

## Hybrid-Core Computing

# Convey Computer Tames Data Deluge

By John Russell

Nothing illustrates the data deluge inundating life sciences better than the staggering output from Next-Generation DNA Sequencers (NGS). Consider that the Illumina HiSeq 2000 — just one of several revolutionary NGS instruments appearing on the market — generates 200 Gigabases of DNA sequence per 8-day run and that both the Wellcome Trust Sanger Institute and the Broad Institute can produce more than 1 Terabase of sequence data per week.

In turn, this data flood poses an enormous computational challenge. BGI (the former Beijing Genomics Institute) when fully provisioned by the end of 2010 will have 164 NGS instruments producing 5 terabytes (TB) of data per day. The data will be processed by 50,000 CPUs (capable of 1000 TFLOPS) with 200 TB of RAM and 1000 TB of storage (1 Petabyte!!).

All across life sciences research, large datasets and rigorous computational demands are the rule, not the exception. Richardson, TX-based Convey Computer understands these challenges. Founded in 2008 by a group of distinguished high-performance computing (HPC) [executives](#) — including IEEE Seymour Cray Award winner, Steven Wallach — Convey Computer's mission is to develop innovative, practical approaches for improving data-intensive computing.

Traditionally, brute force has been the only choice for dealing with the data deluge, generally by throwing ever more racks of commodity servers at the problem. But




Figure 1

this solution adds programming complexity and sends power, cooling, and space requirements and associated costs spiraling higher. (See Figure 1.)

Convey has pioneered a new approach. Its innovative architecture pairs classic Intel® x86 microprocessors with a coprocessor comprised of FPGAs<sup>1</sup> to create the

world's first hybrid-core computer, the Convey HC-1. Particular algorithms — DNA sequence alignment, for example — are optimized and translated into code that's loadable onto the FPGAs at runtime to accelerate the applications that use them. Convey calls these accelerated algorithms "personalities." Personalities and the HC-1 system reduce the complexity around programming applications for the FPGA.

Broadly speaking, personalities are extensions to the x86 instruction set that are implemented in hardware and optimize performance of specific portions of an application. "Basically, this creates hardware that is unique to the algorithm you are running," explains Dr. George Vacek, manager of Convey's Life Sciences business unit. "There is no better way to make things faster than to wire them down into gates on a chip." When an application is running, the extended instructions created specifically for that application are dispatched to the coprocessor. The personality needed for each program is loaded at run-time to reconfigure the coprocessor with optimized instructions for that specific application. The result is flexible, high-performance, application-specific computing that is also cost efficient: "You typically can replace between 5 and 25 traditional servers with a single one of our servers."

Through Convey's proprietary adaptive architecture, servers can dynamically and transparently load different personalities — each optimized for different applications. The ability to adapt the architecture to different applications means the Convey servers can be repurposed "on the fly"— making them extremely flexible in mixed-use environments.

No less important, Convey systems are easy for developers to use because they provide full support of ANSI standard C, C++, and Fortran. The compilers recognize constructs in the source code that generate instructions that will execute on the coprocessor. Alternatively, developers can add—using standard compiler syntax—functions or intrinsics in their source code that will dispatch complete sections of code to the FPGAs.

"Convey systems are 2U rack-mountable systems and run industry-standard Linux. They cluster like other off the shelf servers, are managed like other servers, and connect into your HPC fabric like other servers. The point is, to the outside world, there's nothing exotic about the HC-1. It's just blindingly faster on bioinformatics algorithms than any off-the-shelf server," says Dr. Vacek.

Though still young — the HC-1 has been shipping since June 2009 — Convey's hybrid-core computers

have achieved impressive performance on a variety of bioinformatics applications, allowing researchers to tackle problems previously deemed impractical. For example:

- **Sequencing.** The [Convey implementation](#) of the Smith-Waterman algorithm (used for aligning DNA and protein sequences) is 172x faster than the best software implementation on conventional servers and represents the fastest Smith-Waterman implementation on a single system to date.<sup>2</sup>
- **Proteomics.** University of California, San Diego (UCSD) researchers achieved a roughly 100-fold faster performance of their sophisticated MS/MS database search tool program — [InsPecT](#) — that is able to accurately identify post-translational modifications (PTM).
- **Computational Phylogeny Inference.** The University of South Carolina developed and accelerated MrBayes, a [phylogenetics application](#) able to accurately infer "evolutionary trees," a problem that was previously considered impractical on most computer systems. Performance is significantly faster even than other FPGA implementations.
- **Genomics.** The [Virginia Bioinformatics Institute](#) (VBI) is using Convey hybrid-core systems for its microsatellite analysis work for the [1000 Genomes Project](#), an international effort to sequence the genomes of approximately 2,500 people from about 20 populations around the world.

