

Randomized block Gram-Schmidt process for solution of linear systems and eigenvalue problems.

Oleg Balabanov* and Laura Grigori†

Abstract

We propose a block version of the randomized Gram-Schmidt process for computing a QR factorization of a matrix. Our algorithm inherits the major properties of its single-vector analogue from [Balabanov and Grigori, 2020] such as higher efficiency than the classical Gram-Schmidt algorithm and stability of the modified Gram-Schmidt algorithm, which can be refined even further by using multi-precision arithmetic. As in [Balabanov and Grigori, 2020], our algorithm has an advantage of performing standard high-dimensional operations, that define the overall computational cost, with a unit roundoff independent of the dominant dimension of the matrix. This unique feature makes the methodology especially useful for large-scale problems computed on low-precision arithmetic architectures.

Block algorithms are advantageous in terms of performance as they are mainly based on cache-friendly matrix-wise operations, and can reduce communication cost in high-performance computing. The block Gram-Schmidt orthogonalization is the key element in the block Arnoldi procedure for the construction of Krylov basis, which in its turn is used in GMRES and Rayleigh-Ritz methods for the solution of linear systems and clustered eigenvalue problems. In this article, we develop randomized versions of these methods, based on the proposed randomized Gram-Schmidt algorithm, and validate them on nontrivial numerical examples.

Key words — Gram-Schmidt, QR factorization, randomization, sketching, numerical stability, rounding errors, loss of orthogonality, multi-precision arithmetic, block Krylov subspace methods, Arnoldi iteration, Petrov-Galerkin, Rayleigh-Ritz, generalized minimal residual method.

1 Introduction

Let $\mathbf{W} \in \mathbb{R}^{n \times m}$ be a matrix with a moderately large number of columns, so that $m \ll n$. We consider column-oriented block Gram-Schmidt (BGS) algorithms for QR factorization of \mathbf{W} :

$$\mathbf{W} = \mathbf{Q}\mathbf{R}$$

where $\mathbf{Q} \in \mathbb{R}^{n \times m}$ has ℓ_2 -orthonormal (or very well-conditioned) columns such that $\text{range}(\mathbf{Q}) = \text{range}(\mathbf{W})$, and $\mathbf{R} \in \mathbb{R}^{m \times m}$ is upper triangular with positive diagonal entries. Block algorithms are usually based on matrix-wise BLAS-3 operations allowing proper exploitation of modern cache-based and high-performance computational architectures. The BGS orthogonalization forms a skeleton for block Krylov subspace methods for solving clustered eigenvalue problems as well as linear systems with multiple right-hand sides. It is also used in s -step, enlarged and other communication-avoiding Krylov subspace methods [12, 14]. Please see [9] and the references therein for an extensive overview of BGS variants, and [2, 17, 18, 22] for the underlying block Krylov methods.

A novel randomized Gram-Schmidt (RGS) process based on the random sketching technique (see [15, 20, 21] and the references therein) was recently introduced in [3]. It requires nearly half as many flops and passes over the data than the classical Gram-Schmidt (CGS) process, and is very stable, in the sense that it yields a very well-conditioned \mathbf{Q} factor if $\text{cond}(\mathbf{W}) = \mathcal{O}(u^{-1})$, where u is the unit roundoff. While the RGS algorithm can be already very beneficial under unique precision computations, there are even more benefits that can be gained by working in two precisions: using

*Sorbonne Université, Inria, CNRS, Université de Paris, Laboratoire Jacques-Louis Lions, Paris, France. Email: oleg.balabanov@inria.fr.

†Sorbonne Université, Inria, CNRS, Université de Paris, Laboratoire Jacques-Louis Lions, Paris, France. Email: laura.grigori@inria.fr.

a coarse unit roundoff for expensive high-dimensional operations and a fine unit roundoff elsewhere. In this case the stability of the RGS algorithm can be guaranteed for the coarse roundoff independent of the dominant dimension n of the matrix. This property can be particularly useful for large-scale computations performed on low-precision arithmetic architectures. Another advantage of the RGS algorithm is the ability of efficient a posteriori certification of the factorization without having to estimate the condition number of a large-scale matrix, or to perform any other expensive large-scale computations.

In this paper we propose a block version (RBGS) of the RGS process. It has similar stability guarantees as its single-vector counterpart, and similar flops count. At the same time, thanks to block paradigm, it is better suited to modern computational architectures. In particular, the major operations in RBGS algorithm can be implemented using cache-efficient BLAS-3 subroutines and a reduced number of synchronizations between distributed processors. In addition, our algorithm is only weakly sensitive to the accuracy of the inter-block orthogonalization, and inherits the ability of efficient certification of the factorization.

Furthermore, we address the application of the RBGS algorithm to Krylov methods. We introduce the block version of the randomized Arnoldi algorithm from [3], which is subsequently used to develop a block randomized GMRES and Rayleigh-Ritz methods for solving linear systems and eigenvalue problems. In exact arithmetic, these methods can be interpreted respectively as minimization of a sketched residual norm or imposing a sketched Galerkin orthogonality condition on the residuals. It is shown that, when stable, the randomized GMRES provides a quasi-optimal solution if the sketching matrix satisfies the ε -embedding property for the computed Krylov basis. The rigorous bounds for the residual error of the randomized Rayleigh-Ritz approximation are derived by reformulating the methodology in terms of projection operators. Furthermore, we also discuss an application of RBGS to other Krylov methods such as the full orthogonalization, and s -step methods.

It is noticed that our RBGS algorithm can be augmented with a Cholesky QR step to provide a QR factorization with Q factor that is not just well-conditioned but ℓ_2 -orthogonal to machine precision. The resulting procedure can be readily applied to the classical GMRES and Rayleigh-Ritz approximations, or any other applications.

This article is organized as follows. The basic notations are explained in Section 1.1. We introduce a general BGS process and particularize it to few classical variants in Section 1.2. In Section 1.3 we present the basic ingredients of the random sketching technique and also extend the results from [3] concerning the effect of sketching on rounding errors in a matrix-matrix product. In Section 2, we present novel RBGS algorithm, stability of which is analyzed in Section 3. Section 4 discusses the application of the methodology to solving clustered eigenvalue problems and linear systems, possibly with multiple right-hand sides. For this we develop the randomized block Arnoldi iteration and the associated GMRES and Rayleigh-Ritz methods. Nontrivial numerical experiments in Section 5 demonstrate the great potential of our methodology. In Section 6 we conclude the work and provide an avenue of future research.

1.1 Preliminaries

Throughout this work, we use the following notations, which is an adaptation of the notations from [3] to block linear algebra. As in [3], we denote Euclidean vectors \mathbf{x} by bold lowercase letters. A matrix composed of several column vectors $\mathbf{x}_1, \dots, \mathbf{x}_k$ is denoted with the associated bold capital letter and a subscript specifying the number of columns, i.e., \mathbf{X}_k . If k is constant, this notation can be simplified to \mathbf{X} . The (i, j) -th block of a block matrix \mathbf{X} is denoted by $\mathbf{X}_{(i,j)}$ i.e, we have

$$\mathbf{X} = \begin{bmatrix} \mathbf{X}_{(1,1)} & \mathbf{X}_{(1,2)} & \cdots & \mathbf{X}_{(1,p)} \\ \mathbf{X}_{(2,1)} & \mathbf{X}_{(2,2)} & \cdots & \mathbf{X}_{(2,p)} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{X}_{(l,1)} & \mathbf{X}_{(l,2)} & & \mathbf{X}_{(l,p)} \end{bmatrix},$$

for some p and l . When $l = 1$ (i.e., matrix \mathbf{X} is partitioned column-wise), the notation $\mathbf{X}_{(1,j)}$ can be simplified to $\mathbf{X}_{(j)}$. The sub-block of \mathbf{X} composed of blocks $\mathbf{X}_{(i,j)}$ with $N_1 \leq i \leq N_2$ and $M_1 \leq j \leq M_2$, is denoted by $\mathbf{X}_{(N_1:N_2, M_1:M_2)}$. We denote the matrix $\mathbf{X}_{(N_1:N_2, M_1:M_1)}$ by simply $\mathbf{X}_{(N_1:N_2, M_1)}$. Moreover, if $l = 1$, then matrix $\mathbf{X}_{(1:1, M_1:M_2)}$ is denoted by $\mathbf{X}_{(M_1:M_2)}$, and matrix $\mathbf{X}_{(M_1:M_1)}$ by $\mathbf{X}_{(M_1)}$. The minimal and the maximal singular values of \mathbf{X} are denoted by $\sigma_{min}(\mathbf{X})$ and $\sigma_{max}(\mathbf{X})$, and the condition number by $\text{cond}(\mathbf{X})$. We let $\langle \cdot, \cdot \rangle$ and $\|\cdot\| = \sigma_{max}(\cdot)$ be the ℓ_2 -inner product and ℓ_2 -norm, respectively. $\|\cdot\|_F$ denotes the Frobenius norm. For two matrices (or vectors) \mathbf{X} and \mathbf{Y} , the relation $\mathbf{X} \leq \mathbf{Y}$ indicates that the entries of \mathbf{X} satisfy $x_{i,j} \leq y_{i,j}$. Furthermore, for a matrix (or a vector) \mathbf{X} , we denote by $|\mathbf{X}|$ the matrix \mathbf{Y} with entries

$y_{i,j} = |x_{i,j}|$. We also let \mathbf{X}^T and \mathbf{X}^\dagger respectively indicate the transpose and the Moore–Penrose inverse of \mathbf{X} . Finally, we let \mathbf{I} be the identity matrix of size appropriate to the expression where this notation is used.

A quantity or an arithmetic expression X computed with finite precision arithmetic is denoted by $\text{fl}(X)$ or \hat{X} .

1.2 Block Gram-Schmidt process

Let matrix $\mathbf{W} \in \mathbb{R}^{n \times m}$ be partitioned into p blocks $\mathbf{W}_{(i)} \in \mathbb{R}^{n \times m_p}$, with $1 \leq i \leq p$, $m = m_p p$:

$$\mathbf{W} = \mathbf{W}_{(1:p)} = [\mathbf{W}_{(1)} \mathbf{W}_{(2)} \dots \mathbf{W}_{(p)}].$$

BGS process proceeds recursively, at iteration i , selecting a new block matrix $\mathbf{W}_{(i)}$ and orthogonalizing it with the previously orthogonalized blocks yielding a matrix $\mathbf{Q}'_{(i)}$, followed by orthogonalization of $\mathbf{Q}'_{(i)}$ itself. This procedure is summarized in Algorithm 1.

Algorithm 1 Block Gram-Schmidt process

Given: $n \times m$ block matrix $\mathbf{W} = \mathbf{W}_{(1:p)}$, $m \leq n$

Output: $n \times m$ factor $\mathbf{Q} = \mathbf{Q}_{(1:p)}$ and $m \times m$ upper triangular factor $\mathbf{R} = \mathbf{R}_{(1:p,1:p)}$.

for $i = 1 : p$ **do**

1. Compute a projection $\mathbf{Q}'_{(i)} = \mathbf{\Pi}^{(i-1)} \mathbf{W}_{(i)}$ (also yielding $\mathbf{R}_{(1:i-1,i)}$).

2. Compute QR factorization $\mathbf{Q}_{(i)} \mathbf{R}_{(i,i)} = \mathbf{Q}'_{(i)}$ with suitable efficient routine.

end for

For standard methods, the projector $\mathbf{\Pi}^{(i-1)}$ in Algorithm 1 is taken as an approximation to ℓ_2 -orthogonal projector onto $\text{range}(\mathbf{Q}_{(1:i-1)})^\perp$. For the classical BGS process (BCGS), one chooses $\mathbf{\Pi}^{(i-1)}$ as

$$\mathbf{\Pi}^{(i-1)} = \mathbf{I} - \mathbf{Q}_{(1:i-1)} \mathbf{Q}_{(1:i-1)}^T.$$

Whereas, for the modified BGS (BMGS) process, we have

$$\mathbf{\Pi}^{(i-1)} = (\mathbf{I} - \mathbf{Q}_{(i-1)}(\mathbf{Q}_{(i-1)})^T)(\mathbf{I} - \mathbf{Q}_{(i-2)}(\mathbf{Q}_{(i-2)})^T) \dots (\mathbf{I} - \mathbf{Q}_{(1)}(\mathbf{Q}_{(1)})^T). \quad (1)$$

In infinite precision arithmetic, these two projectors are equivalent, and are exactly equal to the ℓ_2 -orthogonal projector onto $\text{range}(\mathbf{Q}_{(1:i-1)})^\perp$, since, by construction, $\mathbf{Q}_{(1:i-1)}$ is an orthogonal matrix. However, in the presence of rounding errors these projectors may cause instability. In the first case, the condition number of \mathbf{Q} can grow as $\text{cond}(\mathbf{W})^2$ or even worse and thus requires special treatment [8]. While in the second case, $\text{cond}(\mathbf{Q})$ can grow as $\text{cond}(\mathbf{W}) \max_{1 \leq j \leq p} \text{cond}(\mathbf{W}_{(j)})$ unless Step 2 is unconditionally stable [6, 9]. The stability of these processes can be improved by re-orthogonalization, i.e., by running the inner loop twice. In particular, it can be shown that the BCGS process with re-orthogonalization, here denoted by BCGS2, yields an almost orthogonal Q factor as long as the matrix \mathbf{W} is numerically non-singular [7].

Inter-block orthogonalization is an important step in BGS algorithm. For standard versions, this step can be performed with any suitable efficient and stable routine for (ℓ_2 -)QR factorization of tall and skinny matrices, such as TSQR from [11], a Gram-Schmidt QR applied l times, Cholesky QR applied l times, and others. Typically, such QR factorization takes only a fraction of the overall computational cost.

1.3 Random sketching

In this subsection we recall the basic notions of the random sketching technique, and then provide some useful results from [3].

Let $\Theta \in \mathbb{K}^{k \times n}$ with $k \leq n$ be a sketching matrix such that the associated sketched product $\langle \Theta \cdot, \Theta \cdot \rangle$ approximates well the ℓ_2 -inner product between any two vectors in the subspace (or subspaces) of interest $V \subset \mathbb{R}^n$. Such matrix Θ is referred to as a ε -embedding for V , as defined below.

Definition 1.1. For $\varepsilon < 1$, we say that the sketching matrix Θ is a ε -embedding for subspace V (or matrix \mathbf{V} spanning V), if it satisfies the following relation:

$$\forall \mathbf{x}, \mathbf{y} \in V, \quad |\langle \mathbf{x}, \mathbf{y} \rangle - \langle \Theta \mathbf{x}, \Theta \mathbf{y} \rangle| \leq \varepsilon \|\mathbf{x}\| \|\mathbf{y}\|. \quad (2)$$

Furthermore, we assume that Θ is chosen at random from a certain distribution, such that it satisfies (2) for any fixed d -dimensional subspace with high probability (see Definition 1.2).

Definition 1.2. A random sketching matrix Θ is called a (ε, δ, d) oblivious ℓ_2 -subspace embedding, if for any fixed $V \subset \mathbb{R}^n$ of dimension d , it satisfies

$$\mathbb{P}(\Theta \text{ is a } \varepsilon\text{-embedding for } V) \geq 1 - \delta$$

Corollary 1.3 (Corollary 2.4 in [3]). If $\Theta \in \mathbb{R}^{k \times n}$ is a $(\varepsilon, \delta/n, 1)$ oblivious ℓ_2 -subspace embedding, then with probability at least $1 - \delta$, we have

$$\|\Theta\|_F \leq \sqrt{(1 + \varepsilon)n}.$$

Corollary 1.4 (Corollary 2.2 in [3]). If $\Theta \in \mathbb{R}^{k \times n}$ is an ε -embedding for \mathbf{V} , then the singular values of \mathbf{V} are bounded by

$$(1 + \varepsilon)^{-1/2} \sigma_{\min}(\Theta \mathbf{V}) \leq \sigma_{\min}(\mathbf{V}) \leq \sigma_{\max}(\mathbf{V}) \leq (1 - \varepsilon)^{-1/2} \sigma_{\max}(\Theta \mathbf{V}).$$

In recent years, several distributions of Θ have been proposed that satisfy Definition 1.2 and have a small first dimension k that depends at most logarithmically on the dimension n and probability of failure δ . Among them, the most suitable distribution should be selected depending on the problem and computational architecture. The potential of random sketching is here realized on (rescaled) Rademacher matrices and Subsampled Randomized Hadamard Transform (SRHT). The entries of a Rademacher matrix are i.i.d. random variables satisfying $\mathbb{P}(\theta_{i,j} = 1/\sqrt{k}) = \mathbb{P}(\theta_{i,j} = -1/\sqrt{k}) = 1/2$. Rademacher matrices can be efficiently multiplied by vectors and matrices through the proper exploitation of computational resources such as cache or distributed machines. For n , which is a power of 2, SRHT is defined as a product of diagonal matrix of random signs with Walsh-Hadamard matrix, followed by uniform sub-sampling matrix and scaling factor $1/\sqrt{k}$. For a general n , the SRHT has to be combined with zero padding to make the dimension a power of 2. Random sketching with SRHT can reduce the complexity of an algorithm. Products of SRHT matrices with vectors require only $n \log_2 n$ flops using the fast Walsh-Hadamard transform or $2n \log_2(k + 1)$ flops following the methods in [1]. Furthermore, for both distributions, the usage of a seeded random number generator can allow efficient storage and application of Θ . It follows from [4] (based on [19, 21]) that the rescaled Rademacher distribution, and SRHT (possibly with zero padding) respectively are (ε, δ, d) oblivious ℓ_2 -subspace embeddings, if they have a sufficiently large first dimension: if

$$k \geq 7.87\varepsilon^{-2}(6.9d + \log(1/\delta)) \tag{3a}$$

for Rademacher matrices, or if

$$k \geq 2(\varepsilon^2 - \varepsilon^3/3)^{-1} \left(\sqrt{d} + \sqrt{8 \log(6n/\delta)} \right)^2 \log(3d/\delta) \tag{3b}$$

for SRHT.

1.4 Effect of random sketching on rounding errors

Let us now discuss a very useful result from [3, Section 2.2] characterizing the rounding errors in a sketched matrix-vector product, which can be easily extended to matrix-matrix products. In short, this result states that multiplying Θ by a matrix-vector product $\hat{\mathbf{x}} = \text{fl}(\mathbf{Y}\mathbf{z})$ does not increase the rounding error bound of $\hat{\mathbf{x}}$ by more than a small factor. This property is nothing else than a sketched version of the standard “rule of thumb” for randomized algorithms. It was rigorously proven in [3] for the probabilistic rounding model, [10, Model 4.7]. The extension of this result to matrix-matrix products is provided below.

Let us fix a realization of an oblivious ℓ_2 -subspace embedding $\Theta \in \mathbb{R}^{k \times n}$ of sufficiently large size, and consider a matrix-matrix product

$$\mathbf{X} = \mathbf{Y}\mathbf{Z}, \text{ with } \mathbf{Y} \in \mathbb{R}^{n \times m}, \mathbf{Z} \in \mathbb{R}^{m \times l},$$

computed in finite precision arithmetic with unit roundoff $u < 0.01/m$. The following results can be derived directly from [3, Section 2.2] with the observation that each column of \mathbf{X} is a matrix-vector product: $\mathbf{x}_i = \mathbf{Y}\mathbf{z}_i$, $1 \leq i \leq l$.

