

REORTHOGONALIZED PYTHAGOREAN VARIANTS OF BLOCK CLASSICAL GRAM–SCHMIDT*

ERIN CARSON[†], KATHRYN LUND[‡], YUXIN MA[†], AND EDA OKTAY[§]

Abstract. Block classical Gram–Schmidt (BCGS) is commonly used for orthogonalizing a set of vectors X in distributed computing environments due to its favorable communication properties relative to other orthogonalization approaches, such as modified Gram–Schmidt or Householder. However, it is known that BCGS (as well as recently developed low-synchronization variants of BCGS) can suffer from a significant loss of orthogonality in finite-precision arithmetic, which can contribute to instability and inaccurate solutions in downstream applications such as s -step Krylov subspace methods. A common solution to improve the orthogonality among the vectors is reorthogonalization. Focusing on the “Pythagorean” variant of BCGS, introduced in [E. Carson, K. Lund, and M. Rozložník, *SIAM J. Matrix Anal. Appl.*, 42 (2021), pp. 1365–1380], which guarantees an $O(\varepsilon)\kappa^2(X)$ bound on the loss of orthogonality as long as $O(\varepsilon)\kappa^2(X) < 1$, where ε denotes the unit roundoff, we introduce and analyze two reorthogonalized Pythagorean BCGS variants. These variants feature favorable communication properties, with asymptotically two synchronization points per block column, as well as an improved $O(\varepsilon)$ bound on the loss of orthogonality. Our bounds are derived in a general fashion to additionally allow for the analysis of mixed-precision variants. We verify our theoretical results with a panel of test matrices and experiments from a new version of the `BlockStab` toolbox.

Key words. Gram–Schmidt algorithm, low-synchronization, communication-avoiding, mixed precision, multiprecision, loss of orthogonality, stability

MSC codes. 65-04, 65F25, 65G50, 65Y20

DOI. 10.1137/24M1658723

1. Introduction. Interest in low-synchronization variants of the Gram–Schmidt method has been proliferating recently [5, 8, 9, 18, 22, 25, 29, 30, 31]. These methods are part of the more general trend of developing communication-reducing orthogonalization routines with the goal of reducing memory movement between levels of the memory hierarchy or nodes on a network, thereby improving scalability in high-

*Received by the editors May 2, 2024; accepted for publication (in revised form) by J. L. Barlow October 24, 2024; published electronically February 4, 2025.

<https://doi.org/10.1137/24M1658723>

Funding: The first, third, and fourth authors are supported by the European Union (ERC, inEXASCALE, 101075632). Views and opinions expressed are those of the authors only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them. The first and the fourth authors acknowledge support from the Charles University GAUK project 202722 and the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration. The first author additionally acknowledges support from the Charles University Research Centre program UNCE/24/SCI/005. The second author would like to thank the Computational Methods in Systems and Control Theory group at the Max Planck Institute for Dynamics of Complex Technical Systems for funding the fourth author’s visit in March 2023. The fourth author would like to thank the Chemnitz University of Technology for funding the second author’s visit in July 2023.

[†]Department of Numerical Mathematics, Faculty of Mathematics and Physics, Charles University, Praha 8, Czechia (carson@karlin.mff.cuni.cz, yuxin.ma@matfyz.cuni.cz).

[‡]Computational Methods in Systems and Control Theory, Max Planck Institute for Dynamics of Complex Technical Systems, 39106 Magdeburg, Germany, and Computational Mathematics Theme, Building R71, STFC Rutherford Appleton Laboratory, Harwell Oxford, Didcot, Oxfordshire, OX11 0QX, UK (kathryn.lund@stfc.ac.uk).

[§]Corresponding author. Department of Numerical Mathematics, Faculty of Mathematics and Physics, Charles University, Praha 8, Czechia, and Department of Mathematics, Chemnitz University of Technology, 09126 Chemnitz, Germany (oktay@karlin.mff.cuni.cz).

performance, and especially exascale, computing [1, 11, 17]. In this manuscript, we concentrate on block Gram–Schmidt (BGS) methods and, in particular, on reorthogonalized versions of block classical Gram–Schmidt with Pythagorean inner product (BCGS-PIP) from [8], and how they can achieve loss of orthogonality on the order of unit roundoff with only two synchronization points per block of columns. In related work [7], we consider a generalization of BCGS2-type algorithms [3, 4], which have four such synchronization points per block, and we use this generalization to study the stability of a reorthogonalized BGS variant with only one synchronization point per block.

We define a *block vector* $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $m \gg s$ as a concatenation of s column vectors, i.e., a tall-skinny matrix. We are interested in computing an economic QR decomposition for the concatenation of p block vectors

$$\mathbf{X} = [\mathbf{X}_1 \quad \mathbf{X}_2 \quad \cdots \quad \mathbf{X}_p] \in \mathbb{R}^{m \times ps}.$$

We achieve this via a BGS method that takes \mathbf{X} and a block size s as arguments and returns an orthonormal basis $\mathbf{Q} \in \mathbb{R}^{m \times ps}$ along with an upper triangular $\mathbf{R} \in \mathbb{R}^{ps \times ps}$ such that $\mathbf{X} = \mathbf{Q}\mathbf{R}$. Both \mathbf{Q} and \mathbf{R} are computed blockwise, meaning that s new columns of \mathbf{Q} are generated per iteration, as opposed to just one column at a time.

Blocking or batching data is a known technique for reducing the total number of synchronization points, or *sync points*. In a distributed setting, we define a sync point as an operation requiring all nodes to send and receive information to and from one other, such as an `MPI.Allreduce`. In an orthogonalization procedure like BGS, block inner products and *intraorthogonalization* routines (which orthogonalize vectors within a block column) like tall-skinny QR require sync points when block vectors are distributed rowwise across nodes. For the purposes of this manuscript, we will assume that a block inner product $\mathbf{X}^*\mathbf{Y}$ and an intraorthogonalization `I0` each constitute one sync point.

In addition to reducing sync points, we are also concerned with the stability of BGS, which we measure here in terms of the *loss of orthogonality (LOO)*,

$$(1.1) \quad \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\|,$$

where I is the $ps \times ps$ identity matrix and $\bar{\mathbf{Q}} \in \mathbb{R}^{m \times ps}$ denotes the \mathbf{Q} factor computed in floating-point arithmetic. We take $\|\cdot\|$ as the induced matrix 2-norm. As is common in the relevant literature, we regard a rectangular matrix $\mathbf{Q} \in \mathbb{R}^{m \times ps}$ as *orthogonal* when $I - \mathbf{Q}^T \mathbf{Q} = 0$, meaning not only are the columns of \mathbf{Q} orthogonal to one another but also each column has norm one. We regard the term *orthonormal* as synonymous with orthogonal.

Orthogonality is important for a number of downstream purposes, especially eigenvalue approximation; see, e.g., [15, 26] and sources therein. Furthermore, having LOO close to working precision simplifies the backward error analysis of Krylov subspace methods like GMRES; see, e.g., the modular backward stability framework by Buttari et al. [6]. Orthogonality is also stronger than being well-conditioned, i.e., having $\kappa(\bar{\mathbf{Q}}) \approx 1$, where κ denotes the 2-condition number of a matrix, i.e., the ratio between its largest and smallest singular values. Indeed, Gram–Schmidt methods frequently lose orthogonality while producing well-conditioned bases. For all these reasons, we concentrate on LOO as the primary metric of a method’s stability.

We will also consider the standard residual

$$(1.2) \quad \left\| \mathbf{Q}\mathbf{R} - \mathbf{X} \right\|,$$

as well as the Cholesky residual,

$$(1.3) \quad \left\| \mathbf{X}^T \mathbf{X} - \bar{\mathbf{R}}^T \bar{\mathbf{R}} \right\|,$$

where $\bar{\mathcal{R}}$ is the finite precision counterpart of \mathcal{R} . This latter residual measures how close a BGS method is to correctly computing a Cholesky decomposition of $\mathbf{X}^T \mathbf{X}$, which can provide insight into the stability pitfalls of a method; see, e.g., [8, 14, 23].

In the following section, we summarize BCGS-PIP and results from [8] and prove stability bounds on two reorthogonalized variants, BCGS-PIP+ and BCGS-PIPI+, which have $2p$ and $2p-1$ sync points, respectively, or roughly 2 sync points per block vector. Section 3 deals with mixed-precision variants of the new reorthogonalized methods, and section 4 demonstrates the numerical behavior of all methods using the `BlockStab` toolbox. We summarize conclusions and future perspectives in section 5.

A few remarks regarding notation are necessary before proceeding. Generally, uppercase Roman letters (R_{ij}, S_{ij}, T_{ij}) denote $s \times s$ block entries of a $ps \times ps$ matrix, which itself is usually denoted by uppercase Roman script ($\mathcal{R}, \mathcal{S}, \mathcal{T}$). A block column of such matrices is denoted with MATLAB indexing:

$$\mathcal{R}_{1:k-1,k} = \begin{bmatrix} R_{1,k} \\ R_{2,k} \\ \vdots \\ R_{k-1,k} \end{bmatrix}.$$

For simplicity, we also abbreviate standard $ks \times ks$ submatrices as $\mathcal{R}_k := \mathcal{R}_{1:k,1:k}$.

Bold uppercase Roman letters ($\mathbf{Q}_k, \mathbf{X}_k, \mathbf{U}_k$) denote $m \times s$ block vectors, and bold, uppercase Roman script ($\mathcal{Q}, \mathcal{X}, \mathcal{U}$) denotes an indexed concatenation of p such vectors. Standard $m \times ks$ submatrices are abbreviated as

$$\mathcal{Q}_k := \mathcal{Q}_{1:k} = [\mathbf{Q}_1 \quad \mathbf{Q}_2 \quad \cdots \quad \mathbf{Q}_k].$$

Note that when \mathbf{X} is used on its own throughout the text, it denotes a generic block vector. When a subscript is added to \mathbf{X} as \mathbf{X}_k , we are referring to the specific k th block vector of the input matrix \mathbf{X} . Furthermore, \mathcal{X}_k is the concatenation of the first k block vectors of \mathbf{X} , as defined above.

The function $[\mathbf{Q}, R] = \text{IO}(\mathbf{X})$ denotes an intraorthogonalization routine, i.e., a method used to orthogonalize vectors within a generic block vector \mathbf{X} . This can be any number of methods, including Householder, classical Gram–Schmidt, modified Gram–Schmidt, or Cholesky QR.

We use ε to denote the *unit roundoff* of a chosen *working precision*. For example, if we use IEEE double precision arithmetic, then $\varepsilon = 2^{-53} \approx 1.1 \cdot 10^{-16}$, and for IEEE single precision, $\varepsilon = 2^{-24} \approx 6.0 \cdot 10^{-8}$. Throughout the text, we use standard results from [16] (particularly sections 2.2 and 3.5) for rounding-error analysis.

We also make use of big-O notation like $\mathcal{O}(\varepsilon)$, which is essentially $c(m, ps)\varepsilon$, with $c(m, ps)$ denoting a low-degree polynomial in dimensional constants m and ps . To enhance clarity, we employ $\mathcal{O}(\varepsilon)$ to disregard the dimensional factor $c(m, ps)$, although this factor can become significant when m and ps are large.

2. Stability of reorthogonalized variants of BCGS-PIP. BCGS-PIP (Algorithm 2.1) is a corrected version of block classical Gram–Schmidt (BCGS), where the block diagonal entries of the \mathcal{R} factor are computed via the block Pythagorean theorem [8]. This correction stabilizes the algorithm by keeping the relative Cholesky residual (1.3) close to the working precision. However, the overall LOO for BCGS-PIP can be quite high, and the bound

$$\|I - \bar{\mathcal{Q}}^T \bar{\mathcal{Q}}\| \leq \mathcal{O}(\varepsilon) \kappa^2(\mathbf{X})$$

only holds when $\mathcal{O}(\varepsilon) \kappa^2(\mathbf{X}) \leq 1$.

Algorithm 2.1 $[\mathcal{Q}, \mathcal{R}] = \text{BCGS-PIP}(\mathcal{X}, \text{IO})$.

```

1:  $[\mathbf{Q}_1, R_{11}] = \text{IO}(\mathbf{X}_1)$ 
2: for  $k = 2, \dots, p$  do
3:    $[\mathcal{R}_{1:k-1, k}^T]_{P_k} = [\mathbf{Q}_{k-1} \ \mathbf{X}_k]^T \mathbf{X}_k$ 
4:    $R_{kk} = \text{chol}(P_k - \mathcal{R}_{1:k-1, k}^T \mathcal{R}_{1:k-1, k})$ 
5:    $\mathbf{V}_k = \mathbf{X}_k - \mathbf{Q}_{k-1} \mathcal{R}_{1:k-1, k}$ 
6:    $\mathbf{Q}_k = \mathbf{V}_k R_{kk}^{-1}$ 
7: end for
8: return  $\mathcal{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_p]$ ,  $\mathcal{R} = (R_{ij})$ 

```

Algorithm 2.2 $[\mathcal{Q}, \mathcal{R}] = \text{BCGS-PIP+}(\mathcal{X}, \text{IO})$.

```

1:  $[\mathcal{U}, \mathcal{S}] = \text{BCGS-PIP}(\mathcal{X}, \text{IO})$ 
2:  $[\mathcal{Q}, \mathcal{T}] = \text{BCGS-PIP}(\mathcal{U}, \text{IO})$ 
3:  $\mathcal{R} = \mathcal{T}\mathcal{S}$ ;
4: return  $\mathcal{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_p]$ ,  $\mathcal{R} = (R_{ij})$ 

```

Thus as $\kappa(\mathcal{X}) \rightarrow \varepsilon^{-1/2}$, any guarantees on the orthogonality of $\bar{\mathcal{Q}}$ are lost.

A standard strategy for improving $\bar{\mathcal{Q}}$ is running the Gram–Schmidt procedure twice; see, e.g., [4, 24] for analysis of reorthogonalized block variants of BCGS. We do not consider BCGS further in this manuscript; for a detailed analysis and provable bounds on its LOO, see [7], which also analyzes low-sync reorthogonalized versions of BCGS. In the following two subsections, we propose two reorthogonalized versions of BCGS-PIP and prove bounds for their LOO and residuals.

2.1. BCGS-PIP+. We first consider the simple approach of running BCGS-PIP twice in a row. See Algorithm 2.2 for what we call BCGS-PIP+, where + stands for “reorthogonalization.” Note that despite how the pseudocode is written, it is not necessary to store two bases in practice: \mathcal{U} can be stored in place of \mathcal{X} throughout the first step. Similarly, it is possible to build \mathcal{R} by replacing \mathcal{S} gradually in the second step, and there is no need to construct \mathcal{T} explicitly; cf. the second phase of Algorithm 2.3. The pseudocode is written in such a way as to simplify the mathematical analysis.

Under minimal assumptions, proving an $\mathcal{O}(\varepsilon)$ LOO bound for Algorithm 2.2 follows directly from the results of [8]. In particular, we can obtain the following result by applying [8, Theorem 3.4] twice. Note the condition $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$ required for $\text{IO}(\mathbf{X})$: when \mathbf{X}_k is a block vector of \mathcal{X} , this condition holds by [15, Corollary 8.6.3]. Here and elsewhere, it is meant to ensure that IO does not handle block vectors with worse conditioning than that of \mathcal{X} .

COROLLARY 2.1. *Let $\mathcal{X} \in \mathbb{R}^{m \times ps}$ with $\mathcal{O}(\varepsilon) \kappa^2(\mathcal{X}) \leq \frac{1}{2}$ and $\bar{\mathcal{Q}}$ and $\bar{\mathcal{R}}$ be computed by Algorithm 2.2. Assuming that for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$, $[\mathbf{Q}, \bar{R}] = \text{IO}(\mathbf{X})$ satisfy*

$$\begin{aligned} \bar{R}^T \bar{R} &= \mathbf{X}^T \mathbf{X} + \Delta E, & \|\Delta E\| &\leq \mathcal{O}(\varepsilon) \|\mathbf{X}\|^2, \text{ and} \\ \bar{\mathbf{Q}} \bar{R} &= \mathbf{X} + \Delta D, & \|\Delta D\| &\leq \mathcal{O}(\varepsilon) (\|\mathbf{X}\| + \|\bar{\mathbf{Q}}\| \|\bar{R}\|), \end{aligned}$$

then $\bar{\mathcal{Q}}$ and $\bar{\mathcal{R}}$ satisfy

$$(2.1) \quad \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\| \leq \mathcal{O}(\varepsilon), \text{ and}$$

$$(2.2) \quad \bar{\mathbf{Q}} \bar{\mathbf{R}} = \mathbf{X} + \Delta \mathbf{D}, \quad \|\Delta \mathbf{D}\| \leq \mathcal{O}(\varepsilon) \|\mathbf{X}\|.$$

We can obtain a further result on the Cholesky residual of BCGS-PIP+ via the following corollary, which follows by applying [8, Theorem 3.2] twice.

