

1 **THE SINGULAR VALUE DECOMPOSITION: ANATOMY OF**
2 **OPTIMIZING AN ALGORITHM FOR EXTREME SCALE***

3 JACK DONGARRA[†], MARK GATES[‡], AZZAM HAIDAR[‡], JAKUB KURZAK[‡], PIOTR
4 LUSZCZEK[‡], STANIMIRE TOMOV[‡], AND ICHITARO YAMAZAKI[‡]

5 **Abstract.** The computation of the Singular Value Decomposition, or SVD, has a long history,
6 with many improvements over the years, both in implementations and algorithmically. Here, we sur-
7 vey the evolution of SVD algorithms for dense matrices, discussing the motivation and performance
8 impact of changes. There are two main branches of dense SVD methods: bidiagonalization and Ja-
9 cobi. Bidiagonalization methods started with the implementation by Golub and Reinsch in Algol60,
10 which was subsequently ported to Fortran in the EISPACK library, and was later more efficiently
11 implemented in the LINPACK library, targeting contemporary vector machines. To address cache-
12 based memory hierarchies, the SVD algorithm was reformulated to use Level 3 BLAS in the LAPACK
13 library. To address new architectures, ScaLAPACK was introduced to take advantage of distributed
14 computing, and MAGMA was developed for accelerators such as GPUs. Algorithmically, the di-
15 vide and conquer and MRRR algorithms were developed to reduce the number of operations. Still,
16 these methods remained memory bound, so two-stage algorithms were developed to reduce memory
17 operations and increase the computational intensity, with efficient implementations in PLASMA,
18 DPLASMA, and MAGMA. Jacobi methods started with the two-sided method of Kogbetliantz and
19 the one-sided method of Hestenes. They have likewise had many developments, including parallel
20 and block versions, and preconditioning to improve convergence. In this paper, we investigate the
21 impact of these changes by testing various historical and current implementations on a common,
22 modern multicore machine and a distributed computing platform. We show that algorithmic and
23 implementation improvements have increased the speed of the SVD by several orders of magnitude,
24 while using up to 40 times less energy.

25 **Key words.** singular value decomposition, SVD, bidiagonal matrix, QR iteration, divide and
26 conquer, bisection, MRRR, Jacobi method, Kogbetliantz method, Hestenes method

27 **AMS subject classifications.** 15A18 15A23 65Y05

28 **1. Introduction.** The *singular value decomposition*, or SVD, is a very powerful
29 technique for dealing with matrix problems in general. The practical and theoretical
30 importance of the SVD is hard to overestimate, and it has a long and fascinating
31 history. A number of classical mathematicians are associated with the theoretical
32 development of the SVD [107], including Eugenio Beltrami (1835–1899), Camille Jor-
33 dan (1838–1921), James Sylvester (1814–1897), Erhard Schmidt (1876–1959), and
34 Hermann Weyl (1885–1955).

35 In recent years, the SVD has become a computationally viable tool for solving
36 a wide variety of problems that arise in many practical applications. The use of
37 the SVD in these applications is centered on the fact that they require information
38 about the rank of a matrix, or a low rank approximation of a matrix, or orthogonal
39 bases for the row and column spaces of a matrix. Applications are as diverse as least
40 squares data fitting [53], image compression [3], facial recognition [111], principal

*Submitted to the editors February 20, 2017.

Funding: This research is based upon work supported by the National Science Foundation under Grant No. 1339822, and by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. We also received support from NVIDIA and Intel. Furthermore, we would like to thank Intel for access to their distributed computing platform for testing ScaLAPACK and DPLASMA.

[†]University of Tennessee, Oak Ridge National Laboratory, University of Manchester (don-garra@icl.utk.edu)

[‡]University of Tennessee (mgates3@icl.utk.edu, haidar@icl.utk.edu, kurzak@icl.utk.edu, luszczek@icl.utk.edu, tomov@icl.utk.edu, iyamazak@icl.utk.edu)

41 component analysis [92], latent semantic analysis [28], and computing the 2-norm,
 42 condition number, and numerical rank of a matrix.

43 The SVD of an m -by- n matrix A is given by:

44 (1)
$$A = U\Sigma V^T \quad (A = U\Sigma V^H \text{ in the complex case}),$$

45 where U and V are orthogonal (unitary) matrices and Σ is an m -by- n matrix with
 46 real diagonal elements, σ_i , conventionally ordered such that:

47
$$\sigma_1 \geq \sigma_2 \geq \dots \geq \sigma_{\min(m,n)} \geq 0.$$

48 The σ_i are the *singular values* of A and the first $\min(m, n)$ columns of U and V are
 49 the *left* and *right singular vectors* of A , respectively.

50 Theoretically, the SVD can be characterized by the fact that the singular values
 51 are the square roots of the eigenvalues of $A^T A$, the columns of V are the corre-
 52 sponding eigenvectors, and the columns of U are the eigenvectors of AA^T , assuming
 53 distinct singular values. However, this is not a satisfactory basis for computation
 54 because roundoff errors in the formulation of $A^T A$ and AA^T often destroy pertinent
 55 information.

56 The key to using the SVD is the fact that it can be computed very effectively.
 57 There are two dominant categories of SVD algorithms for dense matrices: bidiag-
 58 onalization methods and Jacobi methods. The classical bidiagonalization method
 59 proceeds in the following three stages:

- 60 1. The matrix A is reduced to bidiagonal form: $A = U_1 B V_1^T$ if A is real ($A =$
 61 $U_1 B V_1^H$ if A is complex), where U_1 and V_1 are orthogonal (unitary if A is
 62 complex), and B is real and upper-bidiagonal when $m \geq n$ or lower bidiagonal
 63 when $m < n$, so that B is nonzero on only the main diagonal and either the
 64 first superdiagonal (if $m \geq n$) or the first subdiagonal (if $m < n$).
- 65 2. The SVD of the bidiagonal matrix B is computed: $B = U_2 \Sigma V_2^T$, where U_2 and
 66 V_2 are orthogonal and Σ is diagonal as described above. Several algorithms
 67 exist for the bidiagonal SVD, the original being QR iteration.
- 68 3. If desired, the singular vectors of A are then computed as $U = U_1 U_2$ and
 69 $V = V_1 V_2$.

70 This is the basic, efficient, and stable algorithm as posed by Golub and Kahan in
 71 1965 [53]. Golub and Reinsch [54] realized the first implementation of the SVD al-
 72 gorithm in Algol60, the programming language of the time. Their paper was later
 73 reproduced in the Wilkinson-Reinsch Handbook [117]. Bidiagonalization methods are
 74 covered in [sections 3 to 11](#), with additional tests of accuracy and performance on
 75 various matrix types in [sections 13 and 14](#).

76 In contrast, Jacobi methods apply plane rotations to the entire matrix A , with-
 77 out ever reducing it to bidiagonal form. Two-sided Jacobi methods, first proposed by
 78 Kogbetliantz in 1955 [76], iteratively apply rotations on both sides of A to bring it
 79 to diagonal form, while one-sided Jacobi methods, proposed by Hestenes in 1958 [68],
 80 apply rotations on one side to orthogonalize the columns of A , implicitly bring $A^T A$
 81 to diagonal. While Jacobi methods are often slower than bidiagonalization methods,
 82 there remains interest due to their simplicity, easy parallelization, and potentially bet-
 83 ter accuracy for certain classes of matrices. Jacobi methods are covered in [section 12](#),
 84 with additional tests in [sections 13 and 14](#).

85 This manuscript traces the development of the SVD algorithm over the past 50
 86 years, using various historical implementations. This development includes algo-
 87 rithmic improvements such as blocking, the divide and conquer and MRRR algorithms,

88 a two-stage reduction, as well as adapting to new computer architectures such as distributed
 89 memory, accelerators, and multicore CPUs. We compare the performance of
 90 all the implementations on a common multicore computer. Our focus is on computing
 91 all singular values and optionally singular vectors, for both square and tall dense
 92 matrices. For bisection and MRRR methods we also compute a subset of the singular
 93 values and vectors.

94 **2. Experimental Setup.** To test the various implementations, we ran six different tests:

- 95 1. Square matrices, singular values only (no vectors).
2. Square matrices, singular values and vectors.
3. Tall matrices, $m = 3n$, singular values only (no vectors).
4. Tall matrices, $m = 3n$, singular values and vectors.
5. Tall matrices, $m = 1000n$, singular values only (no vectors).
6. Tall matrices, $m = 1000n$, singular values and vectors.

100 When computing singular vectors, we computed the *reduced SVD* consisting of the
 101 first $\min(m, n)$ columns of U and V , and $\min(m, n)$ rows and columns of Σ . This is
 102 the most useful part computationally, sufficient for many applications such as solving
 103 least squares problems, and we subsequently identify U , V , and Σ with those of the
 104 reduced SVD, which still satisfy (1). For LAPACK, the reduced SVD corresponds to
 105 job=“s” for both U and V . We store U and V separately from A , i.e., they do not
 106 overwrite A . Where applicable, we query for the optimal workspace size; otherwise,
 107 we use the maximum documented workspace size. This ensures that we always use
 108 the “fast” path in codes, including blocking and other optimizations.

109 Unless indicated, matrices have random entries from a uniform distribution on
 110 $(0, 1)$. For some tests, we generate singular values Σ according to one of the distributions
 111 below, then form $A = U\Sigma V^T$ where U and V are random orthogonal matrices
 112 from the Haar distribution [106]. Where given, κ is the condition number of A .

- 113 • Σ random: singular values are random uniform on $(0, 1)$. The condition
 114 number is not determined *a priori*.
- 115 • arithmetic: $\sigma_i = 1 - \frac{i-1}{n-1} \left(1 - \frac{1}{\kappa}\right)$ for $i = 1, \dots, n$.
- 116 • geometric: $\sigma_i = \kappa^{-(i-1)/(n-1)}$ for $i = 1, \dots, n$.
- 117 • log-random: singular values are random in $(\frac{1}{\kappa}, 1)$ such that their logarithms
 118 are random uniform on $(\log \frac{1}{\kappa}, \log 1)$.
- 119 • cluster at $\frac{1}{\kappa}$: $\Sigma = [1, \frac{1}{\kappa}, \dots, \frac{1}{\kappa}]$.
- 120 • cluster at 1: $\Sigma = [1, \dots, 1, \frac{1}{\kappa}]$.

121 All tests were performed in double-precision real arithmetic. Except for PLASMA
 122 and MPI-based implementations, which initialize memory in parallel, we used `numactl`
 123 `--interleave=all` to distribute memory across CPU sockets, and the CPU cache was
 124 flushed before the SVD function call. To avoid repeating minor differences, we shall
 125 generally assume that A is real and $m \geq n$. Operations for complex or $m < n$ are
 126 analogous.

127 We conducted experiments on a two-socket Intel Sandy Bridge Xeon E5-2670
 128 running at 2.6 GHz, with 8 cores per socket, a theoretical double-precision peak of
 129 333 Gflop/s, and 64 GiB of main memory. The measured practical dgemm peak is
 130 313.6 Gflop/s and dgemv peak is 13.9 Gflop/s (55.8 GB/s). The STREAM triad
 131 benchmark [91] measured the memory bandwidth as 57.8 GB/s with 16 OpenMP
 132 threads. All CPU implementations were compiled with gcc and linked against Intel’s
 133 Math Kernel Library (MKL) version 11.2.3 [71].

136 GPU results used an NVIDIA Kepler K40c with 15 multiprocessors, each containing 192 CUDA cores. The theoretical double precision peak performance is
 137 1682 Gflop/s. On the GPU, 12 GiB of device memory can be accessed at a theoretical
 138 bandwidth of 288 GB/s. The measured practical dgemm peak is 1243.1 Gflop/s and
 139 dgemv peak is 45.3 Gflop/s (181.2 GB/s). For the GPU implementation, we used
 140 CUDA version 7.0 [94].
 141

142 **3. EISPACK Implementation.** The EISPACK project was an effort to de-
 143 velop a software library for numerical computation of eigenvalues and eigenvectors of
 144 matrices based on algorithms and ideas that were mainly contained in the Wilkinson-
 145 Reinsch Handbook. EISPACK was a transliteration of these Algol programs into
 146 Fortran. It contains subroutines for calculating the eigenvalues of nine classes of
 147 matrix problems: complex general, complex Hermitian, real general, real symmetric,
 148 real symmetric banded, real symmetric tridiagonal, special real tridiagonal, gener-
 149 alized real, and generalized real symmetric. In addition, it includes subroutines to
 150 perform a singular value decomposition [50]. Some routines were updated to imple-
 151 ment improvements in the numerical accuracy and achieve portability across different
 152 computing systems. However, the basic organization and access to matrix elements
 153 was kept in the Algol style.

154 To arrange multidimensional arrays in linear storage such as memory, Algol uses
 155 row-major order (each row is contiguous in memory), while Fortran uses column-
 156 major order (each column is contiguous in memory). Array layout is critical for
 157 correctly passing arrays between programs written in different languages. It is also
 158 important for performance when traversing an array since accessing array elements
 159 that are contiguous in memory is usually much faster than accessing elements which
 160 are not, due to the structure of the memory cache hierarchy. In the Algol routines,
 161 and subsequently the Fortran routines of EISPACK, matrix elements were referenced
 162 by row, thus causing great inefficiencies in the Fortran EISPACK software on modern
 163 cache based computer systems.

164 Written in standard Fortran 77, with no outside dependencies, EISPACK still
 165 compiles with a modern Fortran compiler. Figure 1 shows its performance results on
 166 a modern computer with the six test problems described in section 2. EISPACK has
 167 no notion of parallelism, so the code runs on only a single core. The operation count
 168 formulas here assume two QR iterations per singular value, and that an initial QR
 169 reduction is not done [23].

170 For square matrices without computing singular vectors, asymptotic performance
 171 is limited to 0.74 Gflop/s for one core, while when computing singular vectors, per-
 172 formance nearly triples to 2.17 Gflop/s. As is common, small sizes perform better
 173 because the entire matrix fits into L2 cache. Performance for the tall 3:1 and 1000:1
 174 cases is less than the square case, but exhibits a similar improvement when computing
 175 singular vectors compared with no vectors. For comparison, the practical peak using
 176 matrix-multiply on one core is 20 Gflop/s.

177 **4. LINPACK Implementation Using BLAS.** In the 1970s, the Level 1
 178 BLAS (Basic Linear Algebra Subroutines) [79] were introduced as a standard set of
 179 interfaces to perform common linear algebra operations. The Level 1 BLAS includes
 180 operations with $O(n)$ floating-point operations (flops), such as vector sum ($y = \alpha x + y$,
 181 called `daxpy`). The LINPACK project [39] reimplemented the SVD algorithm, along
 182 with other linear algebra algorithms, using Level 1 BLAS for efficient execution on
 183 the vector supercomputers of the 1970s and 1980s. It uses Fortran’s native column-
 184 major order, which makes better use of cache and memory bandwidth. However,

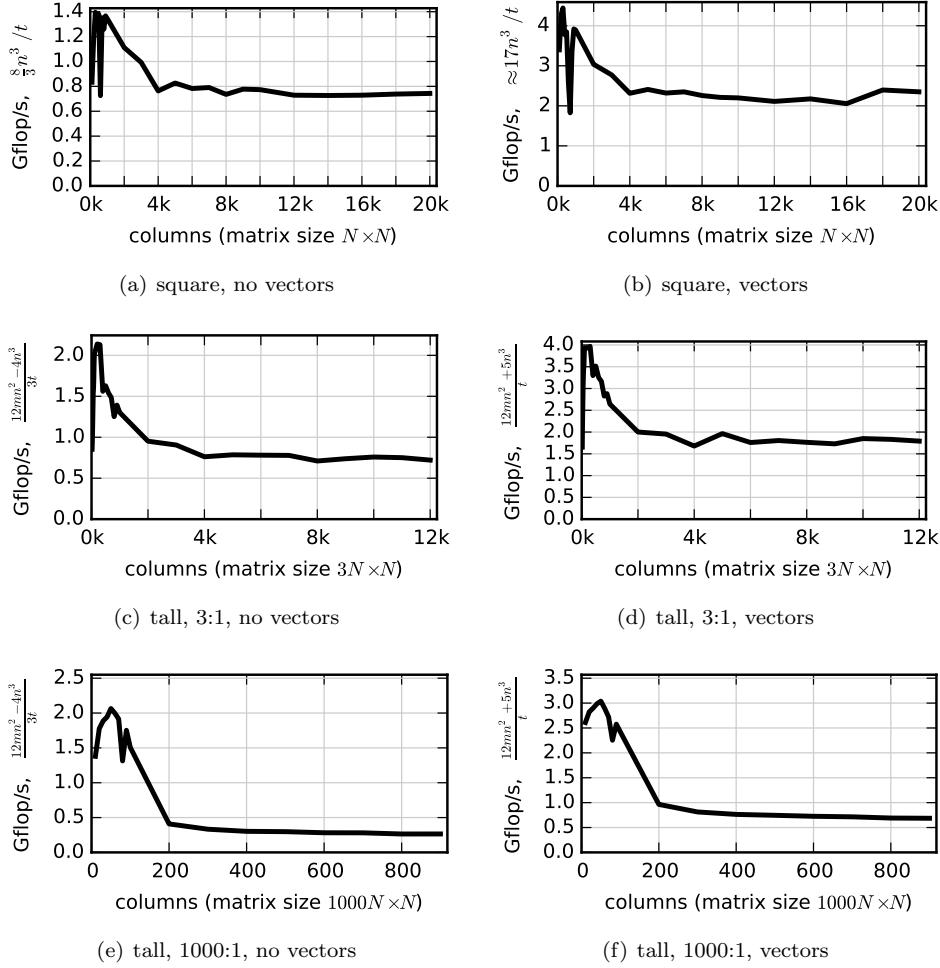


FIG. 1. Results for EISPACK, which uses only one core.

185 using Level 1 BLAS, LINPACK is limited by the memory bandwidth and receives
 186 little benefit from multiple cores. We see in Figure 2 that LINPACK achieves up to
 187 $3.9\times$ speedup over EISPACK for the square, no vectors case, and $2.7\times$ speedup for
 188 the square, vectors case. When computing a tall $m \times n$ matrix with $m = 100n$,
 189 using multithreaded BLAS on 16 cores yields some benefit, with speedups of $22.5\times$
 190 and $13.5\times$ over EISPACK for the no vectors and vectors cases, respectively, compared
 191 with speedups of $7.6\times$ and $3.9\times$, respectively, with single-threaded BLAS. In some
 192 instances, for large matrices such as $n = 16,000$, the code hung, appearing in a “sleep”
 193 state in `ps`, so we were unable to collect all data points.

194 **5. LAPACK Implementation Based on Blocked Householder Transformations.** While successful for vector-processing machines, Level 1 BLAS were not
 195 a good fit for the cache-based machines that emerged later in the 1980s. For cache-
 196 based machines, it is preferable to use higher-level operations such as matrix-matrix
 197 multiply, which is implemented by splitting a matrix into small blocks that fit into
 198 cache memory and performing small matrix-matrix multiplies on these blocks. This

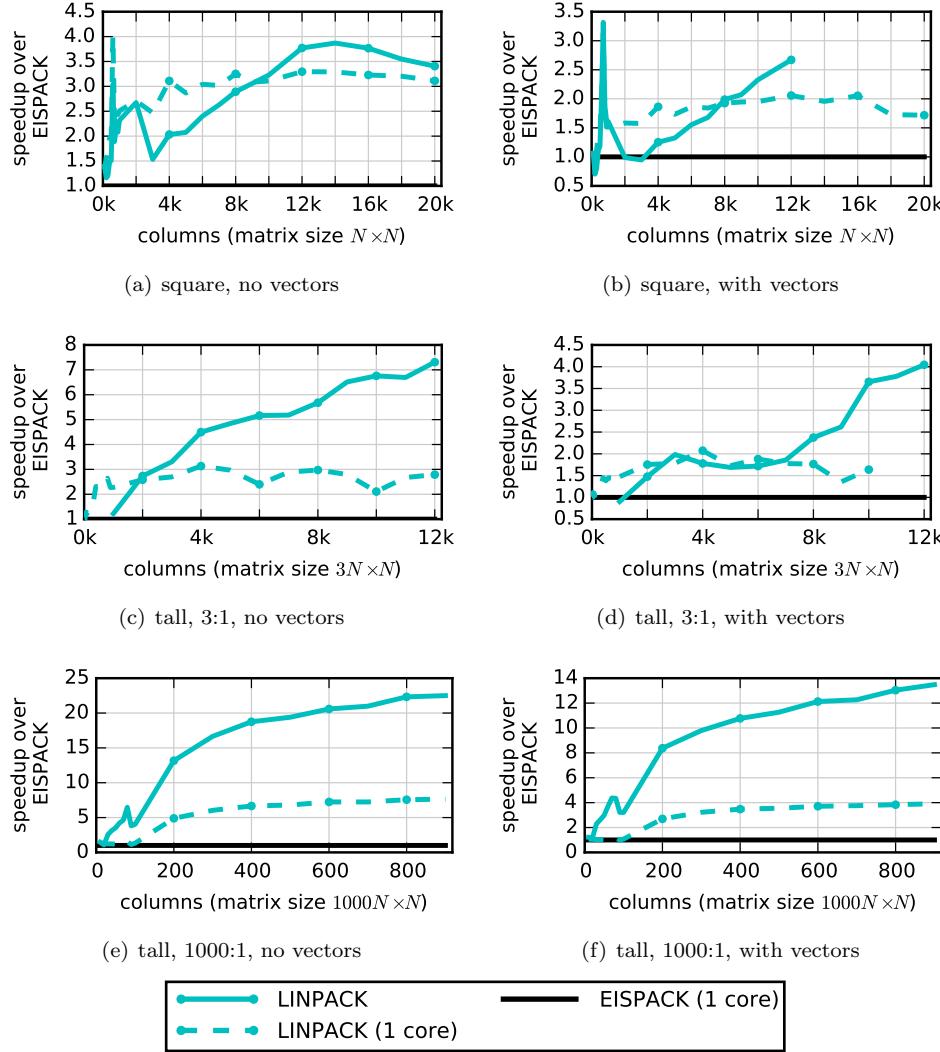


FIG. 2. Comparison of LINPACK to EISPACK.

200 avoids excessive data movement between cache and main memory. This led to the
 201 Level 2 BLAS [41] for operations with $O(n^2)$ flops, such as general matrix-vector
 202 multiply ($y = \alpha Ax + \beta y$, called `dgemv`); and Level 3 BLAS [40] for operations with
 203 $O(n^3)$ flops on $O(n^2)$ data, such as general matrix-matrix multiply ($C = \alpha AB + \beta C$,
 204 called `dgemm`). Level 1 and 2 BLAS access $O(1)$ elements per operation, and are
 205 thus limited in performance by the memory bandwidth. Level 3 BLAS benefit from
 206 the *surface-to-volume* effect of having only $O(n^2)$ elements to access for $O(n^3)$ op-
 207 erations. The performance of Level 1, 2, and 3 BLAS are compared in Figure 3,
 208 showing the significant benefit of Level 3 BLAS. The BLAS provides a means to write
 209 high-level, high-performance, portable numerical software. Optimized BLAS libraries
 210 are available, both from commercial vendors such as the Intel Math Kernel Library
 211 (MKL) [71] and the IBM Engineering and Scientific Subroutine Library (ESSL) [70],

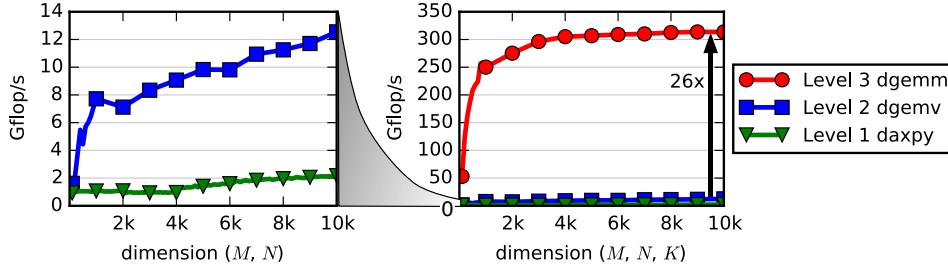


FIG. 3. Comparison of Level 1, 2, and 3 BLAS performance.

212 and in open-source libraries such as OpenBLAS [96] and ATLAS [115]. These math
 213 libraries often also included optimized versions of LAPACK, ScaLAPACK, and other
 214 numerical libraries. Our tests used the optimized routines available in Intel MKL.

215 **5.1. Blocked Householder Transformations.** With the introduction of Level
 216 3 BLAS, algorithms were recast using matrix multiplies, and LINPACK was re-
 217 designed into LAPACK [2] to use Level 3 BLAS where possible. The redesign for
 218 one-sided factorizations such as QR, LU, and Cholesky is relatively easier than re-
 219 ductions for eigenvalue problems and the SVD because the transformations used in
 220 QR, LU, and Cholesky are applied from only the left side [40]. Consecutive elemen-
 221 tary transformations are restricted to a block of columns at a time, referred to as the
 222 panel (depicted in Figure 4(a)), and updates to the rest of the matrix, referred to as
 223 the trailing matrix, are delayed. The transformations used for a panel are blocked
 224 together [14, 104] and applied to the trailing matrix as Level 3 BLAS.

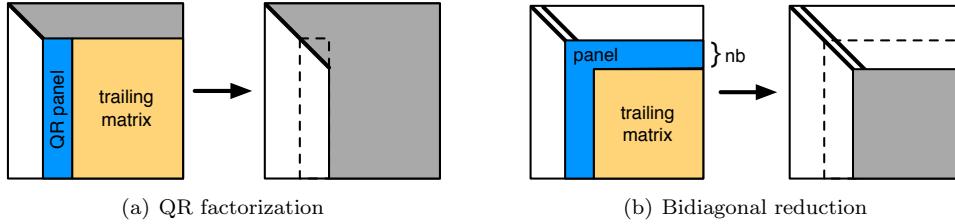


FIG. 4. Comparison of panels and trailing matrix.

225 On the other hand, the reduction of a matrix A to bidiagonal form is done by
 226 applying orthogonal matrices on both the left and right side of A —hence it is called
 227 a “two-sided factorization.” The two-sided transformations create more data depen-
 228 dencies, which make it impossible to entirely remove matrix-vector products involving
 229 the trailing matrix (as in the one-sided factorizations). The panel becomes a block
 230 row and block column, as shown in Figure 4(b), but panel operations also involve
 231 the entire trailing matrix. Dongarra et al. [42] developed the blocked algorithm for
 232 the bidiagonal reduction. The algorithm as implemented in LAPACK is given in
 233 Algorithm 1, and can be summarized as follows.

