An Energy-Efficient Parallel Multigrid Method

Martin Wlotzka, Teresa Beck
Interdisciplinary Center for Scientific Computing (IWR), Heidelberg University
martin.wlotzka@uni-heidelberg.de, teresa.beck@uni-heidelberg.de

Vincent Heuveline
Computing Center, Heidelberg University
vincent.heuveline@uni-heidelberg.de

Abstract

Multigrid solvers belong to the most efficient state of the art numerical methods for solving symmetric positive definite linear systems. The computational complexity is O(n) for sparse systems with n unknowns. This optimal complexity makes multigrid a method of choice for many large-scale applications.

The advances in high-performance computing (HPC) towards exascale machines rise the demand for scalable and energy-efficient techniques. Standard parallelization uses all available computing resources throughout the whole multigrid hierarchy. The overall performance is then limited by the parallel performance on the coarsest level with the smallest problem size. The use of multi-core and many-core platforms and co-processors to improve the parallel performance has been studied in many works, see e.g. [1, 2, 3]. Recently, also the issue of energy consumption has been addressed by means of accelerated smoothers and grid transfer operators [4].

We improve both the parallel performance and the energy efficiency of a multigrid solver by adapting the number of active processors according to the different problem sizes in the multigrid hierarchy. To this end, we introduce a strategy for the dynamic adjustment of the hardware activity during the execution of the multigrid solver. For comparison with the standard parallelization, we assess performance and energy consumption by means of high precision time and power measurements. Our adaption strategy can yield substantial savings both in time and energy to solution in situations where the non-adapted parallelization becomes inefficient. Moreover, we give evidence that the energy savings are not only a consequence of the reduced time to solution, but rather an inherent benefit of our method.

References

1. E. Chow and R.D. Falgout and J.J. Hu and R.S. Tuminaro and U.M. Yang. Parallel Processing for Scientific Computing: A Survey of Parallelization Techniques for Multigrid Solvers. SIAM Series on Software, Environments, and Tools. 2006.
2. H. Anzt and J. Dongarra and M. Gates and S. Tomov. Block-asynchronous multigrid smoothers for GPU-accelerated systems. EMCL Preprint Series No. 15-2011.
3. N. Bell and S. Dalton and L.N. Olson. Exposing fine-grained parallelism in algebraic multigrid methods. SIAM J. Sci. Comput., 34:123-152, 2012.
4. M. Wlotzka and V. Heuveline. GPU-accelerated smoothers and grid transfer operators for parallel multigrid methods on HPC clusters. Submitted to J. Parallel Distrib. Comput., 2015.