COROLLARY 2.2. *Let $\mathbf{X} \in \mathbb{R}^{m \times ps}$ with $\mathcal{O}(\varepsilon) \kappa^2(\mathbf{X}) \leq \frac{1}{2}$ and $\bar{\mathbf{Q}}$ and $\bar{\mathbf{R}}$ be computed by Algorithm 2.2. Assuming that for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathbf{X})$, $[\bar{\mathbf{Q}}, \bar{\mathbf{R}}] = IO(\mathbf{X})$ satisfy*

$$\begin{aligned} \bar{\mathbf{R}}^T \bar{\mathbf{R}} &= \mathbf{X}^T \mathbf{X} + \Delta E, \quad \|\Delta E\| \leq \mathcal{O}(\varepsilon) \|\mathbf{X}\|^2, \text{ and} \\ \bar{\mathbf{Q}} \bar{\mathbf{R}} &= \mathbf{X} + \Delta D, \quad \|\Delta D\| \leq \mathcal{O}(\varepsilon) (\|\mathbf{X}\| + \|\bar{\mathbf{Q}}\| \|\bar{\mathbf{R}}\|), \end{aligned}$$

then $\bar{\mathbf{R}}$ satisfies

$$(2.3) \quad \bar{\mathbf{R}}^T \bar{\mathbf{R}} = \mathbf{X}^T \mathbf{X} + \Delta \mathcal{E}, \quad \|\Delta \mathcal{E}\| \leq \mathcal{O}(\varepsilon) \|\mathbf{X}\|^2.$$

For our mixed-precision analysis, it will also be useful to have a generalized formulation of these bounds that do not rely on a specific precision and that also reveal all the sources of error. From standard rounding-error principles [16], we can write the following bounds for each step of Algorithm 2.2 for constants $\delta_{US}, \omega_U, \delta_{QT}, \omega_Q, \delta_{TS} \in (0, 1)$:

$$(2.4) \quad \bar{\mathbf{U}} \bar{\mathbf{S}} = \mathbf{X} + \Delta \mathcal{D}_{US}, \quad \|\Delta \mathcal{D}_{US}\| \leq \delta_{US} \|\mathbf{X}\|;$$

$$(2.5) \quad \left\| I - \bar{\mathbf{U}}^T \bar{\mathbf{U}} \right\| \leq \omega_U \kappa^2(\mathbf{X});$$

$$(2.6) \quad \bar{\mathbf{Q}} \bar{\mathbf{T}} = \bar{\mathbf{U}} + \Delta \mathcal{D}_{QT}, \quad \|\Delta \mathcal{D}_{QT}\| \leq \delta_{QT} \|\bar{\mathbf{U}}\|;$$

$$(2.7) \quad \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\| \leq \omega_Q \kappa^2(\bar{\mathbf{U}}); \text{ and}$$

$$(2.8) \quad \bar{\mathbf{R}} = \bar{\mathbf{T}} \bar{\mathbf{S}} + \Delta \mathcal{D}_{TS}, \quad \|\Delta \mathcal{D}_{TS}\| \leq \delta_{TS} \|\bar{\mathbf{T}}\| \|\bar{\mathbf{S}}\|.$$

The following theorem summarizes how each step of Algorithm 2.2 influences the LOO and residual per iteration, which will be useful in section 3 when we consider a two-precision variant. We omit dependencies on k from the δ constants, meaning each can be regarded as a maximum over all constants stemming from the same step of the algorithm. We also generally assume that such constants are much smaller than 1 and drop “quadratic” terms resulting from their products.

THEOREM 2.3. *Let $\mathbf{X} \in \mathbb{R}^{m \times ps}$ and suppose $[\bar{\mathbf{Q}}, \bar{\mathbf{R}}] = \text{BCGS-PIP+}(\mathbf{X}, IO)$. Assuming that (2.4)–(2.8) are satisfied with $\omega_U \kappa^2(\mathbf{X}) \leq \frac{1}{2}$ and $\omega_Q \leq \frac{1}{6}$, then $\bar{\mathbf{Q}}$ and $\bar{\mathbf{R}}$ satisfy*

$$(2.9) \quad \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\| \leq 3 \cdot \omega_Q, \text{ and}$$

$$(2.10) \quad \bar{\mathbf{Q}} \bar{\mathbf{R}} = \mathbf{X} + \Delta \mathcal{D}, \quad \|\Delta \mathcal{D}\| \leq (\delta_{US} + \sqrt{6} \cdot \delta_{QT} + 9 \cdot \delta_{TS}) \|\mathbf{X}\|.$$

Proof. From (2.5) and the assumption $\omega_U \kappa^2(\mathbf{X}) \leq \frac{1}{2}$, we obtain

$$(2.11) \quad \|\bar{\mathbf{U}}\|^2 = \left\| \bar{\mathbf{U}}^T \bar{\mathbf{U}} \right\| \leq \left\| I - \bar{\mathbf{U}}^T \bar{\mathbf{U}} \right\| + \|I\| \leq \frac{3}{2},$$

and by the perturbation theory of singular values [15, Corollary 8.6.2], we obtain a lower bound on $\sigma_{\min}^2(\bar{\mathbf{U}})$, namely, $\sigma_{\min}^2(\bar{\mathbf{U}}) \geq \frac{1}{2}$. Consequently, we have the following for $\kappa^2(\bar{\mathbf{U}})$:

$$(2.12) \quad \kappa^2(\bar{\mathbf{U}}) = \frac{\sigma_{\max}^2(\mathbf{U})}{\sigma_{\min}^2(\mathbf{U})} \leq 3.$$

Applying (2.12) to (2.7) immediately confirms (2.9).

To prove (2.10), we combine (2.4), (2.6), and (2.8) to arrive at

$$(2.13) \quad \bar{\mathbf{Q}}\bar{\mathbf{R}} = \mathbf{x} + \underbrace{\Delta\mathcal{D}_{US} + \Delta\mathcal{D}_{QT}\bar{\mathcal{S}} + \bar{\mathcal{Q}}\Delta\mathcal{D}_{TS}}_{=:\Delta\mathcal{D}}.$$

Note that $\|\bar{\mathcal{Q}}\| \leq \frac{\sqrt{6}}{2}$ is derived similarly to (2.11), which together yield

$$(2.14) \quad \begin{aligned} \|\Delta\mathcal{D}\| &\leq \|\Delta\mathcal{D}_{US}\| + \|\Delta\mathcal{D}_{QT}\| \|\bar{\mathcal{S}}\| + \|\bar{\mathcal{Q}}\| \|\Delta\mathcal{D}_{TS}\| \\ &\leq \delta_{US} \|\mathbf{x}\| + \delta_{QT} \|\bar{\mathcal{S}}\| + \frac{\sqrt{6}}{2} \cdot \delta_{TS} \|\bar{\mathcal{T}}\| \|\bar{\mathcal{S}}\|. \end{aligned}$$

By (2.4) and (2.6), we obtain

$$\begin{aligned} \bar{\mathcal{S}}^T \bar{\mathcal{S}} &= \mathbf{x}^T \mathbf{x} + \mathbf{x}^T \Delta\mathcal{D}_{US} + \Delta\mathcal{D}_{US} \mathbf{x} + \bar{\mathcal{S}}^T (I - \bar{\mathbf{U}}^T \bar{\mathbf{U}}) \bar{\mathcal{S}}, \\ \bar{\mathcal{T}}^T \bar{\mathcal{T}} &= \bar{\mathbf{U}}^T \bar{\mathbf{U}} + \bar{\mathbf{U}}^T \Delta\mathcal{D}_{QT} + \Delta\mathcal{D}_{QT}^T \bar{\mathbf{U}} + \bar{\mathcal{T}}^T (I - \bar{\mathcal{Q}}^T \bar{\mathcal{Q}}) \bar{\mathcal{T}}. \end{aligned}$$

Furthermore, from (2.11), (2.12), and the assumptions $\omega_U \kappa^2(\mathbf{x}) \leq \frac{1}{2}$ and $\omega_Q \leq \frac{1}{6}$, it follows that

$$(2.15) \quad \|\bar{\mathcal{S}}\| \leq \sqrt{2 + 4 \cdot \delta_{US}} \|\mathbf{x}\| \leq \sqrt{6} \|\mathbf{x}\|,$$

$$(2.16) \quad \|\bar{\mathcal{T}}\| \leq \sqrt{3 + 6 \cdot \delta_{QT}} \leq 3.$$

Substituting (2.15) and (2.16) into (2.14) proves (2.10). □

2.2. BCGS-PIPI+. One drawback to Algorithm 2.2 is that two for-loops are required to reorthogonalize the entire basis. In practice, we can combine the for-loops without introducing additional sync points. The resulting algorithm is denoted BCGS-PIPI+, where I+ stands for “inner reorthogonalization,” and is provided as Algorithm 2.3. Note that, as in Algorithm 2.2, we construct matrices \mathcal{S} , \mathcal{T} , and \mathbf{U} throughout the algorithm, although none of these need to be explicitly formed in practice.

Combining the for-loops and introducing a general IO for the first orthogonalization step of Algorithm 2.3 creates new challenges in proving its stability bounds, primarily because the first block vector is no longer reorthogonalized. We can therefore no longer directly use the results from [8], as a stricter condition will have to be placed on the choice of IO. In the following, we will first develop a generalized approach depending on small constants stemming from particular sources of error, similarly to Theorem 2.3. We then conclude the section with an application to the uniform-precision case by induction over the generalized results.

For the following, we assume all quantities are computed by Algorithm 2.3 and that at each iteration $k \in \{2, \dots, p\}$, there exists $\omega_{k-1} \in (0, 1)$ such that

$$(2.17) \quad \left\| I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \right\| \leq \omega_{k-1}.$$

Algorithm 2.3 $[\mathcal{Q}, \mathcal{R}] = \text{BCGS-PIPI}+(\mathcal{X}, \text{IO})$.

```

1:  $[\mathbf{Q}_1, R_{11}] = \text{IO}(\mathbf{X}_1)$  ▷  $S_{11} = R_{11}, \mathbf{U}_1 = \mathbf{Q}_1, T_{11} = I$ 
2: for  $k = 2, \dots, p$  do
3:    $[\begin{smallmatrix} S_{1:k-1,k} \\ \Omega_k \end{smallmatrix}] = [\mathbf{Q}_{k-1} \ \mathbf{X}_k]^T \mathbf{X}_k$  ▷ First BCGS-PIP step
4:    $S_{kk} = \text{chol}(\Omega_k - \mathcal{S}_{1:k-1,k}^T \mathcal{S}_{1:k-1,k})$ 
5:    $\mathbf{V}_k = \mathbf{X}_k - \mathbf{Q}_{k-1} \mathcal{S}_{1:k-1,k}$ 
6:    $\mathbf{U}_k = \mathbf{V}_k S_{kk}^{-1}$ 
7:    $[\begin{smallmatrix} \mathcal{T}_{1:k-1,k} \\ P_k \end{smallmatrix}] = [\mathbf{Q}_{k-1} \ \mathbf{U}_k]^T \mathbf{U}_k$  ▷ Second BCGS-PIP step
8:    $T_{kk} = \text{chol}(P_k - \mathcal{T}_{1:k-1,k}^T \mathcal{T}_{1:k-1,k})$ 
9:    $\mathbf{W}_k = \mathbf{U}_k - \mathbf{Q}_{k-1} \mathcal{T}_{1:k-1,k}$ 
10:   $\mathbf{Q}_k = \mathbf{W}_k T_{kk}^{-1}$ 
11:   $\mathcal{R}_{1:k-1,k} = \mathcal{S}_{1:k-1,k} + \mathcal{T}_{1:k-1,k} S_{kk}$  ▷ Finalize  $\mathcal{R}$  entries
12:   $R_{kk} = T_{kk} S_{kk}$ 
13: end for
14: return  $\mathcal{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_p], \mathcal{R} = (R_{ij})$ 

```

Then similarly to (2.11),

$$(2.18) \quad \|\bar{\mathbf{Q}}_{k-1}\| \leq \sqrt{1 + \omega_{k-1}} \leq \sqrt{2}.$$

We can write the following for intermediate quantities computed by BCGS-PIPI+ at each iteration $k \in \{2, \dots, p\}$, where we summarize rounding-error bounds via constants $\delta_* \in (0, 1)$:

$$(2.19) \quad \bar{S}_{1:k-1,k} = \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta \mathbf{S}_k, \quad \|\Delta \mathbf{S}_k\| \leq \delta_{Q^T X} \|\mathbf{X}_k\|;$$

$$(2.20) \quad \bar{\Omega}_k = \mathbf{X}_k^T \mathbf{X}_k + \Delta \Omega_k, \quad \|\Delta \Omega_k\| \leq \delta_{X^T X} \|\mathbf{X}_k\|^2;$$

$$(2.21) \quad \bar{S}_{kk}^T \bar{S}_{kk} = \bar{\Omega}_k - \bar{S}_{1:k-1,k}^T \bar{S}_{1:k-1,k} + \Delta F_k^{(1)} + \Delta C_k^{(1)},$$

$$(2.22) \quad \|\Delta F_k^{(1)}\| \leq \delta_{S^T S} \|\mathbf{X}_k\|^2, \quad \|\Delta C_k^{(1)}\| \leq \delta_{\text{cho1}_1} \|\mathbf{X}_k\|^2;$$

$$(2.23) \quad \bar{\mathbf{V}}_k = \mathbf{X}_k - \bar{\mathbf{Q}}_{k-1} \bar{S}_{1:k-1,k} + \Delta \mathbf{V}_k, \quad \|\Delta \mathbf{V}_k\| \leq \delta_{QS} \|\mathbf{X}_k\|;$$

$$(2.24) \quad \bar{\mathbf{U}}_k \bar{S}_{kk} = \bar{\mathbf{V}}_k + \Delta \mathbf{G}_k^{(1)}, \quad \|\Delta \mathbf{G}_k^{(1)}\| \leq \delta_U \|\bar{\mathbf{U}}_k\| \|\bar{S}_{kk}\|;$$

$$(2.25) \quad \bar{\mathcal{T}}_{1:k-1,k} = \bar{\mathbf{Q}}_{k-1}^T \bar{\mathbf{U}}_k + \Delta \mathbf{T}_k, \quad \|\Delta \mathbf{T}_k\| \leq \delta_{Q^T U} \|\bar{\mathbf{U}}_k\|;$$

$$(2.26) \quad \bar{P}_k = \bar{\mathbf{U}}_k^T \bar{\mathbf{U}}_k + \Delta P_k, \quad \|\Delta P_k\| \leq \delta_{U^T U} \|\bar{\mathbf{U}}_k\|^2;$$

$$(2.27) \quad \bar{T}_{kk}^T \bar{T}_{kk} = \bar{P}_k - \bar{\mathcal{T}}_{1:k-1,k}^T \bar{\mathcal{T}}_{1:k-1,k} + \Delta F_k^{(2)} + \Delta C_k^{(2)},$$

$$(2.28) \quad \|\Delta F_k^{(2)}\| \leq \delta_{T^T T} \|\bar{\mathbf{U}}_k\|^2, \quad \text{and} \quad \|\Delta C_k^{(2)}\| \leq \delta_{\text{cho1}_2} \|\bar{\mathbf{U}}_k\|^2;$$

$$(2.29) \quad \bar{\mathbf{W}}_k = \bar{\mathbf{U}}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathcal{T}}_{1:k-1,k} + \Delta \mathbf{W}_k, \quad \|\Delta \mathbf{W}_k\| \leq \delta_{QT} \|\bar{\mathbf{U}}_k\|; \text{ and}$$

$$(2.30) \quad \bar{\mathbf{Q}}_k \bar{T}_{kk} = \bar{\mathbf{W}}_k + \Delta \mathbf{G}_k^{(2)}, \quad \|\Delta \mathbf{G}_k^{(2)}\| \leq \delta_Q \|\bar{\mathbf{Q}}_k\| \|\bar{T}_{kk}\|.$$

We have applied (2.18) and dropped quadratic error terms throughout the proofs in this section and in (2.19)–(2.30) to simplify the expressions; namely, $\|\mathbf{u}_k\|$ and $\|\mathbf{Q}_k\|$ are baked into the constants, where $\mathbf{u}_k = [\mathbf{U}_1 \ \mathbf{U}_2 \ \dots \ \mathbf{U}_k]$. We also again omit explicit dependence on k for readability. Note as well that each $\Delta F_k^{(i)}$ denotes the floating-point error from the subtraction of the inner product from the previously

computed $\bar{\Omega}_k$ or \bar{P}_k , and $\Delta C_k^{(i)}$ denotes the error from the Cholesky factorization of that result.

In order to apply the Cholesky factorization and obtain (2.21) and (2.27) in the first place, we need that $\bar{\Omega}_k - \bar{S}_{1:k-1,k}^T \bar{S}_{1:k-1,k} + \Delta F_k^{(1)}$ and $\bar{P}_k - \bar{T}_{1:k-1,k}^T \bar{T}_{1:k-1,k} + \Delta F_k^{(2)}$ are symmetric positive definite. We show that $\bar{\Omega}_k - \bar{S}_{1:k-1,k}^T \bar{S}_{1:k-1,k} + \Delta F_k^{(1)}$ is positive definite using the following lemma; symmetry is already clear.

