


Heterogeneous Computing in the Cloud: Crunching Big Data and Democratizing HPC Access for the Life Sciences



A New Paradigm for Life Sciences Computing

The combination of heterogeneous computing and cloud computing is emerging as a powerful new paradigm to meet the requirements for high-performance computing (HPC) and data throughput throughout the life sciences (LS) and healthcare value chains. Of course, neither cloud computing nor the use of innovative computing architectures is new, but the rise of big data as a defining feature of modern life sciences and the proliferation of vastly differing applications to mine the data have dramatically changed the landscape of LS computing requirements.

Heterogeneous cloud computing offers the potential to shift flexibly from one HPC architecture to another in a secure public or private cloud environment. As such, it meets two critical needs for life sciences computing:

- **Crunch more data faster.** As data sets have ballooned, HPC approaches have evolved and diversified to better match specific problems with their most effective HPC solutions. Indeed, some LS problems (for example, *de novo* genome assembly) are essentially intractable unless tackled with special-purpose solutions. Today, no single approach is adequate. Heterogeneous computing embodies the use of multiple approaches to achieve optimal throughput for each big data LS workload.

- **Democratize access.** While the scope and complexity of HPC resources have grown, the ability of research groups to identify, afford, and support them has diminished. Budget constraints are limiting access to necessary compute resources at the very time when the explosive growth in LS data makes access increasingly desirable. HPC-oriented clouds supporting the latest

heterogeneous architectures can provide even small research groups with affordable access to diverse compute resources.

This paper discusses several trends and enablers of affordable, heterogeneous cloud computing for LS, including the new Intel® Xeon Phi™ coprocessor, based on Intel® Many Integrated Core Architecture (Intel® MIC Architecture). The Intel Xeon Phi coprocessor represents a breakthrough in heterogeneous computing by delivering exceptional throughput and energy efficiency without the high costs, inflexibility, and programming challenges that have plagued many previous approaches to heterogeneous computing. This paper also provides brief examples of heterogeneous computing innovators, such as Nimbix and Convey Computer, that are adopting or experimenting with heterogeneous computing approaches.

Rising Demand for Big Data Computing

Genomics is perhaps the clearest LS example where progress is accelerated by access to the right HPC resource and stymied when those resources

Ketan Paranjape
Director, Healthcare
and Life Sciences,
Intel Corporation

Steve Hebert
CEO,
Nimbix

Bob Masson
Marketing Manager,
Convey Computer

Table of Contents

- A New Paradigm for Life Sciences Computing 1**
- Rising Demand for Big Data Computing 1**
- Application-Specific Computing 3**
- Intel Xeon Phi Coprocessor Advances Performance, Energy, and Programmability 3**
- Embracing Heterogeneous Computing: Jackson Laboratory 4**
- Other Advances: Improving Ease of Use and Development for Heterogeneous Computing 5**
- Democratizing HPC Access with Cloud Computing 6**
- Hiding Complexity from the User 7**
- Entering a New Era 8**
- Learn More 8**

are missing. Pioneering computational biologist [Dr. Eric Schadt](#), currently director of the Institute for Genomics and Multiscale Biology at Mt. Sinai Hospital, captured the challenge well in a recent *Nature Reviews Genetics* paper. Dr. Schadt and his colleagues wrote:

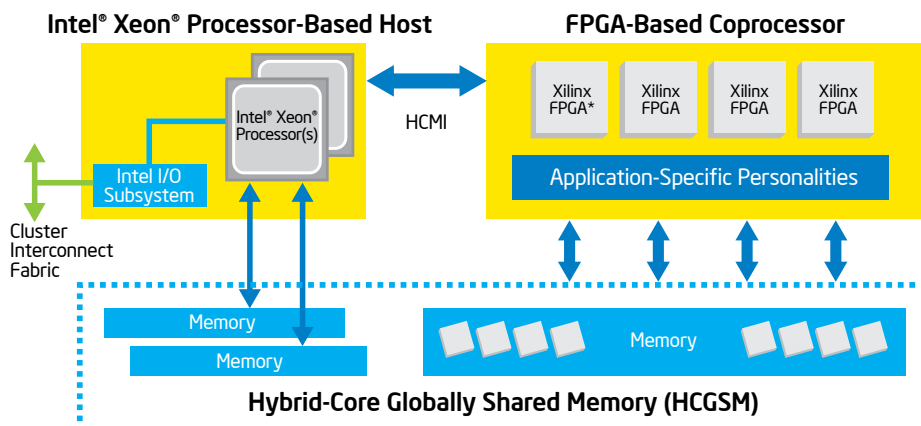
“The astonishing rate of data generation by these low-cost, high-throughput technologies in genomics is being matched by that of other technologies, such as real-time imaging and mass spectrometry-based flow cytometry. Success in the life sciences will depend on our ability to properly interpret the large-scale, high-dimensional data sets that are generated by these technologies, which in turn requires us to adopt advances in informatics...[Scientists must] master different types of computational environments that exist—such as cloud and heterogeneous computing—to successfully tackle our big data problems.

“In under a year genomics technologies will enable individual laboratories to generate terabyte or even petabyte scales of data at a reasonable cost. However, the computational infrastructure that is required to maintain and process these large-scale data sets, and to integrate them with other large-scale sets, is typically beyond the reach of small laboratories and is increasingly posing challenges even for large institutes.”¹

Cloud-based, heterogeneous computing represents a significant step toward solving these problems. Indeed, heterogeneous computing has become virtually a necessity in life sciences, where the output from next-generation sequencing (NGS) instruments represents a data tipping point. This data deluge has outpaced even the steady performance-doubling of Moore’s Law. New approaches based on specialized processors such as field-programmable gate arrays (FPGAs) and general-purpose computing on graphics processing units (GPGPUs), as well as innovative computing strategies such as Apache Hadoop*, are being pressed into service with impressive results.

Not surprisingly, heterogeneous approaches are being embraced by the bioinformatics community. “We just don’t have enough electricity, cooling, floor space, money, etc. for using standard clusters or parallel processing to handle the load,” said Dr. Harold “Skip” Garner, head of the Medical Informatics Systems Division and former executive director of the Virginia Bioinformatics Institute, both part of the Virginia Polytechnic Institute and State University (Virginia Tech). “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.”²

Figure 1: The Convey Computer hybrid-core system provides a reconfigurable, application-specific accelerator, but appears as a standard Intel® architecture-based server to the rest of the computing infrastructure



Application-Specific Computing

It is still early days for heterogeneous HPC cloud computing. Advances in virtualization technology, the porting of key bioinformatics algorithms to special-purpose hardware, and ongoing evolution in standard microprocessor architecture, such as the Intel Xeon Phi coprocessor, are all important enablers of cloud-based, heterogeneous computing. While challenges remain (faster data transfer and data security concerns, for example), heterogeneous HPC in the cloud is demonstrating the potential to broadly enable researchers and clinicians.

GPU-based acceleration was an early success. First developed to speed graphics performance in gaming, GPUs' excellent floating-point capability proved attractive in many applications. GPGPU-based systems were quickly adopted by the oil and gas industry, where they delivered dramatic speedups at reduced cost. Today, general-purpose GPU computing has spread to LS disciplines where floating-point performance is important (molecular modeling, for example, and some alignment algorithms).

More recently, systems based on FPGAs have gained momentum throughout genomics and life sciences. While it's possible to use application-specific integrated circuits (ASICs), LS applications change so quickly that it's impractical to spin up an ASIC for each algorithm for every HPC application. This approach is not only expensive, but it can require years to design and fabricate an ASIC, and the logic is indelibly etched in the semiconductor and unchangeable. While such algorithms would be exceptionally fast by general-purpose processor standards, they're generally impractical.

A reasonable compromise is the use of FPGAs. Programmable "on the fly," FPGAs are a way to achieve hardware-based, application-specific performance without the time and cost of developing an ASIC. FPGAs work well on many bioinformatics applications—for example, those that

do searching and alignment. Such applications rely on many independent and simple operations, and are thus highly parallelizable.

