Multi-Gigabit AdvancedTCA Design Architectures

By: Ed Sayre, NESA (www.nesa.com) and Kathy Breda, F9 Systems (www.f9-systems.com)

The AdvancedTCA (ATCA) specification allows for a wide variety of physical and logical architectures, including single switch star, dual stars, single and dual meshes, and other possible configurations. The designer must implement one based on the system's reliability goals and performance requirements. Typical tradeoffs include logic design complexity, chip set choices, power and thermal performance, materials cost, layout cost and complexity, and ability to provide redundancy and fault tolerance. This article will provide an overview of the architecture choices with their benefits and draw-backs.

ATCA based systems are standardized, multi-gigabit platforms designed specifically for telecom infrastructures. They permit the deployment of "best-in-class" telecom systems, all with common system control interfaces and platforms. However, within the common interfaces and platforms, there are many choices to make. One of the first to consider by the engineering design team is the system architecture.

Architecture Choice

The system architecture is one of the most important, and perhaps the most difficult and sophisticated design choice to make during a project. This is because many product details need to be understood early in the development process to make a considered choice. The types of details the engineering team will discuss as part of the architecture process are the performance and throughput, protocols power and cost. The system protocol choice is driven by the market, a company's existing product lines and future plans. Even though ATCA is protocol agnostic, designers still need to make a protocol choice for their system.

One significant benefit of ATCA is that engineers don't have to design a platform that specifically matches their system protocol resulting in time and project cost savings. Once a protocol is decided upon, a semiconductor chip set is chosen that will support the protocol requirement, along with meeting performance, power and connectivity necessities. This in turn drives the type of architecture that is chosen by the design team. Although this is a simplified example of choosing system architecture, it provides an idea of the type of details that need to be considered during the design process.

The ATCA system interconnect architecture specification has been written in terms of logical slot connections. The most popular configurations are the centralized 14-slot switch and 5-slot full mesh architectures, although the logical definition of the architecture allows designers the choice of where to locate their switch hubs; that is, the designer is free to map the logical architecture into the physical architecture of the chassis and backplane that best suits their reliability and performance requirements. Different combinations of logical and physical configurations can lead to dramatically different physical architectures such as Non-Centralized Single or Redundant Switches and Nodes.

Centralized Switch Architecture

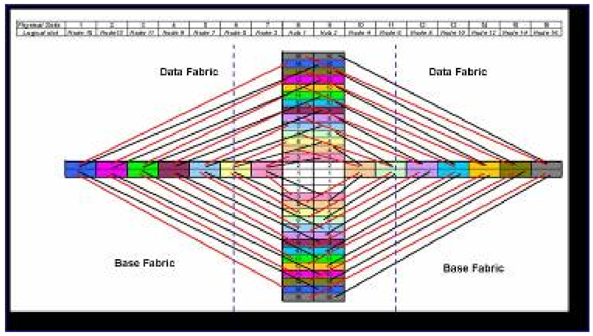

Centralized Switch Architecture is the most common type used in ATCA designs because it is illustrated most clearly in the PICMG 3.0 ATCA specification and is the simplest and probably least expensive to layout and route. It has the added cost and performance benefit of being able to utilize standard FR-4 material (with care) for speeds up to 5 Gbps due to short, direct trace paths. There are many other factors that also contribute (reduce) overall system performance, nonetheless, the centralized switch is clearly the most popular physical architecture.

The greatest drawback of centralized switch architectures is that the two switches constitute a single point of thermal failure. Typically, the switch cards are the highest power boards in the backplane and are located in the two center slots under a single column of fans, creating a single zone of thermal failure. Even though six to nine 5 ¼" fans should cover a 14-slot area, the centralized switch architecture makes the design of an N+1 redundant fan cooling design almost impossible.

