

1       **MIXED-PRECISION ALGORITHM FOR FINDING SELECTED**  
2       **EIGENVALUES AND EIGENVECTORS OF SYMMETRIC AND**  
3       **HERMITIAN MATRICES** \*

4               YAOHUNG M. TSAI<sup>†</sup>, PIOTR LUSZCZEK<sup>‡</sup>, AND JACK DONGARRA<sup>§</sup>

5       **Abstract.** As the new hardware is being equipped with powerful low-precision capabilities  
6 driven primarily by the needs of the burgeoning field of Artificial Intelligence (AI), mixed-precision  
7 algorithms are now showing far greater potential and renewed interest in scientific computing com-  
8 munity. The multi-precision methods commonly follow approximate-iterate scheme by first obtaining  
9 the approximate solution from a low-precision factorization and solve. Then, they iteratively refine  
10 the solution to the desired accuracy that is often as high as what is possible with traditional ap-  
11 proaches. While targeting symmetric and Hermitian eigenvalue problems of the form  $Ax = \lambda x$ , we  
12 revisit the SICE algorithm proposed by Dongarra et al. By applying the Sherman-Morrison formula  
13 on the diagonally-shifted tridiagonal systems, we propose an updated SICE-SM algorithm. By in-  
14 corporating the latest two-stage algorithms from the PLASMA and MAGMA software libraries for  
15 numerical linear algebra, we achieved up to  $3.6\times$  speedup using the mixed-precision eigensolver with  
16 the blocked SICE-SM algorithm for iterative refinement when compared with full double complex  
17 precision solvers for the cases with a portion of eigenvalues and eigenvectors requested.

18       **Key words.** mixed-precision algorithms, eigenvalue solver, hardware accelerators

19       **AMS subject classifications.** 65F15, 65F25, 65Y20, 68N01

20       **1. Introduction.** The symmetric eigenvalue problem is one of the most impor-  
21 tant problems in numerical linear algebra for analysis of invariant subspace. For real  
22 matrices, the objective is to find an eigenvalue  $\lambda$  and the corresponding eigenvector  $x$   
23 such that

24 (1.1)                        $Ax = \lambda x$  where  $A = A^T$ ,  $A \in \mathbb{R}^{n \times n}$

25 The Hermitian eigenvalue problem is to find the eigenvalues and eigenvectors in com-  
26 plex domain. For an Hermitian matrix  $A$ , the conjugate transpose (adjoint) operation  
27 is idempotent:  $A = A^*$  and the eigenvalues are real which implies shared properties  
28 with the symmetric eigenvalue problem in real domain.

29 As mixed-precision algorithms for solving a linear system of equations experi-  
30 enced a substantial interest that resulted in recent developments [10, 11, 25]. These  
31 were mostly driven by the introduction of new hardware platforms that provide in-  
32 creased low-precision performance for AI workloads. However, there was not as much  
33 focus on eigenvalue problems. And with the latest two-stage tridiagonalization ap-  
34 proach [24, 26], the multicore and multi-GPU eigensolvers' algorithms for refining  
35 eigenvalues should be reviewed carefully in order to ascertain the possibility to im-  
36 prove the performance especially on this new hardware.

---

\*This material is also based upon work supported by the National Science Foundation under OAC Grant No. 2004541 and the University of Tennessee grant MSE E01-1315-038 as Interdisciplinary Seed funding. This research used the computational resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725 provided 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. The software library integration work was supported by the National Science Foundation under OAC Grant No. 2004541.

<sup>†</sup>University of Tennessee ([ytsai2@icl.utk.edu](mailto:ytsai2@icl.utk.edu))

<sup>‡</sup>University of Tennessee ([luszczek@icl.utk.edu](mailto:luszczek@icl.utk.edu))

<sup>§</sup>University of Tennessee, Oak Ridge National Laboratory, University of Manchester ([dongarra@icl.utk.edu](mailto:dongarra@icl.utk.edu))

37 **2. Contribution.** The motivation of this paper is to develop a mixed-precision  
38 algorithm for the symmetric/hermitian eigenvalue problem. In this context, we made  
39 the following contributions:

- 40 • We formulated and implemented the SICE-SM algorithm for symmetric and  
41 Hermitian matrices using Sherman-Morrison formula to solve the resulting  
42 tridiagonal systems with rank-one updates based on the SICE algorithm pro-  
43 posed by Dongarra et al. for iteratively refining a pair of eigenvalues and  
44 corresponding eigenvectors.
- 45 • We also developed the blocked SICE-SM algorithm to refine multiple eigen-  
46 values and corresponding eigenvectors simultaneously and improve the per-  
47 formance by aggregating the matrix-vector multiplication operations for im-  
48 proved utilization of the memory hierarchy.
- 49 • We implemented a mixed-precision algorithm by rearranging the tasks in  
50 the 2-stage eigensolver on the machines with heterogeneous architecture by  
51 improving the utilization of both the CPUs and GPUs.
- 52 • We achieved performance improvement of the mixed-precision algorithm with  
53 the iterative refinement for the case that only a portion of eigenpairs are  
54 requested.

### 55 3. Related Work.

56 **3.1. Eigenvalue refinement.** Symm and Wilkinson[41] proposed an algorithm  
57 to determine the error bounds of computed eigenvalues and eigenvectors, which can  
58 also be used to improve the accuracy of a given eigen-pair. Dongarra, Moler, and  
59 Wilkinson[18, 19, 20] later improved the algorithm with reduced computational cost  
60 and provided additional error analysis, including the comparison to Newton’s method[36,  
61 47], numerical results, and discussion of extending the algorithm for ill-conditioned  
62 problems with multiple close eigenvalues. More detail will be reviewed in subsec-  
63 tion 3.4 as it is also the core of the algorithm used in this work.

64 Other related work from Stewart[40] and Chatelin[13] answered the same question  
65 from the point of view of the invariant subspace problem. Demmel[16] later pointed  
66 out that these two methods and the one from Dongarra, Moler, and Wilkinson [18,  
67 19, 20] can all be reduced to solving the same Riccati equation. He also extended the  
68 algorithm for the generalized eigenvalue problem of the form  $Ax = \lambda Bx$ .

69 Alefeld and Spreuer[3] followed the same approach but specifically targeted prob-  
70 lems with doubly-repeated or numerically close eigenvalues. Tisseur[42] did the analy-  
71 sis of Newton’s method under floating-point arithmetic for generalized eigenvalue  
72 problems. Prikopa and Gansterer[37] used the symmetry of the matrix and House-  
73 holder tridiagonalization  $A = QTQ^T$  to reduce the computational cost.

74 Ogita and Aishima[31] proposed a different iterative scheme, which heavily relies  
75 on matrix-matrix multiplication for those applications which require accuracy that is  
76 higher than the base IEEE-754 double precision. The algorithm is applied on the entire  
77 spectrum of eigenvalues but it is capable of improving at the same time the orthogonal-  
78 ity and eigenvalue accuracy. However, it requires high-precision computation for the  
79 most parts of the algorithm, making it costly in practice. Later the authors extended  
80 the algorithm for clustered eigenvalues and singular value decomposition[32, 33].

81 **3.2. Parallel Eigensolvers.** To build an efficient mixed-precision algorithm,  
82 the latest advances in parallel eigensolvers should also be incorporated. The symmet-  
83 ric dense eigensolvers are mainly composed of two phases: tridiagonal reduction and  
84 tridiagonal eigensolver. Firstly, through similarity transformations based on orthog-

85 onal/unitary matrices, the symmetric/Hermitian matrix is reduced to a tridiagonal  
86 form without altering the spectrum in infinite precision or with numerically stable  
87 perturbation in final precision. Then the problem is solved in tridiagonal form with  
88 much less cost than operating on a full matrix by applying different methods which  
89 will be described later in the section. If needed, the eigenvectors can be computed  
90 from the eigenvectors of the tridiagonal system and applying back-transformations of  
91 tridiagonal reduction.

92 **3.2.1. Tridiagonal Reduction.** The first phase is to convert a full dense ma-  
93 trix into upper Hessenberg form, which has zeros below the first subdiagonal. The  
94 real symmetric and complex Hermitian cases result in even better structured form: a  
95 symmetric tridiagonal matrix with only nonzeros on the diagonal, the first superdiag-  
96 onal, and the first subdiagonal. The tridiagonalization of complex Hermitian matrix  
97 is usually chosen to be real tridiagonal symmetric matrix to reduce the computation  
98 cost in following steps. The Householder transformation is a natural choice for the  
99 reduction because of its simplicity and numerical stability. Furthermore, Dongarra et  
100 al.[21] introduced a blocked version of Householder vector application in which the  
101 transformations are aggregated and applied in a blocked fashion, so they can benefit  
102 from the high performance matrix-matrix multiplications rather than be bound by  
103 matrix-vector performance.

104 Bischof et al.[8] proposed the approach based on successive band reduction (SBR).  
105 Each reduction sweep results in a narrower band matrix, and the reduction is done via  
106 a bulge-chasing procedure. The algorithm consists of a series of sweeps: each sweep  
107 will zero-out one column below subdiagonal but create fill-ins down the diagonal as the  
108 transformations are applied to the remaining matrix. Then additional transformations  
109 are applied to zero out the fill-in which was just created and this is repeated all the  
110 way down to the lower-right corner until it disappears from the matrix, hence the  
111 algorithm name: the bulge chasing. The algorithm is naturally parallelizable as the  
112 subsequent sweeps can be chosen to not overlap with each other, making it especially  
113 suitable for multicore CPUs in shared-memory environments.

114 Later work introduced a hybrid 2-stage algorithm[24, 26]. The first stage still  
115 consisted of blocked Householder transformations but it only reduced the matrix to  
116 a band form. Then, the left transformation will only be needed, as the right trans-  
117 formation will not be touching the first block of columns. It thus becomes an  $LQ$   
118 factorization for the block of columns, which is much faster than applying the trans-  
119 formations from both sides (LQ and QR). The second stage uses the bulge-chasing  
120 algorithm from the successive band reductions.

