

SOLUTION BRIEF

Life Sciences
Intel® Scalable System Framework



Science without Constraints

Fueling the Next Great Wave of Data-Driven Innovation in the Life Sciences

We are on the cusp of a revolution in the biological and medical sciences. A whole human genome can now be sequenced in a matter of hours and for as little as USD 1,000,¹ and we are moving quickly toward the USD 100 genome.² Meanwhile, technologies such as Cryo-Electron Microscopy (Cryo-EM) and Molecular Dynamics are helping researchers visualize and understand cellular processes at the molecular level.

These and other technologies are opening a window into the most fundamental processes of life. Biological pathways can be illuminated, disease mechanisms can be identified, and drug discovery can be transformed from a multiyear, multi-billion-dollar, trial-and-error process to an efficient, data-driven workflow. Perhaps most importantly, precision medicine, with molecular-level profiles and personalized treatments, will ultimately transform the way we diagnose and treat injury and disease.

Insights Buried in Massive, Complex Data Sets

Turning this vision into reality will not be easy. One of the biggest challenges will be processing the massive amounts of data generated by high-volume sequencers. Already, organizations must be able to store, secure, and share up to three terabytes of data per day from a single sequencer.³ One study estimates that between 100 million and 2 billion individual genomes will be sequenced by 2025, generating 2 to 40 exabytes of data.⁴

This mountain of data will have to be analyzed in combination with imaging, clinical, personal, and environmental data sets that are also large, complex, and widely distributed. A new generation of hardware and software capabilities is required so that scientists and clinicians will not be constrained by the long delays and extreme costs of moving, integrating, and analyzing all this data.

Systems and networks must be more powerful and affordable. They must also be scalable enough to support ongoing increases in data volumes and algorithmic complexity. As described in the following examples, Intel is working with leading academic, open source, and commercial institutions to overcome these challenges and unleash more and faster discovery.

Opening the Door to Fast Genome Analysis

Order-of-Magnitude and Higher Performance Gains for Key Genomic Algorithms⁵

Complex software algorithms have been developed over many years to analyze genomic data. These algorithms provide amazing functionality, but since they were written largely by genomic researchers, whose focus was primarily on scientific accuracy, the code was never fully optimized for modern computing platforms.

Professor Knut Reinert, PhD, and his team at the Free University of Berlin are collaborating with Intel to accelerate genome analysis by optimizing critical algorithms so they run efficiently on multicore and many-core Intel® processors.

Recent benchmarks using the Intel® Xeon® Gold 6148 processor quantify the benefits. The results show performance gains ranging from 1.6X to 2.7X versus the previous-generation Intel® Xeon® processor E5-2697 v4 (see Figure 1).⁵ The optimized code also provides excellent scalability across large numbers of cores. In benchmarks using the Intel® Xeon Phi™ processor 7250, runtimes were decreased by as much as 55X when utilizing all 68 cores versus the same workload running on a single core (see Figure 2).⁶

Professor Reinert and his team package their optimized algorithms in SeqAn*, an open source software library that is optimized for ease-of-use, maintainability, and portability. Organizations can use SeqAn to quickly assemble complex genomic pipelines that take advantage of the power and scalability of multicore and many-core Intel processors.

Learn more at: <https://software.intel.com/en-us/articles/intel-parallel-computing-center-at-freie-universitat-berlin>

Exploring Biology at Near-Atomic Scale

High-speed image reconstruction for Cryo-EM

Cryo-EM technology allows scientists to explore cellular processes and other complex chemical interactions in three dimensions and at near-atomic scale. However, research has been slowed by long runtimes for image reconstruction—roughly 47 hours for a typical study.

By modernizing the algorithms in open source RELION*, Professor Erik Lindahl of Stockholm University and Professor Sjors Scheres of MRC Laboratory have reduced image reconstruction times by an order of magnitude—to as few as 5.7 hours on a single server and as few as 1.0 hours on an eight-node cluster.³ Their work will help speed research and increase the proliferation of this powerful imaging technique.

