

A new method of volume rendering applied to a seismic dataset of the Barnett shale

Huw James,^{1*} Evgeny Ragoza¹ and Tatyana Kostrova¹ describe how the dramatically improved speed and memory of today's graphics processing units are bringing a new dimension to the rendering of 3D seismic datasets.

Volume rendering is a very intensive graphical process, resulting in relatively slow performance. The manufacturers of 3D graphics cards have been relentlessly improving the bus speed and increasing the size of the graphics memory on these cards. Now the capability of these cards has increased to the point where useful sized volumes of 3D seismic data can be downloaded onto the card and then rendered locally. This provides dramatic improvements in both rendering quality and speed.

We need to describe the technology in order to illustrate that this change is more significant than just another increase in graphics performance. The latest graphics cards have become a significant computing resource that can outperform the central processing units of modern workstations on some specific types of calculations. For interactive applications the graphics card lies between the user and the central processing unit so the choice of which processing unit is 'central' may become a moot point. The images that illustrate this article use seismic data for exploration of the Barnett shale.

The Barnett shale in the Fort Worth Basin of North Texas is a Mississippian marine deposit that lies unconformably on Ordovician limestones of the Viola/Simpson formations and dolomites of the Ellenburger group. The Barnett shale is overlain by carbonates and shales of the Pennsylvanian Marble Falls group. The Barnett shale is one of the most actively pursued shale gas plays in the USA. The depth to the top of the Barnett shale varies from about 2500 ft to 8000 ft and it can be up to 1000 ft thick. Porosity is low, less than 6%, and permeability is in the microdarcy to nanodarcy range. Production is accomplished through horizontal wells with artificial fracturing. The fracturing process may reopen large sealed faults, thus reducing the efficiency of the process. Faults may penetrate either the water rich Ellenburger dolomite below or the Marble Falls limestone above and allow water to be drawn into the well, jeopardizing production. The Ellenburger contains numerous paleo-karsts which have caused the Barnett to sag into the collapsed cavities. These weak zones could also cause easy communication between the Ellenburger dolomite and bore holes in the Barnett. So

an understanding of the historic faults, fractures, and paleo-karsts is important to well placement.

Volume rendering

A volume of 1000 inlines by 1000 crosslines and 1000 time or depth samples can be rendered opaquely as six faces each of a million points for a total of six million graphic operations. Once some transparency is introduced the number of operations required rises to a billion, which represents an increase of more than two orders of magnitude. This makes volume rendering very sensitive to the speed of the graphics system, including the bandwidth available for data transfers from central memory to the graphics board. The speed of rendering is immediately apparent to the interpreter and directly affects ease of use. Until now the speed of data transfer to the board has been the limiting factor for usefully-sized volumes because graphics memory sizes have been relatively small relative to the size of a typical seismic dataset. The speed of graphics rendering on the board is another factor affecting overall performance. Manufacturers have been increasing this speed by adding multiple arithmetic units to their graphics processing units (GPUs) in order to perform the coordinate transformations and lighting calculations necessary for drawing objects in three dimensions. As these arithmetic units have increased in speed and number, the rendered image has become ever more realistic. This realism is evidenced in 3D movies and 3D games that rely on this technology.

Graphics processing units (GPUs)

The latest graphics cards have 4 to 6 GB of graphics memory and can have more than 400 cores. The graphics memory is now sufficient to hold a seismic data volume of 2 GB and still leave plenty of room to provide output buffers for the results of computations as well as room to hold graphics data. Computations on the graphics cards are supported by two computer languages that have been designed to support parallel computation. These are the proprietary language C for CUDA supplied by Nvidia and

¹ Paradigm

* huw.james@pdgm.com

Modelling/Interpretation

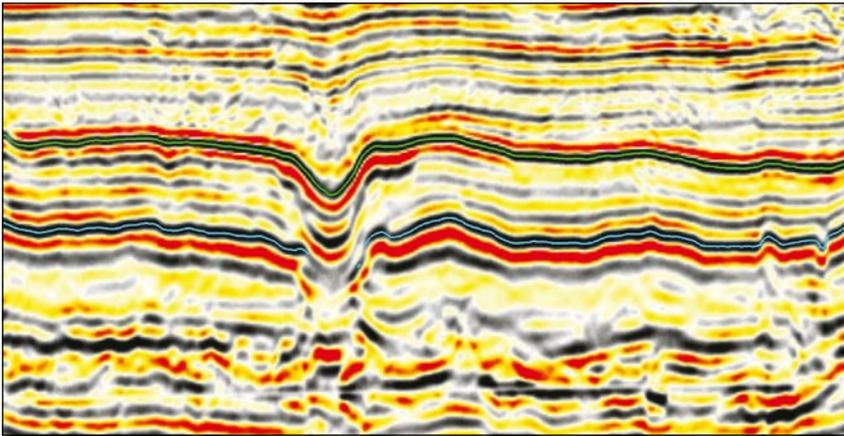


Figure 1 A typical paleo-karst with an induced slump in the Marble Falls formation shown in green.

