Powering the Search for Better Medicine

Georgia Tech chooses AMD processors for Biology Research Supercomputer

Challenges:

- Previous-generation high-performance computing (HPC) cluster with only 4,000 total processing cores was reaching end of lifecycle for conducting compute-intensive drug discovery research at the Center for the Study of Systems Biology
- Limited footprint – the same physical space of their previous-generation system
- Fixed power and cooling budget – 300 kilowatts of power and 80 tons of air conditioning

Solution:

- The next-generation "Myriad" cluster – an HPC solution designed by Penguin Computing, featuring Six-Core AMD Opteron™ processors (formerly codenamed "Istanbul"), with approximately 10,000 total processing cores

Impact:

- Outstanding processor density offers a massive performance increase over the previous generation HPC solution
- Energy-efficient design helps Systems Biology researchers at Georgia Tech operate within their power and cooling budgets

AMD
The future is fusion

"AMD Opteron processors were in the sweet spot of power density, speed-per-job, and aggregate throughput."

— Dr. Jeffrey Skolnick, Director for the Center for the Study of Systems Biology, Georgia Institute of Technology

Universities have a long history of using high-performance computing (HPC) for research and analysis in engineering, physics and mathematics. In recent years, however, high-performance computing clusters have also revolutionized traditionally descriptive sciences such as biology, with the ability to dedicate the power of literally thousands of processing cores to simulations that seek to unlock some of the deepest secrets of our bodies.

Georgia Institute of Technology, better known as Georgia Tech, is a top-five Engineering school.[1] Their Computer Science program ranks in the top ten programs nationally as well.[2] And it's gaining a reputation in Biology as well – thanks in great part to the work of Dr. Jeffrey Skolnick and the Center for the Study of Systems Biology.

Dr. Skolnick came to Georgia Tech in 2006 to conduct compute-intensive research with the assistance of HPC in the field of drug discovery. By studying and simulating the structures of different proteins within the human body, Dr. Skolnick's research strives to accelerate the process of finding new drugs that are better targeted cures with fewer unwanted side effects.

"Whether it's biology or engineering – computation is an extraordinarily invaluable tool. People want to do data analysis, people want to design experiments," says Skolnick. "When you can do simulations and high-throughput data analysis and discovery, it allows you to do a lot of virtual 'what if' scenarios that can accelerate the real-world experimental side."

"Since I have essentially an infinite appetite for an infinite number of computers, for me to be able to do my research, Georgia Tech had to establish a very substantial HPC configuration," says Skolnick. His first-generation HPC system was a 4,000 core cluster named "Razor," built with AMD Opteron™ processors.

By 2009, the Razor cluster was nearing the end of its lifecycle. Thanks to advances in processor design, a next-generation HPC system could offer Dr. Skolnick and his researchers much higher processing core density to help them do more calculations each day.

The new solution, however, would need to exist within rigid limitations in terms of physical space (the second-generation HPC would need to be installed in the same room occupied by Razor), as well as available power (300 kilowatts) and cooling capacity (80 tons of air conditioning).

"Over the years, Georgia Tech, as a lot of universities, has realized that a significant cost is the heating and cooling of these beasts," says Skolnick. "We have fixed square footage, fixed amount of cooling, and fixed amount of power. It was a non-moveable budget."

Higher core density for greater performance

At the Center for the Study of Systems Biology, time is an extremely valuable commodity. "Having a very large number of cores enables us to test a very large number of things simultaneously. So we can quickly dismiss things that are not promising and try to pick up those very few threads which look promising," says Skolnick.

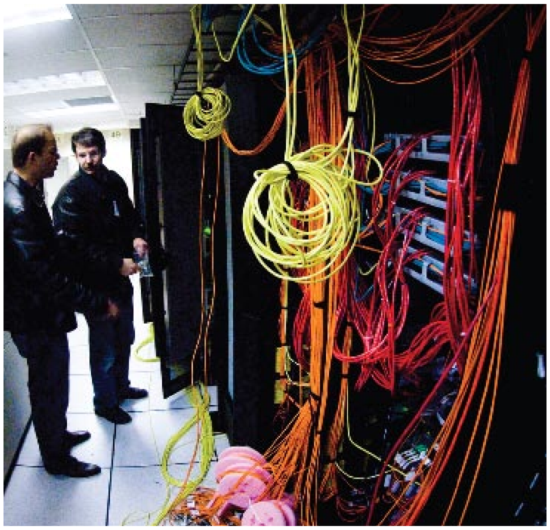

The main challenge that Dr. Skolnick and his department were looking to overcome with their next-generation HPC cluster was increasing performance. Advances in multi-core processor design would allow them to more than double the number of cores in their system, and hopefully, a similar increase in the volume of research they could do each day.

The first step in getting a new HPC system for the Center for the Study of Systems Biology was selecting a processor. Dr. Skolnick and his team did extensive evaluations running their codes on different multi-core processors for performance, efficiency, and cluster configuration capabilities.

"Whatever you choose, if you have 24 cores or an infinite number of cores, you want to keep them all full. Because if the beasts aren't full, it's not cost effective," says Skolnick.

As a result of their testing, Dr. Skolnick's team selected Six-Core AMD Opteron™ processors as the foundation for their new Myriad cluster. "In terms of the per-core performance, the Six-Core AMD Opteron processors seem to perform much better than our previous generation. And we like the fact that we can pack all 24 cores-per-node with jobs and they all run at extremely high utilization."

