THE NON-SYMMETRIC s -STEP LANCZOS ALGORITHM: DERIVATION OF EFFICIENT RECURRENCES AND SYNCHRONIZATION-REDUCING VARIANTS OF BiCG AND QMR

STEFAN FEUERRIEGEL ^{a,*}, H. MARTIN BÜCKER ^b

^aChair for Information Systems Research
University of Freiburg, Platz der Alten Synagoge, 79098 Freiburg, Germany
e-mail: stefan.feuerriegel@is.uni-freiburg.de

^bChair for Advanced Computing
Friedrich Schiller University Jena, Ernst-Abbe-Platz 2, 07743 Jena, Germany
e-mail: martin.buecker@uni-jena.de

The Lanczos algorithm is among the most frequently used iterative techniques for computing a few dominant eigenvalues of a large sparse non-symmetric matrix. At the same time, it serves as a building block within biconjugate gradient (BiCG) and quasi-minimal residual (QMR) methods for solving large sparse non-symmetric systems of linear equations. It is well known that, when implemented on distributed-memory computers with a huge number of processes, the synchronization time spent on computing dot products increasingly limits the parallel scalability. Therefore, we propose synchronization-reducing variants of the Lanczos, as well as BiCG and QMR methods, in an attempt to mitigate these negative performance effects. These so-called s -step algorithms are based on grouping dot products for joint execution and replacing time-consuming matrix operations by efficient vector recurrences. The purpose of this paper is to provide a rigorous derivation of the recurrences for the s -step Lanczos algorithm, introduce s -step BiCG and QMR variants, and compare the parallel performance of these new s -step versions with previous algorithms.

Keywords: synchronization-reducing, s -step Lanczos, s -step BiCG, s -step QMR, efficient recurrences.

1. Rethinking algorithm design

Current large-scale computer systems are sophisticated architectures based on multi- or manycore technology, with deep memory hierarchies and possible heterogeneity in the form of graphic or other coprocessors. For scientific and engineering applications, it is therefore currently challenging to achieve a high performance using these systems. Unfortunately, future extreme-scale computer systems are likely to become even more complex and so it will become increasingly hard to achieve a sustained performance that is somewhere near peak performance. It is widely recognized that there are various intricate challenges for future large-scale computing. Today, there is only a vague idea of how these future platforms will actually be built and how they will be programmed efficiently. Rather than summarizing this ongoing

discussion (Cappello *et al.*, 2009; Davis *et al.*, 2012; Duff, 2012; Shalf *et al.*, 2011), we focus on novel algorithmic techniques that will be required to fully exploit current large-scale and future exascale systems.

Existing technology trends indicate that algorithm designers will have to pay crucial attention in order to reduce data movement at various memory levels and to reduce synchronization at various system levels. While the communication costs associated with data movement have become an important issue in today's parallel algorithm design, the cost associated with synchronization does not currently receive adequate consideration. However, synchronization costs will soon outweigh communication costs as the degree of parallelism increases further. In fact, synchronization dictates the overall performance of various algorithms on current large-scale and, in particular, future exascale systems. Therefore, we address the problem of designing

*Corresponding author

synchronization-reducing variants of popular iterative Krylov subspace methods that are based on the Lanczos algorithm (Lanczos, 1950).

The contribution of the present paper is to derive a new variant of the s -step Lanczos algorithm (Kim and Chronopoulos, 1992) with a normalization scheme that improves numerical stability. In contrast to previous publications, which all lack details on how to obtain the different underlying recurrences, this paper provides a thorough derivation. Based on this s -step Lanczos algorithm, we then introduce synchronization-reducing variants of two Krylov methods for the solution of large sparse non-symmetric systems of linear equations. More precisely, we propose new synchronization-reducing variants of the biconjugate gradient (BiCG) method (Fletcher, 1976) and the quasi-minimal residual (QMR) one (Freund and Nachtigal, 1994), and assess their numerical stability and parallel scalability. This publication is an extended version of our previous conference papers (Feuerriegel and Bücker, 2013a; 2013b) with additional details on the derivation of the recurrences underlying the s -step Lanczos algorithm. We also present some new computational experiments.