the Open Standard OpenCL supported by Apple, IBM, Intel, Microsoft, and both Nvidia and AMD, as well as many other members of the consortium. These languages support a compute paradigm known as single instruction and multiple data (SIMD). This means that a single instruction can be executed on multiple arithmetic units and applied to different data streams. This provides a more efficient use of memory bandwidth for parallel computation than traditional CPUs independently working in parallel.

To make best use of the GPU it is usually necessary to rewrite algorithms in these languages and to pay equal attention to the memory usage of the algorithm as to their usage of computation resources. The cost of rewriting the central algorithms of processing applications has delayed broad adoption of GPUs for high performance computing, i.e., the seismic processing activity in our industry. Adoption has been limited, so far, to small compute kernels of cost-sensitive applications. Another historic cost of GPU computing for non-graphical applications is the need to transfer data from main memory to GPU memory and back over a relative slow interface of 5 GB/sec. For volume rendering we only need to send data one way from main memory to the GPU. Our options are to either send 3D graphics data or seismic data across this interface, so there is no extra cost to using the GPU. This has encouraged us to change from sending 3D graphics data to sending seismic data to the GPU. When we do this we have to perform volume rendering by computing the image in software on the GPU instead of relying on rendering the data via a graphics language.

By performing our own volume rendering we can choose the algorithms for volume rendering rather than depend on algorithms provided by graphics vendors. This is a benefit because the vendor algorithms are not tailored to band limited seismic data. Another advantage of this arrangement is that if we wish to render the volume in a different fashion or over a different extent, we no longer have to retransmit the graphics data across the interface, we can re-render the image from the data already on the

graphics card. Once we have taken this step it becomes attractive to perform other computations on the graphics card. Now that we have access to over 400 cores, fast memory bandwidth and parallel computing languages such as OpenCL or C for CUDA, moving these computations to the GPU becomes an attractive advantage rather than a re-programming cost.

The result of these changes is that processes such as proportional flattening, band pass filtering, gain, attribute calculations, or any other post-stack signal processing can be performed in an instant on useful-sized volumes. Once they are on a level playing field the 400 cores outperform the 12 CPU cores on a typical work station in a fashion that is very empowering.

Volume rendering itself is about eight times faster than rendering 3D graphics on earlier graphics cards. With other increases in speeds this means that useful sized 3D volumes can be interactively volume rendered, viewed, translated, rotated, and scaled. This improves usability and removes some of the skill needed to manipulate variable opacity volume displays. Volume rendering quality is also improved greatly and users can make trade-offs between speed and quality without needing to re-send data from main memory to the graphics cards.

Example from the Barnett Shale

The data set being viewed is relatively small at about 750 in-lines by 750 cross-lines and 500 depth samples, a total of less than 300 million samples. A typical paleo-karst is shown in Figure 1 with a green pick for the Marble Falls formation.

This paleo-karst is about 100 m in diameter and quite obvious from standard displays. Figure 2 shows the dip magnitude of the Marbles Falls surface draped on the structure. The blue colours indicate high values of the dip magnitude. The surface slumps caused by the dissolution of the karsts are plainly apparent. This display also picks out expressions of both major and minor faults for this