These achievements reflect a fundamental change in the scientific computing landscape. "A very important thing has happened in the last couple of years," says Dr. Harold "Skip" Garner, executive director of VBI, Professor in the Departments of Biological Sciences and Computer Science at Virginia Tech, and Professor in the Department of Basic Science at the Virginia Tech Carilion School of Medicine. "We've crossed over from data acquisition being the most expensive [part of research] followed by analysis to just the opposite. Data acquisition is getting cheaper all the time and analysis is becoming more complex and more expensive because of the enormous amount of data, the kinds of problems attacked, and the fact that a lot of the low hanging fruit is gone."

Scaling up standard microprocessor systems to cope with the new challenges won't work, he says. "We just don't have enough electricity, cooling, floor space, money, etc. for using standard clusters or parallel processing to handle the load.

"As a consequence there have been a couple of things that are paradigm shifting. One is the FPGA.

The other is the new general-purpose computation on graphics hardware approach. It is my opinion that future bioinformatics centers, especially large computing centers, will have a mix of technologies in hardware that include standard processors, FPGAs, and GPGPUs, and new jobs will be designed for, implemented on and steered to the appropriate processing environments."

VBI currently has Convey HC-1 machines pounding away on various projects that seek to develop personalities or tackle important science projects such as its work on the 1000 Genomes Project. "Think of us as a super user and Convey collaborator," says Dr. Garner. "We're not the only ones. Convey has gone out seeking early adopters in many science and engineering disciplines.

"Until Convey came along, it required very specialized, expensive programmers to use FPGAs," Dr. Garner says. Convey's architecture and toolkits "allow you to program in more of a traditional environment to create personalities on the FPGAs. This can be done by people who are not absolutely specialists in the sense they used to have to be." Convey's current life-sciences customers comprise a mix of academic, agricultural, genomic, and medical research centers.

## CONVEY TECHNOLOGY

Of course, efforts to accelerate bioinformatics are not new, but they have tended to be I/O board-based solutions. Such solutions typically use I/O slots and, while they certainly achieve some measure of acceleration, they also encounter hard limits around how fast data can be fed to the board.

Convey's hybrid-core technology elegantly solves the data bandwidth issue by providing a cache coherent, shared view of virtual memory between the host and the coprocessor. Physically, the coprocessor supports a large, high-bandwidth (80 GB/s) memory subsystem. That translates into plenty of headroom for accelerating algorithms.

In total, there are 14 FPGAs on the Convey coprocessor (see Figure 2): Four FPGAs serve as the application engines that execute the personality's instruction set; two FPGAs comprise the application engine hub that handles communication to and from the host x86 processor; and eight FPGA-based memory controllers provide the very fast memory.

"Another important feature is the direct data port (DDP), which is basically direct I/O pin access to the application engines," explains Dr. Vacek. "When you need more bandwidth than the frontside bus, the

