

One Supplier's Directions

SKY Computers is participating in the development of all of these specifications. SKY is committed to working with the RapidIO Trade Association to address the problems that we recognize in the specification, and make RapidIO a success.

We are also excited by the capabilities of InfiniBand. We believe that InfiniBand will play an important part in our future product generations. It is critical to make InfiniBand work together with RapidIO for the class of systems that we develop for our customers.

Serial FPDP will be enhanced, and bridges to RapidIO and InfiniBand developed to enable its use in these future systems.

Finally, any computer is little more than an expensive room heater unless there is good software provided to support it. As the hardware becomes more complex and capable, it is increasingly important to provide the software tools to use, diagnose, and repair it. Our customers expend more time and thus more money on software than hardware, and we must provide the software to make them successful. Both InfiniBand and RapidIO support "discovery", where the actual configuration isn't known until the system is started, and the configuration can actually change during operation. Similarly, both specifications require complex error detection and recovery. Error diagnosis and fault isolation on a large system is a complex problem. These issues all add up to significant software support requirements.

where the "corners" of the matrix are effectively flipped. The result is that each processor has a stripe of the data in the azimuth direction to be processed.

4. The processors can now do a portion of the math. The next stage of processing involves going in the range direction. A corner turn is performed, resulting in each processor having a stripe of data in the range direction.
5. In order to complete the image formation, a final stage of processing must be performed in the azimuth direction, requiring yet another corner turn before the processing can occur.

Using today's processors, the corner turn performance is actually the limiting factor in over-all system performance. For every corner turn, almost all of the data leaves one processor to go to another processor. Thus a SAR system requires many processors to achieve real time results. Multiprocessing architectures usually have interconnects which scale with the number of processors, so that one way to achieve the required performance is by adding processors which increases the effective system bandwidth. However, there is usually a limit to how many processors can be added effectively, and this can be a potentially expensive solution.

Interconnect is a key component to system performance. Serial busses are the future of high-performance computing, providing the interconnect speeds needed to effectively build systems of many processors.

Considering the three options (InfiniBand, RapidIO, Serial FPDP), how does a developer solve the application problem? Going back to the SAR radar application, a system might be configured as:

1. Serial FPDP provides the connection to the radar receiver. Separating the transmitter/receiver from the computer system and using a fiber optic cable simplifies packaging and the RFI management problem.
2. RapidIO would be used to build SMP clusters of processors. Each cluster would require only a single board in the system. The low latency and high bandwidth of RapidIO would improve the corner turn performance over today's solutions by 3X or more. The net result would be an increase in throughput or decrease in system cost.
3. The SMP clusters would be connected together using either InfiniBand or RapidIO.
4. InfiniBand would be used to connect the system to the outside and to high performance peripherals. Often there is a requirement to log the input data, for example. InfiniBand could also be used to transmit the resulting image to a workstation for viewing, or another system for further processing.

Serial FPDP

For many applications, Serial FPDP is also important. It is important because it provides an easy, inexpensive way to get high bandwidth data into systems.

The original Front Panel Data Port (FPDP) specification provides a simple 32-bit parallel bus for data acquisition implemented over a ribbon cable. Bandwidth is up to 160 MB/s. It is typically used to connect from A/D converters to the processor(s). The primary limitation is cable length - about 1 meter. By adopting serial technology, cable length can be increased significantly. Using Fiber Optic media, cable lengths can potentially be kilometers in length. Bandwidth is currently at 105 MB/s, but that could easily double by using the latest 2.5 GHz serial technology.

A Real World Example: Synthetic Aperture Radar

Processor performance was once the only metric used to determine the size and success of a multi-processor system. Metrics such as GFLOPS or achieved performance on a key algorithm such as the FFT, were the primary factors in selecting a solution. With the speed of today's processors, bandwidth is now the primary metric for many applications which previously used computational speed as the prime determinant of system size. One real-world example of this is Synthetic Aperture Radar (SAR).

SAR is a technique used to map ground features with very high resolution. A radar system is installed on a vehicle, typically an aircraft or satellite, and is pointed to the side of the vehicle. The radar transmits pulses for some period of time, and the radar returns are collected. Each individual radar pulse fans out as it travels, with the result that each individual pulse has limited resolution. However, because the vehicle is moving, each pulse covers a different portion of the ground, and the pulses have significant overlap. Processing all of these pulses together results in a high resolution image.