121 **3.2.2. Tridiagonal Eigensolvers.** After tridiagonalization completes, a few  
122 standard eigensolver algorithms could be considered. As this is not the main focus of  
123 this work, these will only be reviewed briefly. The **QR algorithm** with shifts[46] is  
124 one of the most popular choices because of its superb stability and cubic convergence  
125 rate in general case. At each iteration, it computes a  $QR$  factorization and multiplies  
126 them back in reverse order:  $Q_k R_k = A_k - \mu_k I$ ;  $A_{k+1} = R_k Q_k + \mu_k I$ . There are other  
127 variants of QR iteration for strategically choosing the shifts  $\mu_k$ .

128 Another algorithm is called **divide and conquer**[14] that observes that with a  
129 rank-1 update, the initial problem can be divided into two independent subproblems  
130 with half the size. This results in repeatably reducing the problem down to the  
131  $1 \times 1$  case which admits a trivial solution. In practice, there is a threshold size and  
132 the implementation switches to another method for below-threshold sizes for better  
133 performance on small problems. The independent problems can easily be parallelized.

134 There are other methods based on the  $LDL^T$  factorization. The **Bisection**  
135 **method**[45] uses a suitable factorization to identify the number of eigenvalues pres-  
136 ent within a section and then it consecutively reduces the size of sections until the  
137 eigenvalues of interest are located with desired accuracy.

138 Finally, **Multiple relatively robust representations (MRRR)**[35] takes the  
139 bisection further by the theoretically estimating the gaps between neighboring eigen-  
140 values. This algorithm divides the whole spectrum into clusters of eigenvalues that  
141 each have a relatively robust representation ( $LDL^T$  factorization).

### 142 3.3. Software Packages for Symmetric/Hermitian Eigenvalue Problems. ■

143 This section provides details on the software packages that are available for numerical  
144 linear algebra and include dense eigensolvers.

145 EISPACK[39] is one the earliest open source software libraries to solve eigen-  
146 problems. It contains subroutines for the following nine classes of matrices: complex  
147 general, complex Hermitian, real general, real symmetric, real symmetric banded, real  
148 symmetric tridiagonal, special real tridiagonal, generalized real, and generalized real  
149 symmetric matrices. Providing performance portability of EISPACK motivated estab-  
150 lishment of Basic Linear Algebra Subprograms (BLAS)[29] as the standard building  
151 blocks for performing basic vector and matrix operations. BLAS was later extended  
152 to include three levels of operations: Level 1 scalar-vector and vector-vector, Level 2  
153 matrix-vector, and Level 3 matrix-matrix. Availability of BLAS proliferated as almost  
154 all hardware vendors provided their own optimized implementations and thus unified  
155 interface for numerical linear algebra software became the de facto standard upon  
156 which more complex methods are implemented including eigensolvers. The vendor  
157 renditions of BLAS for particular architectures include Intel MKL[28] and oneMKL,  
158 IBM ESSL[27], ARM Performance Libraries[7], NVIDIA cuBLAS[30], AMD AOCL[4]  
159 for CPUs and rocBLAS[5] for GPUs. The implementations from academia and open-  
160 source communities also exist and include BLIS[44] and OpenBLAS[34], both of which  
161 build on the success story of portable performance of GotoBLAS[23].

162 LAPACK[6] was designed to utilize Level 3 BLAS routines by introducing blocked  
163 algorithms to bring out the performance from hardware platforms based on then  
164 modern architecture of deep memory hierarchies. LAPACK provides routines for  
165 all the major numerical linear algebra problems, ranging from solving systems of  
166 linear equations, least-squares solutions of linear systems, eigenvalue problems, and  
167 singular value problems. Over the years, the library kept expanding and became the  
168 standard reference for dense numerical linear algebra applications as it includes the  
169 implementations of all the major algorithms in the field.

170 Several software libraries were subsequently developed that aimed to provide sim-  
171 ilar functionality as LAPACK while targeting different kinds of hardware platforms  
172 and environments. ScaLAPACK[9] was designed to scale on distributed-memory ma-  
173 chines by partitioning the matrices into blocks and cyclically distributing the data  
174 across the nodes. Its algorithms were implemented to iterate over these blocks to  
175 achieve parallelism. As the multicore CPUs were emerging, PLASMA[2] took a simi-  
176 lar idea of breaking the matrix down, but instead used smaller submatrices called tiles  
177 that better exploit the hardware structure of these shared-memory multicore systems.  
178 A task-based scheduler was introduced to remove the synchronization points in the  
179 algorithms and replace them with runtime scheduling of small tasks which operate on  
180 the tiles and are tracked based on their data dependences. MAGMA[43] was designed  
181 for heterogeneous architecture settings by exploiting hybrid hardware environment.  
182 These systems were equipped with hardware accelerators, usually GPUs, along with

183 multicore CPUs. As the GPU brought a lot of computational power in terms of  
 184 floating-point operations, the communications between the CPU and GPU remained  
 185 a bottleneck, as the bandwidth between the two continues to be much more limited  
 186 in comparison to internal memory structure of either a CPU or GPU. Thus the im-  
 187 plementations in MAGMA were redesigned to distribute different tasks to the CPU  
 188 and GPU to optimally fit their strengths and at the same time overlap the CPU-GPU  
 189 communication with computations as much as possible. Software for Linear Algebra  
 190 Targeting Exascale (SLATE)[22] aims to replace the venerable ScaLAPACK library.  
 191 As the latest supercomputer installations are commonly accelerated by multiple GPUs  
 192 on every distributed node, it would be hard to modify ScaLAPACK to take advantage  
 193 of such machines. SLATE is designed with this modern HPC hardware in mind and  
 194 features support for multiple computational backends. SLATE also embraces the open  
 195 standards like MPI and OpenMP to promote portability while retaining performance  
 196 and parallel efficiency.

197 **3.4. The SICE Algorithm.** In this section, we review the SICE algorithm by  
 198 Dongarra et al. [18, 19, 20]. Given the base eigenpair  $\lambda, x$  and its nearby eigenpair  
 199  $\lambda + \mu, x + \tilde{y}$ , then based on the original eigenproblem we have:

$$200 \quad (3.1) \quad A(x + \tilde{y}) = (\lambda + \mu)(x + \tilde{y})$$

201 Assuming that  $x$  is normalized in infinite norm:  $|x|_\infty = 1 \equiv x_s$ , we can remove one  
 202 degree of freedom by requiring  $\tilde{y}_s = 0$ . Rearranging Eq. (3.1) we get:

$$203 \quad (3.2) \quad (A - \lambda I)\tilde{y} - \mu x = \lambda x - Ax - \mu\tilde{y}$$

204 The last term is the second order term for the error in  $\lambda$  and  $x$ . By simplify the  
 205 equation, we introduce vector  $y$ , defined as:

$$206 \quad (3.3) \quad y^\top \stackrel{\text{def}}{=} (\tilde{y}_1, \tilde{y}_2, \dots, \tilde{y}_{s-1}, \mu, \tilde{y}_{s+1}, \dots, \tilde{y}_{n-1}, \tilde{y}_n)$$

207 So  $y$  would encode information from both  $\tilde{y}$  and  $\mu$  and thus Eq. (3.2) becomes:

$$208 \quad (3.4) \quad By = r + y_s \tilde{y} = r + \mu \tilde{y}$$

209 where  $r = \lambda x - Ax$  is the residual vector of  $\lambda$  and  $x$  and  $B$  is the matrix  $A - \lambda I$  with  
 210 column  $s$  replaced by  $-x$ .

211 We can also view it as the Newton's method. In particular, by setting  $v = \begin{pmatrix} x \\ \lambda \end{pmatrix}$   
 212 we can be formulate the eigenvalue problem as:

$$213 \quad (3.5) \quad f(v) \equiv \begin{pmatrix} Ax - \lambda x \\ e_s^\top x - 1 \end{pmatrix} = 0$$

214 where  $e_s$  is the  $s$ -th column of the identity matrix of size  $n$ . The Newton's method  
 215 then solves the linear system of the Jacobian matrix:

$$216 \quad (3.6) \quad J \begin{pmatrix} \tilde{y} \\ \mu \end{pmatrix} = \begin{pmatrix} A - \lambda I & -x \\ e_s^\top & 0 \end{pmatrix} \begin{pmatrix} \tilde{y} \\ \mu \end{pmatrix} = \begin{pmatrix} r \\ 0 \end{pmatrix} = f(v)$$

217 Expanding it, we arrive at Eq. (3.2) without the second-order term:

$$218 \quad (3.7) \quad (A - \lambda I)\tilde{y} - \mu x = r$$

219 This is the basic idea of the SICE algorithm: by iteratively solving Eq. (3.4) we  
 220 obtain both the correction to the eigenvalue and to the eigenvector. The original  
 221 algorithm uses Schur decomposition and applies two steps of Givens rotation in order  
 222 to solve Eq. (3.4). For any real matrix  $A$ , there exists an orthogonal matrix  $Q$  and  
 223 an upper quasi-triangular matrix  $T$ , such that

$$224 \quad (3.8) \quad A = QUQ^\top$$

225 where  $U$  is upper quasi-triangular with some  $2 \times 2$  diagonal blocks arising from complex  
 226 conjugate eigenvalue pairs. Here, we define  $Z_\lambda \equiv Z - \lambda I$  and  $z_{\lambda s} \equiv Z_\lambda e_s = (Z - \lambda I)e_s$ .  
 227 By rewriting Eq. (3.4), we get:

$$228 \quad (3.9) \quad [A_\lambda - (x + a_{\lambda s})e_s^\top]y = (A + ce_s^\top)y = r + y_s\tilde{y}$$

229 where  $c = -x - a_{\lambda s}$ . Using the Schur decomposition  $A = QUQ^\top$ , we have:

$$230 \quad (3.10) \quad Q(U_\lambda + Q^\top ce_s^\top Q)Q^\top y = r + y_s\tilde{y}$$

