

Designing Reconfigurable Computing Solutions

The Virtex family of FPGAs is the foundation for Cray XD1 high-performance co-processing solutions.

by Geert C. Wenes, Ph.D.
HPC Architect
Cray, Inc.
geertw@cray.com

Steve Margerm
Senior Hardware Designer
Cray, Inc.
smargerm@cray.com

Sriram R. Chelluri
Marketing Manager
Xilinx, Inc.
sriram.chelluri@xilinx.com

The reconfigurable computing (RC) architecture enables software logic that can be reconfigured or reprogrammed to implement specific functionalities on tunable hardware rather than on a general-purpose processor (GPP). RC can achieve orders-of-magnitude performance improvements on selected applications. An RC solution in a high-performance computing environment should include tight coupling and a transparent interface to a general-purpose processor and its auxiliary resources (storage, I/O, networking).

Because FPGAs are tunable, high-density logic cores that provide high-performance and low-latency features, they are well suited for RC solutions. In this article, we'll describe what RC is, how it fits into the Cray, Inc. XD1 high-performance architecture, and why Cray selected Xilinx as the foundation for high-performance application acceleration.

What's Different about RC?

A GPP is primarily designed to execute sequential instructions. System designers have introduced parallelism in two different ways: off-chip, by simply clustering many GPPs and distributing the workload among these multiple processors; or on-chip, either by duplicating the independent computational units on the chip or (more recently) by doubling the cores on a processor.

In contrast, FPGAs can perform many operations in parallel. Xilinx® FPGAs with embedded PowerPC™ processors can perform fixed-point arithmetic and embedded functions, offloading the main processor to other tasks. Furthermore, standard HDLs

for logic devices (such as VHDL and Verilog) have at their core the notions of parallel execution of statements and event-driven simulation.

FPGAs, on the other hand, do not suffer from serial execution model restrictions found in standard languages such as C for generic CPUs. FPGAs are increasingly used to provide application performance. Industries such as manufacturing, government, research, media, and biosciences are deploying FPGAs as hardware application accelerators that can provide orders-of-magnitude (10 to 100x) performance improvements on selected applications over generic microprocessors. The advantages of such an RC solution are:

- Dramatic overall system performance gains at only incremental co-processing board costs
- The ability to reprogram, utilizing a consistent API to the hardware
- Easily upgrade or retarget hardware

Cray XD1 Architecture

The Cray XD1 high-performance computer is based on the directly connected processor (DCP) architecture. The DCP architecture views the system as a pool of processing, logic, and memory resources interconnected by a high-bandwidth low-latency network. This innovative new computer unifies as many as hundreds of processors into a single, resilient computer. The Cray XD1 system combines both on-chip and off-chip parallelism.

The Cray XD1 architecture includes three key subsystems:

- Compute environment. The Cray XD1 compute subsystem comprises single- or dual-core AMD Opteron 64-bit processors integrated on a single board and six blades constituting a chassis (Figure 1). The operating system is Linux, supporting 32- and 64-bit x86-compatible software.
- Interconnect. The Cray XD1 RapidArray interconnect directly connects blades over high-speed, low-latency pathways (Figure 2). Each fully configured chassis includes two interconnect components:

- An FPGA-based RapidArray communications processor (RAP) is tightly coupled to the AMD Opteron processors and switching fabric to offload and accelerate communications functions from the Opteron processors, freeing the latter to perform core compute tasks and enabling concurrent computing and communication. The FPGA enables interconnect bandwidth on par with memory bandwidth, solving a major system performance bottleneck.

- The RapidArray embedded switching fabric is a 96 GB/s, non-blocking, crossbar switching fabric in each chassis that provides four 2 GB/s links to each node and twenty-four 2 GB/s inter-chassis links.

- Acceleration application modules. The application acceleration subsystem incorporates RC capabilities to deliver substantial performance increases for targeted applications (Figure 3). Each Cray XD1 chassis can be configured with six application acceleration processors (one per blade), originally developed using the Xilinx Virtex™-II Pro device



Figure 1 – Single Cray XD1 chassis




Figure 2 – Cray XD1 compute blade

but soon shipping with Virtex-4 FPGAs that you can program to accelerate key algorithms. The application acceleration processors are tightly integrated with Linux and the AMD Opteron processors and use standard software programming APIs, removing a major obstacle to application development.

The FPGA as Co-Processor Model

Cray views the FPGA as a very tightly coupled application accelerator platform to speed up computations on demanding applications. The FPGAs are designed into an expansion module that connects to the high-speed, low-latency, point-to-point HyperTransport subsystem.

There are three main advantages that an FPGA has over a microprocessor:

- FPGAs have a flexible architecture. You can customize and optimize the logic and manipulate variable length data.
- FPGAs are inherently parallel devices.
- FPGAs can be reprogrammed to execute new applications without having to change the hardware.

Xilinx Solution

Cray selected the Virtex-II Pro and Virtex-4 series of FPGAs as the foundation for its RC solutions because of their industry-leading technical features and support for third-party development.

