

# Accelerating 3D seismic data analysis on desktop workstations

A new range of graphical processors can bring server cluster performance onto the desktop for under \$10 000. Keith Forward reports.

Eine neue Reihe Grafikprozessoren kann für weniger als 10.000 USD Servercluster-Leistung auf Desktops herbeizaubern. Ein Bericht von Keith Forward.

Una nueva gama de procesadores gráficos es capaz de trasladar las prestaciones de un clúster de servidores hasta el mismo escritorio, y por menos de 10.000 dólares. Keith Forward informa.

Geophysicists often say that they could work better and find more oil and gas if they had access to faster computers on their desktop. Until now, intensive processing tasks had to be sent off to a server cluster, returned and the results visualised on a desktop workstation.

Now, for around \$10 000, you can get the speed of a small cluster in a desktop machine, by using the untapped potential of the GPU (Graphics Processing Unit). A machine with four high end GPUs from NVIDIA has up to four teraflops of processing power, equivalent to around 60 of the fastest CPUs.

Having that sort of computing power on your desktop opens up huge possibilities for software developers.

Stephen Purves, Technical Director at Foster Findlay Associates (ffa), a developer of 3D seismic image processing technology for the oil and gas industry, believes the power of the GPU will be a game changer for seismic analysis applications

The latest release of its Windows-based 3D seismic image analysis and visualisation application, SVI Pro 2009, features an Interactive Facies Classification (IFC) module that allows the user to see the results of what they are doing in real time (Fig. 1).

The IFC is a GPU enabled module that performs computations interactively, on time slices on horizons of data. A user can load multiple seismic attributes and run classification jobs very interactively, moving through the 3D volume in real time.

“We’ve been wanting to do this for years but it has never been practical before, the interface was always too clunky,” says Purves. “There was a lot of waiting around with progress bars on the screen.”

“Now we perform calculations in the time it takes to stream data from disk.”

This is the first time that GPU accelerated processing has been used in this way,

interactively visualising the results of the analysis process, rather than just being used in the background to speed up computations.

## Using the GPU

FfA’s latest release of SVI Pro uses the GPU in three different ways:

1. For general visualisation
2. For interactive (visual) computing
3. For GPU accelerated volume processing

A distinction is made between Primary GPUs, which are used for visualisation and are connected to a screen, and one or more Secondary GPUs that are either not connected to a screen or by design cannot connect to one.

General visualisation refers to the 2D and 3D visualisation and volume rendering capabilities that have always been available within SVI Pro. The workstation’s Primary GPU(s) are used for visualisation.

Interactive computing refers to the process of running algorithms directly on the Primary GPU(s) to compute data and send these results directly to the screen for immediate visualisation, such as in the IFC tool.

The IFC and other interactive computing tools are described as being ‘GPU enabled’ meaning results are computed interactively on slices and horizons and displayed immediately in real time.

GPU accelerated volume processing refers to the process of using the GPU to run a SVI Pro process over a 3D volume, recover the results and

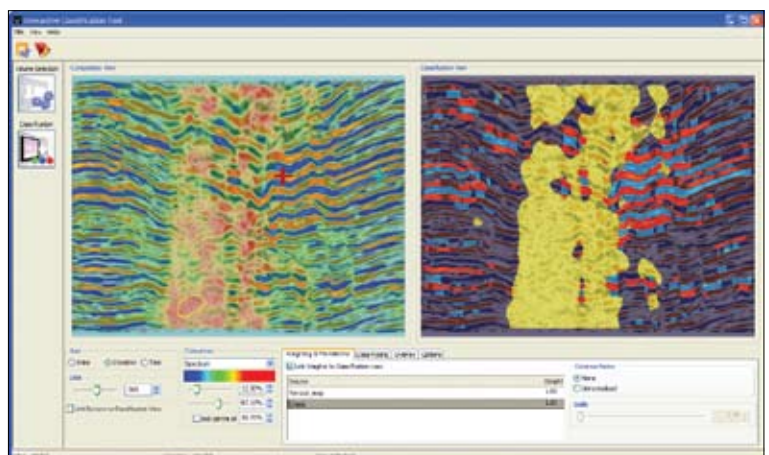


Fig. 1. SVI Pro 2009’s Interactive Facies Classification module is user centric, interpreter driven and makes interactive 3D seismic facies discrimination possible across large seismic datasets.

store these back on disk, just like normal processing in SVI Pro would do.