234 Two orthogonal matrices, U_1 and V_1 , are applied on the left and right side, re-
 235 spectively, of an $m \times n$ matrix A to reduce it to bidiagonal form, $B = U_1^T A V_1$. The

Algorithm 1 LAPACK implementation of bidiagonal reduction. In the $\{\cdot\}_{\dots}$ notation, only the indicated column or row should be computed, not the entire matrix product. y_i and x_i are computed as a series of matrix-vector products by distributing v_i and u_i . In LAPACK, Householder vectors representing V and U overwrite A . Auxiliary function $\text{householder}(x)$ (`dlarfg` in LAPACK) returns τ and v that define a Householder reflector H_i , and the updated vector $\hat{x} = H_i x = [\pm \|x\|, 0, \dots, 0]^T$.

```

// bidiagonal reduction (A is  $m \times n$ ; assumes  $m \geq n$  and  $n$  divisible by  $n_b$ )
function gebrd(  $A$  )
  for  $i = 1 : n$  by  $n_b$ 
     $(V; Y; X; U) = \text{labrd}( A_{i:m, i:n} )$ 
     $A_{i+n_b:m, i+n_b:n} = A_{i+n_b:m, i+n_b:n} - VY^T - XU^T$ 
  end
end function

// panel of bidiagonal reduction (A is  $m \times n$ ; assumes  $m \geq n$ )
function labrd(  $A$  )
   $V, Y, X, U$  initially empty
  for  $i = 1 : n_b$ 
    // compute column  $i$  of  $A_{(i-1)}$  using (4),
    // then compute  $H_i$  to eliminate below diagonal
     $A_{i:m, i} = \{A - V_{i-1}Y_{i-1}^T - X_{i-1}U_{i-1}^T\}_{i:m, i}$ 
     $(\tau_i; v_i; A_{i:m, i}) = \text{householder}( A_{i:m, i} )$ 
     $y_i = \tau_i A_{(i-1)}^T v_i = \tau_i (A - V_{i-1}Y_{i-1}^T - X_{i-1}U_{i-1}^T)^T v_i$ 

    // compute row  $i$  of  $H_i A_{(i-1)}$  using (2) and (4),
    // then compute  $G_i$  to eliminate right of super-diagonal
     $A_{i, i+1:n} = \{A - V_i Y_i^T - X_{i-1} U_{i-1}^T\}_{i, i+1:n}$ 
     $(\pi_i; u_i; A_{i, i+1:n}) = \text{householder}( A_{i, i+1:n} )$ 
     $x_i = \pi_i (A_{(i-1)} - v_i y_i^T) u_i = \pi_i (A - V_i Y_i^T - X_{i-1} U_{i-1}^T) u_i$ 
  end
  return  $(V_{n_b+1:m, 1:n_b}; Y_{n_b+1:n, 1:n_b}; X_{n_b+1:m, 1:n_b}; U_{n_b+1:n, 1:n_b})$ 
end function

```

236 matrices U_1 and V_1 are represented as products of elementary Householder reflectors:

237 $U_1 = H_1 H_2 \dots H_n \quad \text{and} \quad V_1 = G_1 G_2 \dots G_{n-1}.$

238 Each H_i and G_i has the form

239 $H_i = I - \tau_i v_i v_i^T \quad \text{and} \quad G_i = I - \pi_i u_i u_i^T,$

240 where τ_i and π_i are scalars, and v_i and u_i are vectors. H_i eliminates elements below
 241 the diagonal in column i , while G_i eliminates elements right of the superdiagonal in
 242 row i . Let $A_{(i-1)}$ be the reduced matrix A after step $i-1$. Applying H_i on the left
 243 yields

244 (2) $H_i A_{(i-1)} = (I - \tau_i v_i v_i^T) A_{(i-1)} = A_{(i-1)} - v_i y_i^T,$

246 while applying both H_i and G_i yields

$$247 \quad (3) \quad \begin{aligned} A_{(i)} &= H_i A_{(i-1)} G_i = (I - \tau_i v_i v_i^T) A_{(i-1)} (I - \pi_i u_i u_i^T) \\ 248 &= A_{(i-1)} - v_i y_i^T - x_i u_i^T, \end{aligned}$$

249 where $y_i = \tau_i A_{(i-1)}^T v_i$ and $x_i = \pi_i (A_{(i-1)} - v_i y_i^T) u_i$. Blocking together i applications
250 of (3), we obtain

$$251 \quad (4) \quad A_i = H_i \cdots H_1 A G_1 \cdots G_i = A - V_i Y_i^T - X_i U_i^T,$$

253 where $U_i = [u_1, \dots, u_i]$, and similarly with V_i , X_i , and Y_i . Note that it is possible to
254 update just part of A , namely the i -th column and row of A , in order to proceed with
255 the computation of the H_i and G_i . Thus, a delayed update is possible, but at each
256 step we still compute two matrix-vector products involving the entire trailing matrix
257 of A . As a result, if $m = n$, the entire factorization takes approximately $\frac{8}{3}n^3$ flops,
258 with half of the operations in Level 2 BLAS (matrix-vector products), while the other
259 half are in Level 3 BLAS.

260 **5.2. QR Iteration.** After the bidiagonal reduction, LAPACK solves the bidiagonal
261 SVD using QR iteration, similar to EISPACK and LINPACK, or using divide
262 and conquer, which is described later in section 7. The original QR iteration algo-
263 rithm computed singular values to high absolute accuracy, meaning small singular
264 values might be inaccurate. Demmel and Kahan [31] derived the implicit zero-shift
265 QR iteration algorithm and proved that it computes all singular values to high relative
266 accuracy; this is used as needed for accuracy by LAPACK when computing singular
267 vectors. Accuracy is discussed further in section 13.

268 The **qd** (German: quotienten-differenzen) [100] and *differential* **qd** (**dqd**) [101]
269 algorithms proposed by Rutishauser actually predate QR iteration and are among the
270 first algorithms for computing singular values for modern computers. Subsequent to
271 Demmel and Kahan's work, Fernando and Parlett [48] derived a shifted version called
272 **dqds** that allowed using shifts to maintain fast convergence, while still maintaining
273 high relative accuracy. This is used by LAPACK when computing singular values
274 only (no vectors). Quite a few more variants of **qd** can be derived [97].

275 **5.3. Computation of Singular Vectors.** Normally, LAPACK stores orthog-
276 onal matrices in an implicit fashion as a sequence of Householder reflectors, each
277 represented by a scalar τ_i and vector u_i . For QR iteration to accumulate the singular
278 vectors, it first generates U_1 and V_1 explicitly (using **dorgbr**); this is essentially ap-
279 plying block Householder reflectors to an identity matrix, as a series of Level 3 BLAS
280 operations.

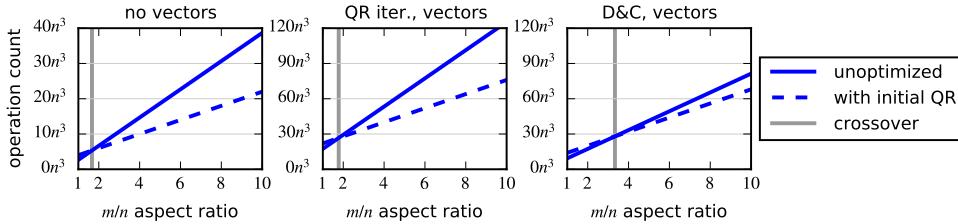
281 The QR iteration algorithm then updates U_1 and V_1 by applying the Givens
282 rotations used to reduce the bidiagonal matrix to diagonal. This is implemented in a
283 Level 2 BLAS-like fashion, where an entire sequence of n Givens rotations is applied
284 to update the entire U and V matrices (using **dlasr**). Recently, Van Zee et al. [113]
285 developed a Level 3 BLAS-like implementation of applying Givens rotations, which
286 they found made the SVD using QR iteration competitive with the SVD using divide
287 and conquer (discussed in section 7).

288 **5.4. Initial QR Factorization.** If $m \gg n$, it is more efficient to first perform a
289 QR factorization of A , and then compute the SVD of the n -by- n matrix R , since if $A =$
290 QR and $R = U\Sigma V^T$, then the SVD of A is given by $A = (QU)\Sigma V^T$. Similarly, if $m \ll$

291 n , it is more efficient to first perform an LQ factorization of A . Chan [23] analyzed this
 292 optimization, showing that it reduces the number of floating point operations. The
 293 operation counts are given in Figure 5, with the theoretical crossover points based on
 294 flops. Figure 6 plots the operation count as the ratio $m:n$ increases, illustrating the
 295 large savings as a matrix becomes taller. The results for tall matrices in Figures 8(c)
 296 to 8(f) show that LAPACK achieves significant speedups, such as $120\times$ compared
 297 with EISPACK. This is a result of the reduced operation count and that much of the
 298 computation is done via Level 3 BLAS in QR factorization, followed by a relatively
 299 small square SVD problem.

	QR iteration, no vectors	QR iteration, with vectors	D&C, with vectors
Unoptimized	$4mn^2 - \frac{4}{3}n^3$	$12mn^2 + \frac{16}{3}n^3$	$8mn^2 + \frac{4}{3}n^3$
With initial QR	$2mn^2 + 2n^3$	$6mn^2 + 16n^3$	$6mn^2 + 8n^3$
Theoretical crossover	$m \geq \frac{5}{3}n$	$m \geq \frac{16}{9}n$	$m \geq \frac{10}{3}n$

FIG. 5. Floating point operation counts.

FIG. 6. Operation counts as ratio $m:n$ increases (i.e., matrix gets taller), showing crossover where doing initial QR factorization is beneficial.

300 Since bidiagonal divide and conquer (D&C, discussed in section 7) always operates
 301 on a square matrix, doing an initial QR factorization with D&C results in less of an
 302 improvement than with QR iteration. Asymptotically, as the ratio $m:n \rightarrow \infty$, the
 303 initial QR factorization, generating Q , and multiplying by Q are responsible for most
 304 of the cost, as shown by the profile of the 1000:1 case in Figure 7. As a result, using
 305 QR iteration or divide and conquer yield the same performance for very tall-skinny
 306 matrices. The crossover points in Figure 5 are based solely on flop counts. Since doing
 307 an initial QR also shifts operations from Level 2 to Level 3 BLAS, the crossovers are
 308 ideally tunable parameters, for instance by LAPACK's `ilaenv` tuning function.

309 **5.5. Results.** An overview of all the phases in the complete SVD is given in
 310 Algorithm 2, with a profile of the time spent in each phase in Figure 7. For the
 311 square, no vectors case, we see that the bidiagonal reduction (blue with \\\ hatching)
 312 takes almost the entire time, while QR iteration (green, no hatching) takes very little
 313 time, as expected since QR iteration is $O(n^2)$ while the bidiagonal reduction costs
 314 $\frac{8}{3}n^3$ flops. When computing singular vectors, the QR iteration time becomes nearly
 315 half of the overall time, due to accumulating $U = U_1 U_2$ and $V = V_1 V_2$. Generating
 316 the explicit U_1 and V_1 matrices (orange with \\\ hatching) is a small portion of the
 317 time, even though together they have nominally the same operation count ($\frac{8}{3}n^3$) as

Algorithm 2 Overview of SVD algorithm using QR iteration (`dgesvd`) for $m \geq n$. Accumulating $U = U_1 U_2$ and $V = V_1 V_2$ occurs during QR iteration. \dagger Marked lines are required only when computing singular vectors.

Description	LAPACK Routine	Cost	Cost for $m \gg n$
if $m \gg n$ then			
$\hat{A} = A$ (QR factorization)	<code>dgeqrf</code>		$2mn^2 - \frac{2}{3}n^3$
$\hat{A} = R$			
else			
$\hat{A} = A$			
end			
$U_1 B V_1^T = \hat{A}$ (bidiagonalization)	<code>dgebrd</code>	$4mn^2 - \frac{4}{3}n^3$	$\frac{8}{3}n^3$
generate explicit U_1	<code>dorgbr</code> \dagger	$2mn^2 - \frac{2}{3}n^3$	$\frac{4}{3}n^3$
generate explicit V_1	<code>dorgbr</code> \dagger	$\frac{4}{3}n^3$	$\frac{4}{3}n^3$
$U_2 \Sigma V_2^T = B$ (QR iteration)	<code>dbdsqr</code>	$O(n^2)$	$O(n^2)$
$U = U_1 U_2$	" " \dagger	$6mn^2$	$6n^3$
$V = V_1 V_2$	" " \dagger	$6n^3$	$6n^3$
if $m \gg n$ then			
generate explicit Q	<code>dorgqr</code> \dagger		$2mn^2 - \frac{2}{3}n^3$
$U = QU$	<code>dgemm</code> \dagger		$2mn^2$
end			
Total cost (with vectors \dagger)		$12mn^2 + \frac{16}{3}n^3$	$6mn^2 + 16n^3$
Total cost (no vectors)		$4mn^2 - \frac{4}{3}n^3$	$2mn^2 + 2n^3$

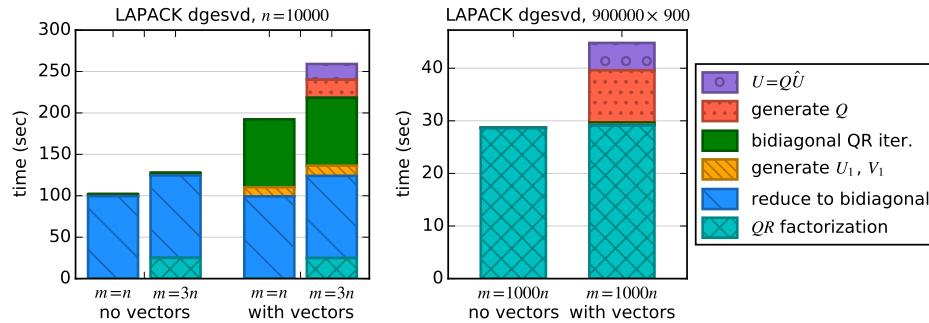


FIG. 7. Profile of LAPACK SVD. Left is 10000×10000 and 30000×10000 problem. QR factorization reduces the 30000×10000 matrix to a 10000×10000 matrix. Right is a 900000×900 problem, where reduction to bidiagonal and QR iteration become vanishingly small.

318 the bidiagonal reduction. This exemplifies the performance difference between Level 2
 319 BLAS, in the bidiagonal reduction, and Level 3 BLAS, in generating U_1 and V_1 .

320 The tall 3:1 matrix first does an initial QR factorization, resulting in a square R
 321 matrix the same size as the square case (10000×10000). Thus the profile for the 3:1
 322 case simply adds the QR factorization, generating Q , and multiplying $U = Q\hat{U}$ steps
 323 to the square case. For the tall 1000:1 matrix, the initial QR factorization dominates
 324 the overall time, with the subsequent bidiagonal reduction and QR iteration becoming
 325 vanishingly small. When vectors are computed, generating Q and multiplying $U =$
 326 $Q\hat{U}$ add significant time, while generating U_1 and V_1 and updating $U = U_1 U_2$ and

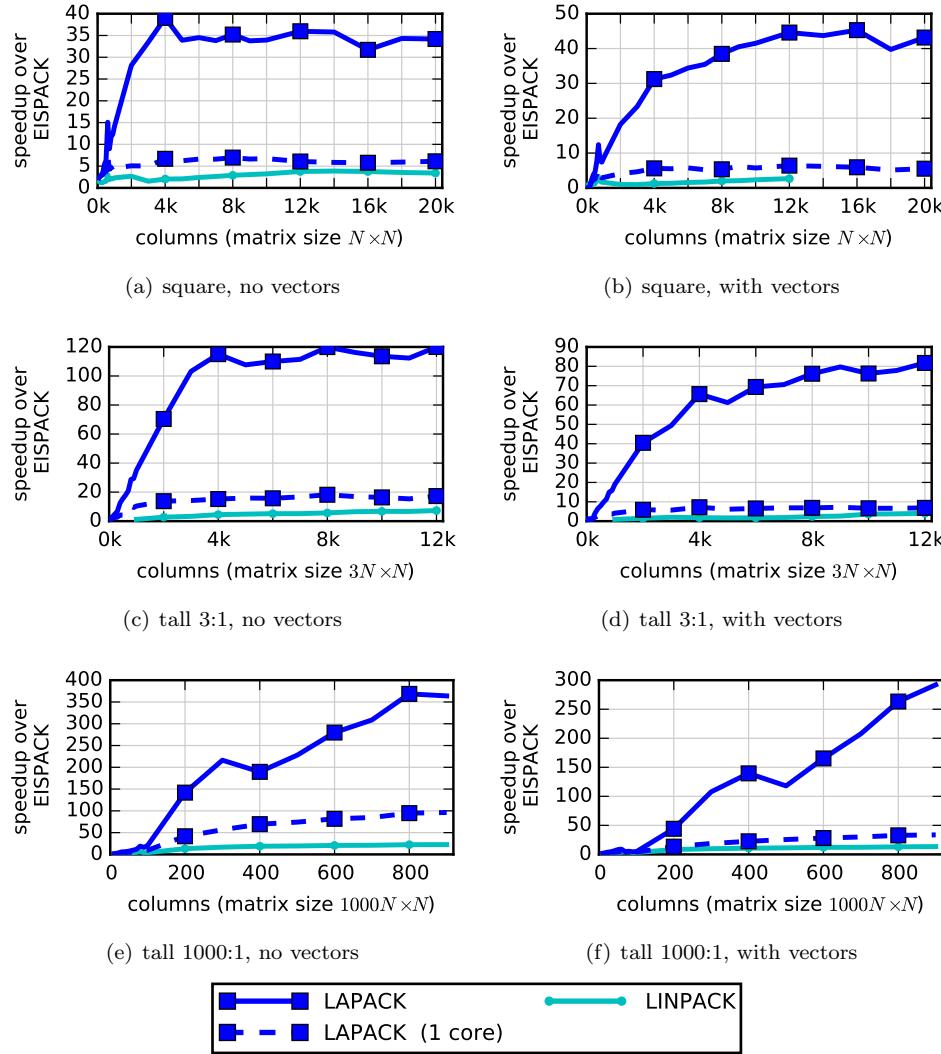


FIG. 8. Comparison of LAPACK, LINPACK, and EISPACK. Solid lines represent 16-core runs; dashed lines represent single core runs.

327 $V = V_1 V_2$ during QR iteration are also vanishingly small. Thus for very tall-skinny
 328 matrices, the performance is dominated by operations rich in Level 3 BLAS.

329 Figure 8 shows the speedup that LAPACK achieves compared with EISPACK.
 330 Even a single-core implementation may achieve over $5\times$ speedup. But the real poten-
 331 tial is shown when using multiple cores (16 in this case)—a $45\times$ speedup is possible for
 332 square matrices and over $350\times$ speedup for tall matrices with a 1000:1 row-to-column
 333 ratio. The square, no vectors case in Figure 8(a) is dominated by the bidiagonalization,
 334 as the subsequent bidiagonal SVD is $O(n^2)$. With Level 3 BLAS being significantly
 335 faster than Level 2 BLAS, and half the operations in Level 2 BLAS, we expect the

336 bidiagonalization to be about $2\times$ the speed of Level 2 BLAS. Modeling the time as

$$337 \quad t = \frac{4n^3}{3r_2} + \frac{4n^3}{3r_3}$$

338 with the Level 2 BLAS rate as $r_2 = 13.4$ Gflop/s and the Level 3 BLAS rate as
 339 $r_3 = 315$ Gflop/s (from [Figure 3](#)), we obtain a theoretical bound of 25.7 Gflop/s.
 340 This yields a speedup of $34.7\times$ over EISPACK—exactly what we see for LAPACK in
 341 [Figure 8\(a\)](#). When computing singular vectors, LAPACK achieves a higher speedup,
 342 up to $45.3\times$, reflecting that computation of U_1 and V_1 uses Level 3 BLAS. The tall
 343 matrix cases achieve even higher speedups because much of the work is done in the
 344 initial QR factorization.

345 **5.6. Level 2.5 BLAS Implementation.** Since the bidiagonalization performance
 346 is limited by the Level 2 BLAS operations, Howell et al. [69] sought to optimize
 347 these operations by observing that several Level 2 operations can be done together,
 348 thus reducing memory transfers by keeping data in cache. This technique of fusing
 349 several Level 2 operations together was called the Level 2.5 BLAS [69, 17]. For
 350 instance, to compute

$$351 \quad x = \beta A^T y + z, \\ 352 \quad w = \alpha A x,$$

354 known as `dgemvt`, A is partitioned into block columns as

$$355 \quad A = [A_1 \ A_2 \ \cdots \ A_k],$$

356 where each A_i has b columns, and is sized such that it fits into cache. Correspondingly,
 357 x and z are partitioned as

$$358 \quad x = \begin{bmatrix} x_1 \\ \vdots \\ x_k \end{bmatrix}, \quad z = \begin{bmatrix} z_1 \\ \vdots \\ z_k \end{bmatrix}.$$

360 The `dgemvt` loops over the A_i blocks, performing two `dgemv` operations with each
 361 block as shown in [Algorithm 3](#). Keeping each A_i in cache for the second `dgemv` cuts
 362 main memory traffic roughly in half, thereby increasing the potential performance.
 363 With some algebraic manipulation, the two products $y_i = \tau_i A^T v_i$ and $x_i = \pi_i A u_i$
 364 from the bidiagonalization panel can be computed together using this `dgemvt`. Tests
 365 that Howell et al. ran showed a $1.2\text{--}1.3\times$ speedup over the existing LAPACK im-
 366 plementation for the bidiagonalization. Van Zee et al. [114] further analyzed these
 367 operations and fused them at the register level, reusing data in registers to also avoid
 368 unnecessary accesses to cache memory, showing potential further speedups. So far,
 369 these results have been for single-threaded implementations, and the speedups do not
 370 carry over when using multithreaded BLAS. If optimized Level 2.5 BLAS become
 371 available for multicore processors, this may become a viable approach, but we don't
 372 pursue this further here.

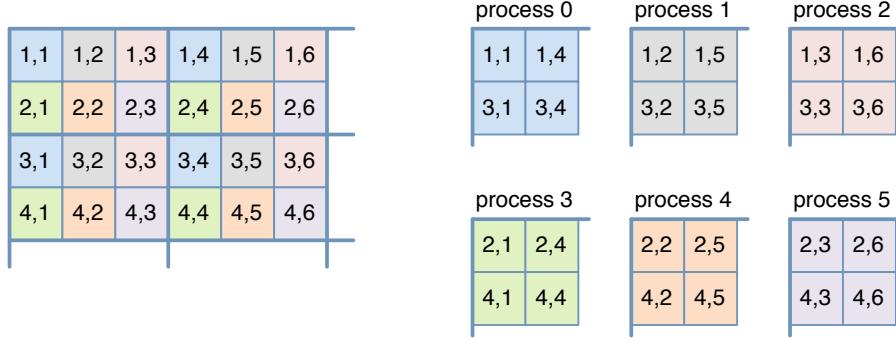
373 **6. ScaLAPACK Implementation.** To use a distributed-memory computer,
 374 the Scalable Linear Algebra Package (ScaLAPACK) [16] extends LAPACK by dis-
 375 tributing the matrices in a 2D block cyclic layout using the prescribed block size n_b
 376 and the pair of parameters (p, q) to define a p -by- q process grid, as illustrated in [Fig-](#)
 377 [ure 9](#). ScaLAPACK parallelizes the LAPACK subroutines using the parallel version of

Algorithm 3 Pseudocode for `dgemvt`

```

 $w = 0$ 
for  $i = 1 : k$ 
   $x_i = \beta A_i^T x_i + z_i$  // dgemv, loads  $A_i$  into cache
   $w = \alpha A_i x_i + w$  // dgemv, reuses  $A_i$  in cache
end

```



(a) Global view of matrix. Each square is an $n_b \times n_b$ block, colored by process that it is distributed to in (b).

(b) Local view of matrix. Each local submatrix is stored in column-major order.

FIG. 9. 2D block cyclic distribution of the matrix A using 2-by-3 processor grid.

378 BLAS (PBLAS) and the Basic Linear Algebra Communication Subprograms (BLACS)
 379 for the interprocess communication, implemented on top of the Message Passing In-
 380 terface (MPI) [93]. For instance, to bidiagonalize the input matrix for computing the
 381 SVD [24], `dgebrd` of LAPACK uses the Level 2 BLAS matrix-vector multiply (`dgemv`)
 382 to perform about half of its total flops. Now, to perform the matrix-vector multiply
 383 on a distributed-memory computer, in `pdgemv` of PBLAS, each process first gathers
 384 all the required nonlocal block rows of the input vector from other processes. After
 385 the completion of this initial interprocess communication, each process independently
 386 computes the matrix-vector multiplication with the local submatrix. Finally, each
 387 process computes the local part of the output vector by gathering and accumulating
 388 the partial results from the other processes in the same row of the process grid. Hence,
 389 ScaLAPACK follows the fork-join parallel programming paradigm and is designed for
 390 the weak parallel scalability of the algorithm. Since PBLAS performs most of its local
 391 computation using BLAS, ScaLAPACK can exploit a NUMA (non-uniform memory
 392 access) architecture using a threaded version of BLAS.

393 Figure 10 compares the performance of ScaLAPACK’s `pdgesvd` with the per-
 394 formance of LAPACK’s threaded `dgesvd` for computing the SVD on our 16-core
 395 shared-memory computer. While, from ScaLAPACK’s perspective, each MPI pro-
 396 cess has its own memory and explicit messages are passed between MPI processes, on
 397 a shared-memory computer the MPI implementation uses an efficient shared-memory
 398 communication layer to copy data. See section 11 for ScaLAPACK’s performance on
 399 a distributed memory computer. The performance of `pdgesvd` was obtained using the
 400 tester included in ScaLAPACK version 2.0.2, which was linked with ScaLAPACK and
 401 the sequential LAPACK/BLAS of Intel MKL. We tested the performance of `pdgesvd`

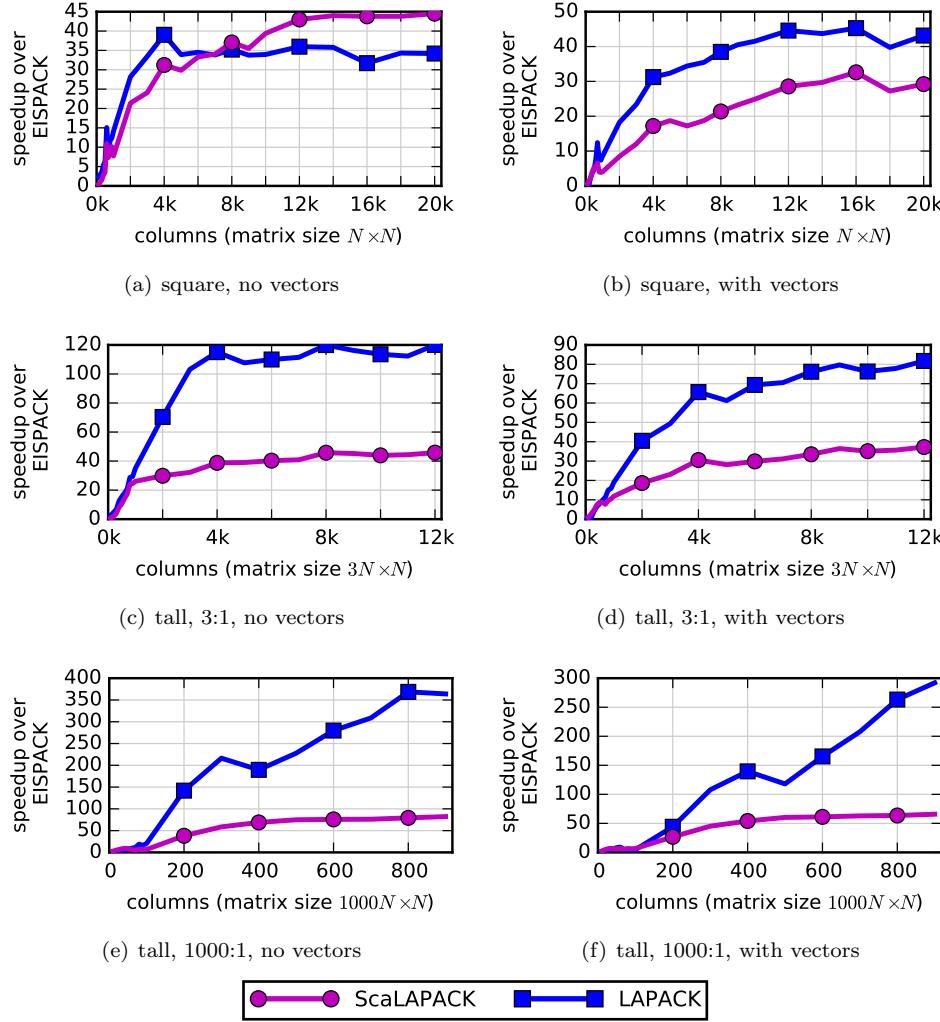
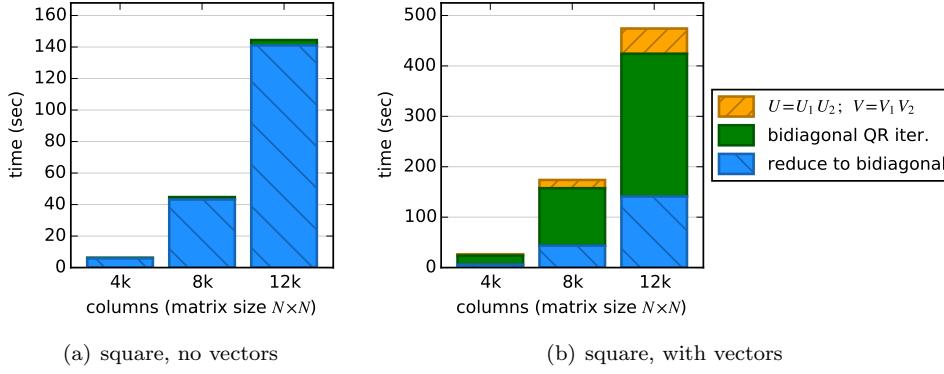


FIG. 10. Comparison of ScaLAPACK to LAPACK.

