Final Scientific/Technical Report
to the US Department of Energy

"Solving large-scale sparse eigenvalue problems and linear systems of equations for accelerator modeling"

Principal Investigator: Gene Golub
Fletcher Jones Professor of Computer Science
Stanford University, Stanford, California 94305

Collaborator: Kwok Ko
SLAC National Accelerator Laboratory
Menlo Park, CA 94025

Table of Contents

Executive Summary

The solutions of sparse eigenvalue problems and linear systems constitute one of the key computational kernels in the discretization of partial differential equations for the modeling of linear accelerators. The computational challenges faced by existing techniques for solving those sparse eigenvalue problems and linear systems call for continuing research to improve on the algorithms so that ever increasing problem size as required by the physics application can be tackled. Under the support of this award, the filter algorithm for solving large sparse eigenvalue problems was developed at Stanford to address the computational difficulties in the previous methods with the goal to enable accelerator simulations on then the world largest unclassified supercomputer at NERSC for this class of problems. Specifically, a new method, the Hemitian skew-Hemitian splitting method, was proposed and researched as an improved method for solving linear systems with non-Hermitian positive definite and semidefinite matrices.

1. Research Objective

One of the most important aspects in accelerator simulations is the accurate calculation of some low frequency electromagnetic fields within accelerator cavities, i.e., accurate eigensolutions of the frequency domain Maxwell's equations (FDME). Discretizations of FDME via edge-based finite element methods result in generalized matrix eigenproblems, $K x = \lambda M x$. In a typical accelerator simulation, the size of matrices K and M is usually large and the eigenvalues of interest are tightly clustered small interior eigenvalues out of a large-eigenvalue dominated spectrum. The Shift-Invert Lanczos method (SIL) is theoretically ideal for the above-mentioned eigenproblems. However, it requires accurate solutions of the shifted linear systems $(K - \sigma M) p = q$, which are large and ill-conditioned in the accelerator modeling and simulation. The research objective of this awarded proposal is to tackle the computational difficulties in solving eigenvalue problems and linear systems arising in large-scale accelerator simulations.

2. Methods and Results

There are two major research activities for the entire period of the project. One is about devising highly parallel algorithms for solving the above-mentioned interior eigenvalue problems with similar convergence properties to those of Shift-Invert Lanczos method but without demanding accurate solutions of any linear systems. The other is to develop novel algorithms for linear systems arising in the discretization of partial differential equations in accelerator simulations.

Two methods have been developed resulting from those research activities. The first method is the filter algorithm [1] for the sparse eigenvalue problem, which consists of two algorithms: an Inexact Shift-Invert Lanczos method for efficiently obtaining good eigenvector approximations and an inexact Newton-type method

to refine the eigenvector approximations to the desired accuracy. The second method is the Hermitian and Skew-Hermitian Splliting method [2-8] for solutions of linear systems.

2.1 The Filter Algorithm

Maxwell's equations in differential form have three independent equations.

$$\nabla \times \overrightarrow{\mathbf{E}} = -\frac{\partial \overrightarrow{\mathbf{B}}}{\partial t}$$

$$\nabla \times \overrightarrow{\mathbf{H}} = \frac{\partial \overrightarrow{\mathbf{D}}}{\partial t} + \overrightarrow{\mathbf{J}}$$

$$\nabla \cdot \overrightarrow{\mathbf{D}} = \rho$$

For a simple medium, the constitutive relations between field quantities are as follows.

$$\overrightarrow{\mathbf{D}} = \varepsilon_0 \varepsilon \overrightarrow{\mathbf{E}}$$

$$\overrightarrow{\mathbf{B}} = \mu_0 \mu \overrightarrow{\mathbf{H}}$$

where ε and μ are the relative electric permittivity and magnetic permeability while ε_0 and μ_0 are the values in the vacuum. In analyzing eigenmodes of electromagnetic cavities, field quantities in Maxwell's equations can be written in the form of harmonically oscillating functions with a single frequency ω. By making only one field quantity (for example E), a vector wave equation is obtained.

$$\nabla \times \left(\frac{1}{\mu} \nabla \times \overrightarrow{\mathbf{E}} \right) - \varepsilon k^2 \overrightarrow{\mathbf{E}} = 0$$

where k is the angular wavenumber ω/c and c the speed of light. A set of hierarchical high-order Nedelec basis functions which provide tangential continuity are used to discretize the electric field. The discretized system is an generalized eigenvalue problem.

