A High-Performance Preconditioned SVD Solver for Accurately Computing Large-Scale Singular Value Problems in PRIMME

Lingfei Wu*, Andreas Stathopoulos (Advisor)*, Eloy Romero (Advisor)*

*College of William and Mary, Williamsburg, Virginia, USA

***Abstract*—The dramatic increase in the demand for solving large scale singular value problems has rekindled interest in iterative methods for the SVD. Unlike the remarkable progress in dense SVD solvers, some promising recent advances in large scale iterative methods are still plagued by slow convergence and accuracy limitations for computing smallest singular triplets. Furthermore, their current implementations in MATLAB cannot address the required large problems. Recently, we presented a preconditioned, two-stage method to effectively and accurately compute a small number of extreme singular triplets. In this research, we present a high-performance software, PRIMME_SVDS, that implements our hybrid method based on the state-of-the-art eigensolver package PRIMME for both largest and smallest singular values. PRIMME_SVDS fills a gap in production level software for computing the partial SVD, especially with preconditioning. The numerical experiments demonstrate its superior performance compared to other state-of-the-art methods and its good scalability performance under strong and weak scaling.**

I. Introduction

We consider the problem of finding a small number of extreme singular values and corresponding left and right singular vectors of a large sparse matrix $A \in \Re^{m \times n}$ $(m \geq n)$,

$$Av_i = \sigma_i u_i, \qquad i = 1, \ldots, k, \quad k \ll n \tag{1}$$

. Such problems arise in many scientific and engineering applications. Many diverse applications require a few of the largest singular triplets that play a critical role in compression and model reduction. A smaller, but increasingly important, set of applications require a few smallest singular triplets, including determination of matrix rank and computation of pseudospectrum [1].

This work is motivated by several considerations. In spite of the remarkable progress of dense SVD solver, an iterative method is still the only mean of addressing these large problems due to unfeasible memory requirements. Over the last decade, dedicated iterative methods have been developed [1]–[3] to effectively compute a few of singular triplets under limited memory. However, current state-of-the-art methods are only available in a MATLAB research implementations that cannot solve large-scale real world problems. In addition, the computation of the smallest singular triplets presents challenges both to the speed of convergence and the accuracy of iterative methods. In our previous work [1], we presented a preconditioned hybrid, two-stage method that achieves both efficiency and accuracy for both largest and smallest singular values under limited memory.

Given the above research activities in SVD algorithms, it is surprising that there is a lack of good quality software for computing the partial SVD, especially with preconditioning as shown in Table 1. Without preconditioning, SVDPACK [4] and PROPACK [5] implement variants of (block) Lanczos methods. In addition, PROPACK implements an implicitly restarted Lanczos bidiagonalization (LBD) method. However, SVDPACK can only compute largest singular triplets while PROPACK has to leverage shift-and-invert techniques to search for smallest. SLEPc offers some limited functionality for computing the partial SVD problem of a large, sparse rectangular matrix using various eigensolvers working on the augmented matrix $B = [0 \ A^T; A \ 0]$ or the normal equations matrix $C = A^T A$ [6]. It also implements a parallel LBD method but focuses mainly on largest singular values. With the growing size and difficulty of real-world problems, there is a clear need for a high quality SVD solver software that allows for additional flexibility, implements state-of-the-art methods, and allows for preconditioning.

In this work we address this need by developing a high quality SVD software, PRIMME_SVDS, based on the state-of-the-art package PRIMME (PReconditioned Iterative Multi-Method Eigensolver) [7]. This paper presents our development in PRIMME in order to provide state-of-the-art robust, high performance SVD solver supporting accurate computation of both largest and smallest singular triplets, for both square and rectangular matrices, with either preconditioning or without. Our numerical experiments show that PRIMME_SVDS can be considerably more efficient than all other methods when computing a few of the largest or smallest singular triplets, even without a preconditioner. We demonstrate the good scalability of PRIMME_SVDS for solving large-scale problems in various real applications under different parameter settings.

II. A Preconditioned Two-Stage SVD Solver

A. *Method for Accurate Computation*

As shown in Fig 1, seeking eigenpairs of B computes the smallest singular values accurately. However, convergence of eigenvalue is very slow since it is a highly interior eigenvalue problem. Computing eigenpairs of C has been theoretically and practically proven to have better convergence compared to all other methods [1], but it suffers an accuracy problem. In our previous work [1], We proposed that accuracy and efficiency can be achieved through a hybrid, two-stage meta-method, as illustrated in Fig 2. In the first stage, the proposed method solves an extreme eigenvalue problem on C up to the user required accuracy or up to the accuracy achievable by the normal equations. If further accuracy is required, the method switches to a second stage where it utilizes the eigenvectors and eigenvalues from C as initial guesses to a Jacobi-Davidson method on B, which has been enhanced by a refined projection method. The appropriate choices for tolerances, transitions, selection of target shifts, and initial guesses are handled automatically by the method. See reference [1] for more details.

