On Level Scheduling for Incomplete LU Factorization Preconditioners on Accelerators

Karl Rupp
Argonne National Laboratory
rupp@mcs.anl.gov

Barry Smith
Argonne National Laboratory
bsmith@mcs.anl.gov

Abstract

The application of the finite element method for the numerical solution of partial differential equations naturally leads tolarge systems of linear equations represented by a sparse system matrix A and right hand side b. These systems are commonly solved using iterative solvers, particularly Krylov subspace methods, which are typically accelerated using preconditioners to obtain good convergence rates [1]. One of the most popular families of preconditioners are incomplete LU factorization (ILU) preconditioners, where the system matrix A is factored approximately into a sparse lower-triangular matrix L and a sparse upper-triangular matrix U. Then, each application of the preconditioner to a residual vector z involves one forward-substitution $Ly = z$ and one backward substitution $Ux = y$. A drawback of ILU preconditioners is the limited amount of parallelism both in the factorization and in the triangular substitutions. This complicates the efficient implementation on parallel computing architectures such as graphics processing units (GPUs) considerably. Our work is based on previous work by Li and Saad using level scheduling for ILU preconditioners on CUDA-enabled GPUs [2]. We refine their approach by considering modifications of reordering algorithms such as Cuthill-McKee or Gibbs-Poole-Stockmeyer for a higher degree of parallelism and thus higher computational efficiency of ILU preconditioners. Furthermore we apply these techniques to block-ILU preconditioners and compare with the Power(q)-method by Heuveline, Lukarski, and Weiss [3]. Results obtained on GPUs from AMD and NVIDIA as well as on INTEL's many-integrated-core (MIC) architecture will be presented. Our implementations are freely available in the open-source library ViennaCL [4], which is currently integrated into the distributed solver package PETSc [5].

References

1. Y. Saad. Iterative Methods for Sparse Linear Systems. 2nd Edition, SIAM (2003).
2. R. Li and Y. Saad. GPU-Accelerated Preconditioned Iterative Linear Solvers. Technical report, University of Minnesota (2010).
3. V. Heuveline and D. Lukarski and J.-P. Weiss. Enhanced Parallel ILU(p)-based Preconditioners for Multi-core CPUs and GPUs - The Power(q)-pattern Method. EMCL Preprint 2011-08 (2011).
4. K. Rupp and F. Rudolf and J. Weinbub. ViennaCL Web Page. http://viennacl.sourceforge.net/.
5. S. Balay and J. Brown and K. Buschelman and W. D. Gropp and D. Kaushik and M. G. Knepley and L. C. McInnes and B. F. Smith and H. Zhang. PETSc Web page. http://www.mcs.anl.gov/petsc/.