

# Implementing Singular Value and Symmetric/Hermitian Eigenvalue Solvers 

Mark Gates<br>Kadir Akbudak<br>Mohammed Al Farhan<br>Ali Charara<br>Jakub Kurzak<br>Dalal Sukkari<br>Asim YarKhan<br>Jack Dongarra<br>Innovative Computing Laboratory

June 26, 2023

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

This research used 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-AC05-000R22725.

| Revision | Notes |
| :--- | :--- |
| $2019-09$ | first publication |
| 2020-04 | added generalized Hermitian definite eigenvalues (Section 2.3) and eigenvectors <br> (Section 2.6) |
| $2021-12$ | added divide and conquer (Chapter 3) and optimization (Chapter 4) |
| $2022-11$ | added details of Hermitian to Hermitian band (Section 2.5) |
| $2023-06$ | added details of divide and conquer |

```
@techreport{gates2019implementing,
    author={Gates, Mark and Akbudak, Kadir and Al Farhan, Mohammed
                and Charara, Ali and Kurzak, Jakub and Sukkari, Dalal
                and YarKhan, Asim and Dongarra, Jack},
    title={{SLATE} Working Note 13:
            Implementing Singular Value and Symmetric/Hermitian Eigenvalue Solvers},
    institution={Innovative Computing Laboratory, University of Tennessee},
    year={2023},
    month={June},
    number={ICL-UT-19-07},
    note={first published 2019-09, revision 2023-06}
}
```


## Contents

Contents ..... ii
List of Figures ..... iii
1 Introduction ..... 1
1.1 Significance of SLATE ..... 1
1.2 Design of SLATE ..... 3
2 Implementation ..... 6
2.1 Singular Value Decomposition ..... 6
2.2 Hermitian Eigenvalue Problem ..... 7
2.3 Generalized Hermitian Definite Eigenvalue Problem ..... 7
2.4 Three Stage Algorithms ..... 8
2.5 Hermitian to Hermitian band reduction (he2hb) ..... 14
2.5.1 Single node ..... 14
2.5.2 Multi-node ..... 15
2.6 Eigenvector Computation ..... 21
2.6.1 Eigenvectors of tridiagonal matrix ..... 21
2.6.2 Second stage back-transformation ..... 22
2.6.3 First stage back-transformation ..... 23
3 Divide and conquer ..... 26
3.1 Cuppen's method ..... 26
3.1.1 Historical note on choice of $\theta$ ..... 27
3.1.2 Secular equation ..... 28
3.1.3 Deflation ..... 29
3.1.4 Back-transformation ..... 30
3.1.5 Summary ..... 33
3.1.6 Examples ..... 34
3.1.7 Example 0: no deflation ..... 36
3.1.8 Example 1: type 1 deflation - serial ..... 37
3.1.9 Example 2: type 2 deflation - serial ..... 39
3.1.10 Example 3: type 1 and 2 deflation - serial ..... 41
3.1.11 Eigenvectors via Löwner Theorem ..... 43
3.1.12 Cost ..... 45
3.2 Routines ..... 46
3.2.1 stedc ..... 46
3.2.2 stedc_solve ..... 47
3.2.3 stedc merge ..... 48
3.2.4 stedc_z_vector ..... 49
3.2.5 stedc_deflate ..... 50
3.2.6 stedc_secular ..... 52
3.2.7 stedc_sort ..... 53
4 Optimization ..... 55
4.1 Hermitian to Hermitian band reduction (he2hb) ..... 55
4.2 Hermitian to Hermitian band reduction (he2hb) ..... 55
4.3 Back-transformation (unmtr_hb2st) ..... 56
5 Performance ..... 59
5.1 Environment ..... 59
5.1.1 Hardware ..... 59
5.1.2 Software ..... 59
5.2 Results ..... 60
Bibliography ..... 63

## List of Figures

1.1 Dependencies of ECP applications on dense linear algebra software. ..... 2
1.2 SLATE in the ECP software stack. ..... 3
1.3 Code size comparison - ScaLAPACK vs SLATE ..... 5
2.1 Three stage Hermitian eigenvalue and SVD algorithms. ..... 11
2.2 One panel of the first stage reduction to band form ..... 11
2.3 Bulge-chasing algorithm. ..... 12
2.4 Hermitian bulge-chasing algorithm. ..... 12
2.5 2D block-cyclic distribution. ..... 15
2.6 Hermitian matrix multiply $W_{a}=A V_{a}$ (hemm). Tiles that have 2 colors (blue/red or green/red) are computed as partial sums on 2 nodes, then reduced. ..... 16
2.7 Hermitian rank $2 k$ (her2k) update, $A_{k}-=V_{a} W_{a}^{H}+W_{a} V_{a}^{H}$ ..... 17
2.8 Hermitian matrix multiply $W_{b}=A V_{b}$ (hemm). ..... 18
2.9 Hermitian rank $2 k$ (her2k) update, $A_{k}-=V_{b} W_{b}^{H}+W_{b} V_{b}^{H}$. ..... 18
2.10 Hermitian matrix multiply $W_{c}=A V_{c}$ (hemm). ..... 19
2.11 Hermitian rank $2 k$ (her2k) update, $A_{k}-=V_{c} W_{c}^{H}+W_{c} V_{c}^{H}$. ..... 19
2.12 Nodes and tiles that are updated by each $Q_{r}$ from panel. ..... 20
2.13 Redistribute 1D block row cyclic distributed matrix using $4 \times 1$ grid into a 2D block cyclic distribution using $2 \times 2$ grid ..... 22
2.14 Second stage back transformation, with $V$ block size $j_{b}=3$ vectors. Block reflector 3 is highlighted to show overlap. ..... 23
2.15 Dependencies allow up to $\left\lceil\frac{m t}{2}\right\rceil$ parallel tasks. ..... 24
4.1 Performance of he2hb Using 1 and 2 nodes on Summit, $1 \times 1$ and $2 \times 2$ process grids. ..... 56
4.2 Overlap the panel factorization with subsequent data movements. ..... 57
4.3 The performance impact by overlapping the panel factorization with subsequent data movements ..... 57
4.4 Performance results of Hermitian to Hermitian band reduction, using 1 node , $1 \times 1$ process grid. $\mathrm{nb}=128,320$, $\mathrm{ib}=16,48$ for $C P U, G P U$ tests, panel-threads $=10$ ..... 57
4.5 Performance of unmtr_hb2st on a host with two 20-core Intel Broadwell Xeon E5-2698 v4 CPUs and one NVIDIA V100 activated. $\mathrm{N}=16384$. As seen in the figure, the device implementation provides up to 6 x speedup.58
5.1 Summit node architecture. ..... 60
5.2 SVD performance comparison. ..... 61
5.3 Generalized to standard eigenvalue performance comparison. ..... 62
5.4 Profile of eigenvalue solver implementations showing each phase for $\mathrm{N}=12288$. One node of Summit is used. ..... 62

## CHAPTER 1

## Introduction

### 1.1 Significance of SLATE

Software for Linear Algebra Targeting Exascale (SLATE) ${ }^{1}$ is being developed as part of the Exascale Computing Project (ECP) ${ }^{2}$, which is a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration (NNSA). SLATE will deliver fundamental dense linear algebra capabilities for current and upcoming distributed-memory systems, including GPU-accelerated systems as well as more traditional multi core-only systems.

SLATE will provide coverage of existing LAPACK and ScaLAPACK functionality, including parallel implementations of Basic Linear Algebra Subroutines (BLAS), linear systems solvers, least squares solvers, and singular value and eigenvalue solvers. In this respect, SLATE will serve as a replacement for LAPACK and ScaLAPACK, which, after two decades of operation, cannot be adequately retrofitted for modern, GPU-accelerated architectures.

Figure 1.1 shows how heavily ECP applications depend on dense linear algebra software. A direct dependency means that the application's source code contains calls to the library's routines. An indirect dependency means that the applications needs to be linked with the library due to another component depending on it. Out of 60 ECP applications, 38 depend on BLAS - either directly on indirectly - 40 depend on LAPACK, and 14 depend on ScaLAPACK. In other words, the use of dense linear algebra software is ubiquitous among the ECP applications.

[^0]| Application | BLAS | LAPACK | SCALAPACK |
| :---: | :---: | :---: | :---: |
| AMPE | Q INDIRECT | © INDIRECT |  |
| AMReX | Q INDIRECT | Q INDIRECT | Q INDIRECT |
| CANDLE | C INDIRECT | O INDIREGT |  |
| CEED-MAGMA | Q DIRECT | O DIRECT |  |
| CEED-MFEM | C INOIRECT | O INOIRECT | O INDIRECT |
| CEED-Nek5000 | Q IDIRECT | O DIRECT |  |
| CEED-OCCA |  |  |  |
| CEED-PUMI |  |  |  |
| Chroma | Q DIRECT | © DIRECT |  |
| Combustion-PELE | Q INDIRECT | O INDIRECT |  |
| CPS |  |  |  |
| Diablo | Q INDIRECT | © INDIRECT |  |
| E3SM-MMF-ACME-MMF |  |  |  |
| EQSIM-SW4 |  |  |  |
| ExaBiome-GOTTCHA |  |  |  |
| ExaBiome-HipMCL |  | O DIRECT |  |
| ExaBiome-MetaHipMer |  |  |  |
| ExaCA | (G) IDIRECT | © INDIRECT |  |
| ExaConstit | Q DIRECT | O DIRECT |  |
| ExaFEL-LUNUS | Q INDIRECT | Q INDIRECT |  |
| ExaFEL-M-TIP | O DIRECT |  | (8) DIRECT |
| ExaFEL-psana | O INDIRECT | Q Indirect |  |
| ExaGraph-AWPM |  | O DIRECT |  |
| ExaGraph-HipMCL |  | © DIRECT |  |
| ExaGraph-Kokkoskernels |  |  |  |
| ExaGraph-Zoltan2 | Q INDIRECT | Q INDIRECT | Q INDIRECT |
| ExaMPM |  |  |  |
| ExaSGD-GOSS |  |  |  |
| ExaSGD-GridPACK | Q INDIRECT | Q INDIRECT | O INDIREGT |
| ExaSGD-PIPS | Q DIRECT | © DIRECT | (6) DIRECT |
| ExaSGD-StructJuMP |  |  |  |
| ExaSky-HACC/CosmoTools |  |  |  |
| ExaSky-Nyx | Q INIRECT | © INDIRECT |  |
| ExaSMD-Nek5000 |  | © DIRECT |  |
| ExaSMR-OpenMC |  |  |  |
| ExaSMR-Shift | Q DIRECT | Q DIRECT |  |
| ExaStar-Castro | Q INDIRECT | Q INDIRECT |  |
| ExaStar-FLASH | Q DIRECT | C INDIRECT |  |
| ExaWind-Nalu | Q INDIRECT | C INDIRECT | CS INDIRECT |
| GAMESS |  | © DIRECT |  |
| LAMMPS | Q DIRECT | © DIRECT |  |
| LATTE | Q DIRECT | © DIRECT |  |
| LIBCCHEM | Q DIRECT | © DIRECT |  |
| MEUMAPPS-SL | C DIRECT |  |  |
| MEUMAPPS-SS | O DIRECT |  |  |
| MFIX-Exa | C INDIRECT | C INDIRECT |  |
| MILS | O DIRECT | © DIRECT |  |
| NWChemEx | Q DIRECT | O DIRECT | © DIRECT |
| ParSplice |  |  |  |
| PICSAR |  |  |  |
| QMCPACK | C DIRECT | © DIRECT | C DIRECT |
| Subsurface-Chombo-Crunch | O DIRECT | O DIRECT | (6) INDIRECT |
| Subsurface-GEOS | Q INDIRECT | Q INDIRECT | Q INDIRECT |
| Truchas-PBF | O INDIRECT | Q INDIRECT |  |
| Tusas | Q DIRECT | © INDIRECT | Q INDIRECT |
| Urban-WRF |  |  |  |
| WarpX | Q INDIRECT | C INDIRECT |  |
| WCMAPP-XGC | C DIRECT | © DIRECT | Q INDIRECT |
| WDMApp-GENE | C DIRECT | C INDIRECT | (G) DIRECT |
| xaFEL-CCTBX |  |  |  |

Figure 1.1: Dependencies of ECP applications on dense linear algebra software.

### 1.2 Design of SLATE

SLATE is built on top of standards, such as MPI and OpenMP, and de facto standard industry solutions such as NVIDIA CUDA and AMD HIP. SLATE also relies on high performance implementations of numerical kernels from vendor libraries, such as Intel oneMKL, IBM ESSL, NVIDIA cuBLAS, and AMD rocBLAS. SLATE interacts with these libraries through a layer of C++ APIs. Figure 1.2 shows SLATE's position in the ECP software stack.


Figure 1.2: SLATE in the ECP software stack.

The following paragraphs outline the foundations of SLATE's design.

Object-Oriented Design: The design of SLATE revolves around the Tile class and the Matrix class hierarchy. The Tile class is intended as a simple class for maintaining the properties of individual tiles and implementing core serial tile operations, such as tile BLAS, while the Matrix class hierarchy maintains the state of distributed matrices throughout the execution of parallel matrix algorithms in a distributed-memory environment. Currently, the classes are structured as follows:

BaseMatrix is an abstract base class for all matrices.
Matrix represents a general $m \times n$ matrix.
BaseTrapezoidMatrix is an abstract base class for all matrices stored as uppertrapezoid or lower-trapezoid. For upper matrices, tiles $A(i, j)$ are stored for $i \leq j$. For lower matrices, tiles $A(i, j)$ are stored for $i \geq j$.
TrapezoidMatrix represents an upper-trapezoid or a lower-trapezoid, $m \times n$ matrix. The opposite triangle is implicitly zero.
TriangularMatrix represents an upper-triangular or a lower-triangular, $n \times n$ matrix.
SymmetricMatrix represents a symmetric, $n \times n$ matrix, with only the upper or lower triangle stored. The opposite triangle is implicitly known by symmetry $\left(A_{j, i}=A_{i, j}\right)$.
HermitianMatrix represents a Hermitian, $n \times n$ matrix, with only the upper or lower triangle stored. The opposite triangle is implicitly known by symmetry $\left(A_{j, i}=\bar{A}_{i, j}\right)$.

Tiled Matrix Layout: The new matrix storage introduced in SLATE is one of its most impactful features. In this respect, SLATE represents a radical departure from other distributed linear algebra software such as ScaLAPACK or Elemental, where the local matrix occupies a contiguous memory region on each process. In contrast, tiles are first class objects in SLATE that can be individually allocated and passed to low-level tile routines. In SLATE, the matrix consists of a collection of individual tiles with no correlation between their positions in the matrix and their memory locations. At the same time, SLATE also supports tiles pointing to data in a traditional ScaLAPACK matrix layout, thereby easing an application's transition from ScaLAPACK to SLATE.

Handling of side, uplo, trans: The classical BLAS takes parameters such as side, uplo, trans (named "op" in SLATE), and diag to specify operation variants. Traditionally, this has meant that implementations have numerous cases. The reference BLAS has nine cases in zgemm and eight cases in ztrmm (times several sub-cases). ScaLAPACK and PLASMA likewise have eight cases in ztrmm. In contrast, by storing both uplo and op within the matrix object itself, and supporting inexpensive shallow copy transposition, SLATE can implement just one or two cases and map all the other cases to that implementation by appropriate transpositions. For instance, SLATE only implements one case for gemm (NoTrans, NoTrans) and handles all other cases by swapping indices of tiles and setting trans appropriately for the underlying tile operations.

Templating of Precisions: SLATE handles multiple precisions by C++ templating, so there is only one precision-independent version of the code, which is then instantiated for the desired precisions. Operations are defined so that they can be applied consistently across all precisions. SLATE's BLAS++ component provides overloaded, precision-independent wrappers for all underlying, node-level BLAS, and SLATE's PBLAS are built on top of these. Currently, the SLATE library has explicit instantiations of the four main data types: float, double, std::complex<float>, and std::complex<double>. The SLATE code should be able to accommodate other data types, such as half, double-double, or quad precision, given appropriate underlying node-level BLAS.

Templating of Execution Targets: Parallelism is expressed in SLATE's computational routines. Each computational routine solves a sub-problem, such as computing an LU factorization (getrf) or solving a linear system given an LU factorization (getrs). In SLATE, these routines are templated for different targets (CPU or GPU), with the code typically independent of the target. The user can choose among various target implementations:

Target::HostTask means multithreaded execution by a set of OpenMP tasks.
Target::HostNest means multithreaded execution by a nested "parallel for" loop.
Target::HostBatch means multithreaded execution by calling a batched BLAS routine.
Target::Devices means (multi-)GPU execution using calls to batched BLAS.
MPI Communication: Communication in SLATE relies on explicit dataflow information. When a tile is needed for computation, it is broadcast to all the processes where it is required. Rather than explicitly listing MPI ranks, the broadcast is expressed in terms of the destination (sub)matrix to be updated. This way, SLATE's messaging layer is oblivious to the mapping of tiles to processes. Also, multiple broadcasts are aggregated to allow for pipelining of MPI messages with transfers between the host and the devices. Since the set
of processes involved in a broadcast is determined dynamically, the use of MPI collectives is not ideal, as it would require setting up a new subcommunicator for each broadcast. Instead, SLATE uses point-to-point MPI communication following a hypercube pattern to broadcast the data.

Node-Level Coherency: For offload to GPU accelerators, SLATE implements a memory consistency model, inspired by the MOSI cache coherency protocol [1, 2], on a tile-by-tile basis. For read-only access, tiles are mirrored in the memories of, possibly multiple, GPU devices and deleted when no longer needed. For write access, tiles are migrated to the GPU memory and returned to the CPU memory afterwards if needed. A tile's instance can be in one of three states: Modified, Shared, or Invalid. Additional flag OnHold can be set along any state, as follows:

Modified (M) indicates that the tile's data is modified. Other instances should be Invalid. The instance cannot be purged.
Shared ( $\mathbf{S}$ ) indicates that the tile's data is up-to-date. Other instances may be Shared or Invalid. The instance may be purged unless it is on hold.
Invalid (I) indicates that the tile's data is obsolete. Other instances may be Modified, Shared, or Invalid. The instance may be purged unless it is on hold.
OnHold (O) is a flag orthogonal to the other three states that indicates a hold is set on the tile instance, and the instance cannot be purged until the hold is released.

Dynamic Scheduling: Dataflow scheduling (omp task depend) is used to execute a task graph with nodes corresponding to large blocks of the matrix. Dependencies are tracked using dummy vectors, where each element represents a block of the matrix, rather than the matrix data itself. For multi-core execution, each large block is dispatched to multiple cores - using either nested tasking (omp task) or batched BLAS. For GPU execution, calls to batched BLAS are used specifically to deliver fast processing of matrix blocks that are represented as large collections of tiles.

One of the main benefits of SLATE's architecture is dramatic reduction in the size of the source code, compared to ScaLAPACK (Figure 1.3). As of August 2019, with more than two thirds of ScaLAPACK's functionality covered, SLATE's source code is $8 \times$ to $9 \times$ smaller than ScaLAPACK's.


Figure 1.3: Code size comparison - ScaLAPACK vs SLATE (numbers from August 2019).

