Geophysical Research Abstracts
Vol. 13, EGU2011-6798, 2011
EGU General Assembly 2011

High-performance unstructured 3D modeling on modern clusters with thousands of CPUs.

Marcin Krotkiewski and Marcin Dabrowski
University of Oslo, Physics of Geological Processes, Oslo, Norway (marcink@fys.uio.no)

Recent years brought a major change in the area of computer architectures. Theoretically, the computational power of modern supercomputers is nowadays large enough to perform numerical simulations on three-dimensional unstructured meshes with billions of mesh nodes, which is required in many applications related to Earth sciences. However, the computer architectures have undergone a significant change. Two fundamental characteristics of modern HPC platforms are the growing number of computational cores and the ever more limited memory bandwidth available per-core. The inherent parallelism calls for the use of parallel algorithms, while the memory bottleneck necessitates a development of new, memory-oriented code optimizations. As a result, many of the existing numerical codes need to be re-written in order to provide adequate performance to handle high-resolution models.

In this work we present an optimized parallel implementation of symmetric sparse matrix - vector product (SpMV) for modern massively parallel architectures [1]. SpMV is a basic building block of iterative solvers of systems of linear equations and one of the most time consuming components of modern numerical simulations software. It is well known that SpMV is a memory bounded algorithm: most implementations exhibit low fractions of FLOPS since it is the memory speed that is the main bottleneck. Thus, to improve the SpMV performance we mainly deal with optimizing the memory access and limiting memory usage. Also, in order to demonstrate the optimality of the implementation we analyze the effective memory bandwidth utilized by the code.

Efficient and parallel SpMV implementation for matrices resulting from unstructured 3D discretizations is especially challenging. The irregular matrix structure and the indirect memory addressing require careful programming in order to achieve high performance on modern cache-based CPUs. Here, matrices with highly variable structure and density arising from unstructured three dimensional FEM discretizations of elasticity and diffusion problems are considered. We demonstrate that for the studied matrices the presented implementation is optimal the Cray XT4, i.e. it utilizes close to 100% of the memory bandwidth available on a single node of the cluster. Moreover, we show that storing only the upper-triangular part of symmetric matrices results in roughly two times speedup because of the reduced memory requirements. Parallel efficiency on 5400 Opteron cores of the Cray XT4 cluster is around 80-90% for problems with approximately 25ˆ3 mesh nodes per core. For a problem with 820 million degrees of freedom the code runs with a sustained performance of 5.2 TeraFLOPs.

1. Parallel symmetric sparse matrix-vector product on scalar multi-core CPUs, M. Krotkiewski, M. Dabrowski, Parallel Computing, Volume 36, Issue 4, April 2010, Pages 181-198