We have,

$$|\mathbf{X} - \hat{\mathbf{X}}| \leq \mathbf{U}, \tag{4}$$

for some matrix \mathbf{U} describing the “worse-case scenario” rounding error. In general, the standard analysis (e.g., see [13]) gives the bound

$$|\mathbf{X} - \widehat{\mathbf{X}}| \leq \frac{mu}{1 - mu} |\mathbf{Y}||\mathbf{Z}| \leq 1.02mu |\mathbf{Y}||\mathbf{Z}|, \quad (5)$$

which means that $\mathbf{U} = 1.02mu |\mathbf{Y}||\mathbf{Z}|$ satisfies (4). Furthermore, as in [3], if \mathbf{YZ} represents a sum of a matrix-matrix product with a matrix, i.e., $\mathbf{YZ} = \mathbf{Y}'\mathbf{Z}' + \mathbf{H}$, then one can take

$$\mathbf{U} = 1.02u(|\mathbf{H}| + m|\mathbf{Y}'||\mathbf{Z}'|).$$

Let us now address bounding the norm of the rounding error after sketching $\widehat{\mathbf{X}}$. We have the following “worse-case scenario” bound

$$\|\Theta(\mathbf{X} - \widehat{\mathbf{X}})\|_{\mathbf{F}} \leq \|\Theta\| \|\mathbf{X} - \widehat{\mathbf{X}}\|_{\mathbf{F}} \leq \|\Theta\| \|\mathbf{U}\|_{\mathbf{F}}, \quad (6)$$

which, combined with Corollary 1.3, implies that if Θ is $(\varepsilon, \delta/n, 1)$ oblivious ℓ_2 -subspace embedding, then $\|\Theta(\mathbf{X} - \widehat{\mathbf{X}})\|_{\mathbf{F}} \leq \sqrt{1 + \varepsilon} \sqrt{n} \|\mathbf{U}\|_{\mathbf{F}}$ holds with probability at least $1 - \delta$. The next important point is that, as was argued in [3], this bound is pessimistic and can be improved by a factor of $\mathcal{O}(\sqrt{n})$ by exploiting statistical properties of rounding errors. We summarize the results from [3, Theorem 2.5 and Corollary 2.6] below.

Corollary 1.5. *Consider probabilistic rounding model, where the rounding errors due to each elementary arithmetic operation are random variables possibly depending on each other, but are independently centered. Furthermore assume that the errors are bounded so that, it holds*

$$|\mathbf{X} - \widehat{\mathbf{X}}| \leq \mathbf{U}.$$

If Θ is a $(\varepsilon/4, l^{-1} \binom{n}{d}^{-1} \delta, d)$ oblivious ℓ_2 -subspace embedding, with $d = 4.2c^{-1} \log(4/\delta)$, where $c \leq 1$ is some universal constant, then

$$\|\Theta(\mathbf{X} - \widehat{\mathbf{X}})\|_{\mathbf{F}} \leq \sqrt{1 + \varepsilon} \|\mathbf{U}\|_{\mathbf{F}} \quad (7)$$

holds with probability at least $1 - 2\delta$.

By using (3) we deduce that for $l \leq n$ the relation (7) is satisfied with probability at least $1 - 2\delta$ if Θ is a Rademacher matrix with $\mathcal{O}(\log(n) \log(1/\delta))$ rows or P-SRHT matrix with $\mathcal{O}(\log^2(n) \log^2(1/\delta))$ rows. This result implies that in practice, the matrix-vector products in randomized algorithms can be performed with a unit roundoff independent of the (high) dimension n .

2 Randomized block Gram-Schmidt process

Consider BGS algorithms with projectors $\Pi^{(i-1)}$ that respect the relation

$$\Pi^{(i-1)} \mathbf{W}_{(i)} = \mathbf{W}_{(i)} - \mathbf{Q}_{(1:i-1)} \mathbf{X},$$

where $\mathbf{X} = \mathbf{R}_{(1:i-1, i)}$ is computed from $\mathbf{W}^{(i)}$ and $\mathbf{Q}_{(1:i-1)}$. Standard algorithms take \mathbf{X} as an approximate solution to the following minimization problem:

$$\min_{\mathbf{Y}} \|\mathbf{Q}_{(1:i-1)} \mathbf{Y} - \mathbf{W}_{(i)}\|_{\mathbf{F}}. \quad (8)$$

For instance, BCGS algorithm approximates the exact solution to (8) by $\mathbf{Q}_{(1:i-1)}^{\mathbf{T}} \mathbf{W}_{(i)}$, whereas BMGS projector (1) improves this approximation under finite precision arithmetic, though it may cause a computational overhead in terms of inter-processor communication and operation with cache/RAM.

As proposed in [3], a reduction of the computational cost and/or improvement of stability of Gram-Schmidt process can be obtained with a projector that gives a \mathbf{Q} factor orthonormal with respect to $\langle \Theta \cdot, \Theta \cdot \rangle$ instead of the ℓ_2 -inner product as in standard methods. In our case, this corresponds to taking \mathbf{X} as an approximate solution to

$$\min_{\mathbf{Y}} \|\Theta \mathbf{Q}_{(1:i-1)} \mathbf{Y} - \Theta \mathbf{W}_{(i)}\|_{\mathbf{F}}, \quad (9)$$

which is a k -dimensional block least-squares problem. Since the sketching dimension satisfies $k \ll n$, hence a very accurate solution to (9) should be more efficient to compute than any, even the BCGS (i.e, the cheapest) solution to (8). Furthermore, with the right choice of random sketching matrices, the precomputation of the sketches of $\mathbf{Q}_{(1:i-1)}$ and $\mathbf{W}_{(i)}$ should also have only a minor cost compared to that of standard high-dimensional operations. If (9) is very small, its solution can be obtained with a direct solver based, for example, on Householder transformations or Givens rotations, requiring a cubic complexity. If the problem has a moderate size, so that the direct solvers have a considerable computational cost, one can use the fact that matrix $\Theta \mathbf{Q}_{(1:i-1)}$ is almost orthonormal, which implies applicability of iterative solvers requiring a quadratic complexity. Several such solvers are provided in Section 2.1. For the QR orthogonalization with respect to $\langle \Theta \cdot, \Theta \cdot \rangle$ we have to also use the corresponding inter-block orthogonalization routine in Step 2 of Algorithm 1, providing a Q factor orthonormal with respect to the sketched inner product. An efficient procedure for this task can be based on a standard ℓ_2 -orthogonalization followed by a generalized Cholesky QR factorization (see Section 2.2). On the other hand, this task can also be performed with the single-vector randomized Gram-Schmidt algorithm from [3]. The RBGS process using the random sketching projector is depicted in Algorithm 2.

Algorithm 2 Randomized block Gram-Schmidt algorithm (RBGS)

Given: $n \times m$ block matrix $\mathbf{W} = \mathbf{W}_{(1:p)}$, and $k \times n$ matrix Θ , $m \leq k \ll n$.
Output: $n \times m$ factor $\mathbf{Q} = \mathbf{Q}_{(1:p)}$ and $m \times m$ upper triangular factor $\mathbf{R} = \mathbf{R}_{(1:p,1:p)}$.
for $i = 1 : p$ **do**
 1. Sketch $\mathbf{W}_{(i)}$: $\mathbf{P}_{(i)} = \Theta \mathbf{W}_{(i)}$. # macheps: u_{fine}
 2. Solve small block least-squares problem:
 $\mathbf{R}_{(1:i-1,i)} = \arg \min_{\mathbf{Y}} \|\mathbf{S}_{(1:i-1)} \mathbf{Y} - \mathbf{P}_{(i)}\|_{\mathbb{F}}$. # macheps: u_{fine}
 3. Compute projection of $\mathbf{W}_{(i)}$: $\mathbf{Q}'_{(i)} = \mathbf{W}_{(i)} - \mathbf{Q}_{(1:i-1)} \mathbf{R}_{(1:i-1,i)}$. # macheps: u_{crs}
 4–5. Compute QR factorization $\mathbf{Q}_{(i)} \mathbf{R}_{(i,i)} = \mathbf{Q}'_{(i)}$ with $\mathbf{Q}_{(i)}$ orthogonal with respect to $\langle \Theta \cdot, \Theta \cdot \rangle$. # macheps: u_{fine}
 Compute sketch of $\mathbf{Q}_{(i)}$: $\mathbf{S}_{(i)} = \Theta \mathbf{Q}_{(i)}$. # macheps: u_{fine}
 6. (Optional) compute $\mathbf{S}'_{(i)} = \Theta \mathbf{Q}'_{(i)}$, $\|\mathbf{I} - \mathbf{S}_{(i)}^T \mathbf{S}_{(i)}\|_{\mathbb{F}}$ and $\frac{\|\mathbf{S}'_{(i)} - \mathbf{S}_{(i)} \mathbf{R}_{(i,i)}\|_{\mathbb{F}}}{\|\mathbf{S}'_{(i)}\|_{\mathbb{F}}}$.
end for
 7. (Optional) compute $\Delta^{(p)} = \|\mathbf{I} - \mathbf{S}_{(1:p)}^T \mathbf{S}_{(1:p)}\|_{\mathbb{F}}$ and $\tilde{\Delta}^{(p)} = \frac{\|\mathbf{P}_{(1:p)} - \mathbf{S}_{(1:p)} \mathbf{R}_{(1:p,1:p)}\|_{\mathbb{F}}}{\|\mathbf{P}_{(1:p)}\|_{\mathbb{F}}}$. # macheps: u_{fine}

At the iteration $i = 1$ in Algorithm 2 we used the notation that $\mathbf{R}_{(1:i-1,i)}$ is a 0 by m_p matrix and $\mathbf{Q}_{(1:i-1)}$ is a n by 0 matrix, so that $\mathbf{Q}'_{(i)} = \mathbf{W}_{(i)}$. Algorithm 2 is presented under multi-precision arithmetic using two unit roundoffs, like its single-vector counterpart from [3]. The working precision is represented by a coarse roundoff u_{crs} . It is used for standard high-dimensional operations in Step 3, which determine the overall computational cost. All other (inexpensive) operations in Algorithm 2 are computed with a fine unit roundoff u_{fine} , $u_{fine} \leq u_{crs}$. It is shown in Section 3 that Algorithm 2 is stable if $u_{crs} \leq \mathcal{O}(\text{cond}(\mathbf{W})m^{-2})$, which is a very mild condition on u_{crs} . The fact that this bound is independent of the high dimension n explains the potential of our methodology for large-scale problems computed on low-precision arithmetic architectures. Furthermore, according to our numerical experiments, the RBGS algorithm can be sufficiently stable even when u_{crs} is larger than $\mathcal{O}(\text{cond}(\mathbf{W})m^{-2})$. In such cases, the stability of the algorithm can be certified by a posteriori bounds given by the quantities $\Delta^{(p)}$, $\tilde{\Delta}^{(p)}$, etc., computed in (optional) Steps 6 and 7. These bounds provide stability certification if $u_{crs} = \mathcal{O}(m^{-3/2})$, which is a milder condition than the one for a priori guarantees, and is independent not only of n but also $\text{cond}(\mathbf{W})$.

The stability guarantees of Algorithm 2 executed in unique precision can be obtained directly from the analysis of the multi-precision algorithm by taking $u_{fine} = u$ and $u_{crs} = F(m, n)u$, where $F(m, n)$ is some polynomial of low degree.

In terms of performance, the SRHT-based RBGS algorithm requires about half the flops and data passes of the cheapest standard block Gram-Schmidt algorithm, which is BCGS. Moreover, its computational cost is defined by p well-parallelizable BLAS-3 operations. As in BCGS, in RBGS only one global synchronization between distributed processors is required per iteration, not counting the cost of inter-block QR.

Remark 2.1. *If necessary, the output of the RBGS algorithm can be post-processed with a Cholesky QR to provide a QR factorization with Q factor, which is not only well-conditioned but is ℓ_2 -orthonormal up to machine precision. More specifically, we can compute an upper-triangular matrix \mathbf{R}' such that $\mathbf{R}'^T \mathbf{R}' = \mathbf{Q}^T \mathbf{Q}$ with Cholesky decomposition, and*

consider

$$\mathbf{Q} \leftarrow \mathbf{Q}\mathbf{R}'^{-1} \text{ and } \mathbf{R} \leftarrow \mathbf{R}'\mathbf{R}.$$

The numerical stability of this procedure follows directly from the fact that the \mathbf{Q} factor produced by the RBGS algorithm (and therefore \mathbf{R}') is very well-conditioned. Its computational cost is dominated by the computation of $\mathbf{Q}^T\mathbf{Q}$ and, possibly, $\mathbf{Q}\mathbf{R}'^{-1}$. In terms of flops, the computational cost of computing $\mathbf{Q}^T\mathbf{Q}$ is similar to the cost of the entire RBGS algorithm. Though as a single BLAS-3 operation, it is better suited for cache-based and parallel computing. The computation of $\mathbf{Q}\mathbf{R}'^{-1}$ can be omitted when the application allows operating with the \mathbf{Q} factor in an implicit form, for instance, in Arnoldi iteration or Rayleigh-Ritz algorithm. Otherwise, this product can be computed by well-parallelizable forward substitution, or even a direct inversion allowing a second BLAS-3 multiplication. In this case, the computation and application of \mathbf{R}'^{-1} should not cause any stability issues, since \mathbf{R}' is a very well-conditioned matrix.

2.1 Solution of block least-squares problem in Step 2

As already pointed out, the stability of Algorithm 2 strongly depends on the stability of the least-squares solver used in Step 2. In particular, in our analysis in Section 3 we require the solution $\widehat{\mathbf{X}} = \widehat{\mathbf{R}}_{(1:i-1,i)}$ to satisfy the following backward-stability condition.

Assumptions 2.2 (Backward-stability of Step 2 of RBGS).

$$\widehat{\mathbf{X}} = \arg \min_{\mathbf{Y}} \left\| (\widehat{\mathbf{S}}_{(1:i-1)} + \Delta\mathbf{S}^{(i-1)})\mathbf{Y} - (\widehat{\mathbf{P}}_{(i)} + \Delta\mathbf{P}^{(i)}) \right\|_{\mathbf{F}}$$

for some matrices $\Delta\mathbf{S}^{(i-1)}$ and $\Delta\mathbf{P}^{(i)}$ such that

$$\|\Delta\mathbf{S}^{(i-1)}\|_{\mathbf{F}} \leq 0.01u_{crs}\|\widehat{\mathbf{S}}_{(1:i-1)}\|, \quad \|\Delta\mathbf{P}^{(i)}\|_{\mathbf{F}} \leq 0.01u_{crs}\|\widehat{\mathbf{P}}_{(i)}\|.$$

This condition can be met by standard direct solvers based on Householder transformation or Givens rotations and a sufficiently large gap between u_{crs} and u_{fine} , which follows from [13, Theorems 8.5, 19.10 and 20.3] and their proofs. However, direct solvers require a cubic complexity and can become too expensive even for relatively small m and k . In such cases, one has to appeal to *iterative* methods that can exploit the approximate orthogonality of $\mathbf{S}_{(1:i-1)}$ to speedup the computations, as is discussed below.

2.1.1 Richardson iterations

Algorithm 3 describes, perhaps, the simplest iterative method for obtaining a solution in Step 2. When the least-squares problem is seen as orthogonalization of columns of $\mathbf{P}_{(i)}$ to $\mathbf{S}_{(1:i-1)}$, this algorithm can be interpreted as nothing more than the classical Gram-Schmidt (CGS) algorithm with l re-orthogonalizations. In more general terms, Algorithm 3 is Richardson method applied to the normal system of equations

$$\left(\mathbf{S}_{(1:i-1)}^T \mathbf{S}_{(1:i-1)} \right) \mathbf{X} = \mathbf{S}_{(1:i-1)}^T \mathbf{P}_{(i)}. \quad (10)$$

Assuming that $l = \mathcal{O}(1)$, it requires $\mathcal{O}(m_p m k)$ flops for computing $\mathbf{R}_{(1:i-1,i)} = \mathbf{X}$, which makes the cost of Step 2 negligible in comparison to other costs of the RBGS algorithm and in particular to the Step 3 of RBGS, which has complexity $\mathcal{O}(m_p m n)$. Finally, we note that the RBGS algorithm based on $l = 1$ Richardson iteration in Step 2 is equivalent to the BCGS process for the orthogonalization with respect to the sketched inner product. In general, there is a strong connection between the RBGS algorithm using l Richardson iterations and BCGS using l re-orthogonalizations.

Let $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_{m_p}]$ be the exact solution to (10) and $\mathbf{R}_{(1:i-1,i)} = \mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_{m_p}]$ be the solution computed with Algorithm 3. In exact arithmetic, we have:

$$\|\mathbf{x}_j - \mathbf{y}_j\| \leq \|\mathbf{I} - \mathbf{S}_{(1:i-1)}^T \mathbf{S}_{(1:i-1)}\|^l \|\mathbf{y}_j\| = (\Delta^{(i-1)})^l \|\mathbf{y}_j\|, \quad 1 \leq j \leq m_p,$$

where $\Delta^{(i-1)} = \|\mathbf{I} - \mathbf{S}_{(1:i-1)}^T \mathbf{S}_{(1:i-1)}\|_{\mathbf{F}}$ measures the orthogonality of the sketch of \mathbf{Q} factor. This result can be easily extended to finite precision arithmetic.

Algorithm 3 Step 2 of RBGS, using CGS with l re-orthogonalizations

Given: $\mathbf{P}^{(i)}, \mathbf{S}_{(1:i-1)}$
Output: $\mathbf{R}_{(1:i-1,i)} = \mathbf{X}$
1. Set $\mathbf{X} = \mathbf{0}_{(1:i-1,1)}$.
for $1 : l$ **do**
 2. $\mathbf{X} \leftarrow \mathbf{X} + \mathbf{S}_{(1:i-1)}^T (\mathbf{P}^{(i)} - \mathbf{S}_{(1:i-1)} \mathbf{X})$.
end for

2.1.2 Iterations of the modified GS process

Another way to compute the solution in Step 2 of Algorithm 2 is to use the modified GS (MGS) process, applied l times. This algorithm can provide an accurate solution in fewer iterations than Algorithm 3, though it can have a higher cost per iteration from a performance standpoint. This drawback can be remedied by appealing to the block version of MGS depicted in Algorithm 4. The RBGS algorithm based on Algorithm 4 can be linked to BMGS algorithm, in which the ℓ_2 -inner products are replaced by $\langle \Theta \cdot, \Theta \cdot \rangle$. In particular, for the case $l = 1$ the two algorithms are essentially equivalent. The numerical analysis of Algorithm 4 is beyond the scope of this manuscript.

Algorithm 4 Step 2 of RBGS, using BMGS with l re-orthogonalizations

Given: $\mathbf{P}^{(i)}, \mathbf{S}_{(1:i-1)}$
Output: $\mathbf{R}_{(1:i-1,i)} = \mathbf{X} = [\mathbf{X}_{(1)}, \dots, \mathbf{X}_{(i-1)}]$
1. Set $\mathbf{X} = \mathbf{0}_{(1:i-1,1)}$.
for $1 : l$ **do**
 for $j = 1 : i - 1$ **do**
 2. $\mathbf{P}^{(i)} \leftarrow \mathbf{P}^{(i)} - \mathbf{S}_{(j)} \mathbf{X}_{(j)}$.
 3. $\mathbf{X}_{(j)} \leftarrow \mathbf{X}_{(j)} + \mathbf{S}_{(j)}^T \mathbf{P}^{(i)}$.
 end for
end for

2.1.3 Other methods

In principle, we can compute \mathbf{X} from the normal equation (10) using any suitable iterative method such as Conjugate Gradient or GMRES. In this case, the normal matrix can be operated with as an implicit map outputting products with vectors using $\mathcal{O}(m_p mk)$ flops, making the cost of the iterative method similar to that of a Richardson iteration.

2.2 Inter-block QR factorization in Steps 4–5

Let us now provide three ways of efficient and stable inter-block orthogonalization in Steps 4–5 of the RBGS algorithm.

