A NEW APPROACH TO DATA INTENSIVE COMPUTING

CONVEY PRESS BACKGROUNDER

CONVEY computer™

The World's First Hybrid-Core Computer.

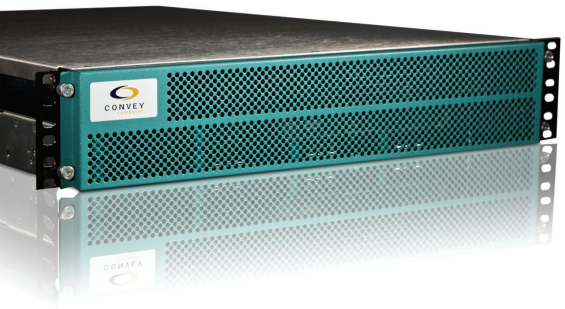

>> Convey's Hybrid-Core Computer in Wonderland:

A New Approach to Data Intensive Computing

'Curiouser and curiouser!' cried Alice (she was so much surprised, that for the moment she quite forgot how to speak good English); 'now I'm opening out like the largest telescope that ever was!'

This kind of uncontrolled growth isn't just confined to Wonderland – the world of high performance computing (HPC) seems to have nibbled on the same magic cake that caused Alice to shoot up like a rocket.

Today, the amount of data that needs to be gathered, analyzed and stored is growing exponentially. Even the most powerful supercomputers and HPC clusters are struggling to keep up.

The problem? Data intensive computing. When contemplating the rapidly growing deluge of data, Steve Wallach, HPC guru, chief scientist and co-founder of Convey Computer, likes to quote Yogi Berra, not Lewis Carroll. "It's like déjà vu all over again," he says.

What Goes Around Comes Around.

A lot of the problems that surfaced 30 years ago are still with us today, he claims. "It's still common for technology buffs to get excited by a new, peak performance micro-architecture," Wallach says. "But once you get past the first flush of enthusiasm, you have to start asking questions about the viability of the supporting software, how well the technology fits into your IT and overall business environment, and whether or not, given your particular set of applications, you can actually realize the potential of the new hardware."

He notes that supercomputer companies like Convex and Cray succeeded because they had a robust, integrated software environment that could be leveraged by programmers working with familiar languages such as FORTRAN and C. Companies with a weak development environment saw their HPC efforts falter and die.

So here we are, three decades later. Despite incredible advances in microprocessor technology price/performance – especially off-the-shelf processors – we are bumping up against both hardware and software problems resulting from the rapid growth of data intensive computing requirements.

The U.S. Department of Energy's Pacific Northwest National Laboratory has a succinct definition of data intensive computing (DIC): "DIC is capturing, managing, analyzing and understanding data at volumes that push the frontiers of current technologies."

Powerful supercomputers and clusters, combined with inexpensive data storage and abundant bandwidth, are churning out huge datasets from modeling and simulation applications, advanced instruments, and system sensors. The challenge, says the Lab, is to manage the explosion of data; extract knowledge from these massive datasets; and reduce the data to facilitate human understanding and response. This is a worthy goal for any organization that relies on HPC – whether it's the government, academia, or business.

HPC's Perfect Storm.

But we are in the middle of what Bruce Toal, Convey's president and co-founder, calls "a perfect storm of technology trends." For example:

- The pressure is on to boost application performance using hardware rather than transferring the burden to programmers (who can cost more annually than the price of a teraflop, high-end HPC system).
- After a decade and a half of dramatic increases, today's single-socket commodity processor clock speeds are flattening—even falling.
- Commodity processors are memory bound. Datasets are either too big, too complicated – or both – to be efficiently pumped through commodity x86 processors.
- Power and cooling requirements are off the charts. It can cost more to power an HPC machine than to buy one. In a modern multi-million dollar data center, power and cooling costs can account for half the operating expenses.

Looking for a Solution.

One tactic to speed up your applications and attempt to cope with data intensive computing is to add more cores to your HPC system. More cores, more speed. But it's not that easy. The downside is that programming becomes even more difficult (consider massively parallel code here), and energy costs can become prohibitive. More processors mean more power, more heat.

One solution has been to enlist the power of specialty processors, known as accelerators. Convey software architect John Leidel calls these custom integrated circuits the Ferrari's of the of the HPC world. "They are very fast in small spurts," he says, "but they aren't designed for the long haul." Like Wallach, he points out the difficulties associated with programming these components and integrating them into the HPC host environment. As Wallach says, "A lot of new acceleration technologies are evolving, but are not dealing with the software environment."

So, what to do? Hybrid-Core Computing to the Rescue.

