Offshore®

Seismic visualization supercomputer brings subsalt data into the light

Complex algorithms require new generation of high-performance hardware

Francisco Ortigosa, Repsol YPF

The Kaleidoscope Project is designed to use new generation seismic imaging algorithms on the next generation of high-performance computers.

The goal is to make it possible to derive the best possible decisions from an ever growing volume of data generated by today's subsalt surveys.

Also, it is a technological response to the competitive business environment in the Gulf of Mexico. The region has a rising volume of seismic data and significant remaining oil reserves, and is of interest to Repsol. Given the cost of a well in the GoM can easily be over $100 million, the new technology provides the company with a competitive advantage.

New seismic imaging algorithms are required for success in the GoM. This new generation of algorithms cannot be processed on the historical computer hardware configuration, so, alongside developing the new seismic interpretation techniques, we needed a new generation of hardware capable of making best use of the new interpretation algorithms. To make the hardware economical, it needed to be based on commodity priced, off-the-shelf technology.

US GOM sea bottom topography.

The context that drives Repsol to create this project also has to do with the challenges of deepwater, subsalt exploration in the GoM, where locating oil reserves below the salt can be extremely difficult.
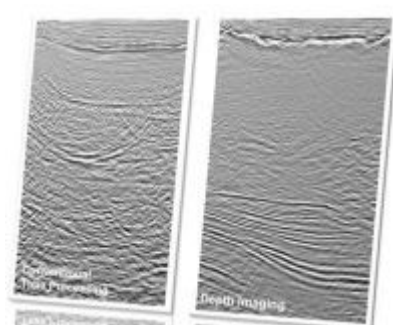
Subsalt imaging in the GOM.

The US GoM extends about 750,000 sq km (289,577 sq mi) between shallow and deep waters. Shallow water exploration in less than 200 m (656 ft) water depth was the focus of exploration from the 1960s through 1980s. But currently, as technology evolves, exploration is continuing into deep and ultra deepwater.

The sea bottom in deep and ultra deepwater in the GoM is characterized by irregular formations. These irregularities are the

sulfate expressions of abnormally thick salt bodies underneath. A northwest to southeast cross-section of the Gulf shows these salt bodies. A representation of the remaining reserves in the Gulf shows them to be mainly subsalt.

Operators in the GoM need to "see" underneath these salt bodies to locate reserves.

The task of trying to image below the salt bodies with conventional time processing technology is impossible. To make it possible to derive useful images, a new generation of algorithms is required.

Many different imaging algorithms exist today. In the late '90s and early 2000s, the standard algorithm used was the Kirchhoff equation. Kirchhoff is good for the salt flanks and for reserves related to high visibility structures on top of salt. However, for reserves below the salt bodies, Kirchhoff algorithms failed to image the structure properly because of the nature of the reflections and of the multi-value pointed away from that average originated below the salt bodies. Wave migration algorithms were developed to properly image this section. However, this was only possible because of the improved processing performance of PC Linux clusters. The increased performance of these clusters enabled the oil and gas industry to shift from Kirchhoff algorithms to the different wave migration algorithms used today.

Kirchhoff vs. wave equation. (Below) It can be even more difficult.
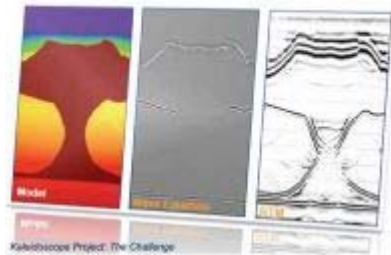
Conventional computing technology has come a long way since Kirchhoff, but even more complex algorithms are required for salt bodies in the new GoM exploration areas. Wave migration algorithms, because they basically are a one-way approximation to the wave equation, failed to image highly steep events in the subsalt. Imaging correctly requires two-way approximations to the wave equation, such reverse time migration (RTM). But RTM requires computing power an order of magnitude higher than wave migration. It is easy to see how computing power and imaging algorithms align in the GoM.

The limitation of PC clusters is due to the "power wall" of the processors. It is impossible for any processor to go beyond 40 gigaflops in frequency flow, roughly.

The Kaleidoscope Project is a solution to this problem. The project works simultaneously in generation of hardware in new platforms with increased capacities for the heavier processing need of the algorithms used for subsalt exploration. At the same time, the project works to develop this new modeling tool to provide seismic detail.

It is important to note the computers that we have used come from mass markets which do not target high-performance computing, in particular from the desktop and the server market. This hardware technology will not take us further because the processor technology based on the '80s processors has reached the power wall. To run the new algorithms we need new hardware technology. To be economic, that technology has to be based in commodity, off-the-shelf markets. That is the Kaleidoscope Project aim.

The project was created by Repsol to promote innovation. It includes a seismic company, 3D GEO; scientific research institution, Spanish Scientific Council for Research; and the Barcelona Supercomputing Center.