2.2.1 Single-vector RGS algorithm

First, the inter-block orthogonalization can be readily performed with the single-vector RGS algorithm from [3] computed with unique roundoff $u = u_{fine} = F(m, n)u_{crs}$, where $F(m, n)$ is a low-degree polynomial. This algorithm provides a QR factorization that satisfies (11a), which follows directly by construction. Furthermore, it should also satisfy (11b), since $\|\Theta\|_F \leq \sqrt{2n}$ holds with high probability (according to Corollary 1.3).

$$|\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)}| \leq 0.1u_{crs} |\widehat{\mathbf{Q}}_{(i)}| |\widehat{\mathbf{R}}_{(i,i)}| \quad (11a)$$

$$\|\Theta(\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)})\|_F \leq 0.1u_{crs} \|\widehat{\mathbf{Q}}_{(i)}\| \|\widehat{\mathbf{R}}_{(i,i)}\| \quad (11b)$$

As shown in [3], if Θ is ε -embedding for $\widehat{\mathbf{Q}}'_{(i)}$, then RGS algorithm satisfies

$$1 - 0.1u_{crs} \text{cond}(\widehat{\mathbf{Q}}'_{(i)}) \leq \sigma_{\min}(\Theta \widehat{\mathbf{Q}}_{(i)}) \leq \sigma_{\max}(\Theta \widehat{\mathbf{Q}}_{(i)}) \leq 1 + 0.1u_{crs} \text{cond}(\widehat{\mathbf{Q}}'_{(i)}). \quad (12)$$

Moreover, in this case Θ is guaranteed to be an ε' -embedding for $\widehat{\mathbf{Q}}_{(i)}$, with $\varepsilon' = 2\varepsilon + u_{crs}$. Though, in practice, this property holds for smaller values ε' , say, $\varepsilon + u_{crs}$.

According to [3], the single-vector RGS algorithm applied to an inter-block has the cost of $2m_p^2n$ flops and $m_p n$ bytes of memory, which can be considered negligible compared to the complexity and memory consumption of other computations such as the standard high-dimensional operations in Algorithm 2. The RGS algorithm, however, requires m_p global synchronizations between distributed processors, which can dominate the computational costs in cache-based or massively parallel architectures. In such cases, it is necessary to appeal to different approaches for the inter-block QR factorization, some of which are described below.

2.2.2 Sketched CholQR factorization

Another way to perform inter-block QR factorization of matrix $\mathbf{Q}'_{(i)}$ with respect to the sketched inner product is to appeal to the following sketched version of the Cholesky QR (CholQR). We can first compute the R factor $\mathbf{R}_{(i,i)}$ by performing a (ℓ_2 -)QR factorization of a small matrix $\mathbf{S}'_{(i)} = \Theta \mathbf{Q}'_{(i)}$ with any suitable routine, possibly relying on Cholesky factorization of $(\mathbf{S}'_{(i)})^T \mathbf{S}'_{(i)}$. Then the Q factor is retrieved by computing $\mathbf{Q}_{(i)} = \mathbf{Q}'_{(i)} (\mathbf{R}_{(i,i)})^{-1}$ with forward substitution (see Algorithm 5). While sketched Cholesky QR is efficient in terms of communication, it can entail stability issues. They can be alleviated by re-orthogonalization (i.e, by running the algorithm l times), or by combination with modern efficient routines for ℓ_2 -orthogonalization.

Algorithm 5 Sketched CholQR

Given: $\mathbf{Q}'_{(i)}, \Theta$.

Output: QR fact. $\mathbf{Q}'_{(i)} = \mathbf{Q}_{(i)} \mathbf{R}_{(i,i)}$ with $\mathbf{Q}_{(i)}$ orthogonal with respect to $\langle \Theta \cdot, \Theta \cdot \rangle$. Sketch $\mathbf{S}_{(i)} = \Theta \mathbf{Q}_{(i)}$.

1. Compute $\mathbf{S}'_{(i)} = \Theta \mathbf{Q}'_{(i)}$.
 2. Compute $\mathbf{R}_{(i,i)}$ as the R factor of ℓ_2 -QR factorization of $\mathbf{S}'_{(i)}$.
 3. Compute $\mathbf{Q}_{(i)} = \mathbf{Q}'_{(i)} (\mathbf{R}_{(i,i)})^{-1}$ with forward substitution (or direct inversion, if $\mathbf{R}_{(i,i)}$ is well conditioned).
 4. Calculate the sketch as $\mathbf{S}_{(i)} = \mathbf{S}'_{(i)} (\mathbf{R}_{(i,i)})^{-1}$ or as $\mathbf{S}_{(i)} = \Theta \mathbf{Q}_{(i)}$.
-

2.2.3 ℓ_2 -QR+sketched CholQR factorization

The idea here is to improve the stability of the sketched CholQR by pre-processing the matrix $\mathbf{Q}'_{(i)}$ with classical routines for ℓ_2 -orthogonalization, as shown in Algorithm 6. In principle, the pre-processing step can be done with any suitable routine for ℓ_2 -orthogonalization of tall-and-skinny matrices such as the classical or modified Gram-Schmidt with l re-orthogonalizations, Householder algorithm, the tall-and-skinny QR from [11] (TSQR) or any others. In each particular situation, the most suitable routine should be selected depending on the computational architecture and programming environment. For instance, because of their popularity and reliability, the Gram-Schmidt and Householder algorithms are often available in scientific libraries as greatly optimized high-level routines. On the other hand, TSQR is favorable in massively parallel environments, since it reduces the amount of messages/synchronizations between processors. It has to be noted that Algorithm 6 permits the usage of algorithms that provide the Q factor in an implicit form, e.g., equal to the first m_p columns of $\mathbf{I} - \mathbf{YLY}^T$, which is the output of the Householder method. The only requirement is that such implicit Q factor is as efficient to multiply by the sketching matrix Θ , and vectors/matrices, as an explicit Q factor would be. In the rest of the paper, however, the matrix $\mathbf{Q}_{(i)}$ is assumed to be computed explicitly.

Algorithm 6 RBGS, Steps 4–5 (ℓ_2 -QR+CholQR factorization of $\mathbf{Q}'_{(i)}$)

Given: $\mathbf{Q}'_{(i)}, \Theta$.

Output: QR fact. $\mathbf{Q}'_{(i)} = \mathbf{Q}_{(i)} \mathbf{R}_{(i,i)}$ with $\mathbf{Q}_{(i)}$ orthogonal with respect to $\langle \Theta \cdot, \Theta \cdot \rangle$. Sketch $\mathbf{S}_{(i)} = \Theta \mathbf{Q}_{(i)}$.

1. Compute ℓ_2 -QR fact. $\mathbf{Q}'_{(i)} = \mathbf{Q}^* \mathbf{R}'$ with any suitable efficient/stable routine.
 2. Use Algorithm 5, taking \mathbf{Q}^* as $\mathbf{Q}'_{(i)}$, to compute sketched CholQR fact. $\mathbf{Q}^* = \mathbf{Q}_{(i)} \mathbf{R}''$ and sketch $\mathbf{S}_{(i)} = \Theta \mathbf{Q}_{(i)}$.
 3. Calculate $\mathbf{R}_{(i,i)} = \mathbf{R}'' \mathbf{R}'$.
-

It can be shown that Algorithm 6 produces a stable QR factorization satisfying (11) and (12), if it is computed with roundoff $u = u_{fine} = F(m, n)u_{crs}$, where $F(m, n)$ is some low-degree polynomial. Let us briefly outline the main arguments that support this fact. First, since we use a sufficiently stable routine in Step 1, we have $\text{cond}(\widehat{\mathbf{Q}}^*) = \mathcal{O}(1)$ and $|\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}^* \widehat{\mathbf{R}}'| = \mathcal{O}(u_{crs} |\widehat{\mathbf{Q}}^*| |\widehat{\mathbf{R}}'|)$. This property, for instance, is satisfied by the classical GS algorithm with $l \geq 2$ re-orthogonalizations, Householder QR, TSQR, and others. Furthermore, since Step 2 is performed independently of Θ we can assume that Θ satisfies the ε -embedding property for $\widehat{\mathbf{Q}}^*$, and therefore we also have $\text{cond}(\Theta \widehat{\mathbf{Q}}^*) \approx \text{cond}(\widehat{\mathbf{Q}}^*) = \mathcal{O}(1)$. These properties imply that $\text{cond}(\widehat{\mathbf{R}}'') = \mathcal{O}(1)$, which in its turn implies the stability of CholQR step. Notice that the fact that $\text{cond}(\widehat{\mathbf{R}}'') = \mathcal{O}(1)$, gives an ability to perform Step 3 in the sketched CholQR algorithm with efficient direct inversion instead of the usual forward substitution. It is concluded that $\Theta \widehat{\mathbf{Q}}_{(i)}$ is approximately orthogonal (i.e., it satisfies (12)) and that $|\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)}| = \mathcal{O}(|\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}^* \widehat{\mathbf{R}}'|) = \mathcal{O}(u_{crs} |\widehat{\mathbf{Q}}^*| |\widehat{\mathbf{R}}'|) = \mathcal{O}(u_{crs} |\mathbf{Q}_{(i)}| |\mathbf{R}_{(i,i)}|)$, which yields (11).

If needed, the output of Algorithm 6 can be certified a posteriori by computing $\|\mathbf{I} - \widehat{\mathbf{S}}_{(i)}^T \widehat{\mathbf{S}}_{(i)}\|_F$ and $\frac{\|\Theta \widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{S}}_{(i)} \widehat{\mathbf{R}}_{(i,i)}\|_F}{\|\Theta \widehat{\mathbf{Q}}'_{(i)}\|_F}$, similarly to the RGS algorithm in [3].

3 Stability analysis

In this section we provide a rigorous stability analysis of the RBGS algorithm. In particular, we show that the RBGS algorithm has similar a priori as well as a posteriori stability guarantees as its single-vector counterpart (see [3, Section 3]).

3.1 Assumptions

Our analysis will be based on the following assumptions. We first assume that Θ has a bounded norm as in (13). This condition is satisfied with probability at least $1 - \delta$, if Θ is $(1/2, \delta/n, 1)$ oblivious subspace embedding (see Corollary 1.3). Let the matrix $\Delta \mathbf{Q}'_{(i)}$ define the rounding error in Step 3:

$$\Delta \mathbf{Q}'_{(i)} = \widehat{\mathbf{Q}}'_{(i)} - \left(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)} \right).$$

Then, the standard worse-case scenario rounding analysis gives (14). Furthermore, we assume that Θ also satisfies (15) which, according to Corollary 1.5, holds under the probabilistic rounding model with probability at least $1 - 4\delta$, if Θ is $(1/8, m^{-1} \binom{n}{d}^{-1} \delta, d)$ oblivious subspace embedding, with $d = \mathcal{O}(\log(m/\delta))$. In its turn this condition is met by Rademacher matrices with $k = \mathcal{O}(\log(n) \log(m/\delta))$ rows or SRHT with $k = \mathcal{O}(\log^2(n) \log^2(m/\delta))$ rows.

Finally, it is assumed that in Steps 4–5, the QR factorization and the sketch of the Q factor satisfies (16) and (17) that can be achieved by using algorithms from Section 2.2.

Assumptions 3.1. *It is assumed that for some $\varepsilon \leq 1/2$,*

$$\|\Theta\|_F \leq \sqrt{1 + \varepsilon} \sqrt{n}. \quad (13)$$

Furthermore, we assume that

$$|\Delta \mathbf{Q}'_{(i)}| \leq 1.02 u_{crs} (|\widehat{\mathbf{W}}_{(i)}| + im_p |\widehat{\mathbf{Q}}_{(1:i-1)}| |\widehat{\mathbf{R}}_{(1:i-1,i)}|) \quad (14)$$

and

$$\|\Theta \Delta \mathbf{Q}'_{(i)}\|_F \leq 1.02 u_{crs} \sqrt{1 + \varepsilon} (|\widehat{\mathbf{W}}_{(i)}| + im_p |\widehat{\mathbf{Q}}_{(1:i-1)}| |\widehat{\mathbf{R}}_{(1:i-1,i)}|) \quad (15)$$

with $1 \leq i \leq p$ and $1 \leq t \leq m_p$. Finally, it is assumed that in in Steps 4–5 of Algorithm 2, we have

$$|\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)}| \leq u_{fine} m (|\widehat{\mathbf{Q}}'_{(i)}| + |\widehat{\mathbf{Q}}_{(i)}| |\widehat{\mathbf{R}}_{(i,i)}|) \leq 0.1 u_{crs} |\widehat{\mathbf{Q}}_{(i)}| |\widehat{\mathbf{R}}_{(i,i)}| \quad (16a)$$

$$\|\Theta(\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)})\|_F \leq u_{fine} m \|\Theta\|_F (|\widehat{\mathbf{Q}}'_{(i)}| + |\widehat{\mathbf{Q}}_{(i)}| |\widehat{\mathbf{R}}_{(i,i)}|) \leq 0.1 u_{crs} \|\Theta\|_F |\widehat{\mathbf{Q}}_{(i)}| |\widehat{\mathbf{R}}_{(i,i)}|, \quad (16b)$$

and, if Θ is ε' -embedding for $\widehat{\mathbf{Q}}'_{(i)}$,

$$1 - 0.1u_{crs}\text{cond}(\widehat{\mathbf{Q}}'_{(i)}) \leq \sigma_{\min}(\Theta\widehat{\mathbf{Q}}_{(i)}) \leq \sigma_{\max}(\Theta\widehat{\mathbf{Q}}_{(i)}) \leq 1 + 0.1u_{crs}\text{cond}(\widehat{\mathbf{Q}}'_{(i)}), \quad (17)$$

and $\|\widehat{\mathbf{Q}}_{(i)}\| \leq (1 - \varepsilon')^{-1/2}\|\Theta\widehat{\mathbf{Q}}_{(i)}\|$.

3.2 Stability guarantees of RBGS algorithm

Our stability analysis will rely on the condition that Θ satisfies the ε -embedding property for $\widehat{\mathbf{W}}$ and $\widehat{\mathbf{Q}}$. See Section 3.3 for a characterization of this property.

3.2.1 A posteriori analysis of RBGS algorithm

Let us first give an a posteriori characterization of the RBGS algorithm. Such characterization can be performed by measuring or bounding coefficients $\Delta^{(p)} = \|\mathbf{I} - \widehat{\mathbf{S}}^T\widehat{\mathbf{S}}\|_F$ and $\tilde{\Delta}^{(p)} = \frac{\|\widehat{\mathbf{P}} - \widehat{\mathbf{S}}\widehat{\mathbf{R}}\|_F}{\|\widehat{\mathbf{P}}\|_F}$, as done in [3]. We have the following result, which essentially extrapolates [3, Theorem 3.1] to the block version of the RGS algorithm.

Theorem 3.2. *Consider Algorithm 2. Assume that*

$$100m^{1/2}n^{3/2}u_{fine} \leq u_{crs} \leq 0.01,$$

along with Assumptions 3.1, possibly excluding (17).

If Θ is an ε -embedding for $\widehat{\mathbf{Q}}$ and $\widehat{\mathbf{W}}$, with $\varepsilon \leq 1/2$, and if $\Delta^{(p)}, \tilde{\Delta}^{(p)} \leq 0.1$, then the following inequalities hold:

$$(1 + \varepsilon)^{-1/2}(1 - \Delta^{(p)} - 0.1u_{crs}) \leq \sigma_{\min}(\widehat{\mathbf{Q}}) \leq \sigma_{\max}(\widehat{\mathbf{Q}}) \leq (1 - \varepsilon)^{-1/2}(1 + \Delta^{(p)} + 0.1u_{crs})$$

and

$$\|\widehat{\mathbf{W}} - \widehat{\mathbf{Q}}\widehat{\mathbf{R}}\|_F \leq 4u_{crs}m^{3/2}\|\widehat{\mathbf{W}}\|_F.$$

Proof. See appendix. □

Remark 3.3. *In Theorem 3.2, we also have*

$$\|\Theta(\widehat{\mathbf{W}} - \widehat{\mathbf{Q}}\widehat{\mathbf{R}})\|_F \leq 5u_{crs}m^{3/2}\|\widehat{\mathbf{W}}\|_F.$$

Proof. See appendix. □

Theorem 3.2 implies the numerical stability of the RBGS algorithm if $\Delta^{(p)}$ and $\tilde{\Delta}^{(p)}$ are ≤ 0.1 . These coefficients can be efficiently computed a posteriori from the sketches $\widehat{\mathbf{S}}$ and $\widehat{\mathbf{P}}$ and the R factor $\widehat{\mathbf{R}}$, thus providing a way for the certification of the solution. Such certification, in particular, does not involve any assumptions on $\text{cond}(\widehat{\mathbf{W}})$ and the accuracy of the least-squares solution in Step 2. Furthermore, we would like to highlight the very mild condition $u_{crs} = \mathcal{O}(m^{-3/2})$ on the working (coarse) unit roundoff to guarantee accuracy of the algorithm, which is in particular independent of the high-dimension n .

3.2.2 A priori analysis of RBGS algorithm

Clearly, to get a priori bounds for $\Delta^{(p)}$ and $\tilde{\Delta}^{(p)}$ we need more assumptions than in Theorem 3.2. In particular, it is necessary to impose a stability condition on the least squares solver used in Step 2. We also need $\widehat{\mathbf{W}}$ to be numerically non-singular, i.e., to satisfy $u_{crs} \leq \mathcal{O}(\text{cond}(\widehat{\mathbf{W}})^{-1})$.

The following a priori guarantee of stability of RBGS algorithm is, in fact, a restating of [3, Theorem 3.2] proposed for the single-vector algorithm.

Theorem 3.4. Consider Algorithm 2 with a backward-stable solver (e.g., based on Richardson iterations) satisfying Assumptions 2.2. Under Assumptions 3.1, if Θ is an ε -embedding for $\widehat{\mathbf{Q}}_{(1:p-1)}$ and $\widehat{\mathbf{W}}$, with $\varepsilon \leq 1/2$, and if

$$\begin{aligned} u_{crs} &\leq 10^{-3} \text{cond}(\widehat{\mathbf{W}})^{-1} m^{-2}, \\ u_{fine} &\leq 10^{-2} m^{-1/2} n^{-3/2} u_{crs}, \end{aligned}$$

then $\Delta^{(p)}$ and $\tilde{\Delta}^{(p)}$ are bounded by

$$\tilde{\Delta}^{(p)} \leq 4.2 u_{crs} m^{3/2} \|\widehat{\mathbf{W}}\|_{\text{F}} / \|\widehat{\mathbf{P}}\|_{\text{F}} \leq 6 u_{crs} m^{3/2}, \quad (18)$$

$$\Delta^{(p)} \leq 20 u_{crs} m^2 \text{cond}(\widehat{\mathbf{W}}). \quad (19)$$

Proof. See appendix. □

Theorem 3.4 states that the RBGS algorithm is stable unless the input matrix $\widehat{\mathbf{W}}$ is numerically singular. This stability guarantee is optimal and is seen in other stable algorithms such as MGS, CGS2, and others. Furthermore, the stability is proven for working unit roundoff independent of the high dimension n . This property of RBGS algorithm is unique and can be especially interesting for large-scale problems solved on low-precision arithmetic architectures.

3.3 Epsilon embedding property

The stability analysis in Section 3 holds if Θ satisfies the ε -embedding property for $\widehat{\mathbf{Q}}$ and $\widehat{\mathbf{W}}$. In this section we analyze this property.

We consider the case when $\widehat{\mathbf{W}}$ and Θ are independent of each other. Then, if Θ is (ε, δ, m) oblivious ℓ_2 -subspace embedding, it satisfies the ε -embedding property for $\widehat{\mathbf{W}}$ with high probability. Below, we show that, in this case Θ will also satisfy an ε -embedding property for $\widehat{\mathbf{Q}}$ with moderately increased value of ε . This result is basically a restatement of [3, Proposition 3.5] proven for the single-vector RGS algorithm.