## CHAPTER 2

## Implementation

### 2.1 Singular Value Decomposition

In linear algebra, the singular value decomposition (SVD) is a factorization of a real or complex matrix $A$ of the form $U \Sigma V^{H}$, where $U$ is an $m \times m$ real or complex unitary matrix, $\Sigma$ is an $m \times n$ rectangular diagonal matrix with non-negative real numbers on the diagonal, and $V$ is $n \times n$ real or complex unitary matrix. The diagonal entries $\sigma_{i}$ of $\Sigma$ are known as the singular values of $A$. The columns of $U$ and the columns of $V$ are known as the left-singular vectors and the right-singular vectors of $A$, respectively. Typically the values $\sigma_{i}$ are ordered such that $\sigma_{1} \geq \sigma_{2} \geq \cdots \geq \sigma_{\min (m, n)} \geq 0$. Typically, only the first $\min (m, n)$ columns of $U$ and rows of $V$ are computed, yielding the "reduced" or "economy-size" SVD, since the remaining columns of $U$ and rows of $V$ are multiplied by the zero part of $\Sigma$ and do not contribute to $A$.

The SVD is the generalization of the eigendecomposition of a positive semidefinite normal matrix to any $m \times n$ matrix via an extension of the polar decomposition. The SVD is related to the eigendecomposition in the following way. The singular values are the square roots of the eigenvalues of $A^{T} A$, the columns of $V$ are the corresponding eigenvectors, and the columns of $U$ are the eigenvectors of $A A^{T}$.

The discovery of the SVD is attributed to four famous mathematicians, who seem to have come across it independently: Eugenio Beltrami (in 1873), Camille Jordan (in 1874), James Joseph Sylvester (in 1889), Léon César Autonneand (in 1915). The first proof of the singular value decomposition for rectangular and complex matrices seems to be by Carl Eckart and Gale J. Young in 1936 [3].
First practical methods for computing the SVD are attributed to Kogbetliantz and Hestenes [4] and resemble closely the Jacobi eigenvalue algorithm, which uses Jacobi (Givens) plane rotations.

These were replaced by the method of Golub and Kahan [5], which uses Householder reflections to reduce to bidiagonal, then plane rotations to continue the reduction to diagonal. The most popular algorithm used today is the variant of the Golub/Kahan algorithm published by Golub and Reinsch [6].

### 2.2 Hermitian Eigenvalue Problem

In linear algebra, an eigendecomposition or spectral decomposition is the factorization of a matrix into a canonical form, where the matrix is represented in terms of its eigenvalues and eigenvectors. An eigenvector or characteristic vector of a linear transformation is a nonzero vector that changes by a scalar factor when that linear transformation is applied to it. That is, a (non-zero) vector $x$ of dimension $n$ is an eigenvector of a square $n \times n$ matrix $A$ if it satisfies the linear equation $A x=\lambda x$. In other words, the eigenvectors are the vectors that the linear transformation $A$ merely elongates or shrinks, and the amount that they elongate/shrink by is the eigenvalue.

A square $n \times n$ matrix $A$ with $n$ linearly independent eigenvectors $q_{i}$ (where $i=1, \ldots, n$ ) can be factored as $A=X \Lambda X^{-1}$ where $X$ is the square $n \times n$ matrix whose $i$ th column is the eigenvector $x_{i}$ of $A$, and $\Lambda$ is the diagonal matrix whose diagonal elements are the corresponding eigenvalues, $\Lambda_{i i}=\lambda_{i}$. Only diagonalizable matrices can be factorized in this way.

Any Hermitian matrix can be diagonalized by a unitary matrix, and the resulting diagonal matrix has only real entries. This implies that all eigenvalues of a Hermitian matrix $A$ with dimension $n$ are real, and that $A$ has $n$ linearly independent eigenvectors. Moreover, a Hermitian matrix has orthogonal eigenvectors for distinct eigenvalues. Given that conjugate transpose of a unitary matrix is also its inverse, the Hermitian eigenvalue problem boils down to $A=X \Lambda X^{H}$. This means $A=X \Lambda X^{T}$ in the case of real symmetric matrices.

Historically, eigenvalues arose in the study of quadratic forms and differential equations. The initial discoveries are attributed to Euler, Lagrange, and Cauchy. The list of mathematicians who contributed to the field includes such famous names as Fourier, Sturm, Hermite, Brioschi, Clebsch, Weierstrass, Liouville, Schwarz, and Poincaré. Generally, Hilbert is credited with using the German word eigen, which means "own", to denote eigenvalues and eigenvectors, though he may have been following a related usage by Helmholtz.

The first numerical algorithm for computing eigenvalues and eigenvectors appeared in 1929, when Von Mises published the power method. One of the most popular methods today, the QR algorithm, was proposed independently by Francis [7] and Kublanovskaya [8] in 1961.

### 2.3 Generalized Hermitian Definite Eigenvalue Problem

The generalized Hermitian definite eigenvalue problem has various types:

- Type 1: $A z=\lambda B z$,
- Type 2: $A B z=\lambda z$,
- Type 3: $B A z=\lambda z$,
where $A$ is Hermitian and $B$ is Hermitian positive-definite.
To solve it, we first reduce it to the standard eigenvalue form, $\hat{A} x=\lambda x$. The reductions for types (2) and (3) are identical; they differ in the back-transformation. First, form the Cholesky factorization of $B$ as either $B=L L^{H}$ with lower triangular $L$, or $B=U^{H} U$ with upper triangular $U$. Then form $\hat{A}$, which overwrites $A$, as:
- Type 1: compute $\hat{A}=L^{-1} A L^{-H}$ or $\hat{A}=U^{-H} A U^{-1}$, as shown in Algorithm 1;
- Type 2 or 3: compute $\hat{A}=L^{H} A L$ or $\hat{A}=U A U^{H}$, as shown in Algorithm 2.

Only the lower or upper triangles of $A, \hat{A}$, and $B$ are stored and computed on, the opposite triangle being known from symmetry. The hegst routine (Hermitian generalized to standard) takes $A$ and the Cholesky factor $L$ or $U$ of $B$ as input; the lower or upper triangle of $\hat{A}$ overwrites the lower or upper triangle of $A$ on output.
After solving the standard eigenvalue problem, $\hat{A} x=\lambda x$, an eigenvector $x$ is back-transformed to be an eigenvector $z$ of the generalized eigenvalue problem as follows:

- Type 1 or 2 : $z=L^{-H} x$ or $z=U^{-1} x$ using trsm;
- Type 3: $z=L x$ or $z=U^{H} x$ using trmm.


### 2.4 Three Stage Algorithms

We solve both the SVD and the Hermitian eigenvalue problem by a three stage algorithm, shown in Figure 2.1:
(1) First stage reduction from full to triangular band (SVD) or Hermitian band (eigenvalue) form, which uses Level 3 BLAS.
(2) Second stage reduction band to real bidiagonal (SVD) or real symmetric tridiagonal (eigenvalue) form. This uses a bulge chasing algorithm.
(3) Third stage reduction to diagonal form, revealing the singular values or eigenvalues. Currently we use QR iteration, but could also use divide and conquer, MRRR, bisection, or other solver.

This is in contrast to the traditional algorithm used in LAPACK and ScaLAPACK that goes directly from full to bidiagonal or symmetric tridiagonal, which uses Level 2 BLAS and is memory-bandwidth limited. If $m \gg n$ (or $m \ll n$ ), the SVD has an optional initial reduction from tall (or wide) to square, using a QR (or LQ) factorization.
(In the literature, this three stage algorithm is often called a two stage algorithm, meaning the reduction from dense to tri/bi-diagonal is two stages. The reduction to diagonal is then a separate phase of the algorithm.)

For the SVD, the first stage proceeds by computing a QR factorization of a block column to annihilate entries below the diagonal, and updating the trailing matrix, as shown in Figure 2.2.

```
Algorithm 1 Reduction to standard form (type 1) pseudocode.
    function hegst(type, \(A, B\) )
        for \(k=1, \ldots, n t / / n t=\) number of block rows in \(A\).
            // \(A(k, k)=B(k, k)^{-1} * A(k, k) * B(k, k)^{-H}\).
            hegst(type, \(A(k, k), B(k, k))\)
            \(/ / A(k+1: n t, k)=A(k+1: n t, k) * B(k, k)^{H}\).
            for \(m=k+1, \ldots, n t\)
                    \(\operatorname{trsm}(B(k, k), A(m, k))\)
            end
            \(/ / A(k+1: n t, k)=[B(k+1: n t, k) * A(k, k)]+A(k+1: n t, k)\).
            for \(m=k+1, \ldots, n t\)
                    \(\operatorname{hemm}(A(k, k), B(m, k), A(m, k))\)
            end
            \(/ / A(k+1: n t, k+1: n t)=[A(k+1: n t, k) * B(k+1: n t, k)]+A(k+1: n t, k+1: n t)\).
            for \(m=k+1, \ldots, n t\)
                for \(n=k+1, \ldots, n t\)
                    \(\operatorname{her} 2 \mathrm{k}(A(m, k), B(m, k), A(m, n))\)
                    end
            end
            \(/ / A(k+1: n t, k)=[B(k+1: n t, k) * A(k, k)]+A(k+1: n t, k)\).
            for \(m=k+1, \ldots, n t\)
                \(\operatorname{hemm}(A(k, k), B(m, k), A(m, k))\)
            end
            \(/ / A(k+1: n t, k)=B(k+1: n t, k+1: n t) * A(k+1: n t, k)\).
            for \(m=k+1, \ldots, n t\)
                for \(n=k+1, \ldots, n t\)
                    \(\operatorname{trsm}(B(m, n), A(m, k))\)
                end
            end
        end
        return \(A\)
    end function
```

```
Algorithm 2 Reduction to standard form (type 2 or 3) pseudocode.
    function hegst(type, \(A, B\) )
        for \(k=1, \ldots, n t / / n t=\) number of block rows in \(A\).
            \(/ / A(k, 1: k)=[A(k, 1: k) * B(1: k, 1: k)]\).
            for \(m=1, \ldots, k\)
                    for \(n=1, \ldots, k\)
                    \(\operatorname{trmm}(B(m, n), A(k, m))\)
                    end
            end
            \(/ / A(k, 1: k)=[A(k, k) * B(k, 1: k)]+A(k, 1: k)\).
            for \(m=1, \ldots, k\)
                \(\operatorname{hemm}(A(k, k), B(k, m), A(k, m))\)
            end
            \(/ / A(1: k, 1: k)=\left[A(k, 1: k)^{H} * B\left(k,(1: k)^{H}\right]+A(1: k, 1: k)\right.\).
            for \(m=1, \ldots, k\)
                for \(n=1, \ldots, k\)
                    \(\operatorname{her} 2 \mathrm{k}(A(k, m), B(k, m), A(m, n))\)
                    end
            end
            \(/ / A(k, 1: k)=[A(k, k) * B(k, 1: k)]+A(k, 1: k)\).
            for \(m=1, \ldots, k\)
                \(\operatorname{hemm}(A(k, k), B(k, m), A(k, m))\)
            end
            \(/ / A(k, 1: k)=\left[B(k, k)^{H} * A(k, 1: k)\right]\).
            for \(m=1, \ldots, k\)
                \(\operatorname{trmm}(B(k, k), A(k, m))\)
            end
            \(/ / A(k, k)=B(k, k)^{H} * A(k, k) * B(k, k)\).
            hegst(type, \(A(k, k), B(k, k))\)
        end
        return \(A\)
    end function
```



| 1. Hermitian | 2. band to | 3. tridiagonal |
| :---: | :---: | :---: |
| to band | tridiagonal | eigenvalue solver |
| (he2hb) | (hb2st) | (sterf, etc.) |


0. tall to
square
(geqrf) 1. general
to band
(ge2tb)
2. band to bidiagonal (tb2bd)
3. bidiagonal SVD solver (bdsqr, etc.)

Figure 2.1: Three stage Hermitian eigenvalue and SVD algorithms. Three stage Hermitian eigenvalue (top) and SVD (bottom) algorithms.


Figure 2.2: One panel of the first stage reduction to band form.

It then computes an LQ factorization of a block row to annihilate entries right of the upper bandwidth, and updates the trailing matrix. It repeats factoring block columns and block rows, until the entire matrix is brought to band form. The width of the block columns and rows is the resulting matrix bandwidth, $n_{b}$.


Figure 2.3: 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.


Figure 2.4: Hermitian bulge-chasing algorithm. Only the lower triangle is accessed; the upper triangle is known implicitly by symmetry.

The second stage reduces the band form to the final bidiagonal form using a bulge chasing technique. It involves $6 n_{b} n^{2}$ operations, so it takes a small percentage of the total operations, which decreases with $n$. The operations are memory bound, but are fused together as Level 2.5 BLAS [9] for cache efficiency. We designed the algorithm to use fine-grained, memory-aware tasks in an out-of-order, data-flow task-scheduling technique that enhances data locality [10, 11].

The second stage proceeds in a series of sweeps, each sweep bringing one row to bidiagonal and chasing the created fill-in elements down to the bottom right of the matrix using successive orthogonal transformations. It uses three kernels. Kernel 1 (yellow task $T_{1,1}$ in Section 2.4) applies a Householder reflector from the right (indicated by the down arrow) to eliminate a row right of the superdiagonal, which also creates a bulge of fill-in beneath the diagonal. It then applies a Householder reflector from the left (indicated by the right arrow) to eliminate the first column of the bulge below the diagonal, and applies the update to the first block column only. The remainder of the bulge is not eliminated, but is instead left for subsequent sweeps to eliminate, as they would reintroduce the same nonzeros.

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

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

For parallelism, once a sweep has finished the first kernel 3, a new sweep can start in parallel. This new sweep is shifted over one column and down one row, as shown in Section 2.4. Before task $i$ in sweep $s$, denoted as $T_{s, i}$, can start, it depends on task $T_{s-1, i+3}$ in the previous sweep being finished, to ensure that kernels do not update the same entries simultaneously. To maximize cache reuse, tasks are assigned to cores based on their data location. Ideally, the band matrix fits into the cores' combined caches, and each sweep cycles through the cores as it progresses down the band.

For the Hermitian eigenvalue problem, the second stage shown in Figure 2.4 is very similar to the SVD second stage. Where the SVD has different reflectors from the right and left, here the same reflector is applied from the left and the right. Symmetry is taken into account, so only entries in the lower triangle are computed, while entries in the upper triangle are known by symmetry.

### 2.5 Hermitian to Hermitian band reduction (he2hb)

### 2.5.1 Single node

Colon notation $i: k=i, \ldots, k$ includes both $i$ and $k$ as in Matlab (unlike Python).

```
Algorithm 3 Hermitian to Hermitian band (he2hb).
    for \(k=0: n t-1\)
        // Denote \(k\)-th trailing sub-matrix \(A_{k}=A_{k+1: n t-1, k+1: n t-1}\)
        // Panel, Householder vectors \(V\), and workspace \(W\) are block columns,
        \(/ / k t \times 1\) tiles for \(k t=n t-k-1\).
        Panel factorization: \(Q R=A_{k+1: n t-1, k}\) where \(Q=I-V T V^{H}\)
        \(W=A_{k} V \quad\) hemm, \(k t \times k t \cdot k t \times 1\)
        \(W=W T=\left(A_{k} V\right) T \quad\) trmm, right, \(k t \times 1 \cdot 1 \times 1\)
        // Symmetric update
        // \(X\) is 1-by-1 Hermitian tile ("TVAVT" in code)
        \(X=V^{H} W=V^{H}\left(A_{k} V T\right)\)
        \(X=T^{H} X=T^{H}\left(V^{H} A_{k} V T\right) \quad\) trmm, left, \(1 \times 1 \cdot 1 \times 1\)
        \(Y=W-\frac{1}{2} V X=A_{k} V T-V\left(T^{H} V^{H} A_{k} V T\right) \quad\) gemm, \(k t \times 1 \cdot 1 \times 1\)
        \(A_{k}=A_{k}-V Y^{H}-Y V^{H} \quad\) her2k, \(k t \times 1 \cdot 1 \times k t\)
    end
```

Derivation of Hermitian 2-sided update.

$$
\begin{aligned}
A & =Q^{H} A Q \\
& =\left(I-V T^{H} V^{H}\right) A\left(I-V T V^{H}\right) \\
& =\left(I-V T^{H} V^{H}\right)\left(A-A V T V^{H}\right) \\
& =\left(I-V T^{H} V^{H}\right)\left(A-W V^{H}\right) \\
& =A-W V^{H}-V T^{H} V^{H} A^{*}+V T^{H} V^{H} W V^{H} \\
& =A-W V^{H}-V W^{H}+V T^{H} V^{H} W V^{H} \\
& =A-\left(V W^{H}-\frac{1}{2} V T^{H} V^{H} W V^{H}\right)-\left(W V^{H}-\frac{1}{2} V T^{H} V^{H} W V^{H}\right) \\
& =A-V\left(W^{H}-\frac{1}{2} T^{H} V^{H} W V^{H}\right)-\left(W-\frac{1}{2} V T^{H} V^{H} W\right) V^{H} \\
& =A-V\left(W-\frac{1}{2} V\left(T^{H} V^{H} W\right)^{*}\right)^{H}-\left(W-\frac{1}{2} V T^{H} V^{H} W\right) V^{H} \\
& =A-V Y^{H}-Y V^{H}
\end{aligned}
$$

where

$$
\begin{aligned}
W & =A V T \\
Y & =W-\frac{1}{2} V\left(T^{H} V^{H} W\right)=W-\frac{1}{2} V\left(T^{H} V^{H} A V T\right) .
\end{aligned}
$$

$A^{*}$ is where we use that $A=A^{H}$. Note that $X=T^{H} V^{H} W=T^{H} V^{H} A V T$ is also Hermitian.

### 2.5.2 Multi-node

For a distributed, multi-node alogrithm, we use a 2D block-cyclic distribution, shown in Figure 2.5. Each tile is labeled with its node (MPI rank). For later illustration purposes, we start the submatrix at the 3 rd block row and column, so the first tile is on node 8 instead of node 0 . Tiles in the lower triangle are stored. Tiles in the upper triangle are not stored and are known by symmetry; the rotated labels indicate the node where the symmetric tile is stored. Nodes on the diagonal of the MPI process grid (here, $\{0,4,8\}$ ) we call diagonal nodes.

(a) Sub-matrix starting at 3rd block row and column.
(b) MPI nodes (ranks) in $3 \times 3$ process grid.

Figure 2.5: 2D block-cyclic distribution.
Summarizing Algorithm 3, there are 4 basic steps in each iteration:

$$
\begin{align*}
Q R & =A_{\text {panel }} & & \text { panel factorization with } Q=I-V T V^{H}  \tag{2.1}\\
W & =A_{k} V T & & \text { hemm, trmm }  \tag{2.2}\\
Y & =W-\frac{1}{2} V\left(T^{H} V^{H} W\right) & & \text { gemm, trmm, gemm }  \tag{2.3}\\
A_{k} & -=V Y^{H}+Y^{H} V & & \text { her2k } \tag{2.4}
\end{align*}
$$

We use CAQR for a multi-node algorithm. In each panel, there is a local QR factorization within each node, giving a local block Householder reflector, then a tree reduction of resulting $R_{i}$ triangles, giving small coupling Householder reflectors. For a $p \times p$ grid, each panel has $p$ local QR factorizations, except the last $p$ panels have fewer local panels. For $p=3$, this results in block Householder vectors $V_{a}, V_{b}$, and $V_{c}$. For column 0 in Figure 2.5, $V_{a}$ is computed on node 6 (dark green), $V_{b}$ on node 7 (medium green), and $V_{c}$ on node 8 (light green).

