supports rear I/O via Rear Transition Modules (RTMs). However, unlike CompactPCI, signals between an AdvancedTCA front board and RTM do not go through the AdvancedTCA backplane. Instead connectors on the RTM plug directly into mating connectors on the associated front board. This overcomes the limitations of previous systems that could only support specific electrical signaling between front and rear boards. Since front and rear boards are specific to a particular application, the vendor of an AdvancedTCA product can select a product-specific connector to use between their front board and RTM. This can include optical, coaxial and other connectors appropriate for the application.

RTMs are optional. Typically they simplify servicing by putting I/O cable assemblies on the RTM and leaving all active components on the front board. This allows front boards to be serviced without disconnecting and reconnecting cable assemblies. In fact the AdvancedTCA specification allows for RTMs that are either printed circuit boards, wiring harnesses or other constructions that fit within the AdvancedTCA RTM envelope.

AdvancedTCA MEZZANINE MODULES

Many proprietary telecom systems can be configured for a range of electrical and optical I/O using mezzanine modules. An open systems market benefits from the availability of these special-purpose sub-modules. PMC modules are the most widely available mezzanine modules today and AdvancedTCA PMC carrier boards are available from multiple manufacturers in diverse configurations.

Emerging telecom requirements, however, are impacting mezzanine modules. AdvancedTCA enables higher speed interfaces, which can be useful on a mezzanine module as well. The deeper (280 mm) AdvancedTCA board allows longer mezzanine modules. Some newer chip technology uses more power and space than is available with PMC. Finally, there is a need for a smaller, hot-swappable module that can minimize the amount of equipment that goes out of service in a single failure.

Responding to these needs, a new PICMG work group is developing a Sub-Module for AdvancedTCA. Of course, it may take several years to evolve from specification development to a rich market with the diversity of the PMC module market. Meanwhile, PMC carrier boards for AdvancedTCA provide the same richness of function that is available in the VME and CompactPCI markets.

REGULATORY ISSUES

Systems that are installed in telecom central offices must meet a set of reliability, safety, emissions and environmental standards that are far more stringent than those seen in typical commercial computer rooms. These requirements vary by region, by country and by network operator. In a multi-vendor open environment, each component must be designed so that when a system is assembled, it will meet all relevant specifications and pass all applicable tests.

Assembly of such a system is ultimately the system integrator's responsibility. But because a global team with many years of experience in telecom systems developed the AdvancedTCA specification, the specification includes extensive design detail and regulatory reference information that can be used by individual vendors and system integrators to assist in meeting regulatory requirements.

APPLICATIONS

With its high-capacity, full-mesh backplane, AdvancedTCA is an ideal platform for next-generation multi-service switches that require high-capacity backplanes that support control traffic plus packet and cell data traffic with well-defined quality of service (QoS) guarantees. Among multi-service switch designers, there are camps that prefer Infiniband, StarFabric or plain Ethernet for the fabric, but there is general agreement on using Ethernet for control traffic and complete agreement on using the PICMG 3.0-mandated infrastructure for system management.

AdvancedTCA is also suited for next-generation DSLAMs. AdvancedTCA boards can support 48 subscriber lines per card using today's electronics. While individual line rates may be 10 Mbps or more, a dual star backplane configuration is more than adequate to support the data traffic. And, AdvancedTCA's test ports (on the Zone 1 power connector) are available for copper loop testing.

AdvancedTCA is also an ideal platform for carrier-grade media servers.[1] This application requires fairly sophisticated media stream packet processing and signal processing, but only modest inter-board data transport that is easily handled by a dual star Gigabit Ethernet configuration.

ADOPTION OF AdvancedTCA

In less than two years, AdvancedTCA has moved from early specification efforts to an early market with a range of announced and available products. Interoperability workshops are running nearly every other month — the fifth session takes place in September 2003. During the third and fourth interoperability workshops, multiple vendors demonstrated management infrastructure interoperability, thus validating one of the most complex and important portions of the AdvancedTCA specification. AdvancedTCA-compliant chassis and backplanes are already available from multiple vendors including Bustronic, CG Mupac, Elma Electronic, Kaperel/Ritel and Schroff. An even larger number of companies have announced and/or demonstrated AdvancedTCA products and discussed roadmaps with their customers, including Artesyn Technologies, Aurora Technologies, Diversified Technology, Force Computers, Intel, Kontron, Motorola Computer Group, NMS Communications, Radisys and Pigeon Point Systems.

More importantly, several major telecom equipment providers are already designing next-generation systems on AdvancedTCA. These are not publicly announced programs, but this activity has attracted the attention of chassis and board vendors. The longer term potential is enormous. Annual worldwide estimates of telecom equipment sales vary depending on what's included in the count, but most estimates run to many hundred million dollars (US) per year. VME and CompactPCI equipment represent a negligible fraction of this business. But with the AdvancedTCA specification driven by telecom vendors' inputs and major equipment manufacturers adopting AdvancedTCA, the new platform appears poised to overtake and far exceed both VME and CompactPCI. As stated recently by Joe Pavlat, president of PICMG, "One telecom supplier has told us they foresee a need for 20 million boards." The remaining question is how rapidly will this market grow? This should become clear over the next 12 months.

For more information on PICMG and the AdvancedTCA specification, visit http://www.picmg.org and http://www.advancedtca.org.

[1] Carrier Grade Media Server Architecture on AdvancedTCA by Charles (Chuck) Linton, CompactPCI Systems Magazine, vol. 7, No. 4, May–June 2003, pp 67–69.