Learn more about Cryo-EM: <https://www.youtube.com/watch?v=BtuAz12zXBs>

Faster SeqAn Performance with More Cores and Advanced Vectorization
Lower is Better

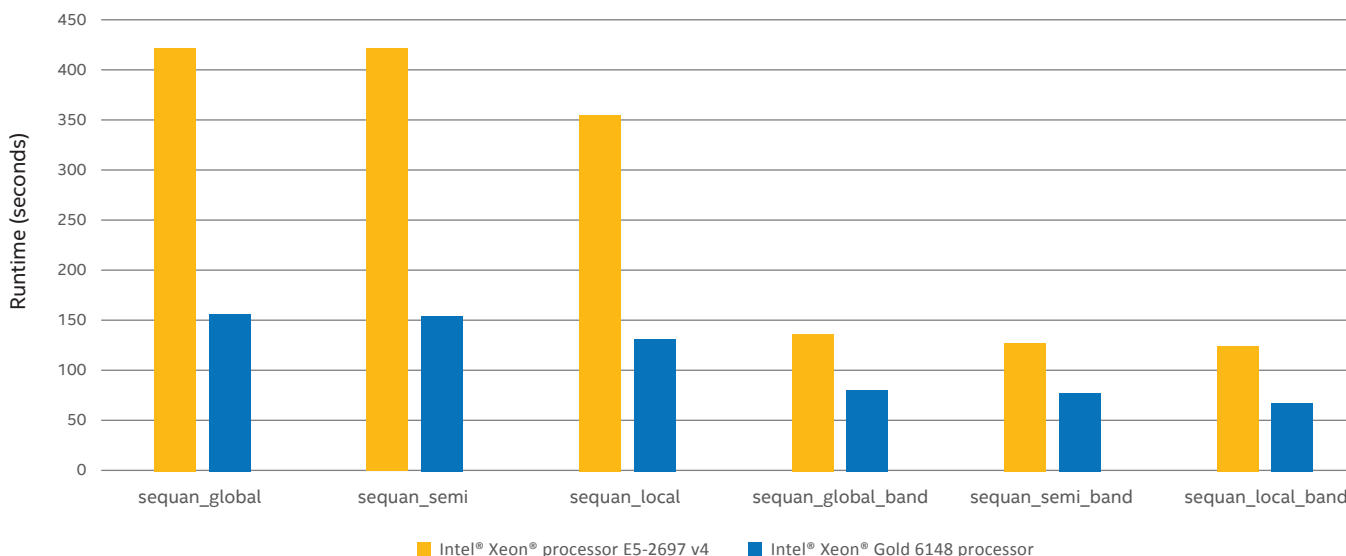


Figure 1. The optimized SeqAn* code takes advantage of advanced vector capabilities in the latest Intel® Xeon® Gold 6148 processor to enable performance gains as high as 1.6X to 2.7X versus previous-generation processors.

Scalable SeqAn Performance on Many-Core Processors (Intel® Xeon® Phi™ processor 7250)