LEMMA 2.4. Fix $k \in \{2, \dots, p\}$ and suppose that (2.17)–(2.20) are satisfied, along with

$$(2.31) \quad \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{R}}_{k-1} = \mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1}, \quad \|\Delta \mathcal{D}_{k-1}\| \leq \delta_X \|\mathbf{X}_{k-1}\|.$$

Assume

$$(2.32) \quad (\delta_X + 2 \cdot \omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + 2\sqrt{2} \cdot \delta_{Q^T X}) \kappa^2(\mathbf{X}_k) < 1.$$

Then it holds that

$$(2.33) \quad \lambda_{\min}(\bar{\Omega}_k - \bar{S}_{1:k-1,k}^T \bar{S}_{1:k-1,k} + \Delta F_k^{(1)}) > 0.$$

Proof. By (2.19) and (2.20), and dropping quadratic error terms, we have

$$(2.34) \quad \begin{aligned} &\bar{\Omega}_k - \bar{S}_{1:k-1,k}^T \bar{S}_{1:k-1,k} + \Delta F_k^{(1)} \\ &= \mathbf{X}_k^T (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k + \Delta G_k, \end{aligned}$$

where

$$\begin{aligned} \Delta G_k &:= \Delta \Omega_k + \Delta F_k^{(1)} - \mathbf{X}_k^T \bar{\mathcal{Q}}_{k-1} \Delta \mathcal{S}_k - \Delta \mathcal{S}_k^T \bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k \\ &\quad + \mathbf{X}_k^T \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k \end{aligned}$$

satisfies the following, thanks to

$$\begin{aligned} \left\| \mathbf{X}_k^T \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k \right\| &= \left\| \mathbf{X}_k^T \bar{\mathcal{Q}}_{k-1} (I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1}) \bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k \right\| \\ &\leq \|\mathbf{X}_k\|^2 \|\bar{\mathcal{Q}}_{k-1}\|^2 \left\| I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \right\| \end{aligned}$$

and (2.18)–(2.20):

$$(2.35) \quad \begin{aligned} \|\Delta G_k\| &\leq \|\mathbf{X}_k\|^2 \|\bar{\mathcal{Q}}_{k-1}\|^2 \left\| I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \right\| + \delta_{S^T S} \|\mathbf{X}_k\|^2 \\ &\quad + \delta_{X^T X} \|\mathbf{X}_k\|^2 + 2\sqrt{1 + \omega_{k-1}} \delta_{Q^T X} \|\mathbf{X}_k\|^2 \\ &\leq ((1 + \omega_{k-1})\omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + 2\sqrt{1 + \omega_{k-1}} \delta_{Q^T X}) \|\mathbf{X}_k\|^2. \end{aligned}$$

Using (2.34) and (2.35), we obtain

$$(2.36) \quad \begin{aligned} &\lambda_{\min}(\bar{\Omega}_k - \bar{S}_{1:k-1,k}^T \bar{S}_{1:k-1,k} + \Delta F_k^{(1)}) \\ &\geq \lambda_{\min}(\mathbf{X}_k^T (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k) - \|\Delta G_k\| \\ &= \sigma_{\min}^2((I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k) - \|\Delta G_k\| \\ &\geq \sigma_{\min}^2((I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k) - ((1 + \omega_{k-1})\omega_{k-1} + \delta_{S^T S} \\ &\quad + \delta_{X^T X} + 2\sqrt{1 + \omega_{k-1}} \delta_{Q^T X}) \|\mathbf{X}_k\|^2. \end{aligned}$$

Then we estimate $\sigma_{\min}((I - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T) \mathbf{X}_k)$. Notice that $(I - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T) \mathbf{X}_k$ can be written as, by assumption (2.31),

$$(2.37) \quad \begin{aligned} \mathbf{X}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k &= \mathbf{X}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathcal{R}}_{k-1} \bar{\mathcal{R}}_{k-1}^{-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k \\ &= [\mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1} \quad \mathbf{X}_k] \mathbf{C}^T \end{aligned}$$

with $\mathbf{C} := [-(\bar{\mathcal{R}}_{k-1}^{-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k)^T \quad I_s] \in \mathbb{R}^{s \times sk}$, and I_s denoting the $s \times s$ identity matrix. Then from the perturbation theory of singular values [15, Corollary 8.6.2], we obtain

$$\begin{aligned} &\sigma_{\min}((I - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T) \mathbf{X}_k) \\ &= \sqrt{\min_{\mathbf{v} \in \mathbb{R}^{sk} \setminus \mathbf{0}} \left(\frac{\|[\mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1} \quad \mathbf{X}_k] \mathbf{C}^T \mathbf{v}\|^2}{\|\mathbf{v}\|^2} \right)} \\ &= \sqrt{\min_{\mathbf{v} \in \mathbb{R}^{sk} \setminus \mathbf{0}} \left(\frac{\|[\mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1} \quad \mathbf{X}_k] \mathbf{C}^T \mathbf{v}\|^2 \|\mathbf{C}^T \mathbf{v}\|^2}{\|\mathbf{C}^T \mathbf{v}\|^2 \|\mathbf{v}\|^2} \right)} \\ &\geq \sqrt{\min_{\mathbf{v} \in \mathbb{R}^{sk} \setminus \mathbf{0}} \left(\frac{\|[\mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1} \quad \mathbf{X}_k] (\mathbf{C}^T \mathbf{v})\|^2}{\|\mathbf{C}^T \mathbf{v}\|^2} \right)} \min_{\mathbf{v} \in \mathbb{R}^{sk} \setminus \mathbf{0}} \left(\frac{\|\mathbf{C}^T \mathbf{v}\|^2}{\|\mathbf{v}\|^2} \right) \\ &\geq \sigma_{\min}([\mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1} \quad \mathbf{X}_k]) \\ &\geq \sigma_{\min}(\mathbf{X}_k) - \|\Delta \mathcal{D}_{k-1}\| \\ &\geq \sigma_{\min}(\mathbf{X}_k) - \delta_X \|\mathbf{X}_{k-1}\|. \end{aligned}$$

Together with (2.36), it holds that

$$(2.38) \quad \begin{aligned} &\lambda_{\min}(\bar{\Omega}_k - \bar{\mathcal{S}}_{1:k-1,k}^T \bar{\mathcal{S}}_{1:k-1,k} + \Delta F_k^{(1)}) \\ &\geq \sigma_{\min}^2(\mathbf{X}_k) - \delta_X \|\mathbf{X}_k\|^2 - ((1 + \omega_{k-1})\omega_{k-1} + \delta_{S^T S} \\ &\quad + \delta_{X^T X} + 2\sqrt{1 + \omega_{k-1}}\delta_{Q^T X}) \|\mathbf{X}_k\|^2 \\ &\geq \sigma_{\min}^2(\mathbf{X}_k) (1 - (\delta_X + (1 + \omega_{k-1})\omega_{k-1} + \delta_{S^T S} \\ &\quad + \delta_{X^T X} + 2\sqrt{1 + \omega_{k-1}}\delta_{Q^T X}) \kappa^2(\mathbf{X}_k)). \end{aligned}$$

By dropping the quadratic error terms and using assumption (2.32), which guarantees

$$(\delta_X + (1 + \omega_{k-1})\omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + 2\sqrt{1 + \omega_{k-1}}\delta_{Q^T X}) \kappa^2(\mathbf{X}_k) < 1,$$

we can therefore conclude that $\lambda_{\min}(\bar{\Omega}_k - \bar{\mathcal{S}}_{1:k-1,k}^T \bar{\mathcal{S}}_{1:k-1,k} + \Delta F_k^{(1)}) > 0$. \square

From Lemma 2.4, we have shown that $\bar{\Omega}_k - \bar{\mathcal{S}}_{1:k-1,k}^T \bar{\mathcal{S}}_{1:k-1,k} + \Delta F_k^{(1)}$ is symmetric positive definite. Before proving that $\bar{P}_k - \bar{\mathcal{T}}_{1:k-1,k}^T \bar{\mathcal{T}}_{1:k-1,k} + \Delta F_k^{(2)}$ is also symmetric positive definite, we need to first bound $\|\bar{\mathbf{U}}_k\|$.

LEMMA 2.5. *Fix $k \in \{2, \dots, p\}$ and suppose that (2.17)–(2.24) are satisfied. Furthermore suppose that \mathbf{Q}_{k-1} and \mathcal{R}_{k-1} satisfy the following:*

$$(2.39) \quad \bar{\mathbf{Q}}_{k-1} \bar{\mathcal{R}}_{k-1} = \mathbf{X}_{k-1} + \Delta \mathcal{D}_{k-1}, \quad \|\Delta \mathcal{D}_{k-1}\| \leq \delta_X \|\mathbf{X}_{k-1}\|$$

with $\delta_X, \omega_{k-1} \in (0, 1)$. Assume

$$(2.40) \quad 8 \cdot \delta_U \kappa^2(\mathcal{X}) \leq 1$$

and

$$(2.41) \quad (\delta_X + 2 \cdot \omega_{k-1} + 2 \cdot \delta_{S^T S} + 2 \cdot \delta_{X^T X} + 2 \cdot \delta_{\text{cho}l_1} + 18 \cdot \delta_{Q^T X} + 8 \cdot \delta_{QS}) \kappa^2(\mathcal{X}_k) \leq \frac{1}{2}.$$

Then it holds that

$$(2.42) \quad \bar{S}_{kk}^T \bar{S}_{kk} = \mathbf{X}_k^T \mathbf{X}_k - \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta S_{kk}, \text{ and}$$

$$(2.43) \quad \bar{\mathbf{U}}_k \bar{S}_{kk} = \mathbf{X}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta \mathbf{U}_k, \text{ where}$$

$$(2.44) \quad \|\Delta S_{kk}\| \leq (\delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho}l_1} + 2\sqrt{2} \cdot \delta_{Q^T X}) \|\mathbf{X}_k\|^2,$$

$$(2.45) \quad \|\bar{S}_{kk}\| \leq 3 \|\mathbf{X}_k\|,$$

$$(2.46) \quad \|\bar{S}_{kk}^{-1}\| \leq \frac{\sqrt{2}}{\sigma_{\min}(\mathcal{X}_k)},$$

$$(2.47) \quad \|\Delta \mathbf{U}_k\| \leq (\delta_U \|\bar{\mathbf{U}}_k\| + \sqrt{2} \cdot \delta_{Q^T X} + \delta_{QS}) \|\mathbf{X}_k\|, \text{ and}$$

$$(2.48) \quad \|\bar{\mathbf{U}}_k\| \leq 2.$$

Proof. From (2.18), (2.19), and (2.20), we obtain

$$(2.49) \quad \begin{aligned} \|\bar{S}_{1:k-1,k}\| &\leq (\sqrt{2} + \delta_{Q^T X}) \|\mathbf{X}_k\| \text{ and} \\ \|\bar{\Omega}_k\| &\leq (1 + \delta_{X^T X}) \|\mathbf{X}_k\|^2. \end{aligned}$$

Substituting (2.19), (2.20), and (2.49) into (2.21) gives (2.42) with

$$\Delta S_{kk} := \Delta \Omega_k + \Delta F_k^{(1)} + \Delta C_k^{(1)} - \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \Delta \mathbf{S}_k - \Delta \mathbf{S}_k^T \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k.$$

From the bounds (2.18), (2.19), (2.20), (2.21), and (2.22), the desired bounds (2.44) and (2.45) then follow immediately.

To find (2.46), we rewrite (2.42) as

$$(2.50) \quad \bar{S}_{kk}^T \bar{S}_{kk} = \mathbf{X}_k^T (I - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T) (I - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T) \mathbf{X}_k + \Delta \tilde{S}_k,$$

where $\Delta \tilde{S}_k := \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T (I - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T) \mathbf{X}_k + \Delta S_{kk}$. Using a similar logic in the proof of Lemma 2.4 together with (2.17) and (2.18) we obtain (2.46) since $\|\bar{S}_{kk}^{-1}\| = \frac{1}{\sigma_{\min}(\bar{S}_{kk})}$ and

$$(2.51) \quad \begin{aligned} \sigma_{\min}^2(\bar{S}_{kk}) &\geq \sigma_{\min}^2(\mathcal{X}_k) - \delta_X \|\mathbf{X}_k\|^2 - (\omega_{k-1} + \delta_{S^T S} \\ &\quad + \delta_{X^T X} + \delta_{\text{cho}l_1} + 2\sqrt{2} \cdot \delta_{Q^T X}) \|\mathbf{X}_k\|^2. \end{aligned}$$

Note that the above reasoning makes sense only if

$$\sigma_{\min}^2(\mathcal{X}_k) > (\delta_X + \omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho}l_1} + 2\sqrt{2} \cdot \delta_{Q^T X}) \|\mathbf{X}_k\|^2,$$

that is, if

$$(\delta_X + \omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho}l_1} + 2\sqrt{2} \cdot \delta_{Q^T X}) \kappa^2(\mathcal{X}_k) < 1,$$

which can be guaranteed by assumption (2.41).

To prove (2.47) we first note from (2.23) and the substitution of (2.19) that

$$(2.52) \quad \begin{aligned} \bar{\mathbf{V}}_k &= \mathbf{X}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta \mathbf{V}_k, \\ \|\Delta \mathbf{V}_k\| &\leq (\sqrt{2} \cdot \delta_{Q^T X} + \delta_{QS}) \|\mathbf{X}_k\|. \end{aligned}$$

Substituting (2.52) into (2.24) leads to (2.43) with

$$\Delta \mathbf{U}_k := \Delta \mathbf{G}_k^{(1)} + \Delta \mathbf{V}_k$$

and the desired bound (2.47) satisfied.

The final bound (2.48) requires a bit more work. Multiplying $\bar{\mathbf{U}}_k \bar{\mathbf{S}}_{kk}$ by its transpose gives

$$(2.53) \quad \begin{aligned} (\bar{\mathbf{U}}_k \bar{\mathbf{S}}_{kk})^T (\bar{\mathbf{U}}_k \bar{\mathbf{S}}_{kk}) &= \mathbf{X}_k^T \mathbf{X}_k - 2 \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k \\ &\quad + \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta H_k \\ &= \mathbf{X}_k^T \mathbf{X}_k - \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k \\ &\quad + \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} (\bar{\mathbf{Q}}_{k-1}^T \bar{\mathbf{Q}}_{k-1} - I) \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta H_k, \end{aligned}$$

where

$$\Delta H_k := \mathbf{X}_k^T \Delta \mathbf{U}_k + (\Delta \mathbf{U}_k)^T \mathbf{X}_k + \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \Delta \mathbf{U}_k + (\Delta \mathbf{U}_k)^T \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k.$$

Applying (2.18) and (2.47) gives

$$(2.54) \quad \|\Delta H_k\| \leq (4 + 2 \cdot \omega_{k-1}) (\delta_U \|\bar{\mathbf{U}}_k\| + \sqrt{2} \cdot \delta_{Q^T X} + \delta_{QS}) \|\mathbf{X}_k\|^2.$$

Substituting (2.42) into (2.53) leads to

$$(\bar{\mathbf{U}}_k \bar{\mathbf{S}}_{kk})^T (\bar{\mathbf{U}}_k \bar{\mathbf{S}}_{kk}) = \bar{\mathbf{S}}_{kk}^T \bar{\mathbf{S}}_{kk} + \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} (\bar{\mathbf{Q}}_{k-1}^T \bar{\mathbf{Q}}_{k-1} - I) \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta H_k - \Delta S_{kk},$$

and then multiplying by S_{kk}^{-T} on the left and S_{kk}^{-1} on the right yields

$$(2.55) \quad \begin{aligned} \bar{\mathbf{U}}_k^T \bar{\mathbf{U}}_k &= I + \bar{\mathbf{S}}_{kk}^{-T} \mathbf{X}_k^T \bar{\mathbf{Q}}_{k-1} (\bar{\mathbf{Q}}_{k-1}^T \bar{\mathbf{Q}}_{k-1} - I) \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k \bar{\mathbf{S}}_{kk}^{-1} \\ &\quad + \bar{\mathbf{S}}_{kk}^{-T} (\Delta H_k - \Delta S_{kk}) \bar{\mathbf{S}}_{kk}^{-1}. \end{aligned}$$

Recalling the assumption (2.40), we note that

$$4(2 + \omega_{k-1}) \delta_U \kappa^2(\mathcal{X}_k) \leq 4(2 + \omega_{k-1}) \delta_U \kappa^2(\mathcal{X}) \leq 1.$$

Applying (2.17), (2.18), (2.54), (2.44), (2.40), along with (2.46), taking norms yields the following:

$$(2.56) \quad \begin{aligned} \|\bar{\mathbf{U}}_k\|^2 &= \|\bar{\mathbf{U}}_k^T \bar{\mathbf{U}}_k\| \\ &\leq 1 + \omega_{k-1} \|\bar{\mathbf{S}}_{kk}^{-1}\|^2 \|\mathbf{X}_k\|^2 + 4(\delta_U \|\bar{\mathbf{U}}_k\| + \sqrt{2} \cdot \delta_{Q^T X} \\ &\quad + \delta_{QS}) \|\bar{\mathbf{S}}_{kk}^{-1}\|^2 \|\mathbf{X}_k\|^2 + (\delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho1}_1} \\ &\quad + 2\sqrt{2} \cdot \delta_{Q^T X}) \|\bar{\mathbf{S}}_{kk}^{-1}\|^2 \|\mathbf{X}_k\|^2 \\ &\leq 1 + 2(\omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho1}_1} + 9 \cdot \delta_{Q^T X} + 4 \cdot \delta_{QS}) \kappa^2(\mathcal{X}_k) \\ &\quad + 8 \cdot \delta_U \|\bar{\mathbf{U}}_k\| \kappa^2(\mathcal{X}_k) \\ &\leq \frac{3}{2} + \|\bar{\mathbf{U}}_k\|. \end{aligned}$$

Solving the quadratic inequality (2.56) gives

$$\|\bar{U}_k\| \leq 2. \quad \square$$

Now using similar logic as in the proof of Lemma 2.4, we can combine the above lemma with (2.25) and (2.26) to conclude that $\bar{P}_k - \bar{T}_{1:k-1,k}^T \bar{T}_{1:k-1,k} + \Delta F_k^{(2)}$ is symmetric positive definite.

The following theorem makes the relationship between each source of error and the bound on the LOO explicit.