Like most elegant solutions, the Convey answer to the dilemmas surrounding data intensive computing is quite simple and straightforward. It acknowledges the futility of trying to cope with the deluge of data by adding hundreds, thousands, or even millions of processors to your HPC systems. Instead, the Convey solution focuses on the lowest common denominator – the individual processor – and makes it run faster.

"Since clock rates are basically flat," says Tony Brewer, Convey chief technology officer and a company co-founder, "you can speed up the individual processor by having one instruction replace 100 or more instructions required by conventional components. The Convey HC-1™ makes tremendous gains in application performance possible by combining an x86 processor with hardware that implements application-specific instructions. Instead of stacking up more pizza box servers in the data center and watching your power, cooling and cabling costs skyrocket, you can get a lot more performance while actually reducing your hardware, energy and cabling expenditures. When I say more performance we are talking about orders of magnitude, not just an incremental tweak or two. We are teaching the x86 new tricks and focusing on existing HPC problems that nobody's found a good way to solve."

Bruce Toal adds, "For key HPC workloads, the Convey HC-1 can drastically reduce the number of servers required. One rack of our servers can replace multiple racks of traditional servers. This means you can realize dramatically lower energy costs. For example, recently one of our customers cut their power and cooling costs by 91%. With hybrid-core computing, they are finding it easy to be green."

The Convey server uses a coprocessor based on standard Field Programmable Gate Arrays (FPGAs). The coprocessor is coupled with a standard, multi-core Intel® Xeon® processor and provides processing elements optimized for data intensive applications.

Because the Convey approach uses ANSI standard C, C++, and FORTRAN languages, from a programmer's point of view working with the hybrid-core computer is exactly like developing in the familiar x86 environment. And there is no need to learn a new parallel computing architecture or language. Also, the host processor and coprocessor share the same cache-coherent view of virtual memory – the FPGA coprocessor is treated like just another processor on the system bus; this too makes the developer's life easier. Applications written in any of the three ANSI-standard languages can be compiled on the Convey open source compiler stack.

Key to the success of the hybrid-core computing model is that it supports multiple instruction sets in a single address space. Within this common space, the off-the-shelf x86 processors execute "normal" (x86_64) instructions, while the coprocessor executes the application-specific instructions.

Financial Analytics Personality Joins the Fray

In the fiercely competitive world of financial services, analysts are using increasing complex financial instruments to get a jump on their competition. As a result, HPC systems built from off-the-shelf commodity components are feeling the strain.

"One of the major problems facing developers of financial applications is that the instruction set of commodity processors doesn't map well to many financial algorithms," explains Tony Brewer, Convey co-founder and chief technology officer. "Even a simple intrinsic function can take hundreds of instructions and, in many of these applications, that intrinsic might be executed millions of times."

This is why Convey has developed a financial analytics personality (FAP). The financial instruction set executes intrinsic operations directly in hardware, greatly accelerating financial algorithms.

By providing more performance per watt, a single rack of Convey HC-1 servers can replace multiple racks of conventional servers, reduce power, decrease floor space, and lower the chance of hardware failure.

Personality Plus.

And this is what really sets the Convey solution apart – its unique personality. Or, to be more exact, its multiple unique personalities.

Personalities are reloadable, application-specific instruction sets implemented in hardware that run on the Convey FPGA coprocessor (see Figure 1). These are bundles of functions and functionalities that are not found on your garden-variety x86 processor.

Personalities are flexible. And they are blank slates. They can range from providing instruction-level acceleration – for example, emulating a standard vector processing programming model – to highly complex algorithms for a wide variety of applications.

As John Leidel points out, "If we need a new personality, we just invent it. We also provide a personality development kit (PDK) that allows our customers to build personalities tailored to their specific needs."

The personality can replace thousands of interactions of standard x86 instructions for a particular application space, all with full FORTRAN, C, and C++ support.

Convey has a number of personalities already available and others in the works. For example, one of Convey personalities implements a vector processing instruction set similar to those found on vector supercomputers. The Convey compiler detects opportunities for vectorization within a source file and generates vector instructions that will execute on the coprocessor. This personality substantially accelerates the processing of long pieces of code with nested loops.

Another Convey personality has been designed specifically for the financial industry. The FAP (financial analytics personality) significantly enhances key financial algorithms such as random number generation and math intrinsics specific to financial analytics (see sidebar: "Financial Analytic Personality Joins the Fray").

