



The Leading Source for Global News and Information Covering the Ecosystem of High Productivity

[Home Page](#) | [Free Subscription](#) | [Advertising](#) | [About HPCwire](#)

Features:

DRC Stakes Claim in Reconfigurable Computing

DRC Computer Corporation is one of just a handful of companies hoping to ride the popularity of Field Programmable Gate Arrays (FPGAs) into the high performance computing realm. While the difficulties of FPGA programming has held back their widespread use for general-purpose applications, their versatility and suitability for compute-intensive codes has made FPGAs a tempting platform for HPC. DRC President and CEO Larry Laurich talks about the company's mission and the nature of the technology they've developed.

HPCwire: Tell us a little bit about the company, how it got started and what it is offering the HPC user?

Laurich: DRC is three years old, after having acquired the IP assets from VCC, a company run by Steve Casselman, now DRC's CTO. Steve is one of the recognized "fathers of reconfigurable computing," and holds some of the earliest and most fundamental patents in the area. DRC has been shipping RPUs (Reconfigurable Processing Units) for almost a year, and launched its second generation product a couple of months ago. With the newest product, the RPU110-L200, DRC provides the HPC user with the most tightly coupled co-processor available with the highest useable memory bandwidth by far of any compute platform.

HPCwire: Compared to other FPGA products targeted for high performance computing, what makes the DRC solution unique?

Laurich: By inserting the RPU directly into a microprocessor socket, the coprocessor gets equivalent access to all the motherboard resources a CPU gets, such as direct HyperTransport (HT) access for CPU to CPU communication, local memory bandwidth, etc. It is DRC's fundamental understanding of system level issues affecting performance that has led to RPU designs with additional simultaneously accessible memories. Since many applications are starved for data, especially once the logic is accelerated, the RPU can provide true application acceleration.

HPCwire: The lack of high-level software tools has been a major hindrance to FPGA adoption in high performance computing in the past. What kind of development environment is supported by the DRC solution?

Laurich: DRC has simplified the most difficult part of moving software to FPGA hardware by providing the RPU Hardware OS. The simple API for this OS allows the programmer access to 80 percent of the FPGA logic for his own code but provides a pre-configured and locked design for all physical pins and design issues. The application programmer no longer has to worry

about timing for the bus and memory interfaces. Those controllers are provided along with DMA, back-pressure or flow control, etc., which allows the application to have an independent clock and assures the data can never overrun the logic or system resources. The remaining programming issues are much more familiar to the application programmer and more easily handled in the C to RTL tools provided by our many partners. Celoxica, Impulse Technologies, and Mitronics have all developed support packages for the DRC RPU.

HPCwire: Besides the software challenge, what else do you think is keeping FPGA technology from going mainstream in high performance computing and which of these elements are addressed by the DRC solution?

Laurich: It is a matter of an early adopter demonstrating what can be done in a given application area or vertical market. Once the advantages are shown in a real production environment, the rest of that industry has an easier time making the decision. The price-performance benefit is there, the "green-technology" or power savings are compelling, and reduction in the number of nodes by five times or more reduces system management and footprint, which is a significant advantage. Mass market adoption, however, is reasonably assured given the support by most all of the big players -- namely Cray, IBM, HP, Intel and AMD -- to hybrid compute platforms incorporating coprocessors or accelerators.

HPCwire: How does reconfigurable computing based on FPGAs stack up against other accelerator technologies that have become available within the past couple of years (e.g., GPUs, ClearSpeed boards, Cell processors)?

Laurich: Each of these new technologies has much the same issue relative to software tools and development flow. In fact, so do multi-core CPUs. All these technologies require programs to be multi-threaded, meaning parallelized for performance. Once the application architects figure out what is necessary to parallelize at least portions of their code, a fine-grained implementation for an FPGA is not much different from a coarse-grained one for CPUs.

The FPGA turns out to be the most flexible architecture that can address the largest cross-section of compute intensive applications. It has logic that can stream or be conditional. The RPU has more memory bandwidth than any of the other technologies. There are multiple vendors supplying and developing tools and libraries.

GPUs will do well in highly streaming threads where no conditional processing is required -- a small but meaningful subset of the co-processor market. Programming GPUs can be even more difficult than FPGAs or Cells, but an extensive library for the streaming applications has helped.

ClearSpeed based technology has continually suffered from memory bandwidth, or the ability to move data through the logic at high speed.

Cells are somewhere in-between, but proprietary since both hardware and any compiler or tools are available only from that vendor.

HPCwire: Are there any early adopter stories you can share with us?

Laurich: We have some demos and proof-of-concepts that we have shown the world. Examples

include everything from a programmed trading example in the financial market where we can give the trader a 30-50X advantage in reduced latency -- which a publication states is worth a minimum of \$100 million per year -- to a seismic imaging application where the user gets the same performance as software running on a large cluster with a quarter the number of nodes and a fifth the amount of power consumed, at half the price.

HPCwire: What's next for DRC?

Laurich: More improvements in the Hardware OS will give future RPUs much more intelligence and system capability. Likewise, the RPU will come in configurations to support newer motherboards with different sockets, more workstations, servers, and blade systems. From an application perspective, more libraries and pre-programmed applications will provide more solutions faster and easier.

Leading HPC Solution Providers



HPCwire Readers' and Editors' Choice Awards