231

$$232 \quad (3.11) \quad (U_\lambda + df^\top)Q^\top y = Q^\top g$$

233 where  $d = Q^\top c$ ,  $f^\top = e_s^\top Q$  and  $g = r + y_s\tilde{y}$ . Matrix  $d \times f^\top$  constitutes a rank-1 update.  
 234 Then two steps of Givens rotation are introduced: the first one  $Q_1$  is constructed so  
 235 that

$$236 \quad (3.12) \quad Q_1 d = (P_2 P_3 \dots P_n) d = \gamma e_1 \quad \text{where } \gamma = \|d\|_2$$

237 and  $P_i$  is the rotation in  $(i-1, i)$  plane that eliminates the  $i$ -th component in  
 238  $P_{i+1} \dots P_n d$ . We also have:

$$239 \quad (3.13) \quad Q_1(U_\lambda + df^\top) = Q_1 U_\lambda + \gamma e_1 f^\top$$

240 The transformation  $Q_1$  introduces one more nonzero element in the subdiagonal di-  
 241 rection of  $U_\lambda$ . The new rank-one update  $\gamma e_1 \times f^\top$  has nonzero elements only in the  
 242 first row, which preserves the original structure. The second step of Givens rota-  
 243 tion  $Q_2$  can be applied subsequently in order to obtain the upper triangular form  
 244  $\bar{U}_\lambda = Q_2 Q_1 (U_\lambda + d \times f^\top)$  in

$$245 \quad (3.14) \quad \bar{U}_\lambda Q^\top y = Q_2 Q_1 Q^\top g$$

246 The triangular solve requires  $O(n^2)$  operations while the remaining steps of the iter-  
 247 ation are only  $O(n)$ . This procedure is shown in Algorithm 3.1.

248 **4. Algorithm and Implementation.** The original SICE algorithm is designed  
 249 for a general real matrices and here we first focus on symmetric ones. The proposed  
 250 algorithm utilizes the tridiagonalization as well as the Sherman–Morrison formula  
 251 to solve the linear system for eigenvalue and eigenvector corrections. The blocked  
 252 version will also be discussed with the implementation details based on PLASMA  
 253 and MAGMA software libraries.

---

**Algorithm 3.1** SICE algorithm

---

```
1: Input: Matrix  $A \in \mathbb{R}^{n \times n}$ . An approximate eigenvalue  $\lambda$  and the corresponding eigen-
vector  $x$ .  $\text{iter}_{\max}$  denotes the maximum number of iterations.
2: Output: Refined eigenvalue  $\lambda$  and its eigenvector  $x$ .
3: function  $[\lambda, x] \leftarrow \text{SICE}(A, \lambda, x, \text{iter})$ 
4:    $[Q, U] \leftarrow \text{schur}(A)$   $\triangleright$  obtain Schur decomposition  $A = QUQ^\top$ ,  $QQ^\top = I$ .
5:    $[m, s] \leftarrow \text{max}(\text{abs}(x)); x \leftarrow x/m$   $\triangleright$  Normalizing  $x$  so that  $\|x\|_\infty = s_x = 1$ .
6:   for  $i$  in  $1 : \text{iter}_{\max}$  do
7:      $r \leftarrow \lambda x - Ax$ 
8:      $c \leftarrow -x - a_{\lambda s}$ 
9:      $d \leftarrow Q^\top c$ 
10:     $f^\top \leftarrow Q(s, :) = e_s^\top Q$   $\triangleright$   $s$ -th row of  $Q$ .
11:     $\bar{U}_\lambda \leftarrow Q_1(U - \lambda I); \bar{d} \leftarrow Q_1 d = \|d\|_2 e_1$   $\triangleright$  Givens rotations  $Q_1$  from Eq. (3.12)
12:     $\bar{U}_\lambda \leftarrow \bar{U}_\lambda + \bar{d}(1)f^\top$ 
13:     $\bar{U}_\lambda \leftarrow Q_2 \bar{U}_\lambda$   $\triangleright$  Givens rotations  $Q_2$  to introduce upper triangular form.
14:    Solve the triangular system  $\bar{U}_\lambda z = Q_2 Q_1 Q^\top r$ 
15:     $y \leftarrow Qy$ 
16:     $\lambda \leftarrow \lambda + y(s)$   $\triangleright$  Update eigenvalue.
17:     $y(s) \leftarrow 0$   $\triangleright$  Set  $y(s)$  to 0.
18:     $x \leftarrow x + y$   $\triangleright$  Update eigenvector.
19:    if desired accuracy is reached then
20:      break
21:    end if
22:  end for
23: end function
```

---

254 **4.1. SICE-SM Algorithm.** For symmetric eigenvalue problems, the matrix  
255  $A$  is first reduced to tridiagonal through unitary similarity transformations:  $T =$   
256  $Q^\top A Q$  where  $QQ^\top = I$  and  $T$  is a symmetric tridiagonal matrix. This corresponds  
257 to LAPACK routines `SSYTRD` and `DSYTRD` for single- and double-precision arithmetic,  
258 respectively. In the same fashion as SICE algorithm in Section 3.4, we start with  
259 Eq. (3.9) and apply the tridiagonal reduction to it. Eqs. (3.10) and (3.11) in this case  
260 become

$$261 \quad (4.1) \quad Q(T_\lambda + Q^\top c e_s^\top Q) Q^\top y = r + y_s \tilde{y}$$

262 and

$$263 \quad (4.2) \quad (T_\lambda + d \times f^\top) Q^\top y = Q^\top g$$

264 the same with  $d = Q^\top c$ ,  $f^\top = e_s^\top Q$  and  $g = r + y_s \tilde{y}$ . Dongarra[18] discussed the  
265 approach of using the Sherman–Morrison formula[38]

$$266 \quad (4.3) \quad (A - uv^\top)^{-1} = A^{-1} - \frac{A^{-1}uv^\top A^{-1}}{1 + v^\top A^{-1}u}$$

267 for solving the rank-one updated system. Eq. (4.2) does not apply since  $T_\lambda = T - \lambda I$   
268 is singular by construction. However, this may not be so in mixed-precision setting.  
269 Consider the scheme that first performs the tridiagonal reduction in single precision  
270 and then solves the tridiagonal eigenvalue problem in double precision. The initial  $\lambda_T$   
271 will be the eigenvalue of  $T$  with double-precision accuracy, but it only approximates  
272  $\lambda_A$ , the eigenvalue of  $A$  with single-precision accuracy. With suitably chosen offset  $\delta$



---

**Algorithm 4.1** SICE-SM algorithm: SICE algorithm with Sherman–Morrison formula

---

```

1: Input: Matrix  $A = A^\top \in \mathbb{R}^{n \times n}$ . An approximate eigenvalue  $\lambda$  and the corresponding
   eigenvector  $x$ .  $\text{iter}_{\max}$  denotes the maximum number of iterations.
2: Output: Refined eigenvalue  $\lambda$  and eigenvector  $x$ .
3: function  $[\lambda, x] \leftarrow \text{SICE\_SM}(A, \lambda, x, \text{iter})$ 
4:    $[Q, T] \leftarrow \text{tridiag}(A)$  ▷ Tridiagonalization  $A = QTQ^\top$ ,  $QQ^\top = I$ .
5:    $[m, s] \leftarrow \text{max}(\text{abs}(x)); x \leftarrow x/m$  ▷ Normalization of  $x$  so that  $\|x\|_\infty = s_x = 1$ .
6:   for  $i$  in  $1 : \text{iter}_{\max}$  do
7:      $r \leftarrow \lambda x - Ax$ 
8:      $c \leftarrow -x - a_{\lambda s}$ 
9:      $d \leftarrow Q^\top c$ 
10:     $f^\top \leftarrow Q(s, :) = e_s^\top Q$  ▷  $s$ -th row of  $Q$ .
11:     $\text{rhs} \leftarrow Q^\top r$ 
12:     $u \leftarrow (T - \lambda I)^{-1} d$ 
13:     $v \leftarrow (T - \lambda I)^{-1} \text{rhs}$ 
14:     $y \leftarrow v - \frac{f^\top v}{1 + f^\top u} u$  ▷ Sherman–Morrison formula
15:     $y \leftarrow Qy$ 
16:     $\lambda \leftarrow \lambda + y(s)$  ▷ Update eigenvalue.
17:    if  $i \neq 1$  then
18:       $y(s) \leftarrow 0$  ▷ Set  $y(s)$  to 0.
19:       $x \leftarrow x + y$  ▷ Update eigenvector.
20:    end if
21:    if desired accuracy reached then
22:      break
23:    end if
24:  end for
25: end function

```

---

273 of order of  $\epsilon_{\text{single}}$ ,  $T - (\lambda + \delta)I$  will no longer be singular in double precision, and the  
274 Sherman–Morrison formula can be applied. The special case in which this would fail  
275 is when  $\|\lambda_T - \lambda_A\| = O(\epsilon_{\text{double}})$ : the initial eigenvalue is also an accurate eigenvalue  
276 of  $A$  in double precision. In such a case, we do not need to refine the eigenvalue and  
277 can simply apply the inverse iteration to find the eigenvector.

278 Applying Sherman–Morrison formula from Eq. (4.3) to Eq. (4.2) we get

$$279 \quad (4.4) \quad Q^\top y = \left( T_\lambda^{-1} - \frac{T_\lambda^{-1} d \times f^\top T_\lambda^{-1}}{1 + f^\top T_\lambda^{-1} d} \right) Q^\top g$$

280 or

$$281 \quad (4.5) \quad Q^\top y = T_\lambda^{-1} Q^\top g - \frac{f^\top (T_\lambda^{-1} Q^\top g)}{1 + f^\top (T_\lambda^{-1} d)} T_\lambda^{-1} d$$

282 These involve solving the tridiagonal system  $T_\lambda$  with two different right hand sides  $d$   
283 and  $Q^\top g$ . It can be easily done with the Thomas algorithm which is a special case of  
284 Gaussian elimination. There are other parallel tridiagonal solvers available and we will  
285 discuss them in Section 4.3.1. We outline the SICE algorithm with Sherman–Morrison  
286 formula in Algorithm 4.1.