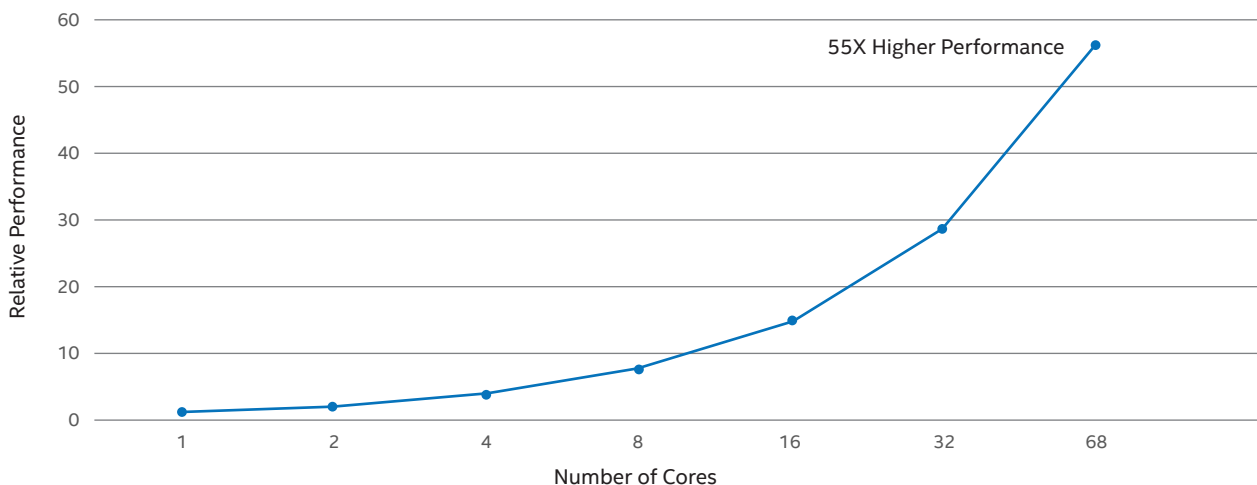


Figure 2. The optimized SeqAn* code also provides excellent scalability across all 68 cores of the many-core Intel® Xeon Phi™ processor 7250.

Providing a New Window into the Brain—at Clinically Useful Speeds

Up to 161X Faster Performance for DCI Image Processing⁷

Traditional CT and MRI scans provide useful information about bones and blood vessels in the brain, but little about soft tissues. Diffusion-Weighted Imaging (DWI) addresses that gap by measuring the natural movements of water molecules to identify microstructures and diagnose the integrity of neural circuits.

This technology is transforming the research and treatment of brain disorders. Simon Warfield, professor of Radiology at Harvard Medical School, took DWI a step further with his Diffusion Compartment Imaging (DCI) technique, which provides more granular and accurate information. His team is using DCI to investigate normal brain development, concussions, multiple sclerosis, autism spectrum disorder, and more.

Tens of gigabytes of data are generated per DCI study, and image processing was initially taking more than 40 hours. Working with Intel® software tools, Professor Warfield and his team optimized their code to provide a 75X performance improvement using Intel® Xeon® processors and a 161X improvement using Intel Xeon Phi processors (see Figure 3).⁷ A complete study can now be processed in just over 16 minutes, which is fast enough to fit into clinical workflows. The optimized code is available in the Insight Segmentation and Registration Toolkit* (ITK) library, which is widely used in applications that process medical imaging data.

Professor Warfield is also working with Intel® tools and libraries to develop artificial intelligence (AI) applications that automatically sort through hundreds of images to pinpoint areas of interest. As these solutions move into the mainstream, our ability to understand, diagnose, and heal the brain will increase dramatically.

Learn more at: <https://software.intel.com/en-us/articles/intel-parallel-computing-center-at-the-computational-radiology-laboratory> and <https://itk.org/>

Scaling into the Future on Fast, Affordable Clusters

Intel® Scalable System Framework for Life Sciences

Improving software performance is a valuable first step, but even higher performance is needed to address the growing demands of bioscience workloads. High performance computing (HPC) clusters offer a clear path to solving these data processing challenges. However, cost and complexity remain barriers to adoption for many organizations. Intel is working with leading research, healthcare, open source, and commercial institutions to eliminate these barriers.

An important result of these efforts is the Intel® Scalable System Framework (Intel® SSF) for Life Sciences. This hardware and software reference architecture brings together all the elements of an HPC system into a balanced, high-performing solution optimized for research and healthcare organizations.

