Higher Frequencies Improve Structural and Stratigraphic Resolution

ION thinks outside the box to provide high-fidelity seismic images

Contributed by ION

Any play we pursue starts with an understanding of the structural elements. Take a complex geologic environment like the Gulf of Mexico. Through advancement in several technologies, notably wide azimuth (WAZ) acquisition and reverse time migration (RTM), we have been able to identify the major subsalt structural features and in turn drill a number of successful wells. However, the challenge that currently faces us is to appraise and develop assets that have been discovered. To do this effectively, we need a robust knowledge not only of the structural elements of the asset, but also of the stratigraphic elements.

Our challenge is to improve the resolution of the seismic images we are using. In the Gulf of Mexico, we have relied on WAZ and RTM to image the subsurface. No other technologies have been as successful. However, running RTM to the frequencies that we need to build effective reservoir models has been problematic — until now.

From a theoretical standpoint, the difference in run time for a 20-Hz RTM versus a 40-Hz RTM is a factor of 16. Similarly, to output an 80-Hz RTM would take 16 times longer than a 40-Hz RTM. Which means, that to output an 80-Hz RTM rather than a 20-Hz RTM, would take 256 times the compute power. This is because the model size for the migration needs to grow while the time step of the wave propagation needs to decrease accordingly to maintain stability and accuracy at the higher frequencies. This brings about the more significant problem — memory.

The typical implementation of RTM is on GPUs. These are very fast but have very limited memory, which means that approximations must be undertaken. ION has implemented its RTM on conventional compute nodes. This means much more memory is available — more than 100 times more memory per core — allowing us to perform high-frequency RTMs without approximations or migration-induced noise.

The benchmark for high-frequency migration has been Kirchhoff. The examples shown here demonstrate that a high-frequency RTM, if implemented correctly, can yield the same fidelity as does a Kirchhoff. The real benefit of this test, however, is that when the geology becomes too complex for Kirchhoff, ION's high-frequency RTM continues to provide a superior high-fidelity image.

Why go to 80 Hz if the data doesn't support it? The answer is quite simply that you wouldn't. The issue is that if the data does support it, you don't want to be restricted to only 20 Hz. In addition, with newer approaches that compensate for attenuation and suppress noise, we have experienced that some data sets actually support much higher frequencies than previously thought.

Finally, higher frequencies also benefit reservoir characterization workflows. By using a higher frequency migration for acoustic and elastic inversions, you improve your prediction of the reservoir and your ability to identify the sweet spots therein. For more information, visit ION in booth 3124. ■

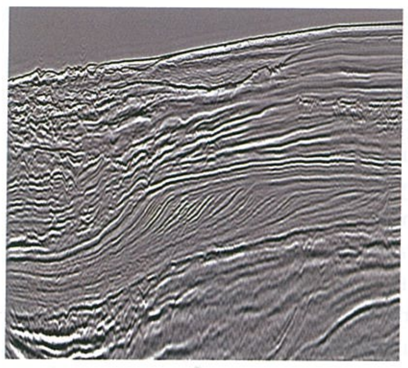

45 Hz Reverse Time Migration (RTM)

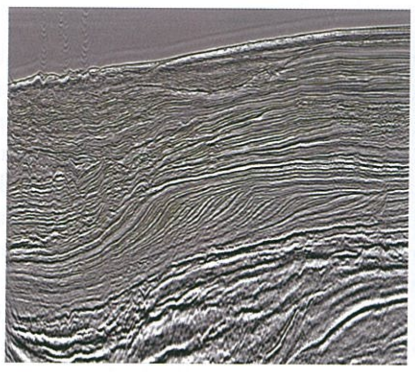

200 Hz Kirchoff Migration

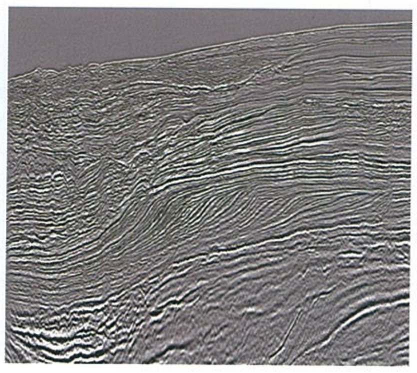

200 Hz RTM

The Anatomy of a Supercomputer Named Bubba

DownUnder Geosolutions, and Bubba the supercomputer, are set to move into new Houston, Texas, office space next month

Contributed by DownUnder Geosolutions

DownUnder GeoSolutions (DUG), an Australian-owned geoscience service company, recently announced its considerable expansion into a new office in the Houston energy corridor. This move, timed for November, will see the company more than triple in size in the region.

