

Dan Martin
July 2014
Revision 1.0

GeoCue's ongoing software development in support of optimized geospatial processing workflows includes research into Graphics Processing Unit (GPU) technology. GPUs, on both local and remote nodes, provide additional processing power for churning through mountains of imagery and LIDAR data. Stating the obvious, GPU technology has been around for a while. But what is less obvious is how GPU technology is best applied to geospatial workflows. It is tempting to think that simply throwing additional GPUs at a long running task will significantly speed up image and LIDAR processing workflows. And the bigger and more expensive the GPU card, the better!

Our research has shown that GPUs can indeed provide a beneficial impact on geospatial workflow processing times. In some cases we've seen 10x reductions in processing times when a processing task is coded to specifically leverage GPU architecture. Stated simply, GPU friendly algorithms allow the GPU to iterate a single operation across a large array of pixels or points. However we have seen in our application development that there is a fine balance between GPU memory use, GPU processing, Central Processing Unit (CPU) memory, CPU processing, and storage I/O operations. These elements must be optimized together to arrive at a meaningful optimization of geospatial workflows.

Examples of algorithms that lend themselves to GPU acceleration include:

- Most Image Processing Algorithms
 - All neighborhood filters such as Gaussian, Median, Bilateral, etc.
 - Histogram Stretching
 - Color Conversions
- Point Matching
- Data Compression/Decompression

Consider too that a GPU, whether a separate card or a component on the CPU motherboard, must also act as a display device and serve images to your screen in support of the computer's operating system. With multiple GPUs in a single machine, one GPU can serve as the display device while any other available GPUs can be treated as dedicated compute devices for image processing. Of course, none of this happens automatically and cards usually have to be configured via device drivers, command line tools, or bios settings to provide dedicated compute services.

Since GeoCue's production management solutions are used in both single machine (standalone) and distributed processing configurations, we are developing solutions that leverage available GPUs in a single machine or across the enterprise. But again, simply adding more GPUs, whether local or on remote nodes, doesn't always provide the decrease in processing times one might expect. With massive amounts of project data being the norm in geospatial workflows, an important aspect of algorithm design for GPU optimization lies in balancing the decision to move data to the processors or add more processors near the data, e.g. on the same machine.

We are seeing great results when enlisting GPUs at discrete steps in image processing. But as in all things in life, there is an optimal balance. GeoCue continues investing in GPU research and development,

Tools, Tips and Workflows

Developing for GPU



determining how best to apply GPU technology to geospatial data processing and building accelerated functionality into our production management solutions.