402 using 8-by-2, 4-by-4, and 2-by-8 processor grids and block sizes of 32, 64, and 128. The
 403 figure shows the optimal performance among these parameter configurations. We see
 404 that the performance of `pdgesvd` was often lower than the performance of LAPACK's
 405 `dgesvd`. This is mainly because several optimizations have not been implemented in
 406 `pdgesvd`. For instance, for a tall-skinny matrix ($m \gg n$), `dgesvd` computes the QR
 407 factorization of the input matrix A , followed by SVD of the resulting upper-triangular
 408 matrix, as described in subsection 5.4. For a tall-skinny matrix A , this greatly reduces
 409 the number of required floating point operations, compared to that of `pdgesvd`, which
 410 directly computes the SVD of the input matrix. As a result, for computing the SVD
 411 of a tall-skinny matrix, `pdgesvd` was slower than `dgesvd`.

412 After the bidiagonalization of the input matrix A , `pdgesvd` computes the SVD
 413 of the bidiagonal matrix using `dbdsqr` of LAPACK. If only the singular values are
 414 requested, `pdgesvd` typically spends an insignificant amount of time in `dbdsqr`. How-

FIG. 11. *Profile of ScaLAPACK reference implementation with $n_b = 32$ and $(p, q) = (4, 4)$.*

ever, if the singular vectors are needed, our performance profile in Figure 11 using the reference implementation of `pdgesvd` revealed that the execution time can be dominated by the time to compute the singular vectors of the bidiagonal matrix. The reason is that `pdgesvd` has all the MPI processes in the same column or row of the processor grid redundantly compute the left or right singular vectors, respectively, of the bidiagonal matrix, that are distributed to the process group. Compared with `pdgesvd`, LAPACK’s `dgesvd` obtained higher performance by using `dbdsqr` with multithreaded BLAS. The reference implementation of `pdgesvd` obtained about the same performance as that of MKL’s `pdgesvd` when linked to MKL BLAS and LAPACK.

Finally, ScaLAPACK supports only the QR iteration algorithm for computing the SVD of the bidiagonal matrix, using LAPACK’s `dbdsqr`, while as shown in section 7, the divide and conquer process in LAPACK’s `dbdsdc` may be faster than `dbdsqr`.

7. Singular Vectors from the Divide and Conquer Process. For solving the bidiagonal SVD subproblem, QR iteration and the related qd algorithms may take as much as 80% of the total time when computing singular vectors of a dense matrix [56]. Gu and Eisenstat introduced the bidiagonal divide and conquer (D&C) [57, 59] algorithm, which may be an order of magnitude faster on some machines [56]. The development of D&C was based on prior work focusing on computing eigenvalues and singular values [4, 25, 52, 58, 74].

The divide and conquer process includes a matrix partitioning step that introduces two large submatrices. The splitting can either occur with “the middle” row [4, 56] or column [57]:

$$B = \begin{bmatrix} B_1 & 0 \\ \alpha_k e_k & \beta_k e_1 \\ 0 & B_2 \end{bmatrix} \quad \text{or} \quad B = \begin{bmatrix} B_1 & \alpha_k e_k & 0 \\ 0 & \beta_k e_1 & B_2 \end{bmatrix}.$$

Note that after the partitioning, B_1 might not be square, even though B was. The fix is to append a zero row or column [57] to obtain the desired shape.

In either row or column case, the process continues recursively to obtain the SVD of B_1 and B_2 , which can be used to decompose B as:

$$B = Q_r M_r W_r \quad \text{or} \quad B = Q_c M_c W_c,$$

with orthogonal matrices Q_r , W_r , Q_c , and W_c . M_r and M_c have a special structure:

444 only the diagonal and either a single row or column is non-zero, respectively:

$$445 \quad M_r = \begin{bmatrix} z_1 & z_2 & \dots & z_n \\ & d_2 & & \\ & & \ddots & \\ & & & d_n \end{bmatrix} \quad \text{or} \quad M_c = \begin{bmatrix} z_1 & & & \\ z_2 & d_2 & & \\ \vdots & & \ddots & \\ z_n & & & d_n \end{bmatrix}.$$

446 Trivially, because matrices Q_r , W_r , Q_c , and W_c are orthogonal, B and M_r or M_c
447 share singular values σ_i . A number of theorems and lemmas [74, 57] lead to a fast
448 and numerically (relatively) stable procedure for computing the SVD of M_r or M_c as
449 $U_m \Sigma_m V_m^T$. The interlacing property sets the bounds and ranges for σ_i :

$$450 \quad 0 \equiv d_1 < \sigma_1 < d_2 < \dots < d_n < \sigma_n < d_n + \|z\|_2$$

451 and the secular equation:

$$452 \quad f(\sigma) = 1 + \sum_{k=1}^n \frac{z_k^2}{d_k^2 - \sigma^2} = 0$$

453 is used for computing the values σ_i with a specifically crafted root finder that ac-
454 counts for floating-point vagaries of past and modern computing systems [80]. The
455 corresponding formulas for the left singular vectors U_m :

$$456 \quad (5) \quad u_i = \left[\frac{z_1}{d_1^2 - \sigma_i^2}, \dots, \frac{z_n}{d_n^2 - \sigma_i^2} \right]^T \Bigg/ \sqrt{\sum_{k=1}^n \frac{z_k^2}{(d_k^2 - \sigma_i^2)^2}},$$

457 and the right singular vectors V_m :

$$458 \quad (6) \quad v_i = \left[-1, \frac{d_2 z_2}{d_2^2 - \sigma_i^2}, \dots, \frac{d_n z_n}{d_n^2 - \sigma_i^2} \right]^T \Bigg/ \sqrt{1 + \sum_{k=2}^n \frac{(d_k z_k)^2}{(d_k^2 - \sigma_i^2)^2}}$$

459 indicate that there could be accuracy problems for the components of either set of
460 vectors, even though the computed singular values $\hat{\sigma}$ are a good approximation of the
461 exact singular values σ , because the ratios $z_k/(d_k^2 - \sigma_i^2)$ in (5) and (6) can be inaccurate.
462 The trick is not to use the same M_r or M_c matrices that were used to compute the
463 approximate singular values $\hat{\sigma}_i$, but instead to construct new \hat{M}_r or \hat{M}_c based on $\hat{\sigma}_i$
464 that improve the accuracy of expressions in Equations (5) and (6) [80, 59]. This can
465 dramatically diminish the departure from orthogonality for both sets of vectors.

466 After computing U_m and V_m , the SVD of B is computed by multiplying $Q_r U_m$
467 and $V_m^T W_r$. This is done for each B matrix in the recursion tree, from the leaf nodes
468 to the root. Most of the cost of D&C is in these matrix multiplications, which are
469 Level 3 BLAS. In particular, most of the cost is at the higher levels of the recursion
470 tree, near the root node, as the matrices get larger.

471 Li et al. [81] recently showed that D&C internally generates matrices with struc-
472 ture that can be exploited. The matrices U_m and V_m , that are the singular vectors
473 of M , have low-rank off-diagonal blocks that can be efficiently compressed with hi-
474 erarchically semiseparable (HSS) matrices. Using HSS improves the speed of matrix
475 multiplies, reducing the cost of the bidiagonal D&C step from $O(n^3)$ to $O(n^2r)$, where
476 r depends on the matrix but usually $r \ll n$ for large n . Li et al. showed over 3×
477 improvement compared to Intel MKL for the bidiagonal D&C step on large matrices.

Algorithm 4 Overview of SVD algorithm using divide and conquer (`dgesdd`) for $m \geq n$. Generating explicit U_2 and V_2 occurs during D&C. \dagger Marked lines are required only when computing singular vectors.

Description	LAPACK Routine	Cost	Cost for $m \gg n$
if $m \gg n$ then			
$\hat{A} = A$ (QR factorization)	<code>dgeqrf</code>		$2mn^2 - \frac{2}{3}n^3$
$\hat{A} = R$			
else			
$\hat{A} = A$			
end			
$U_1 B V_1^T = \hat{A}$ (bidiagonalization)	<code>dgebrd</code>	$4mn^2 - \frac{4}{3}n^3$	$\frac{8}{3}n^3$
$U_2 \Sigma V_2^T = B$ (D&C)	<code>dbdsdc</code>	$O(n^2)$	$O(n^2)$
generate explicit U_2	" " \dagger	$\frac{4}{3}n^3$	$\frac{4}{3}n^3$
generate explicit V_2	" " \dagger	$\frac{4}{3}n^3$	$\frac{4}{3}n^3$
$U = U_1 U_2$	<code>dormbr</code> \dagger	$4mn^2 - 2n^3$	$2n^3$
$V = V_1 V_2$	<code>dormbr</code> \dagger	$2n^3$	$2n^3$
if $m \gg n$ then			
generate explicit Q	<code>dorgqr</code> \dagger		$2mn^2 - \frac{2}{3}n^3$
$U = QU$	<code>dgemm</code> \dagger		$2mn^2$
end			
Total cost (with vectors \dagger)		$8mn^2 + \frac{4}{3}n^3$	$6mn^2 + 8n^3$
Total cost (no vectors)		$4mn^2 - \frac{4}{3}n^3$	$2mn^2 + 2n^3$

478 D&C restructures the SVD algorithm somewhat, as shown in [Algorithm 4](#), com-
479 pared with the QR iteration version in [Algorithm 2](#). D&C directly computes the SVD
480 of the bidiagonal matrix $B = U_2 \Sigma V_2$, and then multiplies $U = U_1 U_2$ and $V = V_1 V_2$
481 afterwards (using `dormbr`), while with QR iteration, LAPACK first generates U_1 and
482 V_1 (using `dorgbr`), then accumulates U_2 and V_2 onto U_1 and V_1 during QR iteration.
483 The profile in [Figure 12](#) shows this difference in the bidiagonal QR iteration (green,
484 no hatching) vs. D&C steps (green, + hatching); and the generate U_1 , V_1 (orange, \|\|
485 hatching) vs. $U = U_1 U_2$, $V = V_1 V_2$ (orange, // hatching) steps. The main advantage
486 of the divide and conquer approach is that it saves nearly half the flops compared to
487 QR iteration when computing singular vectors. For a square matrix, D&C is $\approx 9n^3$
488 flops, compared to $\approx 17n^3$ for QR iteration ([Figure 5](#)). We can observe this as a
489 reduction in time for the steps mentioned above in [Figure 12](#).

490 [Figure 13](#) shows the relative speedup over EISPACK when using a modern mul-
491 ticom system, for both the divide and conquer (D&C) and QR iteration algorithms.
492 We see that for square and tall 3:1 matrices, D&C is consistently faster than QR
493 iteration. Because of the initial QR factorization (described in [subsection 5.4](#)) the
494 advantage decreases as m grows relative to n , so that for a very tall matrix, both
495 methods are nearly the same time, as seen by the 1000:1 case in [Figure 13\(c\)](#). It may
496 be safely assumed that D&C is superior to the QR iteration algorithm for most sce-
497 narios, and the worst case is when both perform at the same speed. When computing
498 only singular values, not singular vectors, LAPACK always uses QR iteration, since
499 in that case both bidiagonal QR iteration and D&C are $O(n^2)$, while the overall time
500 will be dominated by the $O(n^3)$ reduction to bidiagonal.

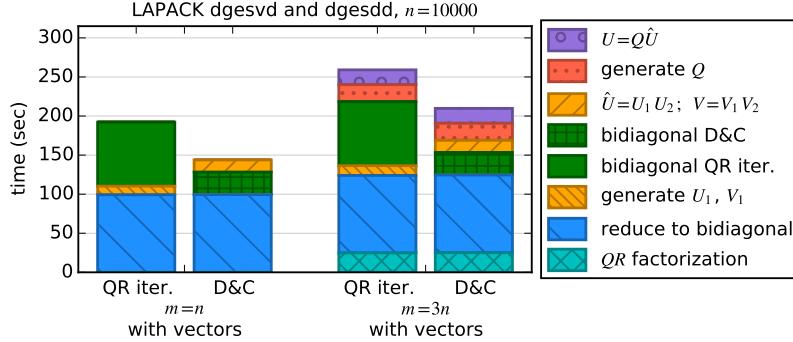


FIG. 12. Profile comparing LAPACK QR iteration (*dgesvd*) and divide and conquer (*dgesdd*) algorithms. For QR iteration, it generates U_1 and V_1 , then updates those with U_2 and V_2 during QR iteration. Divide and conquer generates U_2 and V_2 , then multiplies $\hat{U} = U_1 U_2$ and $V = V_1 V_2$ afterwards.

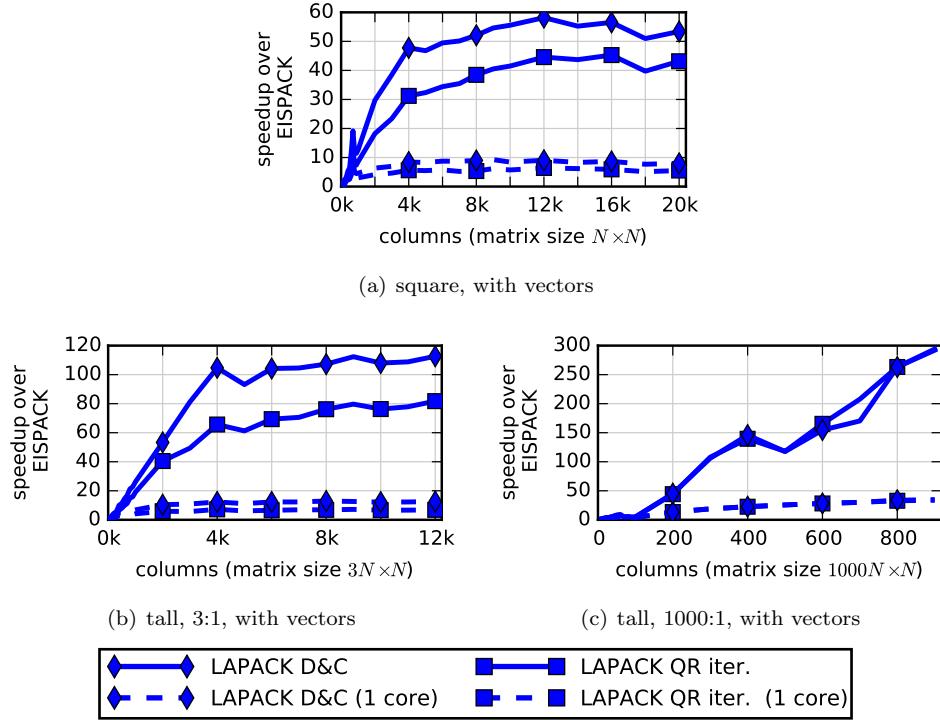


FIG. 13. Comparison of LAPACK divide and conquer (D&C) to QR iteration. Solid lines represent 16-core runs; dashed lines represent single core runs.

501 **8. Bisection and Inverse Iteration.** LAPACK 3.6.0 introduced a bisection
 502 method (*dgesvdx*) to compute all or a subset of the singular values and vectors [86].
 503 Similar to QR iteration (*dgesvd*) and divide and conquer (*dgesdd*), it first reduces
 504 the matrix A to bidiagonal form B . Then it computes the singular values of B
 505 based on bisection and the corresponding singular vectors by inverse iteration, using
 506 *dbdsvdx*. For computing the SVD of B , it converts the bidiagonal matrix B to the

507 Golub-Kahan [53] symmetric tridiagonal matrix T of dimension $2n$,

508 (7)
$$T = \text{tridiag} \begin{pmatrix} b_{1,1} & b_{1,2} & b_{2,2} & b_{2,3} & \dots & b_{n-1,n} & b_{n,n} \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ b_{1,1} & b_{1,2} & b_{2,2} & b_{2,3} & \dots & b_{n-1,n} & b_{n,n} \end{pmatrix},$$

509 whose eigenpairs are $(\pm\sigma_j, z_j)$, where σ_j is the j -th singular value of B . Elements of
 510 u_j and v_j , the corresponding left and right singular vectors of B , are interleaved in the
 511 eigenvector as $z_j = [v_{1,j}, -u_{1,j}, v_{2,j}, -u_{2,j}, \dots, v_{n,j}, -u_{n,j}] / \sqrt{2}$. Instead of developing
 512 new subroutines, `dbdsvdx` relies on the subroutines `dstebz` and `dstein` that compute
 513 the eigenvalues and eigenvectors, respectively, of the symmetric tridiagonal matrix
 514 based on bisection and inverse iteration.

515 The bisection algorithm implemented in `dstebz` uses Sylvester's inertia theorem
 516 to compute the number of eigenvalues within a certain interval. In particular, the
 517 algorithm relies on the LDL^T factorization of the matrix T , where L is a lower-
 518 triangular matrix with unit diagonal and D is a diagonal matrix. For the symmetric
 519 tridiagonal matrix T , the diagonal matrix D can be computed with $O(n)$ flops based
 520 on the simple recurrence formula,

521
$$d_{i,i} = (t_{i,i} - s) - \frac{t_{i-1,i}^2}{d_{i-1,i-1}}.$$

522 Given the LDL^T factorization of the matrix $T - sI$ for a certain shift value s , the
 523 number of negative elements of D is equal to the number of eigenvalues of T smaller
 524 than s . In other words, given the LDL^T factorizations of two shifted matrices, $T - s_1 I$
 525 and $T - s_2 I$, with $s_1 < s_2$, if there are n_1 and n_2 negative entries in their respective
 526 diagonal matrices, then there are $n_2 - n_1$ eigenvalues in the interval $(s_1, s_2]$. In
 527 addition, for the tridiagonal matrix, it can be shown that the LDL^T factorization
 528 without pivoting can be reliably used for counting the number of eigenvalues [34, 75].
 529 Based on these observations, `dstebz` keeps bisecting the initial interval containing
 530 all the desired eigenvalues until it finds a small enough interval for each eigenvalue
 531 such that the computed eigenvalue has the desired accuracy. Each bisection improves
 532 the accuracy of the eigenvalue by one bit, hence the iteration converges linearly. An
 533 advantage of bisection is that it can be naturally adapted to compute a subset of
 534 eigenvalues, which was one of the motivations for introducing `dgesvdx` [86].

535 Given the eigenvalues computed by `dstebz`, `dstein` computes the correspond-
 536 ing eigenvectors based on inverse iteration. Namely, for each computed eigenvalue
 537 λ , it first computes the LU factorization of the shifted matrix $A - \lambda I$ with partial
 538 pivoting. Then the corresponding eigenvector of λ is computed by inverse iteration,
 539 with a starting vector whose entries are random numbers uniformly distributed in
 540 the interval $(-1, 1)$. Given an accurate eigenvalue approximation, inverse iteration
 541 converges quickly [72] (e.g., `dstein` sets the maximum number of iterations to be
 542 five). However, when the eigenvalues are close to each other, inverse iteration may
 543 fail to generate orthogonal eigenvectors. To recover the orthogonality among such
 544 vectors, `dstein` reorthogonalizes the vectors based on the modified Gram-Schmidt
 545 procedure. Unfortunately, when the computed eigenvectors are nearly dependent, the
 546 eigenvectors may not be accurate after the reorthogonalization [33]. In addition, if
 547 many of the eigenvalues are close to each other, this reorthogonalization cost could
 548 become significant with $O(k^2n)$ flops for computing k eigenvectors in the worst case.
 549 As a result, in our experiments shown in Figure 14, we saw that when computing
 550 all the singular values and vectors, bisection can be significantly slower than other

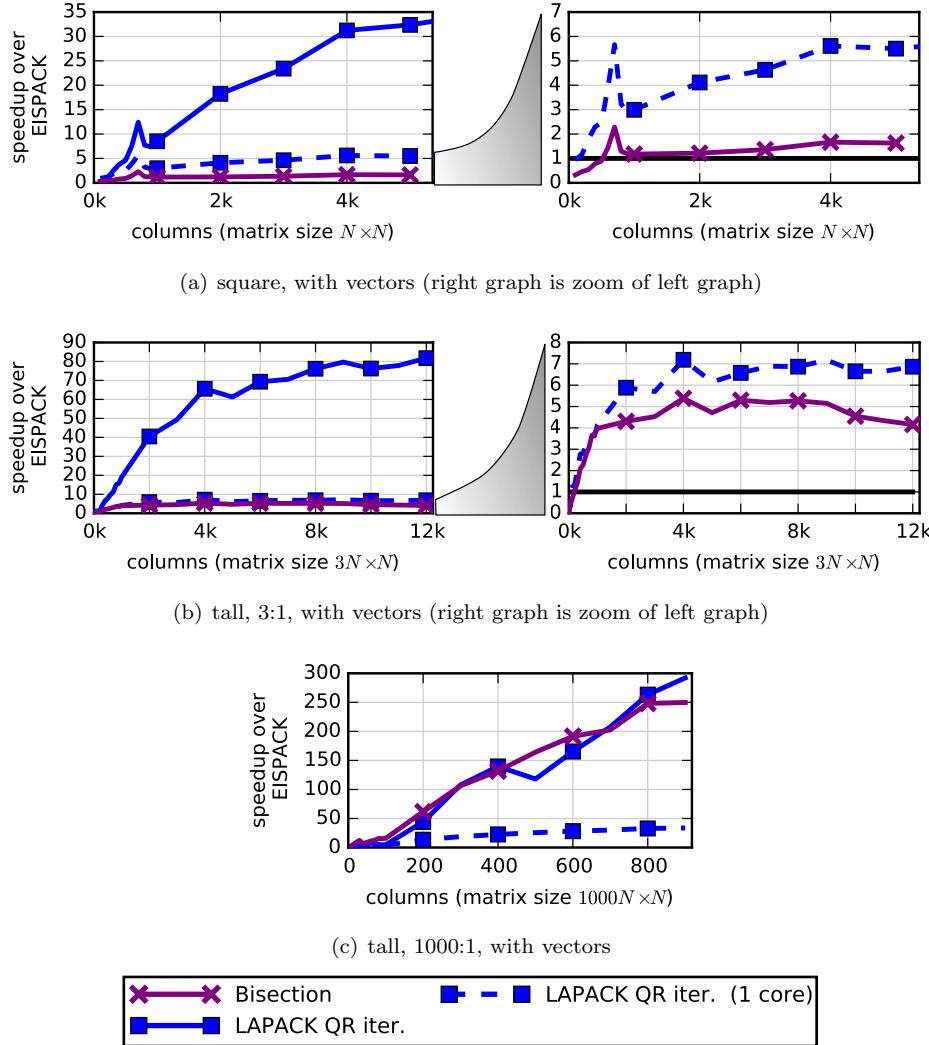


FIG. 14. Comparison of bisection with QR iteration. Solid lines represent 16-core runs; dashed lines represent single core runs.

551 methods. Even with 16 threads available, it is slower than the single-threaded QR
 552 iteration (dashed line). Bisection and inverse iteration are embarrassingly parallel—
 553 each eigenvalue and eigenvector may be computed independently—however, LAPACK
 554 does not currently include such explicit parallelization, instead primarily relying on
 555 parallelism within the BLAS, which is not advantageous in this case. On the other
 556 hand, as seen in Figure 15, when only a subset of k singular values and vectors are
 557 computed, we observed that bisection and inverse iteration (non-hatched bars) can be
 558 up to $2.4\times$ faster than divide and conquer (D&C, black bar and dashed line), which
 559 must compute all the singular values and vectors. Depending on the matrix type,
 560 when computing $k = 400$ or $k = 600$ vectors out of $n = 3000$, it becomes faster to
 561 simply compute all the vectors using D&C. The exception here is the cluster, with
 562 one singular $\sigma_1 = 1$ and all other $\sigma_i = 1/\kappa$. In that case, computing any $k > 1$ vectors

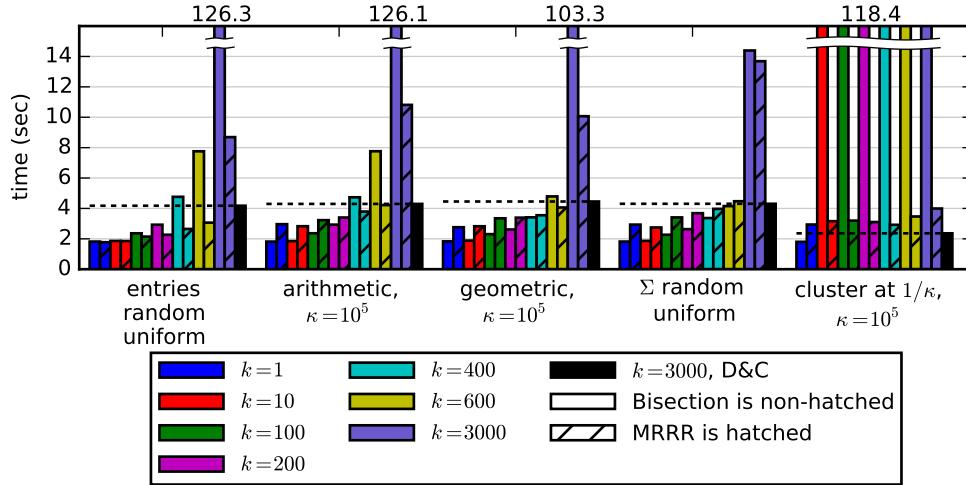


FIG. 15. Results for computing k singular vectors of $n \times n$ matrix, $n = 3000$. Dashed lines shows D&C performance for computing all vectors. Number on top shows time for bars that exceed graph's height.

563 was as slow as computing all vectors with bisection. See [section 2](#) for a description of
 564 the matrices.

565 **9. Multiple Relatively Robust Representations (MRRR).** MRRR [35,
 566 36] was developed to improve both the performance and accuracy of inverse iteration.
 567 Analysis has shown that MRRR can compute the numerically orthogonal eigenvectors
 568 of a symmetric tridiagonal matrix in $O(n^2)$ flops. At the time of preparing this
 569 paper, there was no publicly available software package that implements MRRR for
 570 computing the SVD of a general matrix, but there were at least two software packages
 571 that compute the eigenvalues and eigenvectors of a symmetric tridiagonal matrix using
 572 MRRR: `dstemr` of LAPACK [38], and `dstexr` due to Willems and Lang [118], which
 573 is tailored toward the tridiagonal matrix with zeros on the diagonal, as used in (7) for
 574 the SVD. For our experiments, we replaced the symmetric tridiagonal solver (`dstevx`)
 575 used in `dgesvdx` with `dstexr`. Performance with `dstemr` was generally similar but
 576 somewhat slower.

577 One of the main drawbacks of inverse iteration is that, for the eigenvalues with
 578 small relative gaps, the computed eigenvectors may not be orthogonal to each other.
 579 Hence, reorthogonalization is needed. This increases the computational cost and
 580 potentially leads to loss of accuracy in the computed eigenvectors. To address these
 581 issues, MRRR combines several techniques.

582 First, though the eigenvectors are invariant under a diagonal shift, we can increase
 583 their relative gaps by diagonally shifting the matrix. For instance, let us define the
 584 relative gap between two eigenvalues λ_i and λ_j to be $\frac{|\lambda_i - \lambda_j|}{\max(|\lambda_i|, |\lambda_j|)}$. Then, we can
 585 increase their relative gap by a factor of $\frac{|\lambda|}{|\lambda - \tau|}$ when we diagonally shift the matrix
 586 using a shift τ that is close to λ .

587 Hence, before applying inverse iteration, MRRR recursively refines the approxi-
 588 mation to the eigenvalues and applies appropriate diagonal shifts to a cluster of

589 eigenvalues such that it can guarantee large enough relative gaps between all the
 590 eigenvalues of T to maintain the orthogonality among the computed eigenvectors
 591 without reorthogonalization. For instance, given two approximate eigenvalues λ_i and
 592 λ_j , inverse iteration is used to compute their respective eigenvectors v_i and v_j with
 593 small residual norms, i.e.,

$$594 \quad |Tv_k - \lambda v_k| = O(n\epsilon |T|) \text{ for } k = i \text{ and } j.$$

595 Then, according to [37, 38], a realistic bound on their orthogonality error is given by

$$596 \quad |v_i^T v_j| = O\left(\frac{n\epsilon(|\lambda_i| + |\lambda_j|)}{|\lambda_i - \lambda_j|}\right).$$

597 Therefore, if the gap $|\lambda_i - \lambda_j|$ is of the same order as the eigenvalues, their eigenvectors
 598 are numerically orthogonal, i.e., $|v_i^T v_j| = O(n\epsilon)$.