Higher Performance through Software Optimization

Diffusion Compartment Imaging Estimation (2x2x2 mm³ image)

Software Enhancement	Runtime	Performance Gain
- Intel® Xeon® processor E5-2697 v2:		
• Single core (baseline):	43h06m21s	x14.2
• Parallelization v1:	03h02m02s	x1.76
• Memory optimization:	01h43m31s	x1.12
• Vectorization:	01h43m31s	x1.37
• Intel® TBB filter, dynamic pool:	01h07m26s	x1.22
• Flexible Intel® TBB decomposition:	55m23s	x1.44
• Optimizer improvement:	38m27s	x1.12
• Intel compilation flag:	34m14s	
- Intel® Xeon® Phi™ processor 7210:	16m03s	x2.13

Figure 3. Through code optimization, Professor Warfield and his team were able to reduce DCI estimation runtimes by up to 75X on multicore Intel® Xeon® processors. Running the same code on the many-core Intel® Xeon Phi™ processor 7210 provided additional gains of up to 2.13X.

The goal of Intel SSF for Life Sciences is not only to improve performance and cost models, but also to simplify HPC deployment and to support the full range of workloads on a single platform—molecular dynamics, genomics, molecular imaging, machine learning and AI, data visualization, and more. This flexibility and efficiency makes it easier for organizations to run diverse applications and combine disparate data sets to solve complex problems.

Intel SSF defines a complete, optimized solution stack for cluster computing, bringing together innovations from Intel at every layer to enable higher performance at lower cost.

Powerful, Flexible Processors

Powerful processors provide the essential foundation for any high-quality HPC solution. Choosing the right processors for the right workloads offers additional efficiencies, making it easier for organizations to address the extreme compute and memory requirements of real-world applications. Intel SSF for Life Sciences supports a wide range of Intel processors, including most importantly:

- **Intel® Xeon® Scalable processors.** These powerful, multicore processors provide scalability, optimizations, and substantial performance increases for a wide variety of mainstream and HPC applications. They have demonstrated gains as high as 1.73X for HPC workloads⁸ versus the previous-generation Intel® Xeon® processor E5 v4 family. In addition to providing more cores, cache, and memory bandwidth, they have been re-architected for greater efficiency in data-intensive environments.
- **Intel Xeon Phi processors.** These highly parallel processors provide up to 2.5X more cores and up to 5X more threads than Intel Xeon processors.⁹ They are bootable, available with integrated high-speed memory, and can run all standard x86 applications. The extra execution and memory resources can provide dramatic improvements in performance, scalability, and energy efficiency for highly parallel code, which includes some of today's most demanding life sciences applications.

Efficient, Scalable Clusters

As Intel processor performance has increased over the years, memory and storage technologies have lagged far behind, leading to data access bottlenecks that throttle performance for many applications. Intel SSF brings together multiple technologies to help resolve these bottlenecks and to unleash the full power of Intel processors.

- **A breakthrough in memory capacity and storage performance.** Intel® Optane™ Solid State Drives (SSDs), built with 3D XPoint™ Memory Media, offer a major leap forward in memory capacity and storage performance. These SSDs are designed to deliver 5–8X higher performance than NAND-based SSDs¹⁰ and can be used as ultra-fast storage, storage cache, or extended memory. In combination with NAND-based Intel® Solid State Drives and traditional hard drives, Intel® Optane™ SSDs make it easier to address extreme data access requirements while containing overall costs.

- **Massive scalability for big data.** The open source Lustre* file system is integrated into Intel SSF to provide a next-generation, software-defined storage solution that addresses the key challenges of high-volume analytics environments. By scaling metadata on separate servers and striping object data across multiple storage drives and servers, Lustre enables extreme performance at almost any scale.

- **A fast, affordable fabric.** In an HPC cluster, data must often move as efficiently between nodes as it does within nodes. Intel® Omni-Path Architecture (Intel® OPA) offers the same 100 Gbps link speeds as today's fastest InfiniBand* fabrics, while providing higher port densities per chip for better scalability and cost models.¹¹ Additional cost, density, and performance advantages can be achieved by using Intel Xeon and Intel Xeon Phi processors with integrated Intel OPA controllers.

Intel SSF for Life Sciences combines these technologies into an efficient, balanced cluster architecture that is highly optimized for performance, density, and power efficiency. Systems can be scaled from small workgroup clusters to supercomputers and can be extended and tuned over time by adding compute nodes based on the various Intel processor and accelerator options.

