

The Jacobi-Davidson Eigensolver on GPU Clusters

Jonas Thies

German Aerospace Center (DLR), Simulation and Software Technology

`Jonas.Thies@DLR.de`

Dominik Ernst

Erlangen Regional Computing Center (RRZE)

`dominik.ernst@fau.de`

Abstract

Compared to multi-core processors, GPUs typically offer a higher memory bandwidth, which makes them attractive for memory-bounded codes like sparse linear and eigenvalue solvers. The fundamental performance issue we encounter when implementing such methods for modern GPUs is that the ratio between memory bandwidth and memory capacity is significantly higher than for CPUs. When solving large-scale problems one therefore has to use more compute nodes and is quickly forced into the strong scaling limit. In this paper we consider an advanced eigensolver (the block Jacobi-Davidson QR method [1]), implemented in the PHIST software (<https://bitbucket.org/essex/phist/>). We aim to provide a blueprint and a framework for implementing other iterative solvers like Krylov subspace methods for modern architectures that have relatively small high-bandwidth memory. The techniques we explore to reduce the memory footprint of our solver include mixed precision arithmetic and recalculating quantities ‘on-the-fly’. We use performance models to back our results theoretically and ensure performance portability.

References

1. M. RÖHRIG-ZÖLLNER AND J. THIES AND M. KREUTZER AND A. ALVERMANN AND A. PIEPER AND A. BASERMANN AND G. HAGER AND G. WELLEIN AND H. FEHSKE. Increasing the Performance of the Jacobi-Davidson Method by Blocking. *SIAM J. Sci. Comput.* 37(6), 2015.