

# An efficient interconnect for Advanced Mezzanine Cards

By *Mike Zeile*

One of the many benefits of the AdvancedTCA standards effort is the promise of more systems flexibility due to built-in support for Advanced Mezzanine Cards (AMCs). In this article Mike discusses an efficient interconnection of off-the-shelf silicon that avoids the requirement for added bridge chips on each AMC.

Also known as PICMG 3.0, the AdvancedTCA is one of the most ambitious efforts ever to define the physical, power, and management parameters of next-generation Central Office telecommunications equipment. The resulting specifications are designed to allow manufacturers to lower the cost of equipment through increased use of commercial off-the-shelf components. At the same time, AdvancedTCA builds in more functionality and flexibility to service the demand for new high-speed data oriented services and emerging carrier-deployed IT services.

One of the main goals is to build as much flexibility into system boards (called node cards) as possible. This goal encompasses defining cards with a large board area – 8U by 280mm deep – and a 200W power draw to facilitate dense, processor intensive designs.

Another significant component in board design flexibility is the emerging Advanced Mezzanine Card standard. Figure 1 illustrates that the AMC standard combined with switch interconnect increases board design flexibility. The flexible mezzanine cards can be added to support new physical layer interfaces or new functionality, with the switch enabling configurable any-to-any connectivity. While the AMC working group is still finalizing the details of the standard, it has defined the main components: a carrier card that plugs into the backplane and provides power, system management, interconnect, and other services for up to eight daughter cards. A system designer can design AMCs to support a variety of functions, including:

- Network termination
- Packet processing
- Memory and storage expansion
- Generic processing

One particularly crucial element of an AMC design is the interconnect architecture of the carrier card because interconnect plays a critical role in enabling the flexibility promised by the AdvancedTCA effort. Some industry pundits have predicted that the use of AMCs will lead to the development of the universal access device. Such a device would enable designers to simply replace the interface or add in memory, routing, or other capabilities, thus configuring the hardware to support a new service roll out. This level of intra-card hot swappability has an impact on the interconnect, which must offer high throughput and low latency, any-to-any connectivity, and the ability to manage modularity. In




Figure 1

addition the interconnect requires the ability to sense when a card is not seated, terminate the port, and reconfigure the data flow.

It's almost inevitable that carrier boards will be used to house mezzanine cards not conceived when the system was designed, and may require interaction among onboard processors or other mezzanine cards, including data flowback, which is best achieved with a switch-based interconnect. Switches provide very efficient and high-speed connectivity, with low latency and unlimited data flowback.

While switch interconnects mean a rethinking of interconnect architecture, they can be implemented in many cases without a change in the interconnect interface. The AMC specification mandates a flexible connector and pinout configuration that is compatible with the most popular interface standards, including XAUI, PCI Express, and SPI-4.2. Many of the most popular processors, packet engines, coprocessors, and MAC chips use these standards and can be easily integrated into a board design. In the short term, SPI-4.2 is the most widely supported interface standard in terms of NPUs and other related data path processing devices. However, SERDES-based interfaces such as XAUI and PCI Express are more pin efficient and are expected to be the long term preferred interface choice.

## An SPI-4.2 switch chip optimized for AdvancedTCA

One switch chip that is designed for this application is the PivotPoint FM1010 from Fulcrum Microsystems (see Figure 2). The FM1010 is a six-interface SPI-4.2 switch chip that provides a high-speed interconnect between the devices in the system (and housed on AMCs), such as NPUs, security processors, MAC and PHY devices, custom ASICs and FPGAs, and other related coprocessors.




Figure 2

The device is well suited to the needs of AMC designs because it offers a very efficient interconnection of off-the-shelf silicon that eliminates the need for added bridge chips on each AMC. The FM1010 also manages reconfiguration of the data flow and port termination when AMCs are removed from the system.

At the center of PivotPoint is Nexus, an asynchronous crossbar that provides 192 Gbps of data capacity. The throughput capacity of Nexus is three times as high as the aggregate of the SPI-4.2 interfaces operating at their full 14.4 Gbps data rates, ensuring that full line rate can be maintained to every component on the board.

### Importance of latency

Switch performance is one element in lowering latency, but flow control and port channelization also play important roles. Low latency is increasingly important because a switch interconnect is often inserted into the middle of the data path between two packet-processing devices, and it must look as transparent as possible to the application.

The FM1010 features a tightly coupled flow control system that boosts system efficiency by lowering the latency of the flow control path. Flow control is built into the data path circuitry of PivotPoint. In addition, PivotPoint contains three separate flow control domains:

- Ingress
- Crossbar
- Egress

As Figure 3 illustrates, the ingress and egress flow control domains maintain tight coupling with their respective external link partner. When congestion occurs at the egress, flow control is propagated back through the crossbar and then to the ingress port (if necessary) until congestion subsides.

The FM1010 supports as many as 16 channels per interface. These individually flow-controlled channels, sometimes called ports, may be used to transport separated traffic flows. They have distinct hardware resources, buffers, and flow control mechanisms but share a common physical interface. The ability to map any channel on any interface to any other channel on any other interface is also important to allow complex data flows through multiple devices.

A key issue is Head Of Line (HOL) blocking where a packet is blocked at the ingress port by congestion in the switch and keeps other packets, whose data paths are not affected by the congestion, from progressing. HOL blocking will affect the actual throughput of the switch regardless of the capacity of the switch. Port channelization and per-channel flow control are important features within a chip to overcome HOL blocking.

### Design example

A reference design for an AdvancedTCA line card demonstrates flexibility. The design uses mezzanine cards to offer removable PHY/MAC interfaces for changeable network support. These mezzanine cards can support any physical media or networking technology, such as 10/100/1000 Ethernet or WAN protocols.




Figure 3

The FM1010 sits in the middle of the board connecting an NPU that provides packet classification and services processing, a traffic manager for scheduling and QoS enforcement, a security coprocessor for encryption, and a fabric interface chip for linking to the switch backplane. It provides a high-speed interconnect to all components by mapping the ports to the appropriate interface, enabling efficient, cost-effective packet delivery.

The SPI-4.2 throughput capability can scale from AMCs with low-speed WAN interfaces to those supporting optical 10 Gigabit Ethernet.

The switch offers full connectivity to each processor that can be reconfigured depending upon the application. A configuration that features two LAN-interface mezzanine cards may utilize the traffic manager more heavily than two WAN-interface cards, which will need more access to the security coprocessor.

### Conclusion

This design is emblematic of how a switched interconnect can enable the full flexibility benefits of AMCs in an AdvancedTCA design. PivotPoint serves as one example of how to deliver this flexibility along with the low latency and high performance needed for multigigabit speed designs.

*Mike Zeile is VP, Marketing at Fulcrum Microsystems. Prior to joining Fulcrum, Mike co-founded and managed all aspects of marketing and business development as VP of Business Development for Encryptix, an Internet security infrastructure company.*

For further information, contact Mike at:

**Fulcrum Microsystems**  
 26775 Malibu Hills Road  
 Suite 200  
 Calabasas Hills, CA 91301  
 Tel: 818-871-8164  
 Fax: 818-871-8101  
 E-mail: mzeile@fulcrummicro.com  
 Website: www.fulcrummicro.com