Unique architectures from innovative companies are emerging to help meet the demand. For example, [Convey Computer](#), established in December 2006, has created a hybrid-core system that pairs classic Intel® processors with a coprocessor of FPGAs (see Figure 1). Particular algorithms—DNA sequence assembly, for example—are optimized and translated into code that's loaded onto the FPGAs at runtime. The Convey architecture also features a highly parallel memory subsystem to further increase performance. Convey Computer's approach provides very fast access to random access to memory or single-word access to memory, and is very useful for the hashing functions so widely used in bioinformatics.

Intel Xeon Phi Coprocessor Advances Performance, Energy, and Programmability

Lively debate persists over the right HPC approach for various applications and disciplines. For example, FPGA advocates contend that if you lay out the gates on a chip to execute the specific aspect of an algorithm, it will by definition be more efficient than a series of instructions on a commodity processor. However, if getting an algorithm onto a new architecture requires that you use a new language, then the development time and productivity lost during the development can often outweigh the cost or benefits gained by going down a special-purpose path.

The new Intel Xeon Phi coprocessor, based on Intel MIC architecture, promises to be a game-changer for many LS applications that now run on special-purpose systems (see Figure 2). The Intel MIC architecture preserves the programming model that has been long established for the Intel® architecture, enabling developers and LS users to increase performance without limiting flexibility or investing the time typically needed for earlier

application-specific approaches. It substantially streamlines development time required to create a new application because you do not have to use a new programming paradigm.

The Intel Xeon Phi coprocessor and Intel MIC architecture are designed to tackle HPC applications and numerically intensive operations. Manufactured using next-generation Intel® 22nm 3-D Tri-Gate transistor technology, the Intel Xeon Phi coprocessor utilizes a high degree of parallelism in smaller, lower-power Intel processor cores. Each Intel Xeon Phi coprocessor contains more than 50 cores and a minimum of 8 GB of high-performance, power-sensitive GDDR5 memory. The result is higher performance on highly parallel applications. In addition, the coprocessor's density and energy-efficiency help save space and reduce power and cooling costs in the data center, making it well suited to modern cloud environments and LS applications.

To cite an example from outside the life sciences, the U.S. Department of Energy's National Renewable Energy Laboratory (NREL) is building a petascale HPC system that will use the Intel® Xeon® processor E5 family and Intel Xeon Phi coprocessors. In addition to using the energy-efficient Intel processors and coprocessors, the Department of Energy worked with Intel and HP to optimize its data center design, and expects the system's power usage effectiveness (PUE) rating to be

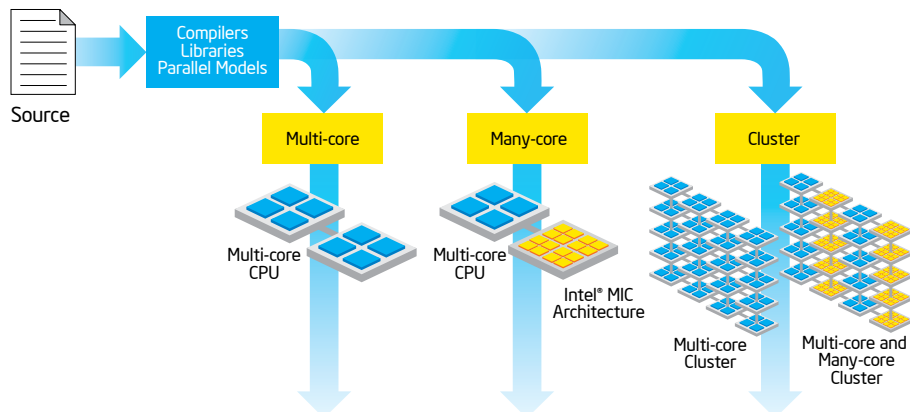
nearly two times more efficient than the average. Illustrating the system's easy programmability, NREL participated in software development for Intel MIC architecture, and needed only a few days to port a half-million lines of code of the Weather Research and Forecasting* (WRF*) application to prepare for taking full advantage of the energy efficiency and performance of Intel Xeon Phi coprocessors. NREL's petascale computer will be dedicated to researching renewable energy and energy efficiency.³

In attacking the power wall, Intel examined hardware and software issues, noted which routines are more energy efficient, and developed a power management functionality that makes it possible to direct how the processor spends its power budget. For example, if a problem needs more numerical computation but doesn't have to move a lot of data around, you can drive up power going to compute circuits to boost performance of the cores and power down data movement. Conversely, you can also power down unused cores and use the extra power budget to accelerate data movement.