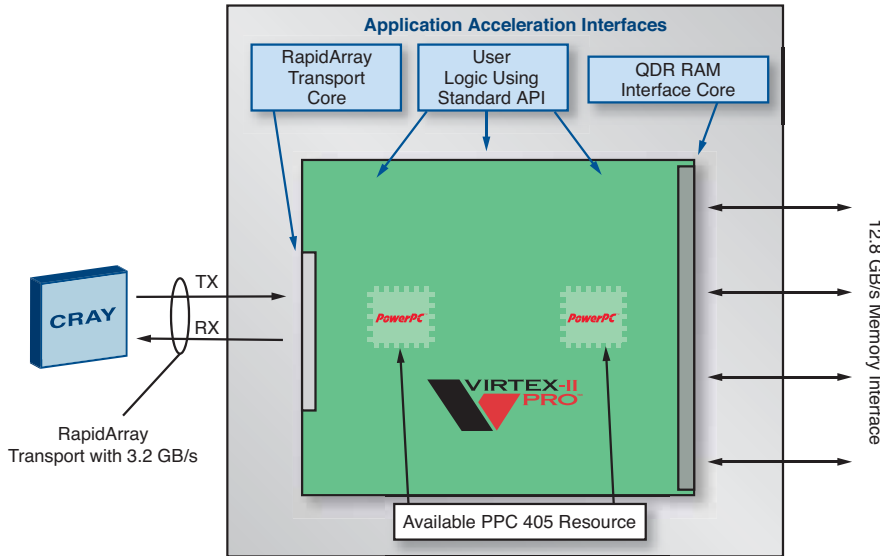


Figure 3 – Xilinx-based application acceleration module

Advantages of Xilinx Virtex FPGAs include:

- Low latency and a high throughput range of 622 Mbps-10.3125 Gbps
- High-speed 500 MHz internal I/O clock
- Low power consumption
- Optional embedded hardware PowerPC processor or a software co-processor
- Built-in scalable RAM

Cray chose the Virtex family of FPGAs because it was developed for high performance, from low- to high-density designs based on IP cores and customized modules.

Development Environment

Many Xilinx and third-party tools simplify the development of co-processing solutions. Software developers can interface with software APIs provided by Cray and use a rich variety of third-party design tools to create their own applications using C and C++.

For instance, Celoxica Ltd. provides a C-based design and synthesis tool (DK Design Suite) for customers who want to use a software design flow to accelerate their applications using FPGAs integrated into the Cray XD1 supercomputer (www.celoxica.com/products/dk/default.asp).

As another example, the Cray XD1 system equipped with Mitronics' Mitrion Virtual Processor and Mitrion Software Development Kit make it possible for supercomputer users to program FPGAs integrated into the Cray XD1 system on a software level, reducing the time and effort required to take advantage of FPGA-based computation (www.mitronics.com/technology.shtml).

For DSP solutions, Xilinx System Generator for DSP enables you to design DSP blocks using commercial tools like the MATLAB package from The MathWorks.

Applications Acceleration

The XD1 has proven to be a successful product. For example, the Naval Research Laboratory (NRL) facility in Washington, D.C., is home to one of the largest Cray XD1 supercomputers ever installed and also employs the largest known number of application acceleration modules in the world. Equipped with 288 AMD Opteron dual-core processors and 144 Virtex-II Pro FPGAs, the 24-chassis machine will provide peak performance of 2.5 teraflops.

The use of FPGAs as application accelerators has been successfully demonstrated in many fields. For instance:

- In encryption/decryption applications, the RC5 cipher-breaking application

runs 1000x faster than on a 2.4 GHz Pentium 4, while for elliptic curve cryptography, speedups of 895 to 1300x compared to (a relatively slow) 1 GHz Pentium III are possible. Encryption algorithms like 3DES have been shown to run at more than 16 Gbps throughput from a high-level abstraction language like Mobius.

- In bioinformatics applications, the well-known Smith-Waterman code performs about 26 times faster than on the AMD Opteron, while applications in proteomics such as thin-spline algorithms for comparing 2D gel contents run more than 20 times faster, reducing analysis times from days to hours.
- Complex, realistic vehicular traffic simulation codes perform 34x faster on Virtex-II Pro devices than on a 2.2 GHz Opteron, and remarkable sustained bandwidths from the FPGA to AMD processor have been observed at more than 1 GB/s, considerably higher than any PCI bandwidth.

Conclusion

Cray's FPGA-based RC solutions are ideal for applications in industries such as manufacturing, government, research, media, and biosciences. Cray selected the Virtex-II Pro and Virtex-4 series because of their high performance, low latency, scalable memory, and an array of software tools and support.

With the Cray XD1 high-performance computer, reconfigurable computing has taken a major step forward by breaking down performance barriers at substantially lowered cost by using off-the-shelf components from Xilinx to solve difficult computational problems.

For more information about Cray RC solutions' HPC architectures, contact geertw@cray.com or visit www.cray.com/products/xd1/index.html.

For more information about Xilinx FPGA-based co-processing solutions, contact srirame@xilinx.com or visit www.xilinx.com/products/design_resources/dsp_central/resource/coprocessing.htm.