Picking Up the Pace of Life Sciences Discovery

NEXT GENERATION ACCELERATION TECHNOLOGY SPEEDS RESEARCH AND REDUCES ELECTRICAL POWER DEMANDS

sgi®

www.sgi.com

mitrionics™

www.mitrionics.com

CONTENTS

Picking Up the Pace of Life Sciences Discovery

Genomics is the key to understanding how genes function and the roles they play and is at the heart of most biological and medical research today. From basic research to drug discovery, scientists increasingly are relying on computational analysis to speed up their genomics research. And given the competitive nature of this research in the life sciences, the need for speed is essential.

Academics doing pure basic research are rushing to publish and make that breakthrough scientific discovery before their peers. In biotech and pharmaceutical companies, competitors often are working on similar drugs to treat or cure the most common diseases. While it takes years to bring a new drug to market, if one company beats another to the patent office by even a single day, it can lock down the exclusive rights to sell a blockbuster drug for years.

That's an extreme example. More realistically, consider the lost revenue potential of a typical blockbuster drug. Every extra day it takes to bring such a drug to market means a day of lost sales which equates to $3M per day in revenue and $1M per day in profit. As a result, everyone is looking for a small edge that can speed up the drug discovery process.

So where do you need speed? Experts agree that many areas of basic biological research and all phases of the drug development process need to be accelerated. One area where many life sciences organizations are using new instrumentation and computing technology is in the drug discovery stage. In particular, many organizations today are investing in new generation of sequencing instruments that will generate up to 1TB of information per instrument per day. They are also investing in high throughput clusters to process that information and conduct much of their bioinformatics and genomic research.

While high-throughput clusters are good for many parts of the drug discovery workflow, they do have limitations. In particular, many algorithms like BLAST are extremely compute intensive and can consume thousands of hours of CPU

time for production runs. While one could simply implement a very large cluster with hundreds or thousands of CPUs, the IT organization would then be faced with a system management challenge, floor space constraints and dramatic increases in power consumption. In addition, applications like BLAST often work with very large datasets which either require the purchase of large amounts of memory per node or extended run times as data is moved in and out of system memory.

ACCELERATION ALTERNATIVES OFFER IMPROVEMENTS

Over the years, there have been a number of attempts based on a variety of technologies to speed up the most commonly used life science applications. Notable among the earlier efforts were Paracel's special high performance systems that used Application-Specific Integrated Circuits (ASICs) to accelerate a number of routines and Timelogic's Field-Programmable Gate Array (FPGA) accelerator boards that helped off-load computationally-intensive tasks from a CPU.

Many of the earlier efforts offered significant performance boosts, but each had their limitations. For instance, ASIC-based systems were often very expensive and lacked flexibility – they could only be used to accelerate just that one algorithm the ASIC was designed to run.

FPGA-based systems were typically hard to program and as a result many were sold as black boxes. These were devices pre-designed and pre-programmed to run perhaps one or a handful of specific applications. Unfortunately, this approach meant that a life sciences organization had to stick to the particular suite of applications the vendor supported. Essentially, this took the "programmability" out of FPGA.

Additionally, many of the earlier add-on products were not tightly integrated with the systems they ran in. This could result in data transfers overwhelming a system that had inadequate I/O or bus speed to keep up with the demands when application processing was accelerated.

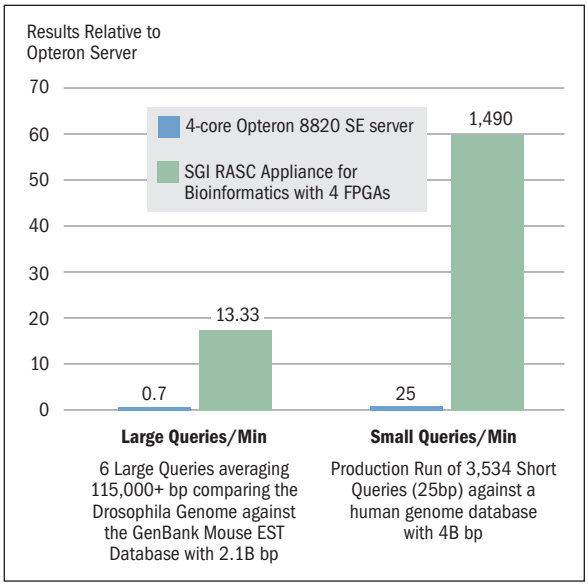

FIGURE 1: BLAST Performance on Large and Small Queries

TIME TO FAST FORWARD: BRINGING NEW TECHNOLOGY TO THE FOREFRONT

SGI has developed a next generation FPGA-based solution for the life sciences that overcomes many of the limitations of earlier systems. The SGI RASC Appliance for Bioinformatics provides the ideal mix of high-performance, an open and easy to program application environment, power efficiency, and simplified system management.