THEOREM 2.6. Fix $k \in \{2, \dots, p\}$ and assume that (2.17)–(2.30) are satisfied with

$$(2.57) \quad \begin{aligned} &2(4 \cdot \omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + \delta_{\text{chol}_1} + 7 \cdot \delta_{QS} + 13 \cdot \delta_{Q^T X} + 14 \cdot \delta_U) \kappa^2(\mathcal{X}) \\ &+ 4(\omega_{k-1} + \delta_{U^T U} + \delta_{T^T T} + \delta_{\text{chol}_2} + 3 \cdot \delta_{Q^T U}) \leq \frac{1}{2}. \end{aligned}$$

Further assume that the assumptions of Lemma 2.5 hold. Then \bar{Q}_k satisfies

$$(2.58) \quad \begin{aligned} \|I - \bar{Q}_k^T \bar{Q}_k\| &\leq 17 \cdot \omega_{k-1} + 30 \cdot \delta_{Q^T U} + 40 \cdot \delta_{QT} \\ &+ 120 \cdot \delta_Q + 8 \cdot \delta_{T^T T} + 8 \cdot \delta_{\text{chol}_2}. \end{aligned}$$

Proof. Writing $\bar{Q}_k = [\bar{Q}_{k-1} \quad \bar{Q}_k]$, we can look at $I - \bar{Q}_k^T \bar{Q}_k$ block-by-block:

$$(2.59) \quad I - \bar{Q}_k^T \bar{Q}_k = \begin{bmatrix} I - \bar{Q}_{k-1}^T \bar{Q}_{k-1} & \bar{Q}_{k-1}^T \bar{Q}_k \\ \bar{Q}_k^T \bar{Q}_{k-1} & I - \bar{Q}_k^T \bar{Q}_k \end{bmatrix}.$$

The induction hypothesis takes care of bounding the upper left block. For the off-diagonals, by (2.48), (2.25), (2.26), and Lemma 2.5, we obtain

$$(2.60) \quad \begin{aligned} \|\bar{T}_{1:k-1,k}\| &\leq 2(\sqrt{1 + \omega_{k-1}} + \delta_{Q^T U}), \\ \|\bar{P}_k\| &\leq 4(1 + \delta_{U^T U}). \end{aligned}$$

Using (2.27), the assumption (2.57), and dropping the quadratic terms, the following bound holds:

$$(2.61) \quad \begin{aligned} \|\bar{T}_{kk}\| &\leq \sqrt{4(1 + \delta_{U^T U}) + 4(\sqrt{1 + \omega_{k-1}} + \delta_{Q^T U})^2 + 4\delta_{T^T T} + 4\delta_{\text{chol}_2}} \\ &\leq \sqrt{\frac{17}{2}} \leq 3. \end{aligned}$$

Using (2.25) and (2.26), we can rewrite (2.27) as

$$(2.62) \quad \begin{aligned} \bar{T}_{kk}^T \bar{T}_{kk} &= \bar{U}_k^T (I - \bar{Q}_{k-1} \bar{Q}_{k-1}^T) (I - \bar{Q}_{k-1} \bar{Q}_{k-1}^T) \bar{U}_k + \Delta \tilde{T}_k, \\ \|\Delta \tilde{T}_k\| &\leq 4(\omega_{k-1} + \delta_{U^T U} + \delta_{T^T T} + \delta_{\text{chol}_2} + 3 \cdot \delta_{Q^T U}). \end{aligned}$$

To bound $\|\bar{T}_{kk}^{-1}\|$ we can use (2.62) to write

$$(2.63) \quad \begin{aligned} \sigma_{\min}^2(\bar{T}_{kk}) &\geq \sigma_{\min}^2((I - \bar{Q}_{k-1} \bar{Q}_{k-1}^T) \bar{U}_k) - \|\Delta \tilde{T}_k\| \\ &\geq \sigma_{\min}^2((I - \bar{Q}_{k-1} \bar{Q}_{k-1}^T) \bar{U}_k) \\ &\quad - 4(\omega_{k-1} + \delta_{U^T U} + \delta_{T^T T} + \delta_{\text{chol}_2} + 3 \cdot \delta_{Q^T U}). \end{aligned}$$

By (2.43), we find that

$$(2.64) \quad \bar{U}_k = \left((I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \mathbf{X}_k + \Delta \mathbf{U}_k \right) \bar{S}_{kk}^{-1}.$$

Multiplying (2.64) by $I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T$ on the left yields

$$(2.65) \quad \begin{aligned} (I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \bar{U}_k &= \bar{U}_k - \bar{\mathcal{Q}}_{k-1} (I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1}) \bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k \bar{S}_{kk}^{-1} \\ &\quad - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T \Delta \mathbf{U}_k \bar{S}_{kk}^{-1}. \end{aligned}$$

Define $\Delta \mathbf{E}_k := \bar{\mathcal{Q}}_{k-1} (I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1}) \bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k \bar{S}_{kk}^{-1} + \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T \Delta \mathbf{U}_k \bar{S}_{kk}^{-1}$. Using (2.46), we then have

$$\|\mathbf{X}_k\| \|\bar{S}_{kk}^{-1}\| \leq \sqrt{2} \cdot \kappa(\mathbf{X}_k).$$

Together with the induction hypothesis and (2.48), we can bound $\|\Delta \mathbf{E}_k\|$ by

$$(2.66) \quad \|\Delta \mathbf{E}_k\| \leq \sqrt{2} (\omega_{k-1} + 2 \cdot \delta_U + \sqrt{2} \delta_{Q^T X} + \delta_{QS}) \kappa(\mathbf{X}_k).$$

Via (2.55), (2.56), (2.65), and (2.66), we find lower bounds on some key singular values:¹

$$\begin{aligned} \sigma_{\min}^2(\bar{U}_k) &\geq 1 - \left\| \bar{S}_{kk}^{-T} \mathbf{X}_k^T \bar{\mathcal{Q}}_{k-1} (I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1}) \bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k \bar{S}_{kk}^{-1} \right\| \\ &\quad - \left\| \bar{S}_{kk}^{-T} (\Delta H_k - \Delta S_{kk}) \bar{S}_{kk}^{-1} \right\| \\ &\geq 1 - 2(\omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho1}} + 6\sqrt{2} \cdot \delta_{Q^T X} + 4 \cdot \delta_{QS} \\ &\quad + 8 \cdot \delta_U) \kappa^2(\mathbf{X}_k) \end{aligned}$$

and

$$(2.67) \quad \begin{aligned} &\sigma_{\min}^2((I - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{Q}}_{k-1}^T) \bar{U}_k) \\ &\geq (\sigma_{\min}(\bar{U}_k) - \|\Delta \mathbf{E}_k\|)^2 \\ &\geq \sigma_{\min}^2(\bar{U}_k) - 2 \|\Delta \mathbf{E}_k\| \|\bar{U}_k\| \\ &\geq \sigma_{\min}^2(\bar{U}_k) - 4\sqrt{2} (\omega_{k-1} + 2 \cdot \delta_U + \sqrt{2} \delta_{Q^T X} + \delta_{QS}) \kappa(\mathbf{X}_k) \\ &\geq 1 - 2(4 \cdot \omega_{k-1} + \delta_{S^T S} + \delta_{X^T X} + \delta_{\text{cho1}} + 13 \cdot \delta_{Q^T X} \\ &\quad + 7 \cdot \delta_{QS} + 14 \cdot \delta_U) \kappa^2(\mathbf{X}_k). \end{aligned}$$

Combining (2.67) with (2.62), (2.63), and using the fact that $\|\bar{T}_{kk}^{-1}\|^2 = \frac{1}{\sigma_{\min}^2(\bar{T}_{kk})}$, we arrive at

$$(2.68) \quad \|\bar{T}_{kk}^{-1}\| = \sqrt{\|\bar{T}_{kk}^{-1}\|^2} \leq \sqrt{2}.$$

With similar logic as in the derivation of (2.24), (2.30) gives

$$(2.69) \quad \bar{\mathcal{Q}}_k \bar{T}_{kk} = \bar{\mathbf{W}}_k + \Delta \mathbf{G}_k^{(2)}, \quad \|\Delta \mathbf{G}_k^{(2)}\| \leq \delta_Q \|\bar{\mathcal{Q}}_k\| \|\bar{T}_{kk}\| \leq 6 \cdot \delta_Q,$$

where we have applied (2.18) and (2.61) and simplified the bound. Multiplying (2.69) by $\bar{\mathcal{Q}}_{k-1}^T$ on the left yields

$$(2.70) \quad \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k \bar{T}_{kk} = \bar{\mathcal{Q}}_{k-1}^T \bar{\mathbf{W}}_k + \bar{\mathcal{Q}}_{k-1}^T \Delta \mathbf{G}_k^{(2)}.$$

¹Again by [15, Corollary 8.6.2].

Then we substitute (2.29) into (2.70), multiply both sides by \bar{T}_{kk}^{-1} , and use (2.25) to obtain

$$\bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k = (I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1}) \bar{\mathcal{T}}_{1:k-1,k} \bar{T}_{kk}^{-1} + \left(\Delta \mathbf{T}_k + \bar{\mathcal{Q}}_{k-1}^T (\Delta \mathbf{G}_k^{(2)} + \Delta \mathbf{W}_k) \right) \bar{T}_{kk}^{-1}.$$

Combining (2.58), (2.48), Lemma 2.5, and bounds (2.25), (2.68), (2.69), and (2.29) leads to

$$(2.71) \quad \left\| \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k \right\| \leq 4 \cdot \omega_{k-1} + 2\sqrt{2} \cdot \delta_{QTU} + 12 \cdot \delta_Q + 4 \cdot \delta_{QT}.$$

To bound the bottom right entry of (2.59), we combine (2.27) and (2.30) and note that

$$(2.72) \quad \begin{aligned} \bar{T}_{kk}^T (I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k) \bar{T}_{kk} &= \bar{T}_{kk}^T \bar{T}_{kk} - \bar{T}_{kk}^T \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \bar{T}_{kk} \\ &= \bar{P}_k - \bar{T}_{1:k-1,k}^T \bar{\mathcal{T}}_{1:k-1,k} + \Delta F_k^{(2)} + \Delta C_k^{(2)} \\ &\quad - (\bar{\mathbf{W}}_k - \Delta \mathbf{G}_k^{(2)})^T (\bar{\mathbf{W}}_k - \Delta \mathbf{G}_k^{(2)}) \\ &= \bar{P}_k - \bar{T}_{1:k-1,k}^T \bar{\mathcal{T}}_{1:k-1,k} - \bar{\mathbf{W}}_k^T \bar{\mathbf{W}}_k \\ &\quad - \underbrace{\bar{\mathbf{W}}_k^T \Delta \mathbf{G}_k^{(2)} - (\Delta \mathbf{G}_k^{(2)})^T \bar{\mathbf{W}}_k + \Delta F_k^{(2)} + \Delta C_k^{(2)}}_{=: \Delta L_k}. \end{aligned}$$

Further substituting (2.26) and (2.29) into (2.72) simplifies to

$$(2.73) \quad \begin{aligned} \bar{T}_{kk}^T (I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k) \bar{T}_{kk} &= -\bar{T}_{1:k-1,k}^T \bar{\mathcal{T}}_{1:k-1,k} + \bar{U}_k^T \bar{\mathcal{Q}}_{k-1} \bar{T}_{1:k-1,k} \\ &\quad + \bar{T}_{1:k-1,k}^T \bar{\mathcal{Q}}_{k-1}^T \bar{U}_k - \bar{T}_{1:k-1,k}^T \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \bar{T}_{1:k-1,k} \\ &\quad - \Delta J_k - \Delta L_k, \end{aligned}$$

where $\Delta J_k := (\bar{U}_k - \bar{\mathcal{Q}}_{k-1} \bar{T}_{1:k-1,k})^T \Delta \mathbf{W}_k - (\Delta \mathbf{W}_k)^T (\bar{U}_k - \bar{\mathcal{Q}}_{k-1} \bar{T}_{1:k-1,k})$. One more substitution of (2.25) into (2.73) and further simplification yields

$$(2.74) \quad \begin{aligned} \bar{T}_{kk}^T (I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k) \bar{T}_{kk} &= \bar{T}_{1:k-1,k}^T (I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1}) \bar{T}_{1:k-1,k} \\ &\quad - (\Delta \mathbf{T}_k)^T \bar{T}_{1:k-1,k} - \bar{T}_{1:k-1,k}^T \Delta \mathbf{T}_k - \Delta J_k - \Delta L_k. \end{aligned}$$

Multiplying (2.74) by \bar{T}_{kk}^{-T} on the left and \bar{T}_{kk}^{-1} on the right, taking norms, and applying nearly all previous bounds along with (2.17) leads to

$$(2.75) \quad \begin{aligned} \left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| &\leq 8 \cdot \omega_{k-1} + 24 \cdot \delta_{QTU} + 32 \cdot \delta_{QT} \\ &\quad + 96 \cdot \delta_Q + 8 \cdot \delta_{T^T T} + 8 \cdot \delta_{\text{chol}_2}. \end{aligned}$$

Finally, using (2.59) along with (2.17) and the bounds (2.71) and (2.75), we see that

$$\begin{aligned} \left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| &= \left\| \begin{bmatrix} I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} & \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k \\ \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_{k-1} & I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \end{bmatrix} \right\| \\ &\leq \left\| \begin{bmatrix} \left\| I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \right\| & \left\| \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k \right\| \\ \left\| \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_{k-1} \right\| & \left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| \end{bmatrix} \right\| \\ &\leq \left\| \begin{bmatrix} \left\| I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \right\| & \left\| \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k \right\| \\ \left\| \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_{k-1} \right\| & \left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| \end{bmatrix} \right\|_{\text{F}} \\ &\leq \left\| I - \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_{k-1} \right\| + 2 \left\| \bar{\mathcal{Q}}_{k-1}^T \bar{\mathcal{Q}}_k \right\| + \left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| \\ &\leq 17 \cdot \omega_{k-1} + 30 \cdot \delta_{QTU} + 40 \cdot \delta_{QT} + 120 \cdot \delta_Q + 8 \cdot \delta_{T^T T} + 8 \cdot \delta_{\text{chol}_2}, \end{aligned}$$

where we have used [13, p. 15.50] as in [8, Theorem 3.1]. \square

If we assume a uniform working precision with unit roundoff ε and apply standard floating-point analysis [16] to Theorem 2.6, it is not hard to show that

$$\delta_{Q^T X}, \delta_{X^T X}, \delta_{S^T S}, \delta_{QS}, \delta_U, \delta_{Q^T U}, \delta_{U^T U}, \delta_{T^T T}, \delta_{QT}, \delta_Q \leq \mathcal{O}(\varepsilon)$$

and

$$\delta_{\text{cho1}_1}, \delta_{\text{cho1}_2} \leq \mathcal{O}(\varepsilon).$$

In contrast to Corollary 2.1, we must impose an I0 condition on I0; HouseQR [16], TSQR [20], Cho1QR++ [28], or ShCho1QR++ [12] should satisfy the requirement, but only TSQR could do so and still maintain a single sync point for the I0 [2].² The following corollaries summarize bounds for BCGS-PIPI+ in uniform precision.

COROLLARY 2.7. *Assume that $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ and that for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$, $[\bar{\mathbf{Q}}, \bar{\mathbf{R}}] = \text{IO}(\mathbf{X})$ satisfy*

$$\begin{aligned} \bar{\mathbf{R}}^T \bar{\mathbf{R}} &= \mathbf{X}^T \mathbf{X} + \Delta E, \quad \|\Delta E\| \leq \mathcal{O}(\varepsilon) \|\mathbf{X}\|^2, \\ \bar{\mathbf{Q}} \bar{\mathbf{R}} &= \mathbf{X} + \Delta \mathbf{D}, \quad \|\Delta \mathbf{D}\| \leq \mathcal{O}(\varepsilon) \|\mathbf{X}\|, \text{ and} \\ \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\| &\leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X})}. \end{aligned}$$

Assume also that $[\bar{\mathbf{Q}}, \bar{\mathcal{R}}] = \text{BCGS-PIPI+}(\mathcal{X}, \text{IO})$ and for all $k \in \{3, \dots, p\}$,

$$\bar{\mathbf{Q}}_{k-1} \bar{\mathcal{R}}_{k-1} = \mathcal{X}_{k-1} + \Delta \mathcal{D}_{k-1}, \quad \|\Delta \mathcal{D}_{k-1}\| \leq \mathcal{O}(\varepsilon) \|\mathcal{X}_{k-1}\|.$$

Then for all $k \in \{1, \dots, p\}$, $\bar{\mathbf{Q}}_k$ satisfies

$$(2.76) \quad \left\| I - \bar{\mathbf{Q}}_k^T \bar{\mathbf{Q}}_k \right\| \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_k)} \leq \mathcal{O}(\varepsilon).$$

Proof. First note that $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ implies that for all $k \in \{1, \dots, p\}$, $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_k) \leq \frac{1}{2}$. Then for the base case, the assumptions on I0 directly give

$$\left\| I - \bar{\mathbf{Q}}_1^T \bar{\mathbf{Q}}_1 \right\| = \left\| I - \bar{\mathbf{Q}}_1^T \bar{\mathbf{Q}}_1 \right\| \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_1)} = \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_1)},$$

which also means that $\omega_1 \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_1)}$. Now assume that (2.76) holds for all $j \in \{1, \dots, k-1\}$, i.e., $\omega_{k-1} \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_{k-1})}$. Noting that the assumption $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ can guarantee (2.57), Theorem 2.6 proves that (2.58) holds for k .

Note that it seems that we prove that $\left\| I - \bar{\mathbf{Q}}_k^T \bar{\mathbf{Q}}_k \right\| \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_{k-1})}$, which is because we omitted $\frac{1}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_{k-1})}$ for simplicity using the assumption $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ in (2.68). Replacing (2.68) in the proof of Theorem 2.6 with

$$(2.77) \quad \left\| \bar{T}_{kk}^{-1} \right\| = \sqrt{\left\| \bar{T}_{kk}^{-1} \right\|^2} \leq \sqrt{\frac{1}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_k)}} \leq \frac{1}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}_k)},$$

we can draw the conclusion. \square

²Note that TSQR can be performed as a single reduction, but the resulting \mathbf{Q} factor is implicitly represented in a tree-based format; if the Householder representation of the \mathbf{Q} factor is desired, a second reduction is required.

A key takeaway from Corollary 2.7, in particular the bound (2.58), is that the LOO of BCGS-PIPI+ depends on the conditioning of \mathcal{X} . We have made the rather artificial assumption that $\mathcal{O}(\varepsilon)\kappa(\mathcal{X}) \leq \frac{1}{2}$. Of course a constant closer to 1 could be used instead, and it would become clear that the constant on $\mathcal{O}(\varepsilon)$ can grow arbitrarily large the closer we let $\kappa(\mathcal{X})$ get to $\frac{1}{\sqrt{\varepsilon}}$. In practice, this edge-case behavior can be quite dramatic, which examples in section 4 demonstrate.

We close this section by bounding both the standard residual (1.2) and Cholesky residual (1.3) for Algorithm 2.3 in uniform precision.