## Computing $W=A V$ (hemm)

After the panel factorization, we start the trailing matrix update by applying $Q_{a}$ on the left and right to update $A_{k}$. We compute $W_{a}=A V_{a} T_{a}$, and $Y_{a}$ by (2.3), then update $A_{k}-=V_{a} Y_{a}^{H}+Y_{a} V_{a}^{H}$. Figure 2.6 depicts $W_{a}=A V_{a}$, which involves nodes $\{0,1,2,3,6\}$, with diagonal node 0 . Tiles that have 2 colors (blue/red or green/red) are computed as partial sums on 2 nodes, then sum-reduced on both nodes, although for some edge tiles only 1 node contributes. Diagonal node 0 (dark red) has only local contributions. Blank tiles in $V_{a}$ are zero, which eliminates contributions from the whited-out regions of $A_{k}$. This yields

$$
\begin{equation*}
W_{i}=\sum_{j \in \text { panel_rank_rows }} A_{i j} V_{j}, \quad \text { for } i=k+1, \ldots, n t-1, \tag{2.5}
\end{equation*}
$$

where panel_rank_rows $=\{1,4,7\}$ is a list of non-zero block-rows in $V_{a}$, which are local block-rows on node 6 . For a tile $A_{i j}$ in the upper triangle of $A_{k}$, we conjugate-transpose the corresponding tile in the lower triangle, $A_{j i}^{H}$, so (2.5) becomes

$$
\begin{equation*}
W_{i}=\sum_{\substack{j \in \text { panel_rank_rows, } \\ i \geq j\left(A_{i j} \text { in lower }\right)}} A_{i j} V_{j}+\sum_{\substack{j \in \text { panel_rank_rows, } \\ i<j\left(A_{j i} \text { in lower }\right)}} A_{j i}^{H} V_{j} . \tag{2.6}
\end{equation*}
$$

The first of these sums is done on the node where the local rows matches panel_rank_rows; the second sum is done on the node where the local columns matches panel_rank_rows. For diagonal nodes, the local rows and local columns are the same, so the entire sum is local. For diagonal tiles, $A_{j j} V_{j}$ is a hemm instead of a gemm. Each node computes and uses $W_{i}$ where $i$ is in its local rows or its local columns.


Figure 2.6: Hermitian matrix multiply $W_{a}=A V_{a}$ (hemm). Tiles that have 2 colors (blue/red or green/red) are computed as partial sums on 2 nodes, then reduced.

## Computing $A-=V W^{H}+W V^{H}$ (her2k)

Figure 2.7 shows the two products $V_{a} W_{a}^{H}$ and $W_{a} V_{a}^{H}$. Due to zero blocks in $V_{a}$, the whited-out rows and columns of $A_{k}$ are not updated, and we need to compute blocks only in the lower triangle, as blocks in the upper triangle are the same by symmetry. Diagonal node 0 (dark red) has non-zero blocks in both products and can be updated using a standard Hermitian rank $2 k$ (her2k) operation. Nodes 3 (blue) and 6 (green) have non-zero blocks only in $V_{a} W_{a}^{H}$, which comes from applying $Q_{a}$ on the left, while nodes 1 (red) and 2 (light red) have non-zero blocks only in $W_{a} V_{a}^{H}$, which comes from applying $Q_{a}$ on the right. Thus we have a custom routine, her2k_offdiag_ranks, that applies either a left or a right update as needed. (We keep the her 2 k name, as the overall operation is a rank $2 k$ update, even though for individual tiles it is only rank $k$.)


Figure 2.7: Hermitian rank $2 k$ (her2k) update, $A_{k}-=V_{a} W_{a}^{H}+W_{a} V_{a}^{H}$.

## Applying $Q_{b}$

After applying $Q_{a}$, we apply $Q_{b}=I-V_{b} T_{b} V_{b}^{H}$, as shown in Figures 2.8 and 2.9. This largely follows the same structure as applying $Q_{a}$, but involves nodes $\{1,3,4,5,7\}$, with diagonal node 4 .

When applying $Q_{a}$ on the right previously, node 1 used $W_{2}, W_{5}, W_{8}$, corresponding to the local rows on node 1 . Here, applying $Q_{b}$ on the left, node 1 uses $W_{1}, W_{4}, W_{7}$, corresponding to the local columns on node 1. This pattern holds for all the nodes: a left update uses $W_{i}$ for $i$ in the node's local rows; a right update uses $W_{i}$ for $i$ in the node's local columns. In a symmetric update the local rows and local columns are the same.


Figure 2.8: Hermitian matrix multiply $W_{b}=A V_{b}$ (hemm).


Figure 2.9: Hermitian rank $2 k$ (her2k) update, $A_{k}-=V_{b} W_{b}^{H}+W_{b} V_{b}^{H}$.

## Applying $Q_{c}$

After applying $Q_{a}$ and $Q_{b}$, we apply $Q_{c}=I-V_{c} T_{c} V_{c}^{H}$, as shown in Figures 2.10 and 2.11. This largely follows the same structure as applying $Q_{a}$ and $Q_{b}$, but involves nodes $\{2,5,6,7,8\}$, with diagonal node 4.


Figure 2.10: Hermitian matrix multiply $W_{c}=A V_{c}$ (hemm).


Figure 2.11: Hermitian rank $2 k$ (her2k) update, $A_{k}-=V_{c} W_{c}^{H}+W_{c} V_{c}^{H}$.

## Participating nodes

Examining the MPI process grid, we see that each $V$ updates one process row and one process column. The diagonal node at the intersection of that row and column does a 2 -sided update (her2k), while nodes in the process row do a left update, and nodes in the process column do a right update, as shown in Figure 2.12. Since nodes 5 and 7 don't participate in $Q_{a}$, they can immediately proceed with applying $Q_{b}$. Similarly, all diagonal nodes, $\{0,4,8\}$, can be updated at the start.

(a) Applying $Q_{a}$ updates MPI row and column 1.

(b) Applying $Q_{b}$ updates MPI row and column 2.

(c) Applying $Q_{c}$ updates MPI row and column 3.

Figure 2.12: Nodes and tiles that are updated by each $Q_{r}$ from panel.

## Multi-node Algorithm

```
Algorithm 4 Hermitian to Hermitian band (he2hb).
    for \(k=0: n t-1\)
        // Denote \(k\)-th trailing sub-matrix as \(A_{k}=A_{k+1: n t-1, k+1: n t-1}\).
        // Panel, Householder vectors \(V\), and workspace \(W\) are block columns,
        // \(k t \times 1\) tiles for \(k t=n t-k-1\).
        local panel factorization: \(Q R=A_{k+1: n t-1, k}\) where \(Q=I-V T V^{H}\)
        triangle panel reduction
        for \(r=1: p\) with \(p\) nodes in panel
            panel_rank \(=\) panel_ranks[ r ]
            if I am in process row \(r\) or column \(r\) then
                \(W=A_{k} V_{r} \quad\) hemm, \(k t \times k t \cdot k t \times 1\)
                \(W=W T_{r}=\left(A_{k} V_{r}\right) T_{r} \quad\) trmm, right, \(k t \times 1 \cdot 1 \times 1\)
                    if I am diagonal node then
                    // Symmetric update
                    // \(X\) is 1-by-1 Hermitian tile ("TVAVT" in code)
                \(X=V_{r}^{H} W \quad\) inner-product gemm, \(1 \times k t \cdot k t \times 1\)
                    \(X=T_{r}^{H} X=T_{r}^{H} V_{r}^{H} A_{k} V_{r} T_{r} \quad\) trmm, left, \(1 \times 1 \cdot 1 \times 1\)
                    \(Y=W-\frac{1}{2} V_{r} X \quad\) gemm, \(k t \times 1 \cdot 1 \times 1\)
                    \(A_{k}=A_{k}-V_{r} Y^{H}-Y V_{r}^{H} \quad\) her2k, \(k t \times 1 \cdot 1 \times k t\)
            else
                // Offdiagonal update
                \(A_{k}-=V_{r} W^{H}\) or \(A_{k}-=W V_{r}^{H} \quad\) her2k_offdiag_ranks, \(k t \times 1 \cdot 1 \times k t\)
            end
            end // if in process row/col
        end // for \(r\)
        apply update from triangle panel reduction (hettmqr)
    end // for k
```


### 2.6 Eigenvector Computation

The three stage Hermitian approach to solve the eigenvalue problem of a dense matrix is to first reduce it to Hermitian band matrix form, $A=Q_{1} B Q_{1}^{H}$ using Householder reflectors, then reduce the banded matrix further into a real symmetric tridiagonal matrix $B=Q_{2} T Q_{2}^{H}$, finally, compute the eigenpairs of the tridiagonal matrix using an iterative method such as QR iteration, or the recursive approach of divide-and-conquer, such that $T=Q_{3} \Lambda Q_{3}^{H}$. The subsequent eigenvectors are then accumulated during the back transformation phase, i.e., $X=Q_{1} Q_{2} Q_{3}$ to calculate the eigenvectors $X$ of the original matrix $A$.

### 2.6.1 Eigenvectors of tridiagonal matrix

Once the tridiagonal reduction is achieved, the implicit QR eigensolver steqr2 calculates the eigenvalues and optionally its associated eigenvectors of the condensed matrix structure. In


Figure 2.13: Redistribute 1D block row cyclic distributed matrix using $4 \times 1$ grid into a 2D block cyclic distribution using $2 \times 2$ grid.

SLATE (and ScaLAPACK), the steqr2 is a modified version of the LAPACK routine steqr which allows each process to perform updates on the distributed matrix $Q_{2}$, and achieve parallelization during this step.

Algorithm 5 shows the call to the tridiagonal eigensolver steqr2. First, a matrix to store the eigenvectors $Q_{3,1 D}$ of the tridaigonal matrix $T$ is created using a 1D block row cyclic with a $n_{p} \times 1$ process grid, where $n_{p}$ is the number of MPI processes. Then each process updates up to $\left(n / n_{b}\right) / n_{p}$ rows of the matrix $Q_{3,1 D}$, where $n$ is the matrix size and $n_{b}$ is the block size used to distribute the rows of $Q_{3,1 D}$. Finally, the matrix of the eigenvectors is redistributed to a 2D block cyclic distribution as illustrated in Figure 2.13.

```
Algorithm 5 Tridiagonal Eigensolver using steqr2 pseudocode.
    function steqr2( \(T, Q_{3}\) )
        // The "1D block row cyclic" grid configuration
        \(1 D=n_{p} \times 1\)
        // Compute the number of rows owned by each processor
        \(n r c=\left(n / n_{b}\right) / n_{p}\)
        // Build SLATE matrix \(Q_{3,1 D}\) using the 1-dim grid
        \(Q_{3,1 D}=\operatorname{Matrix}\left(n r c, n_{b}, n_{p}, 1\right)\)
        // Call steqr2 to compute the eigenpairs of the tridiagonal matrix
        \(\left(Q_{3,1 D} \Lambda Q_{3,1 D}^{H}\right)=\operatorname{steqr2}(T)\)
        // The "2D block cyclic" grid configuration
        \(2 D=p \times q\)
        // Redistribute the 1-dim eigenvector matrix into 2-dim matrix
        \(Q_{3}=\operatorname{redistribute}\left(Q_{3,1 D}\right)\)
    end function
```


### 2.6.2 Second stage back-transformation

The second stage back-transformation multiplies the vectors $Q_{3}$ by $Q_{2}$ from the second stage reduction from band to tridiagonal form ("bulge chasing"), to form $Q_{2} Q_{3}$. SLATE uses a distributed version of the scheme developed by [12]. The Householder vectors generated during the bulge chasing (Figure 2.4) are stored in a matrix $V$, shown in Figure 2.14. Conceptually, the


Figure 2.14: Second stage back transformation, with $V$ block size $j_{b}=3$ vectors. Block reflector 3 is highlighted to show overlap.
vectors from each sweep $i$ are stored in column $i$ of the lower triangular matrix $V$. The vectors are blocked together into parallelograms, as shown in Section 2.6.2, to form block Householder reflectors, $H_{r}=I-V_{r} T_{r} V_{r}^{H}$ where $V_{r}$ is the $r$ th block of $V$, using the compact WY format [13]. Thus $Q_{2}=H_{k} \cdots H_{2} H_{1}$. Application of these $H_{r}$ overlap, illustrated in Section 2.6.2, creating the dependencies between them shown in Section 2.6.2. These dependencies allow up to $\left\lceil\frac{\mathrm{mt}}{2}\right\rceil$ updates to occur in parallel. Figure 2.15 shows these blocks and the corresponding tasks for a $10 \times 10$ block matrix. For instance, all four dark blue tasks update different rows of $Q_{3}$ and so can run in parallel. Using the OpenMP task scheduler makes taking advantage of this parallelism very easy. The routine unmtr_hb2st, outlined in Algorithm 6, applies $Q$ to a matrix $C$; for eigenvectors, $C=Q_{3}$. Application of each $H_{r}$ becomes a single task, with dependencies on the two rows it updates, row $[i]$ and row $[i+1]$. The parallelism in Section 2.6.2 occurs automatically based on these dependencies. Within a row of $C$, updating each tile is independent, so we can use nested parallelism in the parallel for loop.

In SLATE, each parallelogram block $V_{r}$ is $2 n_{b} \times n_{b}$. To ease computation, instead of storing blocks in a lower triangular matrix (Section 2.6.2), each block is stored as one $2 n_{b} \times n_{b}$ tile, with explicit zeros in the upper and lower triangular areas, as shown in Section 2.6.2. This allows us, for instance, to use LAPACK's larft function to compute $T_{r}$ from $V_{r}$, and to use gemm instead of trmm. Normally, $V_{r}$ has unit diagonal. SLATE stores the Householder $\tau$ values on the diagonal of $V_{r}$. During computation, the diagonal is set to 1's, and the $\tau$ values are restored afterwards.

### 2.6.3 First stage back-transformation

The first stage back-transformation multiplies the vectors $\left(Q_{2} Q_{3}\right)$ by $Q_{1}$ from the first stage reduction to band, to form $X=Q_{1}\left(Q_{2} Q_{3}\right)$. The routine unmtr_he2hb applies $Q_{1}$ or $Q_{1}^{H}$ on the left or right of a matrix $C$, which is then overwritten by $Q_{1} C, Q_{1}^{H} C, C Q_{1}$, or $C Q_{1}^{H}$. For eigenvectors, we need only the left, no-transpose case with $C=Q_{2} Q_{3}$, to form the eigenvectors $X=Q_{1}\left(Q_{2} Q_{3}\right)$. It is essentially identical to applying $Q$ from a QR factorization, but shifted by


Figure 2.15: Dependencies allow up to $\left\lceil\frac{m t}{2}\right\rceil$ parallel tasks.

```
Algorithm 6 unmtr_hb2st back-transformation pseudocode. Indices are block rows/cols.
    function unmtr_hb2st( \(V, C\) )
        // \(C\) is \(m t \times n t\) block rows/cols, blocksize \(n_{b} \times n_{b}\)
        \(/ / V\) is \(m t(m t+1) / 2\) blocks, blocksize \(2 n_{b} \times n_{b}\)
        for \(j=m t-1\) to 0
            for \(i=j\) to \(m t-1\)
                task depend in, out on row \([i]\) and \(\operatorname{row}[i+1]\)
                    \(r=i-j+j \cdot m t-j(j-1) / 2\)
                    Broadcast \(V_{r}\)
                    Compute \(T\) from \(V_{r}\) (larft)
                    \(D=V_{r} T\) (gemm or trmm)
                    parallel for \(k=0\) to \(n t-1\)
                    if \(C_{i: i+1, k}\) are local then
                                    \(/ /\) Compute \(Q C=\left(I-V T V^{H}\right) C\)
                                    \(W=V_{r}^{H} C_{i: i+1, k}\)
                                    \(C_{i: i+1, k}=C_{i: i+1, k}-D W\)
                    end
                    end
                    end task
            end
        end
    end function
```

one block-row since we reduced to band form instead of triangular form, as in QR. Thus, as in LAPACK, we can leverage the existing unmqr routine that applies $Q$ from a QR factorization.

## CHAPTER 3

## Divide and conquer

The solution of the tridiagonal eigenvalue decomposition (EVD) can use several different methods: QR iteration, divide and conquer, bisection, or MRRR. Here we derive the divide and conquer algorithm proposed by Cuppen [14], which is significantly faster than QR iteration. Our derivation largely follows that of Tisseur and Dongarra [15], with clarifications and notes related to the implementation in SLATE.