HIGH-PERFORMANCE:

SGI RASC solutions combine the high-performance SGI Altix architecture with leading-edge FPGA technology to deliver a complete platform for accelerating workloads. This solution can be used for a mix of CPU- based and FPGA-accelerated applications in a number of domains, or it can be optimally configured to run one set of applications within a specific domain. An example of an optimally configured solution is the SGI RASC Ap-

INTEGRATED FPGAS

pliance for Bioinformatics, which features Mitrion-accelerated BLAST-N.

From a systems perspective, the SGI RASC solution leverages numerous SGI Altix features like its scalable, high-bandwidth shared memory system architecture. This design enables systems to be configured with virtually any mixture of CPUs, FPGAs, memory and I/O so the configuration can be adapted to the needs of a specific customer. Single systems with over 5 TFLOPS of performance, 100 TB of memory and 10's of GBytes/sec of I/O can be created.

The SGI RASC blade is also optimized for high-performance. First, it integrates two of today's fastest FPGAs (the Xilinx Virtex 4) onto a single board. And because a single RASC blade with two FPGAs has two 6.4GByte/sec direct connections to the shared memory infrastructure, applications using the FPGAs can access enormous amounts of memory much faster than on other system designs. The combination of state-of-the-art FPGA performance, scalable high-performance system architecture and easy to use software development environments that make the most of the FPGA potential means that developers and users are able to achieve sizable real-life performance improvements.

With BLAST-N, performance improvements range from 16x on single queries comprised of 100,000+ base pair to over 60x on production runs of thousands of short queries with 25 base pair per query – demonstrating an extremely powerful performance advantage.

OPEN AND EASY TO PROGRAM:

In order to break out of the "black box" limitations imposed by previous FPGA solutions, SGI has taken an open approach to supporting application development. Applications can be written by developers or end-users in any one of several different application development environments and can be swapped out between application runs in a matter of milliseconds.

In the life-sciences space, Mitrionics has extend-

SGI RASC Appliance for Bioinformatics
8.25" high chassis for standard 19"

ed this open approach by creating the Mitrion-C Open Bio Project. This project recognizes that many critical bioinformatics applications are already open-source and that this encourages their wide-spread use and evolution. Mitrionics has already made great progress in this area. It has already released BLAST-N, has BLAST-P in development, and has TBLAST, Smith-Waterman and hidden-Markov method code under investigation for future implementations.

As with all other FPGA environments, SGI RASC solutions require special "bit-streams" to be downloaded to the FPGA. The traditional approach of writing applications required expert programmers who were familiar with VHDL specific design tools. This limited their deployments in bio-science environments to "black box" solutions.

For years, application developers have been searching for more standard higher level languages in which to write applications for FPGAs. Today's current crop of FPGA application development environments are very "C-like" with extensions that make it possible to take advantage of the parallelism present in FPGAs and allow them to link with the majority of pre-existing CPU based code. SGI and Mitrionics have already proven this approach by working on multiple applications from signal processing, encryption, seismic analysis and bioinformatics. The effect is that existing applications can be enhanced with FPGA capabilities without signifi-

cant reprogramming – shortening the time to solution. For the Mitrion-C implementation of BLAST-N, only about 0.2% of the code was modified to achieve this breakthrough performance.

POWER AND SPACE EFFICIENCY:

Increasingly, the power needed to run HPC clusters and cool data centers is becoming a real issue. Specifically, companies are faced with astronomical costs for data center power and cooling that are expected to grow significantly, all at a time when many companies are trying to be "greener."

In addition to pure acceleration of research results, the SGI RASC solution offers a significant power savings over a high-throughput cluster. These savings come in three ways.

First, the SGI RASC Appliance for Bioinformatics uses extremely efficient power supplies that deliver over 75% of input power to the system. This compares to a 50% average for 1U "pizza-box" servers, and means that from the start the SGI RASC Appliance consumes 1/3 less energy.

Second, FPGA based solutions are significantly faster than standard CPUs running bioinformatics applications. BLAST-N has been shown to be between 16 and 60 times faster than an Opteron CPUs. So a system need only run at peak performance for a shorter period of time.

Third, by consuming significantly less power to begin with, less energy needs to be expended to cool the data center.

If you combine these savings together and factor in the added processing power of the FPGA + CPU as compared to just a CPU, you end up with an overall energy usage for an FPGA solution that is between 90% and 95% less than an equivalent CPU-based only solution.

And this isn't just a theoretical analysis. SGI performed testing of an SGI RASC Appliance for Bioinformatics and a "pizza box server" system with two dual-core Opteron 8220 SE chips running at the maximum rate of 2.8GHz. The impact was startling, with both systems consuming approximately the same amount of power but with the SGI solution out-performing the Opteron solution by a factor of 16x to 60x. This means that the SGI solution was using between 2% and 7% of the energy per query of a high-throughput cluster node.