System software solutions can also be simplified. Maintaining a stable and optimized HPC software environment can be a complex and time-consuming task, even for HPC experts. The OpenHPC* software project aggregates the common ingredients required to deploy and manage HPC Linux* clusters. Intel contributes to this community project to provide up-to-date, open source software ingredients that are optimized for Intel® architecture and make it easier to deploy and operate HPC clusters.

Conclusion

A new generation of computing power is needed to help scientists and clinicians extract meaning more quickly from growing amounts of complex data. Intel offers multiple resources to address this challenge, including:

- **Optimized algorithms and applications** that take better advantage of the parallelism available in modern processors and platforms. These codes can help speed time-to-results, often by an order of magnitude or more, and they tend to scale efficiently on clustered architectures.
- **Intel SSF for Life Sciences**, a blueprint for balanced and efficient HPC clusters that are easier to deploy, manage, and scale, and are designed to efficiently run a wide range of bioscience and medical workloads.

Together, these resources can help organizations accelerate performance and improve their cost models, while simplifying their computing environments so they can focus less on computer science and more on their own science.

Learn More

<https://www.intel.com/content/www/us/en/high-performance-computing/life-sciences.html>



⁴ Tests conducted by Intel as of November 2017. A runtime of 47.6 hours was achieved using the unoptimized version of RELION on a baseline server configured with 2 x Intel® Xeon® processor E5-2697 v4. This was compared with an average runtime of 4.6 hours running the optimized version of RELION on a new server configured with 2 x Intel® Xeon® Gold 6148 processors and a runtime of 1.4 hours on an 8-node cluster of servers, each one configured with 2 x Intel Xeon Gold 6148 processors. Configuration Details:

RELION: Development Build: November 7, 2017 commit on the "cpu-cuda-merge" branch of <https://bitbucket.org/tcblab/relion-devel-tcblab> repository (roughly equivalent to RELION 2.1 beta). This version uses a mix of single- and double-precision where appropriate. It was compiled with GCC* 4.8.5 (GPU runs) and Intel® Cluster Studio 2017 Update 4 (version 2017.4.196 – CPU runs). All benchmarks were run on Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

The Plasmodium ribosome workload is available from http://www2.mrc-lmb.cam.ac.uk/relion/index.php/Benchmarks_%26_computer_hardware

Baseline server configuration: 2 x Intel® Xeon® processor E5-2697 v4 (2.3 GHz, 18 cores/socket, 36 cores, turbo and HT on), SuperMicro X10DRG-H motherboard with American Microtrends Inc. BIOS 2.0a, 128GB total memory, 8 slots / 16 GB / 2400 MT/s / DDR4 RDIMM, 1 x 1TB SATA, Red Hat Enterprise Linux 7.2 kernel 3.10.0-327.

New server configuration: 2 x Intel® Xeon® Gold 6148 processor (2.4 GHz, 40 cores, turbo and HT on), BIOS 86B.01.00.0412, 192GB total memory, 12 slots / 16 GB / 2666 MT/s / DDR4 RDIMM, 1 x 800 GB Intel® SSD SC2BABA0, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327.

¹ Source: National Institute of Health National Human Genome Research Institute, "The Cost of Sequencing a Human Genome," last updated July 6, 2016. <https://www.genome.gov/27565109/the-cost-of-sequencing-a-human-genome/>

² Source: "Illumina introduces the NovaSeq Series—a New Architecture Designed to Usher in the \$100 genome," Illumina press release, January 9, 2017. <https://www.illumina.com/company/news-center/press-releases/press-release-details.html?newsid=2236383>

³ Source: NovaSeq® Series of Sequencing Systems, specification sheet. <https://www.illumina.com/content/dam/illumina-marketing/documents/products/datasheets/novaseq-series-specification-sheet-770-2016-025.pdf>

⁴ Source: "Genome researchers raise alarm over big data," by Erika Check Hayden, Nature, July 7, 2015. <https://www.nature.com/news/genome-researchers-raise-alarm-over-big-data-1.17912>

⁵ Results based on local and global alignments of banded and unbanded cells for 150 bp Illumina reads (2.85 x 1011 unbanded, 3.21 x 1010 banded). Performance claims compare runtimes for a baseline configuration is 2 x Intel® Xeon® processor E5-2697 v4 (2.30 GHz, 18 cores) compared with 2 x Intel® Xeon® Gold processor 6148 (2.4 GHz, 20 cores). Both systems using Linux* 3.10.0-514.21.1.el7.x86_64, and the GNU* compiler 7.2.0 for all tests.

