Medical Imaging Systems: Assuring High Availability Performance With StarFabric

Steve Christo, StarGen and Justin Moll, Bustronic

Introduction

Medical imaging systems are requiring higher bandwidth. With an emphasis on real-time performance and higher resolutions, the amount of data processing may soon reach staggering levels. Using input sensors to receive large amounts of information and Digital Signal Processors to turn analog input into digital data, the systems are leaning towards backplane-based chassis with high performance. Traditionally, many of these systems use standards-based architectures like CompactPCI bus or VMEbus. With backplane performance limitations for standard architectures, medical industry system designers are looking to new switched fabric technologies that offer high performance and high reliability at a reasonable cost.

Medical imaging systems such as MRI, CAT Scan, PET Scan, etc are the most in need of higher performance. The intra-chassis traffic between the processor boards, communications boards, display, device/sensor, and storage is vast. For real-time information and clearer, higher resolution images, the data needs to travel at high speeds without errors. Not only is the data rate important, but the processing must be reliable. With redundancy and inherent reliability, the system can be highly available, with up to 99.999% uptime.

Switched fabrics

A move to a switched fabric based technology is the natural evolution for medical imaging systems. Switched fabric technology inherently provides some key benefits such as design flexibility, physical and bandwidth scalability, reliability, and quality of service. In bus technologies like PCI, only one device has access to the bus at a time and each device must request the bus and an arbitration algorithm is used to grant requests. The duplex nature of a point-to-point switched interconnect allows each device to transmit and receive simultaneously. Through the building of systems with series of end-points and switches a diverse and a flexible array of system topologies can be created. As more connections are added to the system the total bandwidth of the system increases. This networked approach as opposed to bus architecture also provides significant flexibility in the system design topologies that can be created.

With a serial switched architecture, both physical and bandwidth scaling are realizable. Today's buses are typically limited to 12-20 inches on a backplane and 1 to 2 feet over expensive and unreliable ribbon cables. Through the use of serial physical layer technology, these distance limitations can be eliminated. In some medical equipment such as ultrasound machines, the storage and display systems are typically separate from the analog input and image processing arrays. With a serial interconnect, developers would not be limited to having these segments of the system adjacent to each other. In order to scale the bandwidth of a bus-based system, either the bus is widened or the frequency is increased. Both methods have their drawbacks. Increasing the width will require additional routing layers on backplanes and boards, subsequently increasing the total cost of the system. Increasing the frequency on synchronous buses tightens skew and routing requirements, limiting the number of devices supported and increasing reliability concerns. By implementing a serial switched interconnect, increasing the overall bandwidth of a system is as simple as adding another endpoint to the topology.

In addition to scalability, switched interconnects provide a dramatic increase in reliability compared to bus-based designs. In contrast to buses, in which an errant end point can bring down the entire bus, the point-to-point nature of switched interconnects isolates faults to a single end point. The fault at an endpoint could be handled in several different ways. The most basic response would be to remove the faulty FRU in the system and hot swap it with a replacement. Point-to-point connections are inherently friendly to device insertion and removal. So, if there is a problem with the system, MTTR (Mean Time to Repair) is minimal. However, the inaccessibility of a technician could result in equipment downtimes of a few hours resulting in costly rescheduling of patients. Another method would take advantage of the redundancy provided by switched interconnects. For example, a CAT scan system might have several CPUs running different algorithms for processing images. If one of these CPU blades fails in a system, the traffic targeted at that blade could be automatically rerouted to another CPU card using the redundant routing capability of switched interconnects. If the redundant CPU is a standby processor, the system will continue to operate with no performance degradation after a downtime measured in seconds instead of hours. If the redundant CPU is actively running another algorithm, it could expand its processes and run a second algorithm. In this case, the system would run at a reduced performance level until a replacement CPU was installed.