599 There are several parameters that can be tuned to improve the performance [38],
 600 including the accuracy of the eigenvalue approximation computed at each step and
 601 the choice of the algorithm for computing the approximation (e.g., bisection, QR
 602 iteration, or Rayleigh quotient correction).

603 Second, while computing the eigenvalues (e.g., applying the diagonal shift), a
 604 small relative roundoff error in the entry of the tridiagonal matrix could result in
 605 a large relative error in the computed eigenvalues, especially in those with small
 606 magnitudes. To preserve the relatively high accuracy of the computed eigenvalues,
 607 MRRR stores the intermediate matrices in particular representations, referred to as
 608 the Multiple Relatively Robust Representation (MRRR) of the matrices. For instance,
 609 it has been shown that the LDL^T representation of the tridiagonal matrix T , without
 610 pivoting, is relatively robust, even with the presence of the element growth [31].
 611 Hence, MRRR stores the sequence of intermediate matrices with different shifts in
 612 their LDL^T forms.

613 Third, for an eigenvalue with a small relative gap, the cost of inverse iteration
 614 may be high, requiring a few iterations to obtain the eigenvector with a small relative
 615 residual norm. Fortunately, there is at least one starting vector with which inverse
 616 iteration converges in one iteration. For example, when the i -th column of $(T - \lambda I)^{-1}$
 617 has the largest column norm, then with the canonical vector e_i as the starting vector,
 618 one step of inverse iteration computes the approximate eigenvector x such that

$$619 \quad |Tx - \lambda x| \leq \sqrt{n} |\lambda - \bar{\lambda}|,$$

620 where $\bar{\lambda}$ is the exact eigenvalue [72]. Hence, if the eigenvalue is computed to a high
 621 relative accuracy,

$$622 \quad |\lambda - \bar{\lambda}| = O(\epsilon |\bar{\lambda}|),$$

623 (e.g., using bisection with $O(n)$ flops), then the computed eigenpair obtains a small
 624 relative residual norm,

$$625 \quad |Tx - \lambda x| = O(n\epsilon |\bar{\lambda}|).$$

626 There is an algorithm to find the column of $(T - \lambda I)^{-1}$ with largest norm with
 627 $O(n)$ flops [98]. In addition, if a twisted factorization is used to find the starting
 628 vector, then it can be shown that the computed eigenpairs have small residual norm
 629 with respect to the original matrix T [36, 37]. The twisted factorization must be
 630 carefully computed for T with a zero diagonal because the leading dimension of an
 631 odd dimension is singular. To enhance the numerical stability, `dstexr` computes a
 632 block variant of the factorization [118].

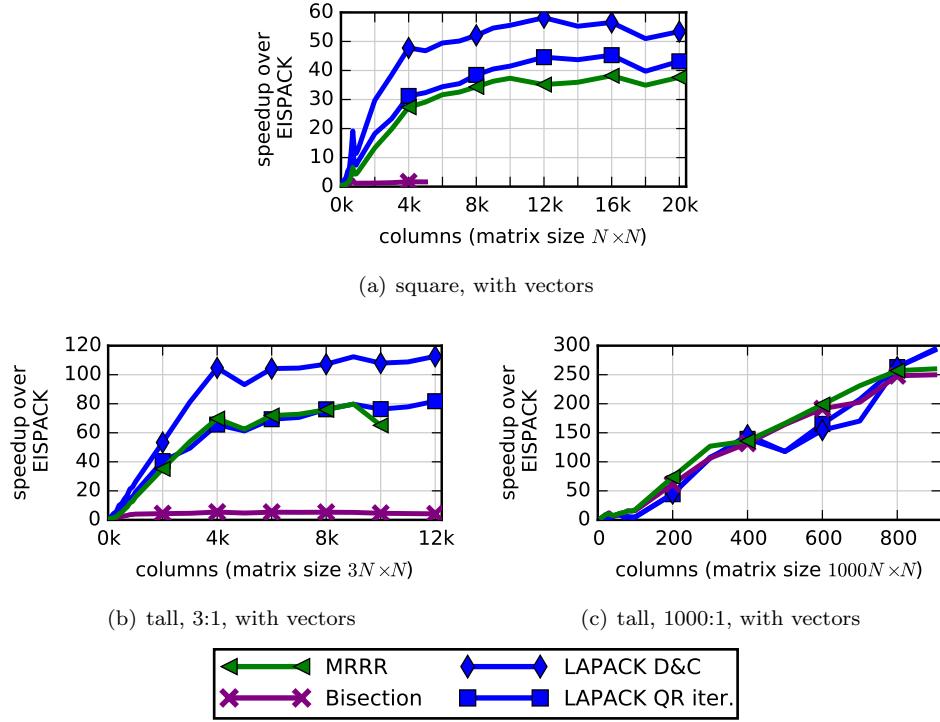


FIG. 16. Comparison of MRRR with QR iteration, divide and conquer, and bisection.

633 As can be seen in Figure 16, by avoiding the reorthogonalization, MRRR can
 634 significantly improve the performance of inverse iteration, making MRRR comparable
 635 to QR iteration. However, divide and conquer is often faster.

636 Especially for large matrices, we noticed numerical issues where the backward er-
 637 rror $\|A - U\Sigma V^T\| / (\min(m, n) \|A\|)$ was large, e.g., 10^{-4} instead of 10^{-16} as expected.
 638 Further tests in section 13 show that, even when the above error is acceptable, MRRR
 639 has poor relative error for the singular values. Marques and Vasconcelos [86] also ob-
 640 served numerical issues with the existing MRRR implementation.

641 When only a subset of k singular vectors are computed, we observe in Figure 15
 642 that inverse iteration can be up to $1.6\times$ faster than MRRR for a small number vectors
 643 ($k = 1$ or 10). For a larger subset of $k = 600$ vectors, MRRR can be up to $1.8\times$
 644 faster than bisection, but in this case, only for the random entries matrix is MRRR
 645 significantly faster ($1.3\times$) than computing all the singular vectors with divide and
 646 conquer. The exception is the cluster matrix, where for $k > 1$, MRRR is $30\times$ faster
 647 than bisection, but always slower than using divide and conquer.

648 **10. MAGMA Implementation for Accelerator Architectures.** Acceler-
 649 tor such as GPUs and the Intel Xeon Phi provide a high degree of parallelism and
 650 a larger memory bandwidth than traditional multicore CPUs. The MAGMA library
 651 was developed to address this new architecture, and accelerates most phases of the
 652 SVD algorithm: reduction to bidiagonal, bidiagonal D&C, and computation of sin-
 653 gular vectors. For tall-skinny matrices, it also accelerates the initial QR factorization
 654 and generating Q .

655 The most prominent place to start is an accelerated version of the bidiagonal

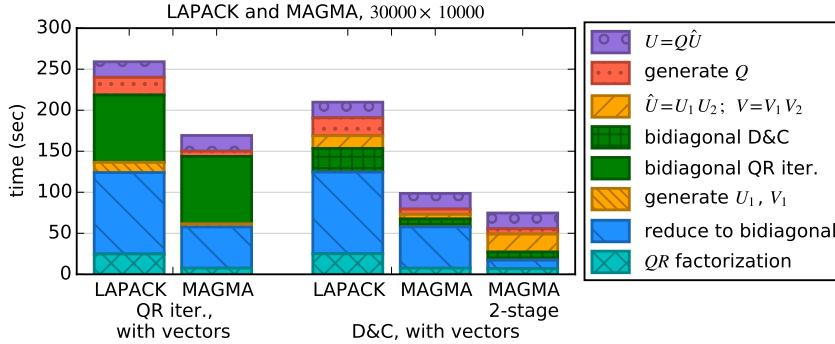


FIG. 17. *Profile comparing LAPACK and MAGMA. Most phases are accelerated using the GPU, except the bidiagonal QR iteration and multiplying $U = Q\hat{U}$. MAGMA 2-stage is described in section 11.*

656 reduction [109]. We have seen in Figures 7 and 12 that this phase (blue tier with \\ 657 hatching) takes from 50% to 70% of the time for a square matrix when computing singular 658 vectors, and 98% of the time when computing only singular values (no vectors). 659 As described in section 5, the bidiagonal reduction has half its flops in Level 2 BLAS 660 and half in Level 3 BLAS. Accelerators are known for achieving very high performance 661 on compute-intensive, Level 3 BLAS operations. On an NVIDIA K40c GPU, cuBLAS 662 achieves 1245 Gflop/s with `dgemm`, compared with 315 Gflop/s using Intel MKL on 663 the multicore CPU. Due to the accelerator’s large memory bandwidth, the memory- 664 bound Level 2 BLAS operations are also significantly faster, achieving 45 Gflop/s with 665 cuBLAS `dgemv`, compared with 14 Gflop/s on the multicore CPU. Therefore, both 666 the trailing matrix-vector product (`dgemv`) and the trailing matrix update (`dgemm`) are 667 performed on the accelerator. The small panel operations—constructing Householder 668 reflectors—are performed on the CPU, which is better at serial operations with more 669 control flow. This incurs CPU-to-GPU communication of a couple of vectors for each 670 `dgemv` operation during the panel. Due to dependencies, the trailing matrix update 671 cannot be overlapped with the next panel, as would occur in a one-sided QR factoriza- 672 tion. Using the accelerator improves the speed of the bidiagonal reduction by about 673 a factor of 2, as shown by the profile in Figure 17 (blue tier with \\ hatching) and 674 by the square, no vectors case in Figure 18(a), which is dominated by the bidiagonal 675 reduction.

676 For the bidiagonal SVD, because D&C is faster than QR iteration, MAGMA 677 will inherently achieve a better overall speedup using D&C. We further implement 678 an accelerated version of D&C [51]. Since most of the operations in D&C are in 679 multiplying $Q_r U_m$ and $V_m^T W_r$ to generate singular vectors, these Level 3 BLAS `dgemm` 680 operations are assigned to the accelerator. The solution of the secular equation to 681 find the singular values of M_c is left on the CPU, since it is a complex iterative 682 algorithm with limited parallelism, as is computing the singular vectors U_m and V_m 683 of M_c . These are parallelized on the CPU using OpenMP. MAGMA achieves about 684 a $3\times$ speedup for the D&C phase compared to LAPACK.

685 For a tall-skinny ($m \gg n$) matrix, we accelerate the initial QR factorization [110]. 686 This is a one-sided factorization, so it doesn’t have the extra dependencies imposed by 687 the two-sided reduction to bidiagonal form. Panel operations are within a simple block 688 column that doesn’t involve the trailing matrix. The panel factorization is performed

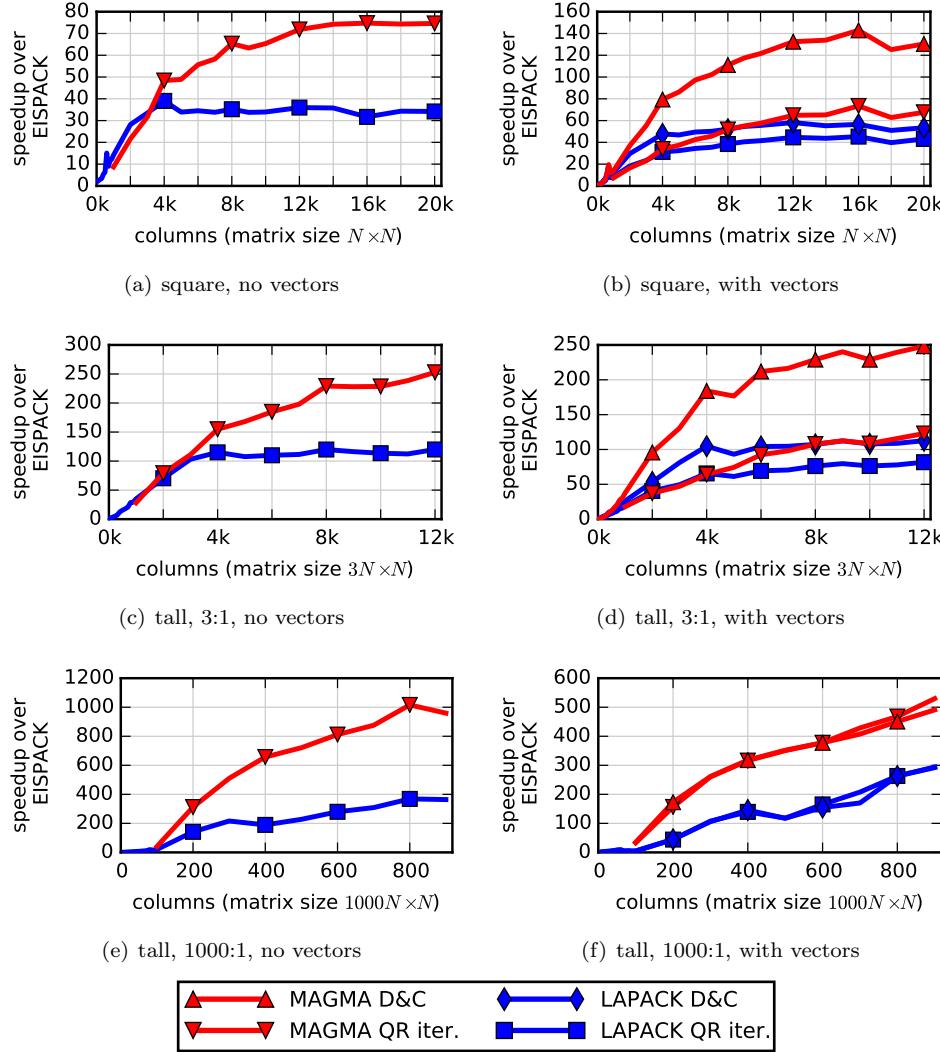


FIG. 18. Comparison of MAGMA with LAPACK.

689 on the CPU, while the trailing matrix update is performed on the accelerator. The
 690 accelerator updates the next panel first and sends it back to the CPU so the CPU can
 691 start factoring it while the accelerator proceeds with the rest of the trailing matrix
 692 update. This overlap allows the factorization to achieve a substantial portion of the
 693 peak `dgemm` speed, up to 970 Gflop/s with an NVIDIA K40c. The QR factorization
 694 phase was up to 3.6 \times faster than on the multicore CPU, as seen in Figure 17 (cyan
 695 tier with \times hatching).

696 There are three routines that are solely applying block Householder reflectors,
 697 which are implemented as a series of Level 3 BLAS matrix multiplies entirely on the
 698 accelerator: (1) for QR iteration, generating explicit U_1 and V_1 matrices (`dorgbr`)
 699 (2) for D&C, multiplying U_1U_2 and V_1V_2 (`dormbr`), and (3) for a tall-skinny matrix,
 700 generating an explicit Q matrix (`dorgqr`). These were up all up to 3.3 \times faster when

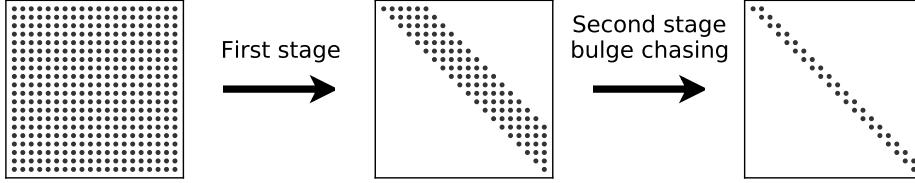


FIG. 19. *Two-stage technique for the reduction phase.*

701 using the accelerator than when using the multicore CPU. We see in Figure 17 that
 702 the time for all three of these phases is substantially reduced.

703 Overall, MAGMA achieves significant improvements using an accelerator for the
 704 SVD problem. Figure 18 shows that it is about $2\times$ faster than LAPACK in most
 705 cases. For the square, vectors case in Figure 18(b), MAGMA’s SVD using D&C is
 706 $2.5\times$ LAPACK’s D&C version, and $2\times$ MAGMA’s SVD using QR iteration, while
 707 MAGMA’s SVD using QR iteration is only $1.6\times$ LAPACK’s QR iteration version,
 708 due to both D&C being inherently faster and having an accelerated version of the
 709 D&C phase. In the tall 1000:1 case in Figure 18(e), MAGMA is $2.6\times$ faster, and
 710 for some sizes as much as $3.5\times$ faster, than LAPACK, and up to $1000\times$ faster than
 711 EISPACK, due to the accelerated QR factorization.

712 **11. Two-stage Reduction.** While all the preceding algorithmic and architec-
 713 tural improvements have greatly increased the speed of the SVD, all these one-stage
 714 methods remain limited by the memory-bound, Level 2 BLAS operations. To over-
 715 come the limitations of the one-stage approach, Großer and Lang [78, 55] introduced
 716 the two-stage bidiagonal reduction, which increases the use of compute-intensive
 717 Level 3 BLAS operations. The idea behind the two-stage algorithm is to split the
 718 original one-stage bidiagonal reduction into a compute-intensive phase (first stage)
 719 and a memory-bound phase (second or *bulge-chasing* stage), as represented in Fig-
 720 ure 19. The first stage reduces the original general dense matrix to a band form
 721 (either upper or lower), and the second stage reduces the band form to bidiagonal
 722 form (again, either upper or lower). The algorithm maps computational tasks to the
 723 strengths of the available hardware components, taking care of the data reuse. It also
 724 uses techniques to mix between dynamic and static scheduling to extract efficiency
 725 and performance. We implemented two-stage algorithms in the PLASMA library for
 726 multicore environments [82, 83, 62, 60], the DPLASMA library for distributed envi-
 727 ronments [19, 18], and the MAGMA library for accelerator architectures [51]. Similar
 728 two-stage reduction [61] and multi-stage successive band reduction (SBR) [13, 6] to
 729 tridiagonal have been used for the symmetric eigenvalue problem. A multi-stage ap-
 730 proach would also work for the bidiagonal reduction, and could be advantageous to
 731 achieve optimal communication costs at each stage. However, when computing sin-
 732 gular vectors, each stage adds cost to the back transformation, making a multi-stage
 733 approach less favorable.

734 **11.1. First Stage: Compute-Intensive and Efficient Kernels.** The first
 735 stage applies a sequence of blocked Householder transformations to reduce the general
 736 dense matrix to an upper (for $m \geq n$) band matrix. This stage uses compute-intensive
 737 matrix-multiply kernels that eliminate the memory-bound matrix-vector products
 738 from the one-stage panel factorization.

739 The first stage proceeds by computing a QR factorization of a block column to
 740 annihilate entries below the diagonal, and updating the trailing matrix, as shown in
 741 [Figure 20](#). It then computes an LQ factorization of a block row to annihilate entries
 742 right of the upper bandwidth, and updates the trailing matrix. It repeats factoring
 743 block columns and block rows, until the entire matrix is brought to band form. The
 744 width of the block columns and rows is the resulting matrix bandwidth, n_b .

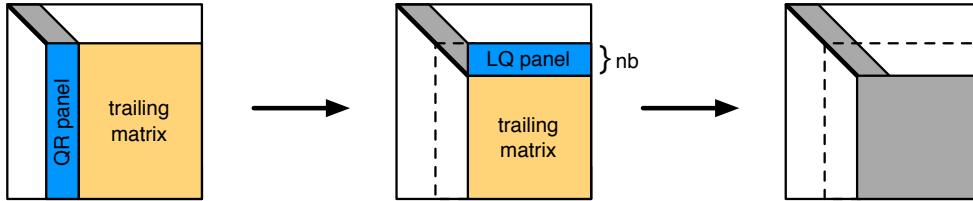


FIG. 20. *One panel of the first stage reduction to band form.*

745 The PLASMA and DPLASMA implementations use a tile algorithm [1] that
 746 makes it highly parallel. The matrix is split into tiles of size $n_b \times n_b$, where n_b is the
 747 matrix bandwidth. Data within each tile is stored contiguously in memory. A panel
 748 factorization is a series of QR or LQ factorizations done between pairs of tiles; once a
 749 pair of tiles has been factored, updates on the corresponding portions of the trailing
 750 matrix can start immediately, before the rest of the panel has finished factoring. This
 751 unlocks a large amount of parallelism very quickly. The algorithm then proceeds
 752 as a collection of interdependent tasks that operate on the tile data layout and are
 753 scheduled in an out-of-order fashion using either the OpenMP runtime for PLASMA
 754 or the powerful PaRSEC distributed runtime system for DPLASMA.

755 The MAGMA implementation uses a standard column-wise layout. It does the
 756 QR and LQ factorizations on the CPU, copies the block Householder reflectors to the
 757 accelerator, and updates the trailing matrix on the accelerator. Unlike in the one-
 758 sided factorizations, it cannot start the next panel until the trailing matrix update is
 759 finished due to data dependencies.

760 The first stage's cost is $\frac{8}{3}n^3$ operations in Level 3 BLAS. As shown in [60], the
 761 performance of this stage is comparable to the performance of the QR factorization
 762 and can reach a high percentage of the machine's peak.

763 **11.2. Second Stage: Cache-Friendly Computational Kernels.** The sec-
 764 ond stage reduces the band form to the final bidiagonal form using a bulge chasing
 765 technique. It involves $6n_b n^2$ operations, so it takes a small percentage of the total
 766 operations, which decreases with n . The operations are memory bound, but are fused
 767 together as Level 2.5 BLAS [69] for cache efficiency. We designed the algorithm to
 768 use fine-grained, memory-aware tasks in an out-of-order, data-flow task-scheduling
 769 technique that enhances data locality [60, 61].

770 The second stage proceeds in a series of sweeps, each sweep bringing one row to
 771 bidiagonal and chasing the created fill-in elements down to the bottom right of the
 772 matrix using successive orthogonal transformations. It uses three kernels. Kernel 1
 773 (yellow task $T_{1,1}$ in [Figure 21\(b\)](#)) applies a Householder reflector from the right (in-
 774 dicated by the down arrow) to eliminate a row right of the superdiagonal, which also
 775 creates a bulge of fill-in beneath the diagonal. It then applies a Householder reflector
 776 from the left (indicated by the right arrow) to eliminate the first column of the
 777 bulge below the diagonal, and applies the update to the first block column only. The

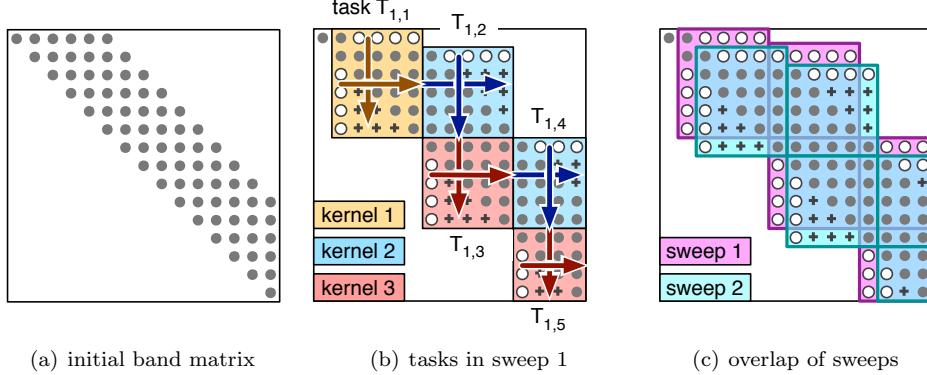


FIG. 21. *Bulge-chasing algorithm.* “o” indicates eliminated elements; “+” indicates fill. Arrows show application of Householder reflector on left (\rightarrow), which update a block row, and on right (\downarrow), which update a block column.

778 remainder of the bulge is not eliminated, but is instead left for subsequent sweeps to
 779 eliminate, as they would reintroduce the same nonzeros.

780 Kernel 2 (blue task $T_{1,2}$) continues to apply the left Householder reflector from
 781 kernel 1 (or kernel 3) to the next block column, creating a bulge above the upper
 782 bandwidth. It then applies a right Householder reflector to eliminate the first row of
 783 the bulge right of the upper bandwidth, updating only the first block row.

784 Kernel 3 (red task $T_{1,3}$) continues to apply the right Householder reflector from
 785 kernel 2, creating a bulge below the main diagonal. As in kernel 1, it then applies a
 786 left Householder reflector to eliminate the first column of the bulge below the diagonal
 787 and updates just the current block column. After kernel 3, kernel 2 is called again
 788 (blue task $T_{1,4}$) to continue application of the left Householder reflector in the next
 789 block column. A sweep consists of calling kernel 1 to bring a row to bidiagonal,
 790 followed by repeated calls to kernels 2 and 3 to eliminate the first column or row of
 791 the resulting bulges, until the bulges are chased off the bottom-right of the matrix.

792 For parallelism, once a sweep has finished the first kernel 3, a new sweep can start
 793 in parallel. This new sweep is shifted over one column and down one row, as shown in
 794 Figure 21(c). Before task i in sweep s , denoted as $T_{s,i}$, can start, it depends on task
 795 $T_{s-1, i+3}$ in the previous sweep being finished, to ensure that kernels do not update
 796 the same entries simultaneously. To maximize cache reuse, tasks are assigned to cores
 797 based on their data location. Ideally, the band matrix fits into the cores’ combined
 798 caches, and each sweep cycles through the cores as it progresses down the band.

799 **11.3. Singular Vectors Computation.** The singular vectors of A are com-
 800 puted from the orthogonal transformations used in the reduction to bidiagonal form
 801 and from the singular vectors of the bidiagonal form. Recall that for the classical
 802 one-stage approach, $A = U_1 B V_1^T$ and $B = U_2 \Sigma V_2^T$. After using D&C to obtain U_2
 803 and V_2 , we multiply $U = U_1 U_2$ and $V = V_1 V_2$, costing $2n^3$ each for U and V (if
 804 $m = n$).

805 In the case of the two-stage approach, the first stage reduces the original matrix
 806 A to a band matrix by applying a two-sided transformation to A such that $A =$
 807 $U_a A_{\text{band}} V_a^T$. Similarly, the second, bulge-chasing stage reduces the band matrix A_{band}
 808 to bidiagonal form by applying a two-sided transformation such that $A_{\text{band}} = U_b B V_b^T$.

809 As a consequence, the singular vectors must be multiplied according to:

810
$$U = U_a U_b U_2 \quad \text{and} \quad V = V_a V_b V_2.$$

812 Hence the two-stage approach introduces a nontrivial amount of extra computation—the application of U_b and V_b —when the singular vectors are needed. The total cost of 813 updating the singular vectors when using the two-stage technique is $2(1 + \frac{i_b}{n_b})n^3 + 2n^3$ 814 each for U and V , where n_b is the bandwidth and i_b is an internal blocking; usually 815 $i_b \leq n_b/4$. This extra cost compared with the one-stage approach reduces the potential 816 speedup, but as it is in Level 3 BLAS, it does not completely negate the large speedup 817 that we gain by the two-stage bidiagonal reduction.

819 **11.4. PLASMA Implementation for Multicore.** The experiments shown in 820 Figure 22 illustrate the superior efficiency of our two-stage SVD solver compared with 821 the optimized LAPACK version from Intel MKL. Figure 22(a) shows that the bidiagonal 822 reduction itself is $6\times$ faster than LAPACK, both using 16 cores, and $2.5\times$ faster 823 than the MAGMA one-stage version. The reason is that LAPACK and MAGMA 824 are bound by the Level 2 BLAS performance, while our two-stage algorithm relies 825 on Level 3 BLAS for most of its computation. When computing singular vectors in 826 Figure 22(b), it is still about $1.8\times$ faster than LAPACK, even though it requires an 827 extra $2 \times 2(1 + \frac{i_b}{n_b})n^3$ operations to multiply by U_b and V_b . Here, the accelerated 828 MAGMA one-stage version is still faster.

829 For the tall 3:1 case in Figure 22(c), both LAPACK and MAGMA fare better, 830 since part of the computation is in the initial QR factorization, which is primarily 831 efficient Level 3 BLAS operations for all three implementations (LAPACK, MAGMA, 832 and PLASMA). For the very tall 1000:1 matrices in Figures 22(e) and 22(f), PLASMA 833 and MAGMA rely on their efficient QR factorization. In PLASMA, this is an imple- 834 mentation of the tall-skinny QR [1, 29], which even beats the accelerated MAGMA 835 implementation.

836 Overall, we expected such an improvement using the two-stage technique, due to 837 its heavy reliance on Level 3 BLAS. Even when performing more operations, it can 838 still have an advantage.