Purves says ffa has released a bunch of GPU enabled tools already for SVI Pro, the Windows version of the software. The Linux version, where the product is known under the name SEA 3D Pro, was also released in June this year.

It has a lot of the functionality in it, including the IFC workflow, and everything will be synchronized by the end of the year, says Purves.

“We now have GPU accelerated volume processing, where we have taken all our algorithms, for calculating dip, azimuth, noise cancellation on seismic reflectivity, even the standard trace attributes like envelope and instantaneous phase.”

“They were multi-threaded within our standard toolkit, and they would run on multiple core, multi-CPU systems. They have now been ported across to run on the NVIDIA CUDA platform. So far we’ve got about a third of our processes GPU enabled.”

SEA 3D Pro also incorporates an interactive link with Landmark’s GeoProbe 3D Volume Interpretation software, providing GeoProbe users with a range of stratigraphic and structural analysis workflows, directly integrated within their current 3D interpretation platform.

There is about a 50-50 split between Windows and Linux amongst ffa’s users, with a lot of the bigger companies still using legacy systems that were run on Unix and now Linux.

Another aspect of the latest release is the addition of interactive colour blending capabilities. This allows users to combine different attributes and blend the results together by controlling the colour and opacity of each aspect interactively (Fig. 2).

This can help users identify structures within the data and pick out the formations of interest, where oil and gas might be found.

“We’ve done a lot to enable the co-rendering of multiple attributes, to compute and work interactively with them and start extracting geological bodies out of them. That was a big theme behind the Spring release. Of course we can do this much more quickly than before with GPU processing,” says Purves.

### The facies classification workflow

FfA’s Interactive Facies Classification (IFC) is a 3D interpreter driven seismic classification module that enables detailed results to be extracted from an entire dataset in real time. With the GPU enabled processes users are able to see the results of every modification they make instantaneously.

The interpreter can work with any set of seismic attributes or inverted volumes loaded within SVI Pro, over the entire dataset or within user selected regions of interest, including previously isolated geological anomalies.

They can define classes by selecting individual seeds, merged seeds, regions of particular character or groupings of existing classes. The interpreter can




Fig. 2. The Opacity and Colour Blend capabilities of SEA 3D Pro 2009 provide a powerful way to analyse seismic reflectivity data.

*Data is courtesy of StatoilHydro.*

interactively slice through the volume to select a specific seed point on different slices (inline, crossline, timeslice or horizon slice) or to check the classification which is updated on the fly at each parameter alteration.

“Often seismic classification takes something of a black box approach, especially if you are applying neural networks or statistical classifiers,” says Purves. “By using GPUs to perform classifications on demand, it’s like stepping into the black box, allowing the user to change classification parameters and see the results immediately.”

### Gas chimney classification

To understand the distribution of the high amplitude peaks and troughs in relation to a gas chimney, the IFC tool was applied to a Terrace Amplitude volume along with an Envelope (instantaneous amplitude) and a Chaos (amplitude independent structural variability) volume (Fig. 3).

The IFC is performed essentially in three stages, all of which take advantage of the GPU enhanced capabilities. A typical workflow, on a 100 km<sup>2</sup> dataset, can be achieved in less than two hours from initial data loading to final classification, says Purves, while the classification by itself can be realised in a matter of minutes.

First, the source attribute data is loaded and there is a pre-processing step to calculate the volumes that is accelerated using CUDA using the parallel processing capabilities of all the GPUs available in the system.

This is followed by the main interactive step where facies classification is performed, using a single visualisation GPU connected to a display. This step uses a graphics programming library called GLSL, which is a forerunner of CUDA.

FfA is working with a company called Mercury that provides the Open Inventor visualisation package used in