When deciding to use this architecture, a typical design trade-off will be made assessing failure risk and cost against performance and time to market requirements. Since Centralized Switch Architecture is such a popular choice, it appears that many designers do not consider the single point of thermal failure to be an overriding design concern. An additional consideration is that most of the platform companies sell this type of backplane off the shelf, since the other design architecture alternatives are for more specialized solutions that require advanced engineering skill to implement.

Isolated Dual Star Switch Architectures

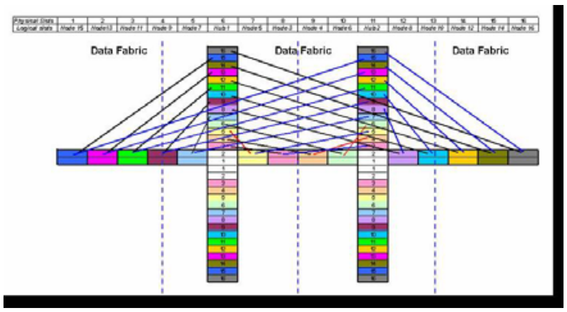

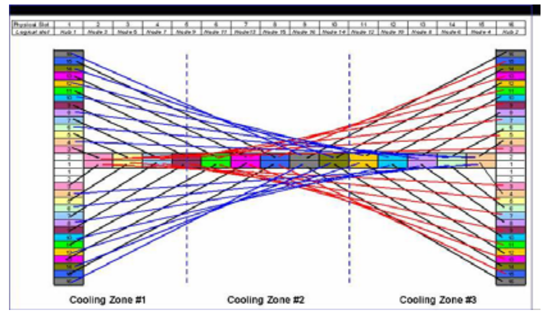

The Isolated Dual Star Switch Architectures are generally implemented as custom ATCA backplane designs. In such systems, the up-time performance and the 5-9's reliability are the overriding requirements. The absence of a single point of thermal failure is the major difference between the standard switch and custom ATCA architectures.

Isolated Dual Star Switch Architectures can have switches located at each end of the backplane (most common) or at intermediate slots. The latter architecture has medium trace lengths with crossovers, making it more complicated to layout and route than a Centralized Switch but there is no single point of thermal failure. However, like the Dual Centralized Switches, this architecture is typically capable of performance up to 5 Gbps over FR-4 PWB material.

When the switch cards are placed in the opposing backplane end slots, a high fault-tolerant solution with no single point of thermal failures for the highest power boards exists. It has balanced, independent thermal zones, and offers a low potential for flow choke.

However, the end slot Dual Switch Architecture has complex Signal Integrity characteristics. It has the longest overall trace lengths, so it is complex to layout and route and requires higher layer count due to crossovers. Careful signal integrity analysis and planning is necessary to properly implement this architecture. Typically, this design approach is chosen by the design team due to the need for a high-reliability physical architecture for applications requiring 5-9's or higher availability. Use of FR-4 means that this architecture is usable to 3.125 Gbps. Higher performance backplane interconnects would require lower loss material and/or switch and node card I/O's with pre-emphasis and/or equalization/decision based receivers.

Conclusion

Choices are good, but trade-offs are never easy. The ATCA standard provides designers with flexibility and a variety of design choices, but at the end of the day, the design decision is up to you. We recommend that you know what your decision parameters are for choosing your system architecture and verify them with your sales and marketing team to ensure they meet the customer and market needs. Then, you must make sure that you understand the impact of the design trade-offs in your product.

The system architecture decision will cascade into many other product features, such as semiconductors, performance, redundancy, applicable target markets, price, etc. For example, think about whether you need some redundancy, or 5-9's redundancy; whether 3.125 Gbps backplane performance will be enough in two years; if you engineering team is up to the task and has the tools for a detailed multi-gigabit design; if you want to use an off the shelf backplane design so that you can focus on your core capabilities such as logic or software development. There are lots of choices, but a well thought out analysis will provide a robust, ATCA system that meets your company's and your customers' needs.

Ed Sayre is the CTO of North East Systems Associates and Kathy Breda is the President of F9 Systems. You can reach them at esayre@nesa.com and kbreda@f9-systems.com, respectively.