In this article, we use the following notation. Given two vectors, their dot product $\mathbf{v}^T \mathbf{w}$ is denoted by $\langle \mathbf{v}, \mathbf{w} \rangle$. The zero vector of dimension n is given by $\mathbf{0}_n$. The symbols $\mathbf{0}_{n,m}$ and $\mathbf{I}_{n,m}$ are used for $n \times m$ zero and identity matrices, respectively. Concatenation of scalar entries that form a row vector is denoted by $[x_1, \dots, x_n]$. Concatenation of vectors or matrices that form a (block) matrix is indicated by $[\mathbf{v}_1 \parallel \dots \parallel \mathbf{v}_n]$.

The paper is organized as follows. After describing related work in Section 2, we review the classical Lanczos method in Section 3. In Section 4, we introduce a novel normalization scheme and the resulting normalized s -step Lanczos algorithm. The derivation of the new underlying recurrences is summarized in Section 5 and detailed in Appendix. Section 6 utilizes the s -step Lanczos algorithm to derive new synchronization-reducing variants of BiCG and QMR. In Section 7, the three s -step variants are compared to their classical versions in terms of both numerical stability and parallel performance.

2. Parallel Krylov methods

Non-symmetric eigenvalue problems arising from computational science and engineering are often large and sparse. When only a few dominant eigenvalues are required, Krylov subspace methods enter the picture. The Lanczos algorithm (Lanczos, 1950) is an archetype of this class of iterative methods. At the same time, it is an important building block of Krylov subspace methods for the solution of large sparse systems of linear equations.

When parallelizing the Lanczos or other Krylov subspace methods on message-passing architectures,

naïve approaches proceed by parallelizing each underlying linear algebra operation individually. However, the resulting parallel performance of such approaches is known to be limited by communication and synchronization. To overcome this impediment to parallel scalability, significant research effort is spent in designing new Krylov algorithms specifically for parallel computers. The long history of these methods is described in several surveys (Saad, 1989; van der Vorst, 1990; Demmel *et al.*, 1993; Duff and van der Vorst, 1999; Bücker, 2002). A broad classification of these parallel iterative methods is as follows:

- i. *Communication-overlapping* algorithms aim to reduce the impact of a communication event by overlapping it with computation and/or other communication (Ghysels *et al.*, 2013; Ghysels and Vanroose, 2014).
- ii. *Communication-avoiding* algorithms rely on blocking to reduce the volume of communication (Mohiyuddin *et al.*, 2009; Hoemmen, 2010; Gustafsson *et al.*, 2012a; Carson *et al.*, 2014).
- iii. *Synchronization-free* algorithms (Fischer and Freund, 1994) do not involve any global synchronization points (GSPs), defined as the locations of an algorithm at which all information local to a process has to be globally available for all processes in order to continue the computation.
- iv. *Synchronization-organizing* algorithms orchestrate synchronization in an attempt to curtail the negative effects caused by global synchronization, for instance, by handling synchronization hierarchically (Curfman McInnes *et al.*, 2014) or using non-blocking all-reduce operations (Kandalla *et al.*, 2012).
- v. *Synchronization-reducing* algorithms try to minimize the number of GSPs (Meurant, 1986; Van Rosendale, 1983; Bücker and Sauren, 1996; 1997; 1999; Zuo *et al.*, 2010; Zhu *et al.*, 2014).