⁶ Results based on Pac Bio Reads for a global alignment of 2.66 x 1013 cells. Tests compared runtimes for the same workload running on from 1 to 68 cores of a single server configured with the Intel® Xeon Phi™ processor 7250 (1.4 GHz, 68 cores, 16 GiBi MCDRAM). All tests run using seqan_global, Linux* 3.10.0-514.21.1.el7.x86_64, and the GNU* compiler 7.2.0.

⁷ Baseline: Unoptimized DCI image processing workload running on Intel® Xeon® processor E5-2697 v2 (2.70 GHz, 24 cores) versus optimized DCI image processing workload running on the same server and on a second server configured with the Intel® Xeon Phi™ processor 7210 (1.30 GHz, 72 cores). The optimized software was able to utilize all available cores on each processor.

⁸ Up to 1.73x claim based on LAMMPS: LAMMPS is a classical molecular dynamics code and an acronym for Large-scale Atomic/Molecular Massively Parallel Simulator. It is used to simulate the movement of atoms to develop better therapeutics, improve alternative energy devices, develop new materials, and more. Performance tests compared a baseline server configured with the Intel® Xeon® processor E5 v4 family with a new server based on the Intel® Xeon® Scalable processor family. Baseline server configuration: 2 x Intel® Xeon® processor E5-2697 v4, 2.3 GHz, 36 cores, Intel® Turbo Boost Technology and Intel® Hyper-Threading Technology (Intel® HT Technology) on, BIOS 86B0271.R00, 8x16 GB 2400 MHz DDR4, Red Hat Enterprise Linux* 7.2 kernel 3.10.0-327. New server configuration: 2 x Intel® Xeon® Gold 6148 processor, 2.4 GHz, 40 cores, Intel® Turbo Boost Technology and Intel HT Technology on, BIOS 86B.01.00.0412.R00, 12x16 GB 2666 MHz DDR4, Red Hat Enterprise Linux 7.2 kernel 3.10.0-327.

⁹ Intel® Xeon Phi™ processors provide up to 72 cores and 4 threads per core, while Intel® Xeon® Scalable processors provide up to 28 cores and 2 threads per core.

¹⁰ Common Configuration: Intel® 2U Server System with 2 x Intel® Xeon® processor E5-2699 v4 (2.20 GHz, 22 cores), 396 GB RAM (DDR @ 2133 MHz), Intel® Optane™ S5D DC P4800X Series 375 GB and Intel® SSD DC P3700 1600 GB, CentOS* 7.2 kernel 3.10.0-327.el7.x86_64. Performance measured under 4K 70-30 workload at QD1-16 using FIO 2.15.

¹¹ Switches based on Intel® Omni-Path Architecture use a 48-port switch chip versus the 36-port chips used in today's InfiniBand® switches. Clusters can potentially be built using fewer switches and fewer switch hops, which helps to optimize cost, scalability, and performance.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown." 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/datacenter>.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

Normalized performance is calculated by assigning a baseline value of 1.0 to one benchmark result, and then dividing the actual benchmark result for the baseline platform into each of the specific benchmark results of each of the other platforms, and assigning them a relative performance number that correlates with the performance improvements reported.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at www.intel.com.

The products described may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

3D XPoint, Intel, the Intel logo, Intel Optane, Intel Xeon Phi, and Xeon are trademarks of Intel Corporation in the U.S. and/or other countries. *Other names and brands may be claimed as the property of others