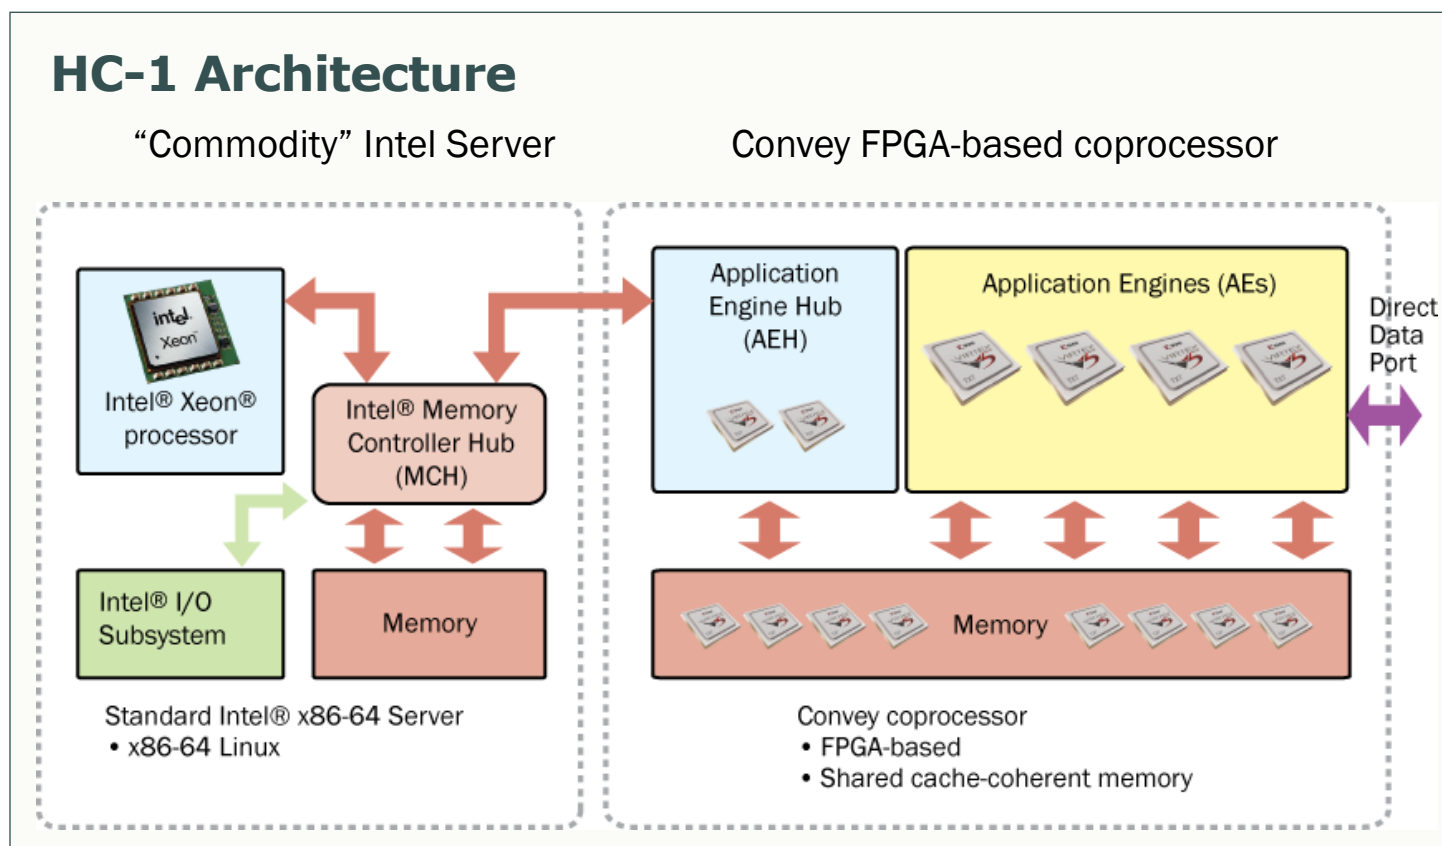


Figure 2

DDP can be used to directly transfer external data to the logic on the FPGAs." For example, signals from sequencing instruments can be piped directly into the application engine for primary analysis. This is important for keeping up with the data-generation rate of the sensors in order to process the signal into base calls in real-time.

Some have argued that even though the growth in life-sciences data is overwhelming microprocessor advances, the fact that bioinformatics data is so parallelizable means you can simply keep adding racks and racks of traditional processors to handle the load. But if you actually look at the economics of the data center, emphasizes Dr. Vacek, that doesn't make sense. It's already the case that it costs more to power and cool a server for a year than to buy the server.

"It's not a question of acquisition costs; it's a question of constraints on power, cooling and floor space. Going forward we see hybrid-core computing as a way to overcome these obstacles," says Dr. Vacek. "You save a lot of money when you can take eight racks of commodity servers and reduce them to one rack of Convey systems."

## BIOINFORMATICS ADVANTAGES

In February 2010, Convey launched a dedicated Life Sciences Business Unit, led by Dr. Vacek. "We're focusing on bioinformatics because it's such a good fit for our architecture. There are a couple of reasons for this. One is that most applications are very data parallel and very amenable to pipelining and parallelism in hardware. Another reason is the data types are typically small."