Embracing Heterogeneous Computing: Jackson Laboratory

The vast majority of bioinformatics and healthcare applications currently run on standard clusters, but that's changing as research organizations hit technical and financial roadblocks in their efforts

Figure 2: The techniques that deliver optimal performance on the widely used Intel® Xeon® processor also apply to the new Intel® Xeon Phi™ coprocessor, providing GPGPU advantages plus simplified development



to obtain the HPC resources they need for deriving value from gush data from experimental instruments. Big data is forcing changes in attitudes, and driving a need for faster, more power-efficient computing. This is especially true of sequencing centers dealing with large data sets. (The Broad Institute's annual sequencing capacity is now about *300,000 billion bases*, and steadily rising.⁴ Worldwide annual sequencing capacity exceeds *13 quadrillion bases*.⁵)

The [Jackson Laboratory \(JAX\)](#), a National Cancer Institute-designated cancer center, is a prominent example of a research organization embracing heterogeneous computing. Long at the forefront of mammalian genetics research, JAX has rapidly increased its use of next-generation sequencing. "Once we could afford whole genome sequencing, we found a significant bottleneck in the time required to process the data," said Laura Reinholdt, Ph.D., a research scientist at JAX. "That's when biologists here began to seek tools and infrastructure to more expediently manage and process the expanding volumes of NGS data."⁶

JAX settled on heterogeneous computing as the solution. "It comes down to power consumption, space, and performance for a fixed amount of dollars," said Glen Beane, senior software engineer at JAX. "We looked at various options for hybrid systems. We found GPUs weren't a good fit for alignment—there are packages that do alignment but the performance isn't that compelling."

JAX chose an FPGA-accelerated system from Convey Computer, and on several key applications has achieved roughly an elevenfold speedup over its 32-core cluster using a single FPGA-based system. JAX didn't abandon the cluster; instead, the lab is trying to divert its various computational tasks to the most effective resource for the problem at hand—whether it's rigorous alignment-seeking to uncover disease gene variants or *de novo* sequencing.

Other Advances: Improving Ease of Use and Development for Heterogeneous Computing

One issue in life sciences and healthcare is that biologists, physicians, and other LS users are usually not IT or HPC experts. For many, it is challenging enough to choose the best application for a given problem, let alone determine which HPC architecture would run it most efficiently. One organization, [The Center for Biotechnology at Bielefeld University](#), has a variety of HPC resources (FPGA, GPU, and others) and is working to set up a suite of applications that essentially know which HPC resources are best. The vision is that users will simply submit a job, which then goes off and finds the most appropriate architecture given the particular application and the job's specified parameters.

Another concern with special-purpose HPC architectures is the need to adapt existing software or develop new software to take advantage of the specific approach, which can consume time and resources. Increasingly, systems makers are tackling this problem and working to ensure their application coverage is attractive. Convey Computer, for example, offers an expanding suite of applications, including key bioinformatics algorithms, optimized for its architecture. User groups are also springing up around particular architectures, developing their own accelerated applications and making them available to others.

The emergence of more high- and intermediate-level tools for FPGAs and GPGPUs is helping to speed and simplify applications development. For example, students at Iowa State University won this year's [MemoCODE Design Contest](#) for rapidly developing an FPGA-based application.⁷ The [2012 challenge](#) was to efficiently locate millions of 100-base-pair short read sequences in a 3-million-base-pair reference genome and, as described by organizers, "Good solutions will combine judicious algorithm design with carefully designed data-handling architecture."⁸ Using a Convey FPGA-based system, the

students' solution achieved the highest overall performance—more than 24 times faster than the second-place finisher. In addition, they were able to develop the algorithm and implement it on an FPGA in roughly one month.

