

# Abstract: Scalable and Portable Supercomputing

## HPEC 2005

Gail Walters, CPU Technology, Inc., Reston, VA ([g.walters@cputech.com](mailto:g.walters@cputech.com))  
Scott Nelson, CPU Technology, Inc., Pleasanton, CA ([s.nelson@cputech.com](mailto:s.nelson@cputech.com))  
Steve Manuel, CPU Technology, Inc., Pleasanton, CA ([s.manuel@cputech.com](mailto:s.manuel@cputech.com))  
(925)224-9920

### Introduction

The demand for small form factor, low power, high performance computing is growing rapidly. Airborne sensor processing, fusion and autonomous operation are examples of the applications that are driving the need for substantially more performance in embedded systems. At the same time, the ability of Moore's law to provide ever faster microprocessors through higher clock frequency is coming to an end. Performance increases will rely on architecture, specifically, on-chip parallel processing, going forward. An emerging general-purpose parallel processing technology called SuperQ™ is under development with a goal of providing scalable supercomputer-class computing while delivering significant advances in reduced form factor, power and weight, i.e., Mobile Supercomputing™.

### Mobile Supercomputing™

Solving many of the most important scientific and engineering problems today requires the use of high performance computers, including supercomputers. The embedded computing community is now more frequently tasked with meeting these needs with innovative computing solutions. However, there are several major issues currently facing the field in order to deliver this level of computing: 1) the large size and weight of the systems, 2) the growing requirements for powering and supporting the systems, 3) the special environment necessary to keep the systems running, 4) highly reliable fault tolerance, and 5) providing advanced security for classified applications.

Successfully resolving these issues enables supercomputing to become mobile so the system can be taken to where it is needed rather than trying to bring the problem to the system. Mobile supercomputing enables the utilization of capability class performance by a broad range of applications which cannot currently be served by existing architectures. The SuperQ was developed specifically for this purpose: bringing supercomputing to portable systems with scalability to teraflops and beyond.

### Introduction to the SuperQ™ Architecture

Developed for critical applications, SuperQ also has the high reliability and environmental robustness required for embedded applications. SuperQ was architected for efficient parallel processing. Each SuperQ System-On-Chip (SOC) is an entire multiprocessor computing node,

including high performance 64-bit scalar and vector processing engines (fixed and floating point), memory, high-bandwidth memory interface and high-bandwidth input/output interfaces. The level of integration makes the SuperQ extremely scalable: chips can be added to a system without any extraneous circuitry. More importantly, the system remains in balance as more processing capacity is added.

In addition to low power and high performance parallel processing, SuperQ is the first commercial instance of ultra low latency inter-processor communication described as Deeply Coupled Computing [1]. As shown in Figure 1, SuperQ facilitates fine grain parallel processing through extremely low latencies between processor interactions, from zero (0) to one (1) microsecond. It also has support for instruction level parallelism and message passing. The ability to perform fine grain parallel processing makes the partitioning and accelerating of programs that require moderate to high level of inter-processor communication feasible. Deeply Coupled Computing also makes programming multiple parallel processors more forgiving since one is not penalized by high overhead/latency.

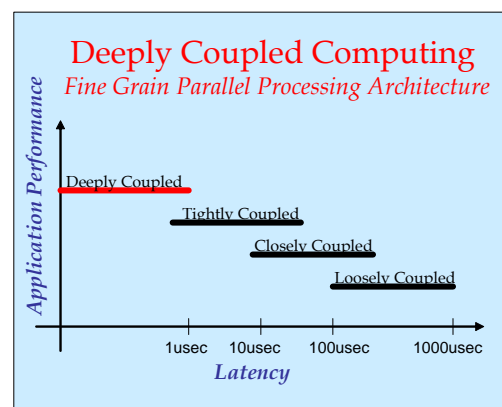


Figure 1: Deeply Coupled Computing.

SuperQ-based systems can be configured with a variety of industry standard bus interfaces and formats. It supports most commercial Operating Systems and will support MPI/OpenMP. Its programming model facilitates the porting of existing software so it can be readily added to existing open systems.

The first multi-core chip configuration of SuperQ will contain:

- (4) 64-bit Scalar Processor(s)
- (16) Double Precision, Floating Point Engines-organized as (4) eight stride vector units
- (4) Communication Processors
- (12) Independent 64-bit Bus Interface Units
- 32 Megabytes RAM.

The system performance benchmarks for the SuperQ system are world-class and the performance/watt benchmarks are superior to most current and near-term commercial microprocessors.

CPU Technology will present a general overview of the SuperQ architecture. In particular, we will highlight the measurable advantages of producing a uniquely balanced and efficient approach to providing significant increases in computational, communication and memory bandwidth within a highly scalable parallel processing environment.

The net result of this balanced scaling of ultra low latency processing nodes offers the opportunity to yield incremental performance through the exploitation of fine grain application partitioning.

## **References**

- [1] CPU Technology, Inc., CPU Tech Announces Breakthrough to Deeply Coupled Computing, November 8, 2004. <http://www.cputech.com/press.html>