



Robust level-3 BLAS Inverse Iteration from the Hessenberg Matrix

ANGELIKA SCHWARZ, Umeå University, Sweden

Inverse iteration is known to be an effective method for computing eigenvectors corresponding to simple and well-separated eigenvalues. In the non-symmetric case, the solution of shifted Hessenberg systems is a central step. Existing inverse iteration solvers approach the solution of the shifted Hessenberg systems with either RQ or LU factorizations and, once factored, solve the corresponding systems. This approach has limited level-3 BLAS potential since distinct shifts have distinct factorizations. This paper rearranges the RQ approach such that data shared between distinct shifts can be exploited. Thereby the backward substitution with the triangular R factor can be expressed mostly with matrix–matrix multiplications (level-3 BLAS). The resulting algorithm computes eigenvectors in a tiled, overflow-free, and task-parallel fashion. The numerical experiments show that the new algorithm outperforms existing inverse iteration solvers for the computation of both real and complex eigenvectors.

CCS Concepts: • **Mathematics of computing** → **Solvers**; *Mathematical software performance*;

Additional Key Words and Phrases: Inverse iteration, shifted Hessenberg systems, overflow-free computation

ACM Reference format:

Angelika Schwarz. 2022. Robust level-3 BLAS Inverse Iteration from the Hessenberg Matrix. *ACM Trans. Math. Softw.* 48, 3, Article 25 (September 2022), 30 pages.
<https://doi.org/10.1145/3544789>

1 INTRODUCTION

Inverse iteration is an established method for computing eigenvectors. When an approximation λ to an eigenvalue of a matrix A is known, inverse iteration approximates an eigenvector by solving

$$(A - \lambda I)z^{(k)} = \rho^{(k)} z^{(k-1)} \quad k \geq 1.$$

Here, $z^{(0)}$ is a unit norm starting vector and $\rho^{(k)}$ is a scalar that normalizes the iterate $z^{(k)}$. A converging sequence of $z^{(k)}$ yields a right eigenvector $z \neq 0$ that is an exact eigenvector of a nearby matrix $A + E$ with $\|E\| = O(\epsilon \|A\|)$, where ϵ denotes the machine precision, implying a small residual $\|Az - \lambda z\| = O(\epsilon \|A\|)$.

This paper concerns the case when A is non-symmetric and real. A standard approach reduces A to an upper Hessenberg matrix $H = Q_0^T A Q_0$, where Q_0 is an orthogonal matrix. Then inverse

Computing resources have been provided by the Swedish National Infrastructure for Computing (SNIC) at High-Performance Computing Center North (HPC2N), Umeå, Sweden, under the grants SNIC 2019/3-311 and SNIC 2020/5-286. Author's address: A. Schwarz, Department of Computing Science, Umeå University, 901 87 Umeå, Sweden; email: angies@cs.umu.se.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM.

0098-3500/2022/09-ART25 \$15.00

<https://doi.org/10.1145/3544789>

iteration approximates an eigenvector $\mathbf{x} \neq \mathbf{0}$ of $\mathbf{H}$ by

$$(\mathbf{H} - \lambda \mathbf{I})\mathbf{x}^{(k)} = \sigma^{(k)} \mathbf{x}^{(k-1)} \quad k \geq 1. \quad (1)$$

The scalar $\sigma^{(k)}$ is again chosen so that $\mathbf{x}^{(k)}$ has unit norm. A computed eigenvector $\mathbf{x}$ is backtransformed into an eigenvector of $\mathbf{A}$ by $\mathbf{z} \leftarrow \mathbf{Q}_0 \mathbf{x}$. Varah [1968] showed that the starting vector $\mathbf{x}^{(0)}$ can be chosen such that a single iteration of (1) suffices. Then the task of computing an eigenvector $\mathbf{z}$ necessitates the efficient solution of shifted Hessenberg systems. This will be the topic of this paper.

Inverse iteration (1) hinges on the availability of good approximations λ to the true eigenvalues. These approximations can be computed through the QR algorithm *without* accumulating the orthogonal transformations [Golub and Van Loan 1996, Section 7.6.1]. This approach is indeed supported by the LAPACK 3.9.0 inverse iteration routine DHSEIN [Anderson et al. 1999]. If a subdiagonal entry of the Hessenberg matrix is zero, the eigenproblem decouples into smaller block-triangular problems. When the QR algorithm processes the blocks separately, DHSEIN exploits the known affiliation of the eigenvalue and the block it belongs to and performs inverse iteration only on the relevant block. This paper is therefore based on two assumptions. First, good approximations to the eigenvalues of $\mathbf{H}$ are available. In other words, each approximation λ is an exact eigenvalue of some matrix $\mathbf{H} + \mathbf{F}$, where $\|\mathbf{F}\| = O(\epsilon)$. Second, the Hessenberg matrix $\mathbf{H}$ is unreduced, i.e., all subdiagonal entries are non-zeros $h_{j+1,j} \neq 0$ for all $j = 1, \dots, n-1$.

This paper proposes a new algorithm for a more efficient computation of (1) if a batch of eigenvectors corresponding to distinct eigenvalues is sought. The new algorithm combines two ideas. The first idea is due to Henry [1994] who addresses the solution of a shifted Hessenberg system through an RQ factorization. In a single sweep, $\mathbf{H}$ is reduced column-by-column to an upper triangular $\mathbf{R}$ such that the newly computed column of $\mathbf{R}$ is immediately used in the backward substitution. The level-3 BLAS potential in this approach is limited since distinct shifts result in distinct RQ factorizations. This problem has been addressed by Bosner et al. [2018, Section 3] and gives the second idea. Level-3 BLAS can be introduced in spite of distinct shifts. A *partially* computed (tiled) RQ factorization preserves data that is shared between distinct shifts. The computation can be arranged such that most of the computation corresponds to matrix–matrix multiplications. Specifically, this paper contributes an inverse iteration algorithm based on solving shifted Hessenberg systems through the RQ approach with the following highlights.

- The RQ approach is revised such that most of the backward substitution corresponds to matrix–matrix multiplications (level-3 BLAS [Dongarra et al. 1990]) in spite of distinct eigenvalues.
- The new algorithm is tiled and can naturally be parallelized with tasks.
- In existing inverse iteration solvers, complex eigenvectors are computationally more expensive than real eigenvectors. The new algorithm supports the computation of real and complex eigenvectors alike, meaning that the computational cost per column is approximately the same.

The rest of this paper is organized as follows. Section 2 reviews the LU approach realized in DHSEIN and the RQ approach by Henry [1994] as well as the tiled RQ factorization by Bosner et al. [2018]. Inspired by these ideas, Section 3 presents the new algorithm for solving shifting Hessenberg system, which is the core of the inverse iteration routine developed in Section 4. Section 5 describes the numerical experiments and presents their results.

2 RELATED WORK

This section reviews approaches to the solution of shifted Hessenberg systems. First, the LU approach implemented in LAPACK and the RQ factorization advocated by Henry [1994] are discussed.

Then, the tiled RQ factorization by Bosner et al. [2018] delivers key ideas used in the new algorithm presented in Section 3.

2.1 LU Factorization

LAPACK includes the inverse iteration routine DLAEIN [Anderson et al. 1999] for computing a single eigenvector. DLAEIN addresses the solution of the first iteration of (1) through an LU factorization with partial pivoting

$$P(H - \lambda I) = LU, \quad Ly = Px^{(0)}, \quad Ux^{(1)} = \alpha y, \quad x \leftarrow \sigma^{(1)}x^{(1)}.$$

Here, U is upper triangular, L is lower unit triangular, and P is a permutation matrix. As the starting vector can be chosen such that a single iteration suffices [Peters and Wilkinson 1971; Varah 1968], DLAEIN executes only one half iteration solving $Ux^{(1)} = \alpha y$. The scalar $\alpha \in (0, 1]$ serves the avoidance of overflow. The vector y is set to a scaled vector of ones e , where the scaling depends on H . In other words, the starting vector is implicitly selected as $x^{(0)} = P^{-1}Le$. If the first initial vector y does not satisfy the convergence criterion in the first half iteration, the initial vector is exchanged rather than computing more iterations. Distinct shifts yield distinct LU factorizations of $H - \lambda I$. Hence, when several eigenvectors corresponding to distinct eigenvalues are sought, a different LU factorization is computed for each eigenvalue. To not overwrite H , the upper triangular factor U is computed in a workspace.

2.2 RQ and UL Factorization

Henry [1994; 1995] approaches the solution of $(H - \lambda I)x = b$ through an RQ factorization. By applying suitable Givens rotations from the right,

$$(H - \lambda I) \underbrace{G_n^T G_{n-1}^T \dots G_2^T}_{Q^T} \underbrace{G_2 \dots G_n}_{Q} x = b,$$

the shifted Hessenberg matrix is transformed into an upper triangular matrix $R = (H - \lambda I)Q^T$. The solution x is obtained by solving $Ry = b$ and backtransforming $x = Q^T y$. The application of the Givens rotations successively computes the columns of R from right to left. Specifically, the Givens rotation G_k^T is constructed to annihilate the subdiagonal entry $h(k, k-1)$ and transforms the columns $k-1$ and k . As soon as a column of R has been computed, it is immediately used in a column-oriented backward substitution and then discarded. Henry uses an auxiliary column to compute and store this single column of R . The full procedure is listed in Algorithm 1. The flop count of Algorithm 1 is $3.5n^2 + O(n)$ for a real shift. By virtue of merging the computation of the R factor and the backward substitution, the matrix H is accessed only once. Since H is left untouched, several shifted Hessenberg systems can be solved simultaneously and benefit from improved temporal locality of accessing columns of H .

When the shift in Algorithm 1 is complex, both the R and the Q factor become complex. The backward substitution with R then relies fully on complex arithmetic and multiplies complex vectors with complex scalars. To exploit that all entries but the diagonal of $H - \lambda I$ are real, Henry [1994, Section 4] employs the RQ approach only for real shifts. Complex shifts, by contrast, are addressed by an UL factorization. The UL approach can benefit from mixed real-complex arithmetic such that real vectors are multiplied with complex scalars. Depending on what Gauss transformation is used, the UL approach requires $3n^2 + O(n)$ or $3.5n^2 + O(n)$ flops.

2.3 Tiled RQ Factorization

Henry's RQ factorization has been generalized to better support the simultaneous solution of many Hessenberg systems. The origin of this generalization is in the computation of the complex

ALGORITHM 1: Solving $(H - \lambda I)\mathbf{x} = \mathbf{b}$ with an RQ decomposition (Henry [1994, Algorithm 2])

```

1  $\mathbf{v} \leftarrow \mathbf{h}(1 : n, n); v(n) \leftarrow v(n) - \lambda; \mathbf{x} \leftarrow \mathbf{b};$ 
2 for  $k \leftarrow n : -1 : 2$  do
3   Determine a Givens rotation  $G_k^T = \begin{bmatrix} c_k & -s_k \\ s_k & c_k \end{bmatrix}$  such that  $[h(k, k-1) \quad v(k)]G_k^T = [0 \quad \phi];$ 
   // Backward substitution
4    $x(k) = x(k)/\phi;$ 
5    $\tau_1 \leftarrow s_k x(k); \tau_2 \leftarrow c_k x(k);$ 
6    $\mathbf{x}(1 : k-2) \leftarrow \mathbf{x}(1 : k-2) - \tau_2 \mathbf{v}(1 : k-2) + \tau_1 \mathbf{h}(1 : k-2, k-1);$ 
7    $x(k-1) \leftarrow x(k-1) - \tau_2 v(k-1) + \tau_1 (h(k-1, k-1) - \lambda);$ 
   // Compute column of the triangular factor
8    $\mathbf{v}(1 : k-2) \leftarrow c_k \mathbf{h}(1 : k-2, k-1) + s_k \mathbf{v}(1 : k-2);$ 
9    $v(k-1) \leftarrow c_k (h(k-1, k-1) - \lambda) + s_k v(k-1);$ 
10  $\tau_1 \leftarrow x(1)/v(1);$ 
    // Backtransform  $\mathbf{x} \leftarrow Q^T \mathbf{x}$ 
11 for  $k \leftarrow 2 : n$  do
12    $\tau_2 \leftarrow x(k);$ 
13    $x(k-1) \leftarrow c_k \tau_1 - s_k \tau_2;$ 
14    $\tau_1 \leftarrow c_k \tau_2 + s_k \tau_1;$ 
15 return  $\mathbf{x};$ 

```

frequency response function (transfer function) $G(\sigma_\ell) = C(A - \sigma_\ell I)^{-1}B$ of a linear time-invariant dynamical system $\dot{\mathbf{x}}(t) = A\mathbf{x}(t) + B\mathbf{u}(t)$, $\mathbf{y}(t) = C\mathbf{x}(t)$. Here, $A \in \mathbb{R}^{n \times n}$ is the system matrix, $B \in \mathbb{R}^{n \times m}$ is the input, $C \in \mathbb{R}^{p \times n}$ is the output, the shifts σ_ℓ are complex and $m, p \ll n$. Beat-tie, Drmač and Gugercin [2012] and Bosner, Bujanović and Drmač [2013; 2018] target the computation of $G(\sigma_\ell)$ through a reduction to the controller-Hessenberg form $\tilde{C}(\tilde{A} - \sigma_\ell I)^{-1}\tilde{B}$. Here, $\tilde{A} = Q^H A Q \in \mathbb{C}^{n \times n}$ is m-Hessenberg (Hessenberg with m subdiagonals), $\tilde{C} = C Q \in \mathbb{C}^{p \times m}$, $\tilde{B} = Q^H B \in \mathbb{C}^{n \times m}$ is upper trapezoidal and $Q \in \mathbb{C}^{n \times n}$ is unitary. A solution can be obtained by substituting the RQ factorization $\tilde{A} - \sigma_\ell I = R_\ell Q_\ell$ into the controller-Hessenberg form. Then $\tilde{C}(R_\ell Q_\ell)^{-1}\tilde{B} = (\tilde{C}Q_\ell^H)(R_\ell^{-1}\tilde{B})$ can be evaluated by solving a triangular system and two matrix-matrix multiplications.

Next we summarize the idea presented in Bosner et al. [2013, 2018]. The authors devise a tile column-oriented algorithm that computes RQ factorizations simultaneously for distinct shifts σ_ℓ and with a significant fraction of level-3 BLAS operations. Given a block size n_b , the RQ factorizations are computed in tile columns of width $m + n_b$, where m is the number of subdiagonals of the m-Hessenberg matrix. For a batch of shifts $\ell = 1, \dots, s$, the current tile column is split into a window on the diagonal and the remaining offdiagonal part. The window is n_b -by- $(m + n_b)$ and m-Hessenberg. It is reduced to triangular shape with m leading zero columns by applying unitary transformations from the right. The unitary transformations are *only* applied to the window and are accumulated into a matrix $Q_\ell \in \mathbb{C}^{(m+n_b) \times (m+n_b)}$. It remains to update the offdiagonal part with Q_ℓ^H . The offdiagonal part is partitioned $[D \quad E_\ell]$. The matrix D is shared across shifts. The matrix E_ℓ holds the shift-distinct columns that have been transformed with previous unitary transformations. The unitary matrix Q_ℓ^H is partitioned conformally as $Q_\ell^H = [U_\ell \quad V_\ell]^H$. The update $[D \quad E_\ell][U_\ell \quad V_\ell]^H$ is executed as a block matrix operation. First, $Z_\ell \leftarrow E_\ell V_\ell^H$ is executed as a batched matrix-matrix multiplication for $\ell = 1, \dots, s$. Then, the matrix-matrix multiplication $[Z_1 \quad \dots \quad Z_s] \leftarrow [Z_1 \quad \dots \quad Z_s] + D[U_1^H \quad \dots \quad U_s^H]$ with the shared D and copying back the tile columns complete the update.