Quality of service (QoS) is another benefit that switched interconnects have over bus-based systems. In a bus-based system, the best you can do to prioritize a specific data stream is to develop an arbitration algorithm that provides preferential treatment to the master generating the high priority data. Using a switched interconnect like StarFabric, specific data streams can be tagged with a higher priority class of service. In platforms where different types of data converge, this capability greatly enhances system performance.

StarFabric

StarFabric is a serial switched interconnect with a physical layer consisting of low cost, high performance 622Mbps low voltage differential signals (LVDS). Four 622Mbps differential pairs form a link with an aggregate bandwidth of 2.5Gb/s in each direction. Each StarFabric link implements both 8B/10B and CRC encoding to protect against bit errors in the data. If either an 8B/10B or CRC error is detected at the receiver, the data will automatically be retransmitted by the hardware. With the 8B/10B encoding, the actual data bandwidth supported by a StarFabric link is 2.0Gb/s or 250MB/s. The serial links are hot plug capable and provide point-to-point connections between chips on a card, across a backplane, and between racks in a room. Chassis-to-chassis interconnects are possible with standard RJ45 connectors and CAT5 cable, eliminating any requirement for expensive high-speed connectors or cables.

StarFabric enables three types of routing methods. Address routing provides 100% compatibility with PCI drivers, BIOS, configuration, etc. Path routing provides advanced features such as QoS, which enables data, voice, and video to be carried over a single interconnect. A typical system would use the asynchronous QoS for data traffic and isochronous for video. Multicast allows single initiators to broadcast messages to multiple destinations. StarFabric also has several high availability features including error detection, correction, notification and isolation. Redundant routing is supported with automatic fail-over to an alternate path when a connection is removed.

PCI and particularly CompactPCI are often used in MRI Systems. The CompactPCI bus provides a rugged, hot-swappable platform for high-reliability medical systems. But, CompactPCI's bandwidth is running out of steam. With speeds up to 528 MB/sec over 66 MHz traffic, the bus is limited to 5 slots. For 33 MHz, the backplane can extend to 8 slots (without bridging). With the many processors, A/D and D/A, data acquisition, and image processing cards, MRI, CAT scan, and PET scan systems require higher slot counts along with the higher bandwidths. The PICMG 2.17 specification based on StarFabric extends the performance of CompactPCI.

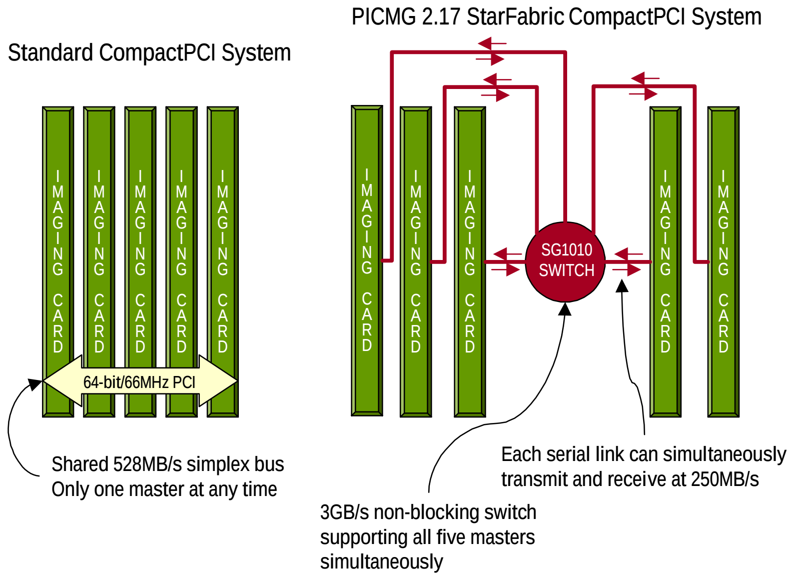

Diagram 1: CompactPCI and StarFabric Imaging Systems

Consider the two systems shown in Diagram #1. The CompactPCI system uses a 64-bit, 66MHz bus connecting five image processing cards. In this system, the total bandwidth shared between all five cards is 528 MB/s. In the PICMG 2.17 system, each card could be transmitting and receiving at the full bandwidth supported by a serial link (250MB/s) and thereby saturating its local 64-bit, 66MHz PCI bus (~500MB/s) with the two-way traffic. Therefore assuming a fully distributed system, the total system bandwidth increases to 2.5GB/s or approximately 5 times more than the bus-based architecture. (see Diagram #2)

