EAGE

HPC01

Scale Out vs. Scale Up for Ultra-Scale Reservoir Simulation

K. Mukundakrishnan* (Stone Ridge Technology), R. Gandham (Stone), K.P. Esler (Stone Ridge Technology), D. Dembeck (Stone Ridge Technology), J. Shumway (Stone Ridge Technology), V. Natoli (Stone Ridge Technology)

Summary

It is an undisputed truth that demand for computational performance for simulating very large models in upstream applications is ever increasing. This demand can be met conceptually in one of two ways. “Scale-out”, implies exploiting additional computational nodes, while “scale-up” implies increasing the computational power, particularly floating point throughput and memory bandwidth of each node. In practice, these two approaches provide opposite bounds on a spectrum of cluster designs, from the use of many relatively weak, “thin” nodes, to a smaller number of powerful, “fat” nodes. The scale-out approach gained increasing dominance in HPC as scalability was prefered over absolute efficiency.

Over the past decade, however, energy efficiency has become the key performance limiter. For applications with significant communication requirements, including reservoir simulation, the use of scale-up fat nodes provides an opportunity to localize communications and minimize interconnect traffic, thereby increasing energy efficiency. However, harnessing fat fat nodes comprising of several extremely high-performance GPUs to achieve performance for implicit simulations requires careful software design and novel algorithmic approaches.

We will first present the algorithmic and computational challenges faced and the approaches needed to efficiently utilize the massive parallelism offered by such scaled-up nodes.

EAGE

Abstract

It is an undisputed truth that demand for computational performance for simulating very large models in upstream applications is ever increasing. This demand can be met conceptually in one of two ways. "Scale-out", implies exploiting additional computational nodes, while "scale-up" implies increasing the computational power, particularly floating point throughput and memory bandwidth of each node. In practice, these two approaches provide opposite bounds on a spectrum of cluster designs, from the use of many relatively weak, "thin" nodes, to a smaller number of powerful, "fat" nodes. The scale-out approach gained increasing dominance in HPC during the 90s and 2000's due the the availability of cheap commodity nodes. HPC applications developed during this period often focused on scalability over absolute efficiency.

Over the past decade, however, energy efficiency has become the key performance limiter. Increased transistor density has dramatically changed the relative energy cost of moving data versus performing operations (flops) on it. For applications with significant communication requirements, including reservoir simulation, the use of scale-up fat nodes provides an opportunity to localize communications and minimize interconnect traffic, thereby increasing energy efficiency. Fortunately, industry trends outside traditional HPC, such as AI and data analytics, are providing a commodity market for scaled-up fat nodes comprising of several extremely high-performance GPUs. Such nodes provide an ideal computation platform for simulating large-scale reservoir models with a fast turnaround with relatively few nodes. However, harnessing these fat nodes to achieve performance for implicit simulations requires careful software design and novel algorithmic approaches.

Results

We will first present the algorithmic and computational challenges faced and the approaches needed to efficiently utilize the massive parallelism offered by such scaled-up nodes. To highlight these, results from the simulation of an ultra-scale reservoir model using a fully GPU-accelerated reservoir simulator and only a very few scaled-up fat nodes will be presented.

Major Conclusions

Using our scale-up approach coupled with current generation GPUs, assisted history matching and uncertainty quantification of workflows even for models with many hundreds of millions of cells is possible with reasonable resources.