Introduction to this special section: High-performance computing

Paul Sava, Colorado School of Mines

It goes without saying that the geoscientist's appetite for more detailed information about the Earth can never be completely satisfied. As geoscientists we want to understand the Earth at higher and higher resolution and on a larger and larger scale. Over the past decades, geoscientists have developed methods for investigating the Earth using complicated theoretical models applicable to huge data sets. The many game-changing applications deployed in recent years include: 3D prestack depth migration for imaging the Earth's reflectivity, 3D wavefield tomography for deriving complex velocity models, sophisticated fluid-flow simulators for understanding oil and gas production, and digital rock simulators for understanding the subtle relationships between the measured fields and the underlying geology. All these applications are linked together by one attribute: They are all extremely compute-intensive.

For many years, the response of the high-performance computing (HPC) community to these challenges were to produce increasingly faster central processing units (CPUs). Over the past decades, CPU speed increased following a trend that is known today as Moore's law named after Intel co-founder Gordon E. Moore who in 1965 wrote:

> The complexity for minimum component costs has increased at a rate of roughly a factor of two per year ... Certainly over the short term this rate can be expected to continue, if not to increase. Over the longer term, the rate of increase is a bit more uncertain, although there is no reason to believe it will not remain nearly constant for at least 10 years.

In fact, this trend remained the norm for several more decades, thus dramatically boosting the compute power provided by commodity computers. In discussing these trends, we often see graphs indicating huge growth rates of, for example, the numbers of transistors on a silicon chip. Unfortunately, such graphs are not at all intuitive to the general public, which does not appreciate how much the computing environment has changed around us. So, here is an analogy which might help some of us. The table below compares the compute power of the Apollo guidance computer with a typical cell phone and a common desktop computer in January 2010.

	Clock speed	Memory	Storage
Apollo lander	2 MHz	0.004 MB	?
Cell phone	600 MHz	256 MB	16 GB
Desktop computer	2500 MHz	>8000 MB	>1000 GB

Would you want to fly to the Moon with the aid of a computer several orders of magnitude less powerful than your cell phone? Well, this has been done, which underscores even more the astonishing achievement of the Apollo crews.

While high emphasis has been placed in the past on CPU speed, this is no longer consistent with recent trends. CPU speed is not increasing at the same rate as in the past, simply because of the physical limits of packing more and more transistors on the surface of a silicon chip. The current trend is to distribute the processing power among many smaller processing units. Today we see many compute "cores" being packed on the same processing unit; therefore, we talk about CPUs with 4 or 8 cores, and certainly there will be even more in the near future. A computer equipped with a multicore CPU will soon be equivalent with a small cluster of no more than 10 years ago. Even more intriguing, geoscientists are beginning to use alternative processing devices like the graphics processing units (GPUs) that drive the graphic cards available in all computers around us, or field programmable gate arrays (FPGAs), which are integrated circuits designed to perform specific tasks efficiently. These new types of processing units deliver huge compute power, sometimes several orders of magnitude larger than that of a conventional CPU, but at the expense of specialization which makes code development more challenging.

An even more significant HPC trend with impact on geoscience applications is the move toward distributed computing on semi-independent processing nodes. Nowadays, supercomputers consisting of tens of thousands of interconnected nodes are common, and they provide the backbone of geophysical data processing, in particular for seismic-related applications like 3D migration. The real HPC challenges today are to link all the compute units together (efficient networking), to transfer large data volumes to and from the processing units (efficient data access), and to process data in small chunks distributed over many processing units (efficient algorithms). Many of those large supercomputers occupy entire buildings, consume an astonishing amount of electric power, and require large cooling units to maintain the processing units at an optimal functioning temperature.

This special section features several papers describing HPC trends and geoscience applications, although this topic is too vast to be covered in just a few journal pages. The articles mainly emphasize the relationships between new computer hardware and their application to imaging, interpretation, and simulation. Camp and Thierry present an overview of the long-term trends in CPU development and how the increasing compute power affects geoscience applications and their feasibility over time. Clapp et al. discuss how alternative processing units like GPUs can be used for reverse time seismic imaging, with emphasis on the relations between algorithms and the hardware necessary to maximize computing efficiency. Kadlec and Dorn discuss the application of GPUs to interactive real-time seismic interpretation and suggest ways in which the huge compute power available on desktop computers will impact interpretation software in the future. Tölke et al. show how GPUs can be applied to fluid-flow simulation in digital rocks, thus providing accurate relations between rock geometry and permeability. Käser et al. discuss the application of HPC to seismic wave simulation using 3D unstructured meshes which can be used for models with complex geometries, like the environment around boreholes. Noble et al. discuss another application of HPC to the problem of velocity model building using traveltime tomography applied to a large data set. Finally, Suh et al. present an application of cluster programming to reverse-time migration and discuss practical implementation issues related to data input/output and to network data transfer.

The race between demanding geoscience applications and the HPC infrastructure will continue into the future. At times, the applications will overpower the computers prompting an increased effort in developing both better hardware and software, together with more efficient applications. It is safe to say that computational geoscience will continue to be an exciting and rewarding field of scientific investigation and to anticipate that it will remain one of the main attractors of new talent to our field. TLE