Proposition 3.5. Consider Algorithm 2 with a backward-stable solver satisfying Assumptions 2.2 and

$$\begin{aligned} u_{crs} &\leq 10^{-3} \text{cond}(\widehat{\mathbf{W}})^{-1} m^{-2}, \\ u_{fine} &\leq 10^{-2} m^{-1/2} n^{-3/2} u_{crs}. \end{aligned}$$

Under Assumptions 3.1, if Θ is an ε -embedding for $\widehat{\mathbf{W}}$, with $\varepsilon \leq 1/4$, then it satisfies the ε' -embedding property for $\widehat{\mathbf{Q}}$, with

$$\varepsilon' = 2\varepsilon + 180 u_{crs} m^2 \text{cond}(\widehat{\mathbf{W}}).$$

Proof. See appendix. □

The ε -embedding property will likely hold even when matrix $\widehat{\mathbf{W}}$ depends on Θ . In such a case the quality of Θ can be certified a posteriori by computing additional sketches $\Phi\widehat{\mathbf{Q}}$ and $\Phi\widehat{\mathbf{W}}$, associated with a new sketching matrix Φ , in addition to the sketches $\Theta\widehat{\mathbf{Q}}$ and $\Theta\widehat{\mathbf{W}}$. Then one may characterize the quality of Θ by measuring the orthogonality of $(\Theta\widehat{\mathbf{W}})\mathbf{X}$ and $(\Theta\widehat{\mathbf{Q}})\mathbf{Y}$, where \mathbf{X} and \mathbf{Y} are inverses of Q factors of $\Phi\widehat{\mathbf{W}}$ and $\Phi\widehat{\mathbf{Q}}$, as is described in [3, Propositions 3.6-3.8] extrapolated from [5]. For efficiency in terms of cache or communication, at each iteration, the products $\Phi\widehat{\mathbf{W}}_{(i)}$ and $\Phi\widehat{\mathbf{Q}}_{(i)}$ can be computed together with $\Theta\widehat{\mathbf{W}}_{(i)}$ in Step 2, and $\Theta\widehat{\mathbf{Q}}_{(i)}$ in Steps 4-5, respectively.

4 Randomized block Krylov methods

In this section we discuss practical applications of the methodology. We particularly focus on the improvement of efficiency of popular block Krylov methods, such as the block GMRES method and the Rayleigh-Ritz method, for solving block linear systems of the form

$$\mathbf{A}\mathbf{X} = \mathbf{B}, \quad (20)$$

and eigenvalue problems of the form

$$\mathbf{A}\mathbf{X} = \mathbf{X}\mathbf{\Lambda}, \quad (21)$$

where \mathbf{A} is large, possibly non-symmetric $n \times n$ matrix, \mathbf{B} is $n \times m_p$ matrix, and $\mathbf{\Lambda}$ is $m_p \times m_p$ diagonal matrix of the extreme eigenvalues of \mathbf{A} .

The block Krylov methods proceed with approximation of \mathbf{X} by a projection $\mathbf{X}^{(j)}$ onto Krylov space $\mathcal{K}^{(j)}(\mathbf{A}, \mathbf{B})$, defined as

$$\mathcal{K}^{(j)}(\mathbf{A}, \mathbf{B}) := \text{span}\{\mathbf{B}, \mathbf{A}\mathbf{B}, \dots, \mathbf{A}^{j-1}\mathbf{B}\},$$

with j being the order of the subspace.

The GMRES method computes $\mathbf{X}^{(j)}$ that minimizes the Frobenius norm of the residual, while the Rayleigh-Ritz method seeks $\mathbf{X}^{(j)}$ that is optimal in the Galerkin sense. Both methods first construct an orthonormal basis for $\mathcal{K}^{(p)}(\mathbf{A}, \mathbf{B})$ with GS orthogonalization, called Arnoldi iteration, and then determine the coordinates of the columns of $\mathbf{X}^{(j)}$ in the computed basis.

4.1 Krylov basis computation: randomized block Arnoldi iteration

The Arnoldi algorithm produces orthonormal matrix $\mathbf{Q} = \mathbf{Q}_{(1:p)}$ satisfying the Arnoldi identity

$$\mathbf{A}\mathbf{Q}_{(1:p-1)} = \mathbf{Q}_{(1:p)}\mathbf{H},$$

where $\mathbf{H} = \mathbf{H}_{(1:p, 1:p-1)}$ is block upper Hessenberg matrix. The Arnoldi algorithm can be viewed as a block-wise QR factorization of matrix $[\mathbf{B}, \mathbf{A}\mathbf{Q}_{(1:p-1)}]$. In this context, the Hessenberg matrix \mathbf{H} can be seen as the R factor $\mathbf{R} = \mathbf{R}_{(1:p, 1:p)}$ without the first column of subblocks, that is, as $\mathbf{R}_{(1:p, 2:p)}$.

Below, we propose a randomized Arnoldi process based on RBGS algorithm. Note that unlike standard methods, Algorithm 7 produces a Krylov basis orthonormal with respect to the sketched product $\langle \Theta \cdot, \Theta \cdot \rangle$.

Algorithm 7 RBGS-Arnoldi algorithm

Given: $n \times n$ matrix \mathbf{A} , $n \times m_p$ matrix \mathbf{B} , $k \times n$ matrix Θ with $k \ll n$, parameter p .

Output: $n \times m$ factor $\mathbf{Q} = \mathbf{Q}_{(1:p)}$ and $m \times m$ upper triangular factor $\mathbf{R} = \mathbf{R}_{(1:p, 1:p)}$.

1. Set $\mathbf{W}_{(1)} = \mathbf{B}$ and perform 1-st iteration of RBGS (Algorithm 2).

for $i = 2 : p$ **do**

2. Compute $\mathbf{W}_{(i)} = \mathbf{A}\mathbf{Q}_{(i-1)}$.

3. Perform i -th iteration of RBGS (Algorithm 2).

end for

4. (Optional) compute $\Delta^{(p)}$ and $\tilde{\Delta}^{(p)}$. Use Theorem 3.2 to certify the output.

In Step 2 of Algorithm 7, the computation of the matrix-vector product can be executed either with roundoff u_{fine} or u_{crs} depending on the situation. The stability guarantees of Algorithm 7 can be obtained directly from Theorems 3.2 and 3.4 with standard stability analysis similar to that from [3, Section 4.1]. In particular, it can be shown that, under the stability conditions of RBGS, the computed $\hat{\mathbf{Q}}$ and $\hat{\mathbf{H}}$ satisfy

$$(\mathbf{A} + \Delta\mathbf{A})\hat{\mathbf{Q}}_{(1:p-1)} = \hat{\mathbf{Q}}_{(1:p)}\hat{\mathbf{H}}, \quad (22)$$

with $\|\Delta\mathbf{A}\|$ close to machine precision, and $\text{cond}(\hat{\mathbf{Q}}) = \mathcal{O}(1)$. We leave the precise analysis of this fact outside of the scope of this manuscript.

4.2 Linear systems: randomized block GMRES method

Let us now discuss solution of block linear systems (20) with GMRES.

The GMRES method computes the approximate solution $\mathbf{X}^{(p-1)} = \mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_{m_p}]$ as

$$\mathbf{u}_i = \widehat{\mathbf{Q}}_{(1:p-1)} \arg \min_{\mathbf{z}} \|\widehat{\mathbf{H}}\mathbf{z} - \widehat{\mathbf{R}}_{(1:p,1)}\mathbf{e}_i\|, \quad (23)$$

performed in sufficient precision, where \mathbf{e}_i is the i -th column of the identity matrix. Let us now characterize the quasi-optimality of such a projection when $\widehat{\mathbf{Q}}$ and $\widehat{\mathbf{H}}$ were obtained with the RBGS-Arnoldi algorithm. Under the stability conditions of RBGS, the Arnoldi identity (22) implies that

$$\|(\mathbf{A} + \Delta\mathbf{A})\widehat{\mathbf{Q}}_{(1:p-1)}\mathbf{z} - \mathbf{b}_i\| = \|\widehat{\mathbf{Q}}_{(1:p)}(\widehat{\mathbf{H}}\mathbf{z} - \widehat{\mathbf{R}}_{(1:p,1)}\mathbf{e}_i)\| \leq \|\widehat{\mathbf{Q}}\|\|\widehat{\mathbf{H}}\mathbf{z} - \widehat{\mathbf{R}}_{(1:p,1)}\mathbf{e}_i\|,$$

and, similarly,

$$\|(\mathbf{A} + \Delta\mathbf{A})\widehat{\mathbf{Q}}_{(1:p-1)}\mathbf{z} - \mathbf{b}_i\| \geq \sigma_{\min}(\widehat{\mathbf{Q}})\|\widehat{\mathbf{H}}\mathbf{z} - \widehat{\mathbf{R}}_{(1:p,1)}\mathbf{e}_i\|.$$

These two relations imply that

$$\|(\mathbf{A} + \Delta\mathbf{A})\mathbf{u}_i - \mathbf{b}_i\| \leq \text{cond}(\widehat{\mathbf{Q}}) \min_{\mathbf{v} \in Q_{p-1}} \|(\mathbf{A} + \Delta\mathbf{A})\mathbf{v} - \mathbf{b}_i\|,$$

with $Q_{p-1} = \text{range}(\widehat{\mathbf{Q}}_{(1:p-1)}) = \mathcal{K}_{p-1}(\mathbf{A} + \Delta\mathbf{A}, \mathbf{B})$. Consequently, the randomized version of the GMRES provides a solution which minimizes the norm of the residual associated with a slightly perturbed matrix, up to a factor of order 1.

Remark 4.1. *Randomized Full Orthogonalization Method.* In contrast to GMRES, the Full Orthogonalization Method (FOM) obtains solution $\mathbf{X}^{(p-1)} = \mathbf{U} = [\mathbf{u}_1, \dots, \mathbf{u}_{m_p}]$ by imposing a Galerkin orthogonality condition on the residuals (in exact arithmetic):

$$\langle \mathbf{v}, \mathbf{r}(\mathbf{u}_i, \mathbf{b}_i) \rangle = 0, \quad \forall \mathbf{v} \in Q_{p-1}, \quad 1 \leq i \leq m_p, \quad (24)$$

where $Q_{p-1} = \text{range}(\mathbf{Q}_{(1:p-1)})$ and $\mathbf{r}(\mathbf{u}_i, \mathbf{b}_i) = \mathbf{A}\mathbf{u}_i - \mathbf{b}_i$ is the residual associated with the i -th right-hand-side in (20). The Galerkin projection may be a more appropriate choice than the minres projection, when the quality of the solution is measured with an energy error rather than the residual error. When $\mathbf{Q}_{(1:p-1)}$ is obtained with the randomized Arnoldi method, the solution to (24) is given by

$$\mathbf{U} = \mathbf{Q}_{(1:p-1)}(\mathbf{H}_{(1:p-1,1:p-1)})^{-1}\mathbf{R}_{(1:p-1,1)}. \quad (25)$$

Suppose now that $\mathbf{Q}_{(1:p-1)}$ was obtained by the randomized Arnoldi method. Then, by noticing that

$$\mathbf{H}_{(1:p-1,1:p-1)} = [\mathbf{I} \ \mathbf{0}_{(1:p-1,p)}]\mathbf{H} = (\Theta\mathbf{Q}_{(1:p-1)})^T\Theta\mathbf{Q}_{(1:p-1)}\mathbf{H} = (\Theta\mathbf{Q}_{(1:p-1)})^T\Theta\mathbf{A}\mathbf{Q}_{(1:p-1)},$$

where $\mathbf{0}_{(1:p-1,p)}$ is a null matrix of size of $\mathbf{H}_{(p,1:p-1)}^T$, we deduce that the solution (25) satisfies

$$\langle \Theta\mathbf{v}, \Theta\mathbf{r}(\mathbf{u}_i, \mathbf{b}_i) \rangle = 0, \quad \forall \mathbf{v} \in Q_{p-1}, \quad 1 \leq i \leq m_p. \quad (26)$$

The relation (26) can be viewed as the sketched version of the Galerkin orthogonality condition. To our knowledge, it was first used in model order reduction community to obtain a reduced-basis solution of parametric linear systems [4]. By using similar considerations as in [4], one can show that (26) preserves the quality of the classical Galerkin projection when Θ is a ε -embedding for Q_p with $\varepsilon\text{cond}(\mathbf{A}) < 1$, though, in practice, this condition can be too pessimistic [4]. We leave the development of the randomized FOM method for future research.

Remark 4.2. It should be noted that the reasoning of this section could be reversed. We could first derive the sketched minres projection $\mathbf{u}_i = \arg \min_{\mathbf{v} \in Q_{p-1}} \|\Theta\mathbf{r}(\mathbf{v}, \mathbf{b}_i)\|$ and sketched Galerkin orthogonality condition (26) by replacing the ℓ_2 -inner products and ℓ_2 -norms in the standard minres equation $\mathbf{u}_i = \arg \min_{\mathbf{v} \in Q_{p-1}} \|\mathbf{r}(\mathbf{v}, \mathbf{b}_i)\|$ and the standard Galerkin orthogonality condition (24) by sketched ones, similarly as was done in [4, 5]. After that, it could be realized that the solution to the sketched minres and Galerkin equations can be obtained by orthogonalizing the Krylov basis with respect to $\langle \Theta\cdot, \Theta\cdot \rangle$ and using the classical identities (23) and (25).

4.3 Eigenvalue problems: randomized Rayleigh-Ritz method

Next, we consider the computation of the extreme eigenpairs of \mathbf{A} . Note that the methodology proposed below can be readily used for finding eigenpairs in the desired region by introducing a shift to the matrix. In addition, our methodology can be extended to inverse methods by replacing \mathbf{A} with \mathbf{A}^{-1} . In this case, the application of the inverse can be readily performed with (possibly preconditioned) randomized block GMRES method from Section 4.2 or other efficient methods.

To simplify the presentation, the analysis of the Rayleigh-Ritz method and its randomized version will be provided only in infinite precision arithmetic. The stability of these approaches follows directly from the stability of the Arnoldi iteration.

Let $\mathbf{Q} = \mathbf{Q}_{(1:p)}$ be the Krylov basis generated with the standard or randomized block-Arnoldi algorithm. The standard Rayleigh-Ritz method, based on the Arnoldi iteration, approximates the extreme eigenpairs (λ, \mathbf{x}) of \mathbf{A} by

$$(\lambda, \mathbf{x}) \approx (\mu, \mathbf{u}) = (\mu, \mathbf{Q}_{(1:p-1)}\mathbf{y}), \quad (27)$$

where (μ, \mathbf{y}) are the corresponding extreme eigenpairs of $\mathbf{H}_{(1:p-1,1:p-1)}$. Furthermore the residual error of an approximate eigenpair (μ, \mathbf{u}) is estimated by $\|\mathbf{H}_{(p,1:p-1)}\mathbf{y}\|$.

For the standard block-Arnoldi algorithm, we have

$$\mathbf{H}_{(1:p-1,1:p-1)} = [\mathbf{I} \mathbf{0}_{(1:p-1,p)}]\mathbf{H} = (\mathbf{Q}_{(1:p-1)})^T\mathbf{Q}\mathbf{H} = (\mathbf{Q}_{(1:p-1)})^T\mathbf{A}\mathbf{Q}_{(1:p-1)},$$

and therefore

$$\langle \mathbf{v}, \mathbf{r}(\mathbf{u}, \mu\mathbf{u}) \rangle = 0, \quad \forall \mathbf{v} \in Q_{p-1}, \quad (28)$$

where $Q_{p-1} := \text{range}(\mathbf{Q}_{(1:p-1)})$ and $\mathbf{r}(\mathbf{u}, \mu\mathbf{u}) = \mathbf{A}\mathbf{u} - \mu\mathbf{u}$ is the residual associated with (μ, \mathbf{u}) . The relation (28) is known as the Galerkin orthogonality condition. At the same time, we have

$$\|\mathbf{r}(\mu, \mathbf{u})\| \geq \sigma_{\min}(\mathbf{Q}) \left\| \mathbf{H}\mathbf{y} - \mu \begin{bmatrix} \mathbf{y} \\ \mathbf{0} \end{bmatrix} \right\| = \sigma_{\min}(\mathbf{Q})\|\mathbf{H}_{(p,1:p-1)}\mathbf{y}\|, \quad (29a)$$

and, similarly,

$$\|\mathbf{r}(\mu, \mathbf{u})\| \leq \sigma_{\max}(\mathbf{Q})\|\mathbf{H}_{(p,1:p-1)}\mathbf{y}\|. \quad (29b)$$

Since, the classical Arnoldi iteration produces an ℓ_2 -orthogonal \mathbf{Q} factor, hence in this case, the quantity $\|\mathbf{H}_{(p,1:p-1)}\mathbf{y}\|$ represents exactly the residual error of (μ, \mathbf{u}) .

On the other hand, the randomized block-Arnoldi algorithm produces a matrix \mathbf{Q} orthogonal with respect to $\langle \Theta \cdot, \Theta \cdot \rangle$. This implies that

$$\mathbf{H}_{(1:p-1,1:p-1)} = [\mathbf{I} \mathbf{0}_{(1:p-1,p)}]\mathbf{H} = (\Theta\mathbf{Q}_{(1:p-1)})^T(\Theta\mathbf{Q}_{(1:p-1)})\mathbf{H} = (\Theta\mathbf{Q}_{(1:p-1)})^T\Theta\mathbf{A}\mathbf{Q}_{(1:p-1)},$$

or equivalently, that (μ, \mathbf{u}) satisfies the following sketched version of the Galerkin orthogonality condition:

$$\langle \Theta\mathbf{v}, \Theta\mathbf{r}(\mathbf{u}, \mu\mathbf{u}) \rangle = 0, \quad \forall \mathbf{v} \in Q_{p-1}, \quad (30)$$

similar to the sketched Galerkin condition for linear systems in Remark 4.1. Unlike for the sketched minres projection in GMRES, the optimality for the sketched Galerkin projection does not trivially follow from the ε -embedding property of Θ (or equivalently from the fact that $\text{cond}(\mathbf{Q}) = 1 + \mathcal{O}(\varepsilon)$). Consequently, for completeness of our paper, in Section 4.3.1 we establish a characterization of the accuracy of this projection. Note that, independently of this work, the sketched Galerkin orthogonality condition was also used by Nakatsukasa and Tropp in their recent paper [16].

Since the ε -embedding property of Θ implies that $\text{cond}(\mathbf{Q}) = 1 + \mathcal{O}(\varepsilon)$, hence according to (29), the quantity $\|\mathbf{H}_{(p,1:p-1)}\mathbf{y}\|$ estimates well the residual error.

The Rayleigh-Ritz algorithm based on RBGS-Arnoldi with restarting is depicted in Algorithm 8.

Remark 4.3. When the classical Galerkin projection is preferable to its sketched variant, say, because of its strong optimality properties for Hermitian operators, the RBGS-Arnoldi algorithm can be applied directly to the classical Galerkin methods. In this case, the ℓ_2 -orthogonal Krylov basis matrix \mathbf{Q} and the Hessenberg matrix \mathbf{H} in (27) can be readily obtained with the RBGS-Arnoldi algorithm followed by an additional Cholesky QR step, as is depicted in Remark 2.1. In particular, this situation can be accounted for in Algorithm 8 by adding the following line between Step 1 and Step 2: “Compute $\mathbf{Q}^T \mathbf{Q}$ and obtain its Cholesky factor \mathbf{R}' . Set $\mathbf{Q} \leftarrow \mathbf{Q} \mathbf{R}'^{-1}$, $\mathbf{R} \leftarrow \mathbf{R}' \mathbf{R}$.” The cost of this additional step is dominated by the computation of product $\mathbf{Q}^T \mathbf{Q}$ (since we can omit computing $\mathbf{Q} \mathbf{R}'^{-1}$), which requires similar number of flops as the entire RBGS-Arnoldi algorithm. Nevertheless, as a single BLAS-3 operation, it is very well suited to modern computational architectures.

Algorithm 8 Randomized Rayleigh-Ritz algorithm for extreme eigenpairs with restarting

Given: $n \times n$ matrix \mathbf{A} , $n \times m_p$ matrix \mathbf{B} , $k \times n$ matrix Θ with $k \ll n$, parameters p and N_{iter} .