COROLLARY 2.8. *Assume that $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ and that for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$, $[\bar{\mathbf{Q}}, \bar{\mathbf{R}}] = \text{IO}(\mathbf{X})$ satisfy*

$$\bar{\mathbf{Q}}\bar{\mathbf{R}} = \mathbf{X} + \Delta\mathbf{D}, \quad \|\Delta\mathbf{D}\| \leq \mathcal{O}(\varepsilon)\|\mathbf{X}\|, \quad \text{and} \quad \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\| \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)\kappa^2(\mathbf{X})}.$$

Then for $[\bar{\mathbf{Q}}, \bar{\mathbf{R}}] = \text{BCGS-PIPI+}(\mathcal{X}, \text{IO})$ and all $k \in \{1, \dots, p\}$,

$$\bar{\mathbf{Q}}_k \bar{\mathbf{R}}_k = \mathbf{x}_k + \Delta\mathbf{D}_k, \quad \|\Delta\mathbf{D}_k\| \leq \mathcal{O}(\varepsilon)\|\mathbf{x}_k\|.$$

Proof. By the assumption on IO, we have for $k = 1$,

$$\bar{\mathbf{Q}}_1 \bar{\mathbf{R}}_1 = \mathbf{x}_1 + \Delta\mathbf{D}_1, \quad \|\Delta\mathbf{D}_1\| \leq \mathcal{O}(\varepsilon)\|\mathbf{x}_1\|.$$

Now we assume that for all $j \in \{2, \dots, k-1\}$, it holds that

$$(2.78) \quad \bar{\mathbf{Q}}_j \bar{\mathbf{R}}_j = \mathbf{x}_j + \Delta\mathbf{D}_j, \quad \|\Delta\mathbf{D}_j\| \leq \mathcal{O}(\varepsilon)\|\mathbf{x}_j\|.$$

For k , it then follows that

$$(2.79) \quad \begin{aligned} \Delta\mathbf{D}_k &:= [\Delta\mathbf{D}_{k-1} \quad \Delta\mathbf{X}_k] \\ &= \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_k - \mathbf{x}_k \\ &= [\bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{k-1} - \mathbf{x}_{k-1} \quad \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{1:k-1,k} + \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_{kk} - \mathbf{x}_k]. \end{aligned}$$

The first element of (2.79) is taken care of by the induction hypothesis. As for the second element, we treat $\bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{1:k-1,k}$ and $\bar{\mathbf{Q}}_k \bar{\mathbf{R}}_{kk}$ separately. There exists $\Delta\mathbf{R}_k$ such that

$$(2.80) \quad \bar{\mathbf{R}}_{1:k-1,k} = \bar{\mathbf{S}}_{1:k-1,k} + \bar{\mathbf{T}}_{1:k-1,k} \bar{\mathbf{S}}_{kk} + \Delta\mathbf{R}_k,$$

$$(2.81) \quad \|\Delta\mathbf{R}_k\| \leq \mathcal{O}(\varepsilon) (\|\bar{\mathbf{S}}_{1:k-1,k}\| + \|\bar{\mathbf{T}}_{1:k-1,k}\| \|\bar{\mathbf{S}}_{kk}\|) \leq \mathcal{O}(\varepsilon)\|\mathbf{x}_k\|,$$

by (2.49), (2.60), and Lemma 2.5. Combining (2.80), (2.81), and (2.19), it follows that

$$(2.82) \quad \begin{aligned} \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{1:k-1,k} &= \bar{\mathbf{Q}}_{k-1} (\bar{\mathbf{S}}_{1:k-1,k} + \bar{\mathbf{T}}_{1:k-1,k} \bar{\mathbf{S}}_{kk} + \Delta\mathbf{R}_k) \\ &= \bar{\mathbf{Q}}_{k-1} \left(\bar{\mathbf{Q}}_{k-1}^T \mathbf{x}_k + \Delta\mathbf{S}_k + \bar{\mathbf{T}}_{1:k-1,k} \bar{\mathbf{S}}_{kk} + \Delta\mathbf{R}_k \right) \\ &= \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{x}_k + \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{T}}_{1:k-1,k} \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_{k-1} (\Delta\mathbf{S}_k + \Delta\mathbf{R}_k) \end{aligned}$$

with

$$\|\bar{\mathbf{Q}}_{k-1} (\Delta\mathbf{S}_k + \Delta\mathbf{R}_k)\| \leq \mathcal{O}(\varepsilon)\|\mathbf{x}_k\|.$$

From line 12 of Algorithm 2.3, (2.45), and (2.61), we can write

$$(2.83) \quad \bar{\mathbf{R}}_{kk} = \bar{\mathbf{T}}_{kk} \bar{\mathbf{S}}_{kk} + \Delta\mathbf{R}_{kk}, \quad \|\Delta\mathbf{R}_{kk}\| \leq \mathcal{O}(\varepsilon)\|\mathbf{x}_k\|.$$

Plugging (2.23), (2.24), (2.29), and (2.30) into (2.83), the term $\bar{\mathbf{Q}}_k \bar{\mathbf{R}}_{kk}$ can be rewritten as

$$\begin{aligned}
 (2.84) \quad \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_{kk} &= \bar{\mathbf{Q}}_k (\bar{\mathbf{T}}_{kk} \bar{\mathbf{S}}_{kk} + \Delta \mathbf{R}_{kk}) \\
 &= \bar{\mathbf{Q}}_k \bar{\mathbf{T}}_{kk} \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_k \Delta \mathbf{R}_{kk} \\
 &= (\bar{\mathbf{W}}_k + \Delta \mathbf{G}_k^{(2)}) \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_k \Delta \mathbf{R}_{kk} \\
 &= (\bar{\mathbf{U}}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{T}}_{1:k-1,k} + \Delta \mathbf{W}_k + \Delta \mathbf{G}_k^{(2)}) \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_k \Delta \mathbf{R}_{kk} \\
 &= \bar{\mathbf{U}}_k \bar{\mathbf{S}}_{kk} - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{T}}_{1:k-1,k} \bar{\mathbf{S}}_{kk} + (\Delta \mathbf{W}_k + \Delta \mathbf{G}_k^{(2)}) \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_k \Delta \mathbf{R}_{kk} \\
 &= \mathbf{X}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{Q}}_{k-1}^T \mathbf{X}_k + \Delta \mathbf{V}_k + \Delta \mathbf{G}_k^{(1)} - \bar{\mathbf{Q}}_{k-1} \Delta \mathbf{S}_k - \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{T}}_{1:k-1,k} \bar{\mathbf{S}}_{kk} \\
 &\quad + (\Delta \mathbf{W}_k + \Delta \mathbf{G}_k^{(2)}) \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_k \Delta \mathbf{R}_{kk}
 \end{aligned}$$

with

$$\left\| \Delta \mathbf{V}_k + \Delta \mathbf{G}_k^{(1)} - \bar{\mathbf{Q}}_{k-1} \Delta \mathbf{S}_k + (\Delta \mathbf{W}_k + \Delta \mathbf{G}_k^{(2)}) \bar{\mathbf{S}}_{kk} + \bar{\mathbf{Q}}_k \Delta \mathbf{R}_{kk} \right\| \leq \mathcal{O}(\varepsilon) \|\mathbf{X}_k\|.$$

Combining (2.82), (2.84), and the induction hypothesis, we can conclude the proof. \square

Remark 2.9. Invoking Lemma 2.5 in the proof above is not circular logic, as it is valid to use with the induction hypothesis (residual in the $(k-1)$ st step), and (2.60), (2.29), and (2.30) follow from standard rounding-error bounds together with Lemma 2.5.

The bound on the Cholesky residual follows directly from Theorems 2.7 and 2.8.

COROLLARY 2.10. *Assume that $\mathcal{O}(\varepsilon) \kappa^2(\mathbf{X}) \leq \frac{1}{2}$ and for all $k \in \{1, \dots, p\}$, $\bar{\mathbf{Q}}_k$ computed by Algorithm 2.3 satisfies*

$$(2.85) \quad \left\| I - \bar{\mathbf{Q}}_k^T \bar{\mathbf{Q}}_k \right\| \leq \frac{\mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon) \kappa^2(\mathbf{X}_k)} \text{ and}$$

$$(2.86) \quad \mathbf{x}_k + \Delta \mathbf{D}_k = \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_k, \quad \|\Delta \mathbf{D}_k\| \leq \mathcal{O}(\varepsilon) \|\mathbf{x}_k\|.$$

Then for all $k \in \{1, \dots, p\}$,

$$\bar{\mathbf{R}}_k^T \bar{\mathbf{R}}_k = \mathbf{x}_k^T \mathbf{x}_k + \Delta \mathcal{E}_k, \quad \|\Delta \mathcal{E}_k\| \leq \mathcal{O}(\varepsilon) \|\mathbf{x}_k\|^2.$$

Proof. From (2.86) we can directly write

$$\bar{\mathbf{R}}_k^T \bar{\mathbf{Q}}_k^T \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_k = \mathbf{x}_k^T \mathbf{x}_k + \Delta M_k, \quad \|\Delta M_k\| \leq \mathcal{O}(\varepsilon) \|\mathbf{x}_k\|^2,$$

where $\Delta M_k = \mathbf{x}_k^T \Delta \mathbf{D}_k + (\Delta \mathbf{D}_k)^T \mathbf{x}_k + (\Delta \mathbf{D}_k)^2$. Rearranging terms and applying the assumption (2.85) yields

$$\begin{aligned}
 (2.87) \quad \bar{\mathbf{R}}_k^T \bar{\mathbf{R}}_k &= \mathbf{x}_k^T \mathbf{x}_k + \underbrace{\Delta M_k + \bar{\mathbf{R}}_k^T (I - \bar{\mathbf{Q}}_k^T \bar{\mathbf{Q}}_k) \bar{\mathbf{R}}_k}_{=:\Delta \mathcal{E}_k}, \\
 \|\Delta \mathcal{E}_k\| &\leq \mathcal{O}(\varepsilon) \left(\|\mathbf{x}_k\|^2 + \|\bar{\mathbf{R}}_k\|^2 \right).
 \end{aligned}$$

Bounding $\|\bar{\mathbf{R}}_k\|$ follows by multiplying (2.86) by $\bar{\mathbf{Q}}_k^T$ on the left and rearranging terms to arrive at

$$\bar{\mathbf{R}}_k = (I - \bar{\mathbf{Q}}_k^T \bar{\mathbf{Q}}_k) \bar{\mathbf{R}}_k + \bar{\mathbf{Q}}_k^T \mathbf{x}_k + \bar{\mathbf{Q}}_k^T \Delta \mathbf{D}_k.$$

Under the assumptions of this lemma and by (2.18), we find

$$\|\bar{\mathcal{R}}_k\| \leq \frac{1 + \mathcal{O}(\varepsilon)}{1 - \mathcal{O}(\varepsilon)} \|\mathcal{X}_k\| \leq \mathcal{O}(1) \|\mathcal{X}_k\|,$$

which, substituted back into (2.87), completes the proof. \square

Remark 2.11. A version of Algorithm 2.3 has been independently developed and its performance studied in [29, Figure 4(b)]; the authors there refer to it as BCGS-PIP2. We became aware of [29] as we were finishing this manuscript. The authors provide a high-level discussion of the stability analysis of BCGS-PIP2, but they erroneously conflate BCGS-PIPI+ (Algorithm 2.3) with BCGS-PIP+ (Algorithm 2.2); furthermore, it is unclear what IO initializes their BCGS-PIP2. Indeed, as our detailed analysis demonstrates, interchanging the for-loops has a nontrivial effect on attainable guarantees for LOO, as well as on conditions for the first IO. In particular, BCGS-PIPI+ requires a stronger IO than BCGS-PIP+ to maintain LOO, because the first block vector is not reorthogonalized. Furthermore, neither their analysis nor ours extends to the “two-stage” algorithm they present, which is essentially a hybrid of BCGS-PIP+ and BCGS-PIPI+ and allows for a larger block size in the reorthogonalization step. We leave the stability analysis of this hybrid algorithm to future work.

Remark 2.12. An immediate consequence of the analysis in this section is that the proven bounds hold trivially for block size $s = 1$ and, unfortunately, the restriction $\mathcal{O}(\varepsilon) \kappa^2(\mathcal{X}) \leq \frac{1}{2}$ cannot be alleviated; indeed, the size of s only affects the (hidden) constants in $\mathcal{O}(\varepsilon)$. As long as $s \ll m$ (i.e., several orders of magnitude smaller than m), we do not expect it to affect the bounds. Indeed, in section 4, we look at examples with $s = 2$ and $s = 10$ and observe no dependence on block size. In high-performance implementations, practical choices for s depend on the application and hardware but typically remain small for a large number of unknowns.

Remark 2.13. Communication-avoiding Krylov subspace methods like s -step GMRES [17] typically use a block Gram-Schmidt orthogonalization scheme and in each outer iteration generate a block vector of the form $[p_0(A)\mathbf{v} \quad p_1(A)\mathbf{v} \quad \cdots \quad p_{s-1}(A)\mathbf{v}]$, where A is a linear operator, \mathbf{v} is some starting vector, and p_i , $i \in \{0, \dots, s-1\}$, are polynomials of degree i , respectively. The methods considered in this manuscript can be used as block skeletons in s -step GMRES, but the analysis of its backward stability is more complicated than simply applying our results, due to the dual role that s plays. In the present manuscript, s denotes a block partitioning of a *fixed* matrix \mathcal{X} and thus does not affect $\kappa(\mathcal{X})$. Conversely, in s -step GMRES, s determines not only the size of block vectors but also the conditioning of each block (and therefore the entire basis), as each is computed from powers of A . Consequently the backward stability of s -step GMRES is sensitive to the choice of s ; for a complete analysis, see [10], especially Figure 1 therein. In particular, note that employing BCGS-PIPI+ for orthogonalization in s -step GMRES may result in a limited backward error for certain examples, as demonstrated in Figures 8 and 9 of [10].

3. Mixed-precision variants. It is possible to use multiple precisions in the implementations of BCGS-PIP+ and BCGS-PIPI+ without affecting the validity of the general results from section 2. We provide pseudocode for two-precision versions of each as Algorithms 3.2 and Algorithm 3.3, respectively. In the pseudocode, ε_ℓ denotes computing and storing a quantity in low precision with the associated unit roundoff, and ε_h likewise for high precision. In particular, $\varepsilon_\ell \geq \varepsilon_h$.

Algorithm 3.1 $[\mathcal{Q}, \mathcal{R}] = \text{BCGS-PIP}^{\text{MP}}(\mathcal{X}, \text{IO})$.

```

1:  $[\mathbf{Q}_1, R_{11}] = \text{IO}(\mathbf{X}_1)$  ▷ compute and return in  $\varepsilon_\ell$ 
2: for  $k = 2, \dots, p$  do
3:    $\mathcal{R}_{1:k-1,k} = \mathbf{Q}_{k-1}^T \mathbf{X}_k$  ▷ compute and return in  $\varepsilon_\ell$ 
4:    $P_k = \mathbf{X}_k^T \mathbf{X}_k$  ▷ compute and return in  $\varepsilon_h$ 
5:    $R_{kk} = \text{chol}(P_k - \mathcal{R}_{1:k-1,k}^T \mathcal{R}_{1:k-1,k})$  ▷ compute and return in  $\varepsilon_h$ ; cast to  $\varepsilon_\ell$  after line 7
6:    $\mathbf{V}_k = \mathbf{X}_k - \mathbf{Q}_{k-1} \mathcal{R}_{1:k-1,k}$  ▷ compute and return in  $\varepsilon_\ell$ 
7:    $\mathbf{Q}_k = \mathbf{V}_k R_{kk}^{-1}$  ▷ compute each  $s \times s$  block locally in  $\varepsilon_h$ ; return in  $\varepsilon_\ell$ 
8: end for
9: return  $\mathcal{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_p]$ ,  $\mathcal{R} = (R_{ij})$ 

```

Algorithm 3.2 $[\mathcal{Q}, \mathcal{R}] = \text{BCGS-PIP}^{\text{MP}}(\mathcal{X}, \text{IO})$.

```

1:  $[\mathbf{U}, \mathcal{S}] = \text{BCGS-PIP}^{\text{MP}}(\mathcal{X}, \text{IO})$  ▷ computed in mixed; returned in  $\varepsilon_\ell$ 
2:  $[\mathcal{Q}, \mathcal{T}] = \text{BCGS-PIP}^{\text{MP}}(\mathbf{U}, \text{IO})$  ▷ computed in mixed; returned in  $\varepsilon_\ell$ 
3:  $\mathcal{R} = \mathcal{T} \mathcal{S}$ ; ▷ compute and return in  $\varepsilon_\ell$ 
4: return  $\mathcal{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_p]$ ,  $\mathcal{R} = (R_{ij})$ 

```

Algorithm 3.3 $[\mathcal{Q}, \mathcal{R}] = \text{BCGS-PIPI}^{\text{MP}}(\mathcal{X}, \text{IO})$.

```

1:  $[\mathbf{Q}_1, R_{11}] = \text{IO}(\mathbf{X}_1)$  ▷ compute and return in  $\varepsilon_\ell$ 
2: for  $k = 2, \dots, p$  do
3:    $\mathcal{S}_{1:k-1,k} = \mathbf{Q}_{k-1}^T \mathbf{X}_k$  ▷ compute and return in  $\varepsilon_\ell$ 
4:    $\Omega_k = \mathbf{X}_k^T \mathbf{X}_k$  ▷ compute and return in  $\varepsilon_h$ 
5:    $S_{kk} = \text{chol}(\Omega_k - \mathcal{S}_{1:k-1,k}^T \mathcal{S}_{1:k-1,k})$  ▷ compute and return in  $\varepsilon_h$ 
6:    $\mathbf{V}_k = \mathbf{X}_k - \mathbf{Q}_{k-1} \mathcal{S}_{1:k-1,k}$  ▷ compute and return in  $\varepsilon_\ell$ 
7:    $\mathbf{U}_k = \mathbf{V}_k S_{kk}^{-1}$  ▷ compute each  $s \times s$  block locally in  $\varepsilon_h$ ; return in  $\varepsilon_\ell$ 
8:    $\mathcal{T}_{1:k-1,k} = \mathbf{Q}_{k-1}^T \mathbf{U}_k$  ▷ compute and return in  $\varepsilon_\ell$ 
9:    $P_k = \mathbf{U}_k^T \mathbf{U}_k$  ▷ compute and return in  $\varepsilon_h$ 
10:   $T_{kk} = \text{chol}(P_k - \mathcal{T}_{1:k-1,k}^T \mathcal{T}_{1:k-1,k})$  ▷ compute and return in  $\varepsilon_h$ 
11:   $\mathbf{W}_k = \mathbf{U}_k - \mathbf{Q}_{k-1} \mathcal{T}_{1:k-1,k}$  ▷ compute and return in  $\varepsilon_\ell$ 
12:   $\mathbf{Q}_k = \mathbf{W}_k T_{kk}^{-1}$  ▷ compute each  $s \times s$  block locally in  $\varepsilon_h$ ; return in  $\varepsilon_\ell$ 
13:   $\mathcal{R}_{1:k-1,k} = \mathcal{S}_{1:k-1,k} + \mathcal{T}_{1:k-1,k} S_{kk}$  ▷ compute and return in  $\varepsilon_\ell$ 
14:   $R_{kk} = T_{kk} S_{kk}$  ▷ compute in  $\varepsilon_h$ ; return in  $\varepsilon_\ell$ 
15: end for
16: return  $\mathcal{Q} = [\mathbf{Q}_1, \dots, \mathbf{Q}_p]$ ,  $\mathcal{R} = (R_{ij})$ 

```

One motivation for using multiple precisions is to attempt to eliminate the restriction on $\kappa(\mathcal{X})$ for the stability bounds and thereby extend stability guarantees for higher condition numbers; see related work in, e.g., [21, 22]. In a uniform working precision, both BCGS-PIP+ and BCGS-PIPI+ require $\mathcal{O}(\varepsilon) \kappa^2(\mathcal{X}) \leq 1$, which practically translates into $\kappa(\mathcal{X}) \leq \mathcal{O}(10^4)$ or $\kappa(\mathcal{X}) \leq \mathcal{O}(10^8)$, for single or double precisions, respectively. It is natural to consider whether using double the working precision in some parts of the algorithm might alleviate this restriction. At the same time, we do not want to increase communication cost by transmitting high-precision data. Therefore Algorithms 3.2–3.3 are formulated so that the data and solutions \mathcal{Q} and \mathcal{R} are stored in low precision, while high precision is used for the local (i.e., on-node)

computation of operations like $\mathbf{V}^T \mathbf{V}$, Cholesky factorization, and inverting Cholesky factors, with the motivation being that the Pythagorean step, based on Cholesky factorization, is ultimately responsible for the condition number restriction. Note that the inner products in line 3 of BCGS-PIP and lines 3 and 7 of BCGS-PIPI+ are now split across two steps to handle different precisions. However, we still regard these as a single sync point, as the synchronization itself just involves the movement of memory, which can of course be handled in multiple precisions.