### 3.1 Cuppen's method

Define tridiagonal matrix $\boldsymbol{T} \in \mathbb{R}^{n \times n}$ with eigenvalue decomposition $\boldsymbol{T}=\boldsymbol{W} \boldsymbol{\Lambda} \boldsymbol{W}^{T}$ where $\boldsymbol{\Lambda}$ is diagonal and $\boldsymbol{W}$ is orthogonal. Split $\boldsymbol{T}$ into the rank-1 update

$$
\boldsymbol{T}=\left[\begin{array}{cc}
\dot{\boldsymbol{T}}_{1} & \Theta \rho \boldsymbol{e}_{k} \boldsymbol{e}_{1}^{T}  \tag{3.1}\\
\Theta \rho \boldsymbol{e}_{1} \boldsymbol{e}_{k}^{T} & \stackrel{\boldsymbol{T}}{2}^{2}
\end{array}\right]=\left[\begin{array}{cc}
\boldsymbol{T}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{T}_{2}
\end{array}\right]+\rho \boldsymbol{v} \boldsymbol{v}^{T} \quad \text { with } \quad \boldsymbol{v}=\left[\begin{array}{c}
\boldsymbol{e}_{k} \\
\Theta \boldsymbol{e}_{1}
\end{array}\right] .
$$

where $\boldsymbol{T}_{1}$ is $n_{1} \times n_{1}, \boldsymbol{T}_{2}$ is $n_{2} \times n_{2}$, and $\Theta \rho=T_{n_{1}+1, n_{1}}$ is the off-diagonal element, with $\Theta= \pm 1$ such that $\rho>0$ (see Section 3.1.1). $\stackrel{\circ}{\boldsymbol{T}}_{1}$ and $\boldsymbol{T}_{1}$ differ in only the bottom-right element, and $\stackrel{\circ}{\boldsymbol{T}}_{2}$ and $\boldsymbol{T}_{2}$ differ in only the top-left element:

$$
\begin{aligned}
\left(\boldsymbol{T}_{1}\right)_{n_{1}, n_{1}} & =\left(\stackrel{\circ}{\boldsymbol{T}}_{1}\right)_{n_{1}, n_{1}}-\rho, \\
\left(\boldsymbol{T}_{2}\right)_{1,1} & =\left(\stackrel{\circ}{\boldsymbol{T}}_{2}\right)_{1,1}-\rho .
\end{aligned}
$$

Recursively solve eigen decompositions for $\boldsymbol{T}_{1}$ and $\boldsymbol{T}_{2}$, yielding

$$
\begin{align*}
& \boldsymbol{T}_{1}=\boldsymbol{Q}_{1} \boldsymbol{D}_{1} \boldsymbol{Q}_{1}^{T},  \tag{3.2}\\
& \boldsymbol{T}_{2}=\boldsymbol{Q}_{2} \boldsymbol{D}_{2} \boldsymbol{Q}_{2}^{T} . \tag{3.3}
\end{align*}
$$

Then the eigen decomposition of $\boldsymbol{T}$ is given by

$$
\begin{aligned}
\boldsymbol{T} & =\boldsymbol{Q} \boldsymbol{D} \boldsymbol{Q}^{T}+\rho \boldsymbol{v} \boldsymbol{v}^{T} \\
& =\boldsymbol{Q}\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right) \boldsymbol{Q}^{T}
\end{aligned}
$$

with

$$
\boldsymbol{Q}=\left[\begin{array}{cc}
\boldsymbol{Q}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{Q}_{2}
\end{array}\right], \quad \boldsymbol{D}=\left[\begin{array}{cc}
\boldsymbol{D}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{D}_{2}
\end{array}\right], \quad \boldsymbol{z}=\boldsymbol{Q}^{T} \boldsymbol{v} .
$$

From the definition of $\boldsymbol{v}, \boldsymbol{z}$ is the last row of $\boldsymbol{Q}_{1}$, and $\Theta$ times the first row of $\boldsymbol{Q}_{2}$. Given Equation (3.2), compute the eigen decomposition

$$
\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}=\boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^{T}
$$

as shown in Section 3.1.2. Finally, the eigen decomposition of $\boldsymbol{T}$ is

$$
\boldsymbol{T}=\boldsymbol{Q}\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right) \boldsymbol{Q}^{T}=\boldsymbol{Q}\left(\boldsymbol{U} \boldsymbol{\Lambda} \boldsymbol{U}^{T}\right) \boldsymbol{Q}^{T}=\boldsymbol{W} \boldsymbol{\Lambda} \boldsymbol{W}^{T}, \quad \boldsymbol{W}=\boldsymbol{Q} \boldsymbol{U}
$$

### 3.1.1 Historical note on choice of $\theta$

In Tisseur and Dongarra [15], for $\rho=\theta T_{n_{1}+1, n_{1}}$ (what they refer to as $\theta \beta$ ), scalar $\theta$ is chosen to avoid cancellation, per Dongarra and Sorensen [16]. This is shown in Algorithm 7, though some details are lacking. See http://www.netlib.org/misc/SEV/ for code. As used here, $\Theta=\theta^{-1}$, but we restrict $\Theta= \pm 1$, so $\Theta=\theta^{-1}=\theta$.

```
Algorithm 7 Choice of \(\theta\) in Dongarra and Sorensen [16]
    if \(\operatorname{sign}\left(\hat{\boldsymbol{T}}_{1}\left(n_{1}, n_{1}\right)\right)==\operatorname{sign}\left(\hat{\boldsymbol{T}}_{2}(1,1)\right)\) then
        \(/ /-\theta \beta\) has same sign as diagonal elements \(\hat{\boldsymbol{T}}_{1}\left(n_{1}, n_{1}\right)\) and \(\hat{\boldsymbol{T}}_{2}(1,1)\).
        \(\theta=-\operatorname{sign}\left(\hat{\boldsymbol{T}}_{1}\left(n_{1}, n_{1}\right)\right) \cdot \operatorname{sign}(\beta)\)
    else
        // Unclear: \(-\theta \beta\) has same sign as one of diagonal elements,
        // and magnitude of \(\theta\) chosen to avoid cancellation when
        // \(\theta^{-1} \beta\) is subtracted from the other diagonal element.
    end
```

Inconsistent with this, (Sca)LAPACK implicitly chooses $\theta=\operatorname{sign}\left(T_{n_{1}+1, n_{1}}\right)$, since it does:

```
!! pdlaed0.f, line 148-149
!! T1_{n1, n1} -= rho
!! T2_{1, 1 } -= rho
D( I-1 ) = D( I-1 ) - ABS( E ( I-1 ) )
D(I ) = D(I ) - ABS (E(I-1 ) )
!! pdlade0.f, line 214-215
!! Pass E( ID+N1-1 ) as rho
CALL PDLAED1( MATSIZ, N1, D( ID ), ID, Q, IQ, JQ, DESCQ, &
                                    E( ID+N1-1 ), WORK, IWORK( SUBPBS+1 ), IINFO )
!! pdlaed2.f, line 206
!! rho = | 2 rho | = | 2 T_{n1+1, n1} |
!! = 2 theta T}{\mp@code{n1+1, n1}, theta = sign( T_{n1+1, n1} ).
!! Factor of 2 comes from normalizing [ z1, z2 ].
RHO = ABS(TWO*RHO )
```

SLATE follows (Sca)LAPACK's convention, since LAPACK's secular equation solver (laed4) requires $\rho>0$.

### 3.1.2 Secular equation

The characteristic equation is $\operatorname{det}(\boldsymbol{A}-\lambda \boldsymbol{I})=0$. In this case,

$$
\operatorname{det}\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}-\lambda \boldsymbol{I}\right)=0
$$

Applying Sylvester's determinant theorem, with non-singular $\boldsymbol{X}=\boldsymbol{D}-\lambda \boldsymbol{I}$, we obtain

$$
\begin{aligned}
& \operatorname{det}\left(\boldsymbol{X}+\rho \boldsymbol{z}^{T}\right)=0 \\
& \operatorname{det}\left(\boldsymbol{X} \boldsymbol{X}^{-1}\left(\boldsymbol{X}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right)\right)=0 \\
& \operatorname{det}(\boldsymbol{X}) \operatorname{det}\left(\boldsymbol{I}+\rho \boldsymbol{X}^{-1} \boldsymbol{z} \boldsymbol{z}^{T}\right)=0 \\
& \operatorname{det}\left(\boldsymbol{I}+\rho\left(\boldsymbol{X}^{-1} \boldsymbol{z}\right) \boldsymbol{z}^{T}\right)=0 \quad \text { assuming } \boldsymbol{X} \text { is non-singular } \\
& 1+\rho \boldsymbol{z}^{T}\left(\boldsymbol{X}^{-1} \boldsymbol{z}\right)=0 \quad \text { Sylvester's determinant theorem. }
\end{aligned}
$$

This is the secular equation, so called because of centuries-long secular (non-periodic) perturbations of planetary orbits (see wikipedia).
Eigenvalues of $\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}$ are roots $\left\{\lambda_{j}\right\}$ of the secular equation

$$
\begin{equation*}
f(\lambda)=1+\rho \boldsymbol{z}^{T}(\boldsymbol{D}-\lambda \boldsymbol{I})^{-1} \boldsymbol{z}=1+\rho \sum_{i=1}^{n} \frac{z_{i}^{2}}{d_{i}-\lambda}, \tag{3.4}
\end{equation*}
$$

with corresponding (unnormalized) eigenvectors

$$
\boldsymbol{u}_{j}=\left(\boldsymbol{D}-\lambda_{j} \boldsymbol{I}\right)^{-1} \boldsymbol{z}=\left[\begin{array}{c}
\frac{z_{1}}{d_{1}-\lambda_{j}}  \tag{3.5}\\
\vdots \\
\frac{z_{n}}{d_{n}-\lambda_{j}}
\end{array}\right], \quad j=1, \ldots, n
$$

Solving Equation (3.4) uses a custom non-linear solver developed by Li [17]. However, direct application of Equation (3.5) to compute eigenvectors can lead to loss of precision and orthogonality; see Section 3.1.11.

The secular equation solver assumes, without loss of generality, that $\rho>0$. To accomplish this, set $\Theta=\operatorname{sign}\left(T_{n_{1}+1, n_{1}}\right)$, which negates $\boldsymbol{z}_{2}$ and $\rho$; see Section 3.1.1.

The secular equation solver further requires that $\boldsymbol{D}$ is sorted in ascending order, accomplished by applying permutation $\boldsymbol{P}_{s}$ represented by vector isort (INDX) ${ }^{1}$ to obtain sorted $\boldsymbol{D}_{s}, \boldsymbol{z}_{s}, \boldsymbol{U}_{s}$ from unsorted $\boldsymbol{D}, \boldsymbol{z}, \boldsymbol{U}$ :

$$
\boldsymbol{P}_{s}\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right) \boldsymbol{P}_{s}^{T}=\boldsymbol{D}_{s}+\rho \boldsymbol{z}_{s} \boldsymbol{z}_{s}^{T}
$$

[^1]
### 3.1.3 Deflation

When an eigenvalue has already converged, it can be deflated from the secular equation, reducing the secular equation size by 1. [15] set tolerance

$$
\eta=\epsilon\left\|\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right\|_{2} \leq \epsilon\left(\|\boldsymbol{D}\|_{2}+|\rho|\left\|\boldsymbol{z} \boldsymbol{z}^{T}\right\|_{2}\right)=\epsilon\left(\max _{j}\left|d_{j}\right|+|\rho|\right)
$$

since $\|z\|_{2}=1$. However, (Sca)LAPACK uses

$$
\eta=8 u \max \left(\max _{j}\left|d_{j}\right|, \max _{j}\left|z_{j}\right|\right)
$$

where $u$ is unit roundoff $(\epsilon / 2)$ (i.e., $u=\operatorname{lamch}($ "e")). It's unclear where this tolerance is derived from; presumably from LAPACK code in [18].

Two types of deflation can occur. The first is if $z_{j}=0$ for some $j$, then $d_{j}$ is an eigenvalue of $\boldsymbol{T}$ with eigenvector $\boldsymbol{e}_{j}$. More generally, when $z_{j}$ is nearly zero, deflate if

$$
\left|\rho z_{j}\right| \leq \eta .
$$

The second type of deflation occurs if $\boldsymbol{D}$ has eigenvalue $d_{j}$ of multiplicity $m>1$, we can rotate to zero out all but one of the corresponding $z_{j}$. More generally, when $d_{i}$ and $d_{j}$ (js1 and js2 (PJ and NJ) ${ }^{1}$ ) are nearly equal, deflate if

$$
\frac{\left|z_{i} z_{j}\right| \cdot\left|d_{i}-d_{j}\right|}{\sqrt{z_{i}^{2}+z_{j}^{2}}} \leq \eta
$$

This is the off-diagonal term when applying $\boldsymbol{G}_{i j}$ to $2 \times 2$ of $d_{i}$ and $d_{j}$ values:

$$
\boldsymbol{G}_{i j}\left[\begin{array}{ll}
D_{i, i} & \\
& D_{j, j}
\end{array}\right] \boldsymbol{G}_{i j}^{T}=\left[\begin{array}{cc}
c^{2} d_{i}+s^{2} d_{j} & -c s d_{i}+c s d_{j} \\
-c s d_{i}+c s d_{j} & s^{2} d_{i}+c^{2} d_{j}
\end{array}\right]
$$

If these off-diagonal terms are negligible, then $\boldsymbol{D}$ is still diagonal. The Givens rotation $\boldsymbol{G}_{i j}$ that applies to rows $i$ and $j$ is defined as

$$
\boldsymbol{G}_{i j}=\left[\begin{array}{cc}
c & s \\
-s & c
\end{array}\right], \quad s=-\frac{z_{i}}{r}, \quad c=\frac{z_{j}}{r}, \quad r=\sqrt{z_{i}^{2}+z_{j}^{2}}, \quad c^{2}+s^{2}=1 \quad \text { (by construction). }
$$

(Formally, $\boldsymbol{G}_{i j}$ is embedded in an $n \times n$ identity matrix.) Let $\boldsymbol{G}_{d}$ be the product of all Givens rotations and let $\boldsymbol{P}_{d}$, represented by ideflate (INDXP) ${ }^{1}$, be the permutation for both sorting and deflating eigenvalues (i.e., $\boldsymbol{P}_{d}$ includes the effect of $\boldsymbol{P}_{s}$ ). Then

$$
\boldsymbol{P}_{d} \boldsymbol{G}_{d}\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right)\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T}=\left[\begin{array}{cc}
\boldsymbol{D}_{\mathrm{sec}}+\rho \boldsymbol{z}_{\mathrm{sec}} \boldsymbol{z}_{\mathrm{sec}}^{T} & \mathbf{0}  \tag{3.6}\\
\mathbf{0} & \boldsymbol{\Lambda}_{\mathrm{df}}
\end{array}\right]+\boldsymbol{E}, \quad\|\boldsymbol{E}\|_{2} \leq c \eta .
$$

where

$$
\boldsymbol{D}_{\mathrm{sec}}+\rho \boldsymbol{z}_{\mathrm{sec}} \boldsymbol{z}_{\mathrm{sec}}^{T}=\boldsymbol{U}_{\mathrm{sec}} \boldsymbol{\Lambda}_{\mathrm{sec}} \boldsymbol{U}_{\mathrm{sec}}^{T}
$$

is the secular equation with its EVD, $\boldsymbol{\Lambda}_{\mathrm{dff}}$ are deflated eigenvalues, and $\boldsymbol{E}$ is the error. Then the EVD of (3.6) is

$$
\boldsymbol{U}_{d} \boldsymbol{\Lambda}_{d} \boldsymbol{U}_{d}^{T}, \quad \text { where } \quad \boldsymbol{U}_{d}=\left[\begin{array}{cc}
\boldsymbol{U}_{\mathrm{sec}} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{I}
\end{array}\right], \quad \boldsymbol{\Lambda}_{d}=\left[\begin{array}{cc}
\boldsymbol{\Lambda}_{\mathrm{sec}} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{\Lambda}_{\mathrm{df}}
\end{array}\right],
$$

yielding

$$
\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}=\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T} \boldsymbol{U}_{d} \boldsymbol{\Lambda}_{d} \boldsymbol{U}_{d}^{T}\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)=\boldsymbol{U} \boldsymbol{\Lambda}_{d} \boldsymbol{U}
$$

### 3.1.4 Back-transformation

The main cost of divide-and-conquer is computing $\boldsymbol{W}=\boldsymbol{Q} \boldsymbol{U}$. With deflation,

$$
\begin{aligned}
\boldsymbol{Q} \boldsymbol{U} & =\boldsymbol{Q}\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T} \boldsymbol{U}_{d} \\
& =\left[\begin{array}{cc}
\boldsymbol{Q}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{Q}_{2}
\end{array}\right] \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}\left[\begin{array}{cc}
\boldsymbol{U}_{\mathrm{sec}} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{I}
\end{array}\right] .
\end{aligned}
$$

The $\boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}$ term can be multiplied on the left into $\boldsymbol{Q}$, or on the right into $\boldsymbol{U}_{d}$.
First, consider multiplying on the right. This leaves the $\boldsymbol{Q}_{1}$ and $\boldsymbol{Q}_{2}$ block sparsity structure in $\boldsymbol{Q}$, but $\boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T} \boldsymbol{U}_{d}$ destroys the block sparsity structure in $\boldsymbol{U}_{d}$, yielding something similar to the non-deflated case, with $n^{3}$ flops:

$$
\begin{aligned}
\boldsymbol{Q U} & =\boldsymbol{Q}\left(\boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T} \boldsymbol{U}_{d}\right) \\
& =\left[\begin{array}{cc}
\boldsymbol{Q}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{Q}_{2}
\end{array}\right]\left(\begin{array}{l}
\left.\boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}\left[\begin{array}{cc}
\boldsymbol{U}_{\mathrm{sec}} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{I}
\end{array}\right]\right) \\
\\
\end{array}=\left[\begin{array}{cc}
\boldsymbol{Q}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{Q}_{2}
\end{array}\right]\left[\begin{array}{l}
\boldsymbol{U}_{1} \\
\boldsymbol{U}_{2}
\end{array}\right]\right. \\
& =\left[\begin{array}{c}
\boldsymbol{Q}_{1} \boldsymbol{U}_{1} \\
\boldsymbol{Q}_{2} \boldsymbol{U}_{2}
\end{array}\right] .
\end{aligned}
$$

Now consider multiplying on the left, $\boldsymbol{Q} \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}$. Add another permutation $\boldsymbol{P}_{t}$, represented by itype (INDX) ${ }^{1}$, to retain some block sparsity structure in $\boldsymbol{Q}$. (In ScaLAPACK, the $\boldsymbol{P}_{t}$ permutation overwrites the $\boldsymbol{P}_{s}$ permutation in INDX.) $\boldsymbol{P}_{t}$ orders cols of $\boldsymbol{Q} \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}$ into 4 column types:

- column type 1 are non-zero in only the top $n_{1}$ rows,
- column type 2 are non-zero in all rows ("dense"),
- column type 3 are non-zero in only the bottom $n_{2}$ rows,
- column type 4 are deflated; they may be non-zero in all rows.
$\boldsymbol{P}_{t}$ operates on only the first $n-n_{d}^{\prime}$ cols of $\boldsymbol{Q} \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}$ and rows of $\boldsymbol{U}_{d}$, leaving unaffected the last
$n_{d}^{\prime}$ cols or rows corresponding to deflated eigenvalues.

$$
\begin{align*}
& \boldsymbol{Q} \boldsymbol{U}=\left(\boldsymbol{Q} \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T}\right) \boldsymbol{U}_{d}  \tag{3.7}\\
& =\left(\boldsymbol{Q} \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T} \boldsymbol{P}_{t}\right)\left(\boldsymbol{P}_{t}^{T} \boldsymbol{U}_{d}\right) \tag{3.8}
\end{align*}
$$

$$
\begin{align*}
& =\left[\begin{array}{l|l}
\left\{\boldsymbol{Q}_{t}\right\}_{1,1: 2}\left\{\boldsymbol{U}_{t}\right\}_{1: 2} & \left\{\boldsymbol{Q}_{t}\right\}_{1,4} \\
\left\{\boldsymbol{Q}_{t}\right\}_{2,2: 3}\left\{\boldsymbol{U}_{t}\right\}_{2: 3} & \left\{\boldsymbol{Q}_{t}\right\}_{2,4}
\end{array}\right]  \tag{3.10}\\
& =\boldsymbol{Q}_{t} \boldsymbol{U}_{t},  \tag{3.11}\\
& \boldsymbol{U}_{t}=\left[\begin{array}{ll}
\left\{\boldsymbol{U}_{t}\right\}_{1} & \mathbf{0} \\
\left\{\boldsymbol{U}_{t}\right\}_{2} & \mathbf{0} \\
\left\{\boldsymbol{U}_{t}\right\}_{3} & \mathbf{0} \\
\hline \mathbf{0} & \boldsymbol{I}
\end{array}\right],
\end{align*}
$$

where

- $\left\{\boldsymbol{Q}_{t}\right\}_{1,1}$ has $n_{s_{1}}$ non-deflated cols from $\boldsymbol{Q}_{1}$ of column type 1 that are non-zero in only the top $n_{1}$ rows. These can be a non-deflated col from type 2 deflation, if both cols involved were from $\boldsymbol{Q}_{1}$, so the Givens rotation doesn't destroy the block structure.
- $\left\{\boldsymbol{Q}_{t}\right\}_{2,3}$ has $n_{s_{3}}$ non-deflated cols from $\boldsymbol{Q}_{2}$ of column type 3 that are non-zero in only the bottom $n_{2}$ rows. Again, these can be a non-deflated col from type 2 deflation, if both cols involved were from $\boldsymbol{Q}_{2}$.
- $\left[\begin{array}{l}\left\{\boldsymbol{Q}_{t}\right\}_{1,2} \\ \left\{\boldsymbol{Q}_{t}\right\}_{2,2}\end{array}\right]$
has $n_{s_{3}}$ non-deflated cols from either $\boldsymbol{Q}_{1}$ or $\boldsymbol{Q}_{2}$ of column type 2 that are nonzero in all $n$ rows, which arise as the non-deflated eigenvector in type 2 deflation when one col is from $\boldsymbol{Q}_{1}$ and one from $\boldsymbol{Q}_{2}$, where the Givens rotation filled in all rows. In parallel, these include some columns from types 1 and 3 , as explained below.
- $\left[\begin{array}{l}\left\{\boldsymbol{Q}_{t}\right\}_{1,4} \\ \left\{\boldsymbol{Q}_{t}\right\}_{2,4}\end{array}\right]$ are cols from deflated eigenvalues (column type 4).

