SILICON GRAPHICS | The Source of Innovation and Discovery™

White Paper

Truly Parallel Processing and Dramatic Application Speed-up with SGI® RASC™ Technology

Table of Contents

1.0 Introduction

Application specific acceleration using hardware methods is not new to computing or to SGI. Technologies such as digital signal processors (DSPs) or custom ASICs have been deployed in several HPC application areas. For example, SGI® Origin® 2000 and 3000 systems tightly integrated MIPS processors along with the SGI Tensor Processing Unit (TPU) which was specifically designed to accelerate applications which require FFT style calculations. The TPU provided approximately forty times (40X) performance improvement.

Reconfigurable solutions built upon field programmable gate array (FPGA) technology have been deployed in such industries as manufacturing, government, research and media. Applications that have benefited from FPGA acceleration have included image recognition, decryption, gene comparison, video format translation, data compression and color intermediaries in the special effects industry. However, these efforts have met with mixed success, due to historical limitations in FPGA performance and host system integration. Traditionally FPGAs also require significant programmer expertise in order to achieve orders of magnitude performance increases.

Recent advances in FPGA capability, power consumption and cost combined with the breakthrough I/O capabilities of SGI® Altix® servers results in extraordinary performance breakthroughs. A greater number and a larger variety of problems can now be addressed by using FPGAs to dramatically increase levels of parallelism and provide orders of magnitude application speed-up.

2.0 FPGAs are Designed for Real-time User Configuration

A FPGA (field programmable gate array) is a semiconductor device containing programmable logic components, or blocks, and programmable interconnects. The programmable logic blocks can be programmed to duplicate the functionality of basic logic gates (such as AND, OR, and NOT) or more complex combinatorial functions such as decoders or mathematical functions. A hierarchy of programmable interconnects allows the logic blocks of an FPGA to be interconnected as needed by the system designer after the manufacturing process so that the FPGA can perform whatever logical function is needed, such as solving a specific algorithm.

3.0 Trends in FPGA Technology Point to New Opportunity for HPC

Prior to the mid-1990's FPGA technology lagged that of ASIC (application specific integrated circuit) microprocessor tech-nology. However, R&D by FPGA industry giants Xilinx and Altera and fabrication improvements by IBM and others have led to faster and more capable FPGAs. In addition, the trends towards a 'wired' world has led to the demand for faster and more capable FPGAs and given them the attractive cost structure of a volume consumer technology. Finally, with the number of gates increasing to one million or more per FPGA, more useful work can be accomplished in an FPGA than was possible five or so years ago.

Current offerings of FPGAs allow:

- Up to a 500 MHz clock rate
- Multi-gigabit serial I/O
- Onboard scalar processor cores
- 90nm process allowing for low power consumption
- Terabyte/sec memory bandwidth, tens of Tops/sec

4.0 Common HPC Industry Approaches for FPGA Deployment Limit their Potential

System Hardware: Although variations on the theme do exist, the typical instantiation of an FPGA in a system is in a co-processor model with the FPGA available via a PCI(X) or VME I/O bus. Scalar processors on a CPU mother board arrangement provide the housekeeping functions and run the operating system and user interface.

As in most co-processor models, data is loaded into the FPGA in a user directed DMA operation and results are loaded back to main memory. This model has several advantages. First of all, PCI and VME buses are ubiquitous through all general-purpose computing platforms and provide a solid and well-known development environment. It is also an inexpensive solution that speeds time to market and leverages existing ASIC simulation software solutions.

The disadvantages to this architecture are many-fold. Because of the slave I/O nature of the interface and I/O bus limitations, latency to the FPGA is high and there is relatively limited bandwidth back to main memory. The limited PCI bandwidth also makes it difficult to integrate the FPGA into high-bandwidth networks and filesystems.

System Software: A variety of software stacks exist in the reconfigurable computing marketplace today. Some are the very same hardware design tools that hardware designers use. In order to utilize the tools, however, programmers need to develop a Hardware Design Language (HDL) which is used to create a bitstream that is downloaded into the FPGA. HDL use is both foreign and difficult for most ISV or end-user programmers who are usually more familiar with C.