287 The main difference between Algorithms 3.1 and 4.1 is the use of the Sher-  
288 man–Morrison formula to solve the system from line 12 to 14 instead of using the



TABLE 1

Performance of  $n \times n$  matrix times  $n \times m$  aggregated vectors on NVIDIA V100-SXM2-32GB GPU, DGEMM routine from cuBLAS v11.0.

Matrix size	Number of vectors	Time (ms)	Performance (GFLOP/s)
20000	1	3.76	212.65
20000	8	3.79	1688.17
20000	32	6.48	3949.32
20000	128	13.57	7544.43

289 Givens rotations for that purpose. It is applied to solving the same tridiagonal sys-  
 290 tem  $T_\lambda$  with two different right hand sides  $d$  and  $Q^\top g$ . The two vector inner products  
 291 are needed to obtain the scalar in order to form the solution. Note that in line 17, we  
 292 only update the eigenvalue at the first iteration and leave the eigenvector unchanged  
 293 because  $T_\lambda$  at the first iteration is nearly singular. Other approaches to this issue  
 294 include manually applying a shift to the initial eigenvalue or using the Ritz value  
 295  $\frac{x^\top Ax}{x^\top x}$  as the starting point. Apart from tridiagonalization, the computational cost for  
 296 algorithm 4.1 is dominated by the matrix-vector multiplications which require  $O(n^2)$   
 297 operations. The remaining steps of the algorithm are all order  $O(n)$  including the  
 298 tridiagonal solve.

299 Alternatively, as described in [37], one can also solve the Jacobian matrix with  
 300 the special structure  $J = \begin{pmatrix} T - \lambda I & y \\ z^\top & 0 \end{pmatrix}$ , which is a tridiagonal system with an extra  
 301 row and column at the bottom and right. However, it is hard to parallelize the  
 302 corresponding solver for this special structure and it is even harder make it scalable.  
 303 This is in stark contrast with the approach of solving the tridiagonal system which  
 304 is well studied and admits several parallel implementations that target a variety of  
 305 computing environments.

306 **4.2. Blocked SICE-SM Algorithm.** The computational cost of Algorithm 4.1  
 307 is dominated by matrix-vector multiplications especially inside the refinement itera-  
 308 tion. In the matrix-vector multiplication, the whole matrix is read once and only a  
 309 single multiplication and addition are performed per each of the fetched elements.  
 310 This results in a low arithmetic intensity of 2, which results in very low inefficient on  
 311 modern hardware including CPU, GPUs, and computational accelerators. To improve  
 312 on this implementation aspect, we can aggregate several eigenpairs simultaneously and  
 313 refine them at the same time while they are cached in higher levels of the memory  
 314 hierarchy. This blocking strategy is common in numerical linear algebra since it was  
 315 introduced in LAPACK[6] and relies on grouping computations so that Level 3 BLAS  
 316 may be utilized to perform operations that are rich in matrix-matrix multilications.  
 317 These operations perform more efficiently as they have higher arithmetic intensity  
 318 resulting from higher data reuse in fast portions of the cache hierarchy. In our case,  
 319 we assume that the matrix size is far greater than the number of eigenpairs to refine.  
 320 Then the matrix-vector multiplication is dominated by the reading of the matrix ele-  
 321 ments. And with the blocked version, it the additional cost of refining extra eigenpairs  
 322 is negligible. In Table 1, we show examples of the performance rates and execution  
 323 times for different numbers of vectors submitted to the DGEMM routine from cuBLAS  
 324 on the NVIDIA V100 GPU. The times for 1 and 8 vectors are almost the same. And  
 325 for 32 or 128 vectors the elapsed time increases  $3.6\times$ .

326 There are a few issues we need to solve while formulating a blocked variant of

327 the algorithm. First, in SICE, the eigenvector is first normalized in infinity norm.  
 328 The index  $s$  is also picked so that  $\|x\|_\infty = s_x = 1$ . If we allow different  $s$  for each of  
 329 the eigenpairs, then we will have to access different columns in  $A$  to construct vector  
 330  $c$ , and also different rows of  $Q$  for vector  $f^\top$ . The row access required for the latter  
 331 is performed in column major layout and results in non-coalescing memory accesses  
 332 which are extremely slow and should be avoided as much as possible due to their  
 333 low utilization of the GPU's memory bandwidth. To show that it is fine to choose  
 334  $s$  arbitrarily, we need to take a closer look at the matrix in Eq. (4.1) and expand it  
 335 without canceling any terms we get

$$336 \quad (4.6) \quad (QT_\lambda Q^\top + QQ^\top ve_s^\top QQ^\top)y = r + y_s \tilde{y}$$

337 Again, for our mixed-precision scheme, we would like to perform the tridiagonaliza-  
 338 tion in single precision. Hence  $QT_\lambda Q^\top$  is only an approximation of  $A$  with precision  
 339  $\epsilon_{\text{single}}$ , i.e.  $\|A_\lambda - QT_\lambda Q^\top\| \sim O(\epsilon_{\text{single}})$ . The same applies to  $QQ^\top$  which is only an  
 340 approximation of  $I$  with  $\|QQ^\top - I\| \sim O(\epsilon_{\text{single}})$ . So no matter which index  $s$  we  
 341 pick, we will always get an error of order  $\epsilon_{\text{single}}$  in the correction of eigenvalue  $y_s$   
 342 coming from the other elements in the solution vector  $y$ . There could be a potential  
 343 problem if the eigenvalue itself is small and the error is preventing the eigenvalue to  
 344 be refined to desire accuracy. This can be remedied by pre-scaling the matrix so that  
 345 the eigenvalues are not too small.

346 The other issue is that by treating the eigenpairs independently they might lose  
 347 their orthogonality. In the worst case, they might all converge to the same eigenpair.  
 348 However, it is easy to reorthogonalize with

$$349 \quad (4.7) \quad X' = X + \frac{1}{2}X(I - X^\top X)$$

350 In practice, we found that it is sufficient to reorthogonalize after the refinement is done.  
 351 Doing so in each iteration would not speed up the convergence. The computation of  
 352  $I - X^\top X$  also lets us detect if they converged to the same eigenvector. By combining  
 353 these considerations, we arrive at Algorithm 4.2.

354 Because a Hermitian matrix can also be tridiagonalized into real matrix, algo-  
 355 rithm 4.2 can easily be extended to be applied on Hermitian matrices. The transfor-  
 356 mation matrix  $Q$  now becomes complex, as well as the intermediate vectors. However,  
 357 the coefficients in  $T - \lambda I$  are all real so it can be optimized to avoid doing all the  
 358 operations in complex space.

359 **4.3. Implementation Details.** In this section, we will describe some of the  
 360 details of our implementation. We implemented the Blocked SICE-SM (Algorithm 4.2)  
 361 in two software packages: PLASMA[2] and MAGMA[43].

362 PLASMA is a dense linear algebra software package targeting multi-core shared-  
 363 memory environments with OpenMP directives. It divides the work into small sub-  
 364 matrices called tiles in order to exploit the parallelism and dynamically schedule tasks  
 365 based on data interdependence. PLASMA used to have a runtime scheduler called  
 366 QUARK but it is now based on OpenMP tasking directives to embrace the open and  
 367 portable standard for runtime scheduling of computational Direct Acyclic Graphs  
 368 (DAGs). OpenMP 4 added the `depend` clause for task dependencies and is able to re-  
 369 solve the task DAGs from PLASMA algorithms. PLASMA has two-stage eigensolver  
 370 implemented in one of its development branches.

371 MAGMA is also a linear algebra software package but it targets heterogeneous  
 372 hardware accelerated with GPUs. Due to the characteristically high floating-point

---

**Algorithm 4.2** Blocked SICE-SM algorithm

---

```
1: Input:  $A = A^T \in \mathbb{R}^{n \times n}$ , initial eigenvectors  $X = [x_1|x_2|\dots|x_\ell] \in \mathbb{R}^{n \times \ell}$  and the corresponding initial eigenvalues  $\Lambda = (\lambda_1, \lambda_2, \dots, \lambda_\ell)^T \in \mathbb{R}^\ell$ .  $\text{iter}_{\max}$  denotes the maximum number of iterations.
2: Output: Refined eigenvectors  $X$  and refined eigenvalues  $\Lambda$ .
3: function  $[X, \Lambda] \leftarrow \text{SICE\_SM\_BLK}(A, X, \Lambda, \text{iter})$ 
4:    $[Q, T] \leftarrow \text{tridiag}(A)$  ▷ Tridiagonalization  $A = QTQ^T$ ,  $QQ^T = I$ .
5:   for  $i$  in  $1 : \text{iter}_{\max}$  do
6:      $s \leftarrow i$ 
7:      $R \leftarrow X \times \text{diag\_matrix}(\Lambda) - A \times X$  ▷ Residual vectors need higher precision.
8:     for  $j$  in  $1 : \ell$  do
9:        $c_j \leftarrow -x_j - A(:, s)$ 
10:    end for
11:    Compose matrix  $C = [c_1|c_2|\dots|c_\ell]$  from column vectors  $c_j$ 
12:     $C(s, :) \leftarrow C(s, :) + \Lambda^T$ 
13:     $D = [d_1|d_2|\dots|d_\ell] \leftarrow Q^T \times C$  ▷ Can be in lower precision.
14:     $RHS = [rhs_1|rhs_2|\dots|rhs_\ell] \leftarrow Q^T \times R$  ▷ Can be in lower precision.
15:     $f \leftarrow Q(s, :)$  ▷ s-th row of  $Q$ .
16:    for  $j$  in  $1 : \ell$  do
17:       $u_i \leftarrow (T - \lambda I)^{-1} d_i$ 
18:       $v_i \leftarrow (T - \lambda I)^{-1} rhs_i$ 
19:       $y_i \leftarrow v_i - \frac{f^T v_i}{1 + f^T u_i} u_i$  ▷ Sherman–Morrison
20:    end for
21:    Compose matrix  $Y = [y_1|y_2|\dots|y_\ell]$  from correction vectors  $y_j$ 
22:     $Y \leftarrow Q \times Y$ 
23:     $\Lambda \leftarrow \Lambda + Y(s, :)^T$  ▷ Update eigenvalues.
24:    if  $i \neq 1$  then
25:       $Y(s, :) \leftarrow 0$  ▷ Set  $y_i(s)$  to 0.
26:       $X \leftarrow X + Y$  ▷ Update eigenvectors.
27:      Normalize eigenvectors  $x_i$  in  $X$ .
28:    end if
29:    if desired accuracy reached then
30:      break
31:    end if
32:  end for
33:   $X \leftarrow X + \frac{1}{2}X(I - X^T X)$  ▷ Orthogonalization.
34: end function
```

