HPC14

Using a GPU Cluster for Multiple Realization Workflows

T. Miller* (Ridgeway Kite Software), G. Bowen (Ridgeway Kite Software) & B. Zineddin (Ridgeway Kite Software)

SUMMARY

There are two long standing trends in reservoir simulation, firstly to increase resolution of models and secondly to run multiple realization workflows to better characterize the inherent uncertainty associated with the sub-surface. These need to be balanced when faced with defined computing resources.

Moore's law has continued to revolutionize the computing environment. Recently it has become a requirement that parallelism be exposed at all levels to exploit computer power. Therefore the GPU has become an important tool for HPC applications.

This paper demonstrates that models (~1M cells, fully implicit, black oil) can be run effectively on a single GPU taking advantage of its parallel nature and more importantly the high bandwidth to memory relative to a CPU. The second parallelism utilizes a cluster of GPUs to handle multiple realization (MR) workflows, where the MR capability is implemented within the simulator itself and is optimized to take full advantage of the hardware.

The "sweet spot" of being able to efficiently run response surface assisted history matching runs and Monte Carlo type analyses as a single batch process holds out the prospect of changing how simulation is used on a day-to-day basis, making it practical to move away from a single deterministic model.

Introduction

There are two long standing trends in reservoir simulation, firstly to increase resolution of models and secondly to run multiple realization workflows to better characterize the inherent uncertainty associated with the sub-surface. Clearly these need to be balanced when faced with defined computing resources at a given point in time.

Moore's law has continued to revolutionize the computing environment, in recent years it has become a requirement that parallelism be exposed at all levels to exploit the development in computing. In particular the GPU has become an important tool for high performance applications.

This paper demonstrates that useful resolution models (~one million cells, fully implicit, black oil) can be run effectively on a single GPU taking advantage of the parallel nature of the GPU and maybe more importantly the high band-width to memory relative to a comparable CPU. The second parallelism demonstrated is to utilize a cluster of GPU to handle multiple realization (MR) workflows, where the MR capability is implemented within the simulator itself and is optimized to take full advantage of the hardware.

The "sweet spot" of being able to efficiently run, for example, response surface assisted history matching runs and Monte Carlo type analyses as a single batch process holds out the prospect of changing how simulation is used on a day-to-day basis, making it practical to move away from a single deterministic model.