At the same time, doubling the precision implies doubling the computational cost and (potentially) the amount of data moved. This overhead is highly dependent on the problem size and it may be negligible in particular cases, such as latency-bound regimes³ and when both precisions are implemented in hardware. When high precision computations are performed locally, such as in line 6 in Algorithm 2.3, the extra overhead may very well be insignificant.

Unfortunately, our intuitive proposals for mixed-precision variants do not achieve the desired stability for either BCGS-PIP+^{MP} or BCGS-PIPI+^{MP}, which we demonstrate in the following sections.

3.1. Two-precision BCGS-PIP and BCGS-PIP+. With Theorem 2.3, we already have generalized bounds on the LOO that will also hold for BCGS-PIP+^{MP}. To determine what the constants $\delta_{US}, \omega_U, \delta_{QT}, \omega_Q$, and δ_{TS} look like in the two-precision case, we need to obtain generalized bounds for BCGS-PIP that do not explicitly rely on a uniform precision but rather express bounds in terms of different constants with subscripts corresponding to the source of error. The following two lemmas take care of this; they generalize [8, Theorems 3.1 and 3.2], respectively. Proofs of Lemmas 3.1 and 3.2 are available in Appendices A.1 and A.2, respectively.

LEMMA 3.1. *Let $\mathbf{x} \in \mathbb{R}^{m \times ps}$ and $[\bar{\mathbf{Q}}, \bar{\mathbf{R}}] = \text{BCGS-PIP}(\mathbf{x}, \mathbf{IO})$ for some \mathbf{IO} . Suppose $\xi \kappa^2(\mathbf{x}) \leq \frac{1}{2}$ and $\rho \kappa^2(\mathbf{x}) \leq \frac{1}{4}$ for constants $\xi, \rho \in (0, 1)$. Furthermore, assume that*

$$(3.1) \quad \bar{\mathbf{R}}^T \bar{\mathbf{R}} = \mathbf{x}^T \mathbf{x} + \Delta \mathcal{E}, \quad \|\Delta \mathcal{E}\| \leq \xi \|\mathbf{x}\|^2, \quad \text{and}$$

$$(3.2) \quad \bar{\mathbf{Q}} \bar{\mathbf{R}} = \mathbf{x} + \Delta \mathcal{D}, \quad \|\Delta \mathcal{D}\| \leq \rho (\|\mathbf{x}\| + \|\bar{\mathbf{Q}}\| \|\bar{\mathbf{R}}\|).$$

Then

$$(3.3) \quad \left\| I - \bar{\mathbf{Q}}^T \bar{\mathbf{Q}} \right\| \leq \frac{(\xi + 11 \cdot \rho) \kappa^2(\mathbf{x})}{1 - \xi \kappa^2(\mathbf{x})};$$

$$(3.4) \quad \|\bar{\mathbf{Q}}\| \leq 3; \quad \text{and}$$

$$(3.5) \quad \|\Delta \mathcal{D}\| \leq 6 \cdot \rho \|\mathbf{x}\|.$$

For the following, we assume that for each $k \in \{2, \dots, p\}$, $\bar{\mathbf{R}}_{k-1}$ and $\bar{\mathbf{Q}}_{k-1}$ are computed by BCGS-PIP and that there exist $\xi_{k-1}, \rho_{k-1} \in (0, 1)$ such that

$$(3.6) \quad \begin{aligned} \bar{\mathbf{R}}_{k-1}^T \bar{\mathbf{R}}_{k-1} &= \mathbf{x}_{k-1}^T \mathbf{x}_{k-1} + \Delta \mathcal{E}_{k-1}, \\ \|\Delta \mathcal{E}_{k-1}\| &\leq \xi_{k-1} \|\mathbf{x}_{k-1}\|^2 \end{aligned}$$

and

$$(3.7) \quad \begin{aligned} \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{k-1} &= \mathbf{x}_{k-1} + \Delta \mathcal{D}_{k-1}, \\ \|\Delta \mathcal{D}_{k-1}\| &\leq \rho_{k-1} (\|\mathbf{x}_{k-1}\| + \|\bar{\mathbf{Q}}_{k-1}\| \|\bar{\mathbf{R}}_{k-1}\|). \end{aligned}$$

³That is, where *latency*, or the time it takes for memory to travel from one node to another across a network or between levels of cache, dominates the runtime of an algorithm.

We can furthermore write the following for intermediate quantities computed by BCGS-PIP, where we omit the explicit dependence on k for readability and each $\delta_* \in (0, 1)$:

$$(3.8) \quad \bar{\mathcal{R}}_{1:k-1,k} = \bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k + \Delta \mathbf{R}_k, \quad \|\Delta \mathbf{R}_k\| \leq \delta_{Q^T X} \|\mathbf{X}_k\|;$$

$$(3.9) \quad \bar{P}_k = \mathbf{X}_k^T \mathbf{X}_k + \Delta P_k, \quad \|\Delta P_k\| \leq \delta_{X^T X} \|\mathbf{X}_k\|^2;$$

$$(3.10) \quad \bar{R}_{kk}^T \bar{R}_{kk} = \bar{P}_k - \bar{\mathcal{R}}_{1:k-1,k}^T \bar{\mathcal{R}}_{1:k-1,k} + \Delta F_k + \Delta C_k,$$

$$(3.11) \quad \|\Delta F_k\| \leq \delta_{R^T R} \|\mathbf{X}_k\|^2, \quad \|\Delta C_k\| \leq \delta_{\text{chol}} \|\mathbf{X}_k\|^2;$$

$$(3.12) \quad \bar{\mathbf{V}}_k = \mathbf{X}_k - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{R}}_{1:k-1,k} + \Delta \mathbf{V}_k, \quad \|\Delta \mathbf{V}_k\| \leq \delta_{QR} \|\mathbf{X}_k\|; \text{ and}$$

$$(3.13) \quad \bar{\mathbf{Q}}_k \bar{R}_k = \bar{\mathbf{V}}_k + \Delta \mathbf{G}_k, \quad \|\Delta \mathbf{G}_k\| \leq \delta_Q \|\bar{\mathbf{Q}}_k\| \|\bar{R}_k\|.$$

Throughout (3.8)–(3.13), we have dropped quadratic error terms and applied Lemma 3.1 to simplify constants (similarly to what we have done in (2.19)–(2.30); the contribution of $\|\bar{\mathcal{Q}}_{k-1}\|$ is absorbed into the constant). Furthermore, ΔF_k denotes the floating-point error from the subtraction of the product $\bar{\mathcal{R}}_{1:k-1,k}^T \bar{\mathcal{R}}_{1:k-1,k}$ from \bar{P}_k , and ΔC_k denotes the error from the Cholesky factorization of that result; note that, similarly to the argument in the proof of [8, Theorem 3.2], $\bar{P}_k - \bar{\mathcal{R}}_{1:k-1,k}^T \bar{\mathcal{R}}_{1:k-1,k}$ should be symmetric positive definite.

LEMMA 3.2. *Let $\mathcal{X} \in \mathbb{R}^{m \times ps}$ and fix $k \in \{2, \dots, p\}$. Assume that (3.6)–(3.13) are satisfied. Then the following hold:*

$$(3.14) \quad \bar{\mathcal{R}}_k^T \bar{\mathcal{R}}_k = \mathcal{X}_k^T \mathcal{X}_k + \Delta \mathcal{E}_k, \quad \|\Delta \mathcal{E}_k\| \leq \xi_k \|\mathcal{X}_k\|^2$$

and

$$(3.15) \quad \bar{\mathbf{Q}}_k \bar{R}_k = \mathcal{X}_k + \Delta \mathcal{D}_k, \quad \|\Delta \mathcal{D}_k\| \leq \rho_k (\|\mathcal{X}_k\| + \|\bar{\mathbf{Q}}_k\| \|\bar{R}_k\|),$$

where

$$\xi_k = \xi_{k-1} + 12 \cdot \rho_{k-1} + 2\sqrt{2} \cdot \delta_{Q^T X} + \delta_{X^T X} + \delta_{\text{chol}} + \delta_{R^T R}$$

and $\rho_k = 6 \cdot \rho_{k-1} + \delta_{QR} + \delta_Q$.

On their own, Lemmas 3.1 and 3.2 do not complete the analysis; they describe the relationship between bounds from one iteration to the next without specifying the precision. Set

$$(3.16) \quad \rho_{\max} := \max_{k \in \{1, \dots, p-1\}} \rho_k \quad \text{and} \quad \xi_{\max} := \max_{k \in \{1, \dots, p-1\}} \xi_k.$$

In the next theorem, we combine these lemmas to obtain bounds on BCGS-PIP^{MP}. A proof of the theorem is available in Appendix A.3.

THEOREM 3.3. *Let $\mathcal{X} \in \mathbb{R}^{m \times ps}$ such that*

$$(3.17) \quad (\rho_{\max} + \rho_p) \kappa^2(\mathcal{X}) \leq \frac{1}{4} \quad \text{and} \quad (\xi_{\max} + \xi_p) \kappa^2(\mathcal{X}) \leq \frac{1}{2}$$

for ρ_p and ξ_p as defined in Lemma 3.2. Suppose $[\bar{\mathbf{Q}}, \bar{R}] = \text{BCGS-PIP}^{\text{MP}}(\mathcal{X}, I0)$, where for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$, $[\bar{\mathbf{Q}}, \bar{R}] = I0(\mathbf{X})$ satisfy

$$(3.18) \quad \bar{R}^T \bar{R} = \mathbf{X}^T \mathbf{X} + \Delta E, \quad \|\Delta E\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathbf{X}\|^2 \quad \text{and}$$

$$(3.19) \quad \bar{\mathbf{Q}} \bar{R} = \mathbf{X} + \Delta D, \quad \|\Delta D\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathbf{X}\|.$$

Then for all $k \in \{1, \dots, p\}$,

$$(3.20) \quad \bar{\mathcal{R}}_k^T \bar{\mathcal{R}}_k = \mathcal{X}_k^T \mathcal{X}_k + \Delta \mathcal{E}_k, \quad \|\Delta \mathcal{E}_k\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathcal{X}_k\|^2;$$

$$(3.21) \quad \bar{\mathcal{Q}}_k \bar{\mathcal{R}}_k = \mathcal{X}_k + \mathcal{D}_k, \quad \|\mathcal{D}_k\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathcal{X}_k\|; \text{ and}$$

$$(3.22) \quad \left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| \leq \mathcal{O}(\varepsilon_\ell) \kappa^2(\mathcal{X}_k).$$

Applying Theorem 3.3 twice to the dual-precision BCGS-PIP+^{MP} shows that

$$(3.23) \quad \delta_{US}, \omega_U, \delta_{QT}, \omega_Q, \delta_{TS} = \mathcal{O}(\varepsilon_\ell),$$

where the constants stem from (2.5)–(2.8). The following corollary is a direct consequence of (3.23) and Theorem 2.3 and demonstrates that BCGS-PIP+^{MP} is dominated by $\mathcal{O}(\varepsilon_\ell)$ error.

COROLLARY 3.4. *Let $\mathcal{X} \in \mathbb{R}^{m \times ps}$ such that $\delta_{\max} \kappa^2(\mathcal{X}) \leq \frac{1}{4}$, where $\delta_{\max} := \max\{\delta_{US}, \omega_U, \delta_{QT}, \omega_Q\}$. Suppose $[\bar{\mathcal{Q}}, \bar{\mathcal{R}}] = \text{BCGS-PIP}^{\text{MP}}(\mathcal{X}, \text{IO})$, where for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$, $[\bar{\mathcal{Q}}, \bar{\mathcal{R}}] = \text{IO}(\mathbf{X})$ satisfy*

$$\begin{aligned} \bar{R}^T \bar{R} &= \mathbf{X}^T \mathbf{X} + \Delta E, \quad \|\Delta E\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathbf{X}\|^2 \text{ and} \\ \bar{\mathcal{Q}} \bar{R} &= \mathbf{X} + \Delta D, \quad \|\Delta D\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathbf{X}\|. \end{aligned}$$

Then for all $k \in \{1, \dots, p\}$,

$$\bar{\mathcal{R}}_k^T \bar{\mathcal{R}}_k = \mathcal{X}_k^T \mathcal{X}_k + \Delta \mathcal{E}_k, \quad \|\Delta \mathcal{E}_k\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathcal{X}_k\|^2;$$

$$\bar{\mathcal{Q}}_k \bar{\mathcal{R}}_k = \mathcal{X}_k + \mathcal{D}_k, \quad \|\mathcal{D}_k\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathcal{X}_k\|; \text{ and}$$

$$\left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| \leq \mathcal{O}(\varepsilon_\ell).$$

3.2. Two-precision BCGS-PIPI+. By standard rounding-error analysis we can show that for (2.19)–(2.30) in BCGS-PIPI+^{MP},

$$\delta_{Q^T X}, \delta_{QS}, \delta_{Q^T U}, \delta_{QT} = \mathcal{O}(\varepsilon_\ell)$$

and

$$\delta_{X^T X}, \delta_{S^T S}, \delta_{\text{cho1}_1}, \delta_U, \delta_{U^T U}, \delta_{T^T T}, \delta_{\text{cho1}_2}, \delta_Q = \mathcal{O}(\varepsilon_h).$$

The following corollary follows directly from these constants and the development in section 2.2, in a manner similar to Corollaries 2.7, 2.8, and 2.10. Consequently, we must conclude that BCGS-PIPI+^{MP} is also dominated by $\mathcal{O}(\varepsilon_\ell)$ error.

COROLLARY 3.5. *Assume that $\mathcal{O}(\varepsilon_\ell) \kappa^2(\mathcal{X}) \leq \frac{1}{2}$ and that for all $\mathbf{X} \in \mathbb{R}^{m \times s}$ with $\kappa(\mathbf{X}) \leq \kappa(\mathcal{X})$, $[\bar{\mathcal{Q}}, \bar{\mathcal{R}}] = \text{IO}(\mathbf{X})$ satisfy*

$$\begin{aligned} \bar{R}^T \bar{R} &= \mathbf{X}^T \mathbf{X} + \Delta E, \quad \|\Delta E\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathbf{X}\|^2, \\ \bar{\mathcal{Q}} \bar{R} &= \mathbf{X} + \Delta D, \quad \|\Delta D\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathbf{X}\|, \text{ and} \end{aligned}$$

$$\left\| I - \bar{\mathcal{Q}}^T \bar{\mathcal{Q}} \right\| \leq \frac{\mathcal{O}(\varepsilon_\ell)}{1 - \mathcal{O}(\varepsilon_\ell) \kappa^2(\mathbf{X})}.$$

Then for $[\bar{\mathcal{Q}}, \bar{\mathcal{R}}] = \text{BCGS-PIPI}^{\text{MP}}(\mathcal{X}, \text{IO})$, the following hold for all $k \in \{1, \dots, p\}$:

$$\bar{\mathcal{R}}_k^T \bar{\mathcal{R}}_k = \mathcal{X}_k^T \mathcal{X}_k + \Delta \mathcal{E}_k, \quad \|\Delta \mathcal{E}_k\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathcal{X}_k\|^2;$$

$$\bar{\mathcal{Q}}_k \bar{\mathcal{R}}_k = \mathcal{X}_k + \Delta \mathcal{D}_k, \quad \|\Delta \mathcal{D}_k\| \leq \mathcal{O}(\varepsilon_\ell) \|\mathcal{X}_k\|;$$

$$\left\| I - \bar{\mathcal{Q}}_k^T \bar{\mathcal{Q}}_k \right\| \leq \frac{\mathcal{O}(\varepsilon_\ell)}{1 - \mathcal{O}(\varepsilon_\ell) \kappa^2(\mathcal{X}_k)} \leq \mathcal{O}(\varepsilon_\ell).$$

Remark 3.6. Our analytical framework is more general than for an arbitrary uniform precision or even two precisions, which we have focused on. Indeed, it is not hard to see that regardless of the number of precisions used in these algorithms, the lowest precision will dominate the LOO and residual bounds. At the same time, numerical experiments in the next section demonstrate that our bounds are rather pessimistic for practical scenarios. More nuanced bounds remain a topic for future work.

4. Numerical experiments.

4.1. BlockStab: A workflow for BGS comparisons. We make use of an updated version of the `BlockStab` code suite [19].⁴ In this version, multiprecision implementations have been added for both the `Advanpix Multiprecision Computing Toolbox`⁵ and the `Symbolic Math Toolbox`,⁶ along with additional low-sync versions of various BGS routines, which are analyzed further in [7]. The procedure for setting up algorithm comparisons is now streamlined to avoid redundant runs and to allow for different choices of Cholesky factorization implementations. Finally, a TeX report is autogenerated with a timestamp, which facilitates sharing experiments with collaborators.

