SPECTRAL TRANSFORMS AND EIGENSOLVERS FOR EXASCALE COMPUTING

Project Description

Eigenvalue problems appear in numerous scientific and engineering fields; for example, calculating electronic structure of molecules or analyzing multivariate problems found in “Big Data.” Solving large eigenvalue problems is computationally intensive, and is often the chief computational bottleneck of these analyses. New computing architectures present challenges for traditional eigenvalue solvers; in particular, communication becomes increasingly expensive. Our goal is to develop methods of avoiding communication to effect more scalable eigenvalue solver software.

Relevance of Work to DOD

Eigenvalue solvers are crucial for material design, data analysis, quantum network design, and many other areas in DOD research. Scalable algorithms are crucial for improved performance on our ever increasingly capable HPC platforms.

Computational Approach

We focus on Lanczos eigensolvers and develop means to reduce the number of collective reduction operations. Collective reductions have been identified as a principle bottleneck, as they not only require communication but also require synchronization. We notice that non-synchronizing communication, such as point-to-point non-blocking communication, is substantially more scalable than blocking collectives. We therefore avoid blocking collectives by applying spectral transforms that trade blocking collectives for non-blocking point-to-point communication. We also notice that blocking collectives are called in asymptotically quadratic proportion to the size of the search subspace that is used, but larger search subspaces result in faster convergence. We solve

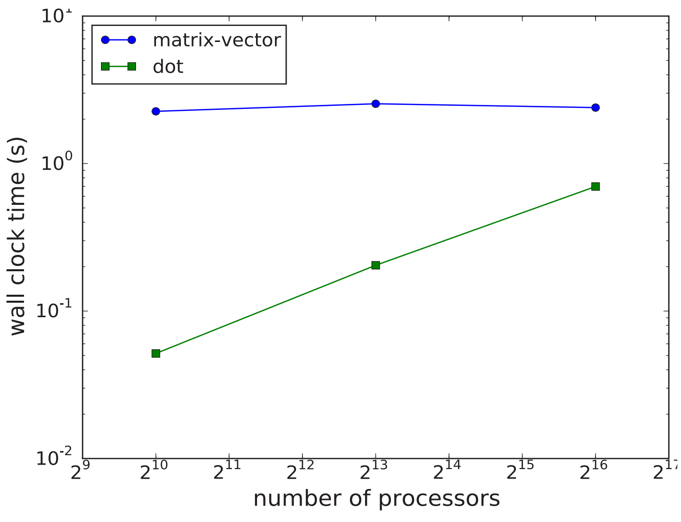

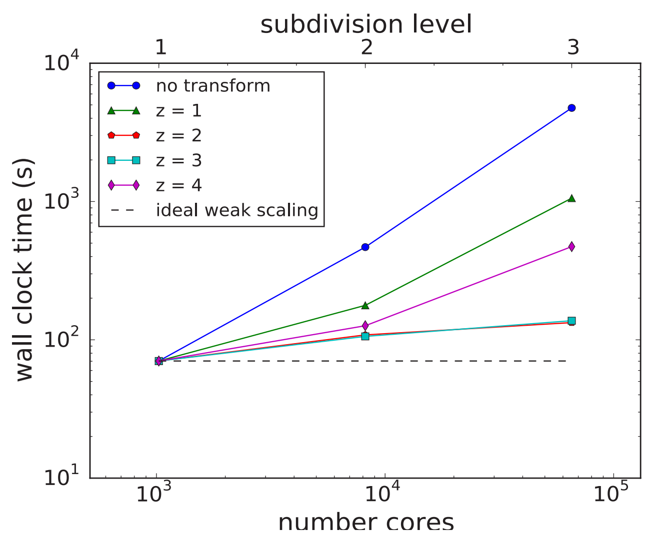

Figure 1: (left) Wall clock times for dot products versus matrix-vector products for a weak scaling experiment.
Figure 2: (right) Applying spectral transforms with degree proportional to dot product scaling flattens weak scaling times. Dot product cost is defined as r := tp/t0 where p is the number of processors and t0 is the time for the zeroth subdivision (in this case,512 processors). The parameter z is a scalar applied to r, such that the transform degree is d := rz.

the slow convergence problem by identifying restart stagnation as the culprit for slow convergence with small search subspaces, and develop stagnation-breaking heuristics. These allow for smaller search spaces to have convergence in nearly as few iterations as larger search subspaces, with a commensurate reduction in blocking collective calls.