Fig. 1. Reduction of the rightmost tile column.

The tiled RQ approach delivers two key concepts that are used for the algorithm developed in the next section. First, the algorithm is designed as a tile column-oriented algorithm such that the shift-specific orthogonal transformations are only applied to the small window on the diagonal. Second, working with a partially reduced Hessenberg matrix and rearranging the computation reveals data that is shared between distinct shift. The next section applies these two concepts to Henry's column-oriented backward substitution algorithm to introduce level-3 BLAS potential.

3 SOLUTION OF SHIFTED HESSENBERG SYSTEMS WITH LEVEL-3 BLAS

In this section we devise a tile column-oriented algorithm that solves $(H - \lambda_\ell I)x_\ell = b_\ell$ simultaneously for many distinct shifts λ_ℓ . For simplicity, we assume λ_ℓ to be real and defer the complex case to Section 4.3. Using ideas by Bosner et al. [2018], we adapt the RQ approach by Henry [1994] such that a large part of the data is shared. This way, the backward substitution phases can be merged for several shifts such that a large fraction of the computation corresponds to matrix-matrix multiplications.

3.1 Simultaneous Backward Substitution of a Batch of Shifts

The RQ approach requires a different sequence of Givens rotations for every shift λ_ℓ , $\ell = 1, \dots, m$. The Given rotation

$$G_{k,\ell}^T = I_{k-2} \oplus \begin{bmatrix} c(k, \ell) & -s(k, \ell) \\ s(k, \ell) & c(k, \ell) \end{bmatrix} \oplus I_{n-k} \in \mathbb{R}^{n \times n}, \quad (2)$$

where $\oplus$ denotes the direct sum, transforms the columns j and $j - 1$ of $H - \lambda_\ell I$. To be available in the backtransformation, the cosine and sine components of all Givens rotations are recorded in the matrices $C = [c(k, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(k, \ell)] \in \mathbb{R}^{n \times m}$.

The RQ approach transforms the Hessenberg matrix $(H - \lambda_\ell I)G_{n,\ell}^T \dots G_{2,\ell}^T = R_\ell$. We apply the Givens rotations in batches and thereby compute a triangular factor R_ℓ tile column by tile column starting from the right.

Rightmost tile column: The first batch of Givens rotations $G_{n,\ell}^T, \dots, G_{k,\ell}^T$ transforms the column range $N := k : n$ of $H - \lambda_\ell I$ into $R_\ell(:, N)$. Figure 1 illustrates the situation; matrix entries that are different for every shift are highlighted. The Givens rotation $G_{k,\ell}^T$ transforms the columns k and $k - 1$, where $k - 1$ is outside of the currently processed tile column. Hence, applying $G_{k,\ell}^T$ yields a cross-over column $\tilde{r}_\ell$.

The diagonal tile has been transformed into a triangular matrix $R_\ell(N, N)$. The system $R_\ell(N, N)x_\ell(N) = b_\ell(N)$ can be solved by backward substitution for every shift (level-2 BLAS). Let $\mathcal{K} := 1 : k - 1$ be the index range of the remaining row indices. In a tile column-oriented




Fig. 2. Reduction of a center tile column.

backward substitution algorithm, the computed solution is used to update $\mathbf{b}_\ell(\mathcal{K}) \leftarrow \mathbf{b}_\ell(\mathcal{K}) - \mathbf{R}_\ell(\mathcal{K}, \mathcal{N})\mathbf{x}_\ell(\mathcal{N})$. If $\mathbf{R}_\ell(\mathcal{K}, \mathcal{N})$ is *not* generated explicitly, the update reads

$$\mathbf{b}_\ell(\mathcal{K}) \leftarrow \mathbf{b}_\ell(\mathcal{K}) - \underbrace{\begin{bmatrix} \mathbf{h}(\mathcal{K}, k-1) - \lambda_\ell \mathbf{e}_{k-1} & \mathbf{H}(\mathcal{K}, \mathcal{N}) \end{bmatrix} \mathbf{G}_{n, \ell}^T(\mathcal{N}^+, \mathcal{N}^+) \cdots \mathbf{G}_{k, \ell}^T(\mathcal{N}^+, \mathcal{N}^+)}_{\begin{bmatrix} \tilde{\mathbf{r}}_\ell & \mathbf{R}_\ell(\mathcal{K}, \mathcal{N}) \end{bmatrix}} \begin{bmatrix} 0 \\ \mathbf{x}_\ell(\mathcal{N}) \end{bmatrix},$$

where $\mathcal{N}^+ := k-1 : n$ and $\mathbf{e}_{k-1}$ is the vector whose $(k-1)$ -th entry is one and all other entries are zero. If the Givens rotations are applied to the right-hand side, the update

$$\mathbf{b}_\ell(\mathcal{K}) \leftarrow \mathbf{b}_\ell(\mathcal{K}) - \underbrace{\begin{bmatrix} \mathbf{h}(\mathcal{K}, k-1) - \lambda_\ell \mathbf{e}_{k-1} & \mathbf{H}(\mathcal{K}, \mathcal{N}) \end{bmatrix} \mathbf{G}_{n, \ell}^T(\mathcal{N}^+, \mathcal{N}^+) \cdots \mathbf{G}_{k, \ell}^T(\mathcal{N}^+, \mathcal{N}^+)}_{\mathbf{z}_\ell} \begin{bmatrix} 0 \\ \mathbf{x}_\ell(\mathcal{N}) \end{bmatrix}$$

can be interpreted as a block matrix multiplication where the second block $\mathbf{H}(\mathcal{K}, \mathcal{N})$ is shared across distinct shifts. Hence, when several right-hand sides $\mathbf{z}_\ell$ are computed simultaneously, the multiplication with the second block $\mathbf{H}(\mathcal{K}, \mathcal{N})$ corresponds to a matrix-matrix multiplication (level-3 BLAS). Only the scalar-vector multiplication with the first block $\mathbf{h}(\mathcal{K}, k-1) - \lambda_\ell \mathbf{e}_{k-1}$ is shift-specific. Note that the computation of $\mathbf{z}_\ell$ is cheap since every Givens rotation only transforms two entries of $\mathbf{x}_\ell$.

Center tile columns: In contrast to the rightmost tile column, center tile columns have shift-dependent cross-over columns and diagonal entries. Figure 2 illustrates the situation. Let $\mathcal{J}_1 := 1 : j-1$, $\mathcal{J}_2 := j : k-1$ and $\mathcal{J}_2^+ := j-1 : k-1$.

The applications of the batch of Givens rotations $\mathbf{G}_{k-1, \ell}^T, \dots, \mathbf{G}_{j, \ell}^T$ yields

$$\left(\begin{bmatrix} \mathbf{h}(\mathcal{J}_1, j-1) & \mathbf{H}(\mathcal{J}_1, j : k-2) & \tilde{\mathbf{r}}_\ell(\mathcal{J}_1) \\ \mathbf{h}(\mathcal{J}_2, j-1) & \mathbf{H}(\mathcal{J}_2, j : k-2) & \tilde{\mathbf{r}}_\ell(\mathcal{J}_2) \end{bmatrix} - \lambda_\ell \begin{bmatrix} \mathbf{e}_{j-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} - \mathbf{e}_{k-1} \mathbf{e}_{k-1}^T \end{bmatrix} \right) \mathbf{G}_{k-1, \ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \cdots \mathbf{G}_{j, \ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) = \begin{bmatrix} \tilde{\mathbf{s}}_\ell & \mathbf{R}_\ell(\mathcal{J}_1, \mathcal{J}_2) \\ \mathbf{0} & \mathbf{R}_\ell(\mathcal{J}_2, \mathcal{J}_2) \end{bmatrix}. \quad (3)$$

The matrix $\mathbf{I} - \mathbf{e}_{k-1} \mathbf{e}_{k-1}^T$ differs from the identity matrix in the last diagonal entry, which is zero. The matrix $\begin{bmatrix} \mathbf{e}_{j-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{I} - \mathbf{e}_{k-1} \mathbf{e}_{k-1}^T \end{bmatrix}$ ensures that the shift affects all diagonal entries of $\mathbf{H}$ in the current tile column, but not $\tilde{\mathbf{r}}_\ell$. The application of the Givens rotations yields $\mathbf{R}_\ell(:, \mathcal{J}_2)$ and a new cross-over column $\tilde{\mathbf{s}}_\ell$.

The diagonal tile is transformed into a triangular matrix $R_\ell(\mathcal{J}_2, \mathcal{J}_2)$. The resulting triangular system $R_\ell(\mathcal{J}_2, \mathcal{J}_2)\mathbf{x}_\ell(\mathcal{J}_2) = \mathbf{b}_\ell(\mathcal{J}_2)$ can be solved with backward substitution for every shift. Its solution $\mathbf{x}_\ell(\mathcal{J}_2)$ is used in a linear update. Instead of generating $R_\ell(\mathcal{J}_1, \mathcal{J}_2)$ explicitly, the Givens rotations can again be applied to the right-hand side. Then the update

$$\mathbf{b}_\ell(\mathcal{J}_1) \leftarrow \mathbf{b}_\ell(\mathcal{J}_1) - \begin{bmatrix} \mathbf{h}(\mathcal{J}_1, j-1) - \lambda_\ell \mathbf{e}_{j-1} & \mathbf{H}(\mathcal{J}_1, j : k-2) & \tilde{\mathbf{r}}_\ell(\mathcal{J}_1) \end{bmatrix} \begin{pmatrix} G_{k-1,\ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \dots G_{j,\ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \begin{bmatrix} 0 \\ \mathbf{x}_\ell(\mathcal{J}_2) \end{bmatrix} \end{pmatrix} \quad (4)$$

reveals the shared data $\mathbf{H}(\mathcal{J}_1, j : k-2)$. If many right-hand sides are computed simultaneously, the computation can benefit from level-3 BLAS. Specifically, when interpreted as a block matrix multiplication, the update comprises a scalar-vector multiplication with $\mathbf{h}(\mathcal{J}_1, j-1) - \lambda_\ell \mathbf{e}_{j-1}$ (level-1 BLAS), a matrix-matrix multiplication with $\mathbf{H}(\mathcal{J}_1, j : k-2)$ (level-3 BLAS) and a shift-specific column update with $\tilde{\mathbf{r}}_\ell(\mathcal{J}_1)$ (level-1 BLAS).

The column range of the center tile is intentionally left open. It is certainly a possibility to have several center tile columns. In this case, a realization can iterate over center tile column from right to left until the leftmost tile column is reached.

Leftmost tile column: It remains to transform the top-left diagonal tile of $\mathbf{H}$ to triangular form. The application of the Givens rotations $G_{j-1,\ell}^T \dots G_{2,\ell}^T$ yields

$$\left(\begin{bmatrix} \mathbf{H}(\mathcal{J}_1, 1 : j-2) & \tilde{\mathbf{s}}_\ell \end{bmatrix} - \lambda_\ell (\mathbf{I} - \mathbf{e}_{j-1} \mathbf{e}_{j-1}^T) \right) G_{j-1,\ell}^T(\mathcal{J}_1, \mathcal{J}_1) \dots G_{2,\ell}^T(\mathcal{J}_1, \mathcal{J}_1) = \mathbf{R}_\ell(\mathcal{J}_1, \mathcal{J}_1).$$

The solution of the triangular system $\mathbf{R}_\ell(\mathcal{J}_1, \mathcal{J}_1)\mathbf{x}_\ell(\mathcal{J}_1) = \mathbf{b}_\ell(\mathcal{J}_1)$ with backward substitution finalizes the backward substitution phase.

Remark. The computational efficiency of (4) hinges on the availability of the cross-over column $\tilde{\mathbf{r}}_\ell$. If $\tilde{\mathbf{r}}_\ell$ is not stored, the right side of (4) is

$$\mathbf{b}_\ell(\mathcal{J}_1) - \underbrace{F \begin{bmatrix} \mathbf{H}(\mathcal{J}_1, j-1 : n) \\ \mathbf{H}(\mathcal{J}_2, j-1 : n) - \lambda_\ell \mathbf{I} \end{bmatrix} G_{n,\ell}^T(j-1 : n, j-1 : n) \dots G_{j,\ell}^T(j-1 : n, j-1 : n)}_{\begin{bmatrix} \tilde{\mathbf{s}}_\ell & \mathbf{R}_\ell(\mathcal{J}_1, j : n) \end{bmatrix}} \begin{bmatrix} 0 \\ \mathbf{x}_\ell(\mathcal{J}_2) \\ \mathbf{0}_{n-k+1} \end{bmatrix},$$

where $F = [\mathbf{I}_{j-1 \times j-1} \quad \mathbf{0}_{j-1 \times n-j+1}]$ extracts the top $j-1$ rows and ensures matching dimensions with $\mathbf{b}_\ell(\mathcal{J}_1)$. If the Givens rotations are applied to the right vector encompassing $\mathbf{x}_\ell(\mathcal{J}_2)$, fill-in is generated and the vector becomes dense. Hence, no matter if the Givens rotations are applied to $\mathbf{H}$ or to $\mathbf{x}_\ell$, the computation depends on the n -th column of $\mathbf{H} - \lambda_\ell \mathbf{I}$. Storing the cross-over column $\tilde{\mathbf{r}}$ effectively prunes the computation of the triangular factor. If the cross-over columns are stored, the reduction and the backward substitution phase can be separated.

3.2 A Tiled Algorithm for Solving Shifted Hessenberg Systems

This section uses the ideas presented in the previous section to derive a tiled algorithm for the simultaneous solution of shifted Hessenberg systems. The Hessenberg matrix $\mathbf{H} \in \mathbb{R}^{n \times n}$ is partitioned into a grid of $N \times N$ tiles. To simplify the presentation of the algorithms, we assume that all tiles are $b_r \times b_r$ and that $n = Nb_r$. The shifts are processed in batches of size b_c . We assume that the total number of shifts m is an integer multiple of b_c , i.e., $m = Mb_c$. The two partitionings induce a partitioning of the right-hand sides $\mathbf{B} \in \mathbb{R}^{n \times m}$ into a grid of $N \times M$ tiles, where every tile is $b_r \times b_c$.