Democratizing HPC Access with Cloud Computing

Moving heterogeneous HPC assets into a cloud computing environment is a natural step. It provides the widely discussed benefits of cloud computing such as lower costs and rapid scalability, and in fact magnifies them, since heterogeneous HPC resources usually entail greater cost, integration, and management challenges than standard cluster-based clouds. Here are a few of cloud computing's substantial benefits:

- **Pay as you go.** With cost-containment an increasing priority, many research organizations are focusing funds on core competencies, preferring to outsource where practical and buy only what's needed and when. Cloud environments also make it possible to rapidly scale jobs up or down as needed, and heterogeneous HPC cloud charges typically vary based on which resources are used.
- **Different budget.** It's often easier to tap variable operations budgets than go through a lengthy approval process for scarce capital equipment funds.
- **Reduced IT support.** Many computer administration costs (and worries) are shifted to the services provider.
- **Technology upgrades.** Pushed by competition and customers, cloud providers can be expected to be earlier adopters of new hardware and software technology. This enables the LS user community to benefit from the latest advances without undertaking costly technology refreshes every couple of years.
- **Public or private.** Cloud providers increasingly offer secure public clouds, used by many clients, or private clouds firewalled off and dedicated 24/7 to a single organization.

For organizations that require diverse HPC resources, turning to a cloud provider may be the only practical choice. Even well-funded genomics shops frequently discover that when they build a system at 120 percent of the capacity they expect for a two-year run, they run out of capacity sooner than expected. Offloading at least some of their computing demand to cloud providers is an attractive idea.

Given the ubiquity of the data deluge in genomics, the number of research organizations investigating cloud-based heterogeneous computing is surging. One example is the Institute of Environmental Science and Research (ESR) at the National Centre for Biosecurity and Infectious Disease (NCBID), New Zealand. ESR has been in midst of a technology build-out for NGS sequencing and analysis capability.⁹ "We'd read a paper about GPU-based acceleration for BLAST* and wanted to explore that," said Jing Wang, scientist (bioinformatics) at ESR, where she is participating in the pilot project Pathogen Discovery (virus identification in various organisms).

Searching the Internet, Wang found the [Nimbix Accelerated Compute Cloud*](#). Nimbix specializes in offering heterogeneous HPC resources including, among other assets, GPU machines. It turned out the version of BLAST available on Nimbix's GPU-accelerators at the time (BLASTp*) wasn't ideal. "We wanted to use BLASTn* because of the nature of our data sets," said Wang.

Nimbix suggested that Wang try the Smith-Waterman* (SW*) algorithm on Convey's FPGA-accelerated computers. "We knew that on a general CPU-based platform, this is not doable, because although the software is quite accurate, it takes too long," Wang said. "We were willing to do a test run with a small data set (approximately 1 million reads)." The FPGA-based approach was so fast, Wang decided to run a much larger data set (approximately 35 million reads).

She emphasized that the cloud approach worked well for ESR and that ramping up a portfolio of diverse HPC assets is beyond ESR's capability. "One challenge that is

quite daunting at the moment is there are so many applications developed for research purposes and to evaluate them is an impossible task for us," Wang said. "I really want to see some recommendations from organizations like Nimbix, which has deep experience with the resources."

Many commercial cloud vendors are adding various HPC elements, but few are focused on offering heterogeneous HPC resources, as Nimbix is. To some extent the "buyer beware" caution applies, as noted in a paper presented at a 2011 IEEE conference on cluster computing: "[S]ince the current cloud computing market evolved from the IT community, it is often not a good match for the needs of technical computing end-users from the high-performance computing (HPC) community. Providers such as Amazon and Rackspace provide users with access to a homogeneous set of commodity hardware...By contrast, technical computing end-users may want to obtain access to a heterogeneous set of resources, such as different accelerators, machine architectures, and network interconnects."¹⁰

Hiding Complexity from the User

As noted earlier, ease of use is an important issue in delivering diverse HPC resources to life science and healthcare workers. Nimbix, as an example, takes pains to mask the heterogeneous computing's complexity, offering users a choice of two approaches (see Figure 3):

- **Web portal.** Using this interface, Nimbix users simply select the desired application for their pipeline, point to their data and any parameters for that particular case that they want to run, and submit the job.