---

373 performance of GPUs and the limited bandwidth between the CPUs and GPUs,  
374 MAGMA algorithms need to be redesigned and refactored to split up the work be-  
375 tween CPU and GPU and to overlap communication and computation. MAGMA  
376 includes both one- and two-stage eigensolvers. And we used them as building blocks  
377 for implementing Algorithm 4.2 for both solvers.

378 The one-stage eigensolver has the following components with its corresponding  
379 LAPACK routine names:

---

**Algorithm 4.3** One stage symmetric eigensolver

---

- 1: DSYTRD: Tridiagonalization via Householder transformations.
  - 2: DSTEDC: Tridiagonal symmetric eigensolver (divide and conquer).
  - 3: DORMTR: back transformation for eigenvectors.
-

380 First the system is transformed to the tridiagonal form via Householder transfor-  
 381 mations. Then the tridiagonal eigensolver is called. We will not discuss the details  
 382 of eigensolvers here, as it is not the focus of this work. After the eigenvalues and  
 383 eigenvectors of the tridiagonal system are computed, the back transformation is ap-  
 384 plied, which is the inverse of the Householder transformations from tridiagonalization  
 385 stage. Because the transformation is orthogonal, the inverse is simply a transpose. If  
 386 only a portion of the eigenvectors are requested, the transform would not be explicitly  
 387 formed for performance reasons. The transform in the form of elementary reflectors is  
 388 directly applied on eigenvectors of the tridiagonal system to obtain the eigenvectors  
 389 for the original matrix.

390 For the mixed-precision eigensolver, we first perform tridiagonalization in single  
 391 precision as it is computationally intensive requiring  $O(n^3)$  operations. After the  
 392 system is transformed to tridiagonal form, the eigensolver is applied. The eigensolver  
 393 operates in double precision as we need to be able to distinguish nearby eigenvalues  
 394 that are closer than  $\epsilon_{single}$  but not closer than  $\epsilon_{double}$ . If single precision is used for  
 395 this case, the eigenvalues are very likely to be considered as repeated, and the returned  
 396 eigenvectors could be an arbitrary orthogonal basis of the eigenspace. For the back  
 397 transformation, the matrix  $Q$  needs to be explicitly formed in order for us to solve  
 398 Eq. (4.2). Then the Blocked SICE-SM (Algorithm 4.2) is used to iteratively refine the  
 399 eigenpairs to the desired accuracy. Most of the operations in the refinement process  
 400 are matrix-matrix operations, which have been developed internally. The batched  
 401 tridiagonal solver in line 16 will be discussed in section 4.3.1.

---

**Algorithm 4.4** Mixed precision one stage symmetric eigensolver with iterative re-  
 refinement

---

- 1: SSYTRD: Tridiagonalization via Householder transformations in single precision.
  - 2: DSTEDC: Tridiagonal symmetric eigensolver (divide and conquer) in double pre-  
 cision.
  - 3: SORGTR: Generate the transformation matrix  $Q$  from elementary reflectors in  
 single precision.
  - 4: Blocked SICE-SM (algorithm 4.2) for iterative refinement.
- 

402 For two-stage algorithms, the structure is similar to the one-stage method but  
 403 both the forward- and back-transformations are split into two steps:

---

**Algorithm 4.5** Two stages symmetric eigensolver

---

- 1: First stage symmetric to band via Householder transformations.
  - 2: Second stage band to tridiagonal via bulge chasing.
  - 3: Tridiagonal symmetric eigensolver (divide and conquer).
  - 4: back transformation for second stage on eigenvectors.
  - 5: back transformation for first stage on eigenvectors.
- 

404 In MAGMA, the first stage is similar to QR factorization with the panel performed  
 405 completely on the CPU and the update of the trailing matrix performed on the GPU.  
 406 The second stage bulge chasing is implemented only for the CPU as the multicore  
 407 architecture with larger cache is a more suitable compared to the GPU. The divide-  
 408 and-conquer eigensolver is also mainly performed on the CPU except for the final step  
 409 of merging with large blocks. Both back transformations are applied on the GPU as  
 410 they are aggregated into matrix-matrix operations.

---

**Algorithm 4.6** Mixed precision two stages symmetric eigensolver with iterative refinement

---

- 1: First stage symmetric to band via Householder transformations in single precision.
  - 2: Second stage band to tridiagonal via bulge chasing in single precision.
  - 3: Tridiagonal symmetric eigensolver (divide and conquer) in double precision.
  - 4: Generate the transformation matrix  $Q$  from first stage in single precision. This can start as soon as 1. finishes.
  - 5: Apply the back transformation for second stage onto  $Q$  in single precision. This can start as soon as both 2. and 4. finish.
  - 6: Blocked SICE-SM (algorithm 4.2) for iterative refinement.
- 

411 Mixed precision for a two-stage eigensolver is actually more problematic performance-  
412 wise. The main reason is that accumulation of the back transformations from the  
413 second stage of bulge chasing is costly: it has a lot of small transformations and is  
414 expensive to apply on a square transform matrix  $Q$  compared to the case of only  
415 computing the eigenvectors. However, we need to explicitly form  $Q$  for the later re-  
416 finement. Here, we exploit the fact that the back transformation is not applied on  
417 the eigenvectors; it can actually start as soon as the first stage is finished. So we  
418 are reversing the order of back transformations to start it first. Similarly, the back  
419 transformation of the second stage can start when both the second stage and the back  
420 transformation of the first stage are completed. This is shown in Algorithm 4.6. For  
421 the case of MAGMA implementation, this would enable more parallelism. The back  
422 transformation of the first stage can be done on the GPU while the second stage of  
423 bulge chasing is done on the CPU. The eigensolver, which is mainly done on the CPU,  
424 can be overlapped with the back-transformation of the second stage on the GPU.

425 **4.3.1. Batched Tridiagonal Solver.** Line 16 in Algorithm 4.2 iterates over  
426 all the eigenvalues and solves the shifted tridiagonal system for each of them. This  
427 kind of computational pattern is suitable for batched interface. The term “batched”  
428 comes from the Batched BLAS[17, 1] that defines the interface for performing identical  
429 operation on multiple matrices independently and simultaneously. In our case, all  
430 the systems are also independent and we can solve them in a batched fashion. On  
431 multicore CPUs, the straightforward and efficient approach is to assign one system  
432 to each thread at a time which is likely bound to a single CPU core. Each thread  
433 can use the Thomas algorithm, which is a special case of Gaussian elimination. But  
434 on the GPU, we need more parallelism to saturate the computational potential of  
435 the hardware. There are previous studies[48, 15, 12] that investigated the solving of  
436 one big tridiagonal system on GPUs. One of the techniques is based on the cyclic  
437 reduction (CR). Consider a tridiagonal system with 8 unknowns:

$$(4.8) \quad \begin{bmatrix} b_1 & c_1 & & & & & & & \\ a_2 & b_2 & c_2 & & & & & & \\ & a_3 & b_3 & c_3 & & & & & \\ & & a_4 & b_4 & c_4 & & & & \\ & & & a_5 & b_5 & c_5 & & & \\ & & & & a_6 & b_6 & c_6 & & \\ & & & & & a_7 & b_7 & c_7 & \\ & & & & & & a_8 & b_8 & \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \\ x_5 \\ x_6 \\ x_7 \\ x_8 \end{bmatrix} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \\ y_6 \\ y_7 \\ y_8 \end{bmatrix}$$

439 By combing all the even-indexed equations with odd-indexed equation, we are

440 able to have an updated system with half of the size:

$$441 \quad (4.9) \quad \begin{bmatrix} b'_1 & c'_1 & & \\ a'_3 & b'_3 & c'_3 & \\ & a'_5 & b'_5 & c'_5 \\ & & a'_7 & b'_7 \end{bmatrix} \begin{bmatrix} x_1 \\ x_3 \\ x_5 \\ x_7 \end{bmatrix} = \begin{bmatrix} y'_1 \\ y'_3 \\ y'_5 \\ y'_7 \end{bmatrix}$$

442 The coefficients of the updated system can be computed with the following for-  
443 mulas:

$$444 \quad (4.10) \quad \begin{aligned} k_1 &= \frac{a_i}{b_{i-1}}, k_2 = \frac{c_i}{b_{i+1}} \\ a'_i &= -a_{i-1}k_1, b'_i = b_i - c_{i-1}k_1 - a_{i+1}k_2 \\ c'_i &= -c_{i+1}k_2, y'_i = y_i - y_{i-1}k_1 - y_{i+1}k_2 \end{aligned}$$

445 By recursively reducing the size of the system by half, it is possible to bring the  
446 size down to a single unknown with a trivial solution. Then, the back-substitutions  
447 follows the same path in reverse order and thus the solution of the full system is  
448 obtained. Alternatively, while reducing the size of systems, we can produce two  
449 independent systems, one with odd-indexed unknowns and the other with the even-  
450 indexed unknowns. Both systems can be solved independently with only its own  
451 coefficients. By repeating the process, we will arrive at trivial systems with a single  
452 unknown  $b''_i x_i = y''_i$  for all of the unknowns  $x_i$ . The back substitutions would not be  
453 needed for this approach, which is called *parallel cyclic reduction* (PCR). The PCR  
454 method exposes more parallelism towards the end but with requires more computation  
455 which represents a design trade-off. For our GPU implementation, we used PCR to  
456 solve one tridiagonal system by each of the thread blocks.

