Practical Design Considerations — Ethernet versus RapidIO

Before committing to Ethernet in board- and chassis-level applications, engineers will find it useful to first compare Ethernet to alternative technologies such as the RapidIO specification.

Greg Shippen, Freescale Semiconductor

Engineers can be thought of as negotiators between the ideal world and the real world. It is their task to apply the theoretical possibilities of technology in a practical manner that balances performance, reliability, cost, and a number of other key considerations, depending upon the application. Many times, engineers are so good at what they do that they are able to carry technology over to applications for which it was never intended. In some cases, the fit is "enough", and the resulting economies of scale are sufficient to make the extended implementation of the technology successful. In other cases, however, the desire to make use of a technology one is already familiar leads engineers to "shoe-horn" it into new applications which stretch the technology beyond its capabilities, unnecessarily complicating design by creating more problems that need to be solved.

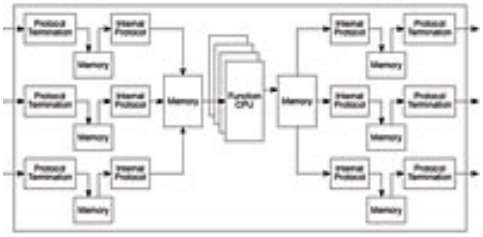

[1]

Expanding the scope of a technology can yield significant system-wide savings. For example, consider the consolidation of the many interconnect levels of a system fabric (see Figure 1). Rather than having to support a new protocol at each level, substantial cost savings as well as simplified system development and management can be realized by collapsing levels into each other through the use a single protocol, thus reducing the number of protocols a system must support.

Given its ubiquitous presence throughout the LAN and WAN, Ethernet is a prime candidate as the convergence protocol for system-level fabrics. Its high volumes and extreme flexibility make it an especially attractive option for many applications, and for these reasons engineers have begun to carry Ethernet from the LAN and WAN down into the chassis- and board-level infrastructure. Before committing to Ethernet in board- and chassis-level applications, however, engineers will find it useful to first compare Ethernet to alternative technologies such as the RapidIO specification that, rather than extending downward as Ethernet does, it pushes up

through the fabric.

Origins

Ethernet was designed to be extremely flexible and to serve in large networks with a great many endpoints, each of which was expected to have a processor available for protocol stack processing. Its applicability becomes increasingly strained as application requirements diverge from Ethernet's base capabilities (see Figure 2). Specifically, while the communications infrastructure of high-speed, embedded board-level and chassis-level applications resemble small networks unto themselves, data transfers at these levels differ in significant ways from the LAN.

RapidIO® technology was designed specifically for embedded in-the-box and chassis control plane applications. The RapidIO protocol emphasizes reliability with minimal latency, limited software dependence, protocol extensibility, and simplified switching providing effective data rates from 667 Mb/s to 30 Gb/s.

Standards and Interoperability

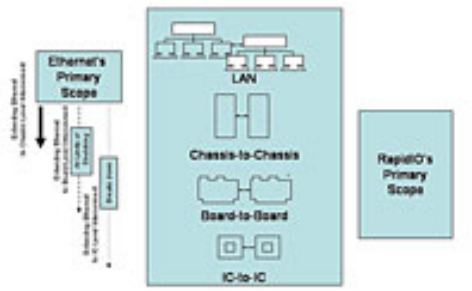

[2]

A key consideration in determining the viability of a technology for use in chassis- and board-level applications is its status as an industry standard. Proprietary standards tie manufacturers to a single source for parts and support, which can leave manufacturers facing a technological dead end.

Even industry standards, however, can run into problems when they are not administered by the same standards body. Ethernet, for example, is supported by a wide range of Layer 3 and higher protocols that implement optional and advanced functionality. While this flexibility expands Ethernet scope in the network, no single unified specification is available that can be uniformly implemented. This results in significantly increased stack complexity, higher cost for processing resources and a higher packet latency.

The RapidIO specification, on the other hand, has more consistent protocol layering and is completely governed by a single standards organization, the RapidIO Trade Association. Rather than providing near unlimited flexibility with corresponding added complexity, the RapidIO specification outlines a more appropriate set of baseline functionality for chassis- and board-level applications than Ethernet. The result is lower implementation cost with reduced overall complexity. And since most of the RapidIO protocol is implemented in hardware, its software drivers are far

simpler than the typical Ethernet TCP/IP stack. Additionally, RapidIO stacks can depend upon the existence and consistency of standardized implementations.

Throughput and Latency

For LAN and WAN applications, Ethernet’s high overhead can be compensated for by large packets. Limited flow control makes switches less complex throughout the network, and the ability to drop packets is a key element for efficient congestion management. These characteristics can be devastating in control plane applications at the chassis and board level. Control plane transactions tend to be limited in size, so high overhead &#151 the TCP/IP header alone adds 40 bytes &#151 directly results in reduced efficiency.

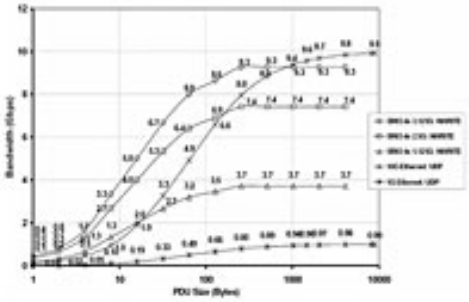

[3]

The RapidIO specification has optimized header size to maximize efficiency for the size of packets typically used in control plane applications. Additionally, the RapidIO protocol provides robust flow control and guaranteed delivery &#151 both essential for maintaining the priority and reliability of control plane transactions &#151 providing higher fabric usage in complex topologies well in excess of 50%. Furthermore, link-level error correction minimizes latency jitter incurred as the result of soft errors. End-to-end latency is substantially lower as well, potentially much less than 500 ns, depending upon the particular implementation, since software stacks at endpoints can be minimized or even eliminated.