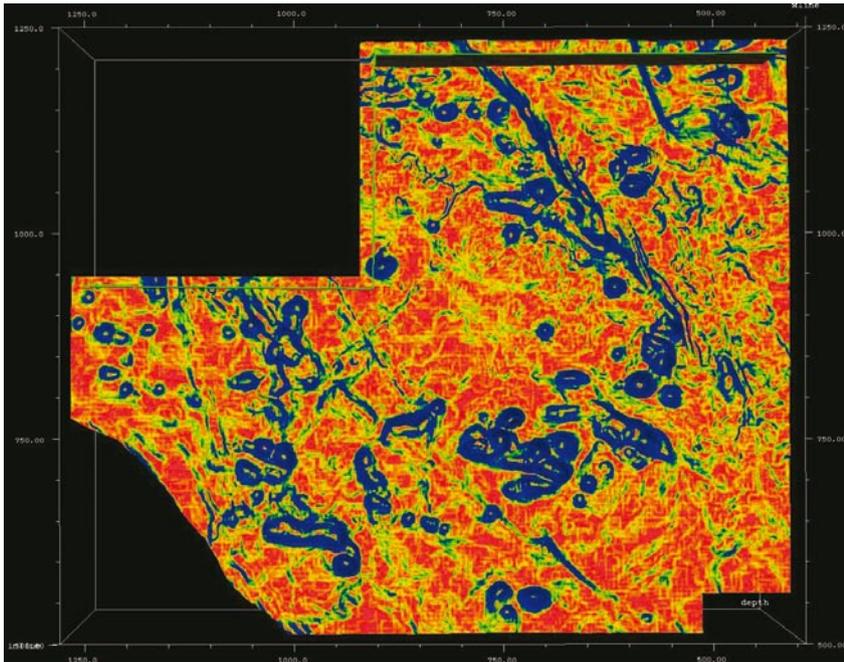


Figure 2 Plan view of the Marble Falls with the dip magnitude attribute draped on the surface. Blue indicates high dip and red low dip. The complexity of the faults can be seen as well as the indications of the paleo-karsts.

interval. The strike direction indicated is generally from the southeast to northwest but covers a range of azimuths with some indications of faults with strike in the orthogonal direction. It may be an illusion but many of the karsts seem to be joined up by faults indicating some relationship between the two.

To get some insight into the current topography of the Marble Falls we show an oblique 3D view of the Marble Falls surface pick with a vertical seismic section. In the example in Figure 3 we have chosen to drape the mean curvature attribute on the surface. This example conveys the rugosity of this surface and the relaying of the fault at the top right. The view in Figure 3 can be compared with the volume rendered view in Figure 4. The quality of this volume rendering is derived from using the GPU to perform

volume rendering. In this example, the whole amplitude volume has been loaded onto the GPU and we are free to operate on it with more than 400 cores to improve the rendering quality at will. Before, when the GPU's memory was too small to hold a volume of data, such improvements in rendering quality would need to be performed on the CPU. This typically resulted in larger volume sizes which in turn then took longer and longer to send across the interface to the graphics card. Most interpreters chose speed over quality in this trade-off. The significance of the improved technology is that this trade-off need not be taken once at volume creation time, it can be taken on-the-fly while viewing the volume. The outcome is dramatically better usability, insofar as most interpreters can now choose quality whenever they wish.

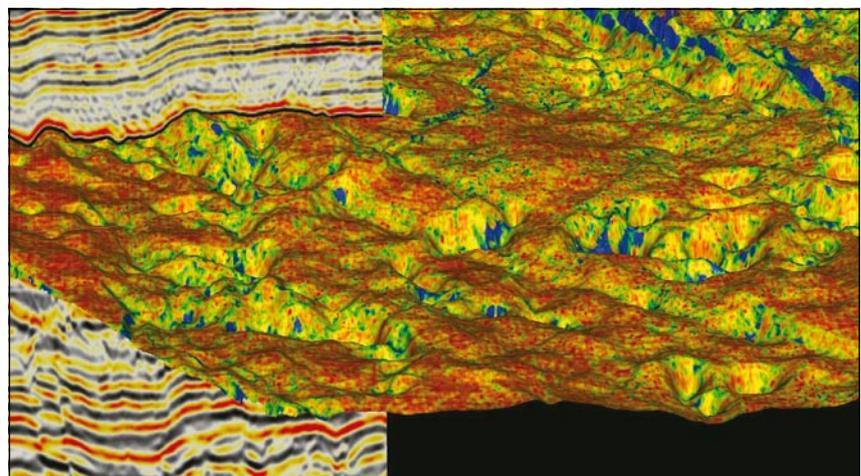


Figure 3 A zoom of the Marble Falls pick with a drape of mean curvature. Three paleo-karsts can be seen at the front lower right. There is also an indication of an acquisition footprint in the red texture.