No block structure is assumed. For type 2 deflation, if one col was from $\boldsymbol{Q}_{1}$ and one from $\boldsymbol{Q}_{2}$, then the Givens rotation filled in all rows, destroying the block structure.

- In the serial algorithm, $n_{d}^{\prime}=n_{d}$ eigenvalues are deflated. In the parallel algorithm, $n_{d}$ eigenvalues are deflated, but the block has size $n_{d}^{\prime} \leq n_{d}$ due to restrictions on $\boldsymbol{P}_{t}$ permuting only locally within a node.

This yields the two gemm operations in (3.10) to compute $\boldsymbol{W}=\boldsymbol{Q} \boldsymbol{U}$.

In parallel, the $\boldsymbol{Q}_{t}$ matrix has a local block structure like this, but the global structure is different. Compared to Tisseur and Dongarra [15], matrices here are renamed from $\overline{\boldsymbol{Q}}$ to $\boldsymbol{Q}_{t}$, and renumbered using the column type as the second index:

$$
\left.\begin{array}{ccc|c}
{\left[\begin{array}{ccc}
\overline{\boldsymbol{Q}}_{1,1} & \overline{\boldsymbol{Q}}_{1,2} & \mathbf{0} \\
\mathbf{0} & \overline{\boldsymbol{Q}}_{2,1} & \overline{\boldsymbol{Q}}_{2,2}
\end{array}\right.} & \overline{\boldsymbol{Q}}_{1,3}
\end{array}\right] \rightarrow\left[\begin{array}{ccc|c}
\left\{\boldsymbol{Q}_{t}\right\}_{1,1} & \left\{\boldsymbol{Q}_{t}\right\}_{1,2} & \mathbf{0} & \left\{\boldsymbol{Q}_{t}\right\}_{1,4} \\
\mathbf{0} & \left\{\boldsymbol{Q}_{t}\right\}_{2,2} & \left\{\boldsymbol{Q}_{t}\right\}_{2,3} & \left\{\boldsymbol{Q}_{t}\right\}_{2,4}
\end{array}\right]
$$

### 3.1.5 Summary

$$
\begin{aligned}
\boldsymbol{T} & =\left[\begin{array}{cc}
\boldsymbol{T}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{T}_{2}
\end{array}\right]+\rho \boldsymbol{v} \boldsymbol{v}^{T} & & \text { with } \quad \boldsymbol{v}=\left[\begin{array}{c}
\boldsymbol{e}_{n_{1}} \\
\boldsymbol{e}_{1}
\end{array}\right], \\
& =\boldsymbol{Q} \boldsymbol{D} \boldsymbol{Q}^{T}+\rho \boldsymbol{v} \boldsymbol{v}^{T} & & \text { with } \quad \boldsymbol{Q}=\left[\begin{array}{cc}
\boldsymbol{Q}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{Q}_{2}
\end{array}\right], \quad \boldsymbol{D}=\left[\begin{array}{cc}
\boldsymbol{D}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{D}_{2}
\end{array}\right], \\
& =\boldsymbol{Q}\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right) \boldsymbol{Q}^{T} & & \text { with } \boldsymbol{z}=\boldsymbol{Q}^{T} \boldsymbol{v}=\left[\begin{array}{c}
\left\{\boldsymbol{Q}_{1}\right\}_{n_{1},:}^{T} \\
\left\{\boldsymbol{Q}_{2}\right\}_{1,:}^{T}
\end{array}\right], \\
& =\boldsymbol{Q}\left(\boldsymbol{U} \Lambda \boldsymbol{U}^{T}\right) \boldsymbol{Q}^{T} & & \\
& =\boldsymbol{W} \Lambda \boldsymbol{W}^{T} & &
\end{aligned}
$$

Applying permutation $\boldsymbol{P}_{\boldsymbol{s}}$ to sort eigenvalues in the deflation routine,

$$
\boldsymbol{D}_{s}=\boldsymbol{P}_{s} \boldsymbol{D}, \quad \boldsymbol{z}_{s}=\boldsymbol{P}_{s} \boldsymbol{z}
$$

Applying permutation $\boldsymbol{P}_{d}$ and Givens rotations $\boldsymbol{G}_{\boldsymbol{d}}$ to sort and deflate eigenvalues,

$$
\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)\left(\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}\right)\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T}=\left[\begin{array}{cc}
\boldsymbol{D}_{\mathrm{sec}}+\rho \boldsymbol{z}_{\mathrm{sec}} \boldsymbol{z}_{\mathrm{sec}}^{T} & \mathbf{0} \\
\mathbf{0} & \Lambda_{\mathrm{dff}}
\end{array}\right]=\boldsymbol{U}_{d} \Lambda_{d} \boldsymbol{U}_{d}^{T}
$$

with

$$
\boldsymbol{U}_{d}=\left[\begin{array}{cc}
\boldsymbol{U}_{\mathrm{sec}} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{I}
\end{array}\right], \quad \Lambda_{d}=\left[\begin{array}{cc}
\Lambda_{\mathrm{sec}} & \mathbf{0} \\
\mathbf{0} & \Lambda_{\mathrm{df}}
\end{array}\right] .
$$

Hence,

$$
\boldsymbol{U}=\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T} \boldsymbol{U}_{d}
$$

Applying permutation $\boldsymbol{P}_{t}$ to group column types, we obtain the eigenvectors

$$
\begin{aligned}
\boldsymbol{W} & =\boldsymbol{Q} \boldsymbol{U}=\boldsymbol{Q}\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T} \boldsymbol{U}_{d}=\boldsymbol{Q}\left(\boldsymbol{P}_{d} \boldsymbol{G}_{d}\right)^{T}\left(\boldsymbol{P}_{t} \boldsymbol{P}_{t}^{T}\right) \boldsymbol{U}_{d} \\
& =\left[\begin{array}{cccc}
\left\{\boldsymbol{Q}_{t}\right\}_{1,1} & \left\{\boldsymbol{Q}_{t}\right\}_{1,2} & \mathbf{0} & \left\{\boldsymbol{Q}_{t}\right\}_{1,4} \\
\mathbf{0} & \left\{\boldsymbol{Q}_{t}\right\}_{2,2} & \left\{\boldsymbol{Q}_{t}\right\}_{2,3} & \left\{\boldsymbol{Q}_{t}\right\}_{2,4}
\end{array}\right]\left[\begin{array}{ccc}
\left\{\boldsymbol{U}_{t}\right\}_{1} & \\
\left\{\boldsymbol{U}_{t}\right\}_{2} & \\
\left\{\boldsymbol{U}_{t}\right\}_{3} & \\
\mathbf{0} & \boldsymbol{I}
\end{array}\right]
\end{aligned}
$$

with

$$
\boldsymbol{Q}_{t}=\boldsymbol{Q} \boldsymbol{G}_{d}^{T} \boldsymbol{P}_{d}^{T} \boldsymbol{P}_{t}, \quad \boldsymbol{U}_{t}=\boldsymbol{P}_{t}^{T} \boldsymbol{U}_{d}
$$

Note for deflation, the transpose $\boldsymbol{P}_{d}^{T}$ is applied to $\boldsymbol{Q}$, but for grouping column types, the untransposed $\boldsymbol{P}_{t}$ (the inverse of $\boldsymbol{P}_{t}^{T}$ ) is applied to $\boldsymbol{Q}$. Or in Matlab notation:

$$
\begin{aligned}
& \boldsymbol{Q}_{d}=\boldsymbol{Q}(:, \text { ideflate }), \\
& \boldsymbol{Q}_{t}(:, \text { itype })=\boldsymbol{Q}_{d}, \quad \operatorname{not} \boldsymbol{Q}_{t}=\boldsymbol{Q}_{d}(:, \text { itype }) .
\end{aligned}
$$

### 3.1.6 Examples

Notes:

- Let $n_{1}=4, n_{2}=4, n=n_{1}+n_{2}=8, n_{b}=2$, npcol $=2$ (number of process columns).
- deflate shows the type of deflation for that column, blank for none.
- pcolumn shows the process column.
- coltype shows the column type:
(1) Unaffected by deflation, part of $\boldsymbol{Q}_{1}$.
(2) Non-deflated eigenvector from type 2 deflation that is non-zero in all rows.
(3) Unaffected by deflation, part of $\boldsymbol{Q}_{2}$.
(4) Deflated, either type 1 or deflated eigenvector in type 2 deflation.
- $Q(1,:)$ is first row of $\boldsymbol{Q}$,
$Q(n,:)$ is last row of $\boldsymbol{Q}$. Together they serve to show the block sparsity of $\boldsymbol{Q}$.
- The decimal in $D$ values ( 0.1 or 0.3 ) denotes the original column type (1 or 3 ).
- $\boldsymbol{z}$ and $\boldsymbol{Q}$ values are not realistic (e.g., not normalized), they are just for ease of tracking how values are permuted. $\boldsymbol{z}=\boldsymbol{D} / 10$.

Each example will show how various vectors and matrices are affected by permutations.
In the output, $\varepsilon$ represents an arbitrary small value that will invoke type 2 deflation.

Original The original $\boldsymbol{D}$ eigenvalues, $\boldsymbol{z}$ vector, and $\boldsymbol{Q}$ matrix of eigenvectors.

Sorted By applying permutation $\boldsymbol{P}_{s}$ represented by vector isort (INDX) ${ }^{1}$, $\boldsymbol{D}=\boldsymbol{P}_{s} \boldsymbol{D}_{0}$ values are sorted. To get the $\boldsymbol{P}_{s}$ matrix, using Matlab notation:

```
I = eye( n, n );
Ps = I( isort, : );
```

We will use these notations interchangeably:

$$
\begin{aligned}
& D_{0}(\text { isort })=P_{s} D_{0}, \\
& Q(:, \text { isort })=Q P_{s}^{T} .
\end{aligned}
$$

Deflated With permutation $\boldsymbol{P}_{\boldsymbol{d}}$ given by ideflate (INDXP) ${ }^{1}$, deflated eigenvectors (col type 4, light orange) are moved to the end and sorted descending. (todo: any reason for sorting deflated eigvals?) Non-deflated eigenvalues are sorted ascending, as required by the secular equation solver. $\boldsymbol{P}_{d}$ includes both deflation and sorting. Columns of $\boldsymbol{Q}$ corresponding to type 2 deflation have lost their block sparsity, becoming dense $(.1 \rightarrow .12$ or .14 and $.3 \rightarrow .32$ or .34 to indicate modified values, where the second digit denotes its new column type).

Locally permuted With permutation $\boldsymbol{P}_{t}$ given by itype (INDX) ${ }^{1}$, within pcolumn 0 (dark blue), col types are sorted, e.g.: 1, 2, 2, 4 ; within pcolumn 1 (light blue), col types are sorted, e.g.: 3, 3, 4, 4 .
$\left[\left\{\boldsymbol{Q}_{t}\right\}_{1,1} \quad\left\{\boldsymbol{Q}_{t}\right\}_{1,2}\right]$ spans columns 1-5 (dark purple), to cover all col type 1 and 2. It also includes a couple col type 3 , but this would not necessarily occur. It could also include col type 4.
$\left[\left\{\boldsymbol{Q}_{t}\right\}_{2,2} \quad\left\{\boldsymbol{Q}_{t}\right\}_{2,3}\right]$ spans columns 2-5 (light purple), to cover all col type 2 and 3. In this case, it does not include any col type 1 or 4 , but it could.

The code copies $\boldsymbol{Q}_{t}(:$, itype $)=\boldsymbol{Q}(:$, ideflate $)$, that is, $\boldsymbol{Q}_{t} \boldsymbol{P}_{t}^{T}=\boldsymbol{Q} \boldsymbol{P}_{d}^{T}$. Thus $\boldsymbol{Q}_{t}(:$, itype $)$ is in the same order as $\boldsymbol{D}$ (ideflate).

Globally permuted If we globally permuted the matrix (e.g., in the serial algorithm), col types would be sorted globally: $1,2,3,4$.
In this case, $\left[\begin{array}{cc}\overline{\overline{\boldsymbol{Q}}}_{11} & \overline{\bar{Q}}_{12}\end{array}\right]$ would span columns 1-3 (dark purple), to cover all col type 1 and 2. No col type 3 would be included.
$\left[\begin{array}{cc}\overline{\bar{Q}}_{22} & \overline{\bar{Q}}_{23}\end{array}\right]$ would span columns 2-5 (light purple), to cover all col type 2 and 3 . No col type 1 would be included.

Again, the code doesn't compute igbar. It computes iglobal such that $Q_{G}=Q_{g}(:, \mathrm{iQ})$. Hence igbar = itype( iQ ).

### 3.1.7 Example 0: no deflation

Using the above setup, with no entries deflated. With no deflation, $P_{t}=P_{d}$, so $P_{t}^{T} P_{d}=I$ and $Q_{t}=Q$.

To run the example in parallel:

| --verbose 1 --ref n --print-precision 2 --print-width 6 \} stedc_deflate |  |  |  |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $D=$ | 2.10 | 4.10 | 6.10 | 8.10 | 3.30 | 6.30 | 9.30 | 12.30 |
| $z=$ | 0.21 | 0.41 | 0.61 | 0.81 | 0.33 | 0.63 | 0.93 | 1.23 |
| $Q(1,:)=$ | 0.21 | 0.41 | 0.61 | 0.81 | 0 | 0 | 0 | 0. |
| $Q(n,:)=$ | 0 | 0 | 0 | 0 | 0.33 | 0.63 | 0.93 | 1.23 |
| pcolumn $=$ | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 |
| coltype $=$ | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| $P_{s}=$ isort $=$ | 0 | 4 | 1 | 2 | 5 | 3 | 6 | 7 |
| $D_{s}=P_{s} D=D($ isort $)=$ | 2.10 | 3.30 | 4.10 | 6.10 | 6.30 | 8.10 | 9.30 | 12.30 |
| $z_{s}=P_{s} z=z($ isort $)=$ | 0.21 | 0.33 | 0.41 | 0.61 | 0.63 | 0.81 | 0.93 | 1.23 |
| $Q(1$, isort $)=$ | 0.21 | 0 | 0.41 | 0.61 | 0 | 0.81 | 0 | 0. |
| $Q(n$, isort $)=$ | 0 | 0.33 | 0 | 0 | 0.63 | 0 | 0.93 | 1.23 |
| pcolumn ( isort ) $=$ | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 |
| coltype ( isort $)=$ | 1 | 3 | 1 | 1 | 3 | 1 | 3 | 3 |
| $P_{d}=$ ideflate $=$ | 0 | 4 | 1 | 2 | 5 | 3 | 6 | 7 |
| $D_{d}=P_{d} D=D($ ideflate $)=$ | 2.10 | 3.30 | 4.10 | 6.10 | 6.30 | 8.10 | 9.30 | 12.30 |
| $D_{\text {secular }}=$ | 2.10 | 3.30 | 4.10 | 6.10 | 6.30 | 8.10 | 9.30 | 12.30 |
| $z_{d}=P_{d} z=z($ ideflate $)=$ | 2.97 | 4.67 | 5.80 | 8.63 | 8.91 | 11.46 | 13.15 | 17.39 |
| $z_{\text {secular }}=$ | 0.21 | 0.33 | 0.41 | 0.61 | 0.63 | 0.81 | 0.93 | 1.23 |
| $Q(1$, ideflate $)=$ | 0.21 | 0 | 0.41 | 0.61 | 0 | 0.81 | 0 | 0. |
| $Q(n$, ideflate $)=$ | 0 | 0.33 | 0 | 0 | 0.63 | 0 | 0.93 | 1.23 |
| $\text { pcolumn ( ideflate })=$ | 0 | 0 | 0 | 1 | 0 | 1 | 1 | 1 |
| $\text { coltype }(\text { ideflate })=$ | 1 | 3 | $1$ | 1 | 3 | 1 | 3 | 3 |
| coltype post-deflation $=$ | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| $P_{t}=$ itype $=$ | 0 | 4 | 1 | 2 | 5 | 3 | 6 | 7 |
| $P_{t}^{T} P_{d}=$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| $P_{t}^{T} P_{d} D=$ | 2.10 | 4.10 | 6.10 | 8.10 | 3.30 | 6.30 | 9.30 | 12.30 |
| $D_{\text {deflated }}=$ | - | - |  | - | - | - |  |  |
| $Q_{t}(1,:)=$ | 0.21 | 0.41 | 0.61 | 0.81 | 0 | 0 | 0 | 0. |
| $Q_{t}(n,:)=$ | 0 | 0 | 0 | 0 | 0.33 | 0.63 | 0.93 | 1.23 |
| $P_{t}^{T} P_{d}$ pcolumn $=$ | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 |
| $P_{t}^{T} P_{d}$ coltype $=$ | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| local $P_{t}^{T} P_{d}$ pcolumn $=$ | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 |
| local $P_{t}^{T} P_{d}$ coltype $=$ | 1 | 1 | 3 | 3 | 1 | 1 | 3 | 3 |
| $P_{g}=$ iglobal $=$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| $P_{g} P_{t}^{T} P_{d} D=$ | 2.10 | 4.10 | 6.10 | 8.10 | 3.30 | 6.30 | 9.30 | 12.30 |
| $P_{g} P_{t}^{T} P_{d}$ coltype $=$ | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |

### 3.1.8 Example 1: type 1 deflation - serial

Using the above setup, then set $d_{0} \approx 0, d_{4} \approx 0, d_{7} \approx 0$, which triggers type 1 deflation for those 3 eigvals. In serial, all 3 are moved to the end of the matrix. Having no type 2 deflation, there are no coltype 2 .