839 **11.5. Energy Consumption.** As we move toward exascale computing, power 840 and energy consumption play increasingly critical roles. Figure 23 shows the power 841 consumption over time during the SVD computation. We observe that PLASMA has 842 the lowest energy consumption, due to its fast execution, despite having the highest 843 power rate, indicative of its high compute intensity using Level 3 BLAS. Its energy 844 consumption is about half that of LAPACK, and $23\times$ less than EISPACK, as shown in 845 Figure 24(b). When computing singular values only, no vectors, the difference is even 846 more remarkable, with PLASMA being $5.6\times$ more energy efficient than LAPACK, 847 and $40\times$ more energy efficient than EISPACK, as shown in Figure 24(a).

848 Interestingly, we can correlate the various phases of the computation with the 849 power consumption. For LAPACK, the long plateau in Figure 23 up to the 105 850 seconds mark is the reduction to bidiagonal, followed by divide and conquer, where 851 the power varies significantly, and ending with the two back transformations by U_1 852 and V_1 from the 130–150 seconds mark. In PLASMA, the reduction to bidiagonal 853 is significantly shorter, up to the 20 seconds mark, followed by divide and conquer, 854 and the back transformations by U_a , U_b , V_a , and V_b , which are twice as long as they 855 are in LAPACK. EISPACK, in contrast, has a very long and steady computation. It 856 uses only one core, and thus has low power consumption; but the computation itself 857 is $48\times$ longer than LAPACK.

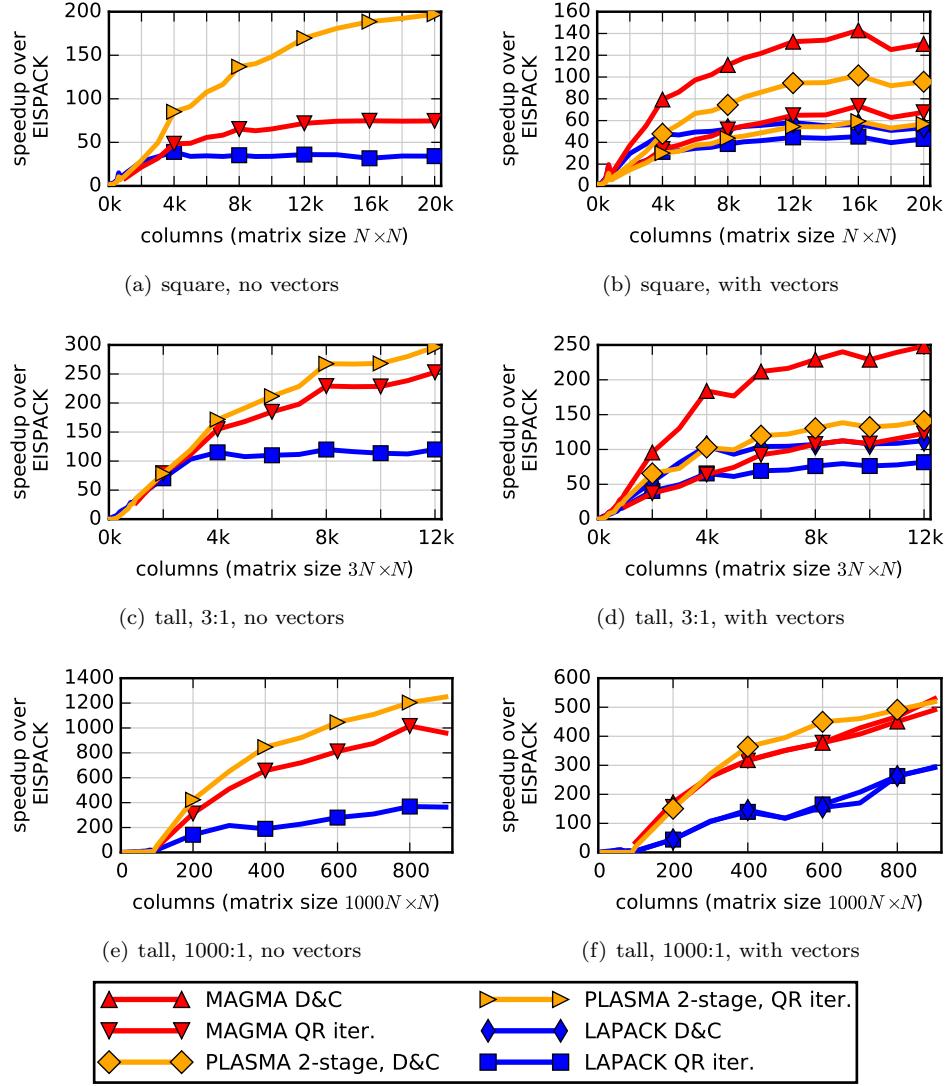


FIG. 22. Comparison of MAGMA, PLASMA, and LAPACK.

858 **11.6. MAGMA Accelerated Two-stage Reduction.** A two-stage algorithm
 859 can also be implemented very effectively using an accelerator. MAGMA accelerates
 860 the first-stage reduction to band form, as described above, and uses PLASMA for the
 861 second-stage reduction from band to bidiagonal. MAGMA also accelerates computa-
 862 tion of singular vectors, both applying the transformations from the second stage
 863 (e.g., $U_b U_2$) and applying the transformations from the first stage (e.g., $U_a (U_b U_2)$).
 864 Other steps are as in the accelerated one-stage MAGMA version. The profile in Fig-
 865 ure 17 shows the difference with the one-stage version: the reduction to bidiagonal
 866 (blue with \hatching) is significantly reduced, but multiplying $U = U_1 U_2 = U_a U_b U_2$
 867 and $V = V_1 V_2 = V_a V_b V_2$ (orange with // hatching) is increased.

868 Figure 25 shows the performance of the MAGMA two-stage implementation

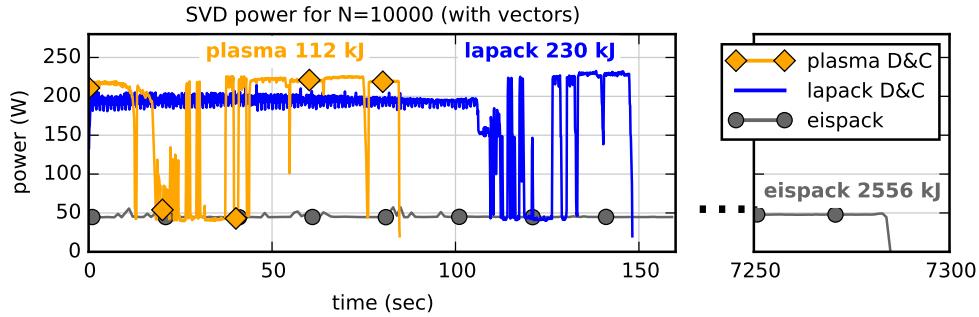


FIG. 23. Comparison of power during SVD computation for PLASMA, LAPACK, and EISPACK, for square matrix of size $n = 10000$. The total energy consumed during the computation is annotated for each.

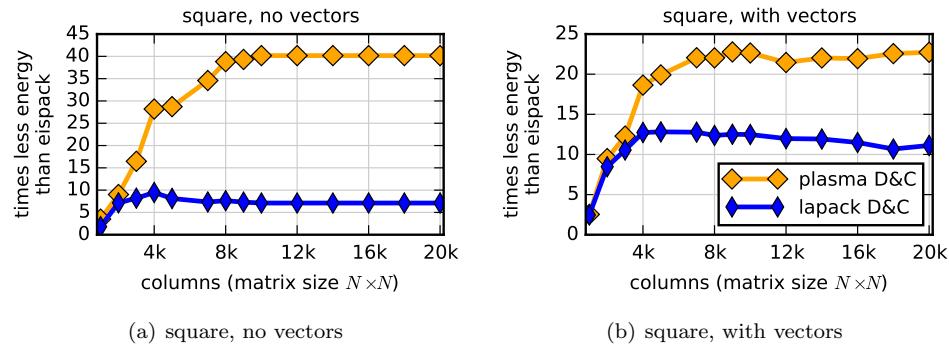


FIG. 24. Reduction in total energy consumption compared to EISPACK.

869 (dashed line), compared with the PLASMA two-stage and MAGMA one-stage im-
 870 plementations. The square, no vectors case in Figure 25(a) shows that for the bidiag-
 871 onal reduction itself, the two-stage MAGMA is up to $2.4\times$ faster than the two-stage
 872 PLASMA and $6.4\times$ faster than the one-stage MAGMA, and nearly $500\times$ faster than
 873 EISPACK. When computing singular vectors, in Figure 25(b), it is again up to $2.4\times$
 874 faster than PLASMA, but only $1.7\times$ faster than the one-stage MAGMA, due to the
 875 extra cost in multiplying by U_b and V_b . It also performs well in the tall 3:1 case, while
 876 for the tall 1000:1 case, its time is dominated by the initial QR factorization, so it
 877 performs similarly to the one-stage MAGMA.

878 **11.7. DPLASMA Implementation for Distributed Memory.** To cover the
 879 distributed memory environment, we also performed a study on a modern, large dis-
 880 tributed system. It is representative of a vast class of supercomputers commonly used
 881 for computationally intensive workloads. The DPLASMA algorithm is the two-stage
 882 algorithm described above for multicore, but implemented using the PaRSEC runtime
 883 engine [19, 18] to exploit the data flow representation, handle all the communication,
 884 and provide asynchronous task execution based on dependency analysis. PaRSEC
 885 employs the dataflow programming and execution model to provide a dynamic plat-
 886 form that can address the challenges posed by distributed hardware resources. The
 887 PaRSEC runtime combines the source program's task and dataflow information with
 888 supplementary information provided by the user—such as data distribution or hints

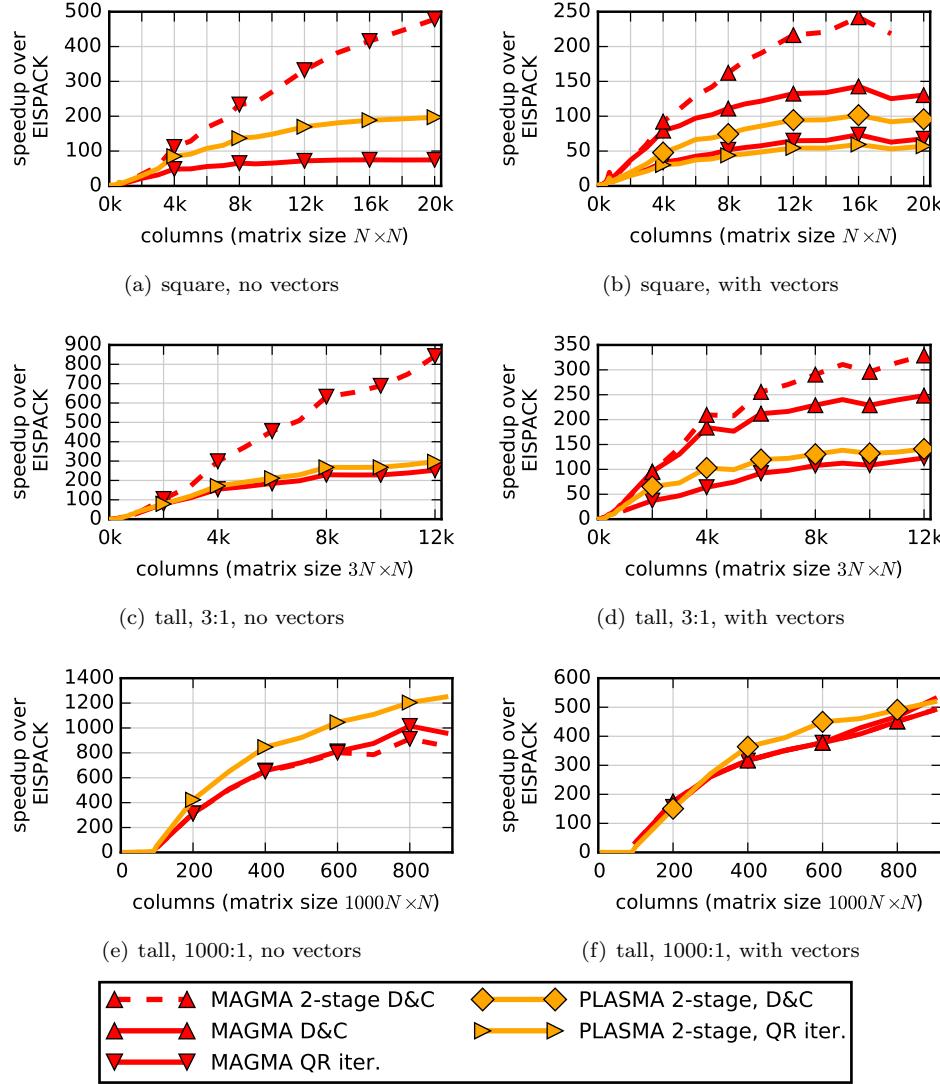


FIG. 25. Comparison of MAGMA 2-stage with MAGMA 1-stage and PLASMA 2-stage.

889 about the importance of different tasks—and orchestrates task execution on the available
 890 hardware. From a technical perspective, ParSEC is an event-driven system.
 891 When an event occurs, such as task completion, the runtime reacts by examining the
 892 dataflow to discover what future tasks can be executed based on the data generated
 893 by the completed task. The runtime handles the data exchange between distributed
 894 nodes, and thus it reacts to the events triggered by the completion of data trans-
 895 fers as well. Thus, communications become implicit and are handled automatically
 896 as efficiently as possible by the runtime. When no events are triggered because the
 897 hardware is busy executing application code, the runtime gets out of the way, allowing
 898 all hardware resources to be devoted to the application code's execution.

899 We benchmarked our two-stage implementation from the DPLASMA library, and

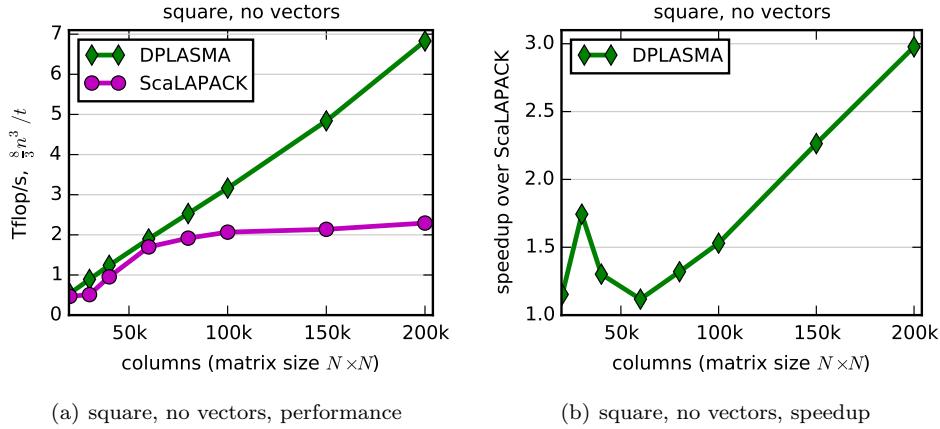


FIG. 26. Comparison of DPLASMA and ScaLAPACK computing singular values only for square matrices on 49 nodes (1764 cores).

900 the ScaLAPACK SVD routine from Intel MKL. Because only the singular values
 901 computation of our two-stage approach is currently implemented in the distributed
 902 DPLASMA library, we limited our tests to the case where only the singular values are
 903 computed. We performed our experiment on a recent hardware system consisting of
 904 49 distributed nodes, where every node has two sockets of 18-core Intel Xeon E5-2697
 905 (Broadwell) processors, running at 2.6 GHz, providing a total of 1764 cores. Each
 906 socket has 35 MiB of shared L3 cache, and each core has a private 3.5 MiB L2 and
 907 448 KiB L1 cache. The system is equipped with 52 GiB of memory per node. When
 908 only singular values are to be computed, the SVD solution consists of the reduction
 909 to bidiagonal and the computation of the singular values using QR iteration. Note
 910 that QR iteration on the bidiagonal matrix is a sequential process and thus it does
 911 not exploit any parallelism for either DPLASMA or ScaLAPACK. Its computational
 912 time is the same on either 1 or 49 nodes, and this time increases quadratically with
 913 the matrix size. Thus, the percentage of time spent in this portion varies with the
 914 matrix size. QR iteration consists of less than 5% of the time for a matrix of size 20k,
 915 while it reaches about 15% for ScaLAPACK and 26% for DPLASMA for a matrix of
 916 size 200k. As a result, the speedup will be affected by this constant:

$$917 \quad \text{speedup} = \frac{\text{time}_{\text{DPLASMA-BRD}} + t_x}{\text{time}_{\text{SCALAPACK-BRD}} + t_x},$$

918 where t_x is the time required to perform the bidiagonal singular value computation.
 919 Figure 26 shows the comparison between our implementation versus the ScaLAPACK
 920 `pdgesvd`. Asymptotically, our code achieves up to a 3 \times speedup for the largest matrices
 921 tested. This is the result of the efficient implementation of the first stage (reduction
 922 to band) using the PaRSEC engine, which enables us to exploit the compute-intensive
 923 nature of this stage, thereby minimizing the communication cost, and also from the
 924 careful design and implementation of the second stage that maps both the algorithm
 925 and the data to the hardware using cache-friendly kernels and data-locality-aware
 926 scheduling. Note that for small matrix sizes (e.g., a matrix of size 20k), there is not
 927 enough parallelism to exploit the 1764 available cores to make our two-stage algorithm
 928 3 \times faster; the tile size is about 160, so there are only about 125 tiles in each direction.

929 **12. Jacobi methods.** In contrast to bidiagonalization methods, Jacobi methods
 930 operate on the entire matrix A , without ever reducing to bidiagonal. This allows
 931 Jacobi methods to attain high relative accuracy, which will be discussed in [section 13](#).
 932 Jacobi first proposed his method in 1848 for solving the symmetric eigenvalue prob-
 933 lem [73] by diagonalizing the matrix A using a sequence of plane rotations:

$$934 \quad A_{(0)} = A, \quad A_{(k+1)} = J_{(k)}^T A_{(k)} J_{(k)}, \quad A_{(k)} \rightarrow \Lambda \text{ as } k \rightarrow \infty.$$

936 Each plane rotation, $J_{(k)} = J_{(k)}(i, j, \theta)$, now called a Jacobi or Givens rotation, is an
 937 orthogonal matrix that differs from the identity only in rows and columns i and j :

$$938 \quad J(i, j, \theta) = \begin{bmatrix} I & & & \\ & c & s & \\ & -s & c & \\ & & & I \end{bmatrix},$$

940 where $c = \cos \theta$ and $s = \sin \theta$. The angle θ is chosen to eliminate the pair a_{ij}, a_{ji}
 941 by applying $J(i, j, \theta)$ on the left and right of A , which can be viewed as the 2×2
 942 eigenvalue problem,

$$943 \quad 944 \quad \hat{J}_{(k)}^T \hat{A}_{(k)} \hat{J}_{(k)} = \begin{bmatrix} c & s \\ -s & c \end{bmatrix}^T \begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix} \begin{bmatrix} c & s \\ -s & c \end{bmatrix} = \begin{bmatrix} d_{ii} & 0 \\ 0 & d_{jj} \end{bmatrix} = \hat{A}_{(k+1)},$$

945 where the notation \hat{A} is the 2×2 submatrix $\begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$ of matrix A . Subsequent
 946 eliminations will fill in the eliminated entry, but at each step the norm of off-diagonal
 947 elements,

$$948 \quad 949 \quad \text{off}(A) = \|A - \text{diag}(A)\|_F = \left(\sum_{i \neq j} a_{ij}^2 \right)^{1/2},$$

950 is reduced until the matrix converges to diagonal form, Λ , revealing the eigenvalues.
 951 Accumulating the plane rotations, $V = J_{(0)} J_{(1)} \dots$, yields the eigenvectors. Originally,
 952 Jacobi chose to eliminate the off-diagonal pair a_{ij}, a_{ji} of largest magnitude at
 953 each step, giving the largest possible reduction in $\text{off}(A)$. This is inefficient as it in-
 954 troduces an $O(n^2)$ search for each rotation of $O(n)$ work. Instead, in modern times
 955 the method was reformulated so that one *sweep* goes over all $n(n-1)/2$ combinations
 956 of (i, j) with $i < j$ in a predetermined order, typically cyclic by rows, i.e.,

$$957 \quad (1, 2), (1, 3), \dots, (1, n); (2, 3), \dots, (2, n); \dots; (n-1, n),$$

959 or cyclic by columns. It converges after a small number of sweeps, typically 5–10.
 960 Wilkinson [116] showed that convergence is ultimately quadratic. Rutishauser [102]
 961 gave a robust implementation in the Wilkinson-Reinsch Handbook.

962 **12.1. Two-sided Jacobi SVD.** Jacobi's eigenvalue method was generalized to
 963 the SVD of a general, nonsymmetric matrix in two different ways. The first way
 964 is the two-sided method due to Kogbetliantz [76], which applies two different plane
 965 rotations, $J(i, j, \theta)$ on the left of A and $K(i, j, \phi)$ on the right of A , to eliminate the
 966 a_{ij} and a_{ji} entries. As before, sweeps are done over the off-diagonal entries until

Algorithm 5 Two-sided Jacobi SVD method for $n \times n$ matrix A .

```

function two_sided_jacobi_svd(  $A$  )
   $U = I$ ;  $V = I$ 
  repeat // loop over sweeps
    for each pair  $(i, j)$ ,  $i < j$ , in prescribed order
      solve  $2 \times 2$  SVD  $\hat{J}^T \hat{A}_{(k)} \hat{K} = \hat{A}_{(k+1)}$ 
       $A = J^T A$  // update rows  $i$  and  $j$ 
       $A = AK$  // update cols  $i$  and  $j$ 
       $U = UJ$ 
       $V = VK$ 
    end
    until  $\text{off}(A) < \text{tol} \|A_0\|_F$ 
    for  $i = 1, \dots, n$ 
       $\sigma_i = |a_{ii}|$ 
      if  $a_{ii} < 0$  then  $u_i = -u_i$ 
    end
    sort  $\Sigma$  and apply same permutation to columns of  $U$  and  $V$ 
    return  $(U, \Sigma, V)$ 
end function

```

967 the norm of off-diagonal entries is below a specified tolerance, revealing the singular
 968 values, Σ , via the iteration:

969 $A_{(0)} = A$, $A_{(k+1)} = J_{(k)}^T A_{(k)} K_{(k)}$, $A_{(k)} \rightarrow \Sigma$ as $k \rightarrow \infty$.

971 Accumulating the left rotations, $U = J_{(0)} J_{(1)} \dots$, gives the left singular vectors, while
 972 accumulating the right rotations, $V = K_{(0)} K_{(1)} \dots$, gives the right singular vectors.
 973 Determining $J(i, j, \theta)$ and $K(i, j, \phi)$ can be viewed as solving a 2×2 SVD problem,

974 (8) $\hat{J}_{(k)}^T \hat{A}_{(k)} \hat{K}_{(k)} = \begin{bmatrix} c_J & s_J \\ -s_J & c_J \end{bmatrix}^T \begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix} \begin{bmatrix} c_K & s_K \\ -s_K & c_K \end{bmatrix} = \begin{bmatrix} d_{ii} & \\ & d_{jj} \end{bmatrix} = \hat{A}_{(k+1)}.$

976 The angles for J and K are not uniquely determined, so various methods have been
 977 derived [22, 49, 76]. Brent et al. [22] proposed the algorithm USVD, which uses one
 978 rotation to symmetrize the 2×2 subproblem, then a second rotation to eliminate the
 979 off-diagonal entries. This produces an unnormalized SVD, where the diagonal entries
 980 are unsorted and may be negative. Post-processing to sort and adjust the signs of
 981 the singular values and singular vectors yields a standard SVD. They also formulated
 982 the normalized rotation/reflection algorithm NSVD that corrects the signs during the
 983 iteration. [Algorithm 5](#) outlines the two-sided Jacobi method.

984 Rectangular matrices can be handled by first doing a QR factorization, optionally
 985 with pivoting, and then doing the SVD of the R matrix, as previously described for
 986 bidiagonalization methods ([subsection 5.4](#)). For Jacobi, this QR factorization has the
 987 added benefit of preconditioning the system to converge faster, as discussed further
 988 in [subsection 12.5](#).

989 Heath et al. [67] developed a variant for computing the SVD of a product of
 990 matrices, $A = B^T C$, without explicitly forming A . Applying rotations $B_{(k+1)} = B_{(k)} J$
 991 and $C_{(k+1)} = C_{(k)} K$ implicitly applies J and K on both sides of A . When $B = C$, it
 992 simplifies to the one-sided Jacobi method, discussed next.

Algorithm 6 One-sided Jacobi SVD method for $m \times n$ matrix A , $m \geq n$

```

function one_sided_jacobi_svd(  $A$  )
   $V = I$ 
  repeat // loop over sweeps
    done = true
    for each pair  $(i, j)$ ,  $i < j$ , in prescribed order
       $b_{ii} = A_i^T A_i = \|A_i\|^2$ 
       $b_{jj} = A_j^T A_j = \|A_j\|^2$ 
       $b_{ij} = A_i^T A_j$ 
      if  $|b_{ij}| \geq \epsilon \sqrt{b_{ii} b_{jj}}$  then
        solve  $2 \times 2$  symmetric eigenvalue problem  $\hat{J}^T \hat{B} \hat{J} = \hat{D}$ 
         $A = A J$  // update cols  $i$  and  $j$ 
         $V = V J$ 
        done = false
      end
    end
  until done
  for  $i = 1, \dots, n$ 
     $\sigma_i = \|a_i\|_2$ 
     $u_i = a_i / \sigma_i$ 
  end
  sort  $\Sigma$  and apply same permutation to columns of  $U$  and  $V$ 
  return  $(U, \Sigma, V)$ 
end function

```

993 **12.2. One-sided Jacobi.** The second way to generalize the Jacobi method to
 994 the SVD is a one-sided method due to Hestenes [68]. Earlier we noted that the
 995 SVD can be solved by computing the eigenvalues of the Gram matrix, $A^T A$, but
 996 that explicitly forming $A^T A$ is undesirable for numerical reasons. Instead, Hestenes
 997 applied plane rotations on only the right side of A to orthogonalize the columns of
 998 A , which implicitly performs the two-sided Jacobi eigenvalue method on $A^T A$. The
 999 columns of A converge to $U \Sigma$, that is, the left singular vectors scaled by the singular
 1000 values:

1001
$$A_{(0)} = A, \quad A_{(k+1)} = A_{(k)} J_{(k)}, \quad A_{(k)} \rightarrow U \Sigma \text{ as } k \rightarrow \infty.$$

1003 This means that, implicitly, $A_{(k)}^T A_{(k)} \rightarrow \Sigma^2$. Accumulating the rotations, $V =$
 1004 $J_{(0)} J_{(1)} \dots$, gives the right singular vectors. Alternatively, V can be solved for after
 1005 the iteration, as described below in [subsection 12.5](#).

1006 The rotations are determined similarly to the Jacobi eigenvalue method, by solv-
 1007 ing the 2×2 eigenvalue problem

1008 (9)
$$\hat{J}_{(k)}^T \begin{bmatrix} b_{ii} & b_{ij} \\ b_{ij} & b_{jj} \end{bmatrix} \hat{J}_{(k)} = \begin{bmatrix} d_{ii} & \\ & d_{jj} \end{bmatrix},$$

1010 where $b_{ij} = a_i^T a_j$ and a_i is the i -th column of $A_{(k)}$. Over the coarse of a sweep, it
 1011 computes the matrix $B = A^T A$, however, J is not applied directly to $A^T A$, but to A
 1012 itself, avoiding the numerical instabilities associated with $A^T A$. [Algorithm 6](#) outlines
 1013 the one-sided Jacobi method.

1014 It skips rotations if $|b_{ij}| < \epsilon\sqrt{b_{ii}b_{jj}}$, indicating that columns a_i and a_j are already numerically orthogonal. It converges when all rotations in a sweep are skipped. Using this formula to check for convergence is required for attaining high relative accuracy [32] (see [section 13](#)). The b_{ii} column norms can be cached rather than recomputed for each pair, which reduces operations when rotations are skipped. Note that the last sweep takes about n^3 flops computing b_{ij} terms to check for convergence, without doing any useful work.

1021 A left-handed version can be defined analogously, by applying rotations on the left to orthogonalize the rows of A [84]. This might be preferred if A is a wide matrix stored row-wise, rather than a tall matrix stored column-wise.

1024 One-sided Jacobi can be applied to a rectangular matrix, but again, preprocessing using a QR factorization, and apply Jacobi on the square R matrix, reduces the operation count and preconditions the system for faster convergence; see [subsection 12.5](#).

