

## Starbridge Hypercomputers Complement Linux Clusters

Posted by [Craig Laurence](#) (Monday May 24 2004 @ 08:35PM EDT) [ ] Starbridge Systems, Inc. has demonstrated technology that can unlock secrets of the human genome. Using a 4U rack-mountable Hypercomputer, Starbridge recently compared the full X (147 million markers) and Y (60 million markers) human chromosomes, a process that took five days for the company's Hypercomputer to complete. Starbridge demonstrated the X/Y comparison as a requirement for purchase by the National Cancer Institute (NCI) in Frederick, Maryland where Dr. Jack Collins and others validated the results. According to Collins, another company performed this comparison on a custom architecture machine, but it required four months of computation time, and necessitated subdividing the data into more than 1000 smaller sequences, which then had to be stitched together again.

Upon delivery to NCI, Starbridge compared a gene sequence of one million markers of the rat genome to one million markers from the human genome. The computation took less than 2 minutes on a single Hypercomputer board. According to Collins, this comparison would typically have been analyzed using BLAST (a less sensitive search), and it would have required at least a day on one of NCI's clusters. NCI purchased a Hypercomputer and is now using it to research challenging, computationally intensive, problems. The goal of NCI is to facilitate genome-scale comparisons.

NCI and the field of bioinformatics aren't the only beneficiaries of Starbridge technology, though. Tricon Geophysics, Inc. and Essential Seismic Solutions (ESS) have turned to Hypercomputing to speed up, and reduce the cost of, the seismic imaging process. The National Aeronautics and Space Administration (NASA) at Langley and the National Security Agency (NSA) are using Starbridge Hypercomputers for research and national security applications.

The Hypercomputer is a reconfigurable computing system that uses field-programmable gate arrays (FPGAs) to deliver significant improvements in computational efficiency. The 4U rack-mountable, server-sized, Hypercomputer consumes the same amount of power as a high-end workstation. Ed McGarr, vice president of sales and customer services at Starbridge points out that "air conditioning costs, space requirements, and power consumption are significantly reduced relative to clustered systems that provide comparable computational capabilities." Starbridge claims Hypercomputers can be a complement to a Linux cluster or SMP environment as a computational co-processor. Driver calls from Windows and Linux environments pass data to and from the Hypercomputer, which acts as a computational co-processor that accelerates computation kernels for a multitude of applications. This allows companies and industries with large amounts of existing C or Fortran code to port only the computationally intensive kernel to the Hypercomputer, leaving the larger portion of code running on the Linux cluster.

Starbridge also provides an FPGA development environment called Viva®, which provides two unique functions: (1) A High Level Graphical Language that captures algorithms in highly parallel and efficient form, and (2) A Dynamic Hardware Synthesizer that customizes the FPGA hardware design for optimal performance of the intended algorithm. The key to Viva's strength lies in the fact that algorithm expression and hardware design are independent of each other, according to McGarr." This allows designers to express algorithms in their most efficient state, without having to force the algorithm to fit into a pre-existing hardware architecture; instead, you can design and synthesize the most optimal hardware architecture for your algorithm."

This is in contrast to the widely accepted and growing practice of building a large Linux cluster or shared-memory supercomputer for general purpose supercomputing and then adapting code to run on it. "Viva allows a developer to create the perfect architecture for an algorithm, instead of force-fitting the algorithm to existing, general-purpose architectures." The idea, claims McGarr, is that "the same piece of hardware can literally be used as custom hardware architecture optimally designed to maximize the performance of any given algorithm."

So what is the significance of Hypercomputing to the Linux cluster market? Rebecca Krull, vice president of marketing at Starbridge explains that with Hypercomputing, an organization can enjoy the general-purpose value of clusters while also benefiting from the increased computational speed of the Hypercomputer. Legacy code remains on the cluster, eliminating the prohibitive task of porting millions of lines of existing code. Execution speed is increased in a staged process, identifying and porting computationally-intensive portions of code one at a time. "With this approach, no wholesale re-write is required. Instead, you can continually tune your cluster by offloading strategic portions of code to the Hypercomputer. Essentially, you can have your cake and eat it, too."