Results

We have shown that the asymptotic reduction in outer Lanczos iterations to solve an eigenvalue problem is proportional to spectral transform degree (Figures 1, 2 and 3). This allows development of a degree-picking heuristic that can overcome non-optimal weak scaling. We have demonstrated a 3x speedup on 65k processor cores with this degree-picking heuristic. We have also shown a 40% reduction in blocking collectives with stagnation breaking over standard thick-restart Lanczos

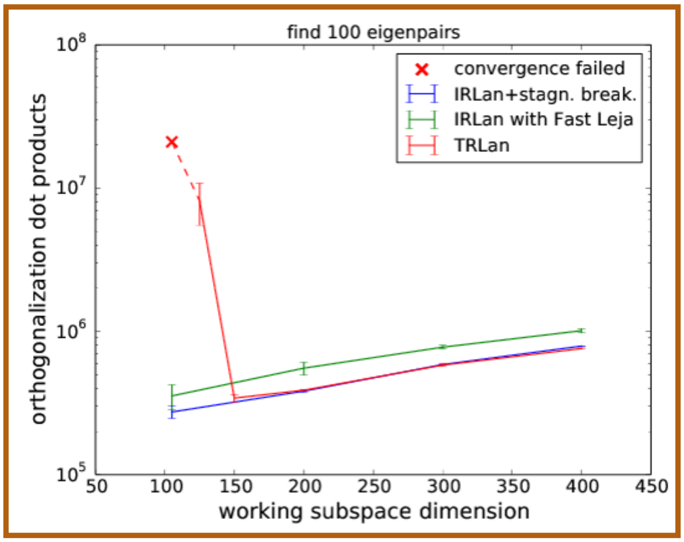

Figure 3: (above)
Stagnation-breaking reduces the number of Lanczos iterations needed to find eigenvalues in a small search space, and allows exact shifts to outperform comparable restart strategies

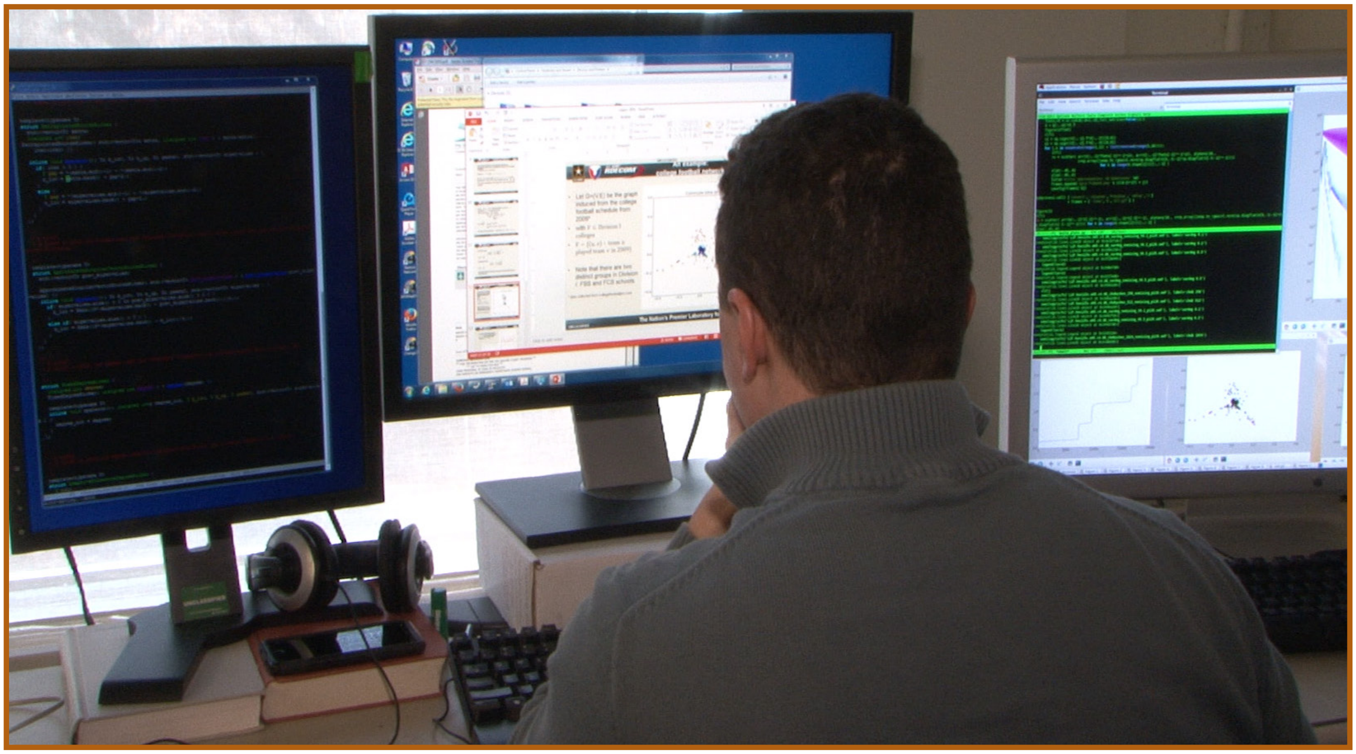

Contact Information

Dr. Alex Breuer
Army Research Laboratory
Computational and Information Sciences Directorate

Co-Investigators