1027 **12.3. Convergence.** For the row and column cyclic orderings, Forsythe and 1028 Henrici [49] proved that all the Jacobi methods (two-sided eigenvalue, one-sided SVD, 1029 and two-sided SVD) converge, provided the rotation angles are bounded below $\pi/2$ 1030 by some b ,

$$1031 \quad (10) \quad |\theta| \leq b < \pi/2.$$

1033 For the two-sided eigenvalue and one-sided SVD methods, θ can always be chosen 1034 to satisfy (10); see [102]. For the two-sided SVD method, however, this condition 1035 may fail to hold. In Forsythe and Henrici's method, the bound is $b = \frac{\pi}{2}$, which 1036 may introduce a cycle interchanging two singular values without converging. For the 1037 methods of Brent et al. [22], NSVD has a bound $b = 3\pi/4$ and USVD has a bound 1038 $b = 5\pi/4$. Proofs for other orderings, particularly for parallel orderings, have been 1039 elusive. Despite failing to satisfy the convergence proof's prerequisites, in practice 1040 Jacobi methods reliably converge. Using a threshold to skip updates to small entries 1041 is a common tactic, especially in the first several sweeps, to accelerate and guarantee 1042 convergence [102, 27, 8].

1043 When applied to triangular matrices, Heath et al. [67] and Hari and Veselić [66] 1044 observed that applying one sweep of the two-sided SVD method with the row-cyclic 1045 ordering (without thresholding) converts an upper triangular matrix to lower triangular, 1046 and vice-versa. Hari and Veselić derived rotation angle formulas in the triangular 1047 case, and prove that the angles are bounded below $\pi/2$, guaranteeing convergence. 1048 Hari and Matejaš [65] later derived more accurate formulas.

1049 Applying column pivoting during the Jacobi iterations can improve convergence. 1050 In the one-sided method, de Rijk [27] follows the row-cyclic ordering, but at the start 1051 of row i , searches columns i, \dots, n for the column of largest norm and pivots it to 1052 column i . Unfortunately, using the row-cyclic ordering makes parallelism difficult. 1053 Zhou and Brent [120] likewise show that sorting column norms improves convergence, 1054 and give a parallel ordering for sorting.

1055 **12.4. Parallel orderings.** In two-sided Jacobi, a pair of rotations applied on 1056 the left and right affect only two rows and two columns. In one-sided Jacobi, each 1057 rotation applied on the right affects only two columns. Therefore, in both cases, 1058 $\lfloor n/2 \rfloor$ rotations can be performed in parallel. However, the row and column cyclic 1059 orderings are not amenable to parallel computation, as they introduce dependencies 1060 between consecutive pairs of elements. Since there are $n(n - 1)/2$ pairs to eliminate, 1061 an optimal parallel ordering would have $n - 1$ steps, with each step eliminating $n/2$

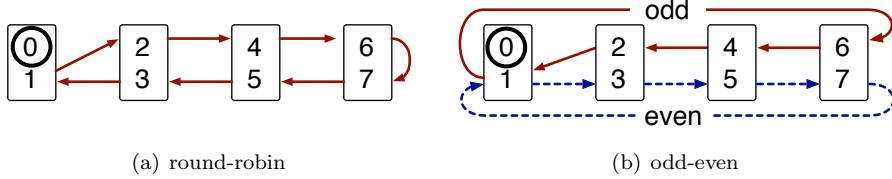


FIG. 27. *Parallel orderings.* Rectangles indicating processors are labeled with their assigned columns. Arrows depict movement of columns between Jacobi sweeps. Circled pivot column is stationary.

1062 pairs in parallel (for n even). Many different parallel Jacobi orderings have been
 1063 devised. While parallel orderings typically lack a proof of convergence, in practice
 1064 they work reliably.

1065 Commonly, for parallel implementations of both one-sided and two-sided Jacobi,
 1066 the matrix is distributed by columns. Early systolic implementations placed two
 1067 columns [21] or a 2×2 submatrix [22] per processor. Later block implementations
 1068 placed two block columns [15, 11] or a 2×2 block submatrix [12] per processor.
 1069 When each processor stores two columns, one-sided Jacobi has the advantage that no
 1070 communication is required during an update, whereas in two-sided Jacobi, the left
 1071 transformations (J 's) must be broadcast in an all-to-all fashion.

1072 Brent and Luk [21] introduced the round-robin ordering, shown in Figure 27(a),
 1073 which had previously been known for chess tournaments. After each Jacobi rotation,
 1074 each node sends and receives two columns, except the pivot node that sends and
 1075 receives one column. Eberlein [46] gave the odd-even ordering in Figure 27(b). After
 1076 each odd sweep, the odd permutation (solid red lines) is used; after even sweeps,
 1077 the even permutation (dashed blue lines) is used. Each node sends and receives one
 1078 column.

1079 Luk and Park [85] studied the equivalence of orderings, demonstrating that many
 1080 orderings are equivalent in the sense that relabeling the columns gives identical or-
 1081 derings. For example, choosing a different pivot column in round-robin will give an
 1082 equivalent ordering. Luk and Park showed that the two main classes of Jacobi order-
 1083 ings are the round-robin and odd-even types. Bečka and Vajterščík [12, 11] compared
 1084 implementations of the round-robin, odd-even, and a butterfly-like ordering inspired
 1085 by the Fast Fourier Transform (FFT), on ring, hypercube, and mesh networks for
 1086 block Jacobi methods.

1087 **12.5. Preconditioning.** Another means to improving the speed of Jacobi meth-
 1088 ods is to precondition the matrix to reduce the number of sweeps required for con-
 1089 vergence. Drmač and Veselić [44] introduced several forms of preconditioning for
 1090 the one-sided Jacobi method. The major ideas are outlined below, with a simplified
 1091 version in Algorithm 7.

1092 First, for a square matrix A , heuristically choose to factor either $X = A$ or
 1093 $X = A^T$. They give the example of $A = DQ$, where D is diagonal and Q is orthogonal.
 1094 One-sided Jacobi applied to A implicitly diagonalizes $Q^T D^2 D$, while applied to A^T ,
 1095 it implicitly diagonalizes D^2 , which is already diagonal. One heuristic they suggest is
 1096 to choose the X that maximizes $\|\text{diag}(X^T X)\|_2$, hence minimizing $\text{off}(X^T X)$. Their
 1097 second heuristic is to choose the X that minimizes the diagonal entropy of $X^T X$,

Algorithm 7 Preconditioned one-sided Jacobi SVD method (simplified)

```

function preconditioned_one_sided_jacobi( A )
  input:  $m \times n$  matrix A,  $m \geq n$ 
  output: U,  $\Sigma$ , V
  transpose = ( $m == n$  and  $\eta_d(AA^T) < \eta_d(A^T A)$ ) // see (11)
  if transpose then
    A = AT
  end
   $Q_r R P_r^T = A$  // QR factorization with column pivoting
   $L Q_l = R$  // LQ factorization
  (Ul,  $\Sigma$ ) = one_sided_jacobi_svd(L) // Algorithm 6; skip V
  U = QrUl
  V = PrQlTL-1(Ul $\Sigma$ ) or V = PrR-1(Ul $\Sigma$ )
  if transpose then
    swap U  $\leftrightarrow$  V
  end
end function

```

1098 defined by

1099 (11)
$$\eta_d(X^T X) = \eta \left(\text{diag}(X^T X) / \text{trace}(X^T X) \right),$$

1101 where the entropy of a vector *p* with $p_i \geq 0$, $\sum_i p_i = 1$, is defined as

1102 (12)
$$\eta(p) = -\frac{1}{\log n} \sum_{i=1}^n p_i \log p_i, \quad \text{with } 0 \log 0 \equiv 0.$$

1104 Both heuristics are $O(n^2)$.1105 The second preconditioning technique is to use a QR factorization with column
1106 pivoting (QRP) of *A*, then factor *R*. This concentrates the mass of the matrix along
1107 the diagonal of *R*^T*R*, reducing the number of Jacobi sweeps. For a rectangular $m \times n$
1108 problem, $m > n$, this also shrinks it to an $n \times n$ problem, as in [subsection 5.4](#).1109 Third, use either an LQ factorization of *R*, or simply let $L = R^T$, then factor
1110 *L*. An LQ factorization further concentrates the mass along the diagonal of *L*^T*L*.
1111 Using LQ is particularly advantageous in the rank deficient case. For a matrix of
1112 rank *r*, QRP generates $R = \begin{bmatrix} R_{11} & R_{12} \\ 0 & R_{22} \end{bmatrix}$ with the $(n-r) \times (n-r)$ block *R*₂₂ being
1113 negligible. Doing an LQ factorization of $[R_{11} \ R_{12}]$ yields a smaller, $r \times r$, full-rank
1114 matrix *L*. Alternatively, simply using $L = R^T$ is an implicit step of Rutishauser's LR
1115 diagonalization applied to *R*^T*R*, again concentrating mass along the diagonal of *L*^T*L*
1116 as compared to *R*^T*R*.1117 Additionally, Drmač and Veselić's error analysis based on using QRP and option-
1118 ally LQ factorization shows that computing *V* by solving with either of the triangular
1119 matrices *L* or *R* is numerically stable and generates an orthogonal matrix; see [Al-](#)
1120 [gorithm 7](#) for specifics. This allows us to skip accumulating *V* during the one-sided
1121 Jacobi iteration, removing some Level 1 BLAS operations, and adding Level 3 BLAS
1122 operations after the iteration, so we can expect a good performance increase. Their
1123 paper gives detailed algorithms that make choices about which preconditioning to use

1124 based on condition estimates. Hari [64] and Bečka et al. [9] also applied QRP and LQ
 1125 preconditioning in the context of parallel one-sided block Jacobi.

1126 In addition to preconditioning, Drmač and Veselić [45] introduced optimizations in
 1127 the one-sided Jacobi iteration, based on the structure of the preconditioned matrix. In
 1128 the first sweep, the zero structure of the triangular matrix can be exploited to reduce
 1129 computation. Second, based on work by Mascarenhas [87], they use a modified row-
 1130 cyclic strategy to more frequently visit diagonal blocks, since those blocks converge
 1131 at a slower rate. Heuristically, based on the expectation that $L^T L$ is diagonally
 1132 dominant, during the first few sweeps, if two rotations in a row are skipped due to
 1133 thresholding, they skip the rest of the row. This avoids computing dot products
 1134 when the rotation will likely be skipped. Finally, they use a tiled row-cyclic strategy
 1135 to improve cache efficiency. All of these improvements combine for a more efficient
 1136 algorithm.

1137 Okša and Vajteršic [95] showed that the same preconditioning techniques, QRP
 1138 factorization optionally followed by LQ factorization, also improve convergence for the
 1139 parallel two-sided block Jacobi method. In their tests, preconditioning concentrated
 1140 more than 99% of the weight of $\|A\|_F$ into the diagonal blocks. Depending on the
 1141 singular value distribution, this gave up to an order-of-magnitude reduction in time.
 1142 This preconditioning was later extended to multiple QR iterations [10].

1143 As noted earlier, two-sided Jacobi preserves the triangular structure when used
 1144 with an appropriate cyclic ordering. Hari and Matejaš [65, 88, 89] use the QRP and
 1145 LQ preprocessing to generate triangular matrices. They prove high relative accuracy
 1146 results for the two-sided Jacobi method on such triangular matrices, and utilize a
 1147 parallel ordering due to Sameh [103] that preserves the triangular structure.

1148 **12.6. Block Jacobi.** In section 5, we saw that blocking was a major improve-
 1149 ment for SVD methods. Blocking can also be favorably applied to Jacobi methods.
 1150 Van Loan [112] and Bischof [15] were among the first to describe two-sided block
 1151 Jacobi SVD methods. The method is very similar to the non-block implementation,
 1152 with plane rotations J and K operating on two rows or columns now becoming or-
 1153 thogonal block rotations operating on two block rows or block columns. For a block
 1154 size n_b , let $N = \lceil n/n_b \rceil$ be the number of blocks. The indices i, j now loop over the
 1155 blocks, $1, \dots, N$. We reinterpret the notation \hat{A} to be the 2×2 block matrix

$$1156 \hat{A} = \begin{bmatrix} A_{ii} & A_{ij} \\ A_{ji} & A_{jj} \end{bmatrix},$$

1157 where each A_{ij} is an $n_b \times n_b$ block. Instead of the 2×2 SVD (8), it computes a 2×2
 1158 block SVD,

$$1159 \hat{J}^T \hat{A} \hat{K} = \hat{J}^T \begin{bmatrix} A_{ii} & A_{ij} \\ A_{ji} & A_{jj} \end{bmatrix} \hat{K} = \begin{bmatrix} D_{ii} & 0 \\ 0 & D_{jj} \end{bmatrix},$$

1160 either recursively using a serial Jacobi method, or using some other SVD method
 1161 such as QR iteration. Each processor now holds two block columns. Block row and
 1162 column updates by the orthogonal matrices J and K are applied as Level 3 BLAS
 1163 matrix multiplies, greatly enhancing the efficiency of the algorithm.

1164 Bischof [15] investigated two methods to solve the SVD subproblem: using QR
 1165 iteration or using a single sweep of two-sided Jacobi. In the later case, using only one
 1166 sweep the block method does not fully annihilate the off-diagonal blocks of the 2×2
 1167 block subproblem, and is in fact simply a reorganization of the non-block method,
 1168 but with updates applied using Level 3 BLAS. He found that using Jacobi to solve

1170 the subproblem was faster than using QR iteration; however, this was prior to the
 1171 fast blocked versions of QR iteration available in LAPACK.

1172 Arbenz and Slapničar [5] gave an early implementation for the one-sided block
 1173 Jacobi SVD method. Again, the block method is very similar to the non-block method,
 1174 with the 2×2 eigenvalue problem (9) being replaced with a 2×2 block eigenvalue
 1175 problem,

$$1176 \quad \hat{J}^T \hat{A} \hat{J} = \hat{J}^T \begin{bmatrix} B_{ii} & B_{ij} \\ B_{ij}^T & B_{jj} \end{bmatrix} \hat{J} = \begin{bmatrix} D_{ii} & 0 \\ 0 & D_{jj} \end{bmatrix},$$

1178 with $B_{ij} = A_i^T A_j$, where A_i as the i -th block column of A . Arbenz and Slapničar used
 1179 the two-sided Jacobi eigenvalue method to solve the subproblem, which is important
 1180 for preserving Jacobi's high relative accuracy. Hari [64] derived an optimization using
 1181 the cosine-sine decomposition as a kind of "fast scaled block-rotation", reducing the
 1182 flop count up to 40%. Boukaram et al. [20] developed batched one-sided Jacobi and
 1183 block Jacobi methods for GPUs, to compute SVD factorizations of a batch of small
 1184 matrices.

1185 Bećka et al. introduced dynamic orderings for the two-sided [8] and one-sided [9]
 1186 Jacobi methods. Instead of using a cyclic ordering such as row-cyclic, round-robin, or
 1187 odd-even, the idea is to find the off-diagonal blocks of maximum norm to eliminate.
 1188 This is Jacobi's original idea, applied on the block level. Using a greedy solution to
 1189 the *maximum-weight perfect matching problem* takes $O(p^2 \log p)$ time for p processors
 1190 and yields a set of $N/2$ subproblems of maximum weight to solve in parallel. Their
 1191 results show significantly improved convergence and time to solution.

1192 **12.7. Performance analysis.** While Jacobi methods have a long history, even
 1193 predating bidiagonalization methods, many implementations have been either research
 1194 codes or for unique systems like the ILLIAC IV [84]. Therefore, we do not have as
 1195 rich a collection of historical implementations to compare as for bidiagonalization
 1196 methods. We tested four current implementations of Jacobi methods:

- 1197 • One-sided Jacobi, available in LAPACK as `dgesvj`, due to Drmač [44].
- 1198 • Preconditioned one-sided Jacobi, available in LAPACK as `dgejsv`, due to
 1199 Drmač [44].
- 1200 • Two-sided Jacobi, available in Eigen 3.3.3 [47].
- 1201 • Preconditioned one-sided block Jacobi, due to Bećka et al. [9].

1202 Jacobi has traditionally trailed bidiagonalization methods in performance for two
 1203 reasons. First, a comparison of flops in Figure 28 shows that for computing singular
 1204 values only (no vectors), Jacobi cannot finish even one sweep in the same flops as
 1205 bidiagonalization ($\frac{8}{3}n^3$). When computing vectors, Jacobi would need to complete
 1206 in two sweeps to have fewer flops than QR iteration, and one sweep to have fewer
 1207 flops than divide and conquer. However, with optimizations to skip rotations and
 1208 take advantage of matrix structure [45, 89], these Jacobi flop counts are significant
 1209 overestimates.

1210 However, as we have repeatedly seen, flops are now a poor metric for performance.
 1211 It matters whether flops are in compute-intensive Level 3 BLAS or not. For Jacobi,
 1212 dot products and plane rotations are Level 1 BLAS, so are memory bandwidth limited.
 1213 For preconditioned Jacobi, QR with column pivoting (QRP) has a mixture of Level 2
 1214 and Level 3 BLAS operations, similar to the traditional one-stage bidiagonalization
 1215 discussed in subsection 5.1, so its performance is also limited by memory bandwidth.
 1216 The triangular solve for V and multiplying QU will both be Level 3 BLAS operations.
 1217 The level of parallelism also matters. The two LAPACK implementations, one-sided

	no vectors	with vectors
QR iteration	$\frac{8}{3}n^3$	$\frac{52}{3}n^3 \approx 17n^3$
divide and conquer	$\frac{8}{3}n^3$	$\frac{28}{3}n^3 \approx 9n^3$
one-sided Jacobi	$5Sn^3$	$7Sn^3$
two-sided Jacobi	$4Sn^3$	$8Sn^3$
preconditioned one-sided Jacobi	$5Sn^3 + \frac{8}{3}n^3$	$5Sn^3 + \frac{17}{3}n^3$
preconditioned two-sided Jacobi	$4Sn^3 + \frac{8}{3}n^3$	$6Sn^3 + \frac{17}{3}n^3$

FIG. 28. Floating point operation counts for square $n \times n$ matrix and S Jacobi sweeps. For Jacobi, fast Givens rotations [63] are assumed. For preconditioned Jacobi, initial QRP and LQ factorizations and triangular solve for V are also assumed.

1218 Jacobi and preconditioned one-sided Jacobi, do not use explicit parallelism. Therefore,
 1219 the only parallelism is within the BLAS, which is very limited for Level 1 BLAS.
 1220 In contrast, the block Jacobi method uses Level 3 BLAS operations and explicit
 1221 parallelism via MPI, so we can expect much better performance.

1222 In Figure 29(a), for square matrices without vectors, both one-sided Jacobi meth-
 1223 ods were about half EISPACK’s speed, while with vectors in Figure 29(b), precon-
 1224 ditioned Jacobi is $2\times$ faster than plain Jacobi, and close to EISPACK’s speed. For
 1225 tall, 3:1 matrices in Figure 29(c), the plain one-sided Jacobi does not do an initial
 1226 QR factorization, so it remains about half of EISPACK’s speed, while the precondi-
 1227 tioned Jacobi improves to about $2\times$ EISPACK’s speed. When computing vectors in
 1228 Figure 29(d), the preconditioned Jacobi version gains even more, being over $3\times$ faster
 1229 than EISPACK.

1230 For the very tall-skinny 1000:1 case in Figures 29(e) and 29(f), the time with
 1231 preconditioned Jacobi is dominated by QRP, which uses more Level 2 and 3 BLAS
 1232 operations, so the performance improves to over $100\times$ EISPACK. LAPACK’s QR it-
 1233 eration uses a regular QR factorization (no pivoting), which is predominantly Level 3
 1234 BLAS, so its performance is significantly faster than Jacobi. However, QRP will gen-
 1235 erate a more accurate factorization than regular QR, especially if A is ill-conditioned.

1236 In most cases, the Jacobi single-core performance was identical to its multi-core
 1237 performance, indicating that the Level 1 BLAS routines do not have appreciable par-
 1238 allelism. For tall matrices, preconditioning gained an advantage when using multiple
 1239 cores, shown by the difference between the solid and dashed green lines in Figures 29(c)
 1240 to 29(f), due to parallelism within QRP, solving for V , and computing QU .

1241 In all of these results, the two-sided Jacobi implementation available in Eigen was
 1242 considerably slower. This can partly be explained because it has to update the matrix
 1243 both row-wise and column-wise, making for poor cache performance. For square
 1244 matrices, it does not do any preconditioning. For tall matrices, it uses QRP, which
 1245 improves its relative performance somewhat. (Note that Eigen can be configured to
 1246 instead call LAPACK’s QR iteration method.)

1247 Figure 30 shows results for the precondition one-sided block Jacobi method. We
 1248 tested two variants of the preconditioning, one using QR factorization with column
 1249 pivoting (QRP), the other using regular QR factorization (no pivoting). In both cases,
 1250 this was followed by an LQ factorization. This implementation has explicit parallelism
 1251 via MPI. It uses ScaLAPACK for the QRP, QR, and LQ factorizations. We see that

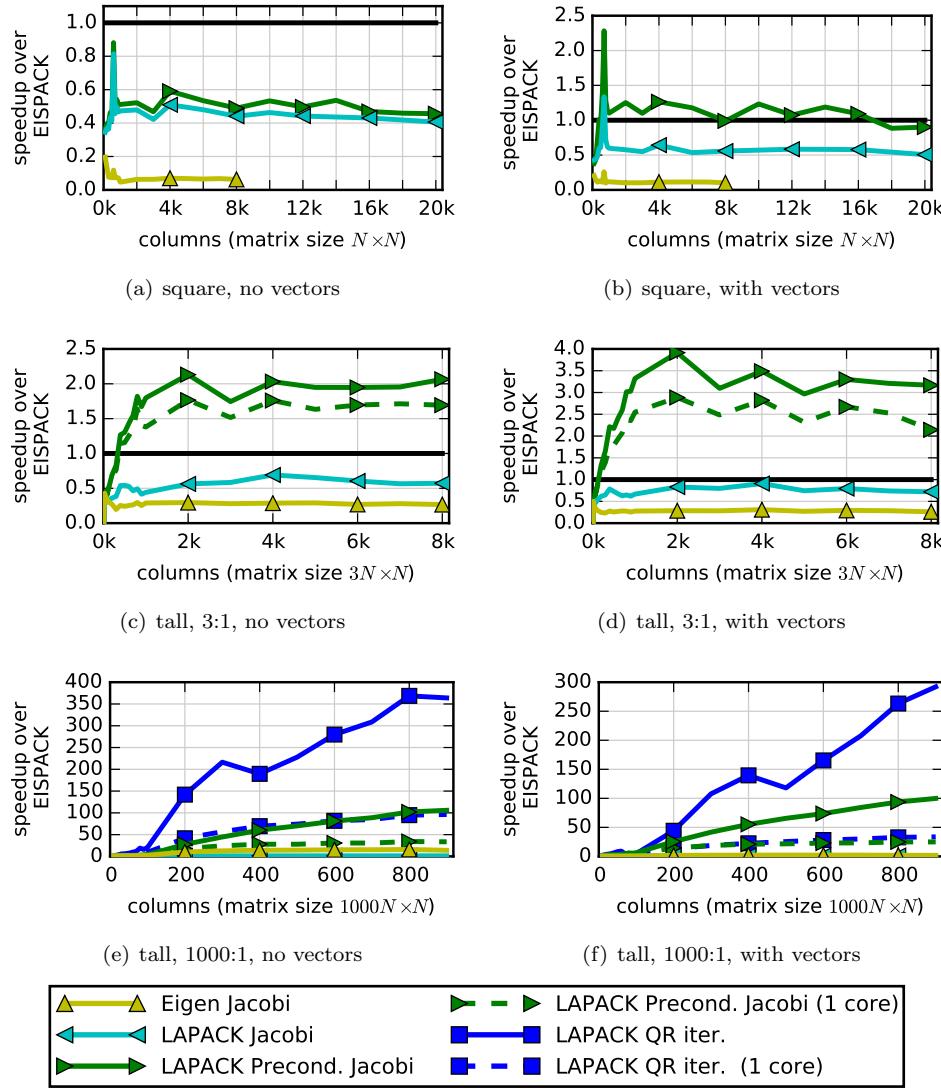


FIG. 29. Comparison of LAPACK's one-sided Jacobi, preconditioned one-sided Jacobi, and Eigen's two-sided Jacobi.

1252 with QRP + LQ, it performed similarly to ScaLAPACK QR iteration, while with QR
 1253 + LQ, it was a bit faster, matching LAPACK's QR iteration in performance for the
 1254 tall, 3:1 case.

1255 **13. Accuracy.** While Jacobi methods have struggled to compete with the per-
 1256 formance of bidiagonalization methods, for some classes of matrices they have a dis-
 1257 tinct advantage in accuracy, which is now their main motivation. In this section, we
 1258 briefly explore the accuracy differences between methods. The traditional perturba-
 1259 tion theory [32] for both bidiagonalization and Jacobi methods shows that

$$\frac{|\sigma_i - \hat{\sigma}_i|}{\sigma_i} \leq O(\epsilon) \kappa(A),$$

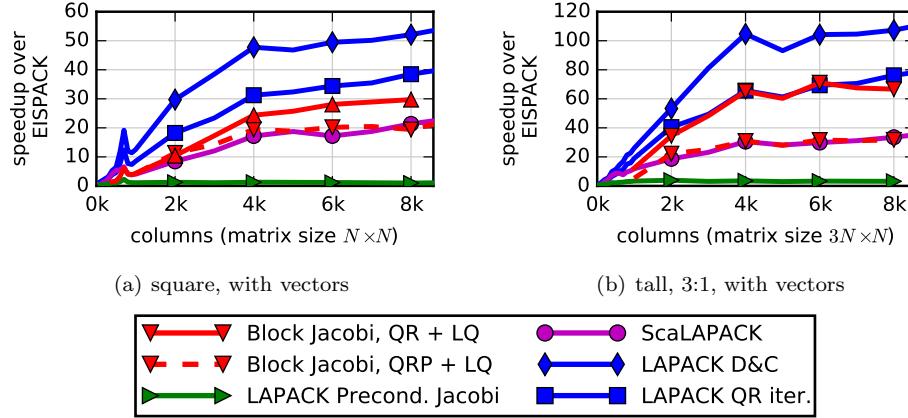


FIG. 30. Comparison of preconditioned one-side block Jacobi, LAPACK's preconditioned one-sided Jacobi, QR iteration, and divide and conquer.

1262 where σ_i and $\hat{\sigma}_i$ are the singular values of A and $A + \delta A$, respectively, with a small
 1263 perturbation δA such that $\|\delta A\|_2 \leq O(\epsilon) \|A\|_2$, and $\kappa(A)$ is the condition number
 1264 of A . This implies that large singular values are computed accurately, but small
 1265 singular values may be totally inaccurate if $\kappa(A)$ is large. For the one-sided Jacobi
 1266 SVD method, this bound can be improved. Specifically, on matrices of the form
 1267 $A = CD$, where C has columns with unit two-norm and D is diagonal, Demmel and
 1268 Veselić [32] proved the bound

$$1269 \quad (13) \quad \frac{|\sigma_i - \hat{\sigma}_i|}{\sigma_i} \leq O(\epsilon)\kappa(C).$$

1271 Crucially, it can be that $\kappa(C) \ll \kappa(A)$, particularly in the instance of a *strongly*
 1272 *scaled* matrix where D is ill-conditioned. If ill-conditioning is artificial, due to poor
 1273 scaling, then one-sided Jacobi will be unaffected by it and will compute even small
 1274 singular values to high relative accuracy. Demmel et al. [30] extended methods of
 1275 computing the SVD with high relative accuracy to a wider class of matrices of the
 1276 form $A = XDY^T$, where D is diagonal, and X and Y are well-conditioned.

1277 Similar results apply for the two-sided Jacobi eigenvalue method with a positive
 1278 definite matrix $A = D^TBD$ [32]. For eigenvalues of an indefinite matrix, though, QR
 1279 iteration may be more accurate than Jacobi [108].

1280 When applied to triangular matrices, Matejaš and Hari [88, 89] proved that the
 1281 two-sided Jacobi SVD method also attains high relative accuracy. One can preprocess
 1282 a general matrix using QRP to obtain such a triangular matrix.