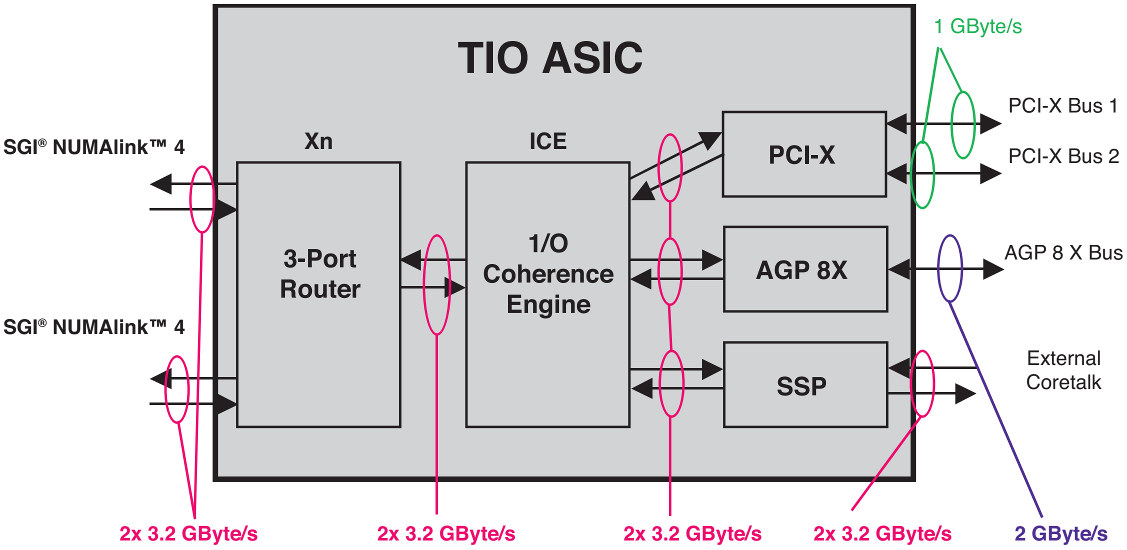

Figure 1. Diagram of SGI's Scalable Systems Port, which is based on the TIO ASIC

For users who do not want to program in HDL, a number of schematic-based tools exist that allow the FPGA programmer to take existing modules for specific computing functions (say an FFT module) and link them together with other modules. An HDL bitstream is created by the tool which can be downloaded to the FPGA.

For scientists, engineers, and other researchers who want to program in higher language a number of C-like languages exist. A compiler or virtual processor converts the C-like constructs into HDL so that a bitstream can be created. The drawback of this routine is that the technology is still in its infancy and need to be optimized further (similar to the early days of C to assembly compilers).

There are a variety of FPGA development tools available in the marketplace today that answer the needs of a variety of development skills from the hardware designer to a scientist. The disadvantage of the current tools is that existing code in higher languages (C, C++, FORTRAN) will need to be re-programmed in either HDL and/or C-like languages. This results in substantial work as well as results in a second binary (the first binary is the one for scalar processors).

5.0 SGI approach to FPGA Based Acceleration –Reconfigurable Application Specific Computing or RASC™ Technology

SGI believes that FPGAs and other reconfigurable processors offer significant – and in some cases orders of magnitude - performance improvements at lower power consumption and heat dissipation than conventional microprocessors when integrated into a system with extraordinary I/O performance.

High performance scalar microprocessors such as IBM POWER and AMD Opteron make compromises in order to adequately fulfill the needs of a wide range of applications. For some applications, these compromises place significant, even dramatic limitations on workflow. For example, what if an application needs to do many 'add' operations? With Intel® Itanium® 2 Processor, users are constrained to the two adder units available. In contrast, FPGAs enables users to create their own unique core to have logic that allow up to twenty adds per cycle – a 10X improvement.