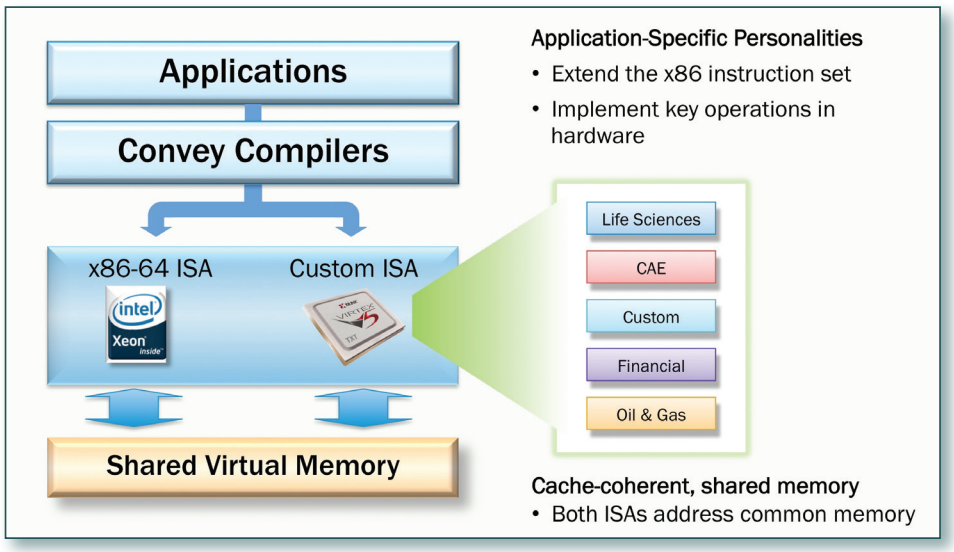

Figure 1.
Personalities on the Convey HC-1 are application specific.

A New Understanding of How Disease Impacts our Cells

The task that researchers at the University of California, San Diego have undertaken – a study of an organism's complete complement of proteins – is complicated and compute-intensive but worth the effort. By understanding how diseases change the proteins in our cells, researchers can help find new diagnostic markers and drug targets in the quest to effectively treat diseases such as Alzheimer's.

Proteins are made up of very large numbers of very small amino acids. Researchers need to identify and catalog the amino acid sequences in the various proteins in different cells in the body. They can then compare the differences between healthy and diseased cells. Not only is the number of requisite comparisons huge, but the data base is a moving target – the proteins keep on changing, a process known as post-translation modifications (PTMs). Until now, PTMs were too complex to be included in the comparisons – a major but inescapable oversight.

The solution came in the form of an application developed at the University known as InsPecT/MS-Alignment, a highly compute-intensive comparison and scoring algorithm. Working with the San Diego researchers, Convey expressed this algorithm as a Convey HC-1 personality. The result? Run times have been accelerated by several orders of magnitude and the elusive PTMs are now part of the comparison and scoring process. An added bonus: One rack of HC-1 servers has replaced eight racks of commodity servers.

A personality has also been developed to assist researchers in the field of proteomics– the study of an organism's complete complement of proteins, one of the most important field of research in the life sciences (see sidebar: "A New Understanding of How Disease Impacts Our Cells"). Scientists at the Computer Science Department at the University of California, San Diego, have developed an application that allows researchers to perform protein database searches that were previously impractical or impossible. The HC-1 personality not only speeds up the time to complete an application, but also reduces the number of servers required and the associated infrastructure costs.

How Hybrid-Core Computing Can Benefit You

One of the main features of the Convey HC-1 platform is its flexibility. You can run all your traditional applications on the hybrid-core system at normal clock speeds. But when you need several orders of magnitude more horsepower to execute a massive data intensive application, the computational capabilities you need are instantly available. You don't have to constantly upgrade or replace your systems to tackle jobs with varying compute requirements. In IT jargon, you have achieved infrastructure reinvestment and reuse. At the same time, by implementing substantially higher per-node performance, your data center requires far less in the way of power, cooling and floor space to meet your computational requirements.

Says Brewer, "When you come up with an approach that makes it possible to boost performance by an order of magnitude or more while using the same or fewer processors, you have a winning combination. Hybrid-core computing is a technology whose time has come."

The moral of the story is clear: Tackling today's data intensive computing environment using clusters and supercomputers powered solely by conventional x86 components is like trying to play croquet using a flamingo for a mallet. Eventually you'll get the job done. But with the right equipment, you'll be a lot more successful.

Convey Computer Corporation
1302 East Collins Boulevard
Richardson, Texas 75081

Phone: 214.666.6024
Fax: 214.576.9848
Toll Free: 866.338.1768

www.conveycomputer.com

Convey's hybrid-core computing solution is the answer.

CONVEY computer™

Convey Computer, the Convey logo, and Convey HC-1 are trademarks of Convey Computer Corporation in the U.S. and other countries. Intel® and Intel® Xeon® are registered trademarks of Intel Corporation in the U.S. and other countries. Xilinx is a registered trademark of Xilinx, Inc.