The amount of computing power per dollar was also a factor in Georgia Tech's selection of AMD processors. "Price-performance ultimately counts. Especially in this day and age, the university doesn't have an infinite amount of money. And we needed to maximize the cost effectiveness," says Skolnick.

The winning proposal for the Myriad cluster was made by Penguin Computing, who completed the final design, system configuration, and physical installation in collaboration with Dr. Bartosz Ilkowski, Senior Research Technologist and the Head of Cluster Computing at the Center for the Study of Systems Biology.

"The Myriad cluster is much denser. At a quarter of the compute node count, it packs over two times as many cores," says Ilkowski. "Hardware-wise it is still a commodity Linux HPC system, but it was built for and optimized according to the requirements of our production applications, rather than being a generic HPC system. That allowed us to obtain significant performance gains."

The installation process itself presented its own challenges. The new Myriad cluster would need to be installed in the same room that housed the previous-generation Razor cluster. The scientists at Georgia Tech, however, needed to continue their research seamlessly through the installation process.

"We had to figure out a way to keep both of them working," says Skolnick. "I couldn't have my research come to a screeching halt while the new stuff is being brought in and the old stuff is being retired. So they had to basically create a transition so that we did not have significant downtime. And Penguin was very successful at doing that. We were very happy."

Energy-efficient design for limited power and cooling

Physical space wasn't the only limiting factor for the new Myriad cluster. While going from 4,000 to 10,000 cores in the same footprint could offer The Center for the Study of Systems Biology a significant advantage in terms of aggregate throughput and increased work done each day, more than double the processing cores could also potentially increase their HPC power consumption.

"I have 300 kilowatts of power. So energy-efficiency is extremely important because we have to live within that power cap. Even if I wanted to add a million processors – they would have to fit within the 300 kilowatt cap," says Skolnick.

The energy-efficient design of AMD Opteron™ processors offered Georgia Tech an HPC system that could live within their power usage limits. And while offering a much denser core count, Myriad wouldn't need additional cooling beyond what was required for their previous generation cluster.

"We only have 80 tons of air conditioning to cool the thing. So power consumption and heat generation were extremely important to the decision because it dictated how many processors I could fit into the square footage," says Skolnick.

"I feel very much that it's an ongoing partnership with Penguin and AMD. It's not that they just drop off the computers on the loading dock and say 'I'll see you in three years.' They have to be involved."

Reliable partners with an eye on the future

With limited resources for system maintenance, reliability and manageability both figured into Georgia Tech's decision to go with the AMD/Penguin solution. "I have one talented person, Dr. Bartosz Ilkowski, handling these 10,000 cores. I don't have the budget of a very large computing center that has 50 people caring for and feeding this thing," says Skolnick.

As with the design and installation, ongoing care for the Myriad cluster is extremely collaborative. Penguin computing works very closely with Dr. Ilkowski to adjust and maintain a highly successful, highly productive HPC system. In the process of developing and configuring the Myriad cluster, Dr. Ilkowski even invented a couple of new products for Penguin.

As an extension of this successful partnership, AMD and Penguin are also working with Georgia Tech to test future computing solutions for the Center for the Study of Systems Biology, including the next generation of multi-core AMD Opteron™ processors as well as an AMD processor-based workstation.

"We're benchmarking now the AMD Opteron™ 6100 Series processors, AMD's 12 core. So far, we've been very happy. They've been very reliable. Nothing but good things to say," says Skolnick.

"I feel very much that it's an ongoing partnership with Penguin and AMD. It's not that they just drop off the computers on the loading dock and say 'I'll see you in three years.' They have to, for the critical things, be involved. People like Bartosz will come up with very constructive suggestions for driving their technology, and it also helps us. And so I view it as a very mutually synergistic relationship," says Skolnick.

In terms of the future of high-performance computing, Dr. Skolnick sees advantages in multiple types of processing cores co-existing on the same chip. "Different types of algorithms take advantage of different types of processors. For graphics rendering and certain types of things, a GPU is great. For other things, a CPU is great. I suspect what's going to happen is that there's going to be a convergence."

The upcoming AMD Fusion™ Family of Accelerated Processing Units (or APUs) combines both CPU and GPU processing in the same chip. AMD's APUs also include a variety of critical system elements, including memory controllers, I/O controllers, specialized video decoders, display outputs, and bus interfaces, but the real appeal of these chips stems from the inclusion of both scalar and vector hardware as full-fledged processing elements.

Dr. Skolnick sees potential for this type of converged chip design to benefit HPC computing and the demands of compute-intensive scientific applications, where the software can select the appropriate type processing core to run a specific algorithm more efficiently.

"That's how we're going to see tremendous acceleration in the calculation, which in the end is all you really care about. If you're talking about next-generation, game-changing technologies, that's what has to happen," says Skolnick.

About AMD

Advanced Micro Devices (NYSE: AMD) is an innovative technology company dedicated to collaborating with customers and technology partners to ignite the next generation of computing and graphics solutions at work, home and play.

For more information, visit http://www.amd.com

1 GIT Engineering Dept rankings: http://en.wikipedia.org/wiki/Georgia_Institute_of_Technology#Rankings
2 GIT Computer Science Dept rankings: http://pages.cs.wisc.edu/~remzi/rank.html

The future is fusion