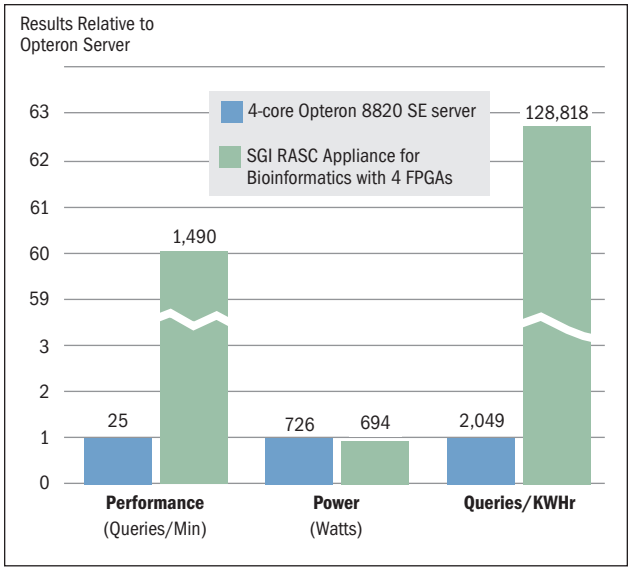

FIGURE 2: Relative Power Efficiency running 3,534 queries of 25 base-pair per query

Naturally, using less power cuts electric bills and is good for the environment, but the benefits don't stop there. The SGI RASC Appliance for Bioinformatics is also much denser than a standard high-throughput cluster, taking up only 6% to 22% of the rack space.

When you consider that industry surveys indicate 96% of computer rooms will run out of capacity within the next 5 years, denser systems that consume less power and are therefore easier to cool can clearly have a big impact on the bottom line.

SIMPLIFIED SYSTEM MANAGEMENT:

While tools for cluster management are improving, many challenges remain in extracting performance from these systems as they grow in size and complexity.

SGI

A single SGI RASC Appliance for Bioinformatics can have up to 24 FPGAs in a 35 inch tall chassis. For production workloads, this single system represents the equivalent of up to 1400 cores worth of a high-throughput Opteron cluster.

Perhaps more important is the fact that the SGI RASC Appliance is a single system. It is managed as a single system, has a single software installation, a single set of users and a single file system. The SGI RASC Appliance will take up to 1/10th of a system administrator while the high-throughput Opteron cluster could take several administrators.

SGI Extends a Rich History of Accelerating Scientific Research

SGI has long been an enabler of scientific research. Early Challenge and PowerChallenge systems brought high performance technical computing to the mass market. For example, in 1996, a 64-CPU Power Challenge Array was used by SGI and EMBL to analyze more than 6000 protein sequences from the genome of yeast (Saccharomyces cerevisiae).

SGI extended that mid-range leadership to the high-end by developing NUMAflex technology which enabled it to grow shared memory systems to up to 1,024 processors with 100s of GBytes/second of memory bandwidth.

As Linux became more commonly used in the scientific community, and particularly in the life sciences community, SGI became the first major vendor dedicated to providing open source solutions. SGI is currently shipping SGI Altix, its fourth generation of scalable shared memory system based on NUMAflex that can grow to 1,024 processor cores, 100 TB of memory and 10s of GBytes/second of I/O. Early life science adopters of the SGI Altix included the National Cancer Institute, the University of Arizona and the Memorial Sloan-Kettering Cancer Center.

In addition to its work on the systems side, SGI has a rich history of partnerships with mathematical analysis and database vendors working with them in the early days of Linux adoption to ensure scientists and researchers could get optimized performance when running their applications. Within the life sciences community, SGI worked with a wide variety of partners such as Gaussian, Schrodinger, SCM, and a number of open-source application providers to maximize performance on SGI systems.

The recent work with Mitrionics on the SGI RASC Appliance for Bioinformatics is an extension of SGI's long-time partnering efforts in the life sciences community.

Mitrionics Fills Acceleration Void

Mitrionics offers the Mitrion Platform that makes programming FPGAs much easier than has been the case in the past. The Mitrion Platform consists of several elements including the Mitrion Virtual Processor, the Mitrion-C programming language and the Mitrion Software Development Kit.

The Mitrion accelerated BLAST-N application uses the NCBI BLAST source code, however, about 1200 lines of computationally-intensive code of the BLAST-N program have been ported to Mitrion-C, to run on the Mitrion Virtual Processor.

This is a good example of the efficiency of the Mitrion Platform. Typically, when reworking a portion of an algorithm that can benefit from FPGA acceleration, a researcher might have to write a few dozen to perhaps a few hundred new lines of code. This could translate into thousands of lines of VHDL code. So in addition to the researcher not having to know VHDL, the Mitrion-C approach, being a higher-level programming language takes significantly fewer lines of code to produce the same result.