Because of the number of computations required to form the image, a multi-processor system is required to produce the result in a reasonable amount of time. The typical processing flow is as follows:

1. Collect data that will be used in the processing.
2. Arrange the data into a matrix, with one axis labeled "range", the other labeled "azimuth". Going along the range axis the data originates from locations farther from the vehicle along a single pulse. Going along the azimuth axis represents successive pulses as the vehicle moves.
3. Image formation begins by processing in the azimuth direction. In order to maximize performance, each row of data in the azimuth direction should be contiguous in a single processor. Unfortunately, the data from the radar arrives in range order, so before processing can start the data must be re-arranged in order to maximize performance. This particular form of re-arranging the data is called a "corner turn",

build very large computer room sized systems composed of hundreds or even thousands of smaller computer systems and supporting peripherals.

As originally architected, InfiniBand was primarily cables connecting boxes. However, the standards have evolved so that a backplane and a packaging definition have been defined as well. The base technology for InfiniBand is 2.5 GHz differential pairs - one transmit and one receive - using 8B/10B encoding to move packets over the connections. Higher bandwidth connections are possible by grouping 4 or 12 of these basic LVDS pairs together, for up to 3 GB/s full-duplex connections. Crossbar switches are used to tie InfiniBand connections together.

RapidIO

RapidIO is newer than InfiniBand, and doesn't have the number of people involved as compared to the masses that have been working on InfiniBand. RapidIO is intentionally smaller and simpler, so the specification should never be as large as InfiniBand. The RapidIO trade association was formed in June 2000 and its membership is growing. Besides SKY Computers, members include: Alcatel, Altera Corp., Thales Computers, Cisco Systems, EMC Corporation, Ericsson, Galileo Technology, Lucent Technologies, Mercury Computers Systems, Motorola, Nortel Networks, PLX Technology, Praesum Communications, RedSwitch, Rydal Research and Development, Seagull Semiconductor, StarGen Technologies, Tundra Semiconductor Corp., V3 Semiconductor, and Xilinx, Inc. With this level of participation it is expected that RapidIO will be rapidly defined.

The target application of RapidIO is different from InfiniBand. InfiniBand is positioned to tie together multiple computer systems, while RapidIO is intended as the fabric to tie multiple processors into a single computer system. RapidIO provides support for cache coherency as well as some simple message passing. On the other hand, InfiniBand is more oriented toward capable message passing, and doesn't concern itself with the level of cache coherency required in a symmetric multi processing (SMP) system.

The base technology for RapidIO is still LVDS, but it is used in a more traditional parallel architecture. The minimum implementation involves 40 wires, with 8 bit paths in each direction along with separate clock and framing signals. These links are used to carry packets, and crossbars are used to switch the packets between links. Data bandwidth ranges from 500 MB/s to perhaps 4 GB/s depending on bus width and clock rate.

There are a number of RapidIO issues that still must be resolved before it makes sense to start an implementation. Some of the more obvious issues involve the lack of backplane and packaging specifications, even though it is supposed to be used in exactly these environments.

Serial Interconnects Support High Performance Computing

By: Steve Paavola, Director of Advanced Development

Introduction

As microprocessors continue to push the limits of performance, busses are running up against the laws of physics in their attempts to keep up. The large voltage swings required for parallel busses limits their achievable peak clock rate and increasing bandwidth to increase the bus width, is impractical because of the pin count and signal skew issues. Parallel busses have been evolving for a long time, but they are running out of steam. Serial interconnects will provide the quantum leap in performance and capability to complement today's processor performance to enable exciting new applications.

The interconnect solution for high performance DSP is coming from the telecommunications industry. By using lower voltage swings and differential signaling (LVDS), it's possible to get very high data rates using a serial connection. Whereas a parallel bus may require 64 data lines plus some parity and control lines, and a clock rate of 133 MHz (PCI-X) to get a half-duplex bandwidth of 1 GB/s, components are becoming available that can provide 250 MB/s bandwidth full-duplex. This delivers an effective bandwidth of .5 GB/s over 4 wires using LVDS serial technology clocked at 2.5 GHz.

The computer industry is building on this serial technology to define and standardize serial busses to replace the parallel busses. There are over 50 proposed interconnects, with the primary public specifications being InfiniBand™ and RapidIO™. Serial FPDP also has some advantages, although it isn't as general purpose as the other two.

InfiniBand

If maturity of a proposed standard specification is measured by weight of the spec, InfiniBand is a very mature specification. This is a very large, complicated specification requiring nearly 2,000 pages. InfiniBand has been under development for two or more years and has the backing of several major companies (Compaq, Dell, Hewlett-Packard, IBM, Intel, Microsoft, and Sun Microsystems to name but a few of the over 200 member companies), with probably over 100 engineers either writing the specifications or implementing product. Another measure of maturity is that chips are already shipping, with many more soon to be released.

InfiniBand was originally designed as a replacement for PCI. It provides performance, scalability and reliability not possible for PCI. InfiniBand allows developers to easily