The algorithm is split into three phases. The *reduction phase* computes the Givens rotations that transform the shifted Hessenberg matrix into a triangular factor. Furthermore, it records the

ALGORITHM 2: Reduction of a shifted Hessenberg tile to triangular form

Data: $H \in \mathbb{R}^{n \times n-1}$ and cross-over columns $\tilde{R} = [\tilde{r}_1, \dots, \tilde{r}_m] \in \mathbb{R}^{n \times m}$ such that $[H \quad \tilde{r}_\ell] \in \mathbb{R}^{n \times n}$ is upper Hessenberg, column $\mathbf{h}^{\text{left}} \in \mathbb{R}^n$, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$

Result: Components $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$ of $(n+1)$ -by- $(n+1)$ Givens rotations $\hat{G}_{k, \ell}^T = I_{k-1} \oplus \begin{bmatrix} c(k, \ell) & -s(k, \ell) \\ s(k, \ell) & c(k, \ell) \end{bmatrix} \oplus I_{n-k}$ whose application yields an upper triangular R_ℓ through $([\mathbf{h}^{\text{left}} \mid H \tilde{r}_\ell + \lambda_\ell \mathbf{e}_n] - [0 \mid \lambda_\ell I_n]) \hat{G}_{n, \ell}^T \dots \hat{G}_{1, \ell}^T = [0 \mid R_\ell]$. (5)

```

1 function REDUCEDIAG( $H, \mathbf{h}^{\text{left}}, \Lambda, \tilde{R}$ )
2   for  $\ell \leftarrow 1 : m$  do
3      $\mathbf{v} \leftarrow \tilde{r}_\ell$ ;
4     for  $k \leftarrow n : -1 : 2$  do
5       Determine and record a Givens rotation such that  $[h(k, k-1) \quad v(k)] \begin{bmatrix} c(k, \ell) & -s(k, \ell) \\ s(k, \ell) & c(k, \ell) \end{bmatrix} =$ 
6          $[0 \quad \star]$ ;
7        $\mathbf{v}(1 : k-2) \leftarrow c(k, \ell) \mathbf{h}(1 : k-2, k-1) + s(k, \ell) \mathbf{v}(1 : k-2)$ ;
8        $\mathbf{v}(k-1) \leftarrow c(k, \ell) (h(k-1, k-1) - \lambda_\ell) + s(k, \ell) \mathbf{v}(k-1)$ ;
9   if  $\mathbf{h}^{\text{left}} \neq \mathbf{0}$  then
10     Determine and record a Givens rotation such that  $[h^{\text{left}}(1) \quad v(1)] \begin{bmatrix} c(1, \ell) & -s(1, \ell) \\ s(1, \ell) & c(1, \ell) \end{bmatrix} = [0 \quad \star]$ ;
11   else
12      $c(1, \ell) \leftarrow 1; s(1, \ell) \leftarrow 0$ ; // Identity matrix (padding)
13   return  $C, S$ ;
```

cross-over columns that allow computing linear updates in the backward substitution as in (4). The *backward substitution phase* solves the triangular systems. The *backtransform phase* transforms the computed solutions into solutions to the original shifted Hessenberg systems.

Reduction phase. The reduction phase computes and records the Givens rotations and cross-over columns required for the backward substitution phase. The computation proceeds tile column by tile column. The transformation of a tile column is split into REDUCEDIAG and REDUCEOFFDIAG. REDUCEDIAG addresses the second block row of (3); REDUCEOFFDIAG concerns the first block row of (3).

The kernel REDUCEDIAG computes the Givens rotations necessary to transform a diagonal tile to triangular form. Algorithm 2 lists the details. It closely resembles Algorithm 1, but omits the backward substitution and adds the cross-over columns. Note that in Algorithm 2 (5) the matrix $[\mathbf{h}^{\text{left}} \mid H \tilde{r}_\ell]$ is n -by- $(n+1)$ and the Givens rotations are $(n+1)$ -by- $(n+1)$ due to the cross-over columns. The case distinction (lines 8–11) aims at supporting the computation of the top-left diagonal tile. For a top-left diagonal tile, the column $\mathbf{h}^{\text{left}}$ does not physically exist. The driver routine TILEDREDUCE discussed further down sets $\mathbf{h}^{\text{left}}$ to a zero column. In this case the computation of the Givens rotation (line 11) is meaningless. The first column of the storage matrices C and S of the Givens rotations is never read and can be viewed as padding. The flop count of Algorithm 2 is $1.5n^2 + O(n)$ per shift.

The kernel REDUCEOFFDIAG concerns the first block row of (3) and is realized in Algorithm 3. It applies the Givens rotations computed in REDUCEDIAG to an offdiagonal tile and records the left cross-over column. Algorithm 3 requires $3nk + O(k)$ flops per shift if the tile being processed is k -by- n .

ALGORITHM 3: Reduction of an offdiagonal Hessenberg tile

Data: $H \in \mathbb{R}^{k \times n}$, cross-over columns $\tilde{R}^{\text{right}} = [\tilde{r}_1^{\text{right}}, \dots, \tilde{r}_m^{\text{right}}] \in \mathbb{R}^{k \times m}$, components $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$ of Givens rotations

$\hat{G}_{k, \ell}^T = I_{k-1} \oplus \begin{bmatrix} c(k, \ell) & -s(k, \ell) \\ s(k, \ell) & c(k, \ell) \end{bmatrix} \oplus I_{n-k}$ as computed by REDUCEDIAG, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$.

Result: Cross-over columns $\tilde{R}^{\text{left}} = [\tilde{r}_1^{\text{left}}, \dots, \tilde{r}_m^{\text{left}}] \in \mathbb{R}^{k \times m}$ satisfying

$$\left[\begin{array}{c|cc} h(1:k, 1) - \lambda_\ell e_k & H(1:k, 2:n) & \tilde{r}_\ell^{\text{right}} \end{array} \right] \hat{G}_{n, \ell}^T \dots \hat{G}_{1, \ell}^T = \left[\begin{array}{c|c} \tilde{r}_\ell^{\text{left}} & \star \end{array} \right].$$

```

1 function REDUCEOFFDIAG( $H, \tilde{R}^{\text{right}}, \Lambda$ )
2   for  $\ell \leftarrow 1 : m$  do
3      $\mathbf{v} \leftarrow \tilde{r}_\ell^{\text{right}}$ ; // Start with previously computed cross-over column
4     for  $j \leftarrow n : -1 : 2$  do
5        $\mathbf{v}(1:k) \leftarrow c(j, \ell)h(1:k, j) + s(j, \ell)\mathbf{v}(1:k)$ ;
        // Compute cross-over column
6      $\tilde{r}_\ell^{\text{left}}(1:k-1) \leftarrow c(1, \ell)h(1:k-1, 1) + s(1, \ell)\mathbf{v}(1:k-1)$ ; // Shift-independent part
7      $\tilde{r}_\ell^{\text{left}}(k) \leftarrow c(1, \ell)(h(k, 1) - \lambda_\ell) + s(1, \ell)\mathbf{v}(k)$ ; // Last shift-dependent entry
8   return  $[\tilde{r}_1^{\text{left}}, \dots, \tilde{r}_m^{\text{left}}]$ ;

```

The partitioning of H into tiles and the kernels REDUCEDIAG and REDUCEOFFDIAG can be combined into a tiled algorithm for the reduction phase. The resulting algorithm TILEDREDUCE is listed in Algorithm 4. In a tiled algorithm, the row range covered by the first block row of (3) is split into smaller tile rows. Since only the tiles directly above the diagonal tiles are affected by the shift, a case distinction is necessary. Line 17 of Algorithm 4 handles the tiles directly above the diagonal tiles. Line 19, by contrast, handles the remaining far-from-diagonal tiles that are unaffected by the shift. The shift is disabled by setting the corresponding parameter to zero.

Backward substitution phase. The backward substitution phase solves $R_\ell \mathbf{x}_\ell = \mathbf{b}_\ell$ for every ℓ . The computation follows a standard pattern of tiled backward substitution. It iterates over tiles of the solution from bottom to top. For each iteration a new tile of the solution is computed with a small backward substitution. The small backward substitution is realized by Algorithm 5 SOLVE. Then the readily computed part of the solution is used in a tile update of above-lying tiles, realized by Algorithm 6 UPDATE. These two kernels are combined to the tiled backward substitution algorithm TILEDsolve listed in Algorithm 7. In the following, we present the details of the algorithms.

SOLVE realizes the small backward substitution. Since our algorithm splits the reduction phase and the backward substitution phase, the relevant part of the triangular factor has to be recomputed. Algorithm 5 gives the details. Analogously to Algorithm 2 REDUCEDIAG, the top-left diagonal tile of $H - \lambda_\ell I$ does not have a physical column $\mathbf{h}^{\text{left}}$ to its left. In that case, the driver routine TILEDsolve discussed below sets $\mathbf{h}^{\text{left}}$ to a zero column and thereby skips the computation of the left cross-over column (line 7). The flop count of Algorithm 5 is $3.5n^2 + O(n)$.

Next we focus on the UPDATE kernel. Recall the index ranges $\mathcal{J}_1 := 1 : j-1$, $\mathcal{J}_2 := j : k-1$ and $\mathcal{J}_2^+ := j-1 : k-1$. Algorithm 6 UPDATE realizes (4). The application of the Givens rotations in (4) to the right transforms

$$G_{k-1, \ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \dots G_{j, \ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \begin{bmatrix} 0 \\ \mathbf{x}_\ell(\mathcal{J}_2) \end{bmatrix} = \begin{bmatrix} \rho \\ \mathbf{z}_\ell \end{bmatrix},$$

ALGORITHM 4: Tiled reduction of a shifted Hessenberg matrix to triangular form and recording of the applied Givens rotations and relevant cross-over columns

Data: Hessenberg matrix $H \in \mathbb{R}^{n \times n}$, tile size b_r with $n = Nb_r$ and $N \geq 2$, shifts

$\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$, tile size b_c with $m = Mb_c$

Result: Givens rotations $G_{j,\ell}^T$ as in (2) and stored as $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$ such that $(H - \lambda_\ell I)G_{n,\ell}^T \dots G_{2,\ell}^T = R_\ell$, cross-over columns $\tilde{R}_\ell \in \mathbb{R}^{n \times N}$

```

1 function TILEDREDUCE( $H, \Lambda$ )
2   for  $\ell \leftarrow 1 : m$  do
3      $\tilde{r}_\ell(:, N) \leftarrow h(:, n); \tilde{r}_\ell(n, N) \leftarrow \tilde{r}_\ell(n, N) - \lambda_\ell;$ 
4   for  $\ell \leftarrow 1 : b_c : m$  do
5     Set the tile column index  $j_{\text{tile}} \leftarrow N - 1;$ 
6      $\mathcal{L} \leftarrow \ell : \ell + b_c - 1;$  // Index range of current batch of shifts
7     for  $j \leftarrow n - b_r + 1 : -b_r : b_r + 1$  do
8        $\mathcal{J} \leftarrow j : j + b_r - 1;$  // Index range of current tile column
9        $\mathcal{J}^- \leftarrow j - 1 : j + b_r - 2;$  // Tile column index range shifted by 1
10      Partition  $\begin{bmatrix} \underbrace{h^{\text{left}}}_{b_r \times 1} & \underbrace{H^{\text{diag}}}_{b_r \times (b_r - 1)} \end{bmatrix} \leftarrow H(\mathcal{J}, \mathcal{J}^-);$ 
11      Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{J}, j_{\text{tile}} + 1), \tilde{r}_{\ell+1}(\mathcal{J}, j_{\text{tile}} + 1), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{J}, j_{\text{tile}} + 1)];$ 
12       $C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}) \leftarrow \text{REDUCEDIAG}(H^{\text{diag}}, h^{\text{left}}, \Lambda(\mathcal{L}), \tilde{R}^{\text{right}});$ 
13      for  $i \leftarrow j - b_r : -b_r : 1$  do
14         $\mathcal{I} \leftarrow i : i + b_r - 1;$  // Index range of current tile row
15        Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{I}, j_{\text{tile}} + 1), \tilde{r}_{\ell+1}(\mathcal{I}, j_{\text{tile}} + 1), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{I}, j_{\text{tile}} + 1)];$ 
16        if  $i + b_r = j$  then
17           $[\tilde{r}_\ell(\mathcal{I}, j_{\text{tile}}), \tilde{r}_{\ell+1}(\mathcal{I}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_r-1}(\mathcal{I}, j_{\text{tile}})] \leftarrow \text{REDUCEOFFDIAG}(H(\mathcal{I}, \mathcal{J}^-),$ 
18             $\tilde{R}^{\text{right}}, \Lambda(\mathcal{L}));$ 
19        else
19           $[\tilde{r}_\ell(\mathcal{I}, j_{\text{tile}}), \tilde{r}_{\ell+1}(\mathcal{I}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_r-1}(\mathcal{I}, j_{\text{tile}})] \leftarrow \text{REDUCEOFFDIAG}(H(\mathcal{I}, \mathcal{J}^-),$ 
20             $\tilde{R}^{\text{right}}, 0);$ 
21       $j_{\text{tile}} \leftarrow j_{\text{tile}} - 1;$ 
22       $\mathcal{J} \leftarrow 1 : b_r;$  // Index range of top-left corner
23      Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{J}, 1), \tilde{r}_{\ell+1}(\mathcal{J}, 1), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{J}, 1)];$ 
24       $C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}) \leftarrow \text{REDUCEDIAG}(H(\mathcal{J}, 1 : b_r - 1), 0, \Lambda(\mathcal{L}), \tilde{R}^{\text{right}});$ 
25   return  $C, S, (\tilde{R}_1, \dots, \tilde{R}_m)$ 

```

which is computed in lines 3 and 9–13. The block matrix multiplication (4) with the left block yields the scalar-vector multiplication $b_\ell \leftarrow b_\ell - \rho(h(\mathcal{J}_1, j - 1) - \lambda_\ell e_{j-1})$ (lines 4–5). The multiplication with the center block $H(\mathcal{J}_1, j : k - 2)$ is executed in line 14. When many vectors z_ℓ are computed simultaneously, this operation corresponds to a matrix–matrix multiplication. Consequently, it can benefit from an efficient implementation of DGEMM available in an optimized BLAS library [Blackford et al. 2002; Dongarra et al. 1990]. The multiplication with the right block $\tilde{r}_\ell$ is realized with the vector-scalar multiplication (line 16). The total flop count of Algorithm 6 is $2mnk + O((m + k)n)$. When several right-hand sides are computed simultaneously, the matrix–matrix multiplication in line 14 dominates the computation.

ALGORITHM 5: Small backward substitution

Data: $H \in \mathbb{R}^{n \times n-1}$ and cross-over columns $\tilde{R} = [\tilde{r}_1, \dots, \tilde{r}_m] \in \mathbb{R}^{n \times m}$ such that $[H \quad \tilde{r}_\ell] \in \mathbb{R}^{n \times n}$ is upper Hessenberg, column h^{left} , $X \in \mathbb{R}^{n \times m}$, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$, $B \in \mathbb{R}^{n \times m}$, Givens rotations computed by `TILEDREDUCE` and stored as $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$.

Result: $X \in \mathbb{R}^{n \times m}$ satisfying $R_\ell x(:, \ell) = b(:, \ell)$ where R_ℓ is given by (5).

```

1 function SOLVE( $H, h^{\text{left}}, \Lambda, \tilde{R}, C, S, B$ )
2    $X \leftarrow B$ ;
3   for  $\ell \leftarrow 1 : m$  do
4      $v \leftarrow \tilde{r}_\ell$ ;
5     Compute  $x(2 : n, \ell)$  using lines 2–9 of Algorithm 1;
6     if  $h^{\text{left}} \neq 0$  then
7        $v(1) \leftarrow v(1)c(1, \ell) - h^{\text{left}}(1)s(1, \ell)$ ;
8        $x(1, \ell) \leftarrow x(1, \ell)/v(1)$ ;
9   return  $X$ ;
```

ALGORITHM 6: Linear tile update

Data: $H \in \mathbb{R}^{m \times k}$, $X \in \mathbb{R}^{k \times n}$, $B \in \mathbb{R}^{m \times n}$, cross-over columns $\tilde{R} = [\tilde{r}_1, \dots, \tilde{r}_n] \in \mathbb{R}^{m \times n}$, $\Lambda = [\lambda_1, \dots, \lambda_n]^T \in \mathbb{R}^n$, Givens rotations computed by `TILEDREDUCE` and stored as $C = [c(j, \ell)] \in \mathbb{R}^{k \times n}$ and $S = [s(j, \ell)] \in \mathbb{R}^{k \times n}$.