The first implication of Mitrion's enhanced software development environment is that it is faster and less complex to write applications and to use application accelerators than it used to be with older Paracel or TimeLogic systems. In addition, since the Mitrion developed BLAST-N software is part of the Mitrion-C Open Bio project, the modified source code is available on SourceForge.net for users that want to modify and build the code themselves, or application support is directly available through SGI.

In the SGI RASC Appliance for Bioinformatics, the initial version of Mitrion-Accelerated BLAST runs BLAST-N searches on a single FPGA up to 60x faster than a single core on a 2.8GHz Opteron 8820 SE processor. This tremendous boost in per stream performance coupled with the ability to have up to 16 FPGAs in a system turns batch analysis into interactive investigation, simplifying the search for drug targets and potential side-effects before reaching the lab bench.

In addition to the BLAST-N, Mitrionics plans to develop accelerated versions of other life sciences applications as part of the Mitrion-C Open Bio Project. For example, it is working on a BLAST-P program. Depending on the priorities of the user and developer communities – who are encouraged to participate in the development program – this could include the Translated BLAST programs, Smith-Waterman and Hidden Markov Models (based on Sean Eddy's HMMer package).

POWERFUL COMBINATION SIGNIFICANTLY CHANGES RESEARCH

From an IT or organizational perspective, the SGI RASC Appliance for Bioinformatics offers simplified management since it is an integrated system.

Additionally, to get the equivalent power using a cluster would require a very large installation, thus requiring more IT support, while taking up large amounts of data center floor space.

In contrast, the SGI RASC Appliance for Bioinformatics delivers the power required to speed genomics and bioinformatics research, making efficient use of data center rack space and consum-

MITRIONICS

ing a fraction of the electricity to operate and cool the system.

The true benefit of combing the energy-efficient, high-performance SGI RASC hardware with Mitrionics Mitrion Platform is that it puts supercomputing power in the hands of individual researchers.

In particular, the SGI RASC Appliance for Bioinformatics lets researchers run jobs without the involvement of IT. In contrast, such jobs would ordinarily need to be submitted to large operational clusters.

Along the same lines, the SGI RASC Appliance for Bioinformatics eases the process for researchers to make changes to their code to tap the processing power of FPGAs. By virtue of the Mitrion-C language and the Mitrion Virtual Processor, this type of operation can be done with minimal training and without the researcher needing to learn hardware-level programming as was the case in the past.

The SGI RASC Appliance for Bioinformatics therefore eliminates delays that would normally occur when a researcher has to wait for an IT department or dedicated programmer to make algorithm changes. This reduction in wait time, combined with the enormous speedups in delivered performance, means researchers can significantly accelerate their genomics investigations and consequently, their scientific and business achievements.

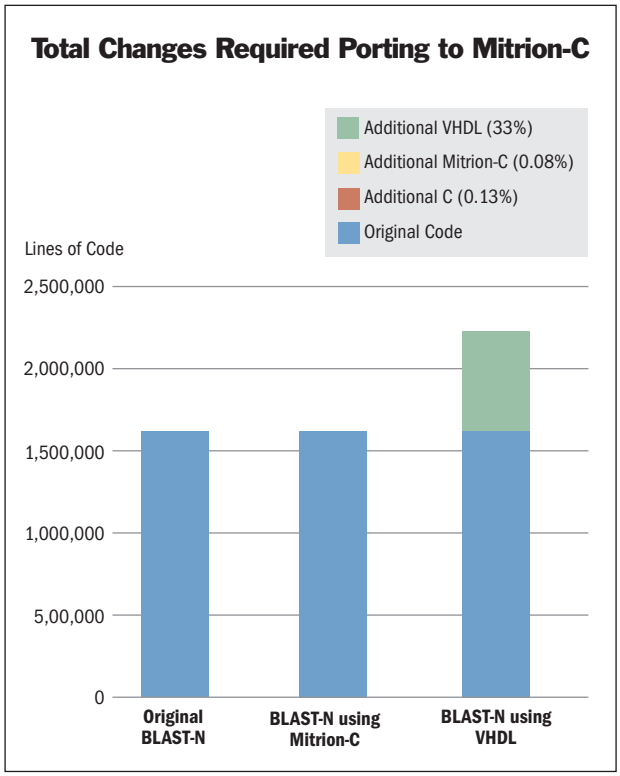

FIGURE 3: The Mitrion Virtual Processor, a massively parallel, soft-core processor, executes the software written in Mitrion-C on the FPGA hardware. What makes the Mitrion Virtual Processor unique is that it is adapted to the Mitrion-C code it is going to execute to make optimal use of the compute resources on the FPGA.