
Eloy Romero Alcalde
**Spectrum Slicing in the Context of Computing Many
Eigenvalues in Large Hermitian Matrices**

Computer Science Department
College of William & Mary
McGlothlin-Street Hall 108
251 Jamestown Rd
Williamsburg
VA 23185
`eloy@cs.wm.edu`
Andreas Stathopoulos

We address the challenge of computing many eigenvalues and eigenvectors of sparse Hermitian matrices so large that neither direct methods nor unrestarted iterative methods such as Lanczos are feasible in terms of memory space. Typically, problems with large dimension are solved with restarted methods such as Thick-Restart Lanczos, Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG), Generalized Davidson with CG restarting (GD+k) and Jacobi-Davidson, which are competitive under limited memory.

However the performance of the restarted methods does not scale linearly with respect to the number of eigenvalues sought because of the computational cost of keeping the new solutions orthogonal to the already converged eigenvectors. As a result, orthogonalization introduces a quadratic term with respect to the number of computed eigenpairs.

We consider spectrum slicing which attempts to improve scalability by splitting the target spectrum region into subintervals and computing all the eigenvalues in each of them independently. The approach avoids orthogonalization against the solutions computed in other subintervals and at the same time introduces a new level of parallelism. However a practical implementation requires 1) a strategy to split the target region that balances the computation cost in each subinterval; for instance splitting the region into subintervals with similar number of eigenvalues; and 2) heuristics to dynamically tune the solver's parameters based on the density of the eigenvalues in the subinterval.

In the last few years there has been renewed interest in this technique encouraged by the emergence of high-performance and scalable solvers for computing all the eigenvalues inside a region in the spectrum. These solvers are based on Contour Integration (e.g., FEAST) and polynomial filtering (e.g., FILTLAN).

In this talk, first we focus on the choice of solver for within individual intervals. We compare FEAST using iterative methods for the solution of the linear

systems and Davidson-type methods available in PRIMME using polynomial filtering. We experiment with various polynomial degrees and various levels of accuracy for the solvers. Our results show that very small polynomial degrees and low linear solver accuracies are better for best performance. They also show that JDMQR or GD+k using polynomial filters outperform FEAST with iterative linear solvers. Second, we show our strategies to automatically tune the solver and choose the subintervals for spectrum slicing. These aim to find the best performance by fine tuning the degree of the polynomial to balance the total number of matrix-vector products versus the cost of orthogonalization.