Result: (4) for $\ell = 1, \dots, n$

```

1 function UPDATE( $H, \tilde{R}, X, \Lambda, C, S, B$ )
2   for  $\ell \leftarrow 1 : n$  do
3      $\rho \leftarrow s(1, \ell)x(1, \ell)$ ;
4      $b(1 : m, \ell) \leftarrow b(1 : m, \ell) + \rho h(1 : m, 1)$ ; // Shift-specific DAXPY
5      $b(m, \ell) \leftarrow b(m, \ell) - \lambda_\ell \rho$ ;
6   Copy  $Z \leftarrow X$ ; // Use workspace to apply Givens rotation to the solution
7   for  $\ell \leftarrow 1 : n$  do
8      $z(1, \ell) \leftarrow c_\ell(1)z(1, \ell)$ ;
9     for  $j \leftarrow 2 : k$  do
10       $\tau_1 \leftarrow z(j-1, \ell)$ ;
11       $\tau_2 \leftarrow z(j, \ell)$ ;
12       $z(j-1, \ell) \leftarrow c(j, \ell)\tau_1 - s(j, \ell)\tau_2$ ;
13       $z(j, \ell) \leftarrow s(j, \ell)\tau_1 + c(j, \ell)\tau_2$ ;
14    $B \leftarrow B - H(1 : m, 2 : k)Z(1 : k-1, 1 : n)$ ; // DGEMM
15   for  $\ell \leftarrow 1 : n$  do
16      $b(1 : m, \ell) \leftarrow b(1 : m, \ell) - \tilde{r}_\ell z(k, \ell)$ ; // Shift-specific DAXPY
17   return  $B$ ;
```

The partitioning into tiles and the routines `SOLVE` and `UPDATE` lead to the tiled backward substitution algorithm `TILEDsolve` listed in Algorithm 7. Offdiagonal tile column updates are split into smaller tile row updates. Only the tiles directly above diagonal tiles are affected by the shifts (line 16). Far-from-diagonal tile updates, by contrast, are unaffected by the shifts (line 18). The computation of z_ℓ is repeated for each tile row. Based on numerical experiments, these additional flops are negligible compared to the gain from task parallelism, which will be introduced in the next section.

ALGORITHM 7: Tiled backward substitution and backtransform

Data: Hessenberg matrix $H \in \mathbb{R}^{n \times n}$, tile size b_r with $n = Nb_r$ and $N \geq 2$, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$, $B \in \mathbb{R}^{n \times m}$, tile size b_c with $m = Mb_c$, Givens rotations $G_{j,\ell}^T$ computed with Algorithm 4 and stored as $C \in \mathbb{R}^{n \times m}$ and $S \in \mathbb{R}^{n \times m}$, cross-over columns $\tilde{R}_\ell \in \mathbb{R}^{n \times N}$ for $\ell = 1, \dots, m$ computed with Algorithm 4.

Result: Solution $X \in \mathbb{R}^{n \times m}$ to $HX = B \text{diag}(\lambda_1, \dots, \lambda_m)$

```

1 function TILEDsolve( $H, \Lambda, C, S, (\tilde{R}_1, \dots, \tilde{R}_m), B$ )
2    $X \leftarrow B$ ;
3   for  $\ell \leftarrow 1 : b_c : m$  do
4      $j_{\text{tile}} \leftarrow N$ ; // Index of current tile column
5      $\mathcal{L} \leftarrow \ell : \ell + b_c - 1$ ; // Index range of current batch of shifts
6     for  $j \leftarrow n - b_r + 1 : -b_r : 1$  do
7        $\mathcal{J} \leftarrow j : j + b_r - 1$ ; // Index range of current tile column
8        $\mathcal{J}^- \leftarrow j - 1 : j + b_r - 2$ ; // Tile column index range shifted by 1
9       Partition  $\begin{bmatrix} \mathbf{h}^{\text{left}} & \mathbf{H}^{\text{diag}} \end{bmatrix} \leftarrow H(\mathcal{J}, \mathcal{J}^-)$ ;
10               $\underbrace{\hspace{1.5cm}}_{b_r \times 1} \quad \underbrace{\hspace{1.5cm}}_{b_r \times (b_r - 1)}$ 
11       Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{J}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{J}, j_{\text{tile}})]$ ;
12        $X(\mathcal{J}, \mathcal{L}) \leftarrow \text{SOLVE}(\mathbf{H}^{\text{diag}}, \mathbf{h}^{\text{left}}, \tilde{R}^{\text{right}}, \Lambda(\mathcal{L}), C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), X(\mathcal{J}, \mathcal{L}))$ ;
13       for  $i \leftarrow j - b_r : -b_r : 1$  do
14          $\mathcal{I} \leftarrow i : i + b_r - 1$ ; // Index range of current tile row
15         Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{I}, j_{\text{tile}}), \tilde{r}_{\ell+1}(\mathcal{I}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{I}, j_{\text{tile}})]$ ;
16         if  $i + b_r = j$  then
17            $X(\mathcal{I}, \mathcal{L}) \leftarrow$ 
18             UPDATE( $H(\mathcal{I}, \mathcal{J}^-), \tilde{R}^{\text{right}}, X(\mathcal{J}, \mathcal{L}), \Lambda(\mathcal{L}), C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), X(\mathcal{I}, \mathcal{L})$ );
19         else
20            $X(\mathcal{I}, \mathcal{L}) \leftarrow \text{UPDATE}(H(\mathcal{I}, \mathcal{J}^-), \tilde{R}^{\text{right}}, X(\mathcal{J}, \mathcal{L}), 0, C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), X(\mathcal{I}, \mathcal{L}))$ ;
21        $j_{\text{tile}} \leftarrow j_{\text{tile}} - 1$ ;
22      $\mathcal{J} \leftarrow 1 : b_r$ ;
23     Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{J}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{J}, j_{\text{tile}})]$ ;
24      $X(\mathcal{J}, \mathcal{L}) \leftarrow \text{SOLVE}(H(\mathcal{J}, 1 : b_r - 1), 0, \Lambda(\mathcal{L}), \tilde{R}^{\text{right}}, C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), X(\mathcal{J}, \mathcal{L}))$ ;
25      $X(:, \mathcal{L}) \leftarrow \text{BACKTRANSFORM}(C(:, \mathcal{L}), S(:, \mathcal{L}), X(:, \mathcal{L}))$ ;
26   return  $X$ ;

```

Backtransform phase. The solution $\mathbf{x}_\ell$ corresponding to a shift λ_ℓ is backtransformed

$$\mathbf{y}_\ell \leftarrow G_{n,\ell}^T \left(\dots \left(G_{2,\ell}^T \mathbf{x}_\ell \right) \right). \quad (5)$$

The components of the Givens rotations $G_{j,\ell}^T$ (2) have been recorded as $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$ in the reduction phase. A routine BACKTRANSFORM realizes (5) for a batch of vectors. A possible implementation is a column-by-column backtransform with the lines 11–14 of Algorithm 1. The flop count for the backtransform is $O(n)$ per shift. Line 23 of Algorithm 7 uses BACKTRANSFORM to backtransform a batch of vectors.

The successive execution of TILEDREDUCE and TILEDsolve yields a tiled solver for the simultaneous solution of shifted Hessenberg systems. This solver is listed in Algorithm 8 and can be viewed as an extension to Algorithm 1. The next section parallelizes Algorithm 8.




Fig. 3. DAG constructed for a single batch of shifts and a tile column count of $n/b_r = 4$ in Algorithm 4 TILEDREDUCE. The task arrangement is based on what tile of H is processed.

ALGORITHM 8: Solving $(H - \lambda_\ell I)x_\ell = b_\ell$ simultaneously in a tiled fashion

- 1 $C, S, (\tilde{R}_1, \dots, \tilde{R}_m) \leftarrow \text{TILEDREDUCE}(H, \Lambda);$
 - 2 $X \leftarrow \text{TILED SOLVE}(H, \Lambda, C, S, (\tilde{R}_1, \dots, \tilde{R}_m), B);$
-

3.3 Parallelization with Tasks

Libraries such as PLASMA [Dongarra et al. 2019] and FLAME [Van Zee et al. 2009] show that the design of dense linear algebra algorithms as tiled algorithms and the parallelization with tasks is an effective way of achieving high performance. A task-parallel algorithm expresses the computation as a **Directed Acyclic Graph (DAG)**. The nodes of the DAG correspond to tasks and represent tile operations. The edges of the DAG represent data dependencies between the tasks. A runtime system asynchronously executes the tasks while satisfying all data dependencies.

This section parallelizes Algorithm 8 with tasks. The implementation follows the usual approach and defines each tile operation as a task. Hence, every function call in Algorithm 4 TILEDREDUCE and Algorithm 7 TILED SOLVE is a task. What remains is the definition of the edges of the DAG. To reduce the amount of edges and the associated overhead, the parallel implementation maintains two separate DAGs. The first DAG represents the computation executed in TILEDREDUCE; the second DAG represents TILED SOLVE. All tasks related to TILEDREDUCE must have been completed before any task of TILED SOLVE is started. The implementation realizes this with a synchronization point. Next we define the edges of each DAG.

The data dependencies of Algorithm 4 TILEDREDUCE are as follows. A task REDUCEDIAG processing the tile (j, j) of H has outgoing dependencies to REDUCEOFFDIAG tasks processing above-lying tiles (i, j) in the same tile column of H , $i = 1, \dots, j - 1$. A REDUCEOFFDIAG task on (i, j) has an outgoing dependence to the task processing the left-lying tile $(i, j - 1)$. This is either a REDUCEDIAG tasks $(j - 1 = i)$ or another REDUCEOFFDIAG task $(j - 1 > i)$. Figure 3 illustrates the task graph for one batch of shifts, i.e., one iteration of line 4 in Algorithm 4. Batches of shifts are independent of each other.

Algorithm 7 TILED SOLVE has three task types whose dependencies are as follows. A SOLVE task on (j, j) has outgoing dependencies to UPDATE tasks (i, j) , $i = 1, \dots, j - 1$. Before a SOLVE on the tile (j, j) can be executed, all updates (i, j) , $j > i$ must have been completed. Once the backward substitution is completed, the BACKTRANSFORM can be executed. Figure 4 shows the DAG for one iteration of line 3 in Algorithm 7. Analogously to the task-parallel execution of TILEDREDUCE, batches of shifts do not have any dependencies.

The implementation adds one simplification to the DAG shown in Figure 4. The computation of the final SOLVE task and the BACKTRANSFORM are merged. This step aims at reducing scheduling overhead.




Fig. 4. DAG constructed for a single batch of shifts and a tile column count of $n/b_T = 4$ in Algorithm 7 TILED SOLVE. The task arrangement is based on what tile of H is processed.

3.4 Further Improvements

The presentation of the algorithms has focused on the algorithmic aspect, namely the introduction of level-3 BLAS operations to the backward substitution phase and the phrasing as tiled algorithms that can be parallelized with tasks. A natural step is to incorporate the ideas proposed for the original RQ approach. This includes improvements of the cache efficiency and the aggregation of a small number of orthogonal transformations for a joint application [Henry 1994, p. 8; Beattie et al. 2012, Sec. 5.2.1] to lower the flop count. The impact of these improvements has been investigated in Schwarz [2020]. The results of the numerical experiments in this work are solely based on the algorithmic improvements presented in this work.

4 ROBUST COMPUTATION OF EIGENVECTORS BY INVERSE ITERATION

This section extends Algorithm 8 for solving shifted Hessenberg systems to the computation of eigenvectors. The algorithm relies on the ability to solve triangular systems. The solution of triangular systems is known to be prone to overflow. This is particularly true if shifted Hessenberg systems are solved as part of an inverse iteration algorithm. In that case $H - \lambda I$ can be expected to be ill-conditioned if λ is close to a true eigenvalue of H . By adding overflow protection, Section 4.1 renders Algorithm 8 robust while preserving the tiled structure. This yields the robust, tiled shifted Hessenberg system solver DHSRQ3. Section 4.2 concerns the convergence criterion and the choice of the starting vector. Together with the robust, tiled shifted Hessenberg solver, this results in a robust tiled inverse iteration algorithm DHSRQ3IN. Section 4.3 discusses the modifications necessary to support complex shifts.

4.1 Ill-conditioned Systems and Overflow Protection

Inverse iteration assumes that a good approximation λ to a true eigenvalue of H is available. In that case λ is an exact eigenvalue of a nearby matrix $H + E$ where $\|E\| = O(\epsilon)$ is of the order of the machine precision. The matrix $H + E - \lambda I$ is exactly singular. Hence, $H - \lambda I$ is close-to-singular and can be expected to have a high condition number. A detailed explanation why the solution of a very ill-conditioned system does not produce completely erroneous solutions is given by Ipsen [1997, Section 6.3], and Parlett [1998, Section 4.3]. In fact, inverse iterations turns the ill-conditioning into an advantage because $\|(H - \lambda I)^{-1}\| \geq 1/\|E\| = 1/\epsilon$ shows that there is at least one vector whose norm grows in the order of $1/\epsilon$. Section 4.2 picks this up for the selection of a starting vector.

While the ill-conditioning is not a problem for the accuracy, the computed solution can be so large that the representational range is exceeded. A mechanism to avoid such a floating-point overflow is necessary. For the RQ approach, once $H - \lambda I$ has been factored, the backward substitution with the corresponding triangular factor can encounter overflow. To eliminate the possibility of floating-point overflow, implementations of inverse iteration [Peters and Wilkinson 1971, p. 435]

and shifted Hessenberg system solvers [Henry 1994, 1995] introduce a scaling factor $\gamma \in (0, 1]$ and solve the scaled triangular linear system $\mathbf{R}\mathbf{x} = \gamma\mathbf{b}$ for the scaled solution $\gamma^{-1}\mathbf{x}$. By virtue of γ , the current representation of the solution can be rescaled such that overflow is avoided. This renders the solution process robust. A robust solver for this scaled triangular system is, for example, DLATRS [Anderson 1991] available in LAPACK 3.9.0.