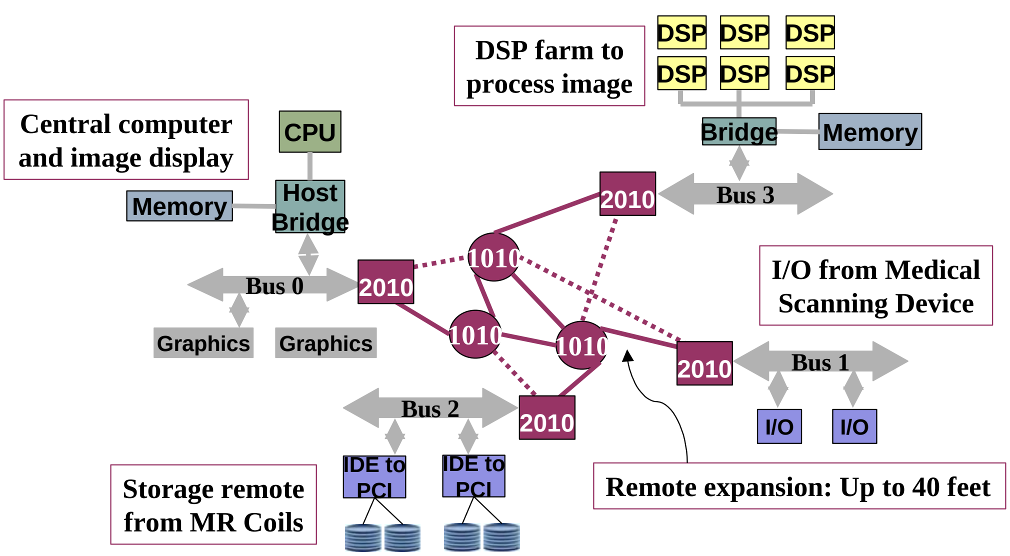

Diagram 2: StarFabric serial interconnect between CPU, I/O, DSP, and storage in a medical imaging system

In a PICMG 2.17 platform, StarFabric traffic is run across P3 (P4 is reserved for H.110 and much of P5 allows user I/O). It is also focused on providing this functionality in an easily adoptable way by not requiring exotic system design in terms of power or signal integrity, and by allowing use of standard cabling and connector technology. Its cost structure is in line with traditional bridging technology. PICMG 2.17 products are now readily available with over ten companies offering PICMG 2.17 compliant boards or backplanes. PICMG 2.17 compliant StarFabric backplanes are available in various slot sizes and configurations. They can be developed in relatively low layer counts, such as a 12-layer controlled impedance stripline design. The CompactPCI busing on P1 and P2 and the H.110 bus on P4 can be routed in 8 layers. Some of the StarFabric links can be routed on the same 8 layers and the remainder on the other four layers. The differential pairs are routed as close together as possible and kept on the same layer. The outside layers are ground for EMI protection and suppression. The signal layers are alternated with power or ground layers for controlled impedance and to minimize crosstalk. Vias are not used on signal traces because they disrupt the impedance of the trace. High and low frequency decoupling capacitors are distributed generously across the backplane. Power and ground planes typically use 2 oz. copper for superior power distribution.