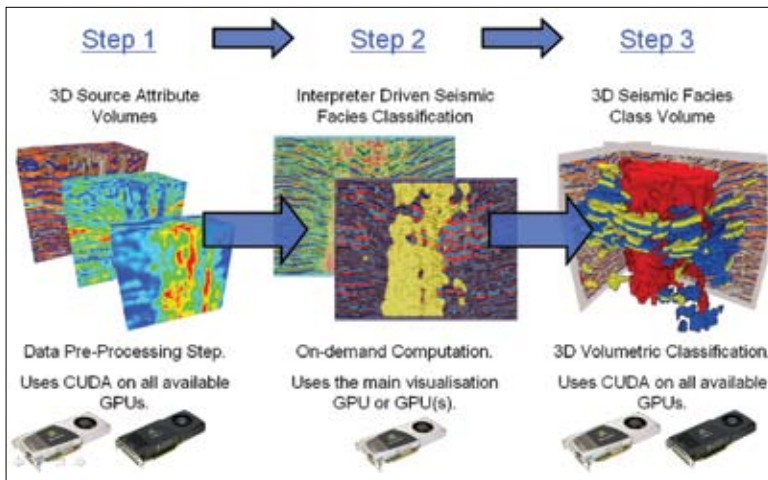


Fig. 3. Using GPUs in the Interactive Facies Classification workflow. Left: three seismic attribute volumes, terrace amplitude, envelope and chaos (top to bottom) are computed using all GPUs. Middle: the interpreter blends the volumes interactively and picks classes using the visualisation GPU. Right: the final facies volume is calculated using all GPUs.

the software to build CUDA into this step. “We expect to have tools out in March next year which uses CUDA in this interactive way. That will allow us to do much more, it’s much easier to write sophisticated algorithms using CUDA,” says Purves.

Finally volume processing, which can often be a slow process on a data set of several gigabytes, is performed again using multiple GPUs to calculate the final volume. The IFC has isolated 3 classes (right): the gas chimney (red), the high amplitude peaks (yellow) and high amplitude troughs (blue).

The complexity of the gas chimney and the distribution of the amplitude anomalies are now easy to interpret from the detailed 3D representation.

### Benchmarks

ffA is working on a series of benchmarks for various common processes, comparing GPU enabled calculations with fast CPUs, and also looking at the speed improvements from upgrading older machines (Fig. 4).

This will give customers a guide as to how much they need to spend to take full advantage of the GPU enabled processes, given their current hardware configuration

Fig. 5 shows the speed improvement for three typical processes: TDiffusion, a complex noise cancellation calculation; DipAzimuth, a structural attribute; and Trace Attribute, a standard trace envelope.

The comparison is with a base system with two quad core Intel Xeon processors and 16Gb of memory. As you add more processor cards, the performance improvements continue to ramp up.

A twelve times improvement might not sound that impressive, but it means for example a 36 minute job coming down to a three minute job.

Once all the algorithms are CUDA enabled, ffA will start to work on optimising them, which should give a

further performance boost for most processes. NVIDIA are also planning to double the speed of their processors every 18 months.

### Using GPUs for spectral inversion

OpenGeoSolutions, a seismic data analysis company that pioneered the use of spectral decomposition and inversion, has reported a 55 times performance increase after upgrading to the Tesla C1060 GPU.

“We are measuring speedups from two hours to two minutes using CUDA and the Tesla C1060,” says James Allison, president of OpenGeoSolutions. “This kind of performance increase is totally unprecedented and in a market where there is great economic value in being able to determine these fine sub-surface details, this is a game changer.”

On a CPU-based cluster, the seismic inversion process could take anywhere from 2 hours to several days, Allison says. In an effort to improve this, the team acquired a workstation equipped with an NVIDIA Tesla C1060 GPU. Over six weeks, the OpenGeoSolutions team converted a key portion of their application to CUDA.

“The Tesla products essentially give us all a personal supercomputer,” says Allison. “Just one Tesla C1060 delivers the same performance as our 64 CPU cluster, and this was a resource we had to share. This is a huge cost and time saving that has transformed our workflow and boosted our productivity.”

### NVIDIA and CUDA

NVIDIA is best known for producing graphics cards for gaming and CAD (Computer Aided Design) applications. But the humble graphics card in many a home gaming PC can be put to use to speed up common computational problems many times over.

GPUs (Graphics Processing Units) are basically multiple core, massively parallel processing chips that can perform many operations at the same time, something that is not possible with the CPU that runs a PC. They also use dedicated fast communications interfaces, meaning the data can get in and out fast enough to take advantage of the processing power.

CUDA is the first C language environment that enables programmers and developers to write software to solve complex computational problems in a fraction of the time by tapping into the many-core parallel processing power of GPUs.

It makes it easier to code sophisticated algorithms to use the graphics processing chip, and performance benefits compared to standard workstation of over one hundred times are not unusual, according to NVIDIA.

It has been developed by NVIDIA to enable wider use of its GPUs for a range of applications, including those outside its traditional gaming and visualisation market.