457 **5. Numerical Results and Performance Experiments.** The numerical ex-  
458 periments in this section will be divided into two parts. The first one examines the  
459 convergence behavior for refining different portions of the eigenvalues and eigenvec-  
460 tors in the spectrum. Then the performance results with PLASMA and MAGMA  
461 software libraries are be given with detailed profiling data to highlighted particular  
462 performance cases.

463 **5.1. Numerical Convergence.** The numerical experiments in this section were  
464 performed in MATLAB version R2020a with implementations of Algorithm 4.2 (blocked  
465 SICE-SM). The expression `A = gallery('randsvd', n, -cond)` was used to generate  
466 symmetric test matrices with a prescribed condition number from random eigenvec-  
467 tors and geometrically distributed eigenvalues in range  $(1, \frac{1}{\text{cond}})$ . The input matrix  
468 is first converted to `single` precision and subsequently tridiagonalized using `[Q,T]`  
469 `= hess(A)` function in single precision. Then converted back to double precision for  
470 finding the eigenvalues and eigenvectors using expression `[V,D] = eig(A)`. The eigen-  
471 vectors in `D` and column eigenvectors in `QV` will be used as the starting point of our  
472 refinement algorithms.

473 Figure 1 shows the convergence of Algorithm 4.2: the blocked SICE-SM. The  
474 input symmetric input matrix had size 100 with geometrically distributed eigenvalues  
475 from 1 to  $10^{-7}$ . The convergence in terms of residual  $\|Ax - \lambda x\|_\infty$  of each eigenvalues  
476 are plotted in different colors from blue as largest eigenvalue 1 to red as the smallest  
477 eigenvalue  $10^{-7}$ . For the first iteration, we only updated the eigenvalues so there  
478 was no initial improvement. For large eigenvalues, the method converges quickly in

479 two iterations. However, for small eigenvalues, that are much closer to each other  
 480 due to the geometrical distribution and thus we observe the resulting slowdown of  
 481 convergence.




FIG. 1. Blocked *SICE-SM* convergence of a  $100 \times 100$  matrix with geometrically distributed eigenvalues from 1 (blue) to  $10^{-7}$  (red).

482 **5.2. Performance Results.** The system we are using has two sockets of Intel(R)  
 483 Xeon(R) CPU E5-2650 v3 CPUs. But only one is being used for more stable  
 484 results. The system is accelerated by a Tesla V100 GPU. The theoretical peak per-  
 485 formance of a V100 is 7.8 TFLOP/s in double precision and 15.6 TFLOP/s in single  
 486 precision. The software stacks was composed of Intel Parallel Studio Cluster 2020.  
 487 (for C and Fortran compilers and BLAS routines from MKL library), NVIDIA CUDA  
 488 v11.0.2, and MAGMA version 2.5.4. The input symmetric matrix  $A \equiv [a_{ij}]$  was gen-  
 489 erated with random elements from a uniform distribution in range  $(0, 1)$ :  $a_{ij} \sim \mathcal{U}(0, 1)$   
 490 and  $a_{ij} = a_{ji}$ . The Hermitian matrix is also generated in the same fashion for its  
 491 imaginary part. The largest eigenvalues in the spectrum were requested. The blocked  
 492 *SICE-SM* algorithm was implemented in both PLASMA and MAGMA.

493 First, we show the profiling results from the PLASMA experiments in left of Fig-  
 494 ure 2. PLASMA was used in a CPU-only mode and no GPUs were used in the system.  
 495 The symmetric input matrix had size  $n = 10000$ . The three stacked bars represent the  
 496 breakdown of time from mixed-precision with refinement, single precision, and double  
 497 precision from the two-stage algorithm, respectively. The time for single precision is  
 498 about half of that of double precision and each of the components take proportion-  
 499 ally the same time for both precisions. The mixed-precision algorithm is slower than  
 500 double precision in this setup because of the requirement of explicitly forming the  
 501 transformation matrices from the first and second stages. They also take much more  
 502 time compared to the double precision algorithm, which only applies transformations  
 503 to the eigenvectors.

504 Figure 3 shows the performance results from the MAGMA. First the solid lines  
 505 are the one-stage algorithm in double, single, and mixed precision (with iterative  
 506 refinement). The input matrix sizes range from 1000 to 20000, and the largest 32  
 507 eigenpairs are requested. Single precision is about  $1.7\times$  faster than double precision






FIG. 2. Breakdown of timings of two-stage eigensolvers with 32 largest eigenpairs requested in PLASMA (left), MAGMA on NVIDIA Volta V100, and MAGMA on NVIDIA GTX1060. The problem sizes are 10 000, 20 000, and 12 000, respectively.

508 and the mixed precision is about  $1.3\times$  faster. The dashed lines represent the two-  
 509 stage algorithm. They are at least  $2\times$  faster than their corresponding single stage  
 510 algorithm in general. The performance improvement over double precision is about  
 511  $1.2\times$ . Figure 4 shows the performance results of complex Hermitian solvers. Complex  
 512 operations has higher arithmetic intensity so the performance gap between single and  
 513 double would also be larger. Mixed precision algorithm can also have greater chance  
 514 to benefit it. On the system with NVIDIA V100, we are observing complex single is  
 515  $2.44\times$  faster than complex double and mixed precision solver is  $1.45\times$

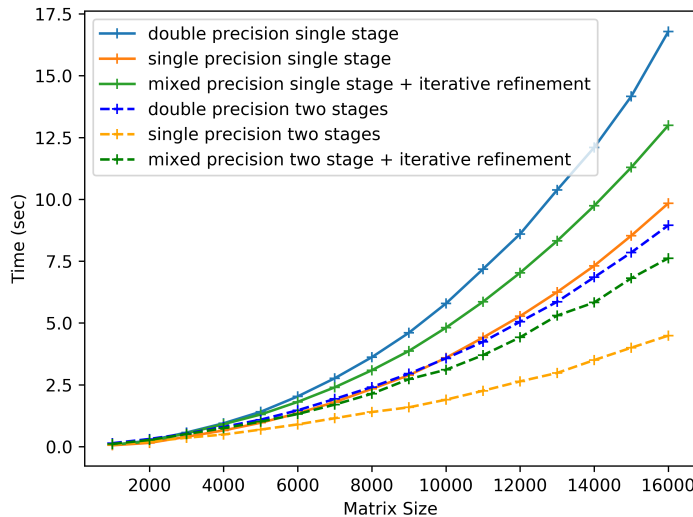


FIG. 3. Performance comparison of single, double, and mixed precision solvers for real symmetric matrix on MAGMA for both single stage and two-stage algorithms on NVIDIA V100 GPU with varying sizes of matrices and fixed number of requested eigenpairs.

516 Figure 5 shows the performance when requesting different numbers of eigenpairs  
 517 with the input matrix size fixed at  $n = 20000$ . Mixed precision is noticeably faster  
 518 than double precision if 64 or fewer eigenpairs are requested. For larger eigenpair  
 519 count, the time in iterative refinement grows linearly with the number of requested  
 520 eigenpairs and it eventually loses its performance advantage.

521 The middle of Figure 2 shows the detailed profile for matrix size  $n = 20000$  and




FIG. 4. Performance of single, double, and mixed precision solvers for complex Hermitian matrix based on MAGMA two-stage algorithm on NVIDIA V100 GPU with varying sizes of matrices and fixed number of requested eigenpairs.

522 32 eigenvalues/eigenvectors requested. The details of computational components were  
 523 explained in Section 4.3. The single precision routine took 60% of time compared to  
 524 double, and the ratios between components across precisions were about the same. For  
 525 mixed precision, there is a 0.5 second overhead at the beginning to convert the whole  
 526 matrix from double to single precision. Then the two-stage reduction is done in single  
 527 precision which is about twice as fast in single precision. The back-transformation of  
 528 the first stage is overlapped with the second stage, and is not shown in the bar. The  
 529 same applies for the eigensolver, which is overlapped with the back-transformation  
 530 from the second stage. Finally, at the top is the timing for the iterative refinement  
 531 stage. As can be easily observed, the back transformation of second stage for mixed  
 532 precision is the bottleneck as it takes almost 40% of the total time in this case.

533 We tested another machine with a drastically different setup by using a consumer-  
 534 grade gaming GPU. It has the same CPUs as the V100 system. The GPU is NVIDIA  
 535 GTX1060 6GB GPU. The theoretical peak performance of GTX1060 is 136.7 GFLOP/s  
 536 in double and 4.375 TFLOP/s in single precision. This is a notable different as the  
 537 gaming maintains 1:32 double-single ratio compared to server-grade NVIDIA V100  
 538 with the ratio being 1:2. Figure ?? shows the performance with different matrix sizes  
 539 on GTX1060 when requesting the largest 32 eigenpairs. The performance of single  
 540 precision is about  $8\times$  better than that of double precision and the mixed precision  
 541 with refinement is about  $2\times$  better than double precision. Figure 6 is the complex  
 542 Hermitian solver and the the speed up over complex double is  $3.6\times$  as In Figure 7 we  
 543 show performance results when the matrix size was fixed at  $n = 12000$  but with varied  
 544 number of requested eigenpairs. The mixed precision solver is still faster than dou-  
 545 ble precision when 128 eigenpairs are requested, but the time in iterative refinement  
 546 became significant if more eigenvalues and eigenvectors were requested.