The PICMG 2.17 specification defines two types of topologies, centralized and distributed, for integrating StarFabric on a CompactPCI backplane. A centralized topology utilizes a dedicated fabric board where all of the node board traffic is aggregated. Each line connecting a node board to a fabric board consists of a StarFabric link with 2.0Gb/s bandwidth in each direction. By adding another fabric board with connections to a second link in each node board, a fully redundant system is possible. The redundant link can be used to increase bandwidth to each node board and/or to increase the reliability of the system. Although a centralized topology provides an effective and scalable solution for a heavily loaded chassis, it may not be necessary in systems with a limited amount of slots. By implementing a distributed topology, slots, which only perform switching, dedicated fabric slots, are no longer required. Instead, the switching functionality is incorporated into each line card. This topology is an efficient implementation for small-scale systems and provides a path for incremental growth as system requirements increase. For either topology, the CompactPCI form factor provides many high availability features such as hot insertion/extraction and IPMI system management. The serial links of StarFabric provide a more robust physical layer than CompactPCI for hot insertion/extraction of boards. In addition, in-band management through StarFabric complements the capabilities of IPMI system management.

Unlike other switched fabric technologies, StarFabric supports several traffic classes including asynchronous classes, isochronous classes, multicast, and high-priority. Asynchronous traffic is traditional data oriented traffic, with large bandwidth requirements but without real-time delivery requirements. Control and signaling traffic are typically asynchronous. Isochronous traffic, including voice and video, requires deterministic real-time delivery. Through use of these traffic classes, StarFabric's technology is ideally suited for applications like MRI and other medical imaging systems where data and video are converging.

Application Example

In many medical applications, latency and throughput are important requirements. In an ultrasound system, latency is critical because the technician uses the real time image to properly position the transducer. Latency is less critical in CAT scan equipment although still an important factor. Both systems require the performance to scale as processing elements are added to the system. If all of the processing boards are contending for a common bus, performance could actually degrade when the number of processing nodes is increased. StarFabric solves these concerns by providing both physical and bandwidth scaling in a low latency interconnect. Consider the ultrasound system shown in diagram #3.

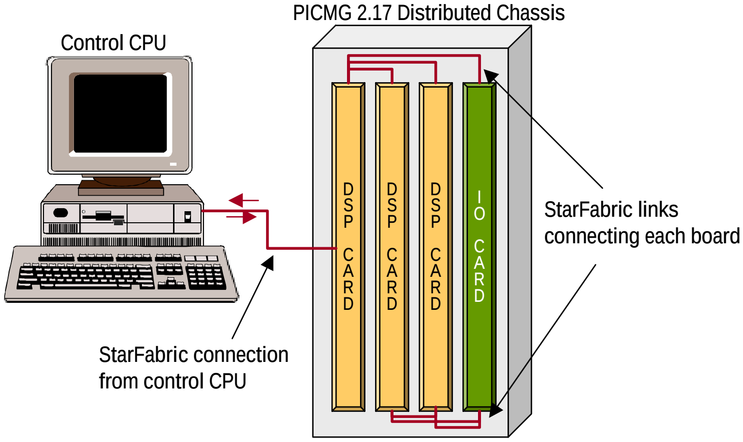

Diagram 3: StarFabric enabled UltraSound System

In today's systems, the control CPU is typically connected to the IO and image processing boards with a expensive custom cable that carries all the PCI bus signals. Due to skew and signal integrity issues, these cables become unreliable if they exceed 1 or 2 feet in length. By replacing the cable with StarFabric serial links, the physical limitations are eliminated as StarFabric can scale up to 10 meters using standard CAT5 cable. In addition, StarFabric provides backward compatibility to legacy PCI while maintaining a very low latency (<1us). For this simple PCI expansion case, StarFabric allows developers to reduce their system cost and improve system reliability without changing the current software.

Inside the PICMG 2.17 CompactPCI chassis, the boards are interconnected by StarFabric links using a distributed or mesh architecture. As image processing requirements increase, DSP boards can be added to the system and the overall bandwidth will scale accordingly. In addition, the multicast capability of StarFabric enables the IO card to broadcast data to each of the DSP cards, which could then perform distinct operations. StarFabric executes these data transfers while adding less than a few microseconds to system latency. If another interconnect such as Ethernet were used, the additional latency would be orders of magnitude greater.