Output: \mathbf{A} and \mathbf{X} .

for $i = 1 : N_{iter}$ **do**

1. Perform RBGS-Arnoldi algorithm (Algorithm 7) returning \mathbf{Q} and \mathbf{R} . Set $\mathbf{H} = \mathbf{R}_{(1:p, 2:p)}$.
2. Compute diagonal matrix $\mathbf{\Lambda}$ with m_p extreme eigenvalues of $\mathbf{H}_{(1:p-1, 1:p-1)}$ on the diagonal, and the matrix of associated eigenvectors \mathbf{Y} .
3. (Optional) Compute the residuals $\|\mathbf{H}_{(p, 1:p-1)} \mathbf{y}\|$ to characterize the approximation error.
4. Compute $\mathbf{B} = \mathbf{Q}_{(1:p-1)} \mathbf{Y}$.

end for

5. Normalize \mathbf{B} with respect to ℓ_2 -norm and return it as \mathbf{X} .
-

4.3.1 Brief analysis of the sketched Galerkin projection

We proceed with reformulation of the methodology in terms of projection operators similarly to [17, Section 4.3] for classical methods. Our analysis will be based on the fact that Θ satisfies an ε -embedding property for some finite collection \mathcal{V} of fixed low-dimensional subspaces V . This property can be satisfied with probability at least $1 - \delta \#\mathcal{V}$, if Θ is an (ε, δ, d) oblivious ℓ_2 -subspace embedding, with $d = \max_{V \in \mathcal{V}} \dim(V)$.

Let \mathcal{K} be an approximation space, which does not necessarily have to be a Krylov space. Let $\Pi_{\mathcal{K}}$ denote the ℓ_2 -orthogonal projector onto \mathcal{K} :

$$\forall \mathbf{w} \in \mathbb{C}^n, \Pi_{\mathcal{K}} \mathbf{w} = \arg \min_{\mathbf{v} \in \mathcal{K}} \|\mathbf{w} - \mathbf{v}\|. \quad (31)$$

Then the Galerkin orthogonality condition (28) (with $\mathcal{K} = Q_{p-1}$) can be expressed as

$$\Pi_{\mathcal{K}} \mathbf{r}(\mathbf{u}, \mu \mathbf{u}) = 0,$$

or, equivalently,

$$\Pi_{\mathcal{K}} \mathbf{A} \Pi_{\mathcal{K}} \mathbf{u} = \mu \mathbf{u}.$$

In other words, the classical Rayleigh-Ritz method can be interpreted as approximation of eigenpairs of \mathbf{A} by the eigenpairs of *approximate* operator $\Pi_{\mathcal{K}} \mathbf{A} \Pi_{\mathcal{K}}$.

Similarly, the sketched Galerkin orthogonality condition (30) (with $\mathcal{K} = Q_{p-1}$) can be expressed as

$$\Pi_{\mathcal{K}}^{\Theta} \mathbf{r}(\mathbf{u}, \mu \mathbf{u}) = 0, \quad (32)$$

or, equivalently,

$$\Pi_{\mathcal{K}}^{\Theta} \mathbf{A} \Pi_{\mathcal{K}}^{\Theta} \mathbf{u} = \mu \mathbf{u},$$

where $\Pi_{\mathcal{K}}^{\Theta}$ is an orthogonal projector onto \mathcal{K} with respect to the sketched inner product $\langle \Theta \cdot, \Theta \cdot \rangle$, i.e.,

$$\forall \mathbf{w} \in \mathbb{C}^n, \Pi_{\mathcal{K}}^{\Theta} \mathbf{w} = \arg \min_{\mathbf{v} \in \mathcal{K}} \|\Theta(\mathbf{w} - \mathbf{v})\|, \quad (33)$$

or, in matrix form,

$$\mathbf{\Pi}_{\mathcal{K}}^{\Theta} = \mathbf{Q}(\Theta\mathbf{Q})^{\dagger}\Theta,$$

where \mathbf{Q} is a matrix whose columns compose a basis for \mathcal{K} . We see that the randomized Rayleigh-Ritz method corresponds to taking the approximate operator as $\mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{A}\mathbf{\Pi}_{\mathcal{K}}^{\Theta}$ instead of $\mathbf{\Pi}_{\mathcal{K}}\mathbf{A}\mathbf{\Pi}_{\mathcal{K}}$. Next we shall provide an upper bound for the residual norm of the exact eigenpair with respect to such approximate operator. For this, it is necessary to first establish a characterization of the sketched orthogonal projector, given in Lemma 4.4.

lemma 4.4. *Let $\mathbf{w} \in \mathbb{C}^n$. If Θ is ε -embedding for $V = \mathcal{K} + \text{span}(\mathbf{w})$, then*

$$\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\mathbf{w}\| \leq \|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}}^{\Theta})\mathbf{w}\| \leq \sqrt{\frac{1+\varepsilon}{1-\varepsilon}}\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\mathbf{w}\|, \quad (34a)$$

$$\frac{1}{1+\varepsilon}(\|\mathbf{\Pi}_{\mathcal{K}}\mathbf{w}\|^2 - 2\varepsilon\|\mathbf{w}\|^2) \leq \|\mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{w}\|^2 \leq \frac{1}{1-\varepsilon}(\|\mathbf{\Pi}_{\mathcal{K}}\mathbf{w}\|^2 + 2\varepsilon\|\mathbf{w}\|^2), \quad (34b)$$

Moreover, if $\mathbf{w} \in \mathcal{K}$, then

$$\mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{w} = \mathbf{w}. \quad (34c)$$

Proof. See appendix. □

Theorem 4.5. *(extension of [17, Theorem 4.3]) Assume that Θ is an ε -embedding for $V = \mathcal{K} + \mathbf{A}\mathcal{K} + \text{span}(\mathbf{x})$. Let*

$$\gamma := \max_{\mathbf{v} \in \mathcal{K} + \text{span}(\mathbf{x}), \|\mathbf{v}\|=1} \|\Theta\mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{A}(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}}^{\Theta})\mathbf{v}\|.$$

Then the sketched residual norms of the pairs $(\lambda, \mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{x})$ and (λ, \mathbf{x}) for the linear operator $\mathbf{A}_m = \mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{A}\mathbf{\Pi}_{\mathcal{K}}^{\Theta}$ satisfy, respectively,

$$\begin{aligned} \|\Theta(\mathbf{A}_m - \lambda\mathbf{I})\mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{x}\| &\leq \gamma\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}}^{\Theta})\mathbf{x}\| \\ \|\Theta(\mathbf{A}_m - \lambda\mathbf{I})\mathbf{x}\| &\leq \sqrt{\lambda^2 + \gamma^2}\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}}^{\Theta})\mathbf{x}\|. \end{aligned}$$

Proof. See appendix. □

Theorem 4.5 together with Lemma 4.4 imply that

$$\begin{aligned} \|(\mathbf{A}_m - \lambda\mathbf{I})\mathbf{\Pi}_{\mathcal{K}}^{\Theta}\mathbf{x}\| &\leq (1 + \mathcal{O}(\varepsilon))\gamma\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\mathbf{x}\| \\ \|(\mathbf{A}_m - \lambda\mathbf{I})\mathbf{x}\| &\leq (1 + \mathcal{O}(\varepsilon))\sqrt{\lambda^2 + \gamma^2}\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\mathbf{x}\|. \end{aligned}$$

Furthermore, according to Lemma 4.4, we also have

$$\gamma \leq \|\mathbf{\Pi}_{\mathcal{K}}\mathbf{A}(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\| + \mathcal{O}(\varepsilon)\|\mathbf{A}\| \leq (1 + \mathcal{O}(\varepsilon))\|\mathbf{A}\|.$$

Thus, Theorem 4.5 guarantees a good approximation if the distance $\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\mathbf{x}\|$ is small and if the approximate eigenproblem is well conditioned. Interestingly, by taking Θ as an identity in Theorem 4.5, we exactly recover Theorem 4.3 in [17].

On the contrary, the following result bounds the residual norm of the approximate eigenpair (μ, \mathbf{u}) in (32) with respect to the exact operator.

Theorem 4.6. *Let (μ, \mathbf{u}) be a solution to (32). If Θ is ε -embedding for $V = \mathcal{K} + \mathbf{A}\mathcal{K}$, then we have*

$$\|(\mathbf{A} - \mu\mathbf{I})\mathbf{u}\| \leq (1 + \mathcal{O}(\varepsilon))\|(\mathbf{I} - \mathbf{\Pi}_{\mathcal{K}})\mathbf{A}\mathbf{\Pi}_{\mathcal{K}}\|\|\mathbf{u}\|,$$

Moreover, if \mathcal{K} is a Krylov space Q_{p-1} , then

$$\|(\mathbf{A} - \mu\mathbf{I})\mathbf{u}\| \leq (1 + \mathcal{O}(\varepsilon))\|(\mathbf{I} - \mathbf{\Pi}_{Q_{p-1}})\mathbf{A}\mathbf{\Pi}_{Q_{p-1}}\|\|(\mathbf{I} - \mathbf{\Pi}_{Q_{p-2}})\mathbf{u}\|,$$

Proof. See appendix. □

Corollary 4.7. *Assume that Θ is an ε -embedding for \mathcal{K} . If \mathcal{K} is invariant under \mathbf{A} then each pair (μ, \mathbf{u}) that satisfies the sketched Galerkin orthogonality condition (32) is an eigenpair of \mathbf{A} .*

Theorem 4.6 implies two important results. The first one is that the sketched Galerkin projection provides an accurate result in terms of the residual error when the approximation space \mathcal{K} approximates well the range of \mathbf{A} or when \mathcal{K} is close to an invariant space. In particular, it follows that if we have $\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{A}\Pi_{\mathcal{K}}\| \leq \tau|\mu_m|$, then the dominant m computed eigenpairs (μ, \mathbf{u}) have relative error less than τ :

$$\frac{\|(\mathbf{A} - \mu\mathbf{I})\mathbf{u}\|}{\|\mu\mathbf{u}\|} \leq (1 + \mathcal{O}(\varepsilon))\tau. \quad (35)$$

At the same time, Theorem 4.6 also proves that if \mathcal{K} represents a Krylov subspace Q_{p-1} , then the convergence of eigenvectors implies convergence of their residuals to zero. More rigorously, if

$$\|\mathbf{u}^{(p)} - \mathbf{u}^{(p-1)}\| \leq \tau|\mu_m|\|\mathbf{u}^{(p)}\|,$$

where $\mathbf{u}^{(p)}$ denotes the sketched Galerkin solution in Q_{p-1} , then (35) holds.

Unfortunately, the fact that the residual norm is small does not guarantee that the eigenpair is accurate due to possibility of bad conditioning of the eigenvalue. If \mathbf{A} is Hermitian, this issue can be circumvented for classical Galerkin approximation, by appealing to the Min-Max principle and the Courant characterization [17]. Perhaps, similar ideas can be also used for the sketched Galerkin projection. The sketched Rayleigh quotient $R(\mathbf{w})$ can be naturally defined as a number that minimizes $\|\Theta(\mathbf{A}\mathbf{w} - R(\mathbf{w})\mathbf{w})\|$. It is easy to see that $R(\mathbf{w}) = \frac{\langle \Theta\mathbf{w}, \Theta\mathbf{A}\mathbf{w} \rangle}{\langle \Theta\mathbf{w}, \Theta\mathbf{w} \rangle}$. Furthermore, it can be shown that the solution (μ, \mathbf{u}) to (32) and the exact solution (λ, \mathbf{x}) satisfy $\mu = R(\mathbf{u})$ and $\lambda = R(\mathbf{x})$. We can also show that $|R(\Pi_{\mathcal{K}}\mathbf{x}) - \lambda|$ is small, if $\|\mathbf{I} - \Pi_{\mathcal{K}}\mathbf{x}\|$ is small. The main difficulty now becomes how to use the Min-Max principle, since the sketched operator $\Pi_{\mathcal{K}}^{\Theta}\mathbf{A}\Pi_{\mathcal{K}}^{\Theta}$ can be non-Hermitian, even when \mathbf{A} is Hermitian.

4.4 Further applications

4.4.1 s -step Krylov methods

The proposed RBGS algorithm can be used to improve the s -step Krylov methods. These methods use efficient matrix power kernels that can output a matrix $F(\mathbf{v})$ of basis vectors for $\mathcal{K}^{(s+1)}(\mathbf{A}, \mathbf{v})$, of the form

$$F(\mathbf{v}) = [p_0(\mathbf{A})\mathbf{v}, p_1(\mathbf{A})\mathbf{v}, \dots, p_s(\mathbf{A})\mathbf{v}],$$

where \mathbf{v} is some input vector, $p_0(\mathbf{A})$ is usually taken as 1, and $p_1(\mathbf{A}), \dots, p_s(\mathbf{A})$ are some suitable polynomials that aim to make $F(\mathbf{v})$ not too badly conditioned. Popular options for $p_0(\mathbf{A}), \dots, p_s(\mathbf{A})$ are the monomials and Newton, or Chebyshev polynomials. A relatively large s (say $s = 10$ or 30 [14]) may lead to stability problems even when a polynomial basis is used. Consequently, to obtain a Krylov basis suitable for practical applications, it becomes necessary to proceed with a block-wise generation, at each iteration, computing a small block of a Krylov basis by using the matrix power kernel and subsequently orthogonalizing it against the previously computed vectors with a BGS approach. In more concrete terms, we have to perform a block-wise QR factorization of the matrix $\mathbf{W} = \mathbf{W}_{(1:p)}$ generated recursively as

$$\mathbf{W}^{(i)} = \begin{cases} \mathbf{b} & \text{if } i = 1 \\ F^*(\mathbf{q}_{s(i-2)+1}) & \text{if } i = 2, \dots, p \end{cases}$$

where $F^*(\mathbf{v})$ corresponds to the matrix $F(\mathbf{v})$ without the first column, and $\mathbf{q}_{s(i-2)+1}$ is the last Krylov basis vector computed at iteration $i - 1$. Such block-wise orthogonalization can readily be done with the RBGS algorithm, as is depicted in Algorithm 9, with all the computational benefits of this algorithm.

An important point is that the random sketching allows not only to effectively construct a well-conditioned Krylov basis, but also to determine the Hessenberg matrix \mathbf{H} , without any additional high-dimensional operations nor the need to appeal to cumbersome recurrence relations. With the output of Algorithm 9, an approximate solution to the linear system can be obtained in a classical way presented in Section 4.2 (by setting $\mathbf{B} = \mathbf{b}$). We leave the numerical analysis of the proposed s -step RBGS-Arnoldi algorithm outside the scope of this manuscript.

Algorithm 9 s -step RBGS-Arnoldi algorithm

Given: $n \times n$ matrix \mathbf{A} , $n \times 1$ vector \mathbf{b} , $k \times n$ matrix Θ with $k \ll n$, parameters p and s .

Output: $n \times m$ Krylov basis matrix $\mathbf{Q} = \mathbf{Q}_{(1:p)}$ and Hessenberg matrix \mathbf{H} .

1. Set $\mathbf{W}_{(1)} = \mathbf{b}$ and perform 1-st iteration of RBGS (Algorithm 2).

for $i = 2 : p$ **do**

2. Compute $\mathbf{W}_{(i)} = F^*(\mathbf{q}_{s(i-2)+1})$ and perform i -th iteration of RBGS (Algorithm 2).

end for

3. Obtain $\mathbf{H} = \mathbf{S}^\dagger \mathbf{P}_{(1:p-1)}$ (possibly with setting the below-subdiagonal elements to zeros) by QR or SVD of \mathbf{S} , or solving a small least-squares problem.

4. (Optional) compute $\Delta^{(p)}$ and $\tilde{\Delta}^{(p)}$. Use Theorem 3.2 to certify the output.

5 Numerical experiments

We test the methodology on three numerical examples: QR factorization of a synthetically generated matrix from [3, Section 5.1], solution of a linear system with block GMRES, and solution of an eigenvalue problem with Rayleigh-Ritz method. The RBGS process is validated by comparison with standard methods such as BCGS, BMGS and BCGS2. Depending on the example, the finite precision arithmetic is performed in float32 or float64 format. To ensure good comparison, we decided to perform the inter-block orthogonalization routines in float64, even in the unique float32 precision algorithms. Furthermore, for standard methods we perform the inter-block orthogonalization with an efficient explicit Householder QR routine while for the RBGS, this routine is combined with the sketched Cholesky QR according to Algorithm 6. Several solvers are considered for the sketched least-squares problems depending on the experiment, but in all of them this step has a negligible complexity and memory requirements. Finally, in all the numerical examples the SRHT and Rademacher matrices give very similar results, so we present the results for SRHT only.

5.1 Orthogonalization of a numerically singular matrix

Take $[\mathbf{W}]_{i,j} = f_{\mu_j}(x_i)$, $1 \leq i \leq n$, $1 \leq j \leq m$, where

$$f_{\mu}(x) = \frac{\sin(10(\mu + x))}{\cos(100(\mu - x)) + 1.1},$$

and x_j and μ_j are chosen as, respectively, $n = 10^6$ and $m = 300$ equally distanced points with $x_0 = \mu_0 = 0$ and $x_{10^6} = \mu_{300} = 1$. The matrix \mathbf{W} is partitioned into blocks of columns of size $m_p = 10$. Then a block-wise QR factorization of \mathbf{W} is performed with the RBGS process (Algorithm 2) executed either in unique float32 precision or in multi-precision, using Θ with $k = 3000$ rows. The least-squares solver in Step 2 of Algorithm 2 is chosen either as the Householder solver, or as 20 iterations of CG applied to the normal equation. Thereafter the stability behavior of RBGS is compared to the standard BCGS, BMGS, and BCGS2 algorithms executed in float32 arithmetic.

We observe a similar picture as in [3, Section 5.1] comparing single-vector Gram-Schmidt algorithms. According to Figure 1a, the BCGS and BCGS2 methods become dramatically unstable respectively at iterations $i \geq 8$, and $i \geq 17$, with the latter iterations corresponding to $\mathbf{W}_{(1:i)}$ being numerically singular. The BMGS algorithm exhibits a better stability than the other two standard approaches. However, it still yields a Q factor with a condition number of two orders of magnitude. The unique precision RBGS algorithm using Householder least-squares solver in Step 2, has similar stability profile as BMGS algorithm. On the other hand, the usage of 20 iterations of CG does not provide a sufficient accuracy in Step 2 and as a consequence, reduces stability of the unique precision RBGS. In contrast to all tested methods, the multi-precision RBGS algorithm remains perfectly stable, even at iterations where $\mathbf{W}_{(1:i)}$ is numerically singular, and outputs a Q factor with the condition number close to $\mathcal{O}(1)$. Unlike in the unique precision algorithm, here the CG solver in Step 2 turns out to be sufficiently accurate. Moreover, Figure 1b examines the behavior of the approximation error $\|\mathbf{W}_{(1:i)} - \mathbf{Q}_{(1:i)}\mathbf{R}_{(1:i,1:i)}\|/\|\mathbf{W}_{(1:i)}\|$. We see that for all tested algorithms, besides BCGS2, this error remains close to machine precision at all iterations.