547 The right of Figure 8 shows the profiling results with timing breakdown for ma-




FIG. 5. Performance comparison of single, double, and mixed precision solvers on top of MAGMA on NVIDIA V100 GPU with varying number of requested eigenpairs and fixed matrix size.

548 trix size  $n = 12000$  and the 32 largest eigenpairs requested. In double precision,  
 549 almost 80% of time was spent at the first stage to reduce the matrix from symmetric to  
 550 band-symmetric form. The operation is compute-bound and relies on GPU’s  
 551 matrix-matrix multiplication efficiency. But the consumer-grade GPU does not have  
 552 hardware to support high-efficiency processing for the double floating-point units and  
 553 consequently extra clock cycles are used to emulate higher precision with single pre-  
 554 cision instructions. The mixed-precision algorithm does the first-stage reduction in  
 555 single precision and does not suffer from the same penalty. The back-transformation  
 556 of second stage is still costly but it is done with single precision on the GPU. Over-  
 557 all, the performance of mixed precision with the iterative refinement algorithm is  $2\times$   
 558 faster over purely double two-stage algorithm.

559 **6. Conclusions and Future Work.** We developed an iterative refinement  
 560 algorithm for symmetric and Hermitian eigenvalue problems based on the initial work  
 561 from the SICE algorithm. By utilizing the Sherman–Morrison formula, our new solver  
 562 has more opportunity to be parallelized compared to the serial Givens rotations in the  
 563 SICE algorithm. The blocked version of the algorithm was also proposed in order to  
 564 refine multiple pairs of eigenvalues and eigenvectors simultaneously for higher utiliza-  
 565 tion of the computational resources with lower demand for memory bandwidth. The  
 566 implementation of the mixed-precision algorithm is based on the two-stage eigen-  
 567 solver in either the PLASMA and MAGMA software libraries for numerical linear  
 568 algebra, which gives our implementation the advantage of both portability and per-  
 569 formance. The computational components inside the mixed-precision algorithm have  
 570 been reordered to create more parallelism at runtime and allow additional overlap to  
 571 computational stages more efficiently. Compared to the double-precision solver, the  
 572 performance benefit has been shown for the cases in which only a portion of eigenval-

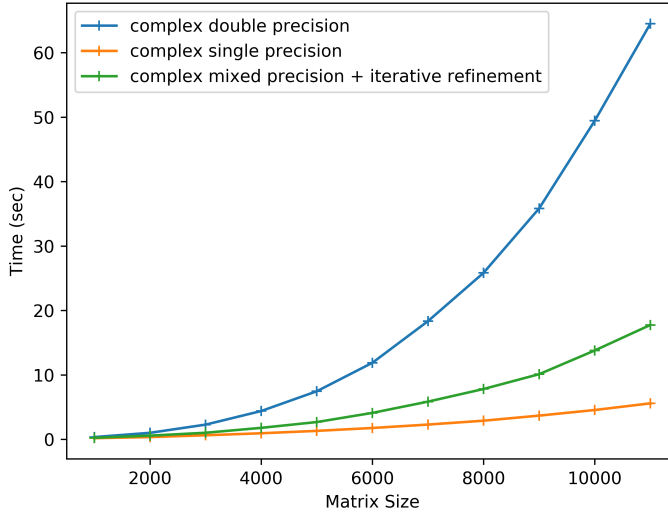


FIG. 6. Performance of single, double, and mixed precision solvers for complex Hermitian matrix based on MAGMA two-stage algorithm on the NVIDIA GTX1060 GPU.

573 ues and corresponding eigenvectors are requested. This remains true across hardware  
 574 with a varying ratio of performance of single and double precision units.

575 As we can see in the profiling result featuring time breakdown of the computa-  
 576 tional tasks, the back-transformation of the second stage that performs bulge chasing  
 577 is slow on either CPU or GPU and becomes the bottleneck for some experiments.  
 578 Although the two stage reduction is a far superior method in terms of performance, if  
 579 only the forward transforms are considered then back-transformations take over the  
 580 performance and must be taken into account while designing mixed-precision algo-  
 581 rithms. One possible approach would be to start aggregating the transformations on  
 582 the GPU as soon as they are generated by GPU-based bulge chasing and not wait  
 583 until all the reductions have been computed.

584 For distributed systems, the matrix is usually too large and it might not be  
 585 feasible to explicitly form the transform matrix. Consequently, the cost of applying  
 586 the transformation  $Q$  during iterative refinement needs to be reevaluated. Also, if  
 587 different eigenpairs are being distributed and refined on different nodes, synchronizing  
 588 and applying  $Q$  to eigenvectors across disparate nodes needs to be designed and  
 589 implemented with care as this is not a usual operation.

590 Another direction is to try different low-precision formats in addition to just  
 591 mixing single and double precisions. The recently released NVIDIA Ampere GPU  
 592 provides TF32 Tensor Cores, which uses all 8 exponent bits and 10 out of 23 mantissa  
 593 bits from the FP32 single precision format, and thus offering  $8\times$  speedup. Because  
 594 our initial eigenpairs and the reduced systems are all coming from the low-precision  
 595 tridiagonalization, the convergence rate of the iterative refinement is affected signifi-  
 596 cantly. Based on our experiments, the FP16 half-precision tensor cores do not provide  
 597 sufficient accuracy and TF32 might appear to be a more promising target with more  
 598 balanced mix of precision and performance.

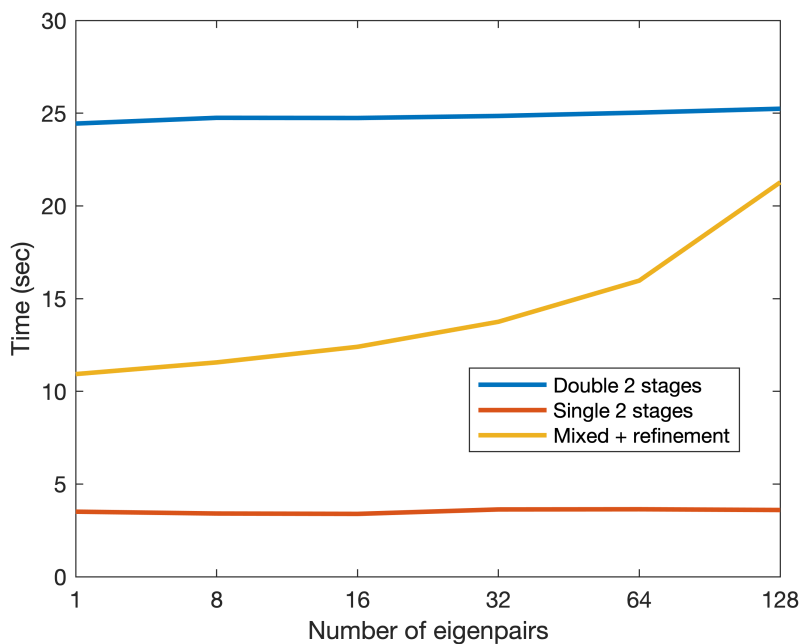


FIG. 7. Performance comparison of single, double, and mixed precision solvers on top of MAGMA on NVIDIA GTX1060 GPU with varying number of requested eigenpairs and fixed matrix size  $n = 12000$ .

599

#### REFERENCES

- 600 [1] A. ABDELFAH, T. COSTA, J. DONGARRA, M. GATES, A. HAIDAR, S. HAMMARLING, N. J.  
 601 HIGHAM, J. KURZAK, P. LUSZCZEK, S. TOMOV, AND M. ZOUNON, *A set of batched basic*  
 602 *linear algebra subprograms and LAPACK routines*, ACM TOMS, 47 (2020), p. 1–23, <https://doi.org/10.1145/3431921>. DOI: 10.1145/3431921.  
 603  
 604 [2] E. AGULLO, J. DONGARRA, B. HADRI, J. KURZAK, J. LANGOU, J. LANGOU, H. LTAIEF,  
 605 P. LUSZCZEK, AND A. YARKHAN, *PLASMA users guide*, tech. report, Technical report,  
 606 ICL, UTK, 2009.  
 607 [3] G. ALEFELD AND H. SPREUER, *Iterative improvement of componentwise errorbounds for invariant*  
 608 *subspaces belonging to a double or nearly double eigenvalue*, Computing, 36 (1986),  
 609 pp. 321–334.  
 610 [4] *AMD optimizing CPU libraries (AOCL)*. <https://developer.amd.com/amd-aocl/>, 2020. [On-  
 611 line; accessed 2-August-2020].  
 612 [5] *rocBLAS: AMD’s library for BLAS on ROCm*. [https://github.com/ROCmSoftwarePlatform/](https://github.com/ROCmSoftwarePlatform/rocBLAS)  
 613 [rocBLAS](https://github.com/ROCmSoftwarePlatform/rocBLAS), 2020. [Online; accessed 2-August-2020].  
 614 [6] E. ANDERSON, Z. BAI, C. BISCHOF, L. S. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DU CROZ,  
 615 A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, ET AL., *LAPACK Users’ guide*, SIAM,  
 616 1999.  
 617 [7] *ARM performance libraries*. [https://developer.arm.com/tools-and-software/server-and-hpc/](https://developer.arm.com/tools-and-software/server-and-hpc/compile/arm-compiler-for-linux/arm-performance-libraries)  
 618 [compile/arm-compiler-for-linux/arm-performance-libraries](https://developer.arm.com/tools-and-software/server-and-hpc/compile/arm-compiler-for-linux/arm-performance-libraries), 2020. [Online; accessed 2-  
 619 August-2020].  
 620 [8] C. H. BISCHOF, B. LANG, AND X. SUN, *A framework for symmetric band reduction*, ACM  
 621 Transactions on Mathematical Software (TOMS), 26 (2000), pp. 581–601.  
 622 [9] L. S. BLACKFORD, J. CHOI, A. CLEARY, E. D’AZEVEDO, J. DEMMEL, I. DHILLON, J. DONGARRA,  
 623 S. HAMMARLING, G. HENRY, A. PETITET, ET AL., *ScaLAPACK users’ guide*, SIAM, 1997.  
 624 [10] E. CARSON AND N. J. HIGHAM, *A new analysis of iterative refinement and its application*  
 625 *to accurate solution of ill-conditioned sparse linear systems*, SIAM Journal on Scientific  
 626 Computing, 39 (2017), pp. A2834–A2856.  
 627 [11] E. CARSON AND N. J. HIGHAM, *Accelerating the solution of linear systems by iterative refine-*  
 628 *ment in three precisions*, SIAM Journal on Scientific Computing, 40 (2018), pp. A817–




FIG. 8. Breakdown of timings of two-stage eigensolvers based on MAGMA on the NVIDIA GTX1060 GPU with size  $n = 12000$  and 32 largest eigenpairs requested.