- **Web services.** It is also possible to use an API call. Nimbix provides the hooks that enable a user of Galaxy*, for example, to allow the pipeline to send certain jobs to the Nimbix cloud.

The underlying assumption is that the applications, tuned for the specific architecture, do indeed exist in the cloud.

Data transmission remains a substantial challenge, particularly in genomics, where the data sets are so large. For truly massive data sets, it is still generally necessary to ship hard-disk drives to cloud centers. However, Nimbix finds that many data sets submitted for individual jobs are below the terabyte range. A user might aggregate up to several terabytes, but each individual work order is often moving in the multi-gigabyte range, which is manageable. With Nimbix, users point to their data, it's accessed via a secure File Transfer Protocol, and Nimbix retrieves it, uploads it, and automatically launches the processing task.

Progress on data transmission rates is also ongoing. Recently, Beijing-based BGI, the largest sequencing center in the world, transferred 24 GB of genomic data from Beijing to the University of California, Davis in less than 30 seconds. (A file of the same size sent over the public Internet a few days earlier took more than 26 hours.) The measured data rate is equivalent to moving more than 100 million megabytes—over 5,400 full Blu-ray Discs*—in a single day.¹¹ The transfer took place on June 22, 2012, as part of an event in Beijing to celebrate a new 10 Gb US-China network connection,

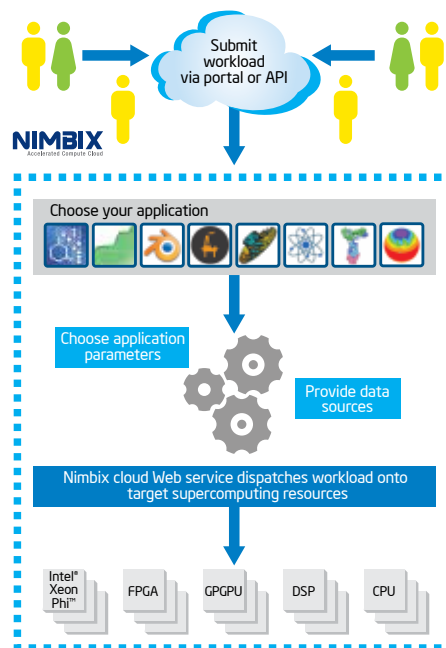


Figure 3: The Nimbix cloud interface: Cloud leaders are working to simplify the user's access to heterogeneous HPC resources

supported by groups including Internet2, the National Science Foundation (NSF), and Indiana University.

Entering a New Era

Looking forward, it's clear that the big data challenge encountered by life sciences and healthcare will only grow. As Dr. Schadt wrote, "[T]he amount of data from large projects such as [1000 Genomes](#) will collectively approach the petabyte scale for the raw information alone. The situation will soon be exacerbated by third-generation sequencing technologies that will enable us to scan entire genomes, microbiomes, and transcriptomes and to assess epigenetic changes directly in just minutes, and for less than USD 100. To this should be added data from imaging technologies, other high-dimensional sensing methods, and personal medical records."¹²

No single HPC architecture is best for managing and analyzing all these workloads and data sets; heterogeneous computing is necessary to proceed productively. Moreover, HPC architectures steadily evolve, and what constitutes the best resource regularly shifts. In this context, moving heterogeneous HPC

resources to the cloud may be the only practical way many LS organizations can afford access to the latest compute power to advance life sciences and healthcare.

Advances such as the Intel Xeon Phi coprocessor, improved development tools, and expanding numbers of applications optimized for special-purpose systems are lowering the barriers to heterogeneous computing—and the benefits are compelling. Speeding throughput enables scientists to publish more quickly, tackle larger data sets and more difficult problems, produce more accurate results by using more rigorous algorithms, and generate critical diagnostics for patients more quickly, to name a few advantages. They can often do this while conserving space and power. Perhaps most exciting, possessing more potent compute resources inevitably spurs new applications and fresh thinking about how to apply the new HPC capability. We are truly on the forefront of a new era in high-performance computing for the life sciences.