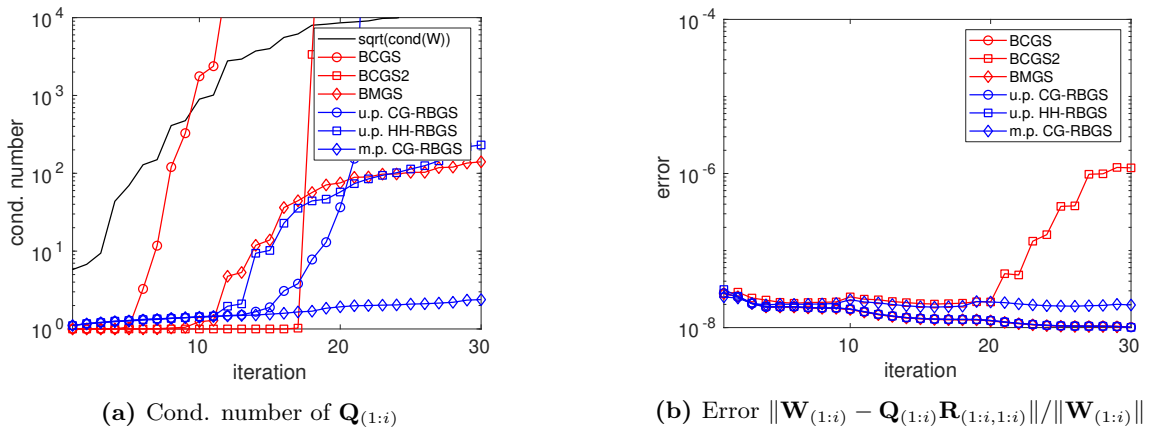


Figure 1: Block QR factorization of the matrix from [3, Section 5.1]. In the plots, the u.p. CG-RBGS and m.p. CG-RBGS respectively refer to the unique precision RBGS and the multi-precision RBGS algorithms, performing Step 2 with 20 iterations of CG. On the other hand, the u.p. HH-RBGS refers to the unique precision RBGS utilizing Householder solver.

5.2 Solution of a linear system with block GMRES

Consider linear system,

$$(\mathbf{A}_{Ga} + \alpha \mathbf{I})\mathbf{X}_{Ga} = \mathbf{B}, \quad (36)$$

where the matrix \mathbf{A}_{Ga} is taken as the “Ga41As41H72” matrix of dimension $n = 268096$ from the SuiteSparse matrix collection, \mathbf{I} is the $n \times n$ identity matrix, and $\alpha = 0.2$ introduced to improve conditioning of \mathbf{A}_{Ga} . The right-hand-side matrix \mathbf{B} is taken as $n \times m_p$ random matrix with entries being i.i.d normal random variables and $m_p = 100$. Solving such system could be a part of the inverse subspace iteration for computing eigenvalues of \mathbf{A}_{Ga} of the smallest magnitude. The shifted “Ga41As41H72” matrix has many clustered, possibly negative, eigenvalues, which can bring stability issues to the solvers.

Furthermore, the system (36) is preconditioned from the right by the incomplete LU factorization \mathbf{P}_{Ga} of $\mathbf{A}_{Ga} + \alpha \mathbf{I}$ with zero level of fill-in and symmetric reverse Cuthill-McKee reordering. With this preconditioner the final system of equations has the following form

$$\mathbf{A}\mathbf{X} = \mathbf{B},$$

where $\mathbf{A} = (\mathbf{A}_{Ga} + \alpha \mathbf{I})\mathbf{P}_{Ga}$ and $\mathbf{X}_{Ga} = \mathbf{P}_{Ga}\mathbf{X}$. Here matrix \mathbf{A} is not computed explicitly but considered as an implicit map outputting product with vectors. This system is approximately solved with GMRES method based on different versions of BGS process. We restart the GMRES every 30 iterations, i.e., when the dimension of the Krylov space becomes $m = 3100$.

Here we examine the behavior of BCGS, BCGS2, BMGS and the unique precision RBGS under float32 arithmetic. There is no need to test the multi-precision RBGS as the unique precision algorithms already are able to provide nearly optimal solution. The products with matrix \mathbf{A} and solutions of reduced GMRES least-squares problems (23) are computed in float64 format. The BGS iterations and other operations are performed in float32 format. In RBGS, the sketched least-squares problems are solved with 5 Richardson iterations as is explained in Section 2.1.

Figure 3a provides convergence of the residual error $\max_{j=1,\dots,100} \|\mathbf{A}\mathbf{u}_j - \mathbf{b}_j\| / \|\mathbf{b}_j\|$. The condition number of the computed Krylov basis $Q_{(1:i-1)}$ at each iteration i is given in Figure 3b. We see that already starting from the first iterations, the BCGS algorithm entails dramatic instability and stagnation of the residual error. At the same time, the BCGS2 algorithm remains stable at all iterations providing a Q factor that is orthonormal up to machine precision. The BMGS algorithm has partial instability resulting in a deteriorated convergence of the error. Finally, the RBGS algorithm is as stable as the BCGS2 algorithm. It provides a well-conditioned Q factor and an optimal convergence of the residual error for all tested sizes of the sketching matrix.




Figure 2: Solution of a linear system with GMRES.

5.3 Solution of an eigenvalue problem with randomized Rayleigh-Ritz

In this test we seek the smallest (negative) eigenpairs of the “Ga41As41H72” matrix \mathbf{A}_{Ga} from Section 5.2. For this we first transform the eigenproblem to the following positive-definite one:

$$(\alpha\mathbf{I} - \mathbf{A}_{Ga})\mathbf{x} = \lambda\mathbf{x}, \quad (37)$$

where $\alpha = 1500 \approx \|\mathbf{A}_{Ga}\|$, and seek the dominant eigenpairs of $\mathbf{A} = \alpha\mathbf{I} - \mathbf{A}_{Ga}$. Then one can use the fact that \mathbf{A}_{Ga} and \mathbf{A} have same eigenvectors and the associated eigenvalues related as $\lambda_{Ga} = \alpha - \lambda$. In this case the first 50 eigenvalues of \mathbf{A} are clustered inside $[1500.9 \ 1501.3]$. To compute the dominant eigenpairs of \mathbf{A} we either apply a subspace iteration or the Rayleigh-Ritz method based on different BGS algorithms. The initial guess matrix \mathbf{B} is taken as an $n \times m_p$ Gaussian matrix of dimension $m_p = 50$. The Arnoldi algorithm is restarted every 50 iterations (with columns of \mathbf{B} chosen as the dominant Ritz vectors) which corresponds to the Krylov space of dimension $m = 2550$. In this numerical example all arithmetic operations are performed in float64. The Step 2 of RBGS algorithm is again performed with 5 Richardson iterations as in the previous example.

We measure the approximation error by the maximum relative residual of the first 80% of computed eigenpairs (i.e., 40 eigenpairs out of 50). The convergence of the approximation error $\max_{j=1, \dots, 40} \|\mathbf{A}\mathbf{u}_j - \mu_j \mathbf{u}_j\| / \|\mu_j \mathbf{u}_j\|$ is depicted in Figure 3a. The condition number of the computed Krylov basis is depicted in Figure 3b. It is revealed that the subspace iteration method yields early stagnation of the error and is unsuitable for this numerical example. On the other hand, the Rayleigh-Ritz method, if stable, provides convergence of the error to machine precision.

It is revealed that the BCGS-based Arnoldi method becomes unstable already at early iterations and produces an ill-conditioned Krylov basis with condition number close to $\mathcal{O}(u^{-1})$. Therefore this method needs to be restarted, say every 5 iterations. Figure 3a plots the error also for this case. It is seen that the early restarting does not help as it causes a dramatic effect on the convergence of the error. In contrast to BCGS, the BCGS2 and RBGS algorithms show perfect stability and yield an approximation that converges to machine precision in 200 iterations. The BMGS method is unstable in the sense that it outputs an ill-conditioned Krylov basis. Despite this, it yields an approximation error that converges to machine precision, though, with somewhat deteriorated rates compared to BCGS2 and RBGS.

6 Concluding remarks

In this work we developed a block version of the randomized Gram-Schmidt process, called RBGS, to compute QR factorization of a large-scale matrix. It was shown that this algorithm inherits main properties of its single-vector analogue from [3], and in particular that it is at least as stable as the single-vector MGS process, and requires nearly half the cost of the classical BGS process in terms of flops and data passes. At the same time, our RBGS algorithm is well suited to cache-based and highly parallel computational architectures because it mainly relies on the matrix-matrix BLAS-3

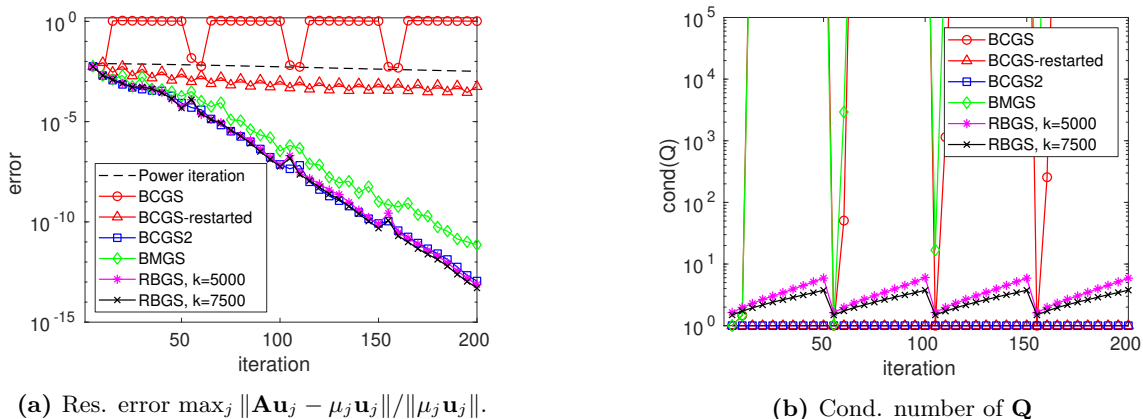


Figure 3: Solution of an eigenvalue problem with Rayleigh-Ritz method. In the plots, BCGS-restarted refers to BCGS-Arnoldi algorithm that restarts every 5 iterations.

operations, in contrast to its single vector counterpart relying on matrix-vector BLAS-2 operations. Like the single-vector RGS algorithm, it can be implemented using multi-precision arithmetic allowing to perform the dominant large-scale operations in precision independent of the dimension of the problem. This can become especially handy for simulations on low-precision arithmetic architectures.

Different strategies for treating the inter-block orthogonalization step, as well as solution of the sketched least-squares problems have been proposed. Special care has been taken to ensure that the computational cost of these steps is negligible compared to other steps in the RBGS algorithm.

The stability of our algorithms was verified in numerical experiments. In particular, it was seen that the multi-precision RBGS can provide a stable QR factorization of a numerically singular matrix, where even the most stable standard algorithms, including BCGS2, do not work. Furthermore, the randomized Arnoldi algorithm based on RBGS showed excellent stability in the context of GMRES and Rayleigh-Ritz approximation in numerical examples where BMGS showed instability and convergence degradation. These factors indicate the robustness of the RBGS algorithm.

Our next goal is to combine RBGS with model order reduction and compression techniques for even more, possibly asymptotic, reduction of the cost of Gram-Schmidt orthogonalization and the associated Krylov methods. Another direction is the application of random sketching to increase the robustness and/or efficiency of Krylov methods that use short recurrences such as CG, BCG, Lanczos, and others. Furthermore, we want to improve not only the Krylov methods but also other methods that involve an orthogonalization of the approximation basis, such as block LOBPCG. As for the theoretical analysis, besides improving the characterization of the sketched Galerkin projection in the randomized Rayleigh-Ritz method, we also plan to investigate the ε -embedding property of the sketching matrix for the computed Krylov space in the presence of rounding errors. Despite the fact that this property has been thoroughly verified in numerical experiments, it still remains unproven for the randomized Arnoldi algorithm.

7 Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation program (grant agreement No 810367).

References

- [1] N. Ailon and E. Liberty. “Fast dimension reduction using Rademacher series on dual BCH codes”. *Discrete & Computational Geometry* 42.4 (2009), p. 615 (cit. on p. 4).

- [2] A. H. Baker, J. M. Dennis, and E. R. Jessup. “On improving linear solver performance: A block variant of GMRES”. *SIAM Journal on Scientific Computing* 27.5 (2006), pp. 1608–1626 (cit. on p. 1).
- [3] O. Balabanov and L. Grigori. “Randomized Gram-Schmidt process with application to GMRES”. *arXiv preprint arXiv:2011.05090* (2020) (cit. on pp. 1, 2, 3, 4, 5, 6, 8, 9, 10, 11, 12, 13, 19, 20, 21).
- [4] O. Balabanov and A. Nouy. “Randomized linear algebra for model reduction. Part I: Galerkin methods and error estimation”. *Advances in Computational Mathematics* 45.5-6 (2019), pp. 2969–3019 (cit. on pp. 4, 14).
- [5] O. Balabanov and A. Nouy. “Randomized linear algebra for model reduction. Part II: minimal residual methods and dictionary-based approximation”. *Advances in Computational Mathematics* 47.2 (2021), pp. 1–54 (cit. on pp. 12, 14).
- [6] J. L. Barlow. “Block Modified Gram–Schmidt Algorithms and Their Analysis”. *SIAM Journal on Matrix Analysis and Applications* 40.4 (2019), pp. 1257–1290 (cit. on p. 3).
- [7] J. L. Barlow and A. Smoktunowicz. “Reorthogonalized block classical Gram–Schmidt”. *Numerische Mathematik* 123.3 (2013), pp. 395–423 (cit. on p. 3).
- [8] E. Carson, K. Lund, and M. Rozložník. “The Stability of Block Variants of Classical Gram–Schmidt”. *SIAM Journal on Matrix Analysis and Applications* 42.3 (2021), pp. 1365–1380 (cit. on p. 3).
- [9] E. Carson, K. Lund, M. Rozložník, and S. Thomas. “An overview of block Gram-Schmidt methods and their stability properties”. *arXiv preprint arXiv:2010.12058* (2020) (cit. on pp. 1, 3).
- [10] M. P. Connolly, N. J. Higham, and T. Mary. “Stochastic Rounding and its Probabilistic Backward Error Analysis” (2020) (cit. on p. 4).
- [11] J. Demmel, L. Grigori, M. Hoemmen, and J. Langou. “Communication-optimal parallel and sequential QR and LU factorizations”. *SIAM Journal on Scientific Computing* 34.1 (2012), A206–A239 (cit. on pp. 3, 9).
- [12] L. Grigori, S. Moufawad, and F. Nataf. “Enlarged Krylov subspace conjugate gradient methods for reducing communication”. *SIAM Journal on Matrix Analysis and Applications* 37.2 (2016), pp. 744–773 (cit. on p. 1).
- [13] N. J. Higham. *Accuracy and stability of numerical algorithms*. 2nd ed. SIAM Publications, Philadelphia, PA, USA, 2002 (cit. on pp. 5, 7).
- [14] M. Hoemmen. *Communication-avoiding Krylov subspace methods*. University of California, Berkeley, 2010 (cit. on pp. 1, 18).
- [15] P.-G. Martinsson and J. A. Tropp. “Randomized numerical linear algebra: Foundations and algorithms”. *Acta Numerica* 29 (2020), pp. 403–572 (cit. on p. 1).
- [16] Y. Nakatsukasa and J. A. Tropp. “Fast & Accurate Randomized Algorithms for Linear Systems and Eigenvalue Problems”. *arXiv preprint arXiv:2111.00113* (2021) (cit. on p. 15).
- [17] Y. Saad. *Numerical methods for large eigenvalue problems: revised edition*. SIAM, 2011 (cit. on pp. 1, 16, 17, 18, 29).
- [18] G. W. Stewart. “A Krylov–Schur algorithm for large eigenproblems”. *SIAM Journal on Matrix Analysis and Applications* 23.3 (2002), pp. 601–614 (cit. on p. 1).
- [19] J. A. Tropp. “Improved analysis of the subsampled randomized Hadamard transform”. *Advances in Adaptive Data Analysis* 3.01n02 (2011), pp. 115–126 (cit. on p. 4).
- [20] R. Vershynin. *High-dimensional probability: An introduction with applications in data science*. Vol. 47. Cambridge university press, 2018 (cit. on p. 1).
- [21] D. P. Woodruff et al. “Sketching as a tool for numerical linear algebra”. *Foundations and Trends® in Theoretical Computer Science* 10.1–2 (2014), pp. 1–157 (cit. on pp. 1, 4).
- [22] Q. Zou. “GMRES algorithms over 35 years”. *arXiv preprint arXiv:2110.04017* (2021) (cit. on p. 1).

Appendix

Here we provide proofs of some propositions and theorems from the paper.

Proof of Theorem 3.2. We have

$$\|\Theta\widehat{\mathbf{Q}} - \widehat{\mathbf{S}}\|_{\mathbf{F}} \leq 1.02u_{fine}n\|\Theta\|_{\mathbf{F}}\|\widehat{\mathbf{Q}}\|_{\mathbf{F}} \leq 1.02\sqrt{1 + \varepsilon}u_{fine}n^{3/2}\|\widehat{\mathbf{Q}}\|_{\mathbf{F}} \leq 0.02u_{crs}\|\widehat{\mathbf{Q}}\|_{\mathbf{F}} := F_1, \quad (38)$$

which implies that

$$\sigma_{\min}(\widehat{\mathbf{S}}) - F_1 \leq \sigma_{\min}(\Theta\widehat{\mathbf{Q}}) \leq \sigma_{\max}(\Theta\widehat{\mathbf{Q}}) \leq \sigma_{\max}(\widehat{\mathbf{S}}) + F_1.$$

By noticing that

$$1 - \Delta^{(p)} \leq \sqrt{1 - \Delta^{(p)}} \leq \sigma_{\min}(\widehat{\mathbf{S}}) \leq \sigma_{\max}(\widehat{\mathbf{S}}) \leq \sqrt{1 + \Delta^{(p)}} \leq 1 + \Delta^{(p)},$$

and Corollary 1.3 we deduce that $\|\widehat{\mathbf{Q}}\|_{\mathbf{F}} \leq 1.57$ and that

$$(1 + \varepsilon)^{-1/2}(1 - \Delta^{(p)} - F_1) \leq \sigma_{\min}(\widehat{\mathbf{Q}}) \leq \sigma_{\max}(\widehat{\mathbf{Q}}) \leq (1 - \varepsilon)^{-1/2}(1 + \Delta^{(p)} + F_1).$$

Next is proven the second statement of the theorem, which is obvious for $p = 1$. Assume that $p \geq 2$ and notice that

$$\begin{aligned} \|\widehat{\mathbf{R}}\|_{\mathbf{F}} &\leq \sigma_{\min}(\widehat{\mathbf{S}})^{-1}\|\widehat{\mathbf{S}}\widehat{\mathbf{R}}\|_{\mathbf{F}} \leq 1.12\|\widehat{\mathbf{S}}\widehat{\mathbf{R}}\|_{\mathbf{F}} \leq 1.12(\|\widehat{\mathbf{P}}\|_{\mathbf{F}} + \|\widehat{\mathbf{P}} - \widehat{\mathbf{S}}\widehat{\mathbf{R}}\|_{\mathbf{F}}) \\ &\leq 1.12(1 + \tilde{\Delta}^{(p)})\|\widehat{\mathbf{P}}\|_{\mathbf{F}} \leq 1.4(1 + \tilde{\Delta}^{(p)})\|\widehat{\mathbf{W}}\|_{\mathbf{F}} \leq 1.6\|\widehat{\mathbf{W}}\|_{\mathbf{F}}, \end{aligned} \quad (39)$$

We also have, for $1 \leq i \leq p$,

$$\widehat{\mathbf{W}}_{(i)} = \widehat{\mathbf{Q}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,i)} + \Delta\mathbf{W}_{(i)},$$

with

$$\begin{aligned} \|\Delta\mathbf{W}_{(i)}\|_{\mathbf{F}} &\leq \|\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)}\widehat{\mathbf{R}}_{(i,i)}\|_{\mathbf{F}} + 1.02u_{crs}(\|\widehat{\mathbf{W}}_{(i)}\|_{\mathbf{F}} + im_p\|\widehat{\mathbf{Q}}_{(1:i-1)}\|_{\mathbf{F}}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}) \\ &\leq 0.1u_{crs}\|\widehat{\mathbf{Q}}_{(i)}\|_{\mathbf{F}}\|\widehat{\mathbf{R}}_{(i,i)}\|_{\mathbf{F}} + 1.02u_{crs}(\|\widehat{\mathbf{W}}_{(i)}\|_{\mathbf{F}} + 1.57i^{3/2}m_p^{3/2}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}) \\ &\leq 1.02u_{crs}(\|\widehat{\mathbf{W}}_{(i)}\|_{\mathbf{F}} + 0.16m_p^{1/2}\|\widehat{\mathbf{R}}_{(i,i)}\|_{\mathbf{F}} + 1.57i^{3/2}m_p^{3/2}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}) \\ &\leq 1.02u_{crs}(1.26m_p^{1/2}\|\widehat{\mathbf{W}}\|_{\mathbf{F}} + 1.57i^{3/2}m_p^{3/2}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}). \end{aligned}$$