While communication-avoiding algorithms successfully reduce the communication volume between processes, they do not directly focus on the synchronization between processes. However, synchronization will increasingly dominate the total execution time of future extreme-scale computer systems, in which the number of processes will be huge. Therefore, we focus on a novel synchronization-reducing Krylov algorithm. Here, a GSP is enforced by dot product-like operations involving a reduction operation on all participating processes. When only a single GSP is enforced for s iterations of the corresponding classical algorithm, this

synchronization-reducing algorithm is referred to as an s -step method (Chronopoulos, 1986; Chronopoulos and Gear, 1989; Chronopoulos and Swanson, 1996). The s -step Lanczos procedure was originally introduced for symmetric matrices (Kim and Chronopoulos, 1991) and later extended to non-symmetric ones (Kim and Chronopoulos, 1992).

3. Classical Lanczos method

The classical Lanczos algorithm (Lanczos, 1950) reduces a non-symmetric $N \times N$ matrix A to a tridiagonal form T_N . At the same time, it also produces two matrices,

$$V_N := [v_1 \parallel \dots \parallel v_N] \in \mathbb{R}^{N \times N} \quad (1)$$

and

$$W_N := [w_1 \parallel \dots \parallel w_N] \in \mathbb{R}^{N \times N}, \quad (2)$$

whose columns v_n and w_n are called Lanczos vectors. For the sake of notational simplicity, we assume here that the iteration proceeds up to step N whereas, in practice, it should stop after significantly fewer steps.

Definition 1. (*Classical Lanczos algorithm*) For a given non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, the Lanczos algorithm generates a tridiagonal matrix $T_N \in \mathbb{R}^{N \times N}$ and matrices $V_N \in \mathbb{R}^{N \times N}$ and $W_N \in \mathbb{R}^{N \times N}$ such that

$$W_N^T V_N = I_{N,N} \quad (\text{biorthonormality}), \quad (3)$$

$$AV_N = V_N T_N, \quad (4)$$

$$A^T W_N = W_N T_N^T. \quad (5)$$

The classical Lanczos algorithm summarized in pseudocode in Algorithm 1 is based on three-term recurrences. That is, the execution of n iterations of this algorithm generates the tridiagonal matrix,

$$T_n := \begin{bmatrix} \alpha_1 & \beta_2 & & \\ \gamma_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_n \\ & & \gamma_n & \alpha_n \end{bmatrix}, \quad (6)$$

which is the $n \times n$ leading principle submatrix of T_N . Thus, (4) indicates that the next Lanczos vector, v_{n+1} , involves a matrix-vector product with the matrix A and computations dependent solely on the two previous Lanczos vectors, v_n and v_{n-1} . The resulting three-term recurrences are given in Step 4 of this algorithm.

In this algorithm, there are two dot products in Steps 3 and 5, both of which enforce a GSP. In other words, each iteration requires two separate synchronizations of all processes that execute this algorithm. In particular, the result of the first GSP, α_n , needs to be available before the computation of the second GSP begins.

Algorithm 1. Classical Lanczos algorithm.

Input: Non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, as well as starting vectors $v_1, w_1 \in \mathbb{R}^N$, with $w_1^T v_1 = 1$.

Output: After n iterations, the algorithm returns a tridiagonal matrix $T_n = \text{tridiag}(\gamma, \alpha, \beta) \in \mathbb{R}^{n \times n}$ with diagonals $\gamma = (\gamma_2, \dots, \gamma_n)$, $\alpha = (\alpha_1, \dots, \alpha_n)$, and $\beta = (\beta_2, \dots, \beta_n)$, as well as the Lanczos basis $V_n = [v_1 \parallel \dots \parallel v_n] \in \mathbb{R}^{N \times n}$.

- 1: Initialize vectors $v_0 \leftarrow 0_N$ and $w_0 \leftarrow 0_N$ and set scalars $\beta_1 \leftarrow 0$ and $\gamma_1 \leftarrow 0$.
- 2: **for** $n = 1$ **until** Convergence **do**
- 3: Compute $\alpha_n = w_n^T A v_n$ with **global synchronization**.
- 4: Compute

$$\tilde{v}_{n+1} = A v_n - \alpha_n v_n - \beta_n v_{n-1},$$

$$\tilde{w}_{n+1} = A^T w_n - \alpha_n w_n - \gamma_n w_{n-1}.$$

- 5: Choose γ_{n+1} and β_{n+1} such that

$$\gamma_{n+1} \beta_{n+1} = \tilde{w}_{n+1}^T \tilde{v}_{n+1}$$

with **global synchronization**.