To run the example in serial:

```
slate/test> mpirun -np 1 ./tester --dim 8 --nb 2 --deflate '0 4 6' \
    --verbose 1 --ref n --print-precision 2 --print-width 6 \
    stedc_deflate
```

| $D=$ | 2.10 | 4.10 | 6.10 | 8.10 | 3.30 | 6.30 | 9.30 | 12.30 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $z=$ | $\varepsilon$ | 0.41 | 0.61 | 0.81 | $\varepsilon$ | 0.63 | $\varepsilon$ | 1.23 |
| $Q(1,:)=$ | 0.21 | 0.41 | 0.61 | 0.81 | 0 | 0 | 0 | 0. |
| $Q(n,:)=$ | 0 | 0 | 0 | 0 | 0.33 | 0.63 | 0.93 | 1.23 |
| pcolumn $=$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| coltype $=$ | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| $P_{s}=$ isort $=$ | 0 | 4 | 1 | 2 | 5 | 3 | 6 | 7 |
| $D_{s}=P_{s} D=D($ isort $)=$ | 2.10 | 3.30 | 4.10 | 6.10 | 6.30 | 8.10 | 9.30 | 12.30 |
| $z_{s}=P_{s} z=z($ isort $)=$ | $\varepsilon$ | $\varepsilon$ | 0.41 | 0.61 | 0.63 | 0.81 | $\varepsilon$ | 1.23 |
| $Q(1$, isort $)=$ | 0.21 | 0 | 0.41 | 0.61 | 0 | 0.81 | 0 | 0. |
| $Q(n$, isort $)=$ | 0 | 0.33 | 0 | 0 | 0.63 | 0 | 0.93 | 1.23 |
| pcolumn( isort ) $=$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| coltype( isort ) $=$ | 1 | 3 | 1 | 1 | 3 | 1 | 3 | 3 |
| $P_{d}=$ ideflate $=$ | 1 | 2 | 5 | 3 | 7 | 6 | 4 | 0 |
| $D_{d}=P_{d} D=D($ ideflate $)=$ | 4.10 | 6.10 | 6.30 | 8.10 | 12.30 | 9.30 | 3.30 | 2.10 |
| $D_{\text {secular }}=$ | 4.10 | 6.10 | 6.30 | 8.10 | 12.30 | - | - | - |
| $z_{d}=P_{d} z=z($ ideflate $)=$ | 5.80 | 8.63 | 8.91 | 11.46 | 17.39 | 13.15 | 4.67 | 2.97 |
| $z_{\text {secular }}=$ | 0.41 | 0.61 | 0.63 | 0.81 | 1.23 | - | - | - |
| $Q(1$, ideflate $)=$ | 0.41 | 0.61 | 0 | 0.81 | 0 | 0 | 0 | 0.21 |
| $Q(n$, ideflate $)=$ | 0 | 0 | 0.63 | 0 | 1.23 | 0.93 | 0.33 | 0. |
| pcolumn ( ideflate ) $=$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| coltype( ideflate $)=$ | 1 | 1 | 3 | 1 | 3 | 4 | 4 | 4 |
| coltype post-deflation $=$ | 4 | 1 | 1 | 1 | 4 | 3 | 4 | 3 |
| $P_{t}=$ itype $=$ | 0 | 1 | 3 | 2 | 4 | 5 | 6 | 7 |
| $P_{t}^{T} P_{d}=$ | 1 | 2 | 3 | 5 | 7 | 6 | 4 | 0 |
| $P_{t}^{T} P_{d} D=$ | 4.10 | 6.10 | 8.10 | 6.30 | 12.30 | 9.30 | 3.30 | 2.10 |
| $D_{\text {deflated }}=$ | - | - | - | - | - | 9.30 | 3.30 | 2.10 |
| $Q_{t}(1,:)=$ | 0.41 | 0.61 | 0.81 | 0 | 0 | 0 | 0 | 0.21 |
| $Q_{t}(n,:)=$ | 0 | 0 | 0 | 0.63 | 1.23 | 0.93 | 0.33 | 0. |
| $P_{t}^{T} P_{d}$ pcolumn $=$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| $P_{t}^{T} P_{d}$ coltype $=$ | 1 | 1 | 1 | 3 | 3 | 4 | 4 | 4 |
| local $P_{t}^{T} P_{d}$ pcolumn $=$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| local $P_{t}^{T} P_{d}$ coltype $=$ | 1 | 1 | 1 | 3 | 3 | 4 | 4 | 4 |
| $P_{g}=$ iglobal $=$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
| $P_{g} P_{t}^{T} P_{d} D=$ | 4.10 | 6.10 | 8.10 | 6.30 | 12.30 | 9.30 | 3.30 | 2.10 |
| $P_{g} P_{t}^{T} P_{d}$ coltype $=$ | 1 | 1 | 1 | 3 | 3 | 4 | 4 | 4 |

## Example 1: type 1 deflation - parallel

In parallel with 2 process columns ( 4 MPI ranks in $2 \times 2$ grid), entries 0 and 4 are moved to the end of process col 0 , so are unfortunately included in $\left\{Q_{t}\right\}_{1,1: 2}$ and $\left\{Q_{t}\right\}_{2,2: 3}$, while entry 7 is moved to the end of process col 1 .

To run the example in parallel:

```
slate/test> mpirun -np 4 ./tester --dim 8 --nb 2 --deflate '0 4 6' \}
    --verbose 1 --ref n --print-precision 2 --print-width 6 \}
    stedc_deflate
```

$\left.\begin{array}{rl}D & = \\ z & 2.10 \\ \varepsilon & 4.10 \\ & 0.41 \\ 6.10 & 0.61 \\ \hline\end{array}\right)$

### 3.1.9 Example 2: type 2 deflation - serial

Using the above setup, deflate pairs $\left\{d_{0}, d_{5}\right\}$ and $\left\{d_{4}, d_{7}\right\}$. The first entry of each pair is deflated (column type 4). Since $d_{0}$ is in $\boldsymbol{Q}_{1}$ and $d_{5}$ is in $\boldsymbol{Q}_{2}, d_{5}$ becomes column type 2. Since $d_{4}$ and $d_{7}$ are both in $\boldsymbol{Q}_{2}, d_{7}$ remains column type 3 .

To run the example in serial:

```
slate/test> mpirun -np 1 ./tester --dim 8 --nb 2 --deflate '0/5 4/7' \}
--verbose 1 --ref n --print-precision 2 --print-width 6 \}
stedc_deflate
```

| $D$ | $=$ |
| ---: | :--- |
| $z$ | 2.10 |
|  | 4.10 |
|  | 6.10 |
|  | 8.10 |
|  | 3.30 |

## Example 2: type 2 deflation - parallel

Lorem ipsum.

To run the example in parallel:

```
slate/test> mpirun -np 4 ./tester --dim 8 --nb 2 --deflate '0/5 4/7' \}
--verbose 1 --ref n --print-precision 2 --print-width 6 \}
stedc_deflate
```



### 3.1.10 Example 3: type 1 and 2 deflation - serial

Using the above setup, deflate $d_{3}$ (type 1) and pairs $\left\{d_{1}, d_{4}\right\}$ and $\left\{d_{2}, d_{5}\right\}$ (type 2).

To run the example in serial:

```
slate/test> mpirun -np 1 ./tester --dim 8 --nb 2 --deflate '3 1/4 2/5' \}
    --verbose 1 --ref \(n\)--print-precision 2 --print-width 6 \}
    stedc_deflate
```

| $\begin{aligned} D & = \\ z & = \\ Q(1,:) & = \\ Q(n,:) & = \\ \text { pcolumn } & =\end{aligned}$ | 2.10 | 4.10 | 6.10 | 8.10 | 4.10 | 6.10 | 9.30 | 12.30 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0.21 | 0.41 | 0.61 | $\varepsilon$ | 0.33 | 0.63 | 0.93 | 1.23 |
|  | 0.21 | 0.41 | 0.61 | 0.81 | 0 | 0 | 0 | 0. |
|  | 0 | 0 | 0 | 0 | 0.33 | 0.63 | 0.93 | 1.23 |
|  | 0 | 0 | 0 | 0 | 0 | , | 0 | 0 |
|  | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| $P_{s}=$ isort $=$ | 0 | 1 | 4 | 2 | 5 | 3 | 6 | 7 |
| $D_{s}=P_{s} D=D($ isort $)=$ | 2.10 | 4.10 | 4.10 | 6.10 | 6.10 | 8.10 | 9.30 | 12.30 |
| $z_{s}=P_{s} z=z($ isort $)=$ | 0.21 | 0.41 | 0.33 | 0.61 | 0.63 | $\varepsilon$ | 0.93 | 1.23 |
| $Q(1$, isort $)=$ | 0.21 | 0.41 | 0 | 0.61 | 0 | 0.81 | 0 | 0. |
| $Q(n$, isort $)=$ | 0 | 0 | 0.33 | 0 | 0.63 | 0 | 0.93 | 1.23 |
| $\operatorname{pcolumn}($ isort $)=$ | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
| coltype ( isort $)=$ | 1 | 1 | 3 | 1 | 3 | 1 | 3 | 3 |
| $\begin{aligned} & P_{d}=\text { ideflate }= \\ & D_{d}=P_{d} D=D(\text { ideflate })= \\ & D_{\text {secular }}= \\ & z_{d}=P_{d} z=z(\text { ideflate })= \\ & z_{\text {secular }}= \\ & Q(1, \text { ideflate })= \\ & Q(n, \text { ideflate })= \\ &\text { pcolumn( ideflate })= \\ &\text { coltype(ideflate })= \\ & \text { coltype post-deflation }= \end{aligned}$ | 0 | 4 | 5 | 6 | 7 | 3 | 2 | 1 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | 8.10 | 6.10 | 4.10 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | - | - | - |
|  | 2.97 | 5.80 | 8.63 | 13.15 | 17.39 | 11.46 | 8.63 | 5.80 |
|  | 0.21 | 0.53 | 0.88 | 0.93 | 1.23 | - | - | - |
|  | 0.21 | 0.32 | 0.42 | 0 | 0 | 0.81 | 0.44 | 0.26 |
|  | 0 | 0.21 | 0.45 | 0.93 | 1.23 | 0 | -0.44 | -0.26 |
|  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
|  | 1 | 2 | 2 | 3 | 3 | 4 | 4 | 4 |
|  | 1 | 4 | 4 | 4 | 2 | 2 | 3 | 3 |
| $\begin{aligned} & P_{t}=\text { itype }= \\ & P_{t}^{T} P_{d}= \\ & P_{t}^{T} P_{d} D= \\ & D_{\text {deflated }}= \\ & Q_{t}(1,)= \\ & Q_{t}(n,:)= \\ &= \\ & P_{t}^{T} P_{d} \text { pcolumn }= \\ & P_{t}^{T} P_{d} \text { coltype }= \\ & \text { local } P_{t}^{T} P_{d} \text { pcolumn }=\end{aligned}$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|  | 0 | 4 | 5 | 6 | 7 | 3 | 2 | 1 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | 8.10 | 6.10 | 4.10 |
|  | - | - | - | - | - | 8.10 | 6.10 | 4.10 |
|  | 0.21 | 0.32 | 0.42 | 0 | 0 | 0.81 | 0.44 | 0.26 |
|  | 0 | 0.21 | 0.45 | 0.93 | 1.23 | 0 | -0.44 | -0.26 |
|  | 0 |  | 0 | 0 | 0 | 0 | 0 | 0 |
|  | 1 | 2 | 2 | 3 | 3 | 4 | 4 | 4 |
|  | 0 | 0 | 0 | 0 | 0 | 0 | 0 | 0 |
|  | 1 | 2 | 2 | 3 | 3 | 4 | 4 | 4 |
| $\begin{aligned} P_{g}=\text { iglobal } & = \\ P_{g} P_{t}^{T} P_{d} D & = \\ P_{g} P_{t}^{T} P_{d} \text { coltype } & = \end{aligned}$ | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | 8.10 | 6.10 | 4.10 |
|  | 1 | 2 | 2 | 3 | 3 | 4 | 4 | 4 |

## Example 3: type 1 and 2 deflation - parallel

Lorem ipsum.

To run the example in parallel:

```
slate/test> mpirun -np 4 ./tester --dim 8 --nb 2 --deflate '3 1/4 2/5' \
--verbose 1 --ref n --print-precision 2 --print-width 6 \
stedc_deflate
```

| $\begin{aligned} D & = \\ z & = \\ Q(1,:) & = \\ Q(n,:) & = \\ \text { pcolumn } & = \\ \text { coltype } & =\end{aligned}$ | 2.10 | 4.10 | 6.10 | 8.10 | 4.10 | 6.10 | 9.30 | 12.30 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0.21 | 0.41 | 0.61 | $\varepsilon$ | 0.33 | 0.63 | 0.93 | 1.23 |
|  | 0.21 | 0.41 | 0.61 | 0.81 | 0 | 0 | 0 | 0. |
|  | 0 | 0 | 0 | 0 | 0.33 | 0.63 | 0.93 | 1.23 |
|  | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 |
|  | 1 | 1 | 1 | 1 | 3 | 3 | 3 | 3 |
| $P_{s}=$ isort $=$ | 0 | 1 | 4 | 2 | 5 | 3 | 6 | 7 |
| $D_{s}=P_{s} D=D($ isort $)=$ | 2.10 | 4.10 | 4.10 | 6.10 | 6.10 | 8.10 | 9.30 | 12.30 |
| $z_{s}=P_{s} z=z($ isort $)=$ | 0.21 | 0.41 | 0.33 | 0.61 | 0.63 | $\varepsilon$ | 0.93 | 1.23 |
| $Q(1$, isort $)=$ | 0.21 | 0.41 | 0 | 0.61 | 0 | 0.81 | 0 | 0. |
| $Q(n$, isort $)=$ | 0 | 0 | 0.33 | 0 | 0.63 | 0 | 0.93 | 1.23 |
| pcolumn $($ isort $)=$ | 0 | 0 | 0 | 1 | 0 | 1 |  | 1 |
| coltype ( isort $)=$ | 1 | 1 | 3 | 1 | 3 | 1 | 3 | 3 |
| $P_{d}=$ ideflate $=$ <br> $D_{d}=P_{d} D=D($ ideflate $)=$ <br> $D_{\text {secular }}=$ <br> $z_{d}=P_{d} z=z($ ideflate $)=$ <br> $z_{\text {secular }}=$ <br> $Q(1$, ideflate $)=$ <br> $Q(n$, ideflate $)=$ <br> pcolumn $($ ideflate $)=$ coltype $($ ideflate $)=$ coltype post-deflation $=$ | 0 | 4 | 5 | 6 | 7 | 3 | 2 | 1 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | 8.10 | 6.10 | 4.10 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | - | - | - |
|  | 2.97 | 5.80 | 8.63 | 13.15 | 17.39 | 11.46 | 8.63 | 5.80 |
|  | 0.21 | 0.53 | 0.88 | 0.93 | 1.23 | - | - | - |
|  | 0.21 | 0.32 | 0.42 | 0 | 0 | 0.81 | 0.44 | 0.26 |
|  | 0 | 0.21 | 0.45 | 0.93 | 1.23 | 0 | -0.44 | -0.26 |
|  | 0 | 0 | 0 | 1 | 1 | 1 | 1 | 0 |
|  | 1 | 2 | 2 | 3 | 3 | 4 | 4 | 4 |
|  | 1 | 4 | 4 | 4 | 2 | 2 | 3 | 3 |
|  | 0 | 1 | 4 | 2 | 3 | 6 | 7 | 5 |
| $\begin{array}{r} P_{t}=\text { itype }= \\ P_{t}^{T} P_{d}= \\ P_{t}^{T} P_{d} D= \\ D_{\text {deflated }}= \end{array}$ | 0 | 4 | 6 | 7 | 5 | 1 | 3 | 2 |
|  | 2.10 | 4.10 | 9.30 | 12.30 | 6.10 | 4.10 | 8.10 | 6.10 |
|  | - | - | - | - | - | 4.10 | 8.10 | 6.10 |
| $Q_{t}(1,:)$$=$$Q_{t}(n,:)$ | 0.21 | 0.32 | 0 | 0 | 0.42 | 0.26 | 0.81 | 0.44 |
|  | 0 | 0.21 | 0.93 | 1.23 | 0.45 | -0.26 | 0 | -0.44 |
| $P_{t}^{T} P_{d}$ pcolumn $=$ | 0 | 0 | 1 | 1 | 0 | 0 | 1 | 1 |
| $P_{t}^{T} P_{d}$ coltype $=$ | 1 | 2 | 3 | 3 | 2 | 4 | 4 | 4 |
| local $P_{t}^{T} P_{d}$ pcolumn $=$ local $P_{t}^{T} P_{d}$ coltype $=$ | 0 | 0 | 0 | 0 | 1 | 1 | 1 | 1 |
|  | 1 | 2 |  |  | 3 | 3 | 4 | 4 |
| $\begin{array}{r} P_{g}=\text { iglobal }= \\ P_{g} P_{t}^{T} P_{d} D= \\ P_{g} P_{t}^{T} P_{d} \text { coltype }= \end{array}$ |  | 1 | 4 | 2 | 3 | 6 | 7 | 5 |
|  | 2.10 | 4.10 | 6.10 | 9.30 | 12.30 | 8.10 | 6.10 | 4.10 |
|  | 1 | 2 |  | 3 | 3 | 4 | 4 | 4 |

### 3.1.11 Eigenvectors via Löwner Theorem

When computing eigenvectors from the computed eigenvalues $\left\{\tilde{\lambda}_{j}\right\}$ of $\boldsymbol{D}+\rho \boldsymbol{z} \boldsymbol{z}^{T}$ using (3.5), repeated here,

$$
\boldsymbol{u}_{j}=\left(\boldsymbol{D}-\lambda_{j} \boldsymbol{I}\right)^{-1} \boldsymbol{z}=\left[\begin{array}{c}
\frac{z_{1}}{d_{1}-\lambda_{j}} \\
\vdots \\
\frac{z_{n}}{d_{n}-\lambda_{j}}
\end{array}\right] \quad \text { for } j=1, \ldots, n,
$$

if $d_{i} \approx \lambda_{j}$, then $z_{i} /\left(d_{i}-\lambda_{j}\right)$ can be inaccurate, causing a loss of orthogonality in $\boldsymbol{U}$. Instead, consider $\left\{\tilde{\lambda}_{j}\right\}$ as exact eigenvalues of the modified system $\boldsymbol{D}+\rho \tilde{\boldsymbol{z}} \tilde{\boldsymbol{z}}^{T}$. See derivation in Tisseur and Dongarra [15]; here each term is multiplied by -1 to match laed4 code. Result: ${ }^{2}$

$$
\tilde{z}_{i}= \pm \sqrt{-\frac{\prod_{j=1}^{n} d_{i}-\tilde{\lambda}_{j}}{\prod_{\substack{j=1 \\ j \neq i}}^{n} d_{i}-d_{j}}},= \pm \sqrt{-\frac{\prod_{j=1}^{n} \delta_{i, j}}{\prod_{\substack{j=1 \\ j \neq i}}^{n} d_{i}-d_{j}}} \quad \text { with } \delta_{i, j}=d_{i}-\tilde{\lambda}_{j},
$$

yielding (unnormalized) eigenvectors

$$
\tilde{\boldsymbol{u}}_{j}=\left(\boldsymbol{D}-\tilde{\lambda}_{j} \boldsymbol{I}\right)^{-1} \tilde{\boldsymbol{z}}=\left[\begin{array}{c}
\frac{\tilde{z}_{1}}{d_{1}-\tilde{\lambda}_{j}} \\
\vdots \\
\frac{\tilde{z}_{n}}{d_{n}-\tilde{\lambda}_{j}}
\end{array}\right]=\left[\begin{array}{c}
\frac{\tilde{z}_{1}}{\delta_{1, j}} \\
\vdots \\
\frac{\tilde{z}_{n}}{\delta_{n, j}}
\end{array}\right]=\tilde{\boldsymbol{z}} \oslash \boldsymbol{\delta}_{j} \quad \text { for } j=1, \ldots, n,
$$

where $\oslash$ denotes element-wise division (Matlab ./).
Both Tisseur and Dongarra [15] and Gu and Eisenstat [19] seem to gloss over the sign of $\tilde{z}_{i}$. (Sca)LAPACK copies the sign from $z_{i}$ in dlaed3.f:
$\mathrm{w}(\mathrm{i})=\operatorname{sign}(\operatorname{sqrt}(-\mathrm{w}(\mathrm{i})), \mathrm{s}(\mathrm{i})$ )
where w (ztilde) contains the "first $k$ (nsecular) values of the final deflation-altered $\boldsymbol{z}$-vector", per dlaed2.

