SiCortex: The Open Source Cluster Systems Company

G. Winfield Treese

December 2006/revised April 2008

The SiCortex family of Linux® cluster systems takes High Performance Technical Computing (HPTC) a step beyond conventional clusters. SiCortex concentrates on power-efficient design and simultaneous tuning of silicon, microcode, and system software to deliver outstanding application performance per dollar, per watt, and per square foot. The Company's product offering includes the following systems:

- The SC5832, which is a 5.8 Teraflop system with up to 8 Terabytes of memory. The SC5832 fits into a single cabinet and draws 18 KW of power.
- The SC1458, which is a 1458 Gigaflop system with up to 1.94 Terabytes of memory. The SC1458 fits in a single cabinet and draws 4 KW of power
- The SC648, which is a 648 Gigaflop system with up to 864 Gigabytes of memory. A single SC648 system draws 2 KW of power.
- The SC072, which is a 72 GigaFlop system with 48 Gigabytes of memory. The low-power desk-side cabinet uses less than 200 watts of power.

Abstract

This paper discusses the SiCortex computer system's integration and enhancement of open source software. In particular, it answers the following questions:

- How has open source software changed the field of technical computing?
- What advantages does open source software offer to customers and computer systems companies?

- How is SiCortex contributing to the evolution of open source software?

Introduction

Open source software has profoundly revolutionized the computer industry, particularly in the area of high-performance technical computing. Scientific and technical applications, which were traditionally the domain of specialized supercomputers, have moved to clusters of hundreds and thousands of commodity PCs. Open source software, usually based on the Linux operating system, has been a critical component of this shift because it can be adapted by its users to the special needs of technical computing. And, because both the hardware and the software are inexpensive, it has become cost effective to build large clusters out of such systems.

However, clusters of commodity PCs running open source software present challenges:

- From a hardware perspective, many of the components that comprise a commodity PC are unnecessary for the cluster. Furthermore, a high-speed, low-latency interconnect for the cluster can nearly double the price of the hardware.
- From a software perspective, assembling all of the pieces that you need to run on a cluster is a major project, and it's a project whose function is simply to make it possible to run the technical application.

Enter SiCortex. At a high level, our philosophy is simple: shrink the important parts of a compute node to a single chip, balance its design for maximum delivered performance on technical applications, and integrate the software into a coherent, supported foundation that gets more computing done, in less time, at lower cost.

The Open Source Revolution in Technical Computing

The most familiar name in Linux® clusters is *Beowulf*, an early project that used inexpensive PCs for technical computing. The original (1993) Beowulf concept was simple: use off-the-shelf, cheap PCs running Linux for parallel applications. Since then, the dramatic increase in the performance of commodity microprocessors means that clusters can be a very cost-effective solution to computation-intensive scientific problems. During the same period, open source software and Linux made rapid progress in functionality and stability, thereby providing an attractive software foundation for technical computing. In addition, programming interfaces such as the standardized Message Passing Interface (MPI) provide a powerful infrastructure that is consistent across different systems.

Of course, open source software is inexpensive to acquire. But for technical computing, there is an even more important advantage: long-term stability of the computing foundation. Technical codes have been ported repeatedly to succeeding generations of supercomputers. For a long time, acquiring almost any new high-performance computer system meant an entirely different operating system, a new set of compilers, and many new tools to learn. Porting an application was rarely an easy task.

Linux clusters have changed that. By providing a stable, yet evolving, set of programming interfaces, Linux lets you roll in new systems and run the same code, without needing to port it. The result, of course, is better performance with newer hardware.

More important, the Linux operating system is available for different processor architectures, hiding most of the differences in the underlying hardware.Therefore, using a different processor architecture no longer means changing operating systems and the way that you code. Instead, it means that you can take advantage of the performance offered by innovative hardware, as long as the programming interfaces are the same. Furthermore, the risk of a new system is small; applications are not locked into a single vendor's system because they can be recompiled easily for another Linux system.

SiCortex Hardware Is Designed to Run Linux Applications

SiCortex is the first *computer systems company* founded since the shift to open source computing. What do we mean by computer systems company? Twenty years ago, it was common for computer companies to design their own processors, write their own compilers and operating systems, and design and build their own systems. The advantage to the customer was the degree of integration in the whole product, and the ability of one company to design a coherent system in which the software could take full advantage of the hardware.

More recently, the computer hardware industry has fragmented into separate companies that make processors, make peripherals, assemble the pieces, and so on. The specialization has been useful in driving down costs and providing commodity computing, but often the beneficial effects of integration and complete system design are lost.

There have been other shifts in the hardware industry as well. In particular, system architects can now license component designs, such as memory controllers, that can be integrated to build a *system-on-a-chip*. This means that systems companies can take advantage of best-of-breed components, while still providing the integration essential to optimizing the whole system.

SiCortex takes this approach in its hardware design: select the best components, innovate and design where necessary to boost the performance of real-world applications, and integrate all of the pieces (except for the DRAM) for a compute node onto a single chip.

SiCortex systems are designed to run Linux. From the programmer's point of view, the SiCortex SC5832 system[1] is just another (very big) Linux cluster. Applications can be ported by recompiling. The hardware, and some of the software internals, have been tuned to run Linux applications fast.

The SiCortex Contribution to the Open Source Community

A fundamental view in the open source community is that the software is the critical infrastructure for computing. We share this view and also recognize that an innovative hardware design (that integrates the open system software model) makes it possible to develop new, powerful, and potentially ground-breaking applications that can solve High Performance Technical Computing (HPTC) challenges. Plus, the stable open source software APIs let

1. For more information about the SiCortex systems, see "*SiCortex Technical Summary*" at http://www.sicortex.com.

existing applications take advantage of our new hardware capabilities without code modifications.

Because our hardware is designed to run with open source software, we benefit from existing open source APIs. By contributing improvements to the open source core base, we support its evolution. By providing new hardware capabilities, we support the use of stable open software APIs that maximize the compute speed and power that SiCortex systems provide to HPTC applications.

Like the process of science itself, the evolution of open source software is driven by cooperation and a community that shares and improves upon its base. We're doing this today. We'll continue to do so.

Trademark and Copyright Information

SiCortex: The Open Source Cluster Systems Company