Ethernet can be implemented in a way that supports control plane applications that cannot tolerate packet loss, but the fabric must be significantly over-provisioned to achieve this. While over-provisioning reduces end-to-end latency and latency jitter, it decimates throughput: at 25% usage, the sustainable effective throughput of Layer 2 traffic for Gigabit Ethernet is ~250 Mb/s and only 2.5 Gb/s for 10 Gigabit Ethernet, depending upon average packet size (see Figure 3). Even with over-provisioning, however, end-to-end latency can still be in the milliseconds.

Power Considerations

From a PHY perspective, both Ethernet and RapidIO interconnects can take advantage of the power efficiencies of a single-lane XAUI-like PHY interface which dissipates anywhere from 70 to 200 mW at 3.125 Gbaud. For Ethernet applications using 1000Base-T PHYs, power dissipation rises to between 640 to 950 mW.

In applications where power consumption must be managed carefully, Ethernet protocol processing can consume more power than is desirable compared to a

RapidIO-based endpoint. This is in large part due to the fact that most Ethernet implementations run a software stack on a high-frequency host processor.

Economics

In general, the economies of scale for any silicon-based device depends upon die size, volumes shipped, and the competitive level of the technology's ecosystem. At first glance, it might appear that Ethernet's economies of scale dominate those of RapidIO devices on every count. In fact, the opposite is true.

RapidIO protocol processing is implemented in hardware. However, one RapidIO endpoint (including messaging support) shipping today is only 25% larger than a Gigabit Ethernet controller on the same processor. Note that this example is slighted in favor of the Ethernet controller since it did not implement a full TCP/IP off-load engine (TOE). With a complete TOE, the Ethernet controller would be at least comparable to or larger in size than the RapidIO controller. When considering PHY cost/size, both Ethernet and RapidIO systems can utilize a XAUI-like PHY, suggesting comparable silicon complexity and areas between the two standards.

From a volume perspective, it would be erroneous to take into account the entire Ethernet market when comparing volumes to RapidIO applications. Ethernet covers a wide range of applications and industries which are similar to high-speed, embedded chassis- and board-level applications but require significantly different hardware and software implementations. Additionally, Ethernet's best price economies today are for 4 to 8 port switches, not the highly aggregated configurations required for efficient backplane configurations.

Finally, one must carefully consider the existing silicon and support ecosystem to understand its true impact on price. For an application such as a 12 to 24 port ATCA-like backplane, an Ethernet controller must support both VLAN QoS and SERDES PHYs. Such devices, however, are offered only by a relatively limited number of vendors in the Ethernet ecosystem. A comparable RapidIO device is offered by more vendors, given its standardized nature, making for a much more competitive marketplace.

Effective Throughput

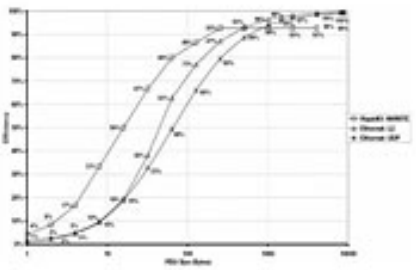

[4]

Given that the cost is potentially comparable, developers need to turn to other considerations such as throughput to determine which technology is best-suited. When all of the factors listed above are taken into account, RapidIO technology

offers 2.5 times more bandwidth per link than Gigabit Ethernet.

The difference is even more pronounced when fabric bandwidth requirements exceed 1 Gb/s. In these cases, the only option for Ethernet is 10 Gb/s which, if an application doesn't need this much throughput, introduces undesirable over-provisioning cost without any corresponding performance or reliability benefits. In addition, today RapidIO devices offer higher effective bandwidth for payloads less than 1024 bytes at a lower cost (see Figure 4). Note that this comparison overlooks the added cost and complexity associated with managing Ethernet stack processing at 10 Gb/s.

Ethernet is clearly the technology of choice for many LAN and WAN applications. While the ability to consolidate interconnect levels by bringing Ethernet down to the chassis and board level, the complexities arising from its derivatives, associated inefficiencies, lower throughput, and smaller ecosystem make Ethernet less appealing for use as a backplane or control plane technology.

The RapidIO protocol, designed from its inception to serve in chassis- and board-level applications, offers the efficiency of a hardware-based protocol processing architecture, superior throughput without extreme over-provisioning, lower overhead, more reliable flow control, more efficient power usage, and a well-established ecosystem. As a result, RapidIO-based implementations can consolidate more levels of the system-level interconnect more efficiently and more reliability than Ethernet, making RapidIO the technology of choice for next-generation high-speed, embedded applications.

About the Author

Greg Shippen is a system architect for Freescale Semiconductor's Digital Systems Division, NCSG, Member, RapidIO Trade Association Technical Working Group and Steering Committee, www.RapidIO.org [5]. Email: Greg.Shippen@Freescale.com [6].

Links:
[1] http://www.wirelessdesignmag.com/sites/wirelessdesignmag.com/files/legacyimages/0711/figure1_600.jpg
[2] http://www.wirelessdesignmag.com/sites/wirelessdesignmag.com/files/legacyimages/0711/figure2_600.jpg
[3] http://www.wirelessdesignmag.com/sites/wirelessdesignmag.com/files/legacyimages/0711/figure3_600.jpg
[4] http://www.wirelessdesignmag.com/sites/wirelessdesignmag.com/files/legacyimages/0711/figure4_600.jpg
[5] http://www.RapidIO.org

[6] mailto:Greg.Shippen@Freescale.com