

GPUs prove their worth in seismic interpretation

Seismic interpreters are encountering record-breaking performance by using GPUs for interactive computation and rendering.

**Huw James, Evgeny Ragoza,
and Tatyana Kostrova, Paradigm**

The latest 3-D graphics cards can have more than 400 compute cores and up to 6 GB of graphics memory. This contrasts with 128 GB of RAM and 16 compute cores on the processor board of a high-end workstation. It normally is assumed that the central processing units (CPUs) are the compute engine in a workstation, but from the numbers above anyone can guess that 400 simple compute cores in a graphics processing unit (GPU) can rival 16 complex cores in the CPUs.

Today 3-D graphics card specifications are the result of regular and fast-paced doubling of the number of cores and memory sizes over several generations. This latest doubling of memory size to 6 GB means that for the first time there is enough memory on the graphics card to store meaningful amounts of seismic data instead of just graphical representations of those data. The power of these 400+ cores can be used for computation and volume rendering instead of just painting graphics on the workstation's screens.


The Barnett shale in the Fort Worth Basin of North Texas currently is one of the most actively pursued shale plays in the US. It is overlain by carbonates and shales of the Pennsylvanian Marble Falls group. Since the Ellenburger below and the Marble Falls above are both water-rich, it is important to avoid generating fractures that penetrate these two formations by opening existing faults or by operating too close to weakened areas over sink-holes.

Computation


The 3-D seismic volume in a reflection amplitude version is 750 inlines, 750 crosslines, and 500 depth samples for a volume size of just less than 300 MB. The Barnett interval lies in the middle of the depth range and is bounded by strong red and blue reflectors at top and bottom. This volume display has some opacity and lighting so the apparent "section" has some depth of a few lines. There is evidence of faulting at the left of the Barnett interval and hints of faulting elsewhere.

This volume easily can be placed completely into the GPU memory, where it can be transformed to instantaneous frequency instantly. The faulted region at the left of the interval is composed of several separate slivers. Faults that cross the whole interval can be identified clearly. The apparent resolution in depth appears to be doubled.


Lighting acts as spatial derivative; the mild opacity allows a few inlines to be averaged, which acts as a smoother. The result of this is a more interpretable structural image, and attribute computation is instantaneous for this volume size. This speed offers interpreters the freedom to experi-



A volume display of reflection amplitude shows the Barnett interval in the middle of the depth range. (Images courtesy of Paradigm)



The volume display of reflection amplitude in the Barnett interval in the middle of the depth range is shown as instantaneous frequency, with lighting and modest opacity.



A detailed volume sculpt of the Barnett interval is rendered with variable opacity.

ment with effects to generate the most interpretable image, which needs to be experienced to get the full impact. Bigger volumes still are quick, just not instantaneous. Volumes larger than GPU memory can be rendered; the rendering just has to be staged.

The results of attribute computation usually are not saved; it is quicker to regenerate the results than to save and restore the results. This sounds wasteful, but it is not. The 400 cores in the GPU have an extremely potent capability, but they need data to process. This makes memory space in the GPU extremely valuable, so good housekeeping is important. Only data that truly are needed to be saved onboard the GPU should be kept, and data that need to be saved anywhere should be returned to the CPU. For the CPU, memory capacity is large, though compute power still is scarce. It will take some time to adapt to this change in the value of resources between the GPU and the CPU.

Volume rendering

Historically, volume rendering has been accomplished by loading the volume to CPU memory and then processing it on the CPU to generate graphics data to send to the GPU across a relatively slow CPU-GPU interface. In such a system, every time the picture content changes, the display data have to be regenerated and retransmitted, leading to slower interaction. With 6 GB of graphics memory, whole volumes or significantly sized trimmed volumes can be placed into GPU memory and then rendered using algo-

rithms tailored to full-bandwidth seismic data rather than general purpose geometry and images. Changes in display content can be accomplished completely onboard the GPU with no need to reload data until a different volume is to be displayed. In this way, traffic across a relatively slow interface can be avoided, and interactivity can be kept high.

This volume can be rotated, scaled, and edited for color and opacity interactively in real time. The whole volume is rendered even though the interval appears limited. This takes a huge number of operations, one for each data point in the volume. With the latest hardware, rendering can be done eight times faster than with the previous generation of graphics cards. The 300-MB volume example was drawn at 15 frames per second. This speed provides ease of use since a mistake in a parameter setting has negligible consequences and can be corrected easily. This reduces user stress, meaning these interactive devices can be handled with a more creative, trial-and-error approach instead of the control required to drive a conventional, slow system. Once a good view is obtained, the quality of the rendering can be improved by using the processor power of the GPU cores.

With data placed into GPU memory, the compute capacity of the GPU cores can be exploited fully. Results can be shown directly to the user with no need to return them to CPU memory. GPUs can be expected to take on more of the processing burden, favoring direct interaction with the interpreter. **ESP**

Acknowledgements

We would like to thank Marathon Oil Corp. for providing the Barnett shale data.