### The latest processors

CUDA can be used on a wide range of NVIDIA products, from home gaming graphics cards to rack-mounted server

Performance improvement for Tesla C1060 per workstation	Dell Precision 690	Dell Precision T7400	HP xw8400	HP z400	WSX 208
Trace Attributes	45.5x	29.1x	30.6x	23.7x	20.1x
Chaos	12.6x	14.8x	10.9x	8.8x	2.2x
DipAzimuth (Gradient)	27.8x	17.7x	15.8x	10.9x	7.8x
DipAzimuth (Phase)	-	-	-	-	-
DA Combine	35.8x	27x	14.7x	12.8x	2.4x
Instantaneous Phase	38.3x	20.7x	21.8x	17.3x	13.2x
RGB Export	-	-	-	-	88.7x
SO FMH	32.9x	13.7x	17.3x	8.5x	11.9x
SO Noise Filter	24.4x	22.5x	20.5x	11.2x	14.9x
TDiffusion	9.7x	6.1x	6.1x	4.8x	3.7x
Tensor	23.7x	15.7x	13.5x	10.8x	3.2x
Volume Combo	23.6x	19.4x	7.6x	9.1x	1.4x
IFC Normalisation	-	-	-	-	7.6x
IFC Volume Classify	-	-	-	-	-

Fig. 4. Comparison of a Tesla C1060 workstation with legacy systems. The table shows the speed improvement of GPU Accelerated processes on a 1GB dataset running on a compute only (no screen connection) Tesla C1060 GPU Card as a Secondary GPU versus the same process running on a range of workstations.

clusters. A typical oil and gas application would use NVIDIA's Tesla range.

The latest chip for workstations, the Tesla C1060, is capable of one teraflop of processing power, and has 4GB of dedicated memory. It has 240 processing cores and uses just 160W (Fig. 6).

A typical workstation could use four GPUs, giving four teraflops of processing power, equivalent to around 60 CPUs.

NVIDIA has just brought out a 1U rack-mount system, the S1070, that features four Tesla GPUs in a single unit,

giving four teraflops for a maximum power consumption of 800W. This can easily be scaled up to cope with even the most intensive processing tasks.

GPU based servers are cheaper than CPUs for equivalent processing power, and have much lower power consumption, reducing running costs.

NVIDIA and Supermicro have also developed a two teraflop 1U server that uses two quad core Intel Xeon processors with two C1060 NVIDIA GPUs. Supermicro claims a 12 times performance boost compared to a traditional quad-core CPU-based 1U server.

Petrobras, the leading Brazilian International Energy company, recently spoke about its reliance on Tesla GPUs to increase the performance of its seismic data processing. Petrobras has invested in a GPU-based cluster consisting of 190 NVIDIA Tesla GPUs.

“With our GPU cluster we are getting performance improvements of 5x to 20x over our traditional multi-core CPU-based cluster,” said Neiva Zago, Geophysical Technology Manager, Petrobras. “We expect that the continued use of GPUs in our business will result in significant reduction in processing time as well as savings in power consumption and datacenter floor space.”

Petrobras expects scalable increases in GPU performance will continue as it builds on its datacenter to deliver more than 400 teraflops. ●

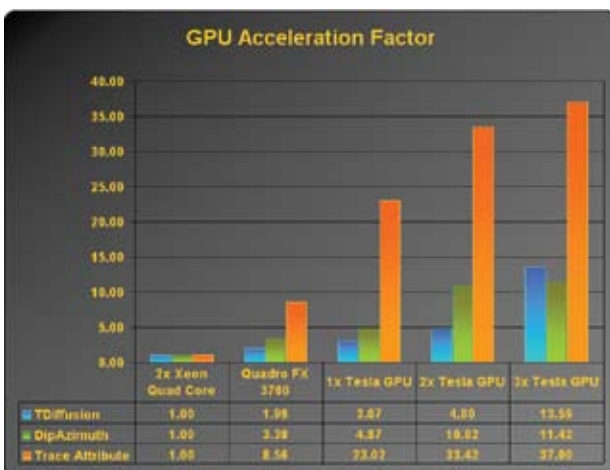


Fig. 5. Significant gains in performance of 3D seismic analysis algorithms are achieved with the Multi-GPU enabled SVI Pro 2009. The graph shows initial benchmark performance improvements measured for a number of SVI Pro's algorithms for noise cancellation, stratigraphic and structural analysis of 3D seismic data, using a modern multiple GPU system. Speed increases of up to 23x are seen on a single GPU and these stack up to a 37x increase when using 3 GPUs in parallel.



Fig. 6. NVIDIA's Tesla C1060 Graphics Processing Unit (GPU) has 1 teraflop of processing power and 4GB memory.