An extension to the solution of scaled triangular linear systems with n right-hand sides has been introduced by Kjelgaard Mikkelsen and Karlsson [2017a] and concerns $\mathbf{R}\mathbf{X} = \mathbf{B} \text{diag}(\gamma_1, \dots, \gamma_n)$. Each right-hand-side is associated with a scaling factor $\gamma_k \in (0, 1]$. If $\mathbf{B}$ is overwritten with the solution $\mathbf{X}$ in a standard (non-robust) tiled backward substitution algorithm, tile updates read $\mathbf{X}_i \leftarrow \mathbf{X}_i - \mathbf{R}_{ij}\mathbf{X}_j, i < j$. These tile updates are rendered robust through segment-wise scaling factors. There is one scaling factor per column per tile. Then robust tile updates are $\mathbf{X}_i \text{diag}(\delta_1^{-1}, \dots, \delta_n^{-1}) \leftarrow \mathbf{X}_i \text{diag}(\beta_1^{-1}, \dots, \beta_n^{-1}) - \mathbf{R}_{ij}(\mathbf{X}_j \text{diag}(\alpha_1^{-1}, \dots, \alpha_n^{-1}))$. Overflow is avoided by bounding the maximum possible growth in each update, computing suited scaling factors $\delta_1^{-1}, \dots, \delta_n^{-1}$ and, if necessary, rescaling the current representation of the solution prior to the matrix-matrix multiplication. This preserves the level-3 BLAS potential of a tile update. Since tile updates during the tiled backward substitution can require different scalings, the final tiled representation of the solution can be inconsistently scaled. A consistent scaling is computed by reducing the segment-wise scaling factors to the global scaling factor γ_k for each column. Then all column segments are rescaled with respect to γ_k . This approach supports task parallelism as demonstrated by Kjelgaard Mikkelsen et al. [2019] and is used to render the algorithms presented in this paper robust.

A robust tiled algorithm for solving shifted Hessenberg systems requires (a) a robust routine for the solution of small shifted Hessenberg systems, (b) a robust tile update, and (c) a robust back-transform. The combination of these three robust kernels yields a robust version of TILEDREDUCE. The robust routines are marked with the prefix R to easily distinguish between the non-robust and the robust version.

Robust solution of small shifted Hessenberg systems. A robust counterpart of Algorithm 5 SOLVE requires a small robust backward substitution routine solving $\mathbf{R}_\ell \mathbf{x}_\ell = \gamma_\ell \mathbf{b}_\ell, \gamma_\ell \in (0, 1]$. If the (small) $\mathbf{R}_\ell$ is generated explicitly, a possible realization is a call to the LAPACK routine DLATRS. This approach is realized by Algorithm 9 RSOLVE. DLATRS returns $\mathbf{x}_\ell$ and γ_ℓ representing $\gamma_\ell^{-1}\mathbf{x}_\ell$. The total scaling of the computed solution is computed in line 13 by multiplying the input scaling factor β_ℓ and γ_ℓ . Together, RSOLVE returns the scaled vector $(\gamma_\ell \beta_\ell)^{-1}\mathbf{x}_\ell$.

Robust tile update. The robust version of the tile update (4) is

$$\delta_\ell^{-1}\mathbf{b}_\ell(\mathcal{J}_1) \leftarrow \beta_\ell^{-1}\mathbf{b}_\ell(\mathcal{J}_1) - \begin{bmatrix} \mathbf{h}(\mathcal{J}_1, j-1) - \lambda_\ell \mathbf{e}_{j-1} & \mathbf{H}(\mathcal{J}_1, j : k-2) & \tilde{\mathbf{r}}_\ell(\mathcal{J}_1) \end{bmatrix} \left(\mathbf{G}_{k-1, \ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \dots \mathbf{G}_{j, \ell}^T(\mathcal{J}_2^+, \mathcal{J}_2^+) \begin{bmatrix} 0 \\ \alpha_\ell^{-1}\mathbf{x}_\ell(\mathcal{J}_2) \end{bmatrix} \right), \quad (6)$$

where $\delta_l \in (0, 1]$ is chosen such that $\mathbf{b}_\ell(\mathcal{J}_1)$ does not exceed the overflow threshold. There are many instantiations of δ_l and $\mathbf{b}_\ell(\mathcal{J}_1)$ that satisfy (4). Algorithm 10 computes one feasible instantiation. Following Kjelgaard Mikkelsen and Karlsson [2017b], the right-hand side tile $\mathbf{X}$ is associated with a vector of scaling factors $\boldsymbol{\alpha} = [\alpha_1, \dots, \alpha_n] \in (0, 1]^n$ and represents the column-wise scaled matrix $\mathbf{X} \text{diag}(\alpha_1^{-1}, \dots, \alpha_n^{-1})$. Similarly, the tile $\mathbf{B}$ is associated with $\boldsymbol{\beta} = [\beta_1, \dots, \beta_n] \in (0, 1]^n$ and represents $\mathbf{B} \text{diag}(\beta_1^{-1}, \dots, \beta_n^{-1})$. To compute a tile update robustly, $\mathbf{X}$ and $\mathbf{B}$ must be consistently scaled. The consistent scaling factor corresponds to the smaller of the two scaling factors (line 5). The remaining computation requires the overflow-free realization of three linear updates.

Each of the three linear updates is guarded by PROTECTUPDATE introduced by Kjelgaard Mikkelsen and Karlsson [2017b, Section 2.2]. PROTECTUPDATE receives $\|\mathbf{B}\|_\infty, \|\mathbf{H}\|_\infty, \|\mathbf{X}\|_\infty$ and

ALGORITHM 9: Small robust backward substitution

Data: $H \in \mathbb{R}^{n \times n-1}$ and cross-over columns $\tilde{R} = [\tilde{r}_1, \dots, \tilde{r}_m] \in \mathbb{R}^{n \times m}$ such that $\begin{bmatrix} H & \tilde{r}_\ell \end{bmatrix} \in \mathbb{R}^{n \times n}$ is upper Hessenberg, column h^{left} , $B \in \mathbb{R}^{n \times m}$ and scaling factors $\beta = [\beta_1, \dots, \beta_m] = (0, 1]^m$ representing $B \text{diag}(\beta_1, \dots, \beta_m)$, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$, Givens rotations computed by **TILEDREDUCE** and stored as $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$.

Result: $X \in \mathbb{R}^{n \times m}$ and $\alpha = [\alpha_1, \dots, \alpha_m] \in (0, 1]^m$ satisfying $R_\ell x(:, \ell) = \alpha_\ell b(:, \ell)$ where R_ℓ is given by (5).

```

1 function RSOLVE( $H, h^{\text{left}}, \tilde{R}, \Lambda, C, S, \beta, B$ )
2    $X \leftarrow B$ ;
3   for  $\ell \leftarrow 1 : m$  do
4     // Recompute the triangular factor
5      $R_\ell \leftarrow 0_{n \times n}$ ;
6      $r_\ell(:, n) \leftarrow \tilde{r}_\ell$ ;
7     for  $k \leftarrow n : -1 : 2$  do
8        $t_\ell \leftarrow h(1 : k, k - 1); t_\ell(k - 1) \leftarrow t(k - 1) - \lambda_\ell$ ;
9        $R_\ell(1 : k, k - 1 : k) \leftarrow \begin{bmatrix} t_\ell(1 : k) & r_\ell(1 : k, k) \end{bmatrix} \begin{bmatrix} c(k, \ell) \\ s(k, \ell) \end{bmatrix}$ ;
10    if  $h^{\text{left}} \neq 0$  then
11       $r_\ell(1, \ell) \leftarrow r_\ell(1, \ell)c(1, \ell) - h^{\text{left}}(1)s(1, \ell)$ ;
12       $\gamma_\ell \leftarrow 1$ ; // Initialize routine-local scaling factor
13      Solve robustly  $R_\ell x(:, \ell) = \gamma_\ell b(:, \ell)$ ; // DLATRS
14       $\alpha_\ell \leftarrow \gamma_\ell \beta_\ell$ ; // Update total scaling
15  return  $\alpha, X$ ;
```

computes a scaling factor $\xi \in (0, 1]$ such that the linear update $\xi B - H(\xi X)$ cannot overflow. Then line 6 computes the column-wise scaling factors required for the first linear update (line 8). The upper bounds are rescaled to account for a consistent scaling of B and X . After rescaling X (line 9), B and X are consistently scaled. Line 12 computes the column-wise scaling factors required for the second linear update (line 15). As line 14 applies the computed scaling factors prior to this linear update, the linear update itself can safely be implemented with a call to **DGEMM**. Line 17 computes the scaling factors necessary for the third linear update (line 19). The scaling of the final output B is δ and corresponds to the product of all scaling factors.

The application of the Givens rotations in line 10 is not guarded by overflow protection logic, but can result in growth that possibly exceeds the overflow threshold. Since the evaluation of overflow protection logic is expensive compared to the cheap Givens transformations, our software lowers the true overflow threshold Ω by some safety margin. This safety margin is set to the tile height b_r , which overestimates the maximum growth possible by Givens transformations within a tile. Working with Ω/b_r is cheap to compute and guarantees that the Givens transformations in line 10 do not trigger overflow.

Robust backtransform. The robust backward substitution returns segment-wise scaled solution vectors. It remains to compute consistently scaled solutions and backtransform these. Recall that the partitioning into tile rows divides the vector evenly into N segments of length b_r and that $n = Nb_r$. The robust counterpart of (5) is then

$$y_\ell = G_{n,\ell}^T \left(\dots \left(G_{2,\ell}^T \alpha_{\min}^{-1} \begin{bmatrix} \frac{\alpha_{\min}}{\alpha_1} x_\ell(1 : b_r) \\ \vdots \\ \frac{\alpha_{\min}}{\alpha_N} x_\ell(n - b_r + 1 : n) \end{bmatrix} \right) \right). \quad (7)$$

ALGORITHM 10: Robust linear tile update

Data: $H \in \mathbb{R}^{m \times k}$, $X \in \mathbb{R}^{k \times n}$, $B \in \mathbb{R}^{m \times n}$, cross-over columns $\tilde{R} = [\tilde{r}_1, \dots, \tilde{r}_n] \in \mathbb{R}^{m \times n}$,
 $\Lambda = [\lambda_1, \dots, \lambda_n]^T \in \mathbb{R}^n$, Givens rotations computed by **TILEDREDUCE** and stored as
 $C = [c(j, \ell)] \in \mathbb{R}^{k \times n}$ and $S = [s(j, \ell)] \in \mathbb{R}^{k \times n}$

Result: (6) for $\ell = 1, \dots, n$

```

1 function RUPDATE( $H, \tilde{R}, \alpha, X, \Lambda, C, S, \beta, B$ )
2   Allocate  $\delta = [\delta_1, \dots, \delta_n] \in (0, 1]^n$ ;
3   for  $\ell \leftarrow 1 : n$  do
4      $t \leftarrow h(1 : m, 1)$ ;  $t(m) \leftarrow t(m) - \lambda_\ell$ ;
5      $\gamma_\ell \leftarrow \min\{\alpha_\ell, \beta_\ell\}$ ;
6     // Compute a scaling factor such that DAXPY in line 8 cannot overflow
7      $\xi_\ell \leftarrow \text{PROTECTUPDATE}\left(\left\|\frac{\gamma_\ell}{\beta_\ell} b(1 : m, \ell)\right\|_\infty, \|t\|_\infty, \left|\frac{\gamma_\ell}{\alpha_\ell} s(1, \ell)x(1, \ell)\right|\right)$ ;
8      $\delta_\ell \leftarrow \xi_\ell \gamma_\ell$ ; // Update global scaling
9      $b(1 : m, \ell) \leftarrow \frac{\delta_\ell}{\beta_\ell} b(1 : m, \ell) - t \left(\frac{\delta_\ell}{\beta_\ell} s(1, \ell)x(1, \ell)\right)$ ; // Shift-specific DAXPY
10    Copy  $Z \leftarrow X \text{diag}\left(\frac{\delta_1}{\alpha_1}, \dots, \frac{\delta_n}{\alpha_n}\right)$ ;
11    Apply Givens rotations to  $Z$  as in lines 7–13 of Algorithm 6;
12    for  $\ell \leftarrow 1 : n$  do
13      // Compute a scaling factor such that DGEMM in line 15 cannot overflow
14       $\xi_\ell \leftarrow \text{PROTECTUPDATE}(\|b(1 : m, \ell)\|_\infty, \|H\|_\infty, \|z(1 : k - 1, \ell)\|_\infty)$ ;
15       $\delta_\ell \leftarrow \delta_\ell \xi_\ell$ ; // Update global scaling
16       $b(1 : m, \ell) \leftarrow \xi_\ell b(1 : m, \ell)$ ;  $z(1 : k, \ell) \leftarrow \xi_\ell z(1 : k, \ell)$ ; // Apply the scaling factors
17     $B \leftarrow B - H(1 : m, 2 : k)Z(1 : k - 1, 1 : n)$ ; // DGEMM
18    for  $\ell \leftarrow 1 : n$  do
19      // Compute a scaling factor such that DAXPY in line 19 cannot overflow
20       $\xi_\ell \leftarrow \text{PROTECTUPDATE}(\|b(1 : m, \ell)\|_\infty, \|\tilde{r}_\ell\|_\infty, |z(k, \ell)|)$ ;
21       $\delta_\ell \leftarrow \delta_\ell \xi_\ell$ ; // Update global scaling
22       $b(1 : m, \ell) \leftarrow \xi_\ell b(1 : m, \ell) - \tilde{r}_\ell (\xi_\ell z(k, \ell))$ ; // Shift-specific DAXPY
23  return  $\delta, B$ ;
```

The application of the Givens rotations can exceed the overflow threshold. To avoid overflow, Henry [1994, Algorithm 4] evaluates the maximum possible growth possible during the backtransform. If overflow *can* occur, the entire vector is rescaled prior to the backtransform. If the computation targets eigenvectors, an alternative strategy is possible. Eigenvectors are commonly normalized. The backtransform can be executed safely if the vector is normalized with respect to the Euclidean norm before the backtransform. The consistency scaling, the normalization, and the backtransform can be computed in two sweeps over the vector. Algorithm 11 **RBACKTRANSFORM** gives the details. The lines 4–12 closely follow the LAPACK 3.9.0 routine **DNRM2**, which computes the Euclidean norm with scaling to avoid overflow. The lines 14–18 simultaneously normalize and backtransform an eigenvector. After the backtransform, the eigenvectors are still normalized.

By replacing all routines with their robust counterparts and adding segment-wise scaling factors, Algorithm 7 **TILEDSOLVE** can be rendered robust. The resulting algorithm **ROBUSTTILEDSOLVE** is listed in Algorithm 12. This leads to **DHSRQ3** listed in Algorithm 13, which solves shifted Hessenberg systems in a tiled, robust fashion. Note that the first part **TILEDREDUCE** is untouched. Only the second part **TILEDSOLVE** is replaced with its robust counterpart. Since the structure is identical to the non-robust version, task parallelism as introduced in Section 3.3 is valid for **DHSRQ3** as well.

