

HARTENERGY

E&P

EPMAG.COM

FEBRUARY 2013

**Reservoir
Characterization**


**Wellhead Pad
Design/Equipment**

**Subsea Processing
Advances**

Enhanced Recovery

Drill Bit Technology

**REGIONAL
REPORT:
South
America**



HYDRAULIC FRACTURING UNDER PRESSURE

**New methodologies
prop open more
opportunities**

GPU usage grows in digital oil field

E&P sector's appetite for data drives IT solutions.


Velda Addison, Associate Online Editor

Experts are touting GPUs as a faster, more economical data solution for use in E&P operations. IT companies are stepping up to meet the challenge of handling large amounts of data, including in remote 3-D visualization environments.

Larger dataset sizes and computational needs driven by techniques to better understand the subsurface have led to a wave of development to meet digital oilfield needs.

E&P companies are finding data solutions by turning to systems developed by companies such as NVIDIA, Cisco, and NetApp. Technical experts from several companies, including Petrobras, participated in Hart Energy's three-part "Big Data and the Cloud" webinar series.

"It's not enough just to have good seismic data. You need to have the content to be able to make better drilling decisions and well path locations. So you have data from electromagnetic surveys to help correlate that," said Ty McKercher, a high-performance computing solution architect for NVIDIA. "And every time you add another data type or a different data type or a different discipline, it introduces delays in the system."



Computationally intensive seismic processing algorithms such as Paradigm's VoxelGeo GPU-based attribute extraction and rendering platform can help reduce delays and allow diverse teams to collaborate and view that data. (Image courtesy of Paradigm)

That is where the GPU system steps in to help reduce delays and allow diverse teams to collaborate and view that data. Such systems are affordable and allow the leverage of existing knowledge to a familiar tool, McKercher said. The technology is being used in complex seismic data processing techniques such as reverse time migration (RTM), wave equation migration, Kirchhoff time-to-depth migration, and multiples elimination.

He noted that it is important for GPUs and CPUs to cooperate via application codes. This enables tasks to be performed simultaneously by the two. For example, data filtering can be done on the CPU while doing computations on the GPU.

Paulo Souza, Geophysical Technology, Petrobras, spoke on the benefits of using hybrid computing for seismic processing. Petrobras started using GPUs in 2006 after moving from mainframes. Now GPUs make up more than 90% of the company's processing power, Souza said.

When using RTM for imaging complex structures on GPUs, Souza said the velocity field is read once per job. Groups of GPUs are used to process a group of shots one shot at a time per group. And a group of shots is stacked in memory before going to disk about every three to six hours.

The process involves breaking the data into small pieces and starting a pipeline to overlap four communication stages.

"We have the GPU calculating the bulk of the model. We are moving data from the GPU to the host. Also, we are sending the data to the neighbor, receiving the data from the neighbor, and copying the data from the GPU."

By using GPUs, Petrobras saw gains of up to 10 times in performance per price and per watt over traditional architectures.

Growing appetite for data

When it comes to making room for common applications such as email and databases along with highly technical applications such as those involving seismic data and well planning, handling massive amounts of data has presented challenges in the digital oil field.

Many are finding answers and learning ways to virtualize the resources with cloud technology as IT companies tout advances such as remote 3-D visualization capabilities and associated supportive systems as well as other applications.

“A lot of the components necessary to create an effective 2-D and 3-D visualization environment have existed separately for quite some time now,” said Stuart Lowery, business development manager of Data Management and Infrastructure for Paradigm. There are high-end work stations, large amounts of memory, high-end graphics cards, and fast network connections to shared network storage devices. “The game-changing event is the evolution of GPUs for both visualization and computation,” he said.

That, coupled with flexible storage and storage virtualization, creates the building blocks to create a remote 3-D visualization environment, he said. “The new challenge is that 3-D graphics capabilities are pervasive and now seen as critical to E&P workflows. Systems and software are available to support this remote 3-D visualization.”


For example, there are widely available graphics capabilities on many platforms, powerful traditional desktops, emerging devices for mobile users, and massive compute power in datacenters, he said. “Now we’re starting to see hybrid systems with coupled CPUs and GPUs and even single chips with both.”

Software is capable of combining data from different repositories from multiple sources, displaying them in 3-D on remote desktops, and allowing them to be shared by multiple users. That could prove useful when handling large amounts of high-resolution data and well log images as well as raw data coming from the field. It has even greater importance considering the industry’s mobility.

“The key is to pool the resources so they can be virtualized and tightly coupled in a secure manner,” Lowery said. “The datacenter of the future includes a shared data repository as reliable storage, compute power, and graphics that can be shared that are in close proximity to the storage network infrastructure that connects those. The real key here is they have to be scalable and flexible in providing the low latency required to support remote visualization in an interactive way for the user.”

Uncovering workflow trends

The trend for GPUs in the workflow is moving from the graphics side for interpretation to seismic processing and flow simulation, and efforts are underway to integrate the two with new cloud technology, said Keith Cockerham, energy business development manager for NVIDIA. He said that hardware to handle large, expensive data has been secured; however, data sizes are outgrowing screen resolution. To overcome the challenge, Cockerham said the company has done compression and color space conversions on the hardware itself and increased algorithms for pixel compression, among other steps. The company also has developed software to handle the increasing size



NVIDIA's Kepler GPU is the first designed specifically for cloud computing. (Image courtesy of NVIDIA)

of datasets, giving power users their own GPUs and memory footprints while others – such as those who do not require much graphic power – share.

As the data volume grows along with the need for collaboration to empower users and handle real-time information, Peter Ferri, energy industry director for NetApp, said the pressure is on IT to get it right. There are operational challenges such as security and time constraints posed by trying to copy large amounts of data to local storage.

“There is a proliferation of siloed infrastructure areas that is making it very difficult for companies to efficiently provision services or make changes across multiple software architectures with unintegrated management tools [along with] the inability to fully leverage assets, which leads to increased costs,” he said. However, the company’s customers are moving from such rigid silos to a more service-oriented infrastructure.

One solution involved working with Cisco on its FlexPod integrated system, which John Thomas, technical solutions architect, Data Center group for Cisco, explained. The system’s features include extended memory for faster rendering, bigger datasets, more desktops per server, and low latency.

Nearly 40% of all large enterprises have virtualized these services, Ferri said.

Using Hess Corp. as an example, Ferri said NetApp has enabled the company to support four times more 3-D seismic interpreters. The company increased efficiency by decreasing backup time, decreasing data loading performance from 20 minutes to one to two minutes, and reducing storage capacity needs by 30%, according to Ferri.

“The key enabler is virtualization,” he said. “It’s an evolutionary process.” **ESP**