

Application of General Purpose HPC Systems in HPEC

David Alexander

Silicon Graphics, Inc.
Phone: (650) 933-1073
Fax: (650) 932-0663
Email: dca@sgi.com

Areas that this paper/presentation will address:

- * Reconfigurable Computing for Embedded Systems
- * High-Speed Interconnect Technologies

Abstract: High performance embedded computing (HPEC) has traditionally been performed by systems designed specifically for the task. Recent years have seen the increasing application of general-purpose high performance computing (HPC) systems in embedded applications. General purpose HPC systems typically have a large user base which results in broad application SW and device driver availability, robust development and debugging tools, and revenue streams which support significant R&D funding of technologies to enhance HPC system performance and reliability.

Various factors have prevented wider adoption of general purpose HPC systems in the embedded space...factors such as lack of dense, ruggedized packaging suitable for embedded applications, lack of real-time capabilities in general purpose operating systems [1], and performance/watt and performance/unit volume advantages that specialized systems have traditionally had over general purpose HPC systems. This presentation details plans for addressing these shortcomings through the deployment of a heterogeneous computing architecture which incorporates FPGA-based reconfigurable computing and I/O elements, system interconnect advancements leveraged from HPC system development, microprocessor and system advancements developed under DARPA's HPCS program, and the mapping of the system into packaging suitable for HPEC applications.

Introduction and System Architectural Review

SGI's ccNUMA (cache coherent non-uniform memory architecture) global shared memory system architecture is the basis of our general-purpose Origin [2] and Altix [3] HPC systems. The presentation will examine ccNUMA's architectural evolution from its commercial introduction with the Origin 2000 in 1996 through current Origin 3000 and Altix system implementations. The use of SGI's current Origin and Altix systems in real-time applications such as the Common Imagery Processor and mobile ground station applications will also be reviewed.

HPEC development and deployment workflows for a global shared memory system will be contrasted with workflows typical of distributed memory systems. Also, the impact of shared memory architectures in HPEC application performance and scalability will be discussed.

System Performance Enhancements for HPEC Applications

High-performance FPGAs represent an important approach to enhancing system performance on HPEC applications. The integration of high performance FPGAs into ccNUMA systems and their application to signal processing algorithm acceleration and integration of unique I/O structures will be discussed. The ultimate goal is highly sustained performance, both in absolute terms and in terms of performance/watt and performance/volume. Also, the status of our investigation on algorithm-to-FPGA mapping will be presented. The benefits of this technology are not limited to HPEC applications; they also benefit and will be used in products targeted at the general HPC market...such as FPGA-based accelerators for application in genomics and seismic processing.

We will present our system interconnect technology development plans starting with a discussion of our current 3.2GB/sec link technology (used on current Origin/Altix systems) and our plans to extend this to 10GB/s and beyond in the near future. Various topology implementations using this interconnect technology and their impact on system performance will also be reviewed.

Performance enhancements in microprocessors and memory systems will also yield significant benefits for HPEC applications. SGI is a participant in DARPA's HPCS (High Productivity Computer System) initiative. This initiative will yield general purpose microprocessors and system enhancements that will enable sustained performance in the tens to hundreds of GFlops per processor through the implementation of enhanced memory architectures, the integration of vector and atomic memory operation units into the microprocessor, and through the deployment of new programming environments and debugging tools. SGI's HPCS research on next generation processor and system architectures will be reviewed, and their relevance to HPEC will be discussed.

Mapping to a Suitable Physical Packaging

Many HPEC applications have packaging constraints and environmental requirements that far exceed those of a typical HPC data-center environment. The mapping of our ccNUMA systems (both Origin and Altix) into a packaging approach appropriate for a broad range of HPEC applications will be discussed. A 6U Eurocard form factor "blade" design will be the basis of the repackaging, combined with a flexible backplane and card cage design which will allow for custom configurations of "blades", and the use of either convective or conductive cooling. The various types of "blades", and the methodology for constructing both standard and custom configurations of these blades for use in various HPEC and HPC applications will be reviewed. The potential mapping of the ccNUMA architecture into MCM (multi-chip module) packaging suitable for extreme HPEC applications will also be discussed.

This packaging approach addresses the physical constraints and environmental requirements of HPEC applications while maintaining architectural and logical equivalency with the HPC data-center systems...enabling application development to occur on HPC data-center systems with

deployment on application appropriate systems without the need for significant application porting and tuning efforts.

Footnotes

- [1] A paper addressing this topic, titled "Low Overhead Real-Time Computing with General Purpose Operating Systems", is being submitted by Michael A. Raymond.
- [2] Origin systems are SGI's ccNUMA systems based on 64-bit MIPS microprocessors and the IRIX operating system.
- [3] Altix systems are SGI's ccNUMA systems based on Intel's Itanium microprocessors and the Linux operating system.

Acknowledgements

Itanium is a trademark of Intel Corporation.

Linux is a registered trademark of Linus Torvalds.