$$\mathbf{K}\mathbf{x} = k^2 \mathbf{M}\mathbf{x}.$$

where the matrices K and M are

$$\mathbf{K}_{ij} = \left(\frac{1}{\mu} \nabla \times \overrightarrow{\mathbf{N}}_i, \nabla \times \overrightarrow{\mathbf{N}}_j \right), \quad and$$

$$\mathbf{M}_{ij} = \left(\varepsilon \overrightarrow{\mathbf{N}}_i, \overrightarrow{\mathbf{N}}_j \right).$$

There are several computational diffficulties associated with the above eigenvalue problems.
 1. Large matrix size: up to tens of millions.
 2. Big matrix bandwidth: typical for 3D finite element problems.
 3. Small eigenvalues out of a large-eigenvalue dominated spectrum: a typical eigenspectrum of the above eigenvalue problems has many eigenvalues with much bigger magnitudes than the eigenvalues of interest. Consequently,

efficient spectral transformations, which often involve solving large, indefinite, and ill-conditioned linear systems, are required.

4. Tightly clustered eigenvalues: the eigenvalues of interest are tightly clustered, due to the characteristic geometry of linear accelerator structures. A few hundred such eigenvalues and their eigenvectors need to be calculated.

5. Accurate solutions: highly accurate solutions are required for accelerator applications.

The combination of the above difficulties poses a significant challenge to many eigenvalue algorithms. The Shift-Invert Lanczos method (SIL) is theoretically ideal for finding tightly clustered interior eigenvalues and their eigenvectors, thanks to the effective spectral transformation achieved by solving the shifted linear systems, $(K-\sigma M)x = b$. However, SIL requires accurate solutions of the shifted linear systems, which are large and ill-conditioned, and hence difficult to be solved accurately by iterative linear solvers. That makes SIL less appealing practically.

For SIL, $(K-\sigma M)^{-1}$ is a effective band pass filter that emphasizes eigencomponents whose frequencies are close to σ. This is illustrated in the figure 1 that shows the spectra of a model matrix problem where M is the identity matrix.
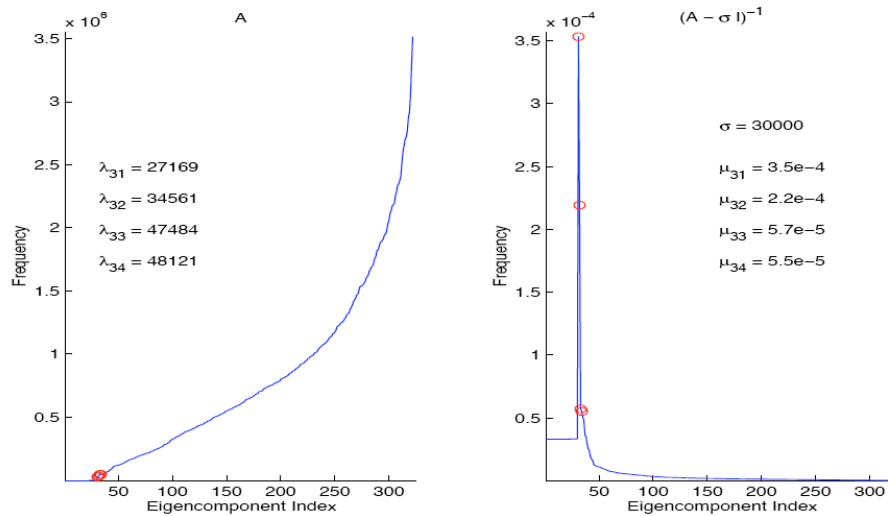

Figure 1. The spectra of A and $(A-\sigma I)^{-1}$ for a model problem.

We developed the filter algorithm [1] for the eigenvalue problem, which consists of two algorithms: an Inexact Shift-Invert Lanczos (ISIL) method for efficiently obtaining good eigenvector approximations and an inexact Newton-type method to refine the eigenvector approximations to the desired accuracy.