The Kaleidoscope Project encompasses a simultaneous innovation of hardware and software to achieve a petascale solution to seismic imaging using off-the-shelf technology. Software research focuses on the quality of algorithms and avoids shortcuts or tradeoffs common because of lack of computing power. Kaleidoscope is to ensure the maximum possible imaging quality regardless of the computer power required. And of course speed and power, which are the two main factors for the project to succeed, are guaranteed while ensuring low cost because it's coming from a massive market.
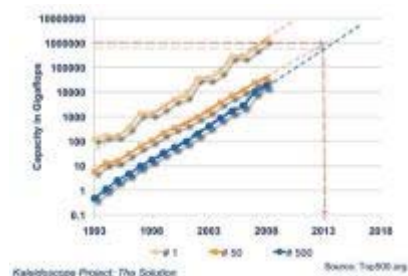
(Above) Top 500
supercomputing evolution.
(Below) The right moment
to invest in imaging.
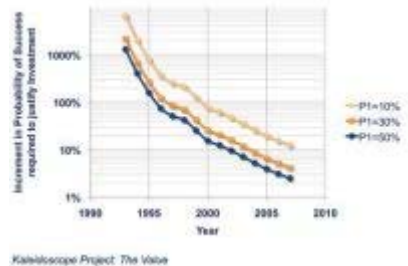
One very important factor in the project is the Mare Nostrum supercomputer. When the project started in 2006, the Mare Nostrum was the fifth largest supercomputer in the world. Repsol has used this supercomputer to test all the algorithms without assembling any circuits. Once the algorithms work, the hardware production phase starts.

Is it reasonable or even possible to reach peta-scale capacity in the oil industry?

There is a continuous linear trend of evolution in computer performance. In 2008 we reached the philosophical and technical barrier of 1 petaflop at the Roadrunner supercomputer in Los Alamos National Laboratory. If we extrapolate the rate of technological development in computing, in 2013 some 50 of the top super computers of the world will be at about 1 petaflop.

This implies that in five years virtually every one will be able to have 1 petaflop supercomputers. The question then turns to whether there will be enough software or enough algorithms that scale to 1 petaflop.

Sustainable peta-scale capacity requires a multi-core processing solution because no single core will reach frequencies higher than 4 gigaflops. So once we were multi-cores, assume it will be a hybrid supercomputer.

Important to this solution is low power consumption. Otherwise, each supercomputer in each office will need a big power source next door. Obviously this is not an economical or a sustainable solution. To flourish and be applied extensively, this technology must be an off-the-shelf commodity. This means it must be developed for a much larger market than the one for high-performance computing (in other words, the consumer electronic market). This also means there will have to be support to adapt our codes to this hardware.

The chip that currently fulfills all these requirements is the cell processor developed by Sony, IBM, and Toshiba. This processor is multi-core in the sense that it has different multi-cell units that work independently. Currently, this chip is aimed at the consumer electronics market, and is in the Sony Playstation 3, in televisions for Toshiba, and by IBM in the PS 20, PS 21, and PS 22.

The chip has a main power processor element and eight sub-units called synergistic-processing elements, which are like accelerators for the main processor. These accelerators are based on technology used in graphic-processor units. The cell processor has 256 gigaflops performance with bandwidth of 75 gigabytes per second of input/output.

The numerical difficulty, for instance, of 2D Kirchhoff PSTM compared to the image intensity in 3D RTM is two orders of magnitude. Other solutions are in testing. Repsol has RTM on production on Cell B platform.

The economic value of Kaleidoscope is in that it gives better images leading to better decisions. During the last 15 years the NPVs of the projects have increased not 10 times but it is not completely linear with the cost of the oil. At the same time, the cost of drilling in the GoM has increased 10 times, and the cost of computing power has decreased 25 times.

If we graph this over time, we see that during the '90s, the incremental probability of success used to justify investment in seismic imaging research was so high as to be unrealistic. That meant that it did not make economic sense for oil companies to invest in this kind of R&D. So, service companies with broad client bases were able to invest and to distribute the costs among the client base.

In the GoM now, it makes sense to invest in this kind of technology or research because a 10% – 15% increase in the probability of success with this technology makes it economic.

Editor's Note: To watch the complete presentation about the Kaleidoscope Project, go to www.offshore-mag.com and look under "More Webcasts".

Why 'Kaleidoscope'?

Everybody always asks why the name "Kaleidoscope". The name has many things to do with the project. It is looking at images, seismic images in particular, but from a different perspective -- like with a kaleidoscope. This perspective uses the same data but with new generation seismic imaging algorithms tailored to the next generation of high-performance computing hardware.

To access this Article, go to:
http://offshore-mag.com/offshore/en-us/index/article-display.articles.offshore.volume-69.issue-6.geology-__geophysics.seismic-visualization.html