Compatible hardware – chassis considerations

In medical systems, the chassis design is an important consideration. Issues such as EMC/EMI, ruggedness, and a flexible design are key issues for achieving high-performance at a reasonable cost. PICMG 2.17-compliant chassis have been developed in a 19-inch rackmount EMC version, ideally suited for medical applications. A good example is Elma's PICMG 2.17 chassis that has accommodations for up to a 21-slot StarFabric backplane with rear I/O capability. It utilizes a push-pull airflow technique using three individually removable plug-in fan trays below the cards, with 90 cfm tube-axial intake fans and dual radial blowers above the cards for exhaust. The intake air is filtered using Bellcore-compliant foam air filters that are easily removable. As the StarFabric architecture will be a migration path from traditional CompactPCI-based bus structure to a switched-fabric backplane architecture, existing packaging concepts can be utilized. This is a significant advantage, as similar packaging solutions in areas like EMI, cooling, and shock and vibration can be incorporated. For example, shielding with EMC gaskets and BeCu "fingers" or contact strips, will help the chassis maintain electrical continuity between mating metallic surfaces (like panels, covers, etc). Conducted emissions can be addressed by incorporating high performance EMI line filters, allowing the chassis to meet FCC/CISPR requirements. Also, using the IEEE-1011.10 specifications makes EMI containment on the front panel/cards easier to achieve.

The cooling requirements for the medical systems can be adequately met with advanced airflow techniques such as the slot air baffle, air plenums, etc. Compact radial blowers or backwards-curved impellers have proved very effective in dissipating heat buildup under high static pressure. Employing fan monitoring alarms is highly encouraged due to the mission-critical nature of medical applications. Some versions of PICMG 2.17 chassis may include positive pressure cooling (fans blowing on the cards), evacuative cooling (exhaust fans), or a combination of both. Incorporating critical features like hot swap and redundant fan trays will play a major role in meeting high availability requirements. Therefore, the chassis accommodate pluggable fan trays, power supplies and system management and monitoring. The continued use of 3.3V, 5V, and 12V makes it possible to find power supply solutions from the multitude of proven vendors in the industry. The PICMG 2.11 Power Supply Interface specification defined a 47-pin Positronic connector for pluggable supplies. Power supplies are load sharing by using internal OR-ing diodes. Power supplies monitor the health of the voltages on the backplane through either third wire or droop current sharing methods. Current droop regulates to within 10-15% and third wire regulates to 5%. The power supplies and fans have indicators in case of failure for ease of diagnosis and replacement. PICMG 2.17-compliant chassis are becoming more widely available in the marketplace. As of this writing, there are various StarFabric backplane configurations in 4, 6, 8, 10, 17, and 21-slot sizes.

Conclusion

Bus-based architectures are running out of bandwidth for today's medical imaging solutions. The migration to switched fabric systems is the natural evolution path. StarFabric is a switched fabric architecture that already has all of the necessary hardware and software components available in the market. It is also compatible to PCI-based systems, so much of the investment in the system can be preserved. An expanding range of PICMG 2.17-based backplane, chassis, and board designs makes it easier to fit most size, space, and configuration requirements.

For more information on StarFabric, visit www.stargen.com. Further details on StarFabric backplanes and chassis can be found at www.nextgenbackplanes.com

Justin Moll
Marketing Manager
Bustronic, An Elma Company
PR Manager
Elma Electronic Inc.
Tel: (510) 490-7388
e-mail: Jmoll@bustronic.com
www.bustronic.com/
www.elma.com/

Stephen Christo
Product Marketing Manager
StarGen Inc.
225 Cedar Hill Street
Marlborough, MA 01752
Tel: (508) 786-9950
e-mail: christo@stargen.com
www.stargen.com/