Consequently,

$$\begin{aligned} \|\Delta\mathbf{W}\|_{\mathbf{F}}^2 &\leq 1.02^2u_{crs}^2 \sum_{1 \leq i \leq p} 2(1.26^2m_p\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2 + 1.57^2i^3m_p^3\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}^2) \\ &\leq 1.02^2u_{crs}^2(1.26^2m\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2 + 1.57^2m^3\|\widehat{\mathbf{R}}\|_{\mathbf{F}}^2) \leq 1.02^2u_{crs}^2(1.26^2m\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2 + 2.52^2m^3\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2) \\ &\leq \left(4u_{crs}m^{3/2}\|\widehat{\mathbf{W}}\|_{\mathbf{F}}\right)^2. \end{aligned}$$

□

Proof of Remark 3.3. We have,

$$\widehat{\mathbf{W}}_{(i)} = \widehat{\mathbf{Q}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,i)} + \Delta\mathbf{W}_{(i)},$$

with

$$\begin{aligned} \|\Theta\Delta\mathbf{W}_{(i)}\|_{\mathbf{F}} &\leq \|\Theta(\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)}\widehat{\mathbf{R}}_{(i,i)})\|_{\mathbf{F}} + 1.02\sqrt{1 + \varepsilon}u_{crs}(\|\widehat{\mathbf{W}}_{(i)}\|_{\mathbf{F}} + im_p\|\widehat{\mathbf{Q}}_{(1:i-1)}\|_{\mathbf{F}}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}) \\ &\leq 1.02u_{crs}(1.25\|\widehat{\mathbf{W}}_{(i)}\|_{\mathbf{F}} + 0.16m_p^{1/2}\|\widehat{\mathbf{R}}_{(i,i)}\|_{\mathbf{F}} + 2i^{3/2}m_p^{3/2}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}) \\ &\leq 1.02u_{crs}(1.51m_p^{1/2}\|\widehat{\mathbf{W}}\|_{\mathbf{F}} + 2i^{3/2}m_p^{3/2}\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_{\mathbf{F}}) \end{aligned}$$

Consequently,

$$\begin{aligned} \|\Theta\Delta\mathbf{W}\|_{\mathbf{F}}^2 &\leq 1.02^2u_{crs}^2(1.51^2m\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2 + 2^2m^3\|\widehat{\mathbf{R}}\|_{\mathbf{F}}^2) \\ &\leq 1.02^2u_{crs}^2(1.51^2m\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2 + 3.2^2m^3\|\widehat{\mathbf{W}}\|_{\mathbf{F}}^2) \leq \left(5u_{crs}m^{3/2}\|\widehat{\mathbf{W}}\|_{\mathbf{F}}\right)^2. \end{aligned}$$

□

Proof of Theorem 3.4. The proof is done by induction on p . Assume that Theorem 3.4 holds for $p = i - 1 \geq 1$. Below, we show that the statement of the theorem then also holds for $p = i$.

Clearly,

$$\|\widehat{\mathbf{S}}_{(1:i-1)}\|_F \leq \sqrt{(i-1)m_p} \sqrt{1 + \Delta^{(i-1)}} \leq 1.01 \sqrt{(i-1)m_p} \quad \text{and} \quad \|\widehat{\mathbf{P}}_{(i)}\| \leq 1.02 \|\Theta \widehat{\mathbf{W}}_{(i)}\| \leq 1.25 \|\widehat{\mathbf{W}}_{(i)}\|. \quad (40)$$

Moreover, we have for $1 \leq j \leq i$,

$$\|\Theta \widehat{\mathbf{Q}}_{(j)} - \widehat{\mathbf{S}}_{(j)}\|_F \leq 1.02 u_{fine} \sqrt{1 + \varepsilon n^{3/2}} \|\widehat{\mathbf{Q}}_{(j)}\|_F \leq 0.02 u_{crs} \|\widehat{\mathbf{Q}}_{(j)}\|$$

From the fact that $\sigma_{min}(\widehat{\mathbf{S}}_{(1:i-1)}) \geq \sqrt{1 - \Delta^{(i-1)}} \geq 0.989$, and Assumptions 2.2 and (40), notice that

$$\sigma_{min}(\widehat{\mathbf{S}}_{(1:i-1)} + \Delta \mathbf{S}^{(i-1)}) \geq \sigma_{min}(\widehat{\mathbf{S}}_{(1:i-1)}) - \|\Delta \mathbf{S}^{(i-1)}\| \geq 0.98. \quad (41)$$

Consequently, we have

$$\|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_F \leq \|\widehat{\mathbf{P}}_{(i)} + \Delta \mathbf{P}^{(i)}\|_F / \sigma_{min}(\widehat{\mathbf{S}}_{(1:i-1)} + \Delta \mathbf{S}^{(i-1)}) \leq 1.4 \|\widehat{\mathbf{W}}_{(i)}\|_F \quad (42)$$

This fact combined with the fact (which holds by the induction hypothesis: $\Delta^{(i-1)} \leq 0.02$ and Theorem 3.2) that

$$\|\widehat{\mathbf{Q}}_{(1:i-1)}\| \leq 1.5 \quad \text{and} \quad \|\widehat{\mathbf{Q}}_{(1:i-1)}\|_F \leq 1.5(i-1)^{1/2} m_p^{1/2}, \quad (43)$$

and (14), leads to the following result:

$$\begin{aligned} \|\Delta \mathbf{Q}'_{(i)}\|_F &= \|\widehat{\mathbf{Q}}'_{(i)} - (\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)})\|_F \leq 1.02 u_{crs} \|\widehat{\mathbf{W}}_{(i)}\| + i m_p \|\widehat{\mathbf{Q}}_{(1:i-1)}\| \|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_F \\ &\leq 1.02 u_{crs} (\|\widehat{\mathbf{W}}_{(i)}\|_F + i m_p \|\widehat{\mathbf{Q}}_{(1:i-1)}\|_F \|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_F) \leq 3.2 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(i)}\|_F, \end{aligned} \quad (44)$$

$$\text{and} \quad \|\widehat{\mathbf{Q}}'_{(i)}\|_F \leq \|\widehat{\mathbf{W}}_{(i)}\|_F + \|\widehat{\mathbf{Q}}_{(1:i-1)}\| \|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_F + \|\Delta \mathbf{Q}'_{(i)}\|_F \leq 3.2 \|\widehat{\mathbf{W}}_{(i)}\|_F. \quad (45)$$

Then by (15),

$$\|\Theta \Delta \mathbf{Q}'_{(i)}\|_F \leq 3.2 \sqrt{1 + \varepsilon} u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(i)}\|_F \leq 4 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(i)}\|_F, \quad (46)$$

$$\text{and} \quad \|\Theta \widehat{\mathbf{Q}}'_{(i)}\|_F \leq \|\Theta \widehat{\mathbf{W}}_{(i)}\|_F + \|\Theta \widehat{\mathbf{Q}}_{(1:i-1)}\| \|\widehat{\mathbf{R}}_{(1:i-1,i)}\|_F + \|\Theta \Delta \mathbf{Q}'_{(i)}\|_F \leq 3.2 \|\widehat{\mathbf{W}}_{(i)}\|_F.$$

Denote residual matrix $\widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,1:i-1)} - \widehat{\mathbf{W}}_{(1:i-1)}$ by $\widehat{\mathbf{B}}^{(i-1)}$. Then, we have

$$\begin{aligned} \|\widehat{\mathbf{B}}^{(i-1)}\| &\leq 4 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(1:i-1)}\|_F \leq 0.01 \sigma_{min}(\widehat{\mathbf{W}}), \quad \text{and} \\ \|\Theta \widehat{\mathbf{B}}^{(i-1)}\| &\leq 5 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(1:i-1)}\|_F \leq 0.01 \sigma_{min}(\widehat{\mathbf{W}}). \end{aligned} \quad (47)$$

Now we are all set to derive the ε' -embedding property of Θ for $\widehat{\mathbf{Q}}'_{(i)}$, which is needed to characterize the quality of the inter-block orthogonalization step. Let us first notice that,

$$\begin{aligned} \sigma_{min}(\widehat{\mathbf{Q}}'_{(i)}) &= \sigma_{min}(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)} + \Delta \widehat{\mathbf{Q}}'_{(i)}) \\ &\geq \sigma_{min}(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)}) - \|\Delta \widehat{\mathbf{Q}}'_{(i)}\| \\ &\geq \sigma_{min}(\widehat{\mathbf{W}}_{(i)} - (\widehat{\mathbf{W}}_{(1:i-1)} + \widehat{\mathbf{B}}^{(i-1)}) \widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)}) - \|\Delta \widehat{\mathbf{Q}}'_{(i)}\| \\ &= \sigma_{min} \left([\widehat{\mathbf{W}}_{(1:i-1)} + \widehat{\mathbf{B}}^{(i-1)}, \widehat{\mathbf{W}}_{(i)}] \begin{bmatrix} -\widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)} \\ \mathbf{I} \end{bmatrix} \right) - \|\Delta \widehat{\mathbf{Q}}'_{(i)}\| \\ &\geq \sigma_{min}([\widehat{\mathbf{W}}_{(1:i-1)} + \widehat{\mathbf{B}}^{(i-1)}, \widehat{\mathbf{W}}_{(i)}]) - \|\Delta \widehat{\mathbf{Q}}'_{(i)}\| \\ &\geq \sigma_{min}(\widehat{\mathbf{W}}) - \|\widehat{\mathbf{B}}^{(i-1)}\| - \|\Delta \widehat{\mathbf{Q}}'_{(i)}\| \\ &\geq 0.98 \sigma_{min}(\widehat{\mathbf{W}}). \end{aligned} \quad (48)$$

Consequently, $\text{cond}(\widehat{\mathbf{Q}}'_{(i)}) \leq 5m_p^{1/2} \text{cond}(\widehat{\mathbf{W}})$. For any $\mathbf{a} \in \mathbb{R}^{m_p}$, it holds

$$\begin{aligned}
& \|\Theta \widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\| = \|\Theta(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \pm \|\Theta \Delta \widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\| \\
& = \|\Theta(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{W}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \pm \|\Theta \widehat{\mathbf{B}}^{(i-1)}\| \|\widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)} \mathbf{a}\| \pm \|\Theta \Delta \widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\| \\
& = \left(1 \pm \sigma_{\min}(\Theta \widehat{\mathbf{W}})^{-1} (\|\Theta \widehat{\mathbf{B}}^{(i-1)}\| + \|\Theta \Delta \widehat{\mathbf{Q}}'_{(i)}\|)\right) \|\Theta(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{W}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \\
& = (1 \pm 0.015) \|\Theta(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{W}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \\
& = (1 \pm 0.015) \sqrt{1 \pm \varepsilon} \|(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{W}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \\
& = (1 \pm 0.015) \sqrt{1 \pm \varepsilon} \left(\|(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \pm \|\widehat{\mathbf{B}}^{(i-1)}\| \|\widehat{\mathbf{R}}_{(1:i-1,1:i-1)}^{-1} \widehat{\mathbf{R}}_{(1:i-1,i)} \mathbf{a}\| \right) \\
& = (1 \pm 0.015) (1 \pm 0.01) \sqrt{1 \pm \varepsilon} \|(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)}) \mathbf{a}\| \\
& = (1 \pm 0.015) (1 \pm 0.01) \sqrt{1 \pm \varepsilon} (\|\widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\| + \|\Delta \widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\|) \\
& = (1 \pm 0.015) (1 \pm 0.01) (1 \pm 0.01) \sqrt{1 \pm \varepsilon} \|\widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\| \\
& = (1 \pm 0.036) \sqrt{1 \pm \varepsilon} \|\widehat{\mathbf{Q}}'_{(i)} \mathbf{a}\|,
\end{aligned}$$

which, combined with the parallelogram identity, implies that Θ is an ε' -embedding for $\widehat{\mathbf{Q}}'_{(i)}$ with $\varepsilon' \leq 1.1\varepsilon + 0.1 \leq 0.65$. By combining this fact with Assumptions 3.1, we obtain:

$$1 - u_{crs} m_p^{1/2} \text{cond}(\widehat{\mathbf{W}}) \leq \sigma_{\min}(\Theta \widehat{\mathbf{Q}}_{(i)}) \leq \sigma_{\max}(\Theta \widehat{\mathbf{Q}}_{(i)}) \leq 1 + u_{crs} m_p^{1/2} \text{cond}(\widehat{\mathbf{W}}),$$

and $\|\widehat{\mathbf{Q}}_{(i)}\| \leq 1.8$.

Since $\|\Theta \widehat{\mathbf{Q}}_{(i)} - \widehat{\mathbf{S}}_{(i)}\|_{\text{F}} \leq 0.02 u_{crs} \|\widehat{\mathbf{Q}}_{(i)}\|$, we have

$$1 - 1.05 u_{crs} m_p^{1/2} \text{cond}(\widehat{\mathbf{W}}) \leq \sigma_{\min}(\widehat{\mathbf{S}}_{(i)}) \leq \sigma_{\max}(\widehat{\mathbf{S}}_{(i)}) \leq 1 + 1.05 u_{crs} m_p^{1/2} \text{cond}(\widehat{\mathbf{W}}).$$

Consequently,

$$\|\mathbf{I} - \widehat{\mathbf{S}}_{(i)}^{\text{T}} \widehat{\mathbf{S}}_{(i)}\|_{\text{F}} \leq m_p^{1/2} \|\mathbf{I} - \widehat{\mathbf{S}}_{(i)}^{\text{T}} \widehat{\mathbf{S}}_{(i)}\| \leq m_p^{1/2} \max_{\mathbf{x}} \frac{|\|\mathbf{x}\|^2 - \|\widehat{\mathbf{S}}_{(i)} \mathbf{x}\|^2|}{\|\mathbf{x}\|^2} \leq 3 u_{crs} m_p \text{cond}(\widehat{\mathbf{W}}).$$

We also have $\|\widehat{\mathbf{R}}_{(i,i)}\|_{\text{F}} \leq 1.02 \|\Theta \widehat{\mathbf{Q}}'_{(i)}\|_{\text{F}} \leq 3.3 \|\widehat{\mathbf{W}}_{(i)}\|_{\text{F}}$, and

$$\sigma_{\min}(\widehat{\mathbf{R}}_{(i,i)}) \geq \sigma_{\min}(\Theta \widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)}) / \|\Theta \widehat{\mathbf{Q}}_{(i)}\| \geq 0.99 (\sigma_{\min}(\Theta \widehat{\mathbf{Q}}'_{(i)}) - 0.1 u_{crs} \|\widehat{\mathbf{Q}}_{(i)}\|) \|\widehat{\mathbf{R}}_{(i,i)}\| \geq 0.58 \sigma_{\min}(\widehat{\mathbf{W}}),$$

and

$$\begin{aligned}
\|\widehat{\mathbf{S}}_{(1:i-1)}^{\text{T}} \widehat{\mathbf{S}}_{(i)}\|_{\text{F}} & = \|\widehat{\mathbf{S}}_{(1:i-1)}^{\text{T}} (\Theta \widehat{\mathbf{Q}}'_{(i)} + \Delta \widehat{\mathbf{S}}'_{(i)}) \widehat{\mathbf{R}}_{(i,i)}^{-1} + \Delta \widehat{\mathbf{S}}''_{(i)}\|_{\text{F}} \\
& \leq (\|\widehat{\mathbf{S}}_{(1:i-1)}^{\text{T}} \Theta \widehat{\mathbf{Q}}'_{(i)}\|_{\text{F}} + 1.01 \|\Delta \widehat{\mathbf{S}}'_{(i)}\|_{\text{F}}) / 0.58 \sigma_{\min}(\widehat{\mathbf{W}}) + \|\Delta \widehat{\mathbf{S}}''_{(i)}\|_{\text{F}},
\end{aligned}$$

where $\Delta \widehat{\mathbf{S}}'_{(i)} = \Theta(\widehat{\mathbf{Q}}_{(i)} \widehat{\mathbf{R}}_{(i,i)} - \widehat{\mathbf{Q}}'_{(i)})$ and $\Delta \widehat{\mathbf{S}}''_{(i)} = \widehat{\mathbf{S}}_{(i)} - \Theta \widehat{\mathbf{Q}}_{(i)}$.

By noting that,

$$\begin{aligned}
\|\Delta \widehat{\mathbf{S}}'_{(i)}\|_{\text{F}} & \leq 0.1 u_{crs} \|\widehat{\mathbf{Q}}_{(i)}\| \|\widehat{\mathbf{R}}_{(i,i)}\| \leq 0.6 u_{crs} \|\widehat{\mathbf{W}}_{(i)}\|_{\text{F}} \\
\|\Delta \widehat{\mathbf{S}}''_{(i)}\|_{\text{F}} & \leq 0.02 u_{crs} \|\widehat{\mathbf{Q}}_{(i)}\| \leq 0.04 u_{crs} \\
\|\widehat{\mathbf{S}}_{(1:i-1)}^{\text{T}} \Theta \widehat{\mathbf{Q}}'_{(i)}\|_{\text{F}} & \leq \|\widehat{\mathbf{S}}_{(1:i-1)}^{\text{T}} \Theta(\widehat{\mathbf{W}}_{(i)} - \widehat{\mathbf{Q}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)})\|_{\text{F}} + 4.04 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(i)}\|_{\text{F}} \\
& \leq \|\widehat{\mathbf{S}}_{(1:i-1)}^{\text{T}} (\widehat{\mathbf{P}}_{(i)} - \widehat{\mathbf{S}}_{(1:i-1)} \widehat{\mathbf{R}}_{(1:i-1,i)})\|_{\text{F}} + 4.2 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(i)}\|_{\text{F}} \\
& \leq 5 u_{crs} i^{3/2} m_p^{3/2} \|\widehat{\mathbf{W}}_{(i)}\|_{\text{F}},
\end{aligned}$$

we deduce that $\|\widehat{\mathbf{S}}_{(1:i-1)}^T \widehat{\mathbf{S}}_{(i)}\|_F \leq 10u_{crs}i^{3/2}m_p^2\text{cond}(\widehat{\mathbf{W}})$.

