SeaMicro Volume Server Power Reduction

Saving 75% of computing energy through an inter-connected array of low-power central processing units (CPUs).

Introduction

Volume servers, the low-cost servers most commonly used for internet services and cloud computing, are estimated to use more than 1% of the nation's electricity. These servers have significant room for energy efficiency improvements; they rely on an older architecture that is no longer well aligned with the most common computational requirements, and this mismatch between server design and most common usage results in wasted energy. The technology used for this project is expected to reduce server computing energy use by 75% compared to conventional volume servers. The project partners at SeaMicro build a system out of simple, replicated blocks of highly efficient technology. Their innovations are in the tying together of these blocks, which is done cheaply, easily, and in such a way that existing software can run unchanged on each block. Their technology enables the disaggregation of highly complex CPUs by enabling small, efficient CPUs to work in parallel without software changes.

The data shows that for the highly partitioned workloads that characterize the internet, simple replicated compute blocks intelligently tied together will outperform large complex highly integrated compute blocks on the dimensions of compute per unit power and compute per unit cost.

Benefits for Our Industry and Our Nation

Data centers and telecommunications facilities consume 3% of our nation's electricity, of which volume servers make up the largest segment of electricity consumers. Accelerating commercialization of this new server architecture creates an alternative solution, capable of achieving a 75% reduction in power use per unit of computation over conventional volume servers. Adopting components originally intended for handheld electronics and intelligently tying them together enables SeaMicro to reduce server energy consumption while increasing overall computing performance.

By addressing the power use of both the CPU and the associated components, the SeaMicro architecture yields energy reductions and reduced cooling needs. In addition, the server takes one-fourth of the space to do the same work as do the best-in-class volume servers.

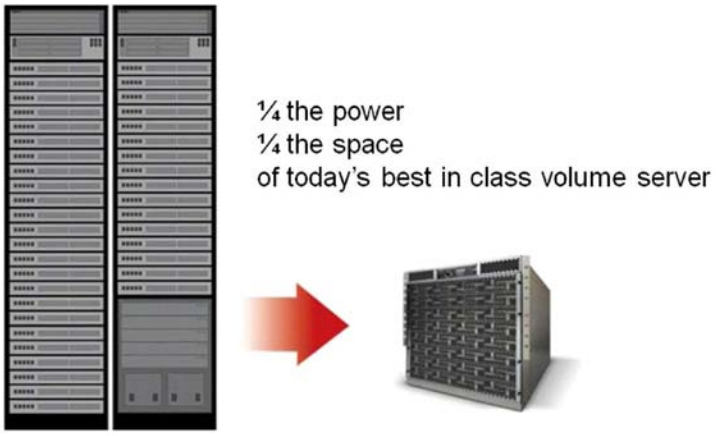

Re-architected server solution for higher energy efficiency.

Illustration courtesy of SeaMicro.

Applications in Our Nation's Industry

Of all server purchases, volume servers (servers that cost less than $25,000) are the most frequently purchased. They were responsible for the majority of electricity used by U.S. servers and related IT equipment in 2006. The server market for volume servers grows aggressively every year and serves multiple business needs, including the following:

- Internet application services and telecommunications transactional services
- Service industry businesses

Project Description

The internet computing model has dramatically increased system requirements for simple, parallel transactions, rather than the complex computations that gave rise to the computer age. Multi-core processors, popular in servers today, are still built around the expectation of complex transactions. This creates significant mismatch between the server and the work to be done, which is manifested in an unnecessarily high power draw.

The new server, based on systems architecture pioneered by Lawrence Berkeley National Laboratory, replaces the sophisticated—yet power hungry—chips found in conventional servers with hundreds of ultra-low-energy chips. A SeaMicro modification reduces the circuit board from the size of a pizza box to the size of a credit card. Server architectures using a massive array of low-power CPUs can increase the efficiency of handling multiple simple transactions in parallel. Similar to High-Performance Computing architectures, this project features technology that spreads the costs of management and connectivity over hundreds of these CPUs.

Actively managing processor power state (on, off, standby) allows for maximum throughput during active times and minimum power use during idle times. The technology in this project uses

server management and load-balancing software to dynamically allocate workloads to specific CPUs. This intelligent work allocation allows the CPUs to operate in the most efficient range, while allowing CPUs that are not being utilized to enter deep sleep.

In a volume server, two-thirds of the power used is consumed by components other than the CPU. SeaMicro realized this and developed input/output (I/O) virtualization technology that allows the removal of all non-CPU/memory components from the motherboard (the main electronics board in the server that houses the CPU and related memory components), while allowing the CPUs to run standard operating systems and software without requiring modification or recompilation. The CPU I/O virtualization technology also allows some common components to be shared across hundreds of CPUs, rather than duplicated on each motherboard.

Hundreds of the modified low-power, credit card-sized computational units are tied together to create a massive array of independent—yet linked—computational units. Work is then distributed over these CPUs via hardware- and software-based load-balancing technology that dynamically directs loads to ensure that each CPU is either in its most efficient performance zone or sleeping. The key technologies reside in three newly designed chips and in the managing, routing, and load-balancing software that directs traffic across this array of CPUs.

Barriers

A change to conventional volume-server architecture requires lengthy, expensive Application-Specific Integrated Circuit (ASIC) fabric-design services. This is why many server system board manufacturers do not depart from a handful of existing pre-fabricated designs. Major barriers include the following:

- Significant upfront financial investment to implement a new circuit board design with new custom-designed chips
- Significant time and tools for performing load testing with willing data center partners

This project is a proving ground for new, energy-efficient processor architectures that can flourish in the market.

Pathways

The DOE Energy-Efficient Information and Communication Technology funding is accelerating product availability by 12 months through development engineering, quality-assurance testing, and enhancing dynamic power-management software to further decrease power usage.

Milestones

- First phase completion of circuit board printing (Completed)
- First phase development and testing of integrated chip components (Completed)
- Delivery of first commercial systems to customers (Completed)
- Development and testing of Version 2.0+ hardware/software

Project Partners

SeaMicro Incorporated
Santa Clara, CA
Principal Investigator: Gary Lauterbach
E-mail: gary@seamicro.com

For additional information, please contact

Gideon Varga
Technology Manager
U.S. Department of Energy
Industrial Technologies Program
Phone: (202) 586-0082
E-mail: Gideon.Varga@ee.doe.gov

Information & Communications Technology

Equipment and Software

Research & Development

U.S. DEPARTMENT OF ENERGY
Energy Efficiency & Renewable Energy

EERE Information Center
1-877-EERE-INFO (1-877-337-3463)
eere.energy.gov/informationcenter

DOE/EE-0503 • May 2011
Printed with a renewable-source ink on paper containing at least 50% wastepaper, including 10% post consumer waste.