[^2]
## Parallelization of eigenvectors

ScaLAPACK computes $n_{s}$ roots on $p$ processes as follows:
$n_{c}=n / n_{\text {pcol }}(\mathrm{klc})$ is the number of eigenvalues each process column computes.
$n_{r}=n_{c} / n_{\text {prow }}(\mathrm{klr})$ is the number of eigenvalues each process row within a process column computes.
Each process computes $n_{r}$ eigenvalues, calling LAPACK's laed4 for each. (Within each pcol, prow 0 (or drow?) computes any remainder. Globally, pcol 0 (or dcol?) computes any remainder.)

After each call to laed4, each process updates its $\overline{\boldsymbol{z}}$ vector of partial products, multiplying one more term:

$$
\bar{z}_{i}=\frac{\prod_{k=1}^{n_{r}}\left(d_{i}-\tilde{\lambda}_{k_{l}}\right)}{\prod_{\substack{k=1 \\ k_{g} \neq i}}^{n_{r}}\left(d_{i}-d_{k_{g}}\right)}=\frac{\prod_{k=1}^{n_{r}} \delta_{i, i}}{\prod_{\substack{k_{r}=1 \\ k_{g} \neq i}}^{n_{r}}\left(d_{k_{g}}-d_{i}\right)} \quad \text { for } i=1 \ldots n,
$$

where $k_{g}$ maps from local index $k$ to global index $k_{g}$.
Then it does a global multiply reduction (in a 2D fashion) of $\overline{\boldsymbol{z}}$, similar to MPI_Reduce. The root node finishes the computation of $\tilde{\boldsymbol{z}}$,

$$
\tilde{z}_{i}=\operatorname{sign}\left(z_{i}\right) \sqrt{-\prod_{\text {process } p} \bar{z}_{i}^{(p)}} \quad \text { for } i=1 \ldots n,
$$

then broadcasts $\tilde{\boldsymbol{z}}$.
It also gathers (in a 2D fashion) the computed eigenvalues, $\tilde{\boldsymbol{\Lambda}}$, to the root node, then broadcasts it back out, similar to MPI_Allgather.

Now, each process computes its local portion of the $\boldsymbol{U}$ matrix, per the 2D block cyclic distribution. For each column $j$ of $\boldsymbol{U}$ that pcol $p$ owns, each process in pcol $p$ redundantly re-computes the corresponding eigenvalue $\lambda_{j}$ and $\boldsymbol{\delta}_{j}$ vector. Each process in pcol then redundantly computes the entire vector $\boldsymbol{u}_{j}=\tilde{\boldsymbol{z}}_{j} \oslash \boldsymbol{\delta}_{j}$, takes its norm, and normalizes and saves just the portion that the process owns. Note this calls laed4 redundantly nprow times, thus limiting its parallel speedup to npcol times. However, it is $O\left(n^{2}\right)$ work, so may not take significant time.

Despite the definition

$$
\delta_{i, j}=d_{i}-\tilde{\lambda}_{j},
$$

using that definition - even with the compute $\tilde{\lambda}_{j}$ - produces inaccurate results (as confirmed in MAGMA). This is unfortunate as recomputing $\boldsymbol{\delta}$ in that fashion would avoid all the redundant laed4 calls. It's also unclear since Gu and Eisenstat [19] seems to use that definition. Neither Rutter [18] nor Li [17] seem to discuss computing $\delta_{i, j}$ in laed4 and why that might be necessary for stability.

Alternatively, it could skip all this redundant computation by saving the $\boldsymbol{\delta}_{j}$ vectors and communicating them, and doing a distributed computation of column norms to normalize the vectors.

### 3.1.12 Cost

Without deflation, flops for multiplying $\boldsymbol{Q U}$ is $n^{3}+O\left(n^{2}\right)$, since it is two gemms of size $\frac{n}{2} \times \frac{n}{2} \times n$ :

$$
\boldsymbol{Q} \boldsymbol{U}=\left[\begin{array}{cc}
\boldsymbol{Q}_{1} & \mathbf{0} \\
\mathbf{0} & \boldsymbol{Q}_{2}
\end{array}\right]\left[\boldsymbol{U}_{1} \boldsymbol{U}_{2}\right]=\left[\begin{array}{l}
\boldsymbol{Q}_{1} \boldsymbol{U}_{1} \\
\boldsymbol{Q}_{2} \boldsymbol{U}_{2}
\end{array}\right] .
$$

The total cost for divide-and-conquer thus satisfies the recursion

$$
t_{n}=n^{3}+2 t_{n / 2}
$$

with solution

$$
t_{n} \approx \frac{4}{3} n^{3}+O\left(n^{2}\right) .
$$

With deflation, flops can be $O\left(n^{2.3}\right)$ on average, or $O\left(n^{2}\right)$ in special cases.

### 3.2 Routines

### 3.2.1 stedc

Symmetric Tridiagonal Eigenvalue Divide \& Conquer solver, top-level routine called from heev.

```
Algorithm 8 Main divide \& conquer driver
    function \(\operatorname{stedc}(\boldsymbol{D}, \boldsymbol{E}, \boldsymbol{Q})\)
    input: real tridiagonal matrix \(\boldsymbol{A}\) represented by diagonal \(\boldsymbol{D}\) and sub-diagonal \(\boldsymbol{E}\) vectors
    output: eigvals \(\boldsymbol{D}\) (sorted) and eigvecs \(\boldsymbol{Q}\) of \(\boldsymbol{A}\)
        scale \(\boldsymbol{A}\) (i.e., \(\boldsymbol{D}\) and \(\boldsymbol{E}\) ) by \(1 /\|\boldsymbol{A}\|\), so it has unit norm
        allocate workspaces \(\boldsymbol{W}, \boldsymbol{U}\)
        // Computing in workspace \(\boldsymbol{W}\) avoids copy in sort, compared to ScaLAPACK.
        stedc_solve \((\boldsymbol{D}, \boldsymbol{E}, \boldsymbol{W}\); workspace \(\boldsymbol{Q}, \boldsymbol{U})\) computes eigvals \(\boldsymbol{D}\) and eigvecs \(\boldsymbol{W}\)
        stedc_sort \((\boldsymbol{D}, \boldsymbol{W}, \boldsymbol{Q})\) sorts eigvals \(\boldsymbol{D}\) and permutes eigvecs \(\boldsymbol{W}\) into \(\boldsymbol{Q}\)
        scale eigvals \(\boldsymbol{D}\) back by original \(\|\boldsymbol{A}\|\)
    end function
```


### 3.2.2 stedc_solve

Main divide \& conquer driver.
Corresponds to ScaLAPACK laedo.

```
Algorithm 9 Main divide \& conquer driver
    function stedc_solve \((\boldsymbol{D}, \boldsymbol{E}, \boldsymbol{Q}\); workspace \(\boldsymbol{W}, \boldsymbol{U})\)
    input: \(\boldsymbol{A}\) represented by diagonal \(\boldsymbol{D}\) and sub-diagonal \(\boldsymbol{E}\) vectors.
    output: \(\boldsymbol{D}\) is (unsorted) eigvals and \(\boldsymbol{Q}\) is eigvecs of \(\boldsymbol{A}\).
        // Tear into subproblems
        for \(i=b, 2 b, \ldots, n-1\) with blocksize \(b\)
            subtract \(\rho=\left|E_{i-1}\right|\) from \(D_{i-1}\) and \(D_{i}\)
        end
        // Solve subproblems
        parallel for each block-col \(i=0, b, \ldots, n\); split over MPI ranks and OpenMP threads
            // todo: In 2DBC, seems highly load imbalanced
            // -only nodes assigned diagonal tiles do work.
            if \(\boldsymbol{Q}_{i, i}\) is local then
                \(i_{2}=i+n b-1\)
                    lapack: :steqr \(\left(\boldsymbol{D}_{i: i_{2}}, \boldsymbol{E}_{i: i_{2}-1}, \boldsymbol{Q}_{i, i}\right)\) or stedc
                        to solve subproblem using serial algorithm
            end
        end
        gather and bcast \(\boldsymbol{D}\) to all nodes
        // Merge subproblems
        for each level in divide \& conquer tree, from leaf to root
            for each pair of subproblems, indexed \(i, \ldots, i_{1}\) and \(i_{1}+1, \ldots, i_{2}\)
            \(\rho=E_{i_{1}} / /\) abs dealt with in merge routine
            stedc_merge \(\left(\rho, D_{i: i_{2}}, Q_{i: i_{2}, i: i_{2}}, W_{i: i_{2}, i: i_{2}}, U_{i: i_{2}, i: i_{2}}\right)\)
            end
        end
    end function
```


### 3.2.3 stedc_merge

Merges two subproblems.
Corresponds to ScaLAPACK laed1.

```
function stedc_merge \(\left(\boldsymbol{D}, \boldsymbol{Q}\right.\); workspace \(\left.\boldsymbol{Q}_{t}, \boldsymbol{U}\right)\)
input: eigvals \(\boldsymbol{D}_{1}\) and \(\boldsymbol{D}_{2}\) in \(\boldsymbol{D}\), eigvecs \(\boldsymbol{Q}_{1}\) and \(\boldsymbol{Q}_{2}\) in \(\boldsymbol{Q}\) of subproblems
output: eigvals \(\boldsymbol{D}\) and eigvecs \(\boldsymbol{Q}\) of merged problem
    stedc_z_vector \((\boldsymbol{Q}, \boldsymbol{z})\) gets \(\boldsymbol{z}\)
    stedc_deflate \(\left(\boldsymbol{D}, \boldsymbol{z}, \boldsymbol{D}_{s}, \boldsymbol{z}_{s}, \boldsymbol{Q}, \boldsymbol{Q}_{t}, \boldsymbol{P}_{t}\right)\) deflates \(n_{d}\) eigvals, leaving \(n_{s}\) secular eqn
eigvals
    stedc_secular \(\left(\boldsymbol{D}_{s}, \boldsymbol{z}_{s}, \boldsymbol{U}, \boldsymbol{P}_{t}\right)\) solves secular equation for eigvals \(\boldsymbol{D}_{s}\) and eigvecs \(\boldsymbol{U}\)
    // todo: merge \(\boldsymbol{D}\) and \(\boldsymbol{D}_{s}\)
    \(\boldsymbol{Q}_{1,1: 2}=\left\{\boldsymbol{Q}_{t}\right\}_{1,1: 2} \boldsymbol{U}_{1: 2}\)
    \(\boldsymbol{Q}_{2,2: 3}=\left\{\boldsymbol{Q}_{t}\right\}_{2,2: 3} \boldsymbol{U}_{2: 3}\)
    Copy with permutation deflated eigvecs \(\left\{\boldsymbol{Q}_{t}\right\}_{1: 2,4}\) to \(\boldsymbol{Q}\)
end function
```


### 3.2.4 stedc_z_vector

Gathers onto all nodes vector $\boldsymbol{z}$ that is last row of $\boldsymbol{Q}_{1}$ and first row of $\boldsymbol{Q}_{2}$,

$$
\boldsymbol{z}=\boldsymbol{Q}^{T}\left[\begin{array}{c}
e_{n_{1}} \\
e_{1}
\end{array}\right]=\left[\begin{array}{c}
\boldsymbol{Q}_{1}^{T} e_{n_{1}} \\
\boldsymbol{Q}_{2}^{T} e_{1}
\end{array}\right]
$$

Corresponds to ScaLAPACK laedz.
This is conceptually like MPI Allgatherv, but due to 2DBC distribution, it doesn't seem a single Allgatherv could do this. Alternatively, each rank could pack its local pieces, then do MPI Gatherv, root unpacks to correct locations, and MPI Bcast; or MPI Allgatherv and everyone unpacks (without bcast).

```
function stedc_z_vector \((\boldsymbol{Q}, \boldsymbol{z})\)
input: eigvecs \(\boldsymbol{Q}_{1}\) and \(\boldsymbol{Q}_{2}\) of subproblems in \(\boldsymbol{Q}\)
output: vector \(\boldsymbol{z}\) is last row of \(\boldsymbol{Q}_{1}\) and first row of \(\boldsymbol{Q}_{2}\)
    for each block-col \(j=0, \ldots, n_{t}-1\)
        if \(j<n_{t 1}\) then
            \(i=n_{t_{1}}-1 \quad / /\) last block-row of \(\boldsymbol{Q}_{1}\)
        else
            \(i=n_{t_{1}} \quad / /\) first block-row of \(\boldsymbol{Q}_{2}\)
        end
        if \(\boldsymbol{Q}_{i, j}\) is local then
            copy last or first row of \(\boldsymbol{Q}_{i, j}\) to \(\boldsymbol{z}_{i: i+b}\)
            MPI_send \(\boldsymbol{z}_{i: i+b}\) to root, if rank \(\neq\) root
        else if rank \(==\) root then
            MPI_recv \(\boldsymbol{z}_{i: i+b}\) from source
        end
    end
    MPI_Bcast \(\boldsymbol{z}\) to all ranks
end function
```


### 3.2.5 stedc_deflate

Deflates eigenvalues where $z_{i}$ is (close to) zero (type 1), or where two eigenvalues are (nearly) the same (type 2), identified by applying a rotation to zero out $z_{i}$. Forms permutation to group columns of $\boldsymbol{Q}$ according to the column type:

- column type 1: non-deflated eigvecs from $\boldsymbol{Q}_{1}$
- column type 2: non-deflated eigvecs updated by deflation
- column type 3: non-deflated eigvecs from $\boldsymbol{Q}_{2}$
- column type 4: deflated eigvecs

Locally within each rank,

$$
\boldsymbol{Q}_{\text {local }}=\left[\begin{array}{cccc}
\boldsymbol{Q}_{1,1} & \boldsymbol{Q}_{1,2} & 0 & \boldsymbol{Q}_{1,4} \\
0 & \boldsymbol{Q}_{2,2} & \boldsymbol{Q}_{2,3} & \boldsymbol{Q}_{2,4}
\end{array}\right]
$$

Corresponds to ScaLAPACK laed2.

```
Algorithm 10 Deflation, part 1
    function stedc_deflate( \(\rho, \boldsymbol{D}, \boldsymbol{D}_{s}, \boldsymbol{z}, \boldsymbol{z}_{s}, \boldsymbol{Q}, \boldsymbol{Q}_{t}\) )
    input: \(\rho\) that tore subproblems,
            eigvals \(\boldsymbol{D}_{1}\) and \(\boldsymbol{D}_{2}\) in vector \(\boldsymbol{D}\),
            \(\boldsymbol{z}_{1}\) and \(\boldsymbol{z}_{2}\) in \(\boldsymbol{z}\),
            eigvecs \(\boldsymbol{Q}_{1}\) and \(\boldsymbol{Q}_{2}\) of subproblems in \(\boldsymbol{Q}\).
    output: \(\boldsymbol{D}\) has \(n_{d}\) deflated eigvals,
            \(\boldsymbol{D}_{s}\) has \(n_{s}\) non-deflated eigvals for secular equation,
            \(\boldsymbol{z}_{s}\) of length \(n_{s}\) is updated \(\boldsymbol{z}\) vector for secular equation,
            \(\boldsymbol{Q}_{t}\) is \(n_{s}\) updated eigvecs permuted into \(\boldsymbol{Q}_{1,1: 2}, \boldsymbol{Q}_{2,2: 3}, \boldsymbol{Q}_{1: 2,4}\),
                another permutation?
        // LAPACK secular equation solver (laed4) requires \(\rho>0\)
        if \(\rho<0\) then
            \(\rho=-\rho\)
            \(\boldsymbol{z}_{2}=-\boldsymbol{z}_{2}\)
    end
    \(/ / \boldsymbol{z}_{1}\) and \(\boldsymbol{z}_{2}\) are normalized; re-normalize so \(\|\boldsymbol{z}\|_{2}=1\)
    \(\rho=2 \rho\)
    \(\boldsymbol{z}=\boldsymbol{z} / \sqrt{2}\)
    compute permutation \(P_{s}\) to sort \(\boldsymbol{D}\)
    // note \(\operatorname{lamch}(\) "e") is unit roundoff \(u=\epsilon / 2\)
    tol \(=4 \epsilon \max \left(\|\boldsymbol{D}\|_{\max },\|\boldsymbol{z}\|_{\max }\right)\)
    if \(\rho\|\boldsymbol{z}\|_{\text {max }}<\) tol then
        \(n_{s}=0\) return
    end
```

```
Algorithm 11 Deflation, part 2
    // Deflate eigvals
    // A candidate eigval is non-negligible (not type 1), but may have type 2 deflation.
    \(/ / j_{s 1}\) is candidate eigval; initially none ( -1 ). \(s\) indicates sorted permutation.
    \(/ / j_{s 2}\) is current eigval under consideration.
    \(j_{s 1}=-1\)
    for \(j=0, \ldots, n-1\)
        \(j_{s 2}=P_{s}[j]\)
        if \(\left|\rho z_{j_{s 2}}\right|<\) tol then
            store \(j_{s 2}\) as deflated eigval (type 1)
        else if not first candidate eigval (i.e., \(j_{s 1} \geq 0\) ) then
            generate Givens rotation \(\boldsymbol{G}\) to zero first entry of \(\left[\begin{array}{l}z_{j_{1}} \\ z_{j_{s} 2}\end{array}\right]\)
            compute \(D_{j_{s 1}, j_{s 2}}\) off-diagonal from applying \(\boldsymbol{G}\left[\begin{array}{cc}D_{j_{s 1}} & 0 \\ 0 & D_{j_{22}}\end{array}\right] \boldsymbol{G}^{T}\)
            if \(D_{j_{s 1}, j_{s 2}}<\) tol then
            update \(\left[\begin{array}{l}z_{j_{s} 1} \\ z_{j_{s} 2}\end{array}\right]=\boldsymbol{G}\left[\begin{array}{l}z_{j_{s} 1} \\ z_{j_{s}} 2\end{array}\right]=\left[\begin{array}{l}0 \\ \tau\end{array}\right]\)
                    update columns \(\boldsymbol{Q}_{j_{s 1}, j_{s 2}}=\boldsymbol{Q}_{j_{s 1}, j_{s 2}} \boldsymbol{G} \quad / /\) involves MPI for remote columns
                    \(\operatorname{update}\left[\begin{array}{ll}D_{j_{s 1}} & \\ & D_{j_{s 2}}\end{array}\right]=\boldsymbol{G}\left[\begin{array}{ll}D_{j_{s 1}} & \\ & D_{j_{s 2}}\end{array}\right] \boldsymbol{G}^{T} \quad / /\) off-diag \(D_{j_{s 1}, j_{s 2}}\) is negligible
                    store \(j_{s 1}\) as deflated eigval (type 2)
            else
                store \(j_{s 1}\) as non-deflated eigval
            end
            \(j_{s 1}=j_{s 2} \quad / / j_{s 2}\) becomes next candidate eigval
        else
            \(j_{s 1}=j_{s 2} \quad / / j_{s 2}\) becomes first candidate eigval
        end
    end
    store \(j_{s 1}\) as non-deflated eigval
```

```
Algorithm 12 Deflation, part 3 (todo)
    // find permutation to order types 1, 2, 3 together locally.
    // find global permutation.
    // find indices of \(Q_{1,1: 2}, Q_{2,2: 3}\), and \(U_{1: 3}\).
    end function stedc_deflate
```