John Hengeveld, director of HPC marketing at Intel, has taken this to heart. "Big data technology combined

with high-performance computing offers the prospect of significantly improved therapies and treatments for cancer patients through the development of cost-effective, personalized genomics," Hengeveld says. "People like me with relatively rare cancers can find hope in this approach as never before."

Learn More

Through its research and development activities and its collaborations with other technology leaders, Intel drives HPC and cloud computing solutions that enhance scientific and technical progress, business efficiency, IT flexibility, and enterprise security. Intel® server, storage, and networking technologies are at the heart of today's most robust clouds, and Intel Science and Technology Centers conduct R&D that shapes tomorrow's cloud and big data infrastructures.

Learn how Intel can help your organization meet its needs for big data and heterogeneous computing in a public or private cloud environment. Talk to your Intel representative, visit us on the Web at www.intel.com/healthcare or www.intel.com/bigdata, or follow us on Twitter: [@IntelHealthIT](https://twitter.com/IntelHealthIT).



¹ E. Schadt et al., *Computational solutions to large-scale data management and analysis*, Nature Reviews: Genetics, Vol 11, Sep 2010.

² *Convey Computer Times Data Deluge*, Convey Computer white paper. http://conveycomputer.com/Resources/Convey_Computer_BioIT_White_Paper.pdf

³ http://newsroom.intel.com/community/intel_newsroom/blog/2012/09/05/intel-xeon-processors-and-xeon-phi-coprocessors-promise-to-power-world-s-most-efficient-hpc-data-center

⁴ 2011 Broad Annual Report, http://www.broadinstitute.org/files/news/media-kit/2011_broad_institute_ar.pdf

⁵ *DNA Sequencing Caught in Deluge of Data*, New York Times, Nov. 30, 2011, http://www.nytimes.com/2011/12/01/business/dna-sequencing-caught-in-deluge-of-data.html?_r=1&ref=science

⁶ *The Jackson Laboratory Chooses Convey's HC-2*, April 24, 2012, <http://conveycomputer.com/Resources/jackson-laboratory-JAX-chooses-convey-hybrid-core-for-genomic-ngs-research.pdf>

⁷ *Iowa State University Students Win MEMOCODE Design Contest Using Convey Computer HC-1*, MarketWire, Jul 12, 2012, <http://www.marketwire.com/press-release/iowa-state-university-students-win-memocode-design-contest-using-convey-computer-hc-1679423.htm>

⁸ MEMOCODE 2012 Hardware/Software Codesign Contest: DNA Sequence Aligner, <http://memocode.irisa.fr/2012/2012-memocode-contest.pdf>

⁹ Author conversation with Jing Wang, scientist, The Institute of Environmental Science and Research (ESR) at the National Centre for Biosecurity and Infectious Disease (NCBID).

¹⁰ Heterogeneous Computing, Steve Crago et al., contributed paper, 2011 IEEE International Conference on Cluster Computing.

¹¹ *BGI Demonstrated Genomic Data Transfer at Nearly 10 Gigabits per Second Between US and China*, <http://bjamericas.com/bgi-demonstrated-genomic-data-transfer-at-nearly-10-gigabits-per-second-between-us-and-china/>

¹² E. Schadt et al., *Computational Solutions*

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors. Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. For more information go to <http://www.intel.com/performance>

Intel does not control or audit the design or implementation of third party benchmark data or Web sites referenced in this document. Intel encourages all of its customers to visit the referenced Web sites or others where similar performance benchmark data are reported and confirm whether the referenced benchmark data are accurate and reflect performance of systems available for purchase.

This document and the information given are for the convenience of Intel's customer base and are provided "AS IS" WITH NO WARRANTIES WHATSOEVER, EXPRESS OR IMPLIED, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, AND NON-INFRINGEMENT OF INTELLECTUAL PROPERTY RIGHTS. Receipt or possession of this document does not grant any license to any of the intellectual property described, displayed, or contained herein. Intel® products are not intended for use in medical, lifesaving, life-sustaining, critical control, or safety systems, or in nuclear facility applications

© 2012, Intel Corporation. All rights reserved. Intel, the Intel logo, Intel Xeon, Intel Xeon Phi, and Xeon inside are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