Consequently, for $i \geq 2$,

$$\begin{aligned} (\Delta^{(i)})^2 &= \|\mathbf{I} - (\widehat{\mathbf{S}}_{(i)})^T \widehat{\mathbf{S}}_{(i)}\|_F^2 + 2\|\widehat{\mathbf{S}}_{(1:i-1)}^T \widehat{\mathbf{S}}_{(i)}\|_F^2 + (\Delta^{(i-1)})^2 \\ &\leq (3u_{crs}m_p\text{cond}(\widehat{\mathbf{W}}))^2 + 2(10u_{crs}i^{3/2}m_p^2\text{cond}(\widehat{\mathbf{W}}))^2 + (20u_{crs}(i-1)^2m_p^2\text{cond}(\widehat{\mathbf{W}}))^2 \\ &\leq (9 + 200i^3 + 400(i-1)^4)(m_p^2u_{crs}\text{cond}(\widehat{\mathbf{W}}))^2 \\ &\leq (400i^4 - 1400i^3 + 2400i^2 - 1600i + 409)(m_p^2u_{crs}\text{cond}(\widehat{\mathbf{W}}))^2 \\ &\leq (20i^2m_p^2u_{crs}\text{cond}(\widehat{\mathbf{W}}))^2 \end{aligned}$$

and,

$$\begin{aligned} \|\widehat{\mathbf{S}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,i)} - \widehat{\mathbf{P}}_{(i)}\|_F &\leq \|\Theta\widehat{\mathbf{Q}}'_{(i)} - (\widehat{\mathbf{P}}_{(i)} - \widehat{\mathbf{S}}_{(1:i-1)}\widehat{\mathbf{R}}_{(1:i-1,i)})\|_F + \|\widehat{\mathbf{S}}_{(i)}\widehat{\mathbf{R}}_{(i,i)} - \Theta\widehat{\mathbf{Q}}'_{(i)}\|_F \\ &\leq \|\Theta\Delta\mathbf{Q}'_{(i)}\|_F + \|\widehat{\mathbf{P}}_{(i)} - \Theta\widehat{\mathbf{W}}_{(i)}\|_F + \|\Theta\widehat{\mathbf{Q}}_{(1:i-1)} - \widehat{\mathbf{S}}_{(1:i-1)}\|_F\|\widehat{\mathbf{R}}_{(1:i-1,i)}\| + \|\Theta\widehat{\mathbf{Q}}_{(i)} - \widehat{\mathbf{S}}_{(i)}\|_F\|\widehat{\mathbf{R}}_{(i,i)}\| + \|\Theta(\widehat{\mathbf{Q}}'_{(i)} - \widehat{\mathbf{Q}}_{(i)}\widehat{\mathbf{R}}_{(i,i)})\|_F \\ &\leq u_{crs}(3.2\sqrt{1 + \varepsilon}i^{3/2}m_p^{3/2} + 0.02 + 0.02 \cdot 1.5 \cdot 1.4 + 0.04 \cdot 3.3 + 0.6)\|\widehat{\mathbf{W}}_{(i)}\|_F \leq 4.2i^{3/2}m_p^{3/2}\|\widehat{\mathbf{W}}_{(i)}\|_F. \end{aligned}$$

Consequently,

$$\begin{aligned} (\Delta^{(i)})^2 &= \|\widehat{\mathbf{S}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)} - \widehat{\mathbf{P}}_{(1:i)}\|_F^2 / \|\widehat{\mathbf{P}}_{(1:i)}\|_F^2 = (\|\widehat{\mathbf{S}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,i)} - \widehat{\mathbf{P}}_{(i)}\|_F^2 + \|\widehat{\mathbf{S}}_{(1:i-1)}\widehat{\mathbf{R}}_{(1:i-1,1:i-1)} - \widehat{\mathbf{P}}_{(1:i-1)}\|_F^2) / \|\widehat{\mathbf{P}}_{(1:i)}\|_F^2 \\ &\leq ((4.2i^{3/2}m_p^{3/2}\|\widehat{\mathbf{W}}_{(i)}\|_F)^2 + (4.2i^{3/2}m_p^{3/2}\|\widehat{\mathbf{W}}_{(1:i-1)}\|_F)^2) / \|\widehat{\mathbf{P}}_{(1:i)}\|_F^2 = (4.2i^{3/2}m_p^{3/2}\|\widehat{\mathbf{W}}_{(1:i)}\|_F)^2 / \|\widehat{\mathbf{P}}_{(1:i)}\|_F^2. \end{aligned}$$

The proof is finished by noting that $\|\widehat{\mathbf{P}}_{(1:i)}\|_F \geq (\sqrt{1 - \varepsilon} - 0.001)\|\widehat{\mathbf{W}}_{(1:i)}\|_F \geq 0.7\|\widehat{\mathbf{W}}_{(1:i)}\|_F$. \square

Proof of Proposition 3.5. We proceed with an induction on p . Clearly, the statement of the proposition holds for $p = 1$. Assume that Θ is an ε -embedding for $\widehat{\mathbf{Q}}_{(1:p)}$, $p = i - 1 \geq 1$. This condition is sufficient for the following results in Theorem 3.4 and its proof to hold for $p = i$:

$$\|\widehat{\mathbf{P}}_{(1:i)} - \widehat{\mathbf{S}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)}\|_F \leq 4.2u_{crs}m_p^{3/2}i^{3/2}\|\widehat{\mathbf{W}}_{(1:i)}\|_F, \quad (49)$$

$$\Delta^{(i)} = \|\mathbf{I} - \widehat{\mathbf{S}}_{(1:i)}^T \widehat{\mathbf{S}}_{(1:i)}\|_F \leq 20u_{crs}m_p^2i^2\text{cond}(\widehat{\mathbf{W}}_{(1:i)}) \leq 0.02. \quad (50)$$

In addition, we have $\|\widehat{\mathbf{R}}_{(1:i,1:i)}\|_F \leq 1.2\|\widehat{\mathbf{W}}_{(1:i)}\|_F$, $\|\widehat{\mathbf{Q}}_{(1:i-1)}\| \leq 1.9$, $\|\widehat{\mathbf{Q}}_{(i)}\| \leq 1.3$, and

$$\begin{aligned} \|\widehat{\mathbf{W}}_{(1:i)} - \widehat{\mathbf{Q}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)}\|_F &\leq \|\Delta\mathbf{Q}'_{(1:i)}\|_F + \left(\sum_{1 \leq j \leq i} \|\widehat{\mathbf{Q}}'_{(j)} - \widehat{\mathbf{Q}}_{(j)}\widehat{\mathbf{R}}_{(j,j)}\|_F^2\right)^{1/2} \\ &\leq \|\Delta\mathbf{Q}'_{(1:i)}\|_F + 0.1u_{crs}\left(\sum_{1 \leq j \leq i} \|\widehat{\mathbf{Q}}_{(j)}\|_F^2\|\widehat{\mathbf{R}}_{(j,j)}\|_F^2\right)^{1/2} \leq \|\Delta\mathbf{Q}'_{(1:i)}\|_F + 0.1u_{crs}(\|\widehat{\mathbf{Q}}_{(i)}\|_F + \|\widehat{\mathbf{Q}}_{(1:i-1)}\|_F)\|\widehat{\mathbf{R}}_{(1:i,1:i)}\|_F \\ &\leq 3.5u_{crs}m_p^{3/2}i^{3/2}\|\widehat{\mathbf{W}}_{(1:i)}\|_F, \end{aligned} \quad (51)$$

which can be proven similarly to (39) and (44) in the proofs of Theorems 3.2 and 3.4. Finally, the following holds

$$\begin{aligned} \|\widehat{\mathbf{S}}_{(1:i)} - \Theta\widehat{\mathbf{Q}}_{(1:i)}\|_F &\leq 1.02u_{fine}\sqrt{1 + \varepsilon}n^{3/2}\|\widehat{\mathbf{Q}}_{(1:i)}\|_F \leq u_{crs}0.02\|\widehat{\mathbf{Q}}_{(1:i)}\|_F, \\ \|\widehat{\mathbf{P}}_{(1:i)} - \Theta\widehat{\mathbf{W}}_{(1:i)}\|_F &\leq 1.02u_{fine}\sqrt{1 + \varepsilon}n^{3/2}\|\widehat{\mathbf{W}}_{(1:i)}\|_F \leq u_{crs}0.02\|\widehat{\mathbf{W}}_{(1:i)}\|_F. \end{aligned} \quad (52)$$

By using the fact that $\|\widehat{\mathbf{S}}_{(1:i)}\| \leq 1.02$ and $\sigma_{\min}(\widehat{\mathbf{P}}_{(1:i)}) \geq \sigma_{\min}(\Theta\widehat{\mathbf{W}}_{(1:i)}) - 0.02u_{crs}\|\widehat{\mathbf{W}}_{(1:i)}\|$, along with (49) and the ε -embedding property of Θ , we get

$$\sigma_{\min}(\widehat{\mathbf{R}}_{(1:i,1:i)}) \geq \frac{1}{\|\widehat{\mathbf{S}}_{(1:i)}\|}(\sigma_{\min}(\widehat{\mathbf{P}}_{(1:i)}) - \|\widehat{\mathbf{P}}_{(1:i)} - \widehat{\mathbf{S}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)}\|_F) \geq 0.7\sigma_{\min}(\widehat{\mathbf{W}}_{(1:i)}). \quad (53)$$

Properties (49), (52) and (53) imply that

$$\begin{aligned} \|\Theta\widehat{\mathbf{W}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)}^{-1} - \widehat{\mathbf{S}}_{(1:i)}\|_F &\leq \|\widehat{\mathbf{P}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)}^{-1} - \widehat{\mathbf{S}}_{(1:i)}\|_F + \|(\widehat{\mathbf{P}}_{(1:i)} - \Theta\widehat{\mathbf{W}}_{(1:i)})\widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}\|_F \\ &\leq (\|\widehat{\mathbf{P}}_{(1:i)} - \widehat{\mathbf{S}}_{(1:i)}\widehat{\mathbf{R}}_{(1:i,1:i)}\|_F + \|\widehat{\mathbf{P}}_{(1:i)} - \Theta\widehat{\mathbf{W}}_{(1:i)}\|_F)\|\widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}\|_F \\ &\leq 6.2u_{crs}m_p^2i^2\text{cond}(\widehat{\mathbf{W}}_{(1:i)}) =: F_1. \end{aligned}$$

Furthermore,

$$1 - \Delta^{(i)} - F_1 \leq \sigma_{\min}(\Theta \widehat{\mathbf{W}}_{(1:i)} \widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}) \leq \sigma_{\max}(\Theta \widehat{\mathbf{W}}_{(1:i)} \widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}) \leq 1 + \Delta^{(i)} + F_1.$$

Due to the fact that Θ is an ε -embedding for $\widehat{\mathbf{W}}_{(1:i)}$, we deduce that

$$(1 + \varepsilon)^{-1/2}(1 - \Delta^{(i)} - F_1) \leq \sigma_{\min}(\widehat{\mathbf{W}}_{(1:i)} \widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}) \leq \sigma_{\max}(\widehat{\mathbf{W}}_{(1:i)} \widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}) \leq (1 - \varepsilon)^{-1/2}(1 + \Delta^{(i)} + F_1). \quad (54)$$

We also have from (51) and (53),

$$\|\widehat{\mathbf{W}}_{(1:i)} \widehat{\mathbf{R}}_{(1:i,1:i)}^{-1} - \widehat{\mathbf{Q}}_{(1:i)}\|_F \leq 3.5u_{crs}m_p^{3/2}i^{3/2}\|\widehat{\mathbf{W}}_{(1:i)}\|_F\|\widehat{\mathbf{R}}_{(1:i,1:i)}^{-1}\| \leq 5u_{crs}i^2m_p^2\text{cond}(\widehat{\mathbf{W}}_{(1:i)}) =: F_2.$$

By combining this property with (54), we get

$$(1 + \varepsilon)^{-1/2}(1 - \Delta_m - F_1) - F_2 \leq \sigma_{\min}(\widehat{\mathbf{Q}}_{(1:i)}) \leq \sigma_{\max}(\widehat{\mathbf{Q}}_{(1:i)}) \leq (1 - \varepsilon)^{-1/2}(1 + \Delta_m + F_1) + F_2.$$

We conclude that

$$(1 + \varepsilon)^{-1/2}(1 - F_3) \leq \sigma_{\min}(\widehat{\mathbf{Q}}_{(1:i)}) \leq \sigma_{\max}(\widehat{\mathbf{Q}}_{(1:i)}) \leq (1 - \varepsilon)^{-1/2}(1 + F_3), \quad (55)$$

where $F_3 := \Delta_m + F_1 + \sqrt{5/4}F_2 \leq 32u_{crs}i^2m_p^2\text{cond}(\widehat{\mathbf{W}}_{(1:i)})$, which in particular implies that $\|\widehat{\mathbf{Q}}_{(1:i)}\| \leq 1.6$ and $\sigma_{\min}(\widehat{\mathbf{Q}}_{(1:i)}) \geq 0.86$. Furthermore, from (52), we get

$$1 - F_4 \leq \sigma_{\min}(\widehat{\mathbf{S}}_{(1:i)}) - 0.1u_{crs} \leq \sigma_{\min}(\Theta \widehat{\mathbf{Q}}_{(1:i)}) \leq \sigma_{\max}(\Theta \widehat{\mathbf{Q}}_{(1:i)}) \leq \sigma_{\max}(\widehat{\mathbf{S}}_{(1:i)}) + 0.1u_{crs} \leq 1 + F_4, \quad (56)$$

where $F_4 := \Delta_m + 0.1u_{crs} \leq 20.1u_{crs}m_p^2i^2\text{cond}(\widehat{\mathbf{W}}_{(1:i)})$.

From (55) and (56) it is deduced that for any vector $\mathbf{a} \in \mathbb{R}^m$,

$$\begin{aligned} & \left| \|\widehat{\mathbf{Q}}_{(1:i)}\mathbf{a}\|^2 - \|\Theta \widehat{\mathbf{Q}}_{(1:i)}\mathbf{a}\|^2 \right| \\ & \leq \|\mathbf{a}\|^2 \max \{ (1 - \varepsilon)^{-1}(1 + F_3)^2 - (1 - F_4)^2, (1 + F_4)^2 - (1 + \varepsilon)^{-1}(1 - F_3)^2 \} \\ & \leq \|\mathbf{a}\|^2 \max \{ 1.45\varepsilon + 2.9F_3 + 2F_4, 1.34\varepsilon + 2.1F_4 + 2.1F_3 \} \\ & \leq 0.74\varepsilon' \|\mathbf{a}\|^2 \leq \varepsilon' \|\widehat{\mathbf{Q}}_{(1:i)}\mathbf{a}\|^2, \end{aligned}$$

where $\varepsilon' = 2\varepsilon + 180u_{crs}m_p^2i^2\text{cond}(\widehat{\mathbf{W}}_{(1:i)})$. By using the parallelogram identity this relation can be brought to form (2) where $V = \text{range}(\widehat{\mathbf{Q}}_{(1:i)})$ and $\varepsilon = \varepsilon'$. It is deduced that Θ is a ε' -embedding for $\widehat{\mathbf{Q}}_{(1:i)}$. \square

Proof of Lemma 4.4. By definitions of Θ , $\Pi_{\mathcal{K}}$ and $\Pi_{\mathcal{K}}^{\Theta}$, we have

$$\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{u}\|^2 \leq \|(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{u}\|^2 \leq \frac{1}{1 - \varepsilon}(\|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{u}\|^2) \leq \frac{1}{1 - \varepsilon}(\|\Theta(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{u}\|^2) \leq \frac{1 + \varepsilon}{1 - \varepsilon}(\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{u}\|^2),$$

which gives the first inequality. Furthermore, we also have the following relation

$$\|\Theta \Pi_{\mathcal{K}}^{\Theta} \mathbf{w}\|^2 = \|\Theta \mathbf{w}\|^2 - \|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{w}\|^2,$$

which implies that

$$\|\Theta \Pi_{\mathcal{K}}^{\Theta} \mathbf{w}\|^2 \leq (1 + \varepsilon)\|\mathbf{w}\|^2 - \|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{w}\|^2 \leq (1 + \varepsilon)\|\mathbf{w}\|^2 - (1 - \varepsilon)\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{w}\|^2 \leq \|\Pi_{\mathcal{K}}\mathbf{w}\|^2 + 2\varepsilon\|\mathbf{w}\|^2.$$

Similarly,

$$\|\Theta \Pi_{\mathcal{K}}^{\Theta} \mathbf{w}\|^2 \geq (1 - \varepsilon)\|\mathbf{w}\|^2 - \|\Theta(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{w}\|^2 \geq (1 - \varepsilon)\|\mathbf{w}\|^2 - (1 + \varepsilon)\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{w}\|^2 \geq \|\Pi_{\mathcal{K}}\mathbf{w}\|^2 - 2\varepsilon\|\mathbf{w}\|^2.$$

By using the relation,

$$\frac{1}{1 + \varepsilon}\|\Theta \Pi_{\mathcal{K}}^{\Theta} \mathbf{w}\|^2 \leq \|\Pi_{\mathcal{K}}^{\Theta} \mathbf{w}\|^2 \leq \frac{1}{1 - \varepsilon}\|\Theta \Pi_{\mathcal{K}}^{\Theta} \mathbf{w}\|^2,$$

we obtain (34b). Finally, if $\mathbf{w} \in \mathcal{K}$, then, by definition, $\Pi_{\mathcal{K}}^{\Theta} \mathbf{w} = \arg \min_{\mathbf{v} \in \mathcal{K}} \|\Theta(\mathbf{w} - \mathbf{v})\| = \mathbf{w}$, which finishes the proof. \square

Proof of Theorem 4.5. The proof directly follows that of [17, Theorem 4.3] replacing ℓ_2 -inner products and norms by the sketched ones. We have,

$$\begin{aligned}\|\Theta(\mathbf{A}_m - \lambda\mathbf{I})\Pi_{\mathcal{K}}^{\Theta}\mathbf{x}\| &= \|\Theta\Pi_{\mathcal{K}}^{\Theta}(\mathbf{A} - \lambda\mathbf{I})(\mathbf{x} - (\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x})\| = \|\Theta\Pi_{\mathcal{K}}^{\Theta}(\mathbf{A} - \lambda\mathbf{I})(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x}\| \\ &= \|\Theta\Pi_{\mathcal{K}}^{\Theta}(\mathbf{A} - \lambda\mathbf{I})(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x}\| \\ &\leq \gamma\|(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x}\|.\end{aligned}$$

Notice that $\mathbf{A}_m(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x} = \mathbf{0}$. Consequently,

$$(\mathbf{A}_m - \lambda\mathbf{I})\mathbf{x} = (\mathbf{A}_m - \lambda\mathbf{I})\Pi_{\mathcal{K}}^{\Theta}\mathbf{x} + (\mathbf{A}_m - \lambda\mathbf{I})(\mathbf{x} - \Pi_{\mathcal{K}}^{\Theta}\mathbf{x}) = (\mathbf{A}_m - \lambda\mathbf{I})\Pi_{\mathcal{K}}^{\Theta}\mathbf{x} + \lambda(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x}.$$

Furthermore, by using the fact that the two vectors on the right hand side are orthogonal with respect to the sketched inner product, we obtain

$$\|\Theta(\mathbf{A}_m - \lambda\mathbf{I})\mathbf{x}\|^2 = \|\Theta(\mathbf{A}_m - \lambda\mathbf{I})\Pi_{\mathcal{K}}^{\Theta}\mathbf{x}\|^2 + |\lambda|^2\|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x}\|^2 \leq (\gamma^2 + |\lambda|^2)\|(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{x}\|^2.$$

□

Proof of Theorem 4.6. We have,

$$\|\Theta\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\|^2 = \|\Theta\Pi_{\mathcal{K}}^{\Theta}\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\|^2 + \|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\|^2 = \|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\|^2.$$

Consequently, by an ε -embedding property of Θ , we obtain

$$\begin{aligned}\|\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\| &\leq \frac{1}{\sqrt{1-\varepsilon}}\|\Theta\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\| = \frac{1}{\sqrt{1-\varepsilon}}\|\Theta(\mathbf{I} - \Pi_{\mathcal{K}}^{\Theta})\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\| \leq \frac{1}{\sqrt{1-\varepsilon}}\|\Theta(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\| \\ &\leq \sqrt{\frac{1+\varepsilon}{1-\varepsilon}}\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{r}(\mathbf{u}, \mu\mathbf{u})\| = \sqrt{\frac{1+\varepsilon}{1-\varepsilon}}\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{A}\mathbf{u}\| = \sqrt{\frac{1+\varepsilon}{1-\varepsilon}}\|(\mathbf{I} - \Pi_{\mathcal{K}})\mathbf{A}\Pi_{\mathcal{K}}\mathbf{u}\|.\end{aligned}$$

which gives theorem's first inequality. Moreover, notice that

$$(\mathbf{I} - \Pi_{Q_{p-1}})\mathbf{A}\Pi_{Q_{p-1}}\mathbf{u} = (\mathbf{I} - \Pi_{Q_{p-1}})\mathbf{A}(\mathbf{I} - \Pi_{Q_{p-2}})\mathbf{u},$$

which gives theorem's second inequality.

□