Reconfigurable Application Specific Computing (SGI® RASC™ technology) is the term SGI has coined in reference to its family of capabilities addressed by FPGA technology. While this paper mainly discusses the FPGA elements in SGI RASC technology, additional application specific computing categories which also benefit from the very high I/O capacity of SGI Altix servers include Graphics Processing Units (GPU), Digital Signal Processors (DSP), math or floating-point processors, and custom ASICs.

SGI Approach to Hardware Architecture: For the high performing RASC solution SGI is using its Scalable Systems Port (SSP) enabled via the TIO chip (Figure 1) to provide a very high bandwidth, low-latency direct interface to each FPGA. This has a number of benefits:

- Higher bandwidth (up to 12.4GB/s) per RASC™ RC100 blade
- Lower latency
- Direct access to the memory coherency domain

The higher bandwidth and lower latency coupling of the FPGA directly to system memory results in dramatic performance improvement. A simplified programming model in virtual memory illustrates how SGI's unique architecture provides FPGA users with distinct performance advantages. Using virtual memory, the user can map a virtual address to access the memory of the FPGA, allowing the use of a global pointer for access. Without virtual memory, the user would need to read(2) and write(2) data back and forth between user and kernel space (similar to a disk drive access). Another added benefit of using virtual memory is that common synchronization constructions such as barriers and semaphores can be used to synchronize the workings of multiple FPGA cards.

SGI Approach to Software Architecture: FPGA adoption in broader-based application segments has been limited by a development environment that is geared more to hardware design engineers than end users. Scientists and engineers who might be interested in deploying FPGAs would be far more comfortable with a development environment based upon language such as C, and debugging tools such as Gnu Debugger (GDB), the open source Linux debugger. Because of this, SGI has designed a user-friendly development environment based on familiar, industry standard elements.

SGI RASC software architecture is best described in layers (Figure 2). The upper layer resides in user space. This layer allows an application to manage an FPGA "device" in the familiar context of a Linux device; run algorithms (which may or may not be FPGA-accelerated); call libraries (which may or may not have FPGA-accelerated routines in them); and debug via an FPGA-aware version of GDB.

The next layer resides in the Linux kernel and is the device driver layer which includes device-specific code to manipulate the FPGA hardware. The kernel layer includes a download driver to download a bitstream into an FPGA. There is another driver to manage the synchronization of data movement to and from the FPGA as well as algorithm execution.

The final layer is the reconfigurable element, commonly called the core, which contains the FPGA logic. Typically, when an FPGA 'routine' is called by a program, and that program is being debugged, the debugger is unable to 'step' through the program and then smoothly into the 'routine' as if it were a non-FPGA (CPU only) routine. SGI has done work with GDB to allow a much more interactive debugging environment including stepping through an FPGA routine.

SGI has also collaborated with several third party tool vendors to make their HLL tools available with RASC technology. These tools include Celoxica Handel-C and DK Design Suite, Starbridge Systems Viva and Mitrionics Mitrion C.

Application Areas: FPGA use lends itself very well to integer and bit data. As FPGA sizes continue to increase, customers are having success designing FPGAs to perform floating point operations and even double precision problems. The bit data capabilities of FPGAs are superior to those of scalar microprocessors and lend themselves to some scientific, filtering, and cryptanalysis applications. Working with early access customers, SGI RASC technology has produced actual bit-matrix-multiply (cryptography) results of 119X over those seen on traditional scalar processors.