- 629 A847.
- 630 [12] L.-W. CHANG, J. A. STRATTON, H.-S. KIM, AND W.-M. W. HWU, *A scalable, numerically stable, high-performance tridiagonal solver using GPUs*, in SC'12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, IEEE, 2012, pp. 1–11.
- 631
- 632
- 633
- 634 [13] F. CHATELIN, *Simultaneous Newton's iteration for the eigenproblem*, in Defect correction methods, Springer, 1984, pp. 67–74.
- 635
- 636 [14] J. J. CUPPEN, *A divide and conquer method for the symmetric tridiagonal eigenproblem*, Numerische Mathematik, 36 (1980), pp. 177–195.
- 637
- 638 [15] A. DAVIDSON, Y. ZHANG, AND J. D. OWENS, *An auto-tuned method for solving large tridiagonal systems on the GPU*, in 2011 IEEE International Parallel & Distributed Processing Symposium, IEEE, 2011, pp. 956–965.
- 639
- 640 [16] J. W. DEMMEL, *Three methods for refining estimates of invariant subspaces*, Computing, 38 (1987), pp. 43–57.
- 641
- 642
- 643 [17] J. DONGARRA, S. HAMMARLING, N. J. HIGHAM, S. D. RELTON, P. VALERO-LARA, AND M. ZOUNON, *The design and performance of batched BLAS on modern high-performance computing systems*, Procedia Computer Science, 108 (2017), pp. 495–504.
- 644
- 645
- 646 [18] J. J. DONGARRA, *Improving the accuracy of computed matrix eigenvalues*, tech. report, Argonne National Lab., Chicago, IL, USA, 1980.
- 647
- 648 [19] J. J. DONGARRA, *Algorithm 589: SICEDR: a FORTRAN subroutine for improving the accuracy of computed matrix eigenvalues*, ACM Transactions on Mathematical Software (TOMS), 8 (1982), pp. 371–375.
- 649
- 650 [20] J. J. DONGARRA, C. B. MOLER, AND J. H. WILKINSON, *Improving the accuracy of computed eigenvalues and eigenvectors*, SIAM Journal on Numerical Analysis, 20 (1983), pp. 23–45.
- 651
- 652 [21] J. 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.
- 653
- 654 [22] M. GATES, J. KURZAK, A. CHARARA, A. YARKHAN, AND J. DONGARRA, *SLATE: design of a modern distributed and accelerated linear algebra library*, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 2019, pp. 1–18.
- 655
- 656 [23] K. GOTO AND R. VAN DE GEIJN, *High-performance implementation of the Level 3 BLAS*, ACM Transactions on Mathematical Software (TOMS), 35 (2008), pp. 1–14.
- 657
- 658 [24] A. HAIDAR, H. LTAIEF, AND J. DONGARRA, *Parallel reduction to condensed forms for symmetric eigenvalue problems using aggregated fine-grained and memory-aware kernels*, in
- 659
- 660
- 661
- 662
- 663

- 664 Proceedings of 2011 International Conference for High Performance Computing, Network-  
665 ing, Storage and Analysis, 2011, pp. 1–11.
- 666 [25] A. HAIDAR, S. TOMOV, J. DONGARRA, AND N. J. HIGHAM, *Harnessing GPU tensor cores for fast*  
667 *fp16 arithmetic to speed up mixed-precision iterative refinement solvers*, in SC18: Inter-  
668 national Conference for High Performance Computing, Networking, Storage and Analysis,  
669 IEEE, 2018, pp. 603–613.
- 670 [26] A. HAIDAR, S. TOMOV, J. DONGARRA, R. SOLCÀ, AND T. SCHULTHESS, *A novel hybrid CPU–*  
671 *GPU generalized eigensolver for electronic structure calculations based on fine-grained*  
672 *memory aware tasks*, The International journal of high performance computing applica-  
673 tions, 28 (2014), pp. 196–209.
- 674 [27] *IBM engineering and scientific subroutine library (ESSL) version 6.3*. [https://www.ibm.com/](https://www.ibm.com/support/knowledgecenter/SSFHY8.6.3/navigation/welcome.html)  
675 [support/knowledgecenter/SSFHY8.6.3/navigation/welcome.html](https://www.ibm.com/support/knowledgecenter/SSFHY8.6.3/navigation/welcome.html), 2020. [Online; accessed  
676 2-August-2020].
- 677 [28] *Intel Math Kernel Library*. [https://software.intel.com/content/www/us/en/develop/tools/](https://software.intel.com/content/www/us/en/develop/tools/math-kernel-library.html)  
678 [math-kernel-library.html](https://software.intel.com/content/www/us/en/develop/tools/math-kernel-library.html), 2020. [Online; accessed 2-August-2020].
- 679 [29] C. L. LAWSON, R. J. HANSON, D. R. KINCAID, AND F. T. KROGH, *Basic linear algebra subpro-*  
680 *grams for fortran usage*, ACM Transactions on Mathematical Software (TOMS), 5 (1979),  
681 pp. 308–323.
- 682 [30] *NVIDIA cuBLAS*. <https://docs.nvidia.com/cuda/cublas/index.html>, 2020. [Online; accessed  
683 2-August-2020].
- 684 [31] T. OGITA AND K. AISHIMA, *Iterative refinement for symmetric eigenvalue decomposition*, Japan  
685 Journal of Industrial and Applied Mathematics, 35 (2018), pp. 1007–1035.
- 686 [32] T. OGITA AND K. AISHIMA, *Iterative refinement for symmetric eigenvalue decomposition II:*  
687 *clustered eigenvalues*, Japan Journal of Industrial and Applied Mathematics, 36 (2019),  
688 pp. 435–459.
- 689 [33] T. OGITA AND K. AISHIMA, *Iterative refinement for singular value decomposition based on*  
690 *matrix multiplication*, Journal of Computational and Applied Mathematics, 369 (2020),  
691 p. 112512.
- 692 [34] *OpenBLAS*. <https://github.com/xianyi/OpenBLAS>, 2020. [Online; accessed 2-August-2020].
- 693 [35] B. N. PARLETT AND I. S. DHILLON, *Relatively robust representations of symmetric tridiagonals*,  
694 Linear Algebra and its applications, 309 (2000), pp. 121–151.
- 695 [36] G. PETERS AND J. H. WILKINSON, *Inverse iteration, ill-conditioned equations and Newton’s*  
696 *method*, SIAM review, 21 (1979), pp. 339–360.
- 697 [37] K. E. PRIKOPA AND W. N. GANSTERER, *On mixed precision iterative refinement for eigenvalue*  
698 *problems*, Procedia Computer Science, 18 (2013), pp. 2647–2650.
- 699 [38] J. SHERMAN AND W. J. MORRISON, *Adjustment of an inverse matrix corresponding to a change*  
700 *in one element of a given matrix*, The Annals of Mathematical Statistics, 21 (1950),  
701 pp. 124–127.
- 702 [39] B. T. SMITH, J. M. BOYLE, B. GARBOW, Y. IKEBE, V. KLEMA, AND C. MOLER, *Matrix eigen-*  
703 *system routines-EISPACK guide*, vol. 6, Springer, 2013.
- 704 [40] G. W. STEWART, *Error and perturbation bounds for subspaces associated with certain eigen-*  
705 *value problems*, SIAM review, 15 (1973), pp. 727–764.
- 706 [41] H. SYMM AND J. H. WILKINSON, *Realistic error bounds for a simple eigenvalue and its associ-*  
707 *ated eigenvector*, Numerische Mathematik, 35 (1980), pp. 113–126.
- 708 [42] F. TISSEUR, *Newton’s method in floating point arithmetic and iterative refinement of gener-*  
709 *alized eigenvalue problems*, SIAM Journal on Matrix Analysis and Applications, 22 (2001),  
710 pp. 1038–1057.
- 711 [43] S. TOMOV, J. DONGARRA, AND M. BABOULIN, *Towards dense linear algebra for hybrid GPU*  
712 *accelerated manycore systems*, Parallel Computing, 36 (2010), pp. 232–240.
- 713 [44] F. G. VAN ZEE AND R. A. VAN DE GEIJN, *BLIS: a framework for rapidly instantiating BLAS*  
714 *functionality*, ACM Transactions on Mathematical Software (TOMS), 41 (2015), pp. 1–33.
- 715 [45] J. WILKINSON, *Calculation of the eigenvalues of a symmetric tridiagonal matrix by the method*  
716 *of bisection*, Numerische Mathematik, 4 (1962), pp. 362–367.
- 717 [46] J. H. WILKINSON, *Global convergene of tridiagonal QR algorithm with origin shifts*, Linear  
718 Algebra and its Applications, 1 (1968), pp. 409–420.
- 719 [47] T. YAMAMOTO, *Error bounds for computed eigenvalues and eigenvectors*, Numerische Mathe-  
720 matik, 34 (1980), pp. 189–199.
- 721 [48] Y. ZHANG, J. COHEN, AND J. D. OWENS, *Fast tridiagonal solvers on the GPU*, ACM Sigplan  
722 Notices, 45 (2010), pp. 127–136.