ALGORITHM 11: Consistency scaling, backtransform and normalization of a single eigenvector

Data: Given rotations $G_{j,\ell}^T$ as in (2) computed by `TILEDREDUCE` and stored as $C = [c(j, \ell)] \in \mathbb{R}^{n \times m}$ and $S = [s(j, \ell)] \in \mathbb{R}^{n \times m}$, matrix of scaling factors $\alpha = [\alpha(i, \ell)] \in (0, 1]^{N \times m}$, segment-wise scaled $X \in \mathbb{R}^{n \times m}$, tile size b_r with $n = Nb_r$

Result: Normalized (7) for $\ell \leftarrow 1, \dots, m$

```

1 function RBACKTRANSFORM( $C, S, \alpha, X$ )
2   for  $\ell \leftarrow 1 : m$  do
3      $\alpha_{\min} \leftarrow \min_{1 \leq h \leq N} \{\alpha(h, \ell)\};$  // Global scaling factor of current column
4      $x_{\max} \leftarrow 0;$ 
5      $t \leftarrow 0;$ 
6     for  $i \leftarrow 1 : n$  do
7        $i_{\text{tile}} \leftarrow \lceil i/b_r \rceil;$  // Index of current segment
8       if  $x(i, \ell) \neq 0$  then
9          $x(i, \ell) \leftarrow \frac{\alpha_{\min}}{\alpha(i_{\text{tile}}, \ell)} x(i, \ell);$  // Consistency scaling
10        if  $x_{\max} < |x(i, \ell)|$  then
11           $t \leftarrow 1 + t \left( \frac{x_{\max}}{|x(i, \ell)|} \right)^2; x_{\max} \leftarrow |x(i, \ell)|;$ 
12        else
13           $t \leftarrow t + \left( \frac{|x(i, \ell)|}{x_{\max}} \right)^2;$ 
14       $x_{\text{nrm}} \leftarrow \alpha_{\min}^{-1} x_{\max} \sqrt{t};$  // Euclidean norm of current column
15      // Simultaneous normalization and backtransform
16       $\tau_1 \leftarrow x(1, \ell)/x_{\text{nrm}};$ 
17      for  $i \leftarrow 2 : n$  do
18         $\tau_2 \leftarrow x(i, \ell)/x_{\text{nrm}};$ 
19         $y(i-1, \ell) \leftarrow c(i) \tau_1 - s(i, \ell) \tau_2;$ 
20         $\tau_1 \leftarrow c(i, \ell) \tau_2 + s(i, \ell) \tau_1;$ 
21  return  $Y;$  //  $\|y(:, \ell)\|_2 = 1$ 

```

4.2 Starting Vector and Convergence Test

It is well understood that the choice of the starting vector is crucial for both the convergence and the performance of inverse iteration (1). This is particularly true because the residual can increase by doing more than one iteration. Ipsen [1997, Sections 2.5, 2.6, 6.2] presents a comprehensive summary of work by Varah, Wilkinson and Peters on choosing a suited starting vector. Furthermore, an example demonstrating an increasing residual when more than one iteration is computed can be found in Section 5.4 in the same reference.

A standard choice for the starting vector is a scaled vector of ones $\mathbf{x}^{(0)} \leftarrow \rho[1, \dots, 1]^T$, $\rho > 0$ [Ipsen 1997, p. 259]. In LAPACK 3.9.0 the inverse iteration routine `DHSEIN` chooses $\rho = \|\mathbf{H}\|_{\infty} \epsilon$ assuming $\|\mathbf{H}\|_{\infty} > 0$. With this choice a single iteration of (1) most frequently leads to convergence [Ipsen 1997, p. 264]. If not, LAPACK tries other starting vectors orthogonal to previous choices rather than computing more iterations of (1). The convergence test is passed if $\|\mathbf{x}^{(1)}\|_1 > 0.1/\sqrt{n}$. LAPACK thereby follows Varah's [1968, p. 786] stopping criterion $\frac{\|\mathbf{x}^{(1)}\|_2}{\|\mathbf{x}^{(0)}\|_2} > \frac{1}{c\epsilon}$, where c is a problem-dependent constant and $\mathbf{H}$ is assumed to be normalized $\|\mathbf{H}\|_2 = 1$. Specifically, since LAPACK does not assume $\|\mathbf{H}\|_2 = 1$, we obtain using $\|\mathbf{x}^{(0)}\|_1 = n\|\mathbf{H}\|_{\infty} \epsilon$, $1/\|\mathbf{x}\|_2 \geq 1/\|\mathbf{x}\|_1$

ALGORITHM 12: Robust tiled backward substitution and backtransform

Data: Hessenberg matrix $H \in \mathbb{R}^{n \times n}$, tile size b_r with $n = Nb_r$, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$, $B \in \mathbb{R}^{n \times m}$, tile size b_c with $m = Mb_c$, Givens rotations $G_{j,\ell}^T$ computed with Algorithm 4 and stored as $C \in \mathbb{R}^{n \times m}$ and $S \in \mathbb{R}^{n \times m}$, cross-over columns $\tilde{R}_\ell \in \mathbb{R}^{n \times N}$ for each shift $\ell = 1, \dots, m$ computed with Algorithm 4.

Result: Solution $X \in \mathbb{R}^{n \times m}$ to $HX = B \text{diag}(\lambda_1, \dots, \lambda_m)$

```

1 function ROBUSTTILEDSOLVE( $H, \Lambda, C, S, (\tilde{R}_1, \dots, \tilde{R}_m), B$ )
2    $\alpha \leftarrow \text{ONES}(N, m);$  // Initialize all tile-local scaling factors with 1
3    $X \leftarrow B;$ 
4   for  $\ell \leftarrow 1 : b_c : m$  do
5      $j_{\text{tile}} \leftarrow N;$  // Index of current tile column
6      $\mathcal{L} \leftarrow \ell : \ell + b_c - 1;$  // Index range of current batch of shifts
7     for  $j \leftarrow n - b_r + 1 : -b_r : 1$  do
8        $\mathcal{J} \leftarrow j : j + b_r - 1;$  // Index range of current tile column
9        $\mathcal{J}^- \leftarrow j - 1 : j + b_r - 2;$  // Tile column index range shifted by 1
10      Partition  $\left[ \underbrace{h^{\text{left}}}_{b_r \times 1} \quad \underbrace{H^{\text{diag}}}_{b_r \times (b_r - 1)} \right] \leftarrow H(\mathcal{J}, \mathcal{J}^-);$ 
11      Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{J}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{J}, j_{\text{tile}})];$ 
12       $\alpha(j_{\text{tile}}, \mathcal{L}), X(\mathcal{J}, \mathcal{L}) \leftarrow$ 
13      RSOLVE( $H^{\text{diag}}, h^{\text{left}}, \tilde{R}^{\text{right}}, \Lambda(\mathcal{L}), C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), \alpha(j_{\text{tile}}, \mathcal{L}), X(\mathcal{J}, \mathcal{L})$ );
14      for  $i \leftarrow j - b_r : -b_r : 1$  do
15         $I \leftarrow i : i + b_r - 1;$  // Index range of current tile row
16         $i_{\text{tile}} \leftarrow (i + b_r - 1)/b_r;$  // Index of current tile row
17        Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(I, j_{\text{tile}}), \tilde{r}_{\ell+1}(I, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_c-1}(I, j_{\text{tile}})];$ 
18        if  $i + b_r = j$  then
19           $\alpha(i_{\text{tile}}, \mathcal{L}), X(I, \mathcal{L}) \leftarrow \text{UPDATE}(H(I, \mathcal{J}^-), \tilde{R}^{\text{right}},$ 
20             $\alpha(j_{\text{tile}}, \mathcal{L}), X(\mathcal{J}, \mathcal{L}), \Lambda(\mathcal{L}), C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), \alpha(i_{\text{tile}}, \mathcal{L}), X(I, \mathcal{L}))$ ;
21        else
22           $\alpha(i_{\text{tile}}, \mathcal{L}), X(I, \mathcal{L}) \leftarrow \text{UPDATE}(H(I, \mathcal{J}^-), \tilde{R}^{\text{right}},$ 
23             $\alpha(j_{\text{tile}}, \mathcal{L}), X(\mathcal{J}, \mathcal{L}), 0, C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), \alpha(i_{\text{tile}}, \mathcal{L}), X(I, \mathcal{L}))$ ;
24       $j_{\text{tile}} \leftarrow j_{\text{tile}} - 1;$ 
25       $\mathcal{J} \leftarrow 1 : b_r;$ 
26      Pack  $\tilde{R}^{\text{right}} \leftarrow [\tilde{r}_\ell(\mathcal{J}, j_{\text{tile}}), \dots, \tilde{r}_{\ell+b_c-1}(\mathcal{J}, j_{\text{tile}})];$ 
27       $\alpha(1, \mathcal{L}), X(\mathcal{J}, \mathcal{L}) \leftarrow \text{RSOLVE}(H(\mathcal{J}, 1 : b_r - 1), 0, \Lambda(\mathcal{L}), \tilde{R}^{\text{right}},$ 
28         $C(\mathcal{J}, \mathcal{L}), S(\mathcal{J}, \mathcal{L}), \alpha(1, \mathcal{L}), X(\mathcal{J}, \mathcal{L}))$ ;
29       $X(:, \mathcal{L}) \leftarrow \text{RBACKTRANSFORM}(C(:, \mathcal{L}), S(:, \mathcal{L}), \alpha(:, \mathcal{L}), X(:, \mathcal{L}));$ 
30  return  $X;$ 

```

and $\|\mathbf{x}\|_2 \geq \|\mathbf{x}\|_1/\sqrt{n}$

$$\frac{\|\mathbf{x}^{(1)}\|_2}{\|\mathbf{x}^{(0)}\|_2} \geq \frac{\|\mathbf{x}^{(1)}\|_1}{\sqrt{nn}\|\mathbf{x}^{(0)}\|_1} = \frac{\|\mathbf{x}^{(1)}\|_1}{\sqrt{nn}\|H\|_{\infty}\epsilon} > \frac{1}{10n^2\|H\|_{\infty}\epsilon}.$$

Our inverse iteration solver chooses $\rho = \|H\|_{\infty}\epsilon$. Since the norm $\|\mathbf{x}^{(1)}\|_2$ is readily available after the backward substitution phase due to the consistency scaling, the check $\|\mathbf{x}^{(1)}\|_2 > 0.1/\sqrt{n}$ lends

ALGORITHM 13: Robust tiled simultaneous solution of shifted Hessenberg systems with real shifts

Data: Hessenberg matrix $H \in \mathbb{R}^{n \times n}$, shifts $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$, $B \in \mathbb{R}^{n \times m}$
Result: Solution $X \in \mathbb{R}^{n \times m}$ to $HX = B \text{diag}(\lambda_1, \dots, \lambda_m)$

```

1 function DHSRQ3( $H, \Lambda, B$ )
2    $C, S, (\tilde{R}_1, \dots, \tilde{R}_m) \leftarrow \text{TILEDREDUCE}(H, \Lambda);$ 
3    $X \leftarrow \text{ROBUSTTILEDSOLVE}(H, \Lambda, C, S, (\tilde{R}_1, \dots, \tilde{R}_m), B);$ 
4   return  $X$ ;
```

ALGORITHM 14: Robust tiled inverse iteration for real eigenvalues

Data: Hessenberg matrix $H \in \mathbb{R}^{n \times n}$, eigenvalues $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{R}^m$
Result: $X \in \mathbb{R}^{n \times m}$ such that $HX = X \text{diag}(\lambda_1, \dots, \lambda_m)$

```

1 function DHSRQ3IN( $H, \Lambda$ )
2   Allocate  $X \in \mathbb{R}^{n \times m}$ ;
3   Choose the starting vector  $\mathbf{x}^{(0)} \leftarrow 1/(\epsilon \|H\|_\infty) [1, \dots, 1]^T$ ;
4   Initialize  $X^{(0)}$  with  $m$  repeated copies of  $\mathbf{x}^{(0)}$ ;
5   while not converged do
6     Compute robustly  $X^{(1)} \leftarrow \text{DHSRQ3}(H, \Lambda, X^{(0)})$ ;
7     Split  $X^{(1)}$  into converged eigenvectors  $X_c$  and non-converged eigenvectors  $X_d \in \mathbb{R}^{n \times k}$ ;
8     Append  $X_c$  to  $X$ ;
9     if  $k = 0$  then
10      Go to line 13;
11     Choose a new starting vector  $\mathbf{x}^{(0)}$  orthogonal to previous choices;
12     Initialize  $X^{(0)}$  with  $k$  repeated copies of  $\mathbf{x}^{(0)}$ ;
13   Sort the columns of  $X$  such that the order matches  $\Lambda$ ;
14   return  $X$ ;
```

itself to a quick convergence test. This convergence test is in line with Varah's stopping criterion where $c = 10n\|H\|_\infty$. The information that an eigenvector has not converged can, for example, be propagated by setting the eigenvector to zero.

The decision on a starting vector and the convergence test combined with the robust backward substitution leads to the inverse iteration routine HSRQ3IN, a routine for computing individual eigenvectors simultaneously by inverse iteration. The core of the routine is the robust, tiled solver for shifted Hessenberg systems introduced as DHSRQ3. Algorithm 14 lists the inverse iteration solver. Following LAPACK, only a single iteration of (1) is computed for a given starting vector. After this single iteration, converged eigenvectors are separated from non-converged eigenvectors. New starting vectors are tried for the non-converged eigenvectors.

4.3 Complex Shifts

The RQ factorization $H - \lambda I = RQ$ has complex factors R and Q if λ is complex. The backward substitution with R then relies fully on complex arithmetic. Due to costly multiplications of complex scalars with complex vectors, Henry preferred the UL factorization over the RQ factorization for complex shifts, see Section 2.2. This section presents two techniques which allow extending the inverse iteration solver DHSRQ3IN to support complex shifts at a reasonable computational cost. The first technique chooses the complex Givens rotation such that the reduction phase avoids multiplications of complex scalars with complex vectors. The second technique lowers the cost of

the backward substitution by exploiting that most entries of $\mathbf{H} - \lambda\mathbf{I}$ are real in spite of a complex shift.

The reduction phase requires complex Givens rotations to compute the unitary $\mathbf{Q}$ factor. This paper adopts the Givens rotations applied by Beattie et al. [2012, p. 6]

$$\mathbf{G}_{j,\ell}^H = \mathbf{I}_{j-2} \oplus \begin{bmatrix} c(j, \ell) & -s(j, \ell) \\ s(j, \ell) & \bar{c}(j, \ell) \end{bmatrix} \oplus \mathbf{I}_{n-j} \in \mathbb{C}^{n \times n}, \quad (8)$$

where $c \in \mathbb{C}$ and $s \in \mathbb{R}$. Thereby most of the reduction phase corresponds to mixed real-complex multiplications. In view of the analysis of Givens rotations in floating-point arithmetic by Bindel et al. [2002], our implementation of this Givens rotation is numerically robust and takes care of underflow and overflow.

Next we discuss the changes to Algorithm 2 REDUCEDIAG and Algorithm 3 REDUCEOFFDIAG. The columns of the $\mathbf{R}$ factor are complex and so are the cross-over columns. We store a complex vector $\mathbf{v} = \mathbf{v}^{\text{re}} + i\mathbf{v}^{\text{im}}$ as adjacent columns $[\mathbf{v}^{\text{re}} \quad \mathbf{v}^{\text{im}}]$. Using the complex Givens rotation (8), the complex version of line 5 in Algorithm 3 REDUCEOFFDIAG

$$\underbrace{\mathbf{v}(1:k)}_{\mathbb{C}^k} \leftarrow \underbrace{c(j, \ell)}_{\mathbb{C}} \underbrace{\mathbf{h}(1:k, j)}_{\mathbb{R}^k} + \underbrace{s(j, \ell)}_{\mathbb{R}} \underbrace{\mathbf{v}(1:k)}_{\mathbb{C}^k}$$

can be realized as

$$\begin{aligned} \mathbf{v}^{\text{re}}(1:k) &\leftarrow c^{\text{re}}(j, \ell)\mathbf{h}(1:k, j) + s(j, \ell)\mathbf{v}^{\text{re}}(1:k) \\ \mathbf{v}^{\text{im}}(1:k) &\leftarrow c^{\text{im}}(j, \ell)\mathbf{h}(1:k, j) + s(j, \ell)\mathbf{v}^{\text{im}}(1:k). \end{aligned}$$

The lines 6–7 in REDUCEDIAG can be realized in a similar fashion. The flop count for the complex versions of REDUCEDIAG and REDUCEOFFDIAG doubles compared to their real counterparts and is $3n^2 + O(n)$ and $6nk + O(n)$, respectively.