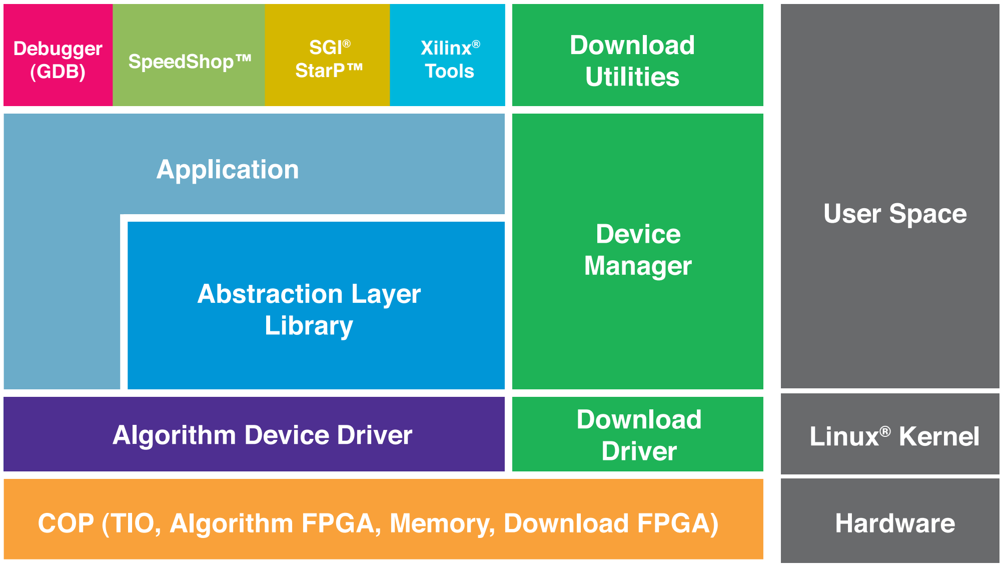

Figure 2. SGI software architecture for RASC technology

Signal processing algorithms such as Fourier transforms, fast Fourier transforms (FFT), convolutions and digital filtering also lend themselves very well to FPGA use. These algorithms are very computationally rich with relatively little decision structures. The advantage of the FPGA approach is that new algorithms can be quickly tried and trade-offs tested in hardware and reprogrammed “in the field.”

The convergence of integer operations performance and signal processing benefits FPGA users in the signals and image processing application areas which include the military and homeland security. Image recognition in the visual, radar and infrared spaces is one of the most popular application areas of FPGA technology today, as well as signals and electronics intelligence 'snooping.' However, there are also some related commercial applications, such as Doppler radar.

Signal processing is also important in the areas of vibration and acoustics which have wide application in a lot of industries. For example, the need to run many different simulations of a particular model to ensure optimal quietness in an automobile puts significant stress on a scalar microprocessor. By using an FPGA, the number of simulations that can be run as well as the complexity of each one can be improved.

FPGAs have seen wide use already in format conversion in the telecommunications industry with an extension into the media broadcast industry. FPGAs can be used to create color intermediaries in digital movie production, a technique to simulate the various effects that used to be available through applying different development techniques to celluloid film. In addition, format conversion from Standard Definition (SD) to High Definition (HD) to 2K to 4K and back, and encryption of valuable digital assets, might be accomplished via an FPGA. The 'dual port' nature of a lot of FPGA designs allows data to "pass through" the FPGA and enables pre- and post-processing of data.

Other application areas that show promise for FPGAs:

- Some seismic processing algorithms
- Data compression and data search of databases
- Format conversion, data encryption, color intermediaries, and digital watermarks in the video and motion picture industry
- Data encryption and decryption in the homeland security area

6.0 Conclusion

Combining the high performance of FPGAs and the extraordinary I/O capability of the SGI Altix servers, SGI RASC technology raises the bar for several classes of important HPC applications. The RASC solution brings this capability to users through a software environment that allows control of powerful yet unfamiliar FPGAs in the friendly terms of C, Fortran and common Linux tools. SGI RASC solution is being deployed today in development mode for such applications as image processing and encryption. The complete RASC solution is emerging as the leading platform for reconfigurable computing, with the capability to deliver extraordinary performance breakthroughs for many HPC problems of critical importance to science, industry and national security.

sgi®

Corporate Office
1500 Crittenden Lane
Mountain View, CA 94043
(650) 960-1980
www.sgi.com

North America +1 800.800.7441
Latin America +55 11.5509.1455
Europe +44 118.912.7500
Japan +81 3.5488.1811
Asia Pacific +1 650.933.3000