### 3.2.6 stedc_secular

Solves secular equation and computes eigenvectors via Löwner theorem.
Corresponds to ScaLAPACK laed3.

```
function stedc_secular \(\left(\boldsymbol{D}, \boldsymbol{z}, \boldsymbol{P}_{u}, \boldsymbol{\Lambda}, \boldsymbol{U}\right)\)
input: deflation-adjusted \(\boldsymbol{D}, \boldsymbol{z}\), global permutation \(\boldsymbol{P}_{u}\)
output: eigvals \(\boldsymbol{\Lambda}\) of merged problem, eigvecs \(\boldsymbol{U}\) of merged problem (before multiplying by
Q)
    // Compute \(\boldsymbol{\Lambda}\) and modified \(\tilde{\boldsymbol{z}}\).
    \(\tilde{z}=1\)
    parallel for \(j=0, \ldots, n_{s}-1\); split over MPI ranks
        lapack::laed4 \(\left(\boldsymbol{D}, \boldsymbol{z}, \lambda_{j}, \boldsymbol{\delta}_{j}\right)\) solves secular equation for \(\lambda_{j}\) and \(\boldsymbol{\delta}_{j}\) vector
        for \(i=0, \ldots, n_{s}-1\) (todo: OpenMP parallel?)
            if \(i==j\) then
                \(\tilde{z}_{i} *=\delta_{i j}\)
            else
                \(\tilde{z}_{i} *=\frac{\delta_{i j}}{D_{i}-D_{j}}\)
            end
        end
    end
    MPI_Allreduce \(\tilde{\boldsymbol{z}}\)
    fix sign \(\tilde{z}_{j}\) to match \(\operatorname{sign} z_{j}\), for \(j=0, \ldots, n-1\)
    MPI_Allgather \(\boldsymbol{\Lambda}\)
    permute \(\boldsymbol{\Lambda}=\boldsymbol{P}_{u} \boldsymbol{\Lambda}\)
    // Compute \(\boldsymbol{U}\) via Löwner theorem.
    // All processes within a process column do this computation redundantly.
    // We could avoid that by communicating \(\boldsymbol{u}_{j}\).
    parallel for \(j=0, \ldots, n_{s}-1\); split over MPI 2DBC process columns
            Re-compute secular equation (laed4) to get \(\boldsymbol{\delta}_{j}\) vector
            \(\boldsymbol{u}_{j}=\tilde{\boldsymbol{z}} \oslash \boldsymbol{\delta}_{j}\) element-wise
            \(\boldsymbol{u}_{j}=\frac{\boldsymbol{u}_{j}}{\left\|\boldsymbol{u}_{j}\right\|}\)
            Store local part of \(\boldsymbol{u}_{j}\)
    end
end function
```


### 3.2.7 stedc_sort

Sorts eigenvalues $\boldsymbol{D}$ and applies same permutation to eigenvectors $\boldsymbol{Q}$. Corresponds to ScaLAPACK lasrt.

```
function stedc_sort( \(\boldsymbol{D}, \boldsymbol{Q}, \boldsymbol{Q}_{\text {out }}\) )
input: eigvals \(\boldsymbol{D}\), eigvecs \(\boldsymbol{Q}\)
output: sorted eigvals \(\boldsymbol{D}\), permuted eigvecs \(\boldsymbol{Q}_{\text {out }}\)
    compute permutation \(\boldsymbol{P}_{s}\) to sort eigvals.
    compute inverse permutation \(\boldsymbol{P}_{s}^{-1}\)
    for each block-col \(j=0, \ldots, n_{t}-1\)
        // todo: these are bad descriptions
        fill pcols[ jj ] with destination process of column \(\boldsymbol{P}_{s}^{-1}(j+j j)\) for \(j j=0, \ldots, j b-1\)
        fill mine[ jj ] with \(\boldsymbol{P}_{s}^{-1}(\mathrm{jj})\) where pcols \([\mathrm{jj}]==\) mycol
        fill pents \([\mathrm{p}]=\) length ( where ( \(\operatorname{pcnts}[\mathrm{jj}]==\mathrm{p})\) ) for \(p=0, \ldots\), npcol
        fill poffset \(=\) prefix_sum ( pent \()\)
        if block-col \(j\) is local then
            for \(j j=0, \ldots, j_{b}-1\)
                \(j_{g}=j+j j\)
                \(k_{g}=\boldsymbol{P}_{s}^{-1}(j+j j)\)
                if local then
                            local copy \(\boldsymbol{Q}\left(:, j_{g}\right) \rightarrow \boldsymbol{Q}_{\text {out }}\left(:, k_{g}\right)\)
                else
                    // Pack into workspace
                    copy \(Q\left(:, j_{g}\right) \rightarrow \operatorname{work}\left(:, \operatorname{poffset}\left(p_{k}\right)\right)\)
                    \(\operatorname{poffset}\left(p_{k}\right)+=1\)
                end
            end
            for \(p=0, \ldots\), npcol -1
                if \(p \neq m e\) and \(p c n t[p]>0\) then
                    MPI_Send pcnt[p] columns at \(\operatorname{work}(:, \operatorname{poffset}(\mathrm{p}))\) to \(\operatorname{rank}(\) myrow, p\()\)
                end
            end
        else
            MPI_Recv pcnt[ mycol ] columns at work from rank( myrow, pj )
            for \(j j=0, \ldots\), length ( mine )
            \(k_{g}=\operatorname{mine}(j j)\)
            copy \(\operatorname{work}(:, \mathrm{jj})\) to \(\boldsymbol{Q}_{\text {out }}\left(:, k_{g}\right)\)
            end
        end
    end
end function
```


## CHAPTER 4

## Optimization

### 4.1 Hermitian to Hermitian band reduction (he2hb)

### 4.2 Hermitian to Hermitian band reduction (he2hb)

In this section, we present the performance optimization of the Hermitian to Hermitian band reduction he2hb. In [20] we show the optimization technique to extend the parallelization of the different steps in he2hb on CPU only. Where we introduce a new internal functions of the various operations of he2hb, such as internal: :he2hb_hemm<HostTask>, internal::he2hb_trmm[Target::HostTask](Target::HostTask), etc. As most of the operations in he2hb are expressed through Level 3 BLAS, there is obviously still room for improvement by doing GPU computations and further hiding the communication overhead by computations. Therefore, we provide a GPU implementation of the new internal functions introduced in he2hb. Figure 4.1 shows the performance of he2hb using CPU and 1 GPU, 2 GPUs on 1 node and 2 nodes. Using single GPU achieves up to $3.2 \times$ and $1.6 \times$ speedup compared to the host test on a 1 node and 2 nodes, respectively.

We generate traces using the Nvidia Nsight Systems viewer nsys to highlight the performance bottlenecks. The traces shows that the panel factorization is the most time consuming and it does not overlap with any of the subsequent computations and data transfer. Therefore, we add new omp tasks to overlap the panel factorization with data movements, as long as the data dependencies satisfied. Figure 4.2 shows the panel factorization overlap with allocating batch arrays and create CUDA streams, and sending the data to the GPU. Figure 4.3 studies the impact of this change on the he2hb performance using different number of GPUs, the new implementaion with enabling the data transfer during the panel factorization achieves up to $20 \%$


Figure 4.1: Performance of he2hb Using 1 and 2 nodes on Summit, $1 \times 1$ and $2 \times 2$ process grids.
improvement compared to the initial implementation.
Figure 4.4 shows the performance of he2hb in time and Gflops using 1 node on Summit.

### 4.3 Back-transformation (unmtr_hb2st)

The initial implementation of the second stage back-transformation (unmtr_hb2st) was sequential. For optimization of the unmtr_hb2st routine presented in Algorithm 6, CPU-only OpenMP parallelism is introduced. Then for further optimization, the gemm operations are moved to GPU. The performance comparison of these two implementations is presented in Figure 4.5. The device implementation of unmtr_hb2st achieves up to 6 x speedup over the CPU-only implementation.


Figure 4.2: Overlap the panel factorization with subsequent data movements.


Figure 4.3: The performance impact by overlapping the panel factorization with subsequent data movements.


Figure 4.4: Performance results of Hermitian to Hermitian band reduction, using 1 node, $1 \times 1$ process grid. $\mathrm{nb}=128,320, \mathrm{i} \mathrm{b}=16,48$ for CPU, GPU tests, panel-threads $=10$


Figure 4.5: Performance of unmtr_hb2st on a host with two 20-core Intel Broadwell Xeon E5-2698 v4 CPUs and one NVIDIA V100 activated. $\mathrm{N}=16384$. As seen in the figure, the device implementation provides up to $6 \times$ speedup.

## CHAPTER 5

## Performance

### 5.1 Environment

### 5.1.1 Hardware

Performance numbers were collected using the Summit system ${ }^{12}$ at the Oak Ridge Leadership Computing Facility (OLCF). Summit is equipped with IBM POWER9 processors and NVIDIA V100 (Volta) GPUs. Each of Summit's nodes contains two POWER9 CPUs (with 22 cores each) and six V100 GPUs. Each node has 512 GB of DDR4 memory, and each GPU has 16 GB of HBM2 memory. NVLink 2.0 provides all-to-all $50 \mathrm{~GB} / \mathrm{s}$ connections for one CPU and three GPUs (i.e., one CPU is connected to three GPUs with $50 \mathrm{~GB} / \mathrm{s}$ bandwidth each, and each GPU is connected to the other two with $50 \mathrm{~GB} / \mathrm{s}$ bandwidth each). The two CPUs are connected with a $64 \mathrm{~GB} / \mathrm{s}$ X Bus. Each node has a Mellanox enhanced-data rate (EDR) InfiniBand network interface controller (NIC) that supports $25 \mathrm{~GB} /$ s of bi-directional traffic. Figure 5.1 shows the hardware architecture of a Summit node.

### 5.1.2 Software

The software environment used for the SVD experiments included:

- GNU Compiler Collection (GCC) 6.4.0,
- NVIDIA CUDA 10.1.105,
- IBM Engineering Scientific Subroutine Library (ESSL) 6.1.0,

[^3]

Figure 5.1: Summit node architecture.

- IBM Spectrum MPI 10.3.0.0,
- Netlib LAPACK 3.8.0, and
- Netlib ScaLAPACK 2.0.2.

For the generalized Hermitian eigenvalues, these were updated to:

- GNU Compiler Collection (GCC) 8.1.1,
- NVIDIA CUDA 10.1.243,
- IBM Engineering Scientific Subroutine Library (ESSL) 6.1.0,
- IBM Spectrum MPI 10.3.1.2,
- Netlib LAPACK 3.8.0, and
- Netlib ScaLAPACK 2.0.2.


### 5.2 Results

Here, we present the results of our preliminary performance experiment with the singular value solve. Figure 5.2 shows the execution time of ScaLAPACK compared to SLATE with and without GPU acceleration. Two MPI ranks are mapped to one node of Summit, i.e., one rank is mapped to one CPU socket ( 22 cores) and three GPU devices. Only singular values are computed in all cases (no vectors).

For a matrix of size $32,768 \times 32,768$, ScaLAPACK took 925 seconds, while SLATE took 324 seconds using CPUs only and 233 seconds with GPU acceleration. That is, SLATE was almost

| $N$ | scalapack:p dgesvd:1nod e | slate:dgesvd :cpus:1node | slate:dgesvd:g pus:1node |
| :---: | :---: | :---: | :---: |
| 4096 | 3 | 4 | 4 |
| 8192 | 19 | 14 | 13 |
| 12288 | 62 | 34 | 29 |
| 16384 | 137 | 65 | 53 |
| 20480 | 248 | 107 | 85 |
| 24576 | 415 | 164 | 126 |
| 28672 | 629 | 237 | 174 |
| 32768 | 925 | 324 | 233 |
| 36864 |  | 431 | 302 |
| 40960 |  | 558 | 380 |
| 49152 |  | 875 | 569 |
| 57344 |  | 1291 | 805 |
| 65536 |  | 1816 | 1090 |
| 73728 |  | 2478 | 1433 |
| 81920 |  | 3252 | 1817 |



Figure 5.2: SVD performance comparison.

3 times faster without acceleration and almost 4 times faster with acceleration. Since the performance gap increases with the problem size, we expect SLATE to be an order of magnitude faster for matrices in the $O(100 K)$ range without acceleration, and further benefit $3 \times$ to $4 \times$ from acceleration.

For the generalized Hermitian definite eigenvalue problem, Figure 5.3 shows the performance for conversion from the generalized form to standard form (hegst). On the CPU host, SLATE closely matches ScaLAPACK's performance, while when using GPUs, SLATE gets a modest acceleration. We will continue to investigate ways to optimize the performance.

Figure 5.4 presents the performance breakdown of eigensolver routines in ScaLAPACK and SLATE. Double precision is used for all routines. ScaLAPACK's pdsyev routine with QR iteration and pdsyevd routine with the D\&C algorithm are used. Both pdsyev and pdsyevd implement 1-stage reduction. SLATE's eigensolver is based on 2-stage reduction and it has the recently-implemented tridiagonal eigensolver with the $\mathrm{D} \& \mathrm{C}$ algorithm. The results belong to only one node of Summit. For ScaLAPACK, 6-by-6 process grid consisting of 36 MPI ranks is used, whereas for SLATE, 2-by-2 process grid having 9 cores and 1 GPU per rank is used. Consequently, both libraries are run on 36 cores for the sake of fair comparison. Both libraries are tuned for various block sizes. The best block sizes are found to be 96 for ScaLAPACK and 224 for SLATE. The default inner blocking size, which is 16 , and 6 panel threads for the QR algorithm are used for SLATE.

As seen in Figure 5.4, the solve part with D\&C algorithm in ScaLAPACK is significantly faster than the solve part with the QR iteration. The red bars in the figure represent times spent for the eigensolver. The most time consuming eigensolver is the one based on QR iteration in ScaLAPACK. When ScaLAPACK and SLATE with the D\&C algorithm are compared, the first stage hermitian to band and the back transformation times are shorter in SLATE since SLATE utilizes GPUs for these computations. In the overall comparison, SLATE is slightly faster than ScaLAPACK. The second stage band to tridiagonal reduction and the recently-implemented D\&C tridiagonal eigenvalue solver in SLATE need further optimization to better utilize the available system resources including GPUs.


Figure 5.3: Generalized to standard eigenvalue performance comparison.


Figure 5.4: Profile of eigenvalue solver implementations showing each phase for $N=12288$. One node of Summit is used.

## Bibliography

[1] Paul Sweazey and Alan Jay Smith. A class of compatible cache consistency protocols and their support by the IEEE futurebus. ACM SIGARCH Computer Architecture News, 14(2): 414-423, 1986.
[2] Daniel J. Sorin, Mark D. Hill, and David A. Wood. A primer on memory consistency and cache coherence. Synthesis Lectures on Computer Architecture, 6(3):1-212, 2011.
[3] Carl Eckart and Gale Young. The approximation of one matrix by another of lower rank. Psychometrika, 1(3):211-218, 1936.
[4] Magnus R Hestenes. Inversion of matrices by biorthogonalization and related results. Journal of the Society for Industrial and Applied Mathematics, 6(1):51-90, 1958.
[5] Gene Golub and William Kahan. Calculating the singular values and pseudo-inverse of a matrix. Journal of the Society for Industrial and Applied Mathematics, Series B: Numerical Analysis, 2(2):205-224, 1965.
[6] Gene H Golub and Christian Reinsch. Singular value decomposition and least squares solutions. In Linear Algebra, pages 134-151. Springer, 1971.
[7] John GF Francis. The QR transformation: a unitary analogue to the LR transforma-tion-part 1. The Computer Journal, 4(3):265-271, 1961.
[8] Vera N Kublanovskaya. On some algorithms for the solution of the complete eigenvalue problem. USSR Computational Mathematics and Mathematical Physics, 1(3):637-657, 1962.
[9] Gary W Howell, James W Demmel, Charles T Fulton, Sven Hammarling, and Karen Marmol. Cache efficient bidiagonalization using BLAS 2.5 operators. ACM Transactions on Mathematical Software (TOMS), 34(3):14, 2008. doi: 10.1145/1356052.1356055.
[10] Azzam Haidar, Jakub Kurzak, and Piotr Luszczek. An improved parallel singular value algorithm and its implementation for multicore hardware. In Proceedings of the International

Conference on High Performance Computing, Networking, Storage and Analysis (SC'13), page 90. ACM, 2013. doi: $10.1145 / 2503210.2503292$.
[11] Azzam Haidar, Hatem Ltaief, and Jack Dongarra. Parallel reduction to condensed forms for symmetric eigenvalue problems using aggregated fine-grained and memory-aware kernels. In Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis (SC'11), pages 8:1-8:11. ACM, 2011. doi: 10.1145/2063384.2063394.
[12] Azzam Haidar, Stanimire Tomov, Jack Dongarra, Raffaele Solca, and Thomas Schulthess. A novel hybrid CPU-GPU generalized eigensolver for electronic structure calculations based on fine-grained memory aware tasks. International Journal of High Performance Computing Applications, 28(2):196-209, 2014. doi: 10.1177/1094342013502097.
[13] Robert Schreiber and Charles Van Loan. A storage-efficient WY representation for products of Householder transformations. SIAM Journal on Scientific and Statistical Computing, 10 (1):53-57, 1989. doi: 10.1137/0910005.
[14] Jan JM Cuppen. A divide and conquer method for the symmetric tridiagonal eigenproblem. Numerische Mathematik, 36(2):177-195, 1980.
[15] Françoise Tisseur and Jack Dongarra. A parallel divide and conquer algorithm for the symmetric eigenvalue problem on distributed memory architectures. SIAM Journal on Scientific Computing, 20(6):2223-2236, 1999.
[16] Jack J Dongarra and Danny C Sorensen. A fully parallel algorithm for the symmetric eigenvalue problem. SIAM Journal on Scientific and Statistical Computing, 8(2):s139-s154, 1987.
[17] Ren-Cang Li. Solving secular equations stably and efficiently. Technical Report LAPACK working note (LAWN) 89, University of California, Berkeley, April 1993.
[18] Jeffery D Rutter. A serial implementation of Cuppen's divide and conquer algorithm for the symmetric eigenvalue problem. Technical Report UCB/CSD 94/799, University of California, Berkeley, February 1994.
[19] Ming Gu and Stanley C Eisenstat. A divide-and-conquer algorithm for the symmetric tridiagonal eigenproblem. SIAM Journal on Matrix Analysis and Applications, 16(1): 172-191, 1995.
[20] Kadir Akbudak, Paul Bagwell, Sebastien Cayrols, Mark Gates, Dalal Sukkari, Asim YarKhan, and Jack Dongarra. SLATE performance improvements: QR and eigenvalues, SWAN no. 17. Technical Report ICL-UT-21-02, Innovative Computing Laboratory, University of Tennessee, 4 2021. URL https://www.icl.utk.edu/publications/swan-017. revision 04-2021.


[^0]:    ${ }^{1}$ http://icl.utk.edu/slate/
    ${ }^{2}$ https://www.exascaleproject.org

[^1]:    ${ }^{1}$ Variable names in SLATE code and ScaLAPACK code (in parenthesis), respectively.

[^2]:    ${ }^{2}$ Tisseur has upper limit $\prod_{\substack{j=1 \\ j \neq i}}^{i-1} ; G u$ and Eisenstat [19] correctly has $\prod_{\substack{j=1 \\ j \neq i}}^{n}$.

[^3]:    ${ }^{1}$ https://www.olcf.ornl.gov/summit/
    ${ }^{2}$ https://en.wikichip.org/wiki/supercomputers/olcf-4