Modelling/Interpretation

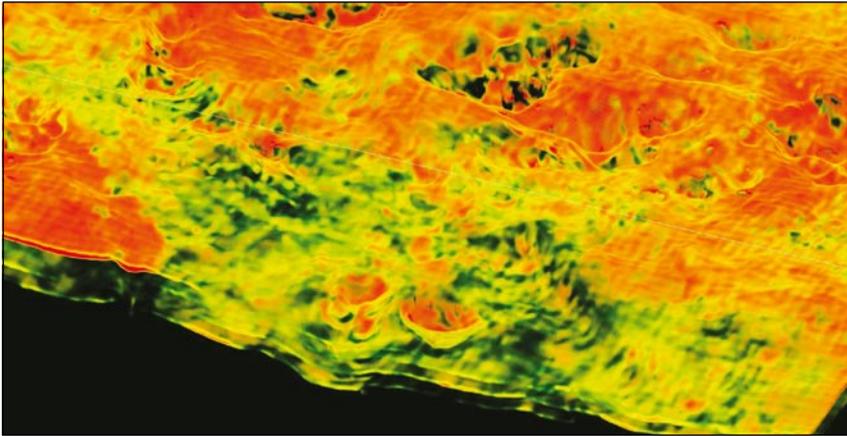


Figure 4 A volume rendered display of an interval of the Barnett Shale, the three small paleo-karsts in Figure 3 can be seen in the front centre as part of a volume display rather than as a surface.

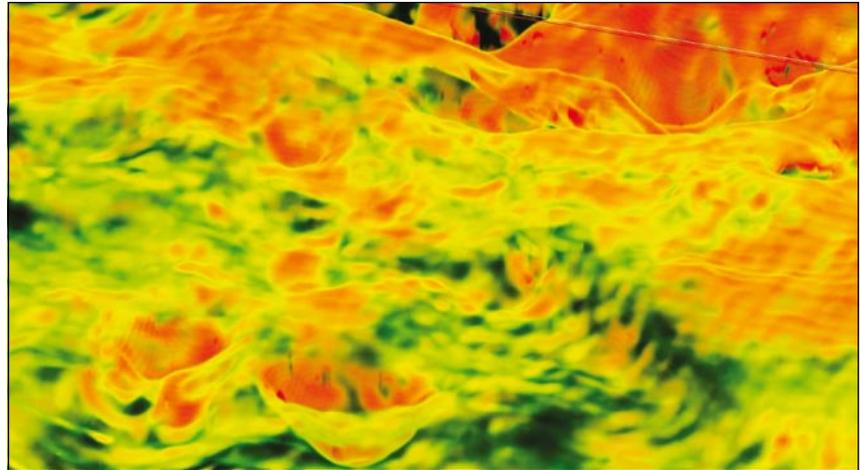


Figure 5 A zoom of Figure 4 showing the volume effect of a paleo-karst in detail as well as the improved quality of volume rendering.

Back to the data: hints of the acquisition footprint can also be seen in this display as lines of increased amplitude in both the in-line and cross-line direction. These can provide better assistance to the interpreter in deciding which lineaments are artefacts and which can be confidently interpreted as faults when the strike direction is close to the in-line or cross-line directions. Figure 4 illustrates both possibilities.

Figure 5 is a zoom of Figure 4. It shows that the volume rendering quality is preserved. No data has been re-transmitted from the CPU to the graphics card. We have simply rendered a smaller volume of the data on the GPU with more fidelity. There is no longer any need to compromise quality for speed. These volumes which are almost 300 MB can be displayed, rotated, and zoomed interactively even with the seismic volume rendered with variable opacity. The interpreter can feel fully in control of the display. Controlling the scene can be accomplished using far less control since errors in control have almost no impact on an interpreter's productivity.

As noted previously, once the volume of data is moved to the GPU we are free to use the GPU cores for data processing

as well as rendering. Complex trace attributes are often used by seismic interpreters. Calculating such attributes is usually a batch job or a batch job initiated interactively. It generates new volumes of seismic data. Now this same processing activity can be performed on the GPU and the results shown to the interpreter and then used without even needing to return the attribute to the CPU's main memory or onto the disk. The large number of cores can make such calculations instantaneous on volumes such as this. To discard these calculations after use sounds like wasteful behaviour but it's not. The GPU is present and the volume of amplitude data is resident so it's quicker to compute these attributes on the fly and to display them rather than to compute them and then save and fetch the results back and forth from the CPU's main memory. If the attributes are needed for some other purpose than display, then they can be returned to the CPU's main memory. The analogy could be that when driving very fast one does not retain a permanent memory of all views, just a notion of the driving path and significant events. If one wants to return to a view it's pretty quick to drive back to see it. This keeps memory from being overloaded with data of little future value.

