


# DU GeoSolutions Leverages Xeon Phi for Improved Seismic Processing

 [www.hpcwire.com/2016/03/23/downunder-geosolutions-leverages-xeon-phi-improved-seismic-processing/](http://www.hpcwire.com/2016/03/23/downunder-geosolutions-leverages-xeon-phi-improved-seismic-processing/)


It's a given that modern oil and gas exploration couldn't exist without increasingly powerful supercomputers. Today's machines enable new capabilities, improve image resolutions, lower costs, and decrease risk. Indeed, many wells, some costing tens of millions of dollars, have missed their targets due to inadequate or incorrect data analysis. This article explores the impact of high-performance computing (HPC) and highly tuned software in DownUnder GeoSolutions (DUG) seismic processing and imaging workflows and how that in turn improves oil and gas exploration.




Fundamentally, seismic processing – turning raw, noisy data into a useful, accurate image of subsurface geology – is a complicated, computation-intensive task. To understand how HPC adds value, it helps to know roughly what the process entails. Our clients use purpose-built ocean or land teams to acquire seismic data. That process uses an energy source (such as an air-gun array or dynamite) to generate acoustic waves, which propagate through the earth, reflect off of geological layers, and are measured at the surface by a recording system.

The raw data can be very noisy, and must be processed to enhance the signal-to-noise ratio of the reflections. Moreover, the true subsurface location of these reflections isn't initially known. A compute-intensive process called migration is used to reposition reflections to their correct subsurface positions. Prior to migration, the apparent position can be literally miles away from the actual subsurface location.

Because our clients subsequently use this data to identify likely hydrocarbon accumulations – and ultimately to decide where to drill wells – accurate subsurface positioning is paramount. Enhanced processing and migration analysis increases the fidelity of the seismic images, which in turn allows for a significantly better understanding of the subsurface geology.



Images on the left show the effect that modern processing and imaging can have on seismic data. Both images were created from the same input data; the DUG result is on the bottom. The hugely increased fidelity allows geoscientists to more accurately evaluate the prospect of hydrocarbons. This kind of processing is necessary for every dataset.



Because seismic data processing is so fundamental to success *and* extraordinarily compute-intensive to analyze, the modern petroleum industry could not exist without supercomputers and the software that drives them. All of the easy hydrocarbons were found a long time ago; we need more-advanced technology to discover and reach the energy reserves that remain. To accurately image complex subsurface geology, our clients acquire more and better-quality data, which then needs to be processed at higher resolution and with faster turnaround. These factors keep pushing compute demands higher, which is why DUG doubles its HPC power every 12-18 months.

Indeed, while DUG is a geosciences company with a broad portfolio of offerings (details below), it is also to an HPC software- and hardware-integration company. We operate a 6 petaflop/s cluster in Perth, another of the same size in Houston, and smaller clusters in London, Kuala Lumpur, Brisbane, and Jakarta – all of which are busy 24 hours a day, 365 days a year. Worldwide, DUG makes extensive use of Intel Xeon processors and Intel Xeon Phi coprocessors and its custom system from SGI is one of the world's larger Intel Xeon Phi coprocessor commercial deployments.

Since 2004 DUG has invested heavily in developing state-of-the-art, highly-parallel processing codes, well-tuned to successive generations of multi-threaded CPUs. We use sophisticated libraries for the basic infrastructure shared amongst all our codes (to manage communication, queues, caches, and so on) along with explicit threading and vectorization, where necessary.

When interacting with the Xeon Phi coprocessors, we begin with Intel's language extensions for offload (LEO); however in order to decouple execution on the host and coprocessor – and to share as much code as possible between them – LEO use is kept to a minimum. It's used initially to spawn persistent threads on the coprocessor along with their associated workflow infrastructure. Subsequently, offload is used purely to feed work packets and input data to this autonomous compute engine on each coprocessor; no useful work occurs within the execution of an offload code block.

From a code structure perspective, the use of offload turns out to be reasonably non-invasive, as most of the offload code segments can be isolated and ignored by algorithm developers. In cases where host memory is at a premium, it's sometimes necessary to subvert the LEO framework and explicitly manage address mappings and data transfers between the host and coprocessor.

The offload model has many practical advantages in terms of software bundling and execution. For example, the single executable and single host model allows existing packaging and scripting tools to remain unchanged. Because our code was already heavily threaded, relatively few changes to the Kirchhoff time migration kernel were required to achieve scalability on the Xeon Phi coprocessor. When the offloaded compute engine is up and running, all 236 hardware thread contexts are fully utilized. In simpler applications, the addition of two OpenMP parallel loop pragmas are enough to fully utilize a whole (four thread) core for each task.