In DNA sequencing, for example, "there are only four nucleotides, which can be represented in just two bits. Because FPGAs aren't hard-wired with any particular size of data element, it means we can do two-bit operations rather than say 32-bit integer calculations, which require more memory, memory bandwidth, and computational power. We're able to make much more efficient use of the hardware to achieve higher orders of parallelization and resulting higher throughput," Dr. Vacek says.

NGS presents another challenge: building up genomes from many short reads using traditional assemblers is difficult, and the new *de novo* assembly models, such as [Velvet](#) and [ABYSS](#), are based on de Bruijn graph algorithms which are memory bound on traditional architectures.

"On an x86 system, when you access a single 64-bit word in memory, you still have to pull a full cache line

of 8 words (64 bytes) at a time, which dramatically reduces net memory bandwidth for random access patterns. Graph algorithms tend not to be compute-constrained but memory-constrained," says Dr. Vacek. "Our system provides very efficient random access to memory. Based on that alone, we should be able to accelerate these graph algorithms significantly and not just for short read assembly but also for other methods, such as systems biology and protein interaction maps."

Convey is building up a portfolio of personalities. For example, they currently offer a Proteomics Personality that implements the kernel of the UCSD InsPecT proteomics search application. Work is also progressing well on a Sequence and Alignment Suite. However, Dr. Vacek emphasizes the HC-1 is a platform on which anyone can develop personalities. "We are talking to software vendors who are creating personalities for the platform. People who buy our platform don't just get the set of whatever personalities we've created, but they can get whatever anyone else develops for it."

## DEPLOYMENT PROCESS

There are two major paths for deploying Convey systems. A customer can load a personality that has already been developed — the InsPecT algorithm, for example. As noted earlier, personalities are basically an instruction set architecture (ISA) that extends the x86 instruction set. When an application is compiled, the compiler generates low-level instructions that can either run on the host or on the coprocessor.

"You can take the source code and do anything with it as long as the application eventually executes the core instructions that invoke the personality on the coprocessor. It's one of the advantages of the Convey system — you're able to incorporate personalities into your application however you see fit. It's not a black box where you are stuck with a fixed set of features," says Dr. Vacek.

Competitive offerings, says Dr. Vacek, tend to be hard-wired, "black box" approaches in which users must adapt their workflow to the solution. With the Convey HC-1, you can work the personality into your workflow.

The second path is when there is no personality implemented yet for the desired task. Personality development "is a multi-step process that involves essentially synthesizing your algorithm at a high level, then at the gate level in some kind of hardware definition language, and deploying it on the HC-1 coprocessor as a personality," explains Dr. Vacek.

Users can tackle this alone or work with Convey. "It can take two weeks or two years," laughs Dr. Vacek,

"depending on what you are trying to do. Right now, we're turning them out in about three months."

## **NEW COMPUTING PARADIGM**

"The real key here," says VBI's Dr. Garner, "is that Convey has created a general purpose, paradigm-shifting machine and software environment. It can be applied anywhere a standard microprocessor can. An HC-1 is easily within the budget of anyone who's buying a small cluster."

As noted earlier, achieving higher speeds at reduced cost is good in itself, yet that's usually just a first step. Faster, more affordable computer technology enables researchers to improve accuracy and tackle projects previously beyond reach because of the compute time

required. VBI's work on the 1000 Genomes Project is a good example. VBI is studying microsatellites — repetitive DNA stretches — and the datasets are enormous. VBI had downloaded on the order of 500 TB of project data recently, with more to come.

Dr. Garner wants to analyze the data, not once, but many times. "Having enough storage, having the lossless compression algorithms that allow us to compress and decompress the data on the fly, and having enough processing power are all necessary to do that." The HC-1 is already providing the needed processing power and rapidly plowing through 1000 Genomes data while Dr. Garner's team works to solve the storage and data-compression problems, which will boost throughput even more.

- 1 Xilinx Field Programmable Gate Array
- 2 Metrics for a single system. According to Convey's internal benchmarking, the Smith-Waterman implementation is 172x faster than SSEARCH in FASTA on an Intel Nehalem core using the SIMD SSE2 instruction set. The company's hybrid-core computer can process 688 billion cell updates per second (GCUPS) as compared to four GCUPS for SSEARCH in FASTA.