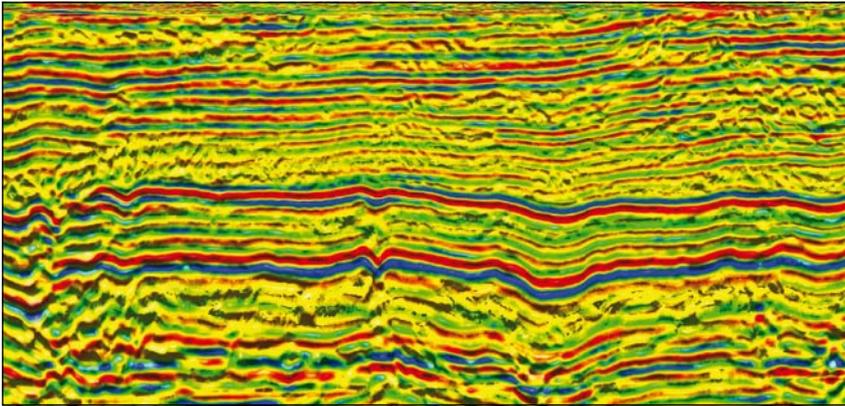


Figure 6 A volume rendered vertical slab of reflection amplitude data with mild opacity and lit equally by three lights from upper front, lower left, and lower right.

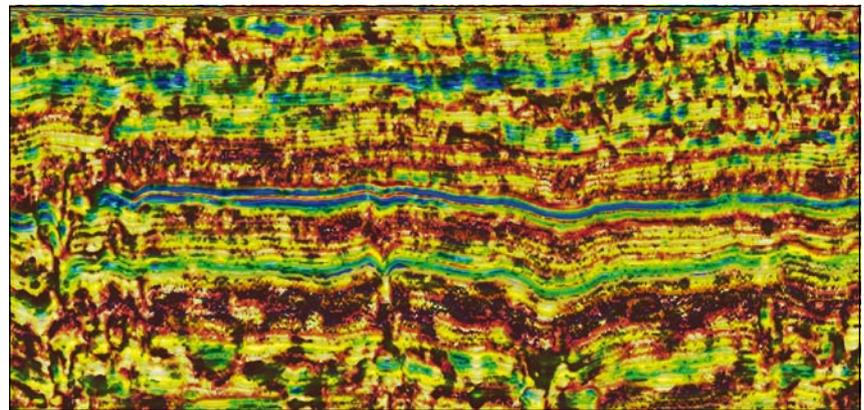


Figure 7 The same slab as in Figure 6 with mild opacity and the same lighting but with instantaneous frequency in place of reflection amplitude. Here the discontinuities in the volume are accentuated.

Figures 6 and 7 are snapshots of the same view of the volume. Opacity has been set to give a relatively shallow view into the volume of a few traces so these are not section displays. Lighting has been set to illuminate the face somewhat equally with three lights all from the front with positions above to the left, above to the right, and below from the centre. Then the display has been toggled from reflection amplitude to instantaneous frequency. This attribute is usually thought to be of value for stratigraphy, here we are using it for structure. The attribute calculation is instantaneous on this size of volume so it can be toggled from one attribute type to another via a keyboard hot key. This offers a different user experience than is typical for the use of attribute calculations.

The Marble Falls event can be seen midway down the apparent 'section'. In these displays the reflection amplitude provides continuity and an easy event to pick. The instantaneous frequency display makes the faults far easier to see, understand, and pick. It should be remembered that these are not simple section displays, the data being viewed has some depth into the screen, and the image and the lighting acts as a differentiator accenting edges. The variable opacity makes the edges more substantial because the edge has to exist several sections deep, then the lighting can illuminate the edge or not. Both displays are using a

rainbow colour scale of blue to red over the range of most of the data in each case.

Conclusions

This latest advance in graphics technology will change the emphasis from computing on the CPU and sending graphics data to the graphics card to an emphasis on sending data to the graphics card and using it for many computations. This change is not merely one of performance, the massive parallelism of the latest GPUs enables orders of magnitude increases in processing speed. This power can be put to great use in volume visualization to provide better quality renderings at faster speeds for greater and more ergonomic usability. The attribute display above is an example that demonstrates that there is more information to extract from today's seismic data if we can find the right tools and displays to do so.

Acknowledgements

We thank Paradigm for permission to publish this work and Marathon for the use of their seismic data.

References

Kidd, G. and Jenson, F. [2010] An integrated workflow for selecting optimal drilling locations. *Hart's E&P*, June.