4.2. Comparisons among newly proposed variants. We perform several numerical experiments to study the stability of `BCGS-PIP`, `BCGS-PIP+`, and `BCGS-PIPI+`, using four classes of matrices available in `BlockStab`, namely, `default`, `glued`, `monomial`, and `piled`. Each matrix class is created using dimensional inputs m, p, s , where m denotes the number of rows, p denotes the number of block vectors, and s denotes the number of columns in each block vector. Descriptions of the matrix classes are as follows:

- *default*: built as $\mathcal{X}_t = \mathbf{U}\Sigma_t\mathbf{V}^T$, where $\mathbf{U} \in \mathbb{R}^{m \times ps}$ is orthonormal, $\mathbf{V} \in \mathbb{R}^{ps \times ps}$ is unitary, and $\Sigma_t \in \mathbb{R}^{ps \times ps}$ is diagonal with entries drawn from the logarithmic interval $10^{[-t, 0]}$.
- *glued*: first introduced in [23] and constructed to cause classical Gram-Schmidt to break down. The matrix is initialized with a `default` matrix \mathcal{X}_t ; then each $m \times s$ block vector is multiplied by $\Sigma_r \tilde{V}$, where $\Sigma_r, \tilde{V} \in \mathbb{R}^{s \times s}$, and \tilde{V} is unitary.
- *monomial*: consists of r block vectors $\mathbf{X}_k = [\mathbf{v}_k \quad A\mathbf{v}_k \quad \cdots \quad A^{t-1}\mathbf{v}_k]$, $k \in \{1, \dots, r\}$, where each \mathbf{v}_k is randomly generated from the uniform distribution and normalized, and A is an $m \times m$ diagonal operator having evenly distributed eigenvalues in $(0.1, 10)$. (Note that r is not necessarily equal to p , and likewise $t \neq s$; however, $rt = ps$. Varying r and t allows for generating matrices with different condition numbers.)
- *piled*: formed as $\mathcal{X} = [\mathbf{X}_1 \quad \mathbf{X}_2 \quad \cdots \quad \mathbf{X}_p]$, where \mathbf{X}_1 is a `default` matrix with a small condition number and for $k \in \{2, \dots, p\}$, $\mathbf{X}_k = \mathbf{X}_{k-1} + \mathbf{Z}_k$, where each \mathbf{Z}_k is also a `default` matrix with the same condition number for all k . Toggling the condition numbers of \mathbf{X}_1 and $\{\mathbf{Z}_k\}_{k=2}^p$ controls the overall conditioning of the test matrix.

We set $m = 100$, $p = 10$, and $s = 2$ for `glued` and `default` matrices; $m = 2000$, $p = 120$, $s = 10$ for `monomial` matrices; and $m = 100$, $p = 10$, $s = 5$ for `piled` matrices.

To illustrate the numerical behavior of the algorithms from sections 2 and 3, we plot the LOO (1.1) and relative Cholesky residual (i.e., (1.3) divided by $\|\mathcal{X}\|^2$) of

⁴<https://github.com/katlund/BlockStab/releases/tag/v2.1.2024>.

⁵Version 4.8.3.14460. <https://www.advanpix.com/>.

⁶Version depends on MATLAB version: <https://mathworks.com/products/symbolic.html>.

each algorithm versus the condition number of the matrix; we refer to these plots as κ -plots, as they are relative to the changing condition number $\kappa(\mathcal{X})$. To observe the effects of the choice of IO , we use `HouseQR` and `CholQR`, where a variant of Cholesky factorization is used to bypass the MATLAB `chol` protocol for halting the computation when a matrix loses numerical positive definiteness. One can regard `HouseQR` as a placeholder for `TSQR`, as their numerical behavior is similar, even though the communication properties would differ in practical distributed computing settings.

Double precision ($\varepsilon = 2^{-53} \approx 10^{-16}$) is used for uniform-precision methods. `Advantpix` is used to simulate quadruple precision ($\varepsilon_h = 2^{-113} \approx 10^{-32}$) in multiprecision algorithms, while the low precision is set to double ($\varepsilon_\ell = \varepsilon$).⁷

All numerical tests are run in MATLAB 2022a. Every test is run on a Lenovo ThinkPad E15 Gen 2 with 8 GB memory and AMD Ryzen 5 4500U CPU with Radeon Graphics. The CPU has 6 cores with 384KiB L1 cache, and 3MiB L2 cache at a clockrate of 1 GHz, as well as 8MiB of shared L3 cache. The script `test_bcgs_pip_reortho.m` can be used for regenerating all plots in this section.

Figure 1 illustrates the stability of reorthogonalized variants compared to BCGS-PIP in uniform precision. We see that BCGS-PIP follows an $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X})$ LOO trend until $\kappa(\mathcal{X}) \approx \frac{1}{\sqrt{\varepsilon}} \approx 10^8$, as expected. Meanwhile the reorthogonalized variants reach nearly 10^{-16} LOO, regardless of the choice of IO , until $\kappa(\mathcal{X}) \approx 10^8$. After this point, we observe breakdowns and an increasing LOO for all methods. Missing points for large condition numbers are due to NaN being computed during the Cholesky factorization, resulting from operations like $\frac{\text{Inf}}{\text{Inf}}$ or $\frac{\text{Inf}}{0}$.

The effect of the choice of IO for uniform-precision methods can be seen in Figure 2. As there is no assumption on the LOO of the IO in Corollary 2.1, we observe that `CholQR` works well for BCGS-PIP+. On the other hand, Theorem 2.6 places a LOO restriction on the IO for BCGS-PIPI+. Indeed, `CholQR` does not satisfy the assumption, so the proven LOO and residual bounds are not guaranteed to hold, and we can see BCGS-PIPI+ \circ `CholQR` failing to reach around 10^{-16} LOO even for small condition numbers. Again, both BCGS-PIP+ and BCGS-PIPI+ exhibit an increasing LOO and residual after $\kappa(\mathcal{X}) > 10^8$.




FIG. 1. κ -plots for *glued* matrices. Note: color appears only in the online article.

⁷Users who don't have access to additional toolboxes can still reproduce the trends in our multiprecision results by replacing `'mp_pair': ['double', 'quad']` in https://github.com/katlund/BlockStab/blob/master/tests/configs/bcgs_pip_reortho.json with `'mp_pair': ['single', 'double']`.

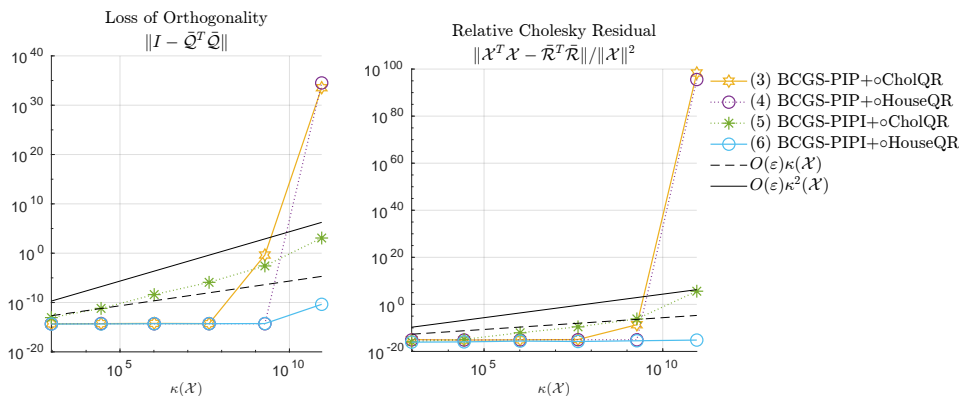


FIG. 2. κ -plots for *monomial* matrices. Note: color appears only in the online article.

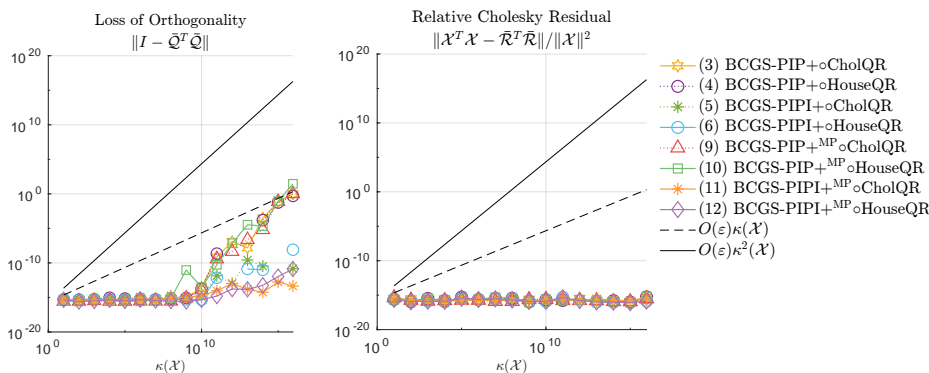


FIG. 3. κ -plots for *default* matrices. Note: color appears only in the online article.

Figure 3 compares two-precision with uniform-precision algorithms. For **default** matrices with or without mixed precision, reorthogonalization keeps the LOO near 10^{-16} as long as $\kappa(\mathcal{X}) \leq 10^8$ and below $\mathcal{O}(\varepsilon)\kappa(\mathcal{X})$ otherwise. Although BCGS-PIPI+^{MP} appears to behave well even for high condition numbers, the **default** matrices are known to be “easy” and may not capture all potential numerical behaviors.

The **glued** matrices are more challenging and reveal worse behavior in Figure 4. After $\kappa(\mathcal{X}) \approx 10^8$, BCGS-PIP+ has stability problems and the relative Cholesky residuals of BCGS-PIPI+ \circ CholQR and BCGS-PIP+ \circ HouseQR become NaN. Mixed precision overcomes this problem for both methods. The LOO of BCGS-PIPI+^{MP} remains below $\mathcal{O}(10^{-16})\kappa(\mathcal{X})$ whereas BCGS-PIP+^{MP} begins to exceed this bound. Notably, the behavior between BCGS-PIPI+^{MP} \circ CholQR and BCGS-PIP+^{MP} \circ HouseQR is very similar, despite the lack of guarantees for CholQR.

Figures 3 and 4 might trick the reader into concluding that BCGS-PIPI+^{MP} \circ CholQR is rather reliable, even without theoretical bounds. The **piled** matrices in Figure 5 should dispel this notion. Neither BCGS-PIPI+ \circ CholQR nor BCGS-PIPI+^{MP} \circ CholQR can attain 10^{-16} LOO for even the smallest condition numbers. And even BCGS-PIPI+^{MP} \circ HouseQR finally manages to lose all orthogonality for $\kappa(\mathcal{X}) < 10^{16}$, which is still numerically nonsingular in double precision. Moreover, we observe rather erratic behavior in the LOO of both variants of BCGS-PIP+^{MP} once $\kappa(\mathcal{X}) \geq 10^8$, which emphasizes the lack of predictability outside of the bounds proven in section 2.




FIG. 4. κ -plots for *glued* matrices. Note: color appears only in the online article.

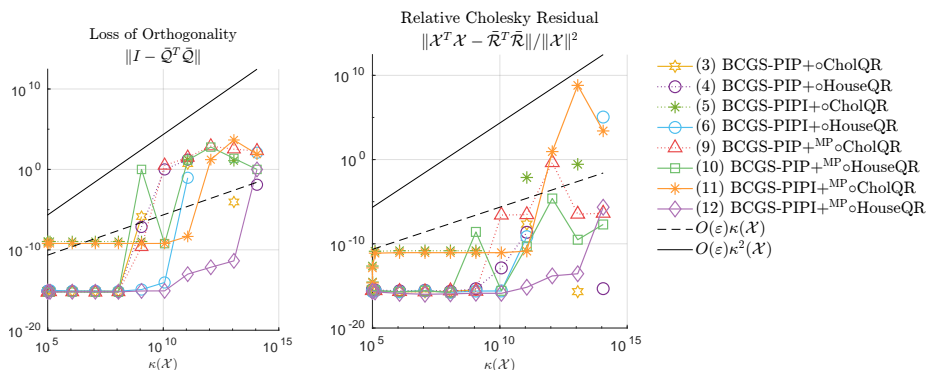


FIG. 5. κ -plots for *piled* matrices. Note: color appears only in the online article.

5. Conclusions. Reorthogonalization is a simple technique for regaining stability in a Gram–Schmidt procedure. We have introduced and examined two reorthogonalized variants of BCGS-PIP, BCGS-PIP+, and BCGS-PIPI+, and demonstrated that both can achieve $\mathcal{O}(\varepsilon)$ LOO under transparent conditions on their intraorthogonalization routines and on the condition number of \mathcal{X} , namely, that $\mathcal{O}(\varepsilon)\kappa^2(\mathcal{X}) \leq 1/2$. We have carried out the analysis in a general enough fashion so that results can be easily extended to multiprecision paradigms, and we have proposed two-precision variants BCGS-PIP+^{MP} and BCGS-PIPI+^{MP}. Numerical experiments verify our findings and demonstrate that despite the lack of theoretical bounds, BCGS-PIPI+^{MP} behaves well for several classes of test matrices and nearly overcomes the restriction on $\kappa(\mathcal{X})$.

At the same time, the restriction on $\kappa(\mathcal{X})$ may not be so problematic when BCGS-PIP+ or BCGS-PIPI+ forms the backbone of a block Arnoldi or GMRES algorithm and can be restarted; see, e.g., [18, 27, 29]. In fact, with the recent modular framework developed for the backward stability analysis of GMRES [6], determining reliable, adaptive restarting heuristics should be quite straightforward. In such scenarios, one usually also has access to preconditioning, which can a priori reduce the conditioning of the basis to be orthogonalized and further improve overall stability. As BCGS-PIP+ or BCGS-PIPI+ both require the same number of sync points as BCGS, but with provably better LOO, they are promising, stable algorithms for a wide variety of applications.

Appendix A. Proofs of the theorems in section 3.

A.1. Proof of Lemma 3.1.

Proof. Assumption (3.1) implies

$$(A.1) \quad \|\bar{\mathcal{R}}\|^2 \leq (1 + \xi) \|\mathcal{X}\|^2,$$

and together with $\xi\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ and the perturbation theory of singular values [15, Corollary 8.6.2], we can derive a bound on $\|\bar{\mathcal{R}}^{-1}\|^2$:

$$(A.2) \quad \|\bar{\mathcal{R}}^{-1}\|^2 = \frac{1}{\sigma_{\min}^2(\bar{\mathcal{R}})} \leq \frac{1}{\sigma_{\min}^2(\mathcal{X})(1 - \xi\kappa^2(\mathcal{X}))}.$$

Multiplying (3.2) on the left by its transpose, then on the left by $\bar{\mathcal{R}}^{-T}$ and on the right by $\bar{\mathcal{R}}^{-1}$ and substituting (3.1) leads to

$$(A.3) \quad \bar{\mathcal{Q}}^T \bar{\mathcal{Q}} = I - \bar{\mathcal{R}}^{-T}(\Delta\mathcal{E} - \mathcal{X}^T \Delta\mathcal{D} - \Delta\mathcal{D}^T \mathcal{X})\bar{\mathcal{R}}^{-1},$$

where we have dropped the quadratic term. Taking the norm and applying bounds (3.1), (3.2), (A.1), and (A.2) to (A.3) leads to a quadratic inequality:

$$(A.4) \quad \begin{aligned} \|\bar{\mathcal{Q}}\|^2 &\leq 1 + \|\bar{\mathcal{R}}^{-1}\|^2 \left(\xi \|\mathcal{X}\|^2 + 2\rho \|\mathcal{X}\| (\|\mathcal{X}\| + \|\bar{\mathcal{Q}}\| \|\bar{\mathcal{R}}\|) \right) \\ &\leq 1 + \frac{(\xi + 2\rho) \|\mathcal{X}\|^2 + 2\rho\sqrt{1+\xi} \|\mathcal{X}\|^2 \|\bar{\mathcal{Q}}\|}{\sigma_{\min}^2(\mathcal{X})(1 - \xi\kappa^2(\mathcal{X}))} \\ &\leq \frac{1 + 2\rho\kappa^2(\mathcal{X})}{1 - \xi\kappa^2(\mathcal{X})} + \frac{2\rho\sqrt{1+\xi}\kappa^2(\mathcal{X})}{1 - \xi\kappa^2(\mathcal{X})} \|\bar{\mathcal{Q}}\|. \end{aligned}$$

Applying the assumptions $\xi\kappa^2(\mathcal{X}) \leq \frac{1}{2}$ and $\rho\kappa^2(\mathcal{X}) \leq \frac{1}{4}$, we see that

$$\frac{1 + 2\rho\kappa^2(\mathcal{X})}{1 - \xi\kappa^2(\mathcal{X})} \leq 3 \text{ and } \frac{2\rho\kappa^2(\mathcal{X})}{1 - \xi\kappa^2(\mathcal{X})} \leq 1,$$

so (A.4) simplifies to

$$(A.5) \quad \|\bar{\mathcal{Q}}\|^2 - \sqrt{1+\xi} \|\bar{\mathcal{Q}}\| - 3 \leq 0,$$

which can be easily solved with the quadratic formula to reveal

$$\|\bar{\mathcal{Q}}\| \leq \frac{1}{2}(\sqrt{1+\xi} + \sqrt{13+\xi}) \leq \frac{1}{2}(\sqrt{2} + \sqrt{14}) \leq 3,$$

thus proving (3.4). Revisiting (A.3), we take the norm of $\|I - \bar{\mathcal{Q}}^T \bar{\mathcal{Q}}\|$ and apply (3.4) along with bounds (3.1), (3.2), (A.1), (A.2) again to arrive at (3.3). Finally, (3.5) follows from (A.1) and (3.4). \square

A.2. Proof of Lemma 3.2.