1283 Applied to a bidiagonal matrix, the implicit zero-shift variant of QR iteration and
 1284 the bisection method have been shown to achieve high relative accuracy for all sin-
 1285 gular values [31]. However, the classical reduction from dense to bidiagonal perturbs
 1286 the singular values so the exact singular values of the bidiagonal matrix no longer
 1287 have high relative accuracy for the original matrix A . Hence, *any method* based on
 1288 an initial reduction to bidiagonal will lose relative accuracy for small singular values
 1289 of an ill-conditioned matrix. To address this deficiency, Barlow [7] developed a more
 1290 accurate bidiagonalization, using QRP followed by a Givens rotation based bidiag-
 1291 onalization. Recently, Drmač [43] demonstrated that preprocessing a matrix with
 1292 QRP (LAPACK's `dgeqp3` routine) is sufficient to make a subsequent QR iteration or

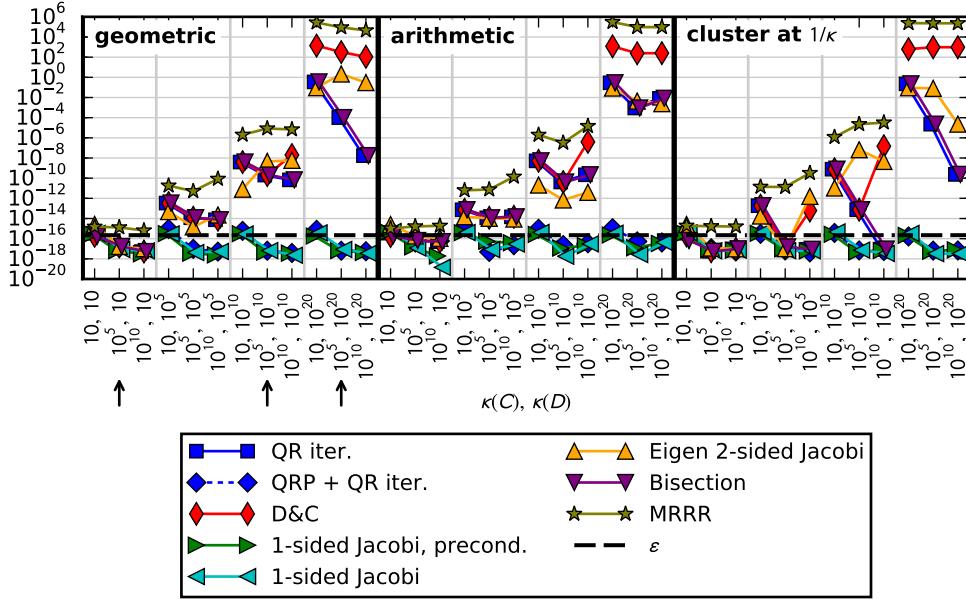


FIG. 31. Maximum relative error in singular values, $\max |\hat{\sigma}_i - \sigma_i| / (\kappa(C)\sigma_i)$, for $i = 1, \dots, 100$, with various test matrices. Figure 32 shows details for three instances indicated by arrows: geometric distribution with $(\kappa(C), \kappa(D)) = (10^5, 10)$; $(10^5, 10^{10})$; $(10^5, 10^{20})$.

1293 bisection have high relative accuracy. (But not divide and conquer, which is not as
 1294 accurate as QR iteration.)

1295 Here we test the accuracy of various methods on matrices with three different
 1296 distributions of singular values: arithmetic, geometric, and a cluster at $1/\kappa(C)$, as
 1297 described in section 2. For each distribution, we generate singular values Σ with
 1298 condition number $\kappa(C)$, scale them so that $\sum \sigma_i^2 = n$, and set $\tilde{C} = U\Sigma V^T$ where
 1299 U and V are random orthogonal matrices from the Haar distribution [106]. To
 1300 satisfy the conditions of (13), we use the method by Davies and Higham [26] to
 1301 make $C = \tilde{C}W$ with columns of unit two-norm, where W is orthogonal. Finally,
 1302 we set $A = CD$, where D is diagonal with entries whose logarithms are random
 1303 uniform on $(\log(1/\kappa(D)), \log(1))$. For each distribution, we set $n = 100$ and vary
 1304 $\kappa(C) \in \{10, 10^5, 10^{10}\}$ and $\kappa(D) \in \{10, 10^5, 10^{10}, 10^{20}\}$. For a reference solution, we
 1305 used MATLAB's [90] variable-precision arithmetic (vpa) with 64 digits.

1306 Figure 31 demonstrates the significant difference between one-sided Jacobi meth-
 1307 ods and bidiagonalization methods (QR iteration, D&C, bisection, MRRR). Both
 1308 one-sided Jacobi methods achieve high relative accuracy for all singular values, at or
 1309 below the dashed line representing machine ϵ . For small scaling, with $\kappa(D) = 10$,
 1310 all methods achieve high accuracy on all the matrices. Most of the bisection meth-
 1311 ods show increased relative errors as the scaling $\kappa(D)$ grows. For $\kappa(D) = 10^{20}$, the
 1312 maximum errors were sometimes larger than 1, i.e., no correct digits in the smallest
 1313 singular values. Among bisection methods, the exception was preprocessing using QR
 1314 with column pivoting, then using QR iteration (QRP + QR iter., blue diamonds),
 1315 which also achieved high relative accuracy, as predicted by Drmač.

1316 QR iteration (blue squares) and bisection (purple down triangles) produce ex-

1317 tremely similar errors, demonstrating that they both accurately compute singular
 1318 values of the bidiagonal matrix, and the error occurs in the reduction to bidiagonal.
 1319 Once the condition number $\kappa(A)$ exceeds $1/\epsilon$, divide and conquer (red diamonds) has
 1320 much worse error than QR iteration. Even with modest scaling, MRRR (stars) has
 1321 the worst error. Eigen's two-sided Jacobi (orange up triangles) also exhibits significant
 1322 error as the scaling increases. Preprocessing with QRP before Eigen's two-sided
 1323 Jacobi (not shown) improved the accuracy, but not to the high relative accuracy of
 1324 one-sided Jacobi. Based on [89], other two-sided Jacobi implementations are expected
 1325 to achieve high relative accuracy.

1326 To explain these results in more detail, we look at three specific cases for the
 1327 geometric distribution with $\kappa(C) = 10^5$ and $\kappa(D) \in \{10, 10^{10}, 10^{20}\}$. In Figure 32,
 1328 the left column shows the actual singular values, in both log and linear scale, while the
 1329 right column shows the relative error in each singular value, σ_i from $i = 1, \dots, 100$. In
 1330 the top row, with minimal scaling ($\kappa(D) = 10$), all the methods achieve high accuracy,
 1331 below ϵ in almost all cases. Eigen has a little higher error for large singular values,
 1332 and MRRR is a little higher for small singular values.

1333 In the middle row, with modest scaling ($\kappa(D) = 10^{10}$), the one-sided Jacobi
 1334 methods and QRP + QR iteration maintain high relative accuracy for all singular
 1335 values. The bidiagonalization methods have high accuracy for the large singular values
 1336 (near σ_1), but the relative error increases for small singular values, losing digits of
 1337 accuracy. Eigen's error also increases.

1338 In the bottom row, with large scaling ($\kappa(D) = 10^{20}$), the error of bidiagonalization
 1339 methods for small singular values grows even more. As seen in the bottom-left graph,
 1340 several methods compute singular values that noticeably diverge from the reference
 1341 solution. For this matrix with $\sigma_{\max} \approx 10^{20}$, MRRR declares all $\sigma_i < 10^7$ to be
 1342 3.27×10^7 , i.e., it cannot resolve smaller singular values. Similarly for D&C, all
 1343 $\sigma_i < 10^3$ are computed as 6.91×10^3 . Eigen also has issues for $\sigma < 10^4$, though it
 1344 does not flatline as MRRR and D&C do. QR iteration and bisection follow the true
 1345 singular values much more closely, but still exhibit significant error for small singular
 1346 values.

1347 **14. Additional test cases.** So far, we have mostly considered the performance
 1348 of random uniform matrices. In this section, we look briefly at additional test cases
 1349 using various distributions of singular values. Our purpose here is to give the reader
 1350 an idea of the variability in performance and how representative the random uniform
 1351 tests are. The distribution of singular values affects the performance of various SVD
 1352 algorithms differently. For QR iteration and divide and conquer, whenever a singular
 1353 value is determined with sufficient accuracy, it can be removed to shrink the problem
 1354 size, a process known as *deflation*, improving the performance. For MRRR, having
 1355 singular values close to one another will cause it to recurse further in the representation
 1356 tree, decreasing its performance [119]. For one-sided Jacobi, matrices that are close
 1357 to orthogonal—i.e., most of the weight of $A^T A$ is near the diagonal—converge faster.

1358 Figure 33 shows results for six methods on various matrices. These all use the
 1359 LAPACK implementations, except MRRR which uses a modification of the bisection
 1360 `dgesvdx` code, as described in section 9. Note that the y -axis scale is different for
 1361 QR iteration, divide and conquer, and MRRR than for Jacobi and bisection. See
 1362 section 2 for a description of the matrix types. For each algorithm, the first, blue bar
 1363 is for a random uniform matrix, matching most of the results elsewhere in this paper.
 1364 The geometric and log-random distributions themselves are similar, so in most cases
 1365 their performance trends are similar, except when using bisection. We see that for

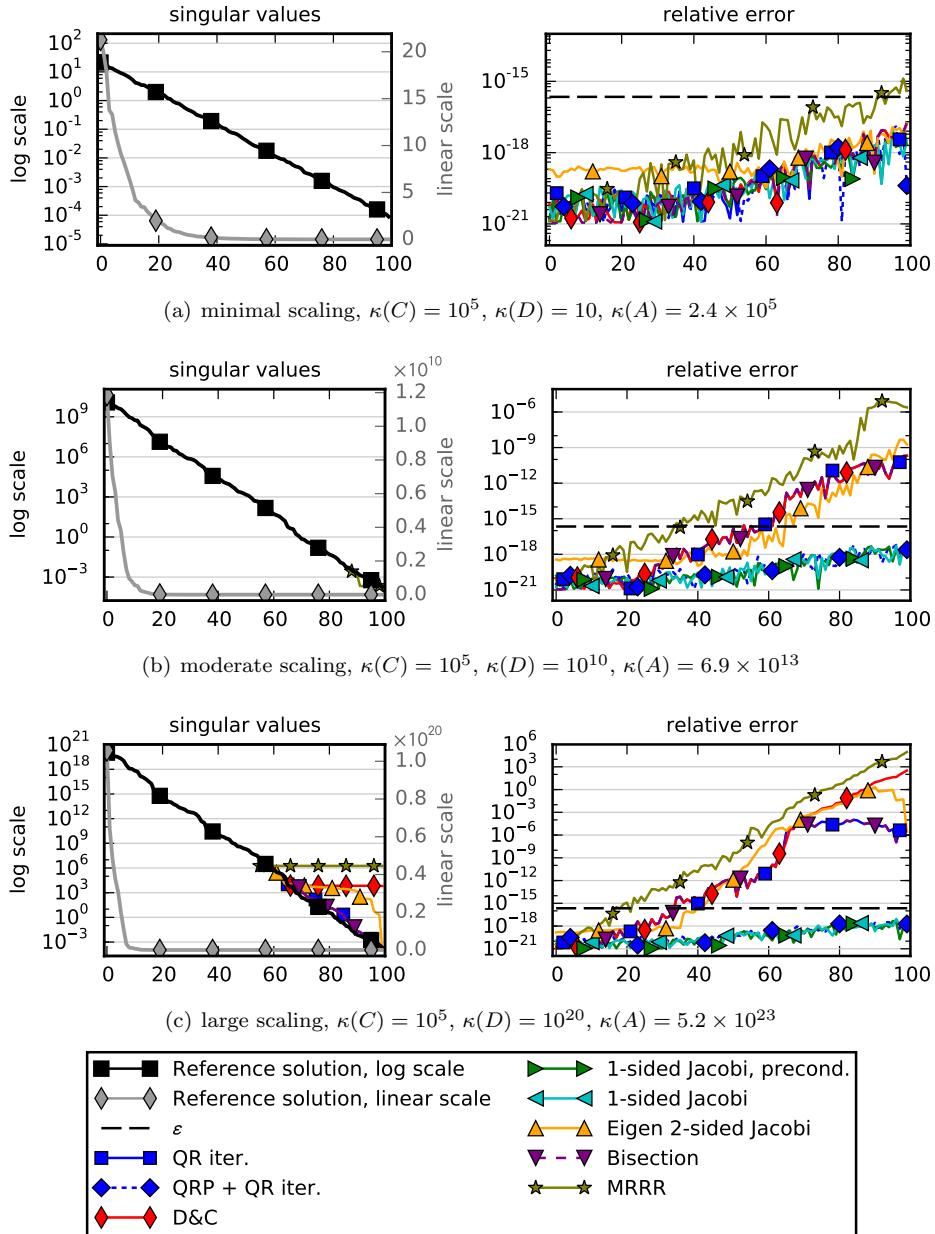


FIG. 32. Singular values of $A = CD$ are plotted twice in left column, once in log scale (black squares), once in linear scale (gray diamonds). In most cases, computed singular values are visually coincident with reference solution (log scale). Right column shows relative error in each singular value, $|\hat{\sigma}_i - \sigma_i| / (\kappa(C)\sigma_i)$. x axis indexes the singular values from largest to smallest, $i = 1, \dots, 100$.

1366 QR iteration, the performance for most matrices is similar to that of random uniform,
1367 with a few being up to 18% slower. For divide and conquer, the arithmetic distribution
1368 (cyan) was up to 24% slower, while log-random (green) was up to 23% faster than
1369 random uniform. The two clusters (red, orange) were 77% and 60% faster, due to
1370 significant deflation. MRRR is more variable, with geometric (purple) and log-random

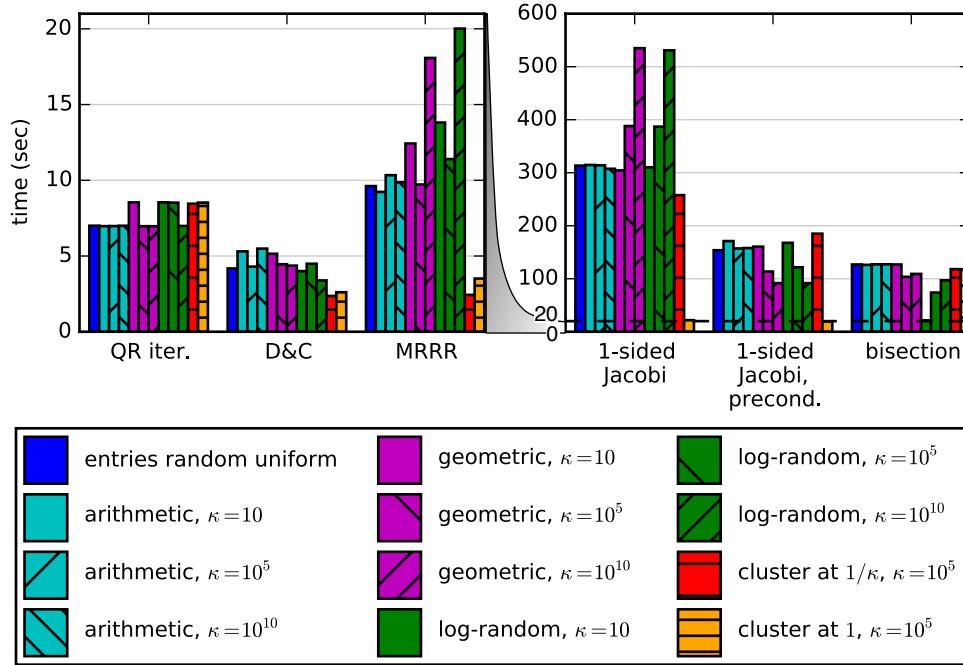


FIG. 33. Time to compute full SVD for $n = 3000$. Note change in y axis; dashed line at $t = 20$ corresponds to y axis in left plot.

1371 (green) being up to 47% and 52% slower on ill-conditioned matrices ($\kappa = 10^{10}$), while
 1372 both clusters of repeated singular values were up to 3 \times faster than random uniform.
 1373 Arithmetic (cyan) was not significantly affected by conditioning.

1374 Because one-sided Jacobi and bisection were significantly slower, they are plotted
 1375 with a different y axis. In all cases, one-sided Jacobi and bisection were slower
 1376 than QR iteration, divide and conquer, and MRRR. The geometric (purple) and log-
 1377 random (green) matrices exhibited opposite behavior for the two Jacobi methods: for
 1378 plain Jacobi, both matrices became slower as the conditioning worsened, while for
 1379 preconditioned Jacobi, both became faster. A cluster at $1/\kappa$ took similar time to
 1380 random. A cluster at 1 was much faster, because $A^T A$ is already nearly diagonal, but
 1381 preconditioning did not further improve it. Bisection was surprisingly 4.9 \times faster for
 1382 a well-conditioned ($\kappa = 10$) log-random matrix, but the speedup decreased for poorer
 1383 conditioning. As we saw earlier in section 8 when computing a subset of vectors,
 1384 clusters were not advantageous.

1385 While the performance does vary for different classes of matrices—sometimes
 1386 substantially—at a high level, our performance conclusions remain valid: divide and
 1387 conquer is the fastest (being tied with MRRR in one case), then QR iteration, then
 1388 MRRR. One-sided Jacobi is the slowest method, with preconditioning generally im-
 1389 proving its speed, often by a factor of 2 \times or more. For computing all vectors, bisec-
 1390 tion is also slow; its main advantage is in computing a subset of vectors, as previously
 1391 shown in section 8.

1392 **15. Conclusions.** As we have seen, algorithms to compute the SVD have con-
 1393 tinually evolved to address changes in computer hardware design, as well as advan-
 1394 tments in mathematics. Early implementations such as EISPACK demonstrated that

computing the SVD stably was feasible. Later implementations focused on improving the performance, first by using Level 1 BLAS for vector computers, then by reformulating the algorithm for Level 3 BLAS to address the emergence of cache-based memory hierarchies. More recently, a two-stage algorithm shifted even more operations from Level 2 to Level 3 BLAS. These changes have addressed the growing gap between memory bandwidth and computational speed, as well as enabled greater use of parallel hardware. Implementations have also taken advantage of different architectures such as distributed memory computers and accelerators. Mathematical advancements have been important in reducing the number of operations performed. For tall-skinny problems, using an initial QR factorization can eliminate a quarter to half of the operations. For square matrices, the divide and conquer algorithm reduces operations by nearly half. For Jacobi methods, preconditioning has been vital to improve convergence, while at the same time making computation of singular vectors more efficient. Block Jacobi methods with dynamic selection of subproblems have become competitive with some bidiagonalization methods. Improvements in algorithms used to preprocess a matrix, such as using a CAQR factorization [29] for tall-skinny matrices, or future improvements to QRP methods, are immediately applicable to benefit SVD performance.

As we build the next generation of linear algebra software targeting exascale computers [77], the goal is to integrate these techniques—such as the two-stage reduction to bidiagonal, accelerators, and distributed computing—into a scalable SVD solver. While the techniques have been demonstrated to work, the challenge is being able to hide communication latencies in large distributed machines. Bottlenecks due to Amdahl’s law, such as solving the bidiagonal SVD, will also be crucial to resolve. Improving algorithms to remove communication and memory bandwidth limitations becomes critically important.

For certain classes of matrices that are strongly scaled, classical methods based on reduction to bidiagonal will not accurately compute small singular values. In these cases, one should turn to Jacobi methods or preprocessing the matrix using QR factorization with column pivoting (QRP) to attain high relative accuracy.

We have focused on solving dense systems. There are of course different techniques for solving SVD problems with sparse linear systems. Also, if one is concerned with only an approximate, low-rank solution, then using a randomized SVD algorithm [99] may be another avenue to pursue. This is often the case for huge systems arising from big data problems.

Here we have compared implementations on a common, modern architecture. To give some historical perspective, in 1977, EISPACK took 0.79 seconds (1.7 Mflop/s) to compute singular values for $n = 80$ on an IBM 370/195 [105]. Today, the same EISPACK code achieves 0.74 Gflop/s on large problems, yielding over two orders-of-magnitude advancement in single core hardware speed. On top of this, we have shown an additional two orders-of-magnitude improvement going from EISPACK to PLASMA (146 Gflop/s) on a multicore architecture, and four orders of magnitude to DPLASMA (6.8 Tflop/s) on a distributed-memory machine—while moving from solving systems of dimension 100 to over 100,000—yielding over six orders-of-magnitude performance improvement in 40 years.

Acknowledgments. We thank Martin Bečka, Gabriel Okša, and Marián Vajteršic for use of their block Jacobi code; Osni Marques for assistance with the MRRR code; and the anonymous reviewers for feedback to improve the quality and scope of this work.

1444 **References.**

1445 [1] E. AGULLO, B. HADRI, H. LTAIEF, AND J. DONGARRRA, *Comparative study*
 1446 *of one-sided factorizations with multiple software packages on multi-core hard-*
 1447 *ware*, in Proceedings of the Conference on High Performance Computing Net-
 1448 *working, Storage and Analysis (SC'09)*, ACM, 2009, p. 20, [https://doi.org/10.](https://doi.org/10.1145/1654059.1654080)
 1449 [1145/1654059.1654080](https://doi.org/10.1145/1654059.1654080).

1450 [2] E. ANDERSON, Z. BAI, C. BISCHOF, S. BLACKFORD, J. DONGARRA,
 1451 J. DU CROZ, A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, AND
 1452 D. SORENSEN, *LAPACK users' guide*, SIAM, Philadelphia, third ed., 1999,
 1453 <https://doi.org/10.1137/1.9780898719604>.

1454 [3] H. ANDREWS AND C. PATTERSON, *Singular value decomposition (SVD) image*
 1455 *coding*, IEEE Transactions on Communications, 24 (1976), pp. 425–432, <https://doi.org/10.1109/TCOM.1976.1093309>.

1456 [4] P. ARBENZ AND G. GOLUB, *On the spectral decomposition of Hermitian matri-*
 1457 *ces modified by low rank perturbations with applications*, SIAM Journal on Ma-
 1458 *trix Analysis and Applications*, 9 (1988), pp. 40–58, <https://doi.org/10.1137/0609004>.

1459 [5] P. ARBENZ AND I. SLAPNIČAR, *An analysis of parallel implementations of the*
 1460 *block-Jacobi algorithm for computing the SVD*, in Proceedings of the 17th In-
 1461 *ternational Conference on Information Technology Interfaces ITI*, vol. 95, 1995,
 1462 pp. 13–16, <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.53.4595>.

1463 [6] G. BALLARD, J. DEMMEL, AND N. KNIGHT, *Avoiding communication in suc-*
 1464 *cessive band reduction*, ACM Transactions on Parallel Computing, 1 (2015),
 1465 p. 11, <https://doi.org/10.1145/2686877>.

1466 [7] J. L. BARLOW, *More accurate bidiagonal reduction for computing the singular*
 1467 *value decomposition*, SIAM Journal on Matrix Analysis and Applications, 23
 1468 (2002), pp. 761–798, <https://doi.org/10.1137/S0895479898343541>.

1469 [8] M. BEČKA, G. OKŠA, AND M. VAJTERŠIC, *Dynamic ordering for a parallel*
 1470 *block-Jacobi SVD algorithm*, Parallel Computing, 28 (2002), pp. 243–262, [https://doi.org/10.1016/S0167-8191\(01\)00138-7](https://doi.org/10.1016/S0167-8191(01)00138-7).

1471 [9] M. BEČKA, G. OKŠA, AND M. VAJTERŠIC, *New dynamic orderings for the*
 1472 *parallel one-sided block-Jacobi SVD algorithm*, Parallel Processing Letters, 25
 1473 (2015), p. 1550003, <https://doi.org/10.1142/S0129626415500036>.

1474 [10] M. BEČKA, G. OKŠA, M. VAJTERŠIC, AND L. GRIGORI, *On iterative QR pre-*
 1475 *processing in the parallel block-Jacobi SVD algorithm*, Parallel Computing, 36
 1476 (2010), pp. 297–307, <https://doi.org/10.1016/j.parco.2009.12.013>.

1477 [11] M. BEČKA AND M. VAJTERŠIC, *Block-Jacobi SVD algorithms for distributed*
 1478 *memory systems I: hypercubes and rings*, Parallel Algorithms and Application,
 1479 13 (1999), pp. 265–287, <https://doi.org/10.1080/10637199808947377>.

1480 [12] M. BEČKA AND M. VAJTERŠIC, *Block-Jacobi SVD algorithms for distributed*
 1481 *memory systems II: meshes*, Parallel Algorithms and Application, 14 (1999),
 1482 pp. 37–56, <https://doi.org/10.1080/10637199808947370>.

1483 [13] C. BISCHOF, B. LANG, AND X. SUN, *Algorithm 807: The SBR Toolbox—*
 1484 *software for successive band reduction*, ACM Transactions on Mathematical
 1485 Software (TOMS), 26 (2000), pp. 602–616, <https://doi.org/10.1145/365723.365736>.

1486 [14] C. BISCHOF AND C. VAN LOAN, *The WY representation for products of*
 1487 *Householder matrices*, SIAM Journal on Scientific and Statistical Computing, 8
 1488 (1987), pp. 2–13, <https://doi.org/10.1137/0908009>.

1489 [15] C. H. BISCHOF, *Computing the singular value decomposition on a distributed*

1494 *system of vector processors*, Parallel Computing, 11 (1989), pp. 171–186, [https://doi.org/10.1016/0167-8191\(89\)90027-6](https://doi.org/10.1016/0167-8191(89)90027-6).

1495 [16] L. S. BLACKFORD, J. CHOI, A. CLEARY, E. D’AZEVEDO, J. DEMMEL,
 1496 I. DHILLON, J. DONGARRA, S. HAMMARLING, G. HENRY, A. PETITET,
 1497 ET AL., *ScaLAPACK users’ guide*, SIAM, Philadelphia, 1997, <https://doi.org/10.1137/1.9780898719642>.

1500 [17] L. S. BLACKFORD, A. PETITET, R. POZO, K. REMINGTON, R. C. WHALEY,
 1501 J. DEMMEL, J. DONGARRA, I. DUFF, S. HAMMARLING, G. HENRY, ET AL.,
 1502 *An updated set of basic linear algebra subprograms (BLAS)*, ACM Transactions
 1503 on Mathematical Software (TOMS), 28 (2002), pp. 135–151, <https://doi.org/10.1145/567806.567807>.

1505 [18] G. BOSILCA, A. BOUTEILLER, A. DANALIS, M. FAVERGE, A. HAIDAR,
 1506 T. HERAULT, J. KURZAK, J. LANGOU, P. LEMARINIER, H. LTAIEF, ET AL.,
 1507 *Flexible development of dense linear algebra algorithms on massively parallel*
 1508 *architectures with DPLASMA*, in 2011 IEEE International Symposium on Parallel
 1509 and Distributed Processing Workshops and Phd Forum (IPDPSW), IEEE,
 1510 2011, pp. 1432–1441, <https://doi.org/10.1109/IPDPS.2011.299>.

1511 [19] G. BOSILCA, A. BOUTEILLER, A. DANALIS, T. HERAULT, P. LEMARINIER,
 1512 AND J. DONGARRA, *DAGuE: A generic distributed DAG engine for high perfor-*
 1513 *mance computing*, Parallel Computing, 38 (2012), pp. 37–51, <https://doi.org/10.1016/j.parco.2011.10.003>.

1515 [20] W. H. BOUKARAM, G. TURKIYYAH, H. LTAIEF, AND D. E. KEYES, *Batched*
 1516 *QR and SVD algorithms on GPUs with applications in hierarchical matrix com-*
 1517 *pression*, Parallel Computing, (2017), <https://doi.org/10.1016/j.parco.2017.09.001>.

1519 [21] R. P. BRENT AND F. T. LUK, *The solution of singular-value and symmetric*
 1520 *eigenvalue problems on multiprocessor arrays*, SIAM Journal on Scientific and
 1521 Statistical Computing, 6 (1985), pp. 69–84, <https://doi.org/10.1137/0906007>.

1522 [22] R. P. BRENT, F. T. LUK, AND C. VAN LOAN, *Computation of the singular*
 1523 *value decomposition using mesh-connected processors*, Journal of VLSI and com-
 1524 *puter systems*, 1 (1985), pp. 242–270, <http://maths-people.anu.edu.au/~brent/pd/rpb080i.pdf>.