Next we discuss the changes to the backward substitution phase. Systems with a complex shift yield a complex solution. Aiming for a robust backward substitution, every complex solution vector is associated with a single scaling factor. In other words, the real and the imaginary part are scaled alike. Then a robust SOLVE task addresses $\mathbf{R}_\ell \mathbf{y}_\ell = \gamma_\ell \mathbf{b}_\ell$ where $\gamma_\ell \in (0, 1]$ and all other quantities are complex. If the (small) triangular system matrices $\mathbf{R}_\ell$ are computed explicitly, a conversion to a complex datatype allows the robust solution of this system with a call to ZLATRS for every complex right-hand side. ZLATRS is available in LAPACK 3.9.0. and is the complex counterpart of DLATRS, see Section 4.1. Thereby Algorithm 9 naturally generalizes to complex arithmetic. The flop count is $15n^2 + O(n)$ per right-hand side.

An analysis of the complex counterpart of the linear update (6)

$$\begin{aligned} \delta_\ell^{-1} \mathbf{b}_\ell(\mathcal{J}_1) &\leftarrow \beta_\ell^{-1} \mathbf{b}_\ell(\mathcal{J}_1) - \left[\mathbf{h}(\mathcal{J}_1, j-1) - \lambda_\ell \mathbf{e}_{j-1} \quad \mathbf{H}(\mathcal{J}_1, j:k-2) \quad \tilde{\mathbf{r}}_\ell(\mathcal{J}_1) \right] \\ &\quad \left(\mathbf{G}_{k-1,\ell}^H(\mathcal{J}_2^+, \mathcal{J}_2^+) \cdots \mathbf{G}_{j,\ell}^H(\mathcal{J}_2^+, \mathcal{J}_2^+) \begin{bmatrix} 0 \\ \alpha_\ell^{-1} \mathbf{x}_\ell(\mathcal{J}_2) \end{bmatrix} \right) \end{aligned}$$

reveals the potential for mixed real-complex arithmetic. Analogously to the real linear update, the block structure of the system matrix suggests three block operations. The first block $\mathbf{h}(\mathcal{J}_1, j-1) - \lambda_\ell \mathbf{e}_{j-1}$ and the third block $\tilde{\mathbf{r}}_\ell(\mathcal{J}_1)$ issue multiplications of a complex scalar with a complex vector and can be emulated using real arithmetic. The second block $\mathbf{H}(\mathcal{J}_1, j:k-2)$ is real and requires the multiplication of a real matrix with a complex vector. If complex quantities are stored in interleaved storage, that is, the real and imaginary parts of a complex vector are stored in adjacent columns, the multiplication of a real matrix $\mathbf{H}$ and a complex vector $\mathbf{z} = \mathbf{u} + i\mathbf{v}$ can be realized as $\mathbf{H}[\mathbf{u} \quad \mathbf{v}]$. Hence, the computationally expensive part of the linear update can be realized with a wide DGEMM

Table 1. Flops Approximations for a Single Shift/Eigenvector Segment of Length n and Assuming Square Tiles

Routine	Real eigenvector	Complex eigenvector
REDUCEDIAG	$1.5n^2 + O(n)$	$3n^2 + O(n)$
REDUCEOFFDIAG	$3n^2 + O(n)$	$6n^2 + O(n)$
SOLVE	$3.5n^2 + O(n)$	$15n^2 + O(n)$
UPDATE	$2n^2 + O(n)$	$4n^2 + O(n)$
BACKTRANSFORM	$6(n - 1)$	$20(n - 1)$

Flops due to overflow protection (norm computations, overflow protection logic, numerical scaling) are disregarded.

operation if many right-hand sides are computed simultaneously. The flop count of a linear update with an m -by- n matrix H is approximately $4mn + O(m + n)$ per complex right-hand side.

The complex solution $\mathbf{x}_\ell$ is backtransformed analogously to (5) using complex arithmetic. The flop count sums to $20n + O(1)$. An overview of all flop counts is listed in Table 1. The computational cost per column is approximately the same for a real eigenvector and a complex eigenvector comprising two columns.

Coupling the complex routines of all task types results in ZHSRQ3, the complex counterpart of DHSRQ3. Analogously to LAPACK, the starting vector and the convergence criterion are chosen identically for real and complex eigenvalues. ZHSRQ3 allows generalizing the inverse iteration routine DHSRQ3IN to handle *any* selection of eigenvalues, see Algorithm 15. The selection of eigenvalues is split into real eigenvalues and complex eigenvalues. This requires additional tracking of the affiliation between computed eigenvectors and selected eigenvalues. Recall that complex eigenvectors occur in complex conjugate pairs. Hence, for each selected complex conjugate pair of eigenvalues, only one eigenvector has to be computed. The other one can be obtained for free by complex conjugation. HSRQ3IN assumes that the provided eigenvalues exploit this.

The real/complex eigenvectors are computed by successive calls to the real/complex version of HSRQ3IN. The successive computation of real/complex eigenvectors is justified because the storage requirement quickly limits the problem sizes solvable with HSRQ3IN. An analysis of the storage requirement is given in the next section.

4.4 Storage Requirement Analysis

Algorithm 14 records the Givens rotations and the cross-over columns during the reduction phase. If the Hessenberg matrix is n -by- n , n Givens rotations including the first padded entry have to be stored for every eigenvector. The storage requirement of the cross-over columns depends on the tiling of H . The partitioning of H into an N -by- N grid of $b_r \times b_r$ tiles requires the storage of N (potentially complex) cross-over columns. The storage requirement of the cross-over columns sums to $\sum_{i=1}^N (ib_r) = O(nN)$ for every eigenvector. If m columns (1 column per real or 2 columns per complex eigenvector) are computed, the storage of the cross-over columns is $O(nNm)$ and can quickly exceed the memory available on a compute node. This problem can be addressed by computing the eigenvectors in groups. Algorithm 15 realizes this. Only workspace necessary for storing the cross-over columns and the Givens rotations of a single group has to be allocated. Once a group has been computed, the workspace can be reused for the next group. Since complex eigenvectors and cross-over column are stored in interleaved storage, the storage requirement is doubled compared to the real computation. Line 4 calculates the group size for the complex case requiring two columns per cross-over column. To fully harness the available workspace, the real computation doubles the group size (lines 6–8).

ALGORITHM 15: Robust tiled inverse iteration

Data: Hessenberg matrix $H \in \mathbb{R}^{n \times n}$, tile size b_r with $n = Nb_r$, eigenvalues $\Lambda = [\lambda_1, \dots, \lambda_m]^T \in \mathbb{C}^m$
 such that only one complex eigenvalue of a complex conjugate pair is included in Λ

Result: $X \in \mathbb{C}^{n \times m}$ such that $HX = X \text{diag}(\lambda_1, \dots, \lambda_m)$

```

1 function HSRQ3IN( $H, \Lambda$ )
2   Sort  $\Lambda$  into  $\Lambda_{\text{sorted}} = [\Lambda_{\text{real}} \ \Lambda_{\text{complex}}]$  so that the first  $m_1$  eigenvalues are real;
3   Define the maximum workspace size for the cross-over columns as  $w_{\text{max}}$ ;
4    $g \leftarrow \lfloor w_{\text{max}} / (2nN) \rfloor$ ; // Group size
5   Allocate workspaces for the Givens rotations and the cross-over columns;
6   for  $\ell \leftarrow 1 : 2g : m_1$  do
7      $\mathcal{L} \leftarrow \ell : \max\{\ell + 2g, m\}$ ;
8     Task-parallel  $X(\mathcal{L}) \leftarrow \text{DHSRQ3IN}(H, \Lambda_{\text{sorted}}(\mathcal{L}))$ ; // Real eigenvectors
9   for  $\ell \leftarrow m_1 + 1 : g : m$  do
10     $\mathcal{L} \leftarrow \ell : \max\{\ell + g, m_2\}$ ;
11    Task-parallel  $X(\mathcal{L}) \leftarrow \text{ZHSRQ3IN}(H, \Lambda_{\text{sorted}}(\mathcal{L}))$ ; // Complex eigenvectors
12   Sort the columns of  $X$  so that the order matches  $\Lambda$ ;
13   return  $X$ ;
```

5 NUMERICAL EXPERIMENTS

This section describes how the numerical experiments were set up and executed and presents the results.

5.1 Execution Environment

Hardware. The experiments are run on an Intel Xeon Gold 6132 (Skylake) node where dynamic frequency scaling is enabled. This node has 2 NUMA islands with 14 cores each. In double-precision arithmetic the theoretical peak performance is 83.2 Gflops per core and 2329.6 Gflops per node. The L1 data cache and L2 cache are 32 KB and 1 MB, respectively, and the shared L3 cache is 19.25 MB. The available memory is 192 GB RAM. The memory bandwidth was measured at 12.7 GB/s for one core and 162 GB/s for a full node using the STREAM triad benchmark.

Software and configuration. The software is built with the Intel compiler 19.0.1.144 where the optimization level is set to -O2, AVX-512 instructions are enabled and interprocedural optimizations -ipo are activated. We link against the MKL 2019.1.144 BLAS implementation. OpenMP threads are bound to physical processing units by setting KMP_AFFINITY to compact.

In the following we describe the routines and their configuration used in the numerical experiments. The first routine targets shifted Hessenberg systems and solves $(H - \lambda_\ell I)x_\ell = \alpha_\ell e$. The matrix H is real and upper Hessenberg, the shift λ_ℓ is real or complex, e is the vector with all ones and $\alpha_\ell \in (0, 1]$ is a scaling factor. The next four routines target the computation of eigenvectors by inverse iteration. The routines are supplied with $\|H\|_\infty e$ as the starting vector, see Section 4.2. This starting vector leads to convergence in one iteration in all of our numerical experiments.

- {Z,D}HSRQ3. This robust routine was introduced in this work and solves $(H - \lambda_\ell I)x_\ell = \alpha_\ell e$. It generalizes the RQ approach originally proposed by Henry. The reduction phase records the Givens rotations necessary to compute the RQ factorizations for all shifts. The backward substitution phase utilizes level-3 BLAS for the linear updates.
- RQIN (Henry). This routine extends the shifted Hessenberg system solver by Henry [1994, Algorithm 2] to the computation of eigenvectors. It solves $(H - \lambda_\ell I)x_\ell = \|H\|_\infty e$ through an RQ factorization for every eigenvalue and normalizes every eigenvector after the

backtransform. The core of this routine corresponds to Algorithm 1. While Henry uses the RQ decomposition only for real shifts, the numerical experiments use the RQ approach both for real and complex shifts. The matrix H is not overwritten.

- ULIN (Henry). This routine solves $(H - \lambda_\ell I)x_\ell = \|H\|_\infty \epsilon e$ through an UL factorization for every shift and, in a final step, normalizes the computed eigenvectors. Henry [1994, Section 4] introduced the UL approach for complex shifts to avoid costly complex-complex multiplications. Since the UL approach overwrites H , the original matrix H is copied when more than one system is solved.
- DHSEIN. LAPACK 3.9.0 contains the driver routine DHSEIN for successively computing selected left and/or right eigenvectors of a real upper Hessenberg matrix. DHSEIN calls DLAIEIN for computing a single eigenvector by inverse iteration. The routine is supplied with ϵe as the user-defined starting vector (INITV = 'U'). For each shift, the matrix $B = H - \lambda I$ is explicitly constructed in a workspace. Then $Bx = \alpha(\|H\|_\infty \epsilon e)$ is solved through an LU factorization with partial pivoting. In all numerical experiments conducted here, this initial guess leads to convergence in the first iteration. In other words, DHSEIN effectively solves a single shifted Hessenberg system through an LU decomposition with partial pivoting for each eigenvector.
- HSRQ3IN. This driver routine, listed in Algorithm 15, splits real and complex eigenvalues and computes the corresponding eigenvectors by successive calls to DHSRQ3IN and ZSRQ3IN. It solves $(H - \lambda_\ell I)x_\ell = \alpha_\ell(\|H\|_\infty \epsilon e)$ and normalizes the eigenvectors with respect to the Euclidean norm before the backtransform.

5.2 Test Problems

The numerical experiments use two test problems. The **first test problem** is designed to have known, well-separated eigenvalues and computes the corresponding eigenvectors. This experiment controls the ratio of real/complex eigenvalues and allows us to examine the cost of complex arithmetic. For this purpose, a quasi-triangular matrix $T \in \mathbb{R}^{n \times n}$ is constructed where the eigenvalues are placed as 1-by-1 or 2-by-2 blocks on the diagonal of T . If the k -th eigenvalue is real, then the 1-by-1 diagonal block is $t(k, k) = k$. Complex eigenvalues occur in complex conjugate pairs and correspond to 2-by-2 blocks. Such a 2-by-2 block is set to

$$T(k : k + 1, k : k + 1) = \begin{bmatrix} k & k \\ -k & k \end{bmatrix} \in \mathbb{R}^{2 \times 2} \quad (9)$$

and corresponds to the eigenvalues $k + ik$ and $k - ik$. This choice of diagonal blocks ensures that all eigenvalues are well-separated. In particular, the case with 100% real eigenvalues yields an upper triangular matrix with eigenvalues $1, 2, \dots, n$. Assuming that n is even, the case with 100% complex eigenvalues yields a matrix with only 2-by-2 blocks on the diagonal and eigenvalues $1 \pm i, 3(1 \pm i), \dots, (n - 1)(1 \pm i)$. The remaining superdiagonal entries are random in $(0, 1]$.

The matrix T is then transformed into a Hessenberg matrix through an orthogonal similarity transformation. For this purpose, two orthogonal transformations are applied. First, a random Householder matrix is constructed, $Q_0 = I - 2vv^T$, where v is a random unit norm vector. By applying the (symmetric) Householder matrix, a dense matrix $A = Q_0 T Q_0$ is computed. Second, A is reduced to Hessenberg form through the LAPACK routine DGEHRD $H_1 \leftarrow Q_1^T A Q_1$. Together, H_1 is given by $H_1 \leftarrow (Q_0 Q_1)^T T (Q_1 Q_0)$. The numerical routines receive the exact eigenvalues as input parameter.

The **second test problem** solves shifted Hessenberg systems and aims at quantifying the overhead from overflow protection. Two systems are constructed: the “bad” system requires frequent numerical scaling system, whereas the “good” system never requires numerical scaling.

