

Requirements for Scalable Application Specific Processing in Commercial HPEC

Steven Miller

Silicon Graphics, Inc.

Phone: 650-933-1899

Email Address: scm@sgi.com

Abstract: More and more High Performance Embedded Computing (HPEC) leverages technology from commercial high performance computing systems. To date, HPEC has only tapped the lower end of commercial high performance computing technology. As more of the advanced commercial technology moves into the embedded space, this presents a unique opportunity to change the fundamentals of how HPEC solutions are addressed.

Within HPEC, two types of application specific processing elements, reconfigurable and custom are being used. Reconfigurable elements are comprised of Field Programmable Gate Array (FPGA) technology and custom elements are comprised of various devices as Digital Signal Processors (DSP), DARPA Polymorphic Computing Architectures (PCA) [1] and others. The integration of these devices presents significant challenges both to the system architecture and to the programming models. This presentation will describe a set of system requirements and methods to not only include these application specific processing devices but to allow effecting scaling of application specific processing devices.

Introduction and System Architectural Review

SGI's ccNUMA (cache coherent non-uniform memory architecture) [2] global shared memory system architecture is the basis for our general-purpose Origin and Altix HPC systems. The presentation will explore the architectural features used within both Origin and Altix systems which allows scaling in excess of one thousand commercial high performance processors. The architecture of SGI's systems which included 192 custom application specific processors (Tensor Processor Units) and 128 general-purpose processors will also be described.

System Architectural Features for Scalability

High-performance FPGAs represent an important architectural tool to provide total system performance while keeping both power and space requirements to a minimum. The use of FPGA elements have been limited to only the most computationally dense algorithms by the amount of data which can be pass through the device. But even with these limitations, significant performance and power improvements have been demonstrated [3]. This presentation will describe methods to increase the bandwidth to these devices and the communication primitives required to scale to hundreds of devices with out sacrificing performance.

Systems that include hundreds of application specific processing devices pose a significant software challenge. This presentation will describe how to effectively allocate, manage, and decommission under changing workloads various application specific processing elements. Both software methods and APIs will be presented on ways to interface application specific processing targeted at HPEC applications.

References

- [1] <http://www.darpa.mil/ipto/programs/pca/index.htm>
- [2] <http://www.sgi.com/servers/altix/>
- [3] **W.D. Smith and A.R. Schnore. Towards an RCC-based Accelerator for Computational Fluid Dynamics Applications. In T.P. Plaks, eds., Proceedings of the International Conference of Engineering of Reconfigurable Systems and Algorithms, pp. 222-231. CSREA Press, Las Vegas, Nevada, 2003.**