We found that the Xeon processors and Xeon Phi coprocessors were a natural combination when looking to increase our per-node HPC compute density. As important as high-thread-count scalability is to driving performance, it's really the HPC-oriented 512 bit Single Instruction Multiple Data (SIMD) instruction set on the coprocessors that removes many constraints associated with GPUs and enables non-trivial algorithms and optimizations.

For example, efficient seismic processing involves seismic and model data interacting on different resolution (generally non-uniform) grids. The gather-scatter instructions are a natural fit for these vector

operations, but they don't greatly relieve memory traffic overheads. Often a better fit for these algorithms are generic permutation vector instructions that can be used to express vectorization over both grids simultaneously. In the latter instance, mapping and interactions between the grids can occur entirely within vector registers, which significantly reduces memory traffic and improves performance.

### **Moving to the Intel Xeon Phi Coprocessor**

Even before we considered using coprocessors, the DUG software suite was heavily optimized for modern NUMA architectures with relatively high core counts. Most of the classical HPC optimization techniques used in DUG software still apply and, indeed, the benefits tend to be amplified on the Intel Xeon Phi because of the scale of the coprocessor.


As an example of classical tuning techniques at work, we achieved large performance gains by ensuring that the 30 MB L2 cache of each coprocessor actually holds close to 30 MB of unique data. Careful data reordering and vectorization allowed us to keep the memory footprint of even very complex nested loops within the 128 KB of L2 available to each thread.

Other features of the vector instruction set simplified the existing optimized code. For example, the availability of fast 23-bit-accurate floating-point square root and division instructions, which are endemic to seismic processing, remove the need to choose between fast or accurate results. We also make use of the comprehensive set of instructions for floating-point rounding and conversion to/from fixed-point, to express operations such as interpolation.

The abundance of existing sophisticated library infrastructure in C and C++ required virtually no changes to function on the coprocessor. The integrated compiler support and execution model means that the deployment is reasonably transparent.

One noticeable difference with the coprocessor implementation of Kirchhoff time migration was the need to rely heavily on up-front allocations of large memory blocks and manual memory management. Between the overheads of pinning memory for DMA transfers, and the Linux kernel bottlenecks exposed by 236 threads simultaneously allocating and deallocating, it was important to remove as much dynamic allocation as possible. In practice, we use most of the coprocessor memory as a cache.

DUG has a number of applications, including the workhorse Kirchhoff time migration, which run up to 1,000-way thread parallelism per compute node, on hundreds of cluster nodes concurrently. Within a single compute node, adaptive load balancing across multiple coprocessors achieves essentially linear performance gains with multiple coprocessors. Each coprocessor delivers between 1.5 to 2 times the performance of the high-end hosts (which typically have 2 x 10-core Intel Xeon E5-2660 v2 processors), allowing us to increase of the compute density of our



nodes by roughly 8x.

## Future

Accurate subsurface imaging, made possible by highly specialized HPC solutions, are a fundamental part of oil and gas exploration and production. In these extremely difficult times for the industry, our success is due in large part to our total focus on very high-density computing in which Intel Xeon Phi coprocessors play an important role. The more FLOPs that we can pack into a single compute node at full bandwidth, the better the results, faster turnaround times, and better prices that we can offer to our clients. We're always working on cutting-edge petroleum exploration issues, and supercomputing is as fundamental a part of those solutions as geophysics.

## Company Description:

DownUnder GeoSolutions (DUG) is a geosciences company headquartered in Australia with offices in Perth, Brisbane, London, Houston, Kuala Lumpur, and Jakarta. Services include seismic acquisition design and implementation, seismic data processing, depth imaging, petrophysical processing and interpretation, quantitative interpretation services, geostatistical depth conversion, and a complete range of DUG software. DUG operates a 6 petaflop/s cluster in Perth, another of the same size in Houston, and smaller clusters in London, Kuala Lumpur, Brisbane, and Jakarta – all of which are busy 24 hours a day, 365 days a year. Worldwide, DUG makes extensive use of Intel Xeon processors and Intel Xeon Phi coprocessors.

## Author Bio:

Phil Schwan is Head of Software at [DownUnder GeoSolutions](#).



## Share this:

[Tweet](#)

## Related Posts

- [SGI's HPC GM Gabriel Broner on 2016 Trends](#)
- [Podcast: Industry Leaders on the Promise & Peril of NSCI](#)
- [Intel® Omni-Path Architecture and a Flat Universe](#)
- [Colfax International Raises the Bar for Code Modernization](#)
- [Strip-Mining for Vectorization to Achieve Order of Magnitude Improvement](#)
- [The Future of High-Performance Computing Has Arrived](#)





Intel® Scalable