One of the move's more exciting developments is the installation of the company's supercomputer, on site, within the office environment. The Houston supercomputer, affectionately named "Bubba" by the geos at DUG, will grow in size and power from its current 1.6 petaflops to a massive 5 petaflops of power.

DUG is no stranger to supercomputers of this size. The company's Perth-based supercomputer, "Bruce," also boasts greater than 5 petaflops of power and currently is being used to process one of the oil industry's largest-ever multiclient seismic projects. The processing and imaging of the Capreolus Multi-Client 3D survey includes fast-track and TTI PreSDM outputs in a total time of 12 months, a mere three months after survey completion. The survey consists of double-density data (12.5-m shots, continuously recorded) over 24,000 square km across the Roebuck Basin. The raw dataset is greater than half a petabyte of data. This is a joint Polarcus/DUG project. BHP Billiton's petroleum president, Tim Cutt, has said beneath their licensed Capreolus data they have already identified four leads "that each have the potential of over 400 million barrels recoverable."

Bruce is touted as the largest supercomputer in the Southern Hemisphere: five times larger than the better-known scientific installation at the Pawsey Centre in Australia and three times larger than the proposed CSIRO Bureau of Meteorology system to be built in 2016. At 5 petaflops capacity (that's 5,000,000,000,000,000 floating-point operations per second), Bubba likewise will be a significant machine in the Northern Hemisphere — especially considering that the computer room (and the chillers) in the new DUG Houston facility are set up to allow Bubba to be expanded to 20 petaflops and beyond.

It is possible that these machines process seismic data faster than any other machine on earth. Not only do they have the capacity to process one of the single largest-ever acquired seismic surveys (Capreolus), but they also process many other significant geophysical projects at the same time.

The quest for a more reliable, cost-effective, and energy-efficient cooling solution has led DUG to introduce revolutionary oil cooling technology. Many parts of Bruce and Bubba are submerged in a cooled mineral oil bath, which makes them not just cool but also two of the greenest machines around. All the cool technology enables the supercomputers to be 40% bigger within the same power envelope and offers far greater reliability under extreme temperature conditions. They can continue to process data, uninterrupted, even when the temperature surpasses 110 degrees Fahrenheit.

With a high-speed, low-latency, full-bisection-bandwidth 10 Gb network, each part of Bubba will have the ability to communicate with the others without contention. This, combined with a Lustre-based cluster file system running on Netapp hardware, enables the machine to process data from disk at a massive 35GB/s (35,000,000,000 bytes per second).

DUG's new supercomputer in Houston is no ordinary machine. Bubba's brain consists of hundreds of thousands of processing units packed up in thousands of Intel Xeon Phi Co-processors. The computer uses more Xeon Phi's than any other commercial supercomputer on earth and runs highly optimized software to extract every last flop. In fact, through DUG's recoding of their software to optimize the power of the Phi, they are able to take advantage of 95% of the Xeon Phi's theoretical power.

Not content with just a good brain, Bubba has excellent memory. Not a single bit goes unnoticed. His ECC memory ensures exactly the same result every time — unlike commodity GPU's, for example, which do not use ECC RAM and therefore sacrifice accuracy for speed and price.

Storage and compute nodes in traditional water chilled racks.

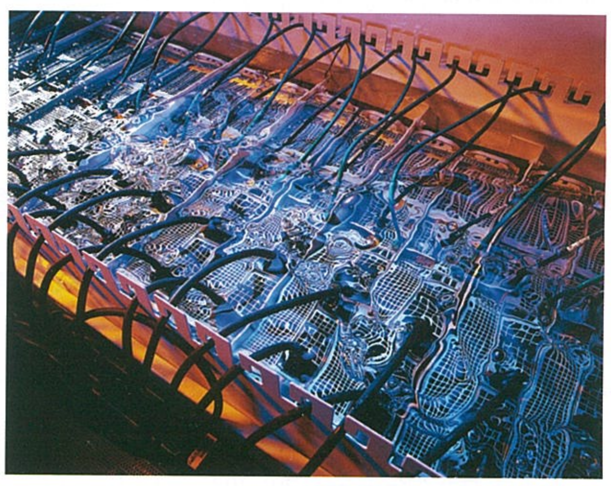

Compute nodes submerged in an oil bath. This saves up to 40% of the power costs.

Bubba does not eat quiche! He is big, brawny, and brainy. And he will be helping DUG turn around large projects for the Houston oil industry at speeds like never before. He's Bubba by name, big by nature, and has a lot more growing to do. To learn more, visit DownUnder Geosolutions in booth 2034. ■