B. A Parallel Implementation of PRIMME_SVDS

The main contribution of this work is to provide a high quality SVD solver software, which implements state-of-the-art methods for large-scale SVD problems. Our goal is to enable practitioners to solve a variety of large, sparse singular value problems with unprecedented efficiency, robustness, and accuracy. PRIMME is a high-performance parallel library that provides many state-of-the-art algorithms for the solution of large, sparse Hermitian and real symmetric eigenvalue problems [7]. Therefore, we have chosen to build our proposed PRIMME_SVDS method in [1] on top of PRIMME. To support PRIMME_SVDS, we have implemented many enhancements, such as a specialized refined projection method, a dynamic tolerance adjusting scheme and a two-stage meta-method in PRIMME. All these implementations follow the original design philosophy of PRIMME, where a parallel SPMD application provides the local vector dimensions on each processor. A global summation of the reduced value is demanded by dot products. Furthermore, the user must provide a parallel matrix-vector and preconditioning functions. The parallel characteristics of PRIMME_SVDS are listed in Table 2.

III. Experiments and Results

For our performance test we consider seeking small number of smallest and largest SVD problems from various real applications including ill-conditioned least-squares problem, DNA electrophoresis, and graph partitioning and clustering. To test the weak scaling performance, we used laplacian matrices due to easy manipulation of matrix size. We perform our experiments on the NERSC's Edison and a cluster SciClone at college of William and Mary. Basic information of matrices and parameter settings are listed in Table 3.

Figure 3 reveals the comparison of different methods on SciClone. When seeking smallest singular values, PRIMME_SVDS is substantially faster than JD, Krylov-Shur, and TRLAN by a factor of 3, 10, and 60 respectively. For largest singular values, PRIMME_SVDS is still faster than other methods for seeking a few. When seeking many largest singular triplets, we expect Krylov method will be a winner due to its better global convergence.

Figures 4, and 5 show the scalability performance of PRIMME_SVDS on Edison when seeking a small number of extreme singular triplets. In Figure 4, PRIMME_SVDS can achieve near-ideal speedup until 256 processes on largest rectangular matrix in the Florida Sparse Matrix Collection. When reaching 512 processes, the communication cost becomes close to computation cost due to too small local matrix dimension in each process. For difficult smallest SVD problems, a good preconditioner is a necessity for solving real large application. The good speedup largely relies on the efficiency of a preconditioner and heavy sparsity of a matrix. Figure 5 illustrates the same good scalability performance when seeking largest singular triplets without preconditioning. Finally, we also demonstrate its good performance under weak scaling when adjusting the nodes number of a 3D laplacian matrix from 80 to 200.

Acknowledgment

The author would also like to thank the Scientific Data Management Group for generously help on running experiments on Edison. This work is supported by NSF under grants No. CCF 1218349 and ACI SI2-SSE 1440700, and by DOE under a grant No. DE-FC02-12ER41890.

References

[1] L. Wu and A. Stathopoulos, "A preconditioned hybrid SVD method for computing accurately singular triplets of large matrices," *accepted for publication in SIAM J. Sci. Comput.*, preprint: http://arxiv.org/abs/1408.5535, April 2015.

[2] Z. Jia and D. Niu, "A refined harmonic Lanczos bidiagonalization method and an implicitly restarted algorithm for computing the smallest singular triplets of large matrices," *SIAM J. Sci. Comput.*, vol. 32, no. 2, pp. 714–744, 2010.

[3] M. E. Hochstenbach, "A Jacobi–Davidson type SVD method," *SIAM J. Sci. Comput.*, vol. 23, no. 2, pp. 606–628, 2001.

[4] M. W. Berry, "Large-scale sparse singular value computations," *Int. J. Supercomputer. Appl.*, vol. 6, no. 1, pp. 13–49, 1992.

[5] R. M. Larsen, "Lanczos bidiagonalization with partial reorthogonalization," *DAIMI Report Series*, vol. 27, no. 537, 1998.

[6] V. Hernandez, J. E. Roman, and V. Vidal, "Slepc: A scalable and flexible toolkit for the solution of eigenvalue problems," *ACM Trans. Math. Softw.*, vol. 31, no. 3, pp. 351–362, 2005.

[7] A. Stathopoulos and J. R. McCombs, "PRIMME: Preconditioned iterative multimethod eigensolver: Methods and software description," *ACM Trans. Math. Softw.*, vol. 37, no. 2, pp. 21:1–21:30, Apr. 2010.