The bad system constructs the Hessenberg matrix $H_2 = R_2 Q_2 + \gamma I$ where $R_2 \in \mathbb{R}^{n \times n}$ and $Q_2 \in \mathbb{R}^{n \times n}$ are given by

$$r_2(i, j) = \begin{cases} n - i + 1 & i = j \\ -n & j > i \\ 0 & i > j \end{cases}, \quad q_2(i, j) = \begin{cases} -1 & (i - 1 = j) \text{ or } (i = 1 \text{ and } j = n) \\ 0 & \text{otherwise.} \end{cases}$$

As an example, consider how $H_2 - \gamma I$ is constructed for $n = 6$ by

$$H_2 - \gamma I = \begin{bmatrix} 6 & 6 & 6 & 6 & 6 & -6 \\ -5 & 6 & 6 & 6 & 6 & 0 \\ & -4 & 6 & 6 & 6 & 0 \\ & & -3 & 6 & 6 & 0 \\ & & & -2 & 6 & 0 \\ & & & & -1 & 0 \end{bmatrix}$$

$$= R_2 Q_2 = \begin{bmatrix} 6 & -6 & -6 & -6 & -6 & -6 \\ & 5 & -6 & -6 & -6 & -6 \\ & & 4 & -6 & -6 & -6 \\ & & & 3 & -6 & -6 \\ & & & & 2 & -6 \\ & & & & & 1 \end{bmatrix} \begin{bmatrix} 0 & 0 & 0 & 0 & 0 & -1 \\ -1 & 0 & 0 & 0 & 0 & 0 \\ & -1 & 0 & 0 & 0 & 0 \\ & & -1 & 0 & 0 & 0 \\ & & & -1 & 0 & 0 \\ & & & & -1 & 0 \end{bmatrix}.$$

The numerical experiments solve $H_2 X = E \Gamma$ where E is the matrix with all ones and $\Gamma = \text{diag}(2, \dots, 2)$. Hence, the same shifted system is solved repeatedly without exploiting that the shift is shared. Kjelgaard Mikkelsen [2020] has shown that the matrix R_2 introduces quick growth to the solution vectors during the backward substitution.

The good system constructs the Hessenberg matrix $H_3 = R_3 Q_2 + \gamma I$ where $R_3 \in \mathbb{R}^{n \times n}$ is given by

$$r_3(i, j) = \begin{cases} n + 1 - i & i = j \\ \frac{1}{2} & j > i \\ 0 & i > j. \end{cases}$$

The numerical experiments solve $H_3 X = E \Gamma$. Following Kjelgaard Mikkelsen [2020], the solution of $R_3 y = e$ with backward substitution never requires numerical scaling to avoid overflow.

5.3 Validation

The experiments are validated by comparing the relative normwise backward error with 10^{-12} . Specifically, the first test problem evaluates $\frac{\|H_1 x - \lambda x\|_F}{\|H_1\|_F \|x\|_F + |\lambda| \|x\|_F}$ for each computed eigenvector. The second test problem solves the same system repeatedly. The evaluation computes $\frac{\|Ax - \gamma e\|_F}{\|A\|_F \|x\|_F + |\gamma| \|e\|_F}$ only for a single vector, where A corresponds to H_2 or H_3 , respectively.

5.4 Results

This section presents the results of four numerical experiments. The first experiment concerns the sequential runtimes of the inverse iteration solvers and compares the existing approaches DHSEIN (LAPACK), RQIN/ULIN (Henry) with HSRQ3IN introduced in this paper. The second experiment aims at identifying bottlenecks in the implementation and analyzes what fraction each computational phase contributes to the total runtime. The third experiment addresses the parallel scalability of HSRQ3IN. The fourth experiment quantifies the cost of robustness.




Fig. 5. Sequential runtimes on H_1 computing two columns (left) and 1500 columns (right) for $n = 10000$.

Serial Comparison. The first experiment compares the sequential runtimes of the inverse iteration routines DHSEIN, RQIN, ULIN and HSRQ3IN on H_1 where $n = 10000$. Numerical scaling is never triggered and a single iteration suffices to satisfy the convergence criterion for every eigenvector. Figure 5 (left) shows the runtimes for computing two columns, either two real vectors or a single complex vector (storing the real and the imaginary part in two columns). In the latter case, the UL approach overwrites the input matrix and, hence, does not require any data copies. For all solvers the computation of the complex vector is at least as expensive as the computation of two real vectors.

Figure 5 (right) compares the runtimes for computation of 1,500 columns, either 1,500 real vectors or 750 complex vectors. The timing of the UL approach includes the overhead of 749 copies of the input matrix. The runtime gap between DHSEIN and the RQ factorization-based solvers widens with an increasing number of right-hand sides. A large number of right-hand sides allows reusing data. Since the RQ decomposition-based solvers do not overwrite the input matrix, the computation may benefit from temporal locality. HSRQ3IN outperforms RQIN, which can be attributed to the matrix-matrix multiplications (level-3 BLAS) in the backward substitution phase. Complex eigenvectors are more expensive than real eigenvectors for DHSEIN and RQIN. HSRQ3IN, by contrast, performs similarly. The next experiment aims at investigating the underlying cause.

Analysis. This experiment decomposes the computational cost of HSRQ3IN and RQIN and thereby analyzes the ratio of the three computational phases (reduction, backward substitution, backtransform). The experiment setup is identical to the one in Figure 5 (right) and uses H_1 with $n = 10000$ for the computation of either 1,500 real or 750 complex eigenvectors.

A run of RQIN on 1,500 real eigenvectors spends 46% of the runtime in the reduction phase (lines 3, 8–9 in Algorithm 1) and 53% of the runtime in the backward substitution phase (lines 4–7 in Algorithm 1). This ratio is approximately in line with the flop distribution of these two phases. When RQIN computes 750 complex eigenvectors, the reduction phase constitutes 36% and the backward substitution phase 63% of the runtime. Thus, the majority of the time is spent on backward substitution.

The runtime decompositions of HSRQ3IN are shown in Figure 6. The runtime is split into the contribution of each task type to the total compute time for a sequential (left) and a parallel (right) run. Between 53% and 68% of the runtime is spent on REDUCE tasks. The runtime difference between the real and the complex runs is due to the different amount of Givens rotations computed




Fig. 6. Runtime decomposition of a sequential (left) and parallel (right) run on H_1 for $n = 10000$.

during the run. Since a series of Givens rotation is computed per eigenvector, the complex run computes only half the number of Givens rotations, but requires mixed real-complex multiplications. The backward substitution phase contributes with SOLVE and UPDATE tasks. The complex runs spend approximately double the time on SOLVE tasks than the real runs. This can be attributed to the complex-complex multiplications during the small backward substitutions. The runtime difference of UPDATE tasks is due to the application of the Givens rotations. Analogously to the REDUCE tasks, the complex runs apply only half the number of Givens compared to the real runs. The backtransform phase makes a negligible contribution to the total runtime. The runtime decompositions of RQIN and HSRQ3IN suggest that the reduction phase has become the new bottleneck of the revised RQ approach.

The parallel speedup ranges in 14–18 for all task types. Due to dynamic frequency scaling, the best possible parallel speedup on the test node is 18.4. Idle cores during the parallel runs contribute to the overhead/idle time when there are not enough tasks available for being scheduled.

Parallel scalability. The third experiment analyzes the strong scalability of HSRQ3IN. Strong scaling concerns the speedup for a fixed problem size subject to an increasing number of processing units. The used test system is H_1 where $n \in \{10000, 40000\}$. The eigenvalue selection ratio is chosen as 5%, 15%, or 25%. Then, for example, the experiment with $n = 10000$ and 5% selected eigenvalues computes 500 columns (500 real eigenvectors or 250 complex eigenvectors). The eigenvalues are either all real or all complex.

Figure 7 extends the sequential experiment shown in Figure 5. It reports the runtime of a strong scaling experiment with up to 28 cores. The speedup over the sequential run is approximately by factor or 14. The best attainable speedup is 18.4. Figures 5 and 7 show that the new inverse iteration solver is clearly faster than the existing inverse iteration solvers and has reasonable parallel scalability.

Figure 8 displays the performance results with respect to the machine capabilities. It includes the multithreaded implementation of DGEMM as an upper bound. The plot assumes $3.5n^2$ flops per column, which is a lower bound of the true flop count. The real and the complex experiments attain a similar fraction of the theoretical peak performance and achieve a similar parallel speedup.

Overhead due to numerical scaling. The fourth experiment evaluates the cost of overflow avoidance. For this purpose, a run of DHSRQ3 on the bad system H_2 and the good system H_3 are compared. Recall that the bad system requires frequent numerical scaling, whereas the good system




Fig. 7. Parallel runtimes on H_1 using $n = 10000$.




Fig. 8. Parallel scalability on H_1 using $n = 10000$ (left column) and $n = 40000$ (right column). For reference, DGEMM (MKL) executes the same amount of flops as the 5% real shifts case and writes to an output matrix of size n -by- n .

never requires numerical scaling. The experiments solve shifted Hessenberg systems rather than an eigenvector problem because the triangular solve associated with H_2 and H_3 has been shown as the best and worst case, respectively [Kjelgaard Mikkelsen 2020]. As a consequence, the cost of overflow protection for inverse iteration is in between the two extremes. Note that both runs include the cost associated with the evaluation of overflow protection logic. The runtime difference between the two runs is solely the cost of numerical rescaling.

Figure 9 illustrates the measurements using $n = 10000$ and 1,500 real right-hand sides. The runs on the bad system are circa 6% slower than the ones on the good system. Numerical scaling only affects the backward substitution phase, i.e., SOLVE and UPDATE tasks. The slowdown of these two task types is approximately 20% alike. The parallel speedup is comparable for both the good and the bad system. This suggests that robustness does not affect the parallel scalability.

6 CONCLUSION

This paper revises the RQ approach for solving shifted Hessenberg systems $(H - \lambda_\ell I)\mathbf{x}_\ell = \alpha_\ell \mathbf{b}_\ell$ robustly for a large number of shifts. By rearranging the computation of the partial RQ factorization, matrix-matrix multiplications (level-3 BLAS) are introduced to the backward substitution phase. Since the solution of shifted Hessenberg systems is the most compute-intensive step in the computation of eigenvectors by inverse iteration, the revised RQ approach leads to a new inverse




Fig. 9. Runtime decomposition of a sequential (left) and parallel (right) run computing 1500 real right-hand sides on H_2 (bad system) and H_3 (good system) for $n = 10000$.

iteration solver. The numerical experiments show that the new inverse iteration solver outperforms existing inverse iteration solvers.

ACKNOWLEDGMENTS

The author thanks Lars Karlsson for initiating and supporting this project. Furthermore, the author is grateful for the valuable comments by Martin Berggren, Lars Karlsson, Carl Christian Kjelgaard Mikkelsen and the anonymous reviewers.

REFERENCES

- Edward Anderson. 1991. *Robust Triangular Solves for Use in Condition Estimation*. Technical Report. LAWN 36.
- Edward Anderson, Zhaojun Bai, Christian Bischof, L. Susan Blackford, James Demmel, Jack Dongarra, Jeremy Du Croz, Anne Greenbaum, Sven Hammarling, and Alan McKenney. 1999. *LAPACK Users' Guide* (3rd ed.). SIAM.
- Christopher Beattie, Zlatko Drmač, and Serkan Gugercin. 2012. A note on shifted Hessenberg systems and frequency response computation. *ACM Trans. Math. Softw.* 38, 2, Article 12 (Jan. 2012), 16 pages.
- David Bindel, James Demmel, William Kahan, and Osni Marques. 2002. On computing Givens rotations reliably and efficiently. *ACM Trans. Math. Softw.* 28, 2 (June 2002), 206–238.
- L. Susan Blackford, James Demmel, Jack Dongarra, Iain Duff, Sven Hammarling, Greg Henry, Michael Heroux, Linda Kaufman, Andrew Lumsdaine, Antoine Petitet, Roldan Pozo, Karin Remington, and R. Clint Whaley. 2002. An updated set of basic linear algebra subprograms (BLAS). *ACM Trans. Math. Software* 28, 2 (2002), 135–151.
- Nela Bosner, Zvonimir Bujanović, and Zlatko Drmač. 2013. Efficient generalized Hessenberg form and applications. *ACM Transactions on Mathematical Software (TOMS)* 39, 3 (2013), 1–19.
- Nela Bosner, Zvonimir Bujanović, and Zlatko Drmač. 2018. Parallel solver for shifted systems in a hybrid CPU–GPU framework. *SIAM Journal on Scientific Computing* 40, 4 (2018), C605–C633.
- Jack Dongarra, Mark Gates, Azzam Haidar, Jakub Kurzak, Piotr Luszczek, Panruo Wu, Ichitaro Yamazaki, Asim YarKhan, Maksims Abalenkovs, Negin Bagherpour, et al. 2019. PLASMA: Parallel linear algebra software for multicore using OpenMP. *ACM Transactions on Mathematical Software (TOMS)* 45, 2 (2019), 1–35.
- Jack J. Dongarra, Jeremy Du Croz, Sven Hammarling, and Iain S. Duff. 1990. A set of level 3 basic linear algebra subprograms. *ACM Trans. Math. Software* 16, 1 (1990), 1–17.
- Gene H. Golub and Charles F. Van Loan. 1996. *Matrix Computations* (3rd ed.). John Hopkins University Press.
- Greg Henry. 1994. *The Shifted Hessenberg System Solve Computation*. Technical Report. Cornell University, NY, USA.
- Greg Henry. 1995. A parallel unsymmetric inverse iteration solver. In *Proceedings of the Seventh SIAM Conference on Parallel Processing for Scientific Computing*. 546–551.
- Ilse C. F. Ipsen. 1997. Computing an eigenvector with inverse iteration. *SIAM Rev.* 39, 2 (1997), 254–291.
- Carl Christian Kjelgaard Mikkelsen. 2020. Well-conditioned eigenvalue problems that overflow. (2020). [arXiv:2005.05356](https://arxiv.org/abs/2005.05356)
- Carl Christian Kjelgaard Mikkelsen and Lars Karlsson. 2017a. Blocked algorithms for robust solution of triangular linear systems. In *International Conference on Parallel Processing and Applied Mathematics*. Springer, 68–78.
- Carl Christian Kjelgaard Mikkelsen and Lars Karlsson. 2017b. Robust Solution of Triangular Linear Systems, NLAFFET Working Note 9. (May 2017).

- Carl Christian Kjelgaard Mikkelsen, Angelika B. Schwarz, and Lars Karlsson. 2019. Parallel robust solution of triangular linear systems. *Concurrency and Computation: Practice and Experience* 31, 19 (2019), e5064.
- Beresford N. Parlett. 1998. *The Symmetric Eigenvalue Problem*. SIAM.
- Gwendoline Peters and James H. Wilkinson. 1971. The calculation of specified eigenvectors by inverse iteration. In *Handbook for Automatic Computation*. Springer, 418–439.
- Angelika Schwarz. 2020. *Comparing the Performance of Computing Eigenvectors from the Schur Matrix and by Inverse Iteration from the Hessenberg Matrix*. UMINF 21.08. <https://webapps.cs.umu.se/uminf/reports/2021/008/part1.pdf>.
- Field Van Zee, Ernie Chan, Robert A. van de Geijn, Enrique S. Quintana-Ortí, and Gregorio Quintana-Ortí. 2009. Introducing: The libflame library for dense matrix computations. *Computing in Science & Engineering* (2009).
- J. M. Varah. 1968. The calculation of the eigenvectors of a general complex matrix by inverse iteration. *Math. Comp.* 22, 104 (1968), 785–791.

Received September 2020; revised December 2021; accepted May 2022