In the ISIL, we solve the shifted linear system inexactly, e.g., up to residual reduction of 0.1. The computed inexact solutions are not good enough for SIL to work but they are good enough to be considered as outputs of a band pass filter similar to the SIL $(K-\sigma M)^{-1}$ band pass filter. After applying the ISIL band pass

filter to create a new vector, we employ the Lanczos three-term recurrence process with complete re-orthogonalization on the newly created vector to expand the subspace. Finally we perform an explicit subspace projection to detect convergence and extract solutions. The following is the schematic description of the ISIL algorithm for a generalized eigenvalue problem Kx=λMx with a shift σ.

$$
\begin{aligned}
&\beta_0 q_0 = 0; \ \|q_1\|_M = 1 \\
&\texttt{For } k = 1, 2, 3 \ldots \texttt{ do} \\
&\quad \texttt{solve } (K - \sigma M)w = Mq_k \ \texttt{ for } \eta \leq 0.1 \\
&\quad \alpha_k = (Mq_k)^T w \\
&\quad w = w - \alpha_k q_k - \beta_{k-1} q_{k-1} \\
&\quad \texttt{for } i = 1, \ldots, k-2 \\
&\quad\quad h = (Mq_i)^T w \\
&\quad\quad w = w - h q_i \\
&\quad \texttt{end} \\
&\quad \beta_k = \|w\|_M \\
&\quad q_{k+1} = w/\beta_k \\
&\texttt{End} \\
&Q = [q_1, q_2, \ldots, q_k, q_{k+1}] \\
&S = Q^T K Q; \ T = Q^T M Q \\
&[\lambda, y] = eig(S, T); \ x = Qy
\end{aligned}
$$

Given a good eigenvector approximation calculated by the ISIL algorithm, we use the inexact JOCC (Jacobi Orthogonal Component Correction) method refine the eigenvector approximations to the desired accuracy. In the method, at each iteration, a JOCC equation is solved approximately to get a correction vector for the current eigenvector approximation. The following gives a schematic description of the inexact JOCC method.

$$
\begin{aligned}
&\texttt{On entry, } \theta_1 = \theta; \ q_1 = q; \ r_1 = r; \ j = 1 \\
&\texttt{While } q_j \texttt{ has not converged,} \\
&\quad \texttt{solve the JOCC equation approximately (e.g., } \eta = 0.1) \\
&\quad (I - (Mq_j)q_j^T)(K - \theta_j M)(I - (Mq_j)q_j^T)t_j = -(I - (Mq_j)q_j^T)r_j \\
&\quad q_{j+1} = q_j + t_j; \ q_{j+1} = \frac{q_{j+1}}{\|q_{j+1}\|_M}; \ a_{j+1} = Kq_{j+1}; \ \theta_{j+1} = q_{j+1}^T a_{j+1}; \\
&\quad b_{j+1} = Mq_{j+1}; \ r_{j+1} = a_{j+1} - \theta_{j+1} b_{j+1}; \ j = j + 1 \\
&\texttt{On exit, } \lambda = \theta_j; \ x = q_j
\end{aligned}
$$

We use both the SIL and the filter algorithm to calculate an eigenvalue near 30,000 for the model problem. The convergence histories are plotted in the Figure 2. As shown in the figure, during the first 4 iterations, the ISIL converges almost as fast as the exact SIL; the ISIL stagnates at the 5th iteration and switches to the inexact JOCC. The inexact JOCC converges at a rate similar to exact SIL for the last 3 iterations.
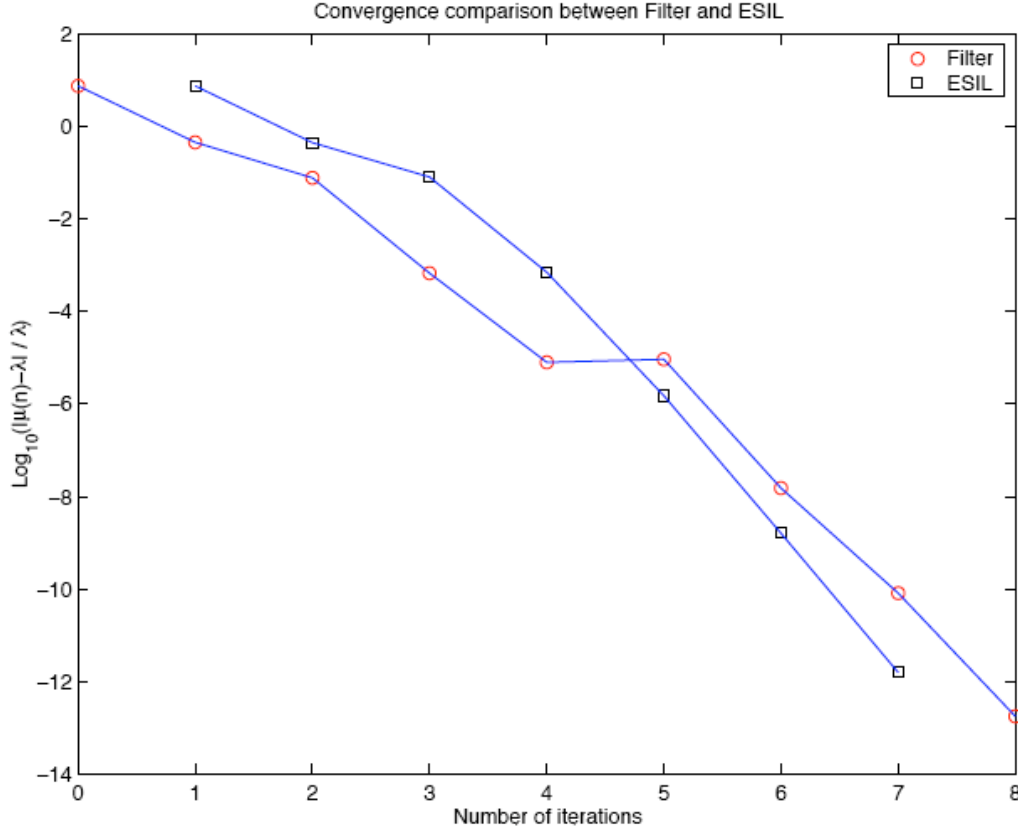
Figure 2. The convergence comparison of the SIL and filter algorithms for a model eigenvalue problem.