1526 [23] T. F. CHAN, *An improved algorithm for computing the singular value decompo-*
 1527 *sition*, ACM Transactions on Mathematical Software (TOMS), 8 (1982), pp. 72–
 1528 83, <https://doi.org/10.1145/355984.355990>.

1529 [24] J. CHOI, J. DONGARRA, AND D. W. WALKER, *The design of a par-*
 1530 *allel dense linear algebra software library: reduction to Hessenberg, tri-*
 1531 *diagonal, and bidiagonal form*, Numerical Algorithms, 10 (1995), pp. 379–399,
 1532 <https://doi.org/10.1007/BF02140776>.

1533 [25] J. J. M. CUPPEN, *A divide and conquer method for the symmetric tridiagonal*
 1534 *eigenproblem*, Numerische Mathematik, 36 (1980), pp. 177–195, <https://doi.org/10.1007/BF01396757>.

1536 [26] P. I. DAVIES AND N. J. HIGHAM, *Numerically stable generation of correlation*
 1537 *matrices and their factors*, BIT Numerical Mathematics, 40 (2000), pp. 640–651,
 1538 <https://doi.org/10.1023/A:102238421>.

1539 [27] P. P. M. DE RIJK, *A one-sided Jacobi algorithm for computing the singular*
 1540 *value decomposition on a vector computer*, SIAM Journal on Scientific and Sta-
 1541 *tistical Computing, 10 (1989), pp. 359–371, <https://doi.org/10.1137/0910023>*.

1542 [28] S. DEERWESTER, S. T. DUMAIS, G. W. FURNAS, T. K. LANDAUER, AND
 1543 R. HARSHMAN, *Indexing by latent semantic analysis*, Journal of the Ameri-

1544 can society for information science, 41 (1990), p. 391, [https://doi.org/10.1002/\(SICI\)1097-4571\(199009\)41:6<391::AID-ASI1>3.0.CO;2-9](https://doi.org/10.1002/(SICI)1097-4571(199009)41:6<391::AID-ASI1>3.0.CO;2-9).

1545 [29] J. DEMMEL, L. GRIGORI, M. HOEMMEN, AND J. LANGOU, *Communication-optimal parallel and sequential QR and LU factorizations*, SIAM Journal of Scientific Computing, 34 (2012), pp. A206–A239, <https://doi.org/10.1137/080731992>.

1546 [30] J. DEMMEL, M. GU, S. EISENSTAT, I. SLAPNIČAR, K. VESELIĆ, AND Z. DRMAČ, *Computing the singular value decomposition with high relative accuracy*, Linear Algebra and its Applications, 299 (1999), pp. 21–80, [https://doi.org/10.1016/S0024-3795\(99\)00134-2](https://doi.org/10.1016/S0024-3795(99)00134-2).

1547 [31] J. DEMMEL AND W. KAHAN, *Accurate singular values of bidiagonal matrices*, SIAM Journal on Scientific and Statistical Computing, 11 (1990), pp. 873–912, <https://doi.org/10.1137/0911052>.

1548 [32] J. DEMMEL AND K. VESELIĆ, *Jacobi's method is more accurate than QR*, SIAM Journal on Matrix Analysis and Applications, 13 (1992), pp. 1204–1245, <https://doi.org/10.1137/0613074>.

1549 [33] J. W. DEMMEL, *Applied Numerical Linear Algebra*, SIAM, Philadelphia, 1997, <https://doi.org/10.1137/1.9781611971446>.

1550 [34] J. W. DEMMEL, I. DHILLON, AND H. REN, *On the correctness of some bisection-like parallel eigenvalue algorithms in floating point arithmetic*, Electronic Transactions on Numerical Analysis, 3 (1995), pp. 116–149, <http://emis.ams.org/journals/ETNA/vol.3.1995/pp116-149.dir/pp116-149.pdf>.

1551 [35] I. S. DHILLON, *A New $O(n^2)$ Algorithm for the Symmetric Tridiagonal Eigenvalue/Eigenvector Problem*, PhD thesis, EECS Department, University of California, Berkeley, 1997, <http://www.dtic.mil/docs/citations/ADA637073>.

1552 [36] I. S. DHILLON AND B. N. PARLETT, *Multiple representations to compute orthogonal eigenvectors of symmetric tridiagonal matrices*, Linear Algebra and its Applications, 387 (2004), pp. 1–28, <https://doi.org/10.1016/j.laa.2003.12.028>.

1553 [37] I. S. DHILLON AND B. N. PARLETT, *Orthogonal eigenvectors and relative gaps*, SIAM Journal on Matrix Analysis and Applications, 25 (2004), pp. 858–899, <https://doi.org/10.1137/S0895479800370111>.

1554 [38] I. S. DHILLON, B. N. PARLETT, AND C. VÖMEL, *The design and implementation of the MRRI algorithm*, ACM Transactions on Mathematical Software (TOMS), 32 (2006), pp. 533–560, <https://doi.org/10.1145/1186785.1186788>.

1555 [39] J. DONGARRA, J. R. BUNCH, C. B. MOLER, AND G. W. STEWART, *LINPACK users' guide*, SIAM, Philadelphia, 1979, <https://doi.org/10.1137/1.9781611971811>.

1556 [40] J. DONGARRA, J. DU CROZ, S. HAMMARLING, AND I. S. DUFF, *A set of level 3 basic linear algebra subprograms*, ACM Transactions on Mathematical Software (TOMS), 16 (1990), pp. 1–17, <https://doi.org/10.1145/77626.79170>.

1557 [41] J. DONGARRA, J. DU CROZ, S. HAMMARLING, AND R. J. HANSON, *An extended set of FORTRAN basic linear algebra subprograms*, ACM Transactions on Mathematical Software (TOMS), 14 (1988), pp. 1–17, <https://doi.org/10.1145/42288.42291>.

1558 [42] J. DONGARRA, D. C. SORENSEN, AND S. J. HAMMARLING, *Block reduction of matrices to condensed forms for eigenvalue computations*, Journal of Computational and Applied Mathematics, 27 (1989), pp. 215–227, [https://doi.org/10.1016/0377-0427\(89\)90367-1](https://doi.org/10.1016/0377-0427(89)90367-1). Special Issue on Parallel Algorithms for Numerical Linear Algebra.

1559 [43] Z. DRMAČ, *Algorithm 977: A QR-preconditioned QR SVD method for comput-*

1594 *ing the SVD with high accuracy*, ACM Transactions on Mathematical Software
 1595 (TOMS), 44 (2017), p. 11, <https://doi.org/10.1145/3061709>.

1596 [44] Z. DRMAČ AND K. VESELIĆ, *New fast and accurate Jacobi SVD algorithm, I*,
 1597 SIAM Journal on Matrix Analysis and Applications, 29 (2008), pp. 1322–1342,
 1598 <https://doi.org/10.1137/050639193>.

1599 [45] Z. DRMAČ AND K. VESELIĆ, *New fast and accurate Jacobi SVD algorithm, II*,
 1600 SIAM Journal on Matrix Analysis and Applications, 29 (2008), pp. 1343–1362,
 1601 <https://doi.org/10.1137/05063920X>.

1602 [46] P. EBERLEIN, *On one-sided Jacobi methods for parallel computation*, SIAM
 1603 Journal on Algebraic Discrete Methods, 8 (1987), pp. 790–796, <https://doi.org/10.1137/0608064>.

1604 [47] EIGEN, *Eigen 3.3.3*, 2017, <http://eigen.tuxfamily.org/>.

1605 [48] K. V. FERNANDO AND B. N. PARLETT, *Accurate singular values and differ-
 1606 ential qd algorithms*, Numerische Mathematik, 67 (1994), pp. 191–229, <https://doi.org/10.1007/s002110050024>.

1607 [49] G. E. FORSYTHE AND P. HENRICI, *The cyclic Jacobi method for computing
 1608 the principal values of a complex matrix*, Transactions of the American Mathe-
 1609 matical Society, 94 (1960), pp. 1–23, <https://doi.org/10.2307/1993275>.

1610 [50] B. S. GARBOW, J. M. BOYLE, C. B. MOLER, AND J. DONGARRA, *Ma-
 1611 trix eigensystem routines – EISPACK guide extension*, vol. 51 of Lecture
 1612 Notes in Computer Science, Springer, Berlin, 1977, <https://doi.org/10.1007/3-540-08254-9>.

1613 [51] M. GATES, S. TOMOV, AND J. DONGARRAA, *Accelerating the SVD two stage
 1614 bidiagonal reduction and divide and conquer using GPUs*, Parallel Computing,
 1615 74 (2018), pp. 3–18, <https://doi.org/10.1016/j.parco.2017.10.004>.

1616 [52] G. GOLUB, *Some modified matrix eigenvalue problems*, SIAM Review, 15
 1617 (1973), pp. 318–334, <https://doi.org/10.1137/1015032>.

1618 [53] G. GOLUB AND W. KAHAN, *Calculating the singular values and pseudo-inverse
 1619 of a matrix*, SIAM Journal on Numerical Analysis (Series B), 2 (1965), pp. 205–
 1620 224, <https://doi.org/10.1137/0702016>.

1621 [54] G. GOLUB AND C. REINSCH, *Singular value decomposition and least squares
 1622 solutions*, Numerische Mathematik, 14 (1970), pp. 403–420, <https://doi.org/10.1007/BF02163027>.

1623 [55] B. GROSSER AND B. LANG, *Efficient parallel reduction to bidiagonal
 1624 form*, Parallel Computing, 25 (1999), pp. 969–986, [https://doi.org/10.1016/S0167-8191\(99\)00041-1](https://doi.org/10.1016/S0167-8191(99)00041-1).

1625 [56] M. GU, J. DEMMEL, AND I. DHILLON, *Efficient computation of the singular
 1626 value decomposition with applications to least squares problems*, Tech. Report
 1627 LBL-36201, Lawrence Berkeley Laboratory, September 29 1994, http://www.cs.utexas.edu/users/inderjit/public_papers/least_squares.pdf.

1628 [57] M. GU AND S. C. EISENSTAT, *A divide and conquer algorithm for the
 1629 bidiagonal SVD*, Tech. Report YALEU/DCS/TR-933, Department of Com-
 1630 puter Science, Yale University, November 1992, <http://cpsc.yale.edu/research/technical-reports/1992-technical-reports>.

1631 [58] M. GU AND S. C. EISENSTAT, *A stable and efficient algorithm for the rank-
 1632 one modification of the symmetric eigenproblem*, SIAM Journal on Matrix
 1633 Analysis and Applications, 15 (1994), pp. 1266–1276, <https://doi.org/10.1137/S089547989223924X>.

1634 [59] M. GU AND S. C. EISENSTAT, *A divide-and-conquer algorithm for the bidi-
 1635 agonal SVD*, SIAM Journal on Matrix Analysis and Applications, 16 (1995),

1644 pp. 79–92, <https://doi.org/10.1137/S0895479892242232>.

1645 [60] A. HAIDAR, J. KURZAK, AND P. LUSCZEK, *An improved parallel singular*
 1646 *value algorithm and its implementation for multicore hardware*, in Proceedings
 1647 of the International Conference on High Performance Computing, Network-
 1648 Working, Storage and Analysis (SC’13), ACM, 2013, p. 90, <https://doi.org/10.1145/2503210.2503292>.

1649 [61] A. HAIDAR, H. LTAIEF, AND J. DONGARRA, *Parallel reduction to condensed*
 1650 *forms for symmetric eigenvalue problems using aggregated fine-grained and*
 1651 *memory-aware kernels*, in Proceedings of 2011 International Conference for High
 1652 Performance Computing, Networking, Storage and Analysis (SC’11), ACM,
 1653 2011, pp. 8:1–8:11, <https://doi.org/10.1145/2063384.2063394>.

1654 [62] A. HAIDAR, H. LTAIEF, P. LUSCZEK, AND J. DONGARRA, *A comprehen-*
 1655 *sive study of task coalescing for selecting parallelism granularity in a two-stage*
 1656 *bidiagonal reduction*, in 2012 IEEE 26th International Parallel and Distributed
 1657 Processing Symposium (IPDPS), IEEE, 2012, pp. 25–35, <https://doi.org/10.1109/IPDPS.2012.13>.

1658 [63] S. HAMMARLING, *A note on modifications to the Givens plane rotation*, IMA
 1659 *Journal of Applied Mathematics*, 13 (1974), pp. 215–218, <https://doi.org/10.1093/imamat/13.2.215>.

1660 [64] V. HARI, *Accelerating the SVD block-Jacobi method*, Computing, 75 (2005),
 1661 pp. 27–53, <https://doi.org/10.1007/s00607-004-0113-z>.

1662 [65] V. HARI AND J. MATEJAŠ, *Accuracy of two SVD algorithms for 2×2 triangular*
 1663 *matrices*, Applied Mathematics and Computation, 210 (2009), pp. 232–257,
 1664 <https://doi.org/10.1016/j.amc.2008.12.086>.

1665 [66] V. HARI AND K. VESELIĆ, *On Jacobi methods for singular value decomposi-*
 1666 *tions*, SIAM Journal on Scientific and Statistical Computing, 8 (1987), pp. 741–
 1667 754, <https://doi.org/10.1137/0908064>.

1668 [67] M. HEATH, A. LAUB, C. PAIGE, AND R. WARD, *Computing the singular*
 1669 *value decomposition of a product of two matrices*, SIAM Journal on Scientific
 1670 and Statistical Computing, 7 (1986), pp. 1147–1159, <https://doi.org/10.1137/0907078>.

1671 [68] M. R. HESTENES, *Inversion of matrices by biorthogonalization and related re-*
 1672 *sults*, Journal of the Society for Industrial and Applied Mathematics, 6 (1958),
 1673 pp. 51–90, <https://doi.org/10.1137/0106005>.

1674 [69] G. W. HOWELL, J. W. DEMMEL, C. T. FULTON, S. HAMMARLING, AND
 1675 K. MARMOL, *Cache efficient bidiagonalization using BLAS 2.5 operators*, ACM
 1676 Transactions on Mathematical Software (TOMS), 34 (2008), p. 14, <https://doi.org/10.1145/1356052.1356055>.

1677 [70] IBM CORPORATION, *ESSL Guide and Reference*, 2016, <http://publib.boulder.ibm.com/epubs/pdf/a2322688.pdf>.

1678 [71] INTEL CORPORATION, *User’s Guide for Intel Math Kernel Library for Linux*
 1679 *OS*, 2015, <http://software.intel.com/en-us/mkl-for-linux-userguide>.

1680 [72] I. C. F. IPSEN, *Computing an eigenvector with inverse iteration*, SIAM Review,
 1681 (2006), pp. 254–291, <https://doi.org/10.1137/S0036144496300773>.

1682 [73] C. G. J. JACOBI, *Über ein leichtes verfahren die in der theorie der*
 1683 *säcularstörungen vorkommenden gleichungen numerisch aufzulösen.*, Journal
 1684 für die reine und angewandte Mathematik, 30 (1846), pp. 51–94, <http://eudml.org/doc/147275>.

1685 [74] E. JESSUP AND D. SORENSEN, *A divide and conquer algorithm for computing*
 1686 *the singular value decomposition*, in Proceedings of the Third SIAM Conference

1694 on Parallel Processing for Scientific Computing, Philadelphia, PA, 1989, SIAM,
 1695 pp. 61–66.

1696 [75] W. KAHAN, *Accurate eigenvalues of a symmetric tri-diagonal matrix*, tech. report,
 1697 Stanford University, Stanford, CA, USA, 1966, <http://www.dtic.mil/docs/citations/AD0638796>.

1698 [76] E. KOGBETLIANTZ, *Solution of linear equations by diagonalization of coefficients matrix*, Quarterly of Applied Mathematics, 13 (1955), pp. 123–132,
 1700 <http://www.ams.org/journals/qam/1955-13-02/S0033-569X-1955-88795-9/S0033-569X-1955-88795-9.pdf>.

1702 [77] J. KURZAK, P. WU, M. GATES, I. YAMAZAKI, P. LUSZCZEK, G. RAGGHI-
 1704 ANTI, AND J. DONGARRA, *Designing SLATE: Software for linear algebra targeting exascale*, SLATE Working Note 3, Innovative Computing Laboratory, Uni-
 1706 versity of Tennessee, Sep 2017, <http://www.icl.utk.edu/publications/swan-003>.

1708 [78] B. LANG, *Parallel reduction of banded matrices to bidiagonal form*, Parallel
 1709 Computing, 22 (1996), pp. 1–18, [https://doi.org/10.1016/0167-8191\(95\)00064-X](https://doi.org/10.1016/0167-8191(95)00064-X).

1710 [79] C. L. LAWSON, R. J. HANSON, D. R. KINCAID, AND F. T. KROGH, *Basic
 1711 linear algebra subprograms for FORTRAN usage*, ACM Transactions on Math-
 1712 ematical Software (TOMS), 5 (1979), pp. 308–323, <https://doi.org/10.1145/355841.355847>.

1713 [80] R.-C. LI, *Solving secular equations stably and efficiently*, Tech. Report
 1714 UCB//CSD-94-851, University of California Berkeley, Computer Science Di-
 1716 vision, 1994, <http://www.netlib.org/lapack/lawns/>. Also: LAPACK Working
 1717 Note 89.

1718 [81] S. LI, M. GU, L. CHENG, X. CHI, AND M. SUN, *An accelerated divide-and-
 1719 conquer algorithm for the bidiagonal SVD problem*, SIAM Journal on Matrix
 1720 Analysis and Applications, 35 (2014), pp. 1038–1057, <https://doi.org/10.1137/130945995>.

1722 [82] H. LTAIEF, J. KURZAK, AND J. DONGARRA, *Parallel two-sided matrix re-
 1723 duction to band bidiagonal form on multicore architectures*, IEEE Transac-
 1724 tions on Parallel and Distributed Systems, 21 (2010), pp. 417–423, <https://doi.org/10.1109/TPDS.2009.79>.

1726 [83] H. LTAIEF, P. LUSZCZEK, AND J. DONGARRA, *High-performance bidiagonal
 1727 reduction using tile algorithms on homogeneous multicore architectures*, ACM
 1728 Transactions on Mathematical Software (TOMS), 39 (2013), pp. 16:1–16:22,
 1729 <https://doi.org/10.1145/2450153.2450154>.

1730 [84] F. T. LUK, *Computing the singular-value decomposition on the ILLIAC IV*,
 1731 ACM Transactions on Mathematical Software (TOMS), 6 (1980), pp. 524–539,
 1732 <https://doi.org/10.1145/355921.355925>.

1733 [85] F. T. LUK AND H. PARK, *On parallel Jacobi orderings*, SIAM Journal on
 1734 Scientific and Statistical Computing, 10 (1989), pp. 18–26, <https://doi.org/10.1137/0910002>.

1736 [86] O. MARQUES AND P. B. VASCONCELOS, *Computing the bidiagonal SVD
 1737 through an associated tridiagonal eigenproblem*, in International Conference
 1738 on Vector and Parallel Processing (VECPAR), Springer, 2016, pp. 64–74,
 1739 https://doi.org/10.1007/978-3-319-61982-8_8.

1740 [87] W. F. MASCARENHAS, *On the convergence of the Jacobi method for arbi-
 1741 trary orderings*, SIAM Journal on Matrix Analysis and Applications, 16 (1995),
 1742 pp. 1197–1209, <https://doi.org/10.1137/S0895479890179631>.

1743 [88] J. MATEJAŠ AND V. HARI, *Accuracy of the Kogbetliantz method for scaled diag-*

1744 *onally dominant triangular matrices*, Applied Mathematics and Computation,
 1745 217 (2010), pp. 3726–3746, <https://doi.org/10.1016/j.amc.2010.09.020>.

1746 [89] J. MATEJAŠ AND V. HARI, *On high relative accuracy of the Kogbetlantz*
 1747 *method*, Linear Algebra and its Applications, 464 (2015), pp. 100–129, <https://doi.org/10.1016/j.laa.2014.02.024>.

1748 [90] MATHWORKS, *MATLAB*, 2017, <http://www.mathworks.com/products/matlab.html>.

1749 [91] J. D. MCCALPIN, *A survey of memory bandwidth and machine balance in*
 1750 *current high performance computers*, IEEE Computer Society Technical Com-
 1751 mittee on Computer Architecture (TCCA) Newsletter, (1995), pp. 19–25,
 1752 http://tab.computer.org/tcca/NEWS/DEC95/dec95_mccalpin.ps.

1753 [92] B. MOORE, *Principal component analysis in linear systems: Controllability,*
 1754 *observability, and model reduction*, IEEE Transactions on Automatic Control,
 1755 26 (1981), pp. 17–32, <https://doi.org/10.1109/TAC.1981.1102568>.

1756 [93] MPI FORUM, *MPI: A message-passing interface standard: Version 3.1*, June
 1757 2015, <http://www.mpi-forum.org/>.

1758 [94] NVIDIA CORPORATION, *CUDA Toolkit 7.0*, March 2015, <http://developer.nvidia.com/cuda-zone>.

1759 [95] G. OKŠA AND M. VAJTERŠIC, *Efficient pre-processing in the parallel block-*
 1760 *Jacobi SVD algorithm*, Parallel Computing, 32 (2006), pp. 166–176, <https://doi.org/10.1016/j.parco.2005.06.006>.

1761 [96] OPENBLAS, *OpenBLAS User Manual*, 2016, <http://www.openblas.net/>.

1762 [97] B. N. PARLETT, *The new qd algorithms*, Acta Numerica, 4 (1995), pp. 459–491,
 1763 <https://doi.org/10.1017/S0962492900002580>.

1764 [98] B. N. PARLETT AND I. S. DHILLON, *Fernando’s solution to Wilkinson’s prob-*
 1765 *lem: An application of double factorization*, Linear Algebra and its Applications,
 1766 267 (1997), pp. 247–279, [https://doi.org/10.1016/S0024-3795\(97\)80053-5](https://doi.org/10.1016/S0024-3795(97)80053-5).

1767 [99] V. ROKHLIN, A. SZLAM, AND M. TYGERT, *A randomized algorithm for prin-*
 1768 *cipal component analysis*, SIAM Journal on Matrix Analysis and Applications,
 1769 31 (2009), pp. 1100–1124, <https://doi.org/10.1137/080736417>.

1770 [100] H. RUTISHAUSER, *Derquotienten-differenzen-algorithmus*, Zeitschrift für ange-
 1771 *wandte Mathematik und Physik ZAMP*, 5 (1954), pp. 233–251, <https://doi.org/10.1007/BF01600331>.

1772 [101] H. RUTISHAUSER, *Solution of eigenvalue problems with the LR-transformation*,
 1773 National Bureau of Standards Applied Mathematics Series, 49 (1958), pp. 47–
 1774 81.

1775 [102] H. RUTISHAUSER, *The Jacobi method for real symmetric matrices*, in Handbook
 1776 for Automatic Computation: Volume II: Linear Algebra, vol. 186 of Grundlehren
 1777 Der Mathematischen Wissenschaften, Springer-Verlag, New York, NY, 1971,
 1778 pp. 202–211, <https://doi.org/10.1007/978-3-642-86940-2>.

1779 [103] A. H. SAMEH, *On Jacobi and Jacobi-like algorithms for a parallel computer*,
 1780 Mathematics of Computation, 25 (1971), pp. 579–590, <https://doi.org/10.1090/S0025-5718-1971-0297131-6>.

1781 [104] R. SCHREIBER AND C. VAN LOAN, *A storage-efficient WY representation for*
 1782 *products of Householder transformations*, SIAM Journal on Scientific and Sta-
 1783 *tistical Computing, 10 (1989), pp. 53–57, <https://doi.org/10.1137/0910005>*.

1784 [105] B. T. SMITH, J. M. BOYLE, J. DONGARRA, B. S. GARBOW, Y. IKEBE, V. C.
 1785 KLEMA, AND C. B. MOLER, *Matrix Eigensystem Routines – EISPACK Guide*,
 1786 *Second Edition*, vol. 6 of Lecture Notes in Computer Science, Springer, Berlin,
 1787 1976, <https://doi.org/10.1007/3-540-07546-1>.

1788

1789

1790

1791

1792

1793

1794 [106] G. W. STEWART, *The efficient generation of random orthogonal matrices with*
 1795 *an application to condition estimators*, SIAM Journal on Numerical Analysis,
 1796 17 (1980), pp. 403–409, <https://doi.org/10.1137/0717034>.

1797 [107] G. W. STEWART, *On the early history of the singular value decomposition*,
 1798 SIAM Review, 35 (1993), pp. 551–566, <https://doi.org/10.1137/1035134>.

1799 [108] G. W. STEWART, *QR sometimes beats Jacobi*, Tech. Report CS-TR-3434, Uni-
 1800 versity of Maryland, 1995, <http://drum.lib.umd.edu/handle/1903/709>.

1801 [109] S. TOMOV, R. NATH, AND J. DONGARRA, *Accelerating the reduction to upper*
 1802 *Hessenberg, tridiagonal, and bidiagonal forms through hybrid GPU-based com-*
 1803 *puting*, Parallel Computing, 36 (2010), pp. 645–654, <https://doi.org/10.1016/j.parco.2010.06.001>.

1804 [110] S. TOMOV, R. NATH, H. LTAIEF, AND J. DONGARRA, *Dense linear alge-*
 1805 *bra solvers for multicore with GPU accelerators*, in 2010 IEEE International
 1806 on Parallel and Distributed Processing Symposium, Workshops and Phd Fo-
 1807 rum (IPDPSW), IEEE, 2010, pp. 1–8, <https://doi.org/10.1109/IPDPSW.2010.5470941>.

1808 [111] M. A. TURK AND A. P. PENTLAND, *Face recognition using eigenfaces*, in
 1809 Computer Vision and Pattern Recognition, 1991. Proceedings CVPR’91., IEEE
 1810 Computer Society Conference on, IEEE, 1991, pp. 586–591, <https://doi.org/10.1109/CVPR.1991.139758>.

1811 [112] C. VAN LOAN, *The block Jacobi method for computing the singular value decom-*
 1812 *position*, Tech. Report TR 85-680, Cornell University, 1985, <https://ecommons.cornell.edu/handle/1813/6520>.

1813 [113] F. G. VAN ZEE, R. A. VAN DE GEIJN, AND G. QUINTANA-ORTÍ, *Restruc-*
 1814 *turing the tridiagonal and bidiagonal QR algorithms for performance*, ACM
 1815 Transactions on Mathematical Software (TOMS), 40 (2014), p. 18, <https://doi.org/10.1145/2535371>.

1816 [114] F. G. VAN ZEE, R. A. VAN DE GEIJN, G. QUINTANA-ORTÍ, AND G. J.
 1817 ELIZONDO, *Families of algorithms for reducing a matrix to condensed form*,
 1818 ACM Transactions on Mathematical Software (TOMS), 39 (2012), p. 2, <https://doi.org/10.1145/2382585.2382587>.

1819 [115] R. C. WHALEY AND J. DONGARRA, *Automatically tuned linear algebra soft-*
 1820 *ware*, in Proceedings of the 1998 ACM/IEEE Conference on Supercomputing,
 1821 IEEE Computer Society, 1998, pp. 1–27, <https://doi.org/10.1109/SC.1998.10004>.

1822 [116] J. H. WILKINSON, *Note on the quadratic convergence of the cyclic Jacobi pro-*
 1823 *cess*, Numerische Mathematik, 4 (1962), pp. 296–300, <https://doi.org/10.1007/BF01386321>.

1824 [117] J. H. WILKINSON AND C. REINSCH, *Handbook for Automatic Computation:*
 1825 *Volume II: Linear Algebra*, vol. 186 of Grundlehren Der Mathematischen Wis-
 1826 *senschaften*, Springer-Verlag, New York, NY, 1971, <https://doi.org/10.1007/978-3-642-86940-2>.

1827 [118] P. R. WILLEMS AND B. LANG, *A framework for the MR³ algorithm: theory and*
 1828 *implementation*, SIAM Journal on Scientific Computing, 35 (2013), pp. A740–
 1829 A766, <https://doi.org/10.1137/110834020>.

1830 [119] P. R. WILLEMS, B. LANG, AND C. VÖMEL, *Computing the bidiagonal SVD us-*
 1831 *ing multiple relatively robust representations*, SIAM Journal on Matrix Analysis
 1832 and Applications, 28 (2006), pp. 907–926, <https://doi.org/10.1137/050628301>.

1833 [120] B. B. ZHOU AND R. P. BRENT, *A parallel ring ordering algorithm for effi-*
 1834 *cient one-sided Jacobi SVD computations*, Journal of Parallel and Distributed
 1835

