

# CARMA: A Comprehensive Management Framework for High-Performance Reconfigurable Computing

Ian A. Troxel<sup>1</sup>, Aju M. Jacob, Alan D. George, Raj Subramaniyan, and Matthew A. Radlinski  
High-performance Computing and Simulation (HCS) Research Laboratory  
Department of Electrical and Computer Engineering, University of Florida  
Gainesville, FL 32611-6200

Programmable logic devices are rapidly becoming a cornerstone architecture technology for a broad and increasing range of critical applications. One relatively new direction for exploitation of PLD technologies is as the basis in fundamentally new architectures and systems for computationally challenging applications that require high-performance computing. Whether serving as the basis for high-speed embedded systems, or featured in the next generation of supercomputers, the advantages of PLD technologies in performance and versatility will be exploited in adaptive computing systems to solve problems that require extremely high efficiency of computation, communication, and data access. Based on commodity parts and emerging technologies, reconfigurable computing systems hold the potential to achieve performance levels far exceeding those of conventional systems while maintaining the advantages of cost, interoperability, etc. associated with systems based on industry standards and off-the-shelf components. Following the trends in high-performance computing, powerful reconfigurable systems are likely to be heterogeneous in nature, featuring a diverse set of processing, communication, and storage technologies. In these systems, speedup and offload of various computation and communication tasks will be provided through dynamic hardware reconfiguration, with FPGA and CPU devices working in tandem to solve key tasks in a collaborative fashion.

One of the limiting factors in building and exploiting such systems has been the lack of a simple and effective management framework within which applications, system services, programming models, and middleware can be developed and ported to support a broad range of platforms and tools. Such frameworks exist in the allied fields of cluster computing and grid computing, but to date significantly less attention has been paid to such issues in reconfigurable systems for high-end computing. While initial algorithm development on large-scale reconfigurable systems has achieved significant performance improvements for select applications, one of the biggest challenges confronting system designers and users is a lack of comprehensive, powerful, and vendor-independent management services for such systems. This presentation will address this issue and focus on research activities and new results in the design and analysis of open-system infrastructure for hardware-adaptive, reconfigurable systems for high-performance computing.

In this presentation, we propose a new framework known as the Comprehensive Approach to Reconfigurable Management Architecture (CARMA). CARMA provides the basic infrastructure to develop and integrate key components for reconfigurable high-performance computing systems. Some of the many examples of components that fit within this framework include board-independent application mapping, dynamic and robust job scheduling, distributed configuration management, scalable performance monitoring into the hardware, and board-independent interface modules for the reconfigurable fabric. With CARMA, we seek to build a unified, fault-tolerant, scalable tool to specifically address key issues such as dynamic reconfigurable fabric discovery and configuration management, coherent multitasking in a versatile multi-user environment, scheduling and management of heterogeneous jobs including both hardware-reconfigurable and software-reconfigurable (i.e. conventional) tasks, breadth-wise and depth-wise performance monitoring across networked nodes and into the FPGA fabric for both debug and performance management, and vendor-independent middleware and programming models. This presentation will introduce the CARMA framework and showcase the initial design and development of modules within the framework with a focus on design decisions, tradeoffs, and lessons learned. Included will be a performance analysis of the initial top-to-bottom prototype of CARMA using several case studies to show cost versus performance tradeoffs performed to tune the tool to specific systems. Finally, projections will be presented to illustrate robustness and scalability for large-scale systems.

---

<sup>1</sup> Corresponding author, email: troxel@hcs.ufl.edu, telephone: 352-392-9046.