Block version of the filter algorithms along with deflation techniques are used to deal with the tightly clustered eigenvalues. A thick restarted Lanczos process is employed to limit memory usage. The details of the techniques are documented in reference [1]. The parallel filter algorithm is implemented in Omega3P, a finite-element eigenmode analysis for accelerator cavities. It ran on then the world largest unclassified parallel computers at National Energy Research Scientific Computing Center (NERSC) at Lawrence Berkeley National Laboratory (LBNL), and exhibited desirable parallel performance that scaled up to hundreds of processors.

2.2 The Hermitian and Skew-Hermitian Splliting (HSS) Method

To solve the numerical solutions of sparse linear systems of equations is one of the most common computational kernels in the simulations arising in the discretization of partial differential equations such as computational accelerator design.

We consider the solution of large sparse system of linear equations

$$Ax = b, \quad A \in \mathbb{C}^{n \times n} \text{ non-singular, and } x, b \in \mathbb{C}^n,$$

Where A is a non-Hermitian and positive definite matrix. Because the coefficient matrix A naturally possesses a Hermitian/skew-Hermitian (HS) splitting A = H + S, with

$$H = \frac{1}{2}(A + A^*) \quad \text{and} \quad S = \frac{1}{2}(A - A^*).$$

Collaborating with Bai et al, Golub proposed the following Hermitian skew-Hermitian splitting (HSS) method [2] in 2003 to iteratively compute a reliable and accurate approximate solution for the above system of linear equations.

The HSS iteration method. Given an initial guess $x^{(0)}$. For k = 0, 1, 2, . . . until $\{x^{(k)}\}$ converges, compute

$$\begin{cases} (\alpha I + H)x^{(k+\frac{1}{2})} & = (\alpha I - S)x^{(k)} + b, \\ (\alpha I + S)x^{(k+1)} & = (\alpha I - H)x^{(k+\frac{1}{2})} + b, \end{cases}$$

where α is a given positive constant.

Theoretical results show that the HSS iteration converges unconditionally to the unique solution of the original system of linear equations. The upper bounds of its contraction factor in a special weighted norm and its asymptotic convergence rate are only dependent on the spectrum of the Hermitian part H, but are independent of the spectrum of the skew-Hermitian part S as well as the eigenvectors of the matrices H, S and A. In addition, the optimal value of the parameter α can be determined by the lower and the upper eigenvalue bounds of the matrix H. The details of the convergence study and the convergence rate were discussed in the reference [2] as well as a model problem analysis for a discrete three-dimensional convection-diffusion equation.

Working with Benzi [3,5] and Bai [4], Golub proposed a preconditioning strategy based on the symmetric skew-symmetric splitting of the coefficient matrix for a 2x2 block positive semi-definite linear systems such as generalized saddle point problems. The optimal choice of the involved iteration parameter and the corresponding asymptotic convergence rate can be computed exactly [3,9]. Numerical examples confirm the theory and the effectiveness of the method [3,4,5]. The preconditioned HSS method was also applied to a discrete convection-diffusion equation and demonstrated its superiority with respect to the existing techniques [7].