Proof. Writing $\bar{\mathcal{R}}_k$ as

$$\bar{\mathcal{R}}_k = \begin{bmatrix} \bar{\mathcal{R}}_{k-1} & \bar{\mathcal{R}}_{1:k-1,k} \\ 0 & \bar{R}_k \end{bmatrix}$$

implies

$$(A.6) \quad \bar{\mathcal{R}}_k^T \bar{\mathcal{R}}_k = \begin{bmatrix} \bar{\mathcal{R}}_{k-1}^T \bar{\mathcal{R}}_{k-1} & \bar{\mathcal{R}}_{k-1}^T \bar{\mathcal{R}}_{1:k-1,k} \\ \bar{\mathcal{R}}_{1:k-1,k}^T \bar{\mathcal{R}}_{k-1} & \bar{\mathcal{R}}_{1:k-1,k}^T \bar{\mathcal{R}}_{1:k-1,k} + \bar{R}_k^T \bar{R}_k \end{bmatrix}.$$

The upper diagonal block of (A.6) can be handled directly from (3.6). Lemma 3.1 holds for $k - 1$ via the assumptions (3.6) and (3.7). Then for the off-diagonals, from (3.4) and (3.8), we can write

$$(A.7) \quad \|\bar{\mathcal{R}}_{1:k-1,k}\| \leq \|\bar{\mathcal{Q}}_{k-1}\| \|\mathbf{X}_k\| + \|\Delta \mathbf{R}_k\| \leq (3 + \delta_{Q^T X}) \|\mathbf{X}_k\| \leq 4 \|\mathbf{X}_k\|.$$

Multiplying (3.8) by $\bar{\mathcal{R}}_{k-1}^T$ from the left together with (3.7) leads to

$$(A.8) \quad \begin{aligned} \bar{\mathcal{R}}_{k-1}^T \bar{\mathcal{R}}_{1:k-1,k} &= \bar{\mathcal{R}}_{k-1}^T (\bar{\mathcal{Q}}_{k-1}^T \mathbf{X}_k + \Delta \mathbf{R}_k) \\ &= \mathcal{X}_{k-1}^T \mathbf{X}_k + \Delta \mathcal{E}_{1:k-1,k}, \end{aligned}$$

where $\Delta \mathcal{E}_{1:k-1,k} := \Delta \mathcal{D}_{k-1}^T \mathbf{X}_k + \bar{\mathcal{R}}_{k-1}^T \Delta \mathbf{R}_k$. From (3.6), (3.7), Lemma 3.1, and (3.8) it follows that

$$(A.9) \quad \begin{aligned} \|\Delta \mathcal{E}_{1:k-1,k}\| &\leq 6 \cdot \rho_{k-1} \|\mathcal{X}_{k-1}\| \|\mathbf{X}_k\| + \delta_{Q^T X} \sqrt{1 + \xi_{k-1}} \|\mathcal{X}_{k-1}\| \|\mathbf{X}_k\| \\ &\leq (6 \cdot \rho_{k-1} + \sqrt{2} \cdot \delta_{Q^T X}) \|\mathcal{X}_{k-1}\| \|\mathbf{X}_k\|. \end{aligned}$$

As for the bottom diagonal block of (A.6), we can combine (3.9)–(3.11) and rearrange terms to find

$$(A.10) \quad \begin{aligned} \bar{\mathcal{R}}_{1:k-1,k}^T \bar{\mathcal{R}}_{1:k-1,k} + \bar{R}_k^T \bar{R}_k &= \bar{P}_k + \Delta F_k + \Delta C_k \\ &= \mathbf{X}_k^T \mathbf{X}_k + \underbrace{\Delta P_k + \Delta F_k + \Delta C_k}_{=: \Delta E_k}, \end{aligned}$$

where

$$(A.11) \quad \|\Delta E_k\| \leq (\delta_{X^T X} + \delta_{R^T R} + \delta_{\text{chol}}) \|\mathbf{X}_k\|^2.$$

Combining (3.6), (A.8), (A.9), and (A.11) we can write

$$\bar{\mathcal{R}}_k^T \bar{\mathcal{R}}_k = \underbrace{\begin{bmatrix} \mathcal{X}_{k-1}^T \mathcal{X}_{k-1} & \mathcal{X}_{k-1}^T \mathbf{X}_k \\ \mathbf{X}_k^T \mathcal{X}_{k-1} & \mathbf{X}_k^T \mathbf{X}_k \end{bmatrix}}_{=: \mathcal{X}_k^T \mathcal{X}_k} + \underbrace{\begin{bmatrix} \Delta \mathcal{E}_{k-1} & \Delta \mathcal{E}_{1:k-1,k} \\ \Delta \mathcal{E}_{1:k-1,k}^T & \Delta E_k \end{bmatrix}}_{=: \Delta \mathcal{E}_k},$$

where applying [13, p. 15.50] leads to

$$\begin{aligned} \|\Delta \mathcal{E}_k\| &\leq \|\Delta \mathcal{E}_{k-1}\| + 2 \|\Delta \mathcal{E}_{1:k-1,k}\| + \|\Delta E_k\| \\ &\leq \xi_{k-1} \|\mathcal{X}_{k-1}\|^2 + 2(6 \cdot \rho_{k-1} + \sqrt{2} \cdot \delta_{Q^T X}) \|\mathcal{X}_{k-1}\| \|\mathbf{X}_k\| \\ &\quad + (\delta_{X^T X} + \delta_{\text{chol}} + \delta_{R^T R}) \|\mathbf{X}_k\|^2 \\ &\leq (\xi_{k-1} + 12 \cdot \rho_{k-1} + 2\sqrt{2} \cdot \delta_{Q^T X} + \delta_{X^T X} + \delta_{\text{chol}} + \delta_{R^T R}) \|\mathcal{X}_k\|^2, \end{aligned}$$

which proves (3.14).

To prove (3.15), combining (3.12) and (3.13) yields

$$\bar{\mathcal{Q}}_k \bar{R}_k = \mathbf{X}_k - \bar{\mathcal{Q}}_{k-1} \bar{\mathcal{R}}_{1:k-1,k} + \Delta \mathbf{D}_k,$$

where $\Delta \mathbf{D}_k = \Delta \mathbf{V}_k + \Delta \mathbf{G}_k$ and

$$(A.12) \quad \begin{aligned} \|\Delta \mathbf{D}_k\| &\leq \delta_{QR} \|\mathbf{X}_k\| + \delta_Q \|\bar{\mathbf{Q}}_k\| \|\bar{\mathbf{R}}_k\| \\ &\leq (\delta_{QR} + \delta_Q) (\|\mathbf{X}_k\| + \|\bar{\mathbf{Q}}_k\| \|\bar{\mathbf{R}}_k\|). \end{aligned}$$

Writing

$$\begin{aligned} \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_k &= [\bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{k-1} \quad \bar{\mathbf{Q}}_{k-1} \bar{\mathbf{R}}_{1:k-1,k} + \bar{\mathbf{Q}}_k \bar{\mathbf{R}}_k] \\ &= \underbrace{[\mathbf{X}_{k-1} \quad \mathbf{X}_k]}_{=\mathbf{X}_k} + \underbrace{[\Delta \mathcal{D}_{k-1} \quad \Delta \mathbf{D}_k]}_{=:\Delta \mathcal{D}_k}, \end{aligned}$$

together with (3.7), Lemma 3.1, and (A.12) gives

$$\begin{aligned} \|\Delta \mathcal{D}_k\| &\leq \|\Delta \mathcal{D}_{k-1}\| + \|\Delta \mathbf{D}_k\| \\ &\leq 6 \cdot \rho_{k-1} \|\mathbf{X}_{k-1}\| + (\delta_{QR} + \delta_Q) (\|\mathbf{X}_k\| + \|\bar{\mathbf{Q}}_k\| \|\bar{\mathbf{R}}_k\|) \\ &\leq (6 \cdot \rho_{k-1} + \delta_{QR} + \delta_Q) (\|\mathbf{X}_k\| + \|\bar{\mathbf{Q}}_k\| \|\bar{\mathbf{R}}_k\|), \end{aligned}$$

which completes the proof. \square

A.3. Proof of Theorem 3.3.

Proof. We start with the base case, $k = 1$. By assumptions (3.18) and (3.19), (3.20) and (3.21) follow trivially. Consequently, by setting $\xi_1 = \mathcal{O}(\varepsilon_\ell)$, $\rho_1 = \mathcal{O}(\varepsilon_\ell)$ in (3.11) and (3.7), respectively, and by applying (3.16) and (3.17), it holds that $\xi_1 \kappa^2(\mathbf{X}_1) \leq \frac{1}{2}$ and $\rho_1 \kappa^2(\mathbf{X}_1) \leq \frac{1}{4}$. Then we can apply Lemma 3.1 to conclude (3.22) for $k = 1$.

Now, by following standard rounding-error analysis from [16] (particularly [16, Lemma 6.6, Theorem 8.5, and Theorem 10.3]), we find that for all $k \in \{2, \dots, p\}$, there exist constants $\delta_{Q^T X}, \delta_{QR}, \delta_{X^T X}, \delta_{R^T R}, \delta_{\text{cho1}}, \delta_Q$ such that

$$(A.13) \quad \delta_{Q^T X}, \delta_{QR} = \mathcal{O}(\varepsilon_\ell) \quad \text{and} \quad \delta_{X^T X}, \delta_{R^T R}, \delta_{\text{cho1}}, \delta_Q = \mathcal{O}(\varepsilon_h),$$

and (3.8)–(3.13) hold.

Consider just $k = 2$ for a moment. Then Lemma 3.2 can be applied (because we already have the bound (3.22) for $k = 1$) to conclude

$$\begin{aligned} \bar{\mathbf{R}}_2^T \bar{\mathbf{R}}_2 &= \mathbf{X}_2^T \mathbf{X}_2 + \Delta \mathcal{E}_2, \quad \|\Delta \mathcal{E}_2\| \leq \xi_2 \|\mathbf{X}_2\|^2 \quad \text{and} \\ \bar{\mathbf{Q}}_2 \bar{\mathbf{R}}_2 &= \mathbf{X}_2 + \mathcal{D}_2, \quad \|\mathcal{D}_2\| \leq \rho_2 (\|\mathbf{X}_2\| + \|\bar{\mathbf{Q}}_2\| \|\bar{\mathbf{R}}_2\|), \end{aligned}$$

where $\xi_2 = \mathcal{O}(\varepsilon_\ell)$ and $\rho_2 = \mathcal{O}(\varepsilon_\ell)$ because $\xi_1 = \mathcal{O}(\varepsilon_\ell)$ and the bounds (A.13). Thus (3.20) and (3.21) hold for $k = 2$, and Lemma 3.1 can subsequently be applied to prove (3.22) for $k = 2$.

Proceeding by strong induction and identical logic as for $k = 2$, we can then conclude (3.20)–(3.22) hold for all $k \in \{2, \dots, p\}$. \square

REFERENCES

- [1] G. BALLARD, E. CARSON, J. DEMMEL, M. HOEMMEN, N. KNIGHT, AND O. SCHWARTZ, *Communication lower bounds and optimal algorithms for numerical linear algebra*, Acta Numer. 2011, (2014), pp. 1–155, <https://doi.org/10.1017/S0962492914000038>.
- [2] G. BALLARD, J. DEMMEL, L. GRIGORI, M. JACQUELIN, N. KNIGHT, AND H. NGUYEN, *Reconstructing Householder vectors from Tall-Skinny QR*, J. Parallel Distrib. Comput., 85 (2015), pp. 3–31, <https://doi.org/10.1016/j.jpdc.2015.06.003>.

- [3] J. L. BARLOW, *Reorthogonalized block classical Gram–Schmidt using two Cholesky-based TSQR algorithms*, SIAM J. Matrix Anal. Appl., 45 (2024), pp. 1487–1517, <https://doi.org/10.1137/23M1605387>.
- [4] J. L. BARLOW AND A. SMOKTUNOWICZ, *Reorthogonalized block classical Gram–Schmidt*, Numer. Math., 123 (2013), pp. 395–423, <https://doi.org/10.1007/s00211-012-0496-2>.
- [5] D. BIELICH, J. LANGOU, S. THOMAS, K. ŚWIRYDOWICZ, I. YAMAZAKI, AND E. G. BOMAN, *Low-synch Gram–Schmidt with delayed reorthogonalization for Krylov solvers*, Parallel Comput., 112 (2022), 102940, <https://doi.org/10.1016/j.parco.2022.102940>.
- [6] A. BUTTARI, N. J. HIGHAM, T. MARY, AND B. VIEUBLÉ, *A Modular Framework for the Backward Error Analysis of GMRES*, preprint, hal-04525918, <https://hal.science/hal-04525918> (2024).
- [7] E. CARSON, K. LUND, Y. MA, AND E. OKTAY, *On the Loss of Orthogonality in Low-Synchronization Variants of Reorthogonalized Block Classical Gram–Schmidt*, preprint, arXiv:2408.10109, <https://doi.org/10.48550/arXiv.2408.10109> (2024).
- [8] E. CARSON, K. LUND, AND M. ROZLOŽNÍK, *The stability of block variants of classical Gram–Schmidt*, SIAM J. Matrix Anal. Appl., 42 (2021), pp. 1365–1380, <https://doi.org/10.1137/21M1394424>.
- [9] E. CARSON, K. LUND, M. ROZLOŽNÍK, AND S. THOMAS, *Block Gram–Schmidt algorithms and their stability properties*, Linear Algebra Appl., 638 (2022), pp. 150–195, <https://doi.org/10.1016/j.laa.2021.12.017>.
- [10] E. CARSON AND Y. MA, *On the backward stability of s-step GMRES*, preprint, arXiv:2409.03079, <https://doi.org/10.48550/arXiv.2409.03079> (2024).
- [11] E. C. CARSON, *Communication-Avoiding Krylov Subspace Methods in Theory and Practice*, Ph.D. thesis, Department of Computer Science, University of California at Berkeley, Berkeley, CA, 2015, <http://escholarship.org/uc/item/6r91c407>.
- [12] T. FUKAYA, R. KANNAN, Y. NAKATSUKASA, Y. YAMAMOTO, AND Y. YANAGISAWA, *Shifted Cholesky QR for computing the QR factorization of ill-conditioned matrices*, SIAM J. Sci. Comput., 42 (2020), pp. A477–A503, <https://doi.org/10.1137/18M1218212>.
- [13] S. R. GARCIA AND R. A. HORN, *A Second Course in Linear Algebra*, Cambridge University Press, Cambridge, 2017, <https://doi.org/10.1017/9781316218419>.
- [14] L. GIRAUD, J. LANGOU, M. ROZLOŽNÍK, AND J. VAN DEN ESHOF, *Rounding error analysis of the classical Gram–Schmidt orthogonalization process*, Numer. Math., 101 (2005), pp. 87–100, <https://doi.org/10.1007/s00211-005-0615-4>.
- [15] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 4th ed., Johns Hopkins Stud. Math. Scie., Johns Hopkins University Press, Baltimore, MD, 2013.
- [16] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, 2002.
- [17] M. HOEMMEN, *Communication-Avoiding Krylov Subspace Methods*, Ph.D. thesis, Department of Computer Science, University of California at Berkeley, Berkeley, CA, 2010, <http://www2.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-37.pdf>.
- [18] K. LUND, *Adaptively restarted block Krylov subspace methods with low-synchronization skeletons*, Numer. Algorithms, 93 (2023), pp. 731–764, <https://doi.org/10.1007/s11075-022-01437-1>.
- [19] K. LUND, E. OKTAY, E. CARSON, AND Y. MA, *BlockStab*, 2024, <https://github.com/katlund/BlockStab>.
- [20] D. MORI, Y. YAMAMOTO, AND S. L. ZHANG, *Backward error analysis of the AllReduce algorithm for householder QR decomposition*, Jpn. J. Ind. Appl. Math., 29 (2012), pp. 111–130, <https://doi.org/10.1007/s13160-011-0053-x>.
- [21] E. OKTAY, *Mixed-Precision Computations in Numerical Linear Algebra*, Ph.D. thesis, Faculty of Mathematics and Physics, Charles University, Prague, Czechia, 2024, <http://hdl.handle.net/20.500.11956/191480>.
- [22] E. OKTAY AND E. CARSON, *Using mixed precision in low-synchronization reorthogonalized block classical Gram–Schmidt*, PAMM Proc. Appl. Math. Mech., 23 (2023), e202200060, <https://doi.org/10.1002/pamm.202200060>.
- [23] A. SMOKTUNOWICZ, J. L. BARLOW, AND J. LANGOU, *A note on the error analysis of classical Gram–Schmidt*, Numer. Math., 105 (2006), pp. 299–313, <https://doi.org/10.1007/s00211-006-0042-1>.
- [24] G. W. STEWART, *Block Gram–Schmidt orthogonalization*, SIAM J. Sci. Comput., 31 (2008), pp. 761–775, <https://doi.org/10.1137/070682563>.
- [25] S. THOMAS, E. CARSON, M. ROZLOŽNÍK, A. CARR, AND K. ŚWIRYDOWICZ, *Iterated Gauss–Seidel GMRES*, SIAM J. Sci. Comput., 46 (2024), pp. S254–S279, <https://doi.org/10.1137/22M1491241>.

- [26] L. N. TREFETHEN AND D. I. BAU, *Numerical Linear Algebra*, SIAM, Philadelphia, 1997.
- [27] Z. XU, J. J. ALONSO, AND E. DARVE, *A numerically stable communication-avoiding S-Step GMRES algorithm*, SIAM J. Matrix Anal. Appl., 45 (2024), pp. 2039–2074.
- [28] Y. YAMAMOTO, Y. NAKATSUKASA, Y. YANAGISAWA, AND T. FUKAYA, *Roundoff error analysis of the Cholesky QR2 algorithm*, Electron. Trans. Numer. Anal., 44 (2015), pp. 306–326, <http://www.emis.de/journals/ETNA/vol.44.2015/pp306-326.dir/pp306-326.pdf>.
- [29] I. YAMAZAKI, A. J. HIGGINS, E. G. BOMAN, AND D. B. SZYLD, *Two-Stage Block Orthogonalization to Improve Performance of s-step GMRES*, in 2024 IEEE International Parallel and Distributed Processing Symposium (IPDPS), San Francisco, CA, IEEE, Piscataway, NJ, 2024, pp. 26–37, <https://doi.org/10.1109/IPDPS57955.2024.00012>.
- [30] I. YAMAZAKI, S. THOMAS, M. HOEMMEN, E. G. BOMAN, K. ŚWIRYDOWICZ, AND J. J. ELLIOT, *Low-synchronization orthogonalization schemes for s-step and pipelined Krylov solvers in Trilinos*, in Proceedings of the 2020 SIAM Conference on Parallel Processing for Scientific Computing (PP), SIAM, Philadelphia, 2020, pp. 118–128, <https://doi.org/10.1137/1.9781611976137.11>.
- [31] Q. ZOU, *A flexible block classical Gram–Schmidt skeleton with reorthogonalization*, Numer. Linear Algebra Appl., 30 (2023), e2491, <https://doi.org/10.1002/nla.2491>.