In [6], Bai, Golub, Lu, and Yin further generalized the concept of Hermitian (or normal) and skew-Hermitian splitting for a non-Hermitian and positive-definite matrix nd introduce a new splitting, called positive-definite and skew-Hermitian splitting (PSS), and then established a class of PSS methods similar to the

Hermitian (or normal) and skew-Hermitian splitting (HSS or NSS) method for iteratively solving the positive-definite systems of linear equations.

In [8], Golub along with Bai and Li, derived necessary and sufficient conditions for guaranteeing the unconditional convergence of the preconditioned Hermitian and skew-Hermitian splitting iteration methods for the non-Hermitian and positive semidefinite systems of linear equations.

In [2] Bai, Golub and Ng also proposed to solve the linear systems with coefficient matrices $\alpha I + H$ and $\alpha I + S$ inexactly by iterative methods, e.g., solving the linear systems with coefficient matrix $\alpha I + H$ by the conjugate gradient (CG) method and those with coefficient matrix $\alpha I + S$ by the Lanczos or the conjugate gradient for normal equations (CGNE) method, to some prescribed accuracies, and obtained two special but quite practical inexact Hermitian skew-Hermitian splitting (IHSS) iterations. In [10], Bai, Golub and Ng studied the convergence properties of both IHSS(CG, Lanczos) and IHSS(CG, CGNE) in depth and investigated the optimal numbers of inner iteration steps in detail by considering both global convergence speed and overall computation workload.

In summary, the funding provided by this award supported fruitful research on the Hermitian skew-Hermitian Splitting method and its variants for solving non-Hermitian positive definite and semi-definite matrices. And 9 journal articles were published on the subject.

3. List of Publications
See the reference section below.

Reference

[1] Yong Sun, The filter algorithm for solving large-scale eigenproblems from accelerator simulations, Ph.D. Thesis, Stanford University, 2003.

[2] Z.-Z. Bai, G.H. Golub and M.K. Ng, Hermitian and skew-Hermitian splitting methods for non-Hermitian positive definite linear systems, SIAM J. Matrix Anal. Appl. 24 (2003), pp. 603–626.

[3] M. Benzi, M. Gander and G.H. Golub, Optimization of the Hermitian and skew-Hermitian splitting iteration for saddle-point problems, BIT 43 (2003), pp. 881–900.

[4] Z.-Z. Bai, G.H. Golub and J.-Y. Pan, Preconditioned Hermitian and skew-Hermitian splitting methods for non-Hermitian positive semidefinite linear systems, Numer. Math. 98 (2004), pp. 1–32.

[5] M. Benzi and G.H. Golub, A preconditioner for generalized saddle point problems, SIAM J. Matrix Anal. Appl. 26 (2004), pp. 20–41.

[6] Zhong-Zhi Bai, Gene H. Golub, Lin-Zhang Lu, and Jun-Feng Yin, Block Triangular and Skew-Hermitian Splitting Methods for Positive-Definite Linear Systems, SIAM J. Sci. Comput. Volume 26, Issue 3, pp. 844-863 (2005).

[7] D. Bertaccini, G.H. Golub, S. Serra Capizzano and C. Tablino Possio, Preconditioned HSS method for the solution of non-Hermitian positive definite linear systems and applications to the discrete convection–diffusion equation, Numer. Math. 99 (2005), pp. 441–484.

[8] Zhong-Zhi Bai, Gene H. Golub, And Chi-Kwong Li, Convergence Properties Of Preconditioned Hermitian And Skew-Hermitian Splitting Methods For Non-Hermitian Positive Semidefinite Matrices, Mathematics Of Computation, Volume 76, Number 257, January 2007, Pages 287–298

[9] Zhong-Zhi Bai, Gene H. Golub, And Chi-Kwong Li, Optimal Parameter In Hermitian And Skew-Hermitian Splitting Method For Certain Two-By-Two Block Matrices, Siam J. Sci. Comput., Vol. 28, No. 2, pp. 583–603, 2006.

[10] Zhong-Zhi Bai, Gene H. Golub, and Michael K. Ng, On inexact hermitian and skew-Hermitian splitting methods for non-Hermitian positive definite linear systems, Linear Algebra and its Applications, Volume 428, Issues 2-3, 15 January 2008, Pages 413-440.