- 6: Scale the Lanczos basis via

$$v_{n+1} = \frac{1}{\gamma_{n+1}} \tilde{v}_{n+1} \quad \text{and} \quad w_{n+1} = \frac{1}{\beta_{n+1}} \tilde{w}_{n+1}.$$

- 7: **end for**

In practice, different versions of the classical Lanczos algorithm are typically preferred over Algorithm 1. One of the reasons is that Algorithm 1 allows the scaling of only one of the sequences of Lanczos vectors, either v_n or w_n . However, to control the numerical stability, one would like to scale both to, say,

$$\|v_n\|_2 = 1 \quad \text{and} \quad \|w_n\|_2 = 1. \quad (7)$$

This is accomplished by replacing the identity in (3) by a diagonal matrix whose nonzero elements are used to scale the second sequence of Lanczos vectors. In addition, there is another version of this algorithm that involves an LU decomposition of the tridiagonal matrix (6) leading to coupled two-term recurrences (Gutknecht, 1997).

4. s -Step Lanczos method

A single block iteration of the non-symmetric s -step Lanczos algorithm introduced by Kim and Chronopoulos (1992) generates s iterations of the classical Lanczos algorithm using only a single GSP. Rather than computing a pair of individual Lanczos vectors v_k and w_k , the k -th block iteration of the s -step Lanczos algorithm computes

a pair of blocks of s Lanczos vectors denoted by

$$\bar{V}_k := [\mathbf{v}_k^1 \parallel \dots \parallel \mathbf{v}_k^s] \in \mathbb{R}^{N \times s} \quad (8)$$

and

$$\bar{W}_k := [\mathbf{w}_k^1 \parallel \dots \parallel \mathbf{w}_k^s] \in \mathbb{R}^{N \times s}. \quad (9)$$

After presenting an overview of the algorithm, we introduce a normalization scheme and present the results of the underlying orthogonalization process.

4.1. High-level overview. The s -step Lanczos method proceeds in two steps. First, *relaxed* Lanczos vectors are computed in a block-wise fashion. In each block iteration, a new block containing s of these vectors as columns is computed. Second, a back transformation is applied to these vectors. The overall structure of the algorithm is schematically depicted as follows:

The following definition summarizes important facts of the s -step Lanczos algorithm using the block formulation of the Lanczos vectors (8) and (9), as well as a corresponding block notation for matrices.

Definition 2. (*s-Step Lanczos algorithm*) Let $n = sk$ with $1 \leq n \leq N$. For a given non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, the s -step Lanczos algorithm generates an upper Hessenberg matrix $\tilde{T}_n \in \mathbb{R}^{n \times n}$, as well as two additional matrices $\tilde{V}_n \in \mathbb{R}^{N \times n}$ and $\tilde{W}_n \in \mathbb{R}^{N \times n}$, such that

$$\tilde{W}_n^T \tilde{V}_n = \text{block biorthogonal}, \quad (10)$$

$$A \tilde{V}_n = \tilde{V}_n \tilde{T}_n + f_{k+1} \mathbf{v}_{k+1}^1 [0, \dots, 0, 1], \quad (11)$$

$$A^T \tilde{W}_n = \tilde{W}_n \tilde{T}_n + f_{k+1} \mathbf{w}_{k+1}^1 [0, \dots, 0, 1]. \quad (12)$$

The upper Hessenberg matrix $\tilde{T}_n$ is block tridiagonal,

$$\tilde{T}_n := \begin{bmatrix} \bar{G}_1 & \bar{E}_2 & & \\ \bar{F}_2 & \bar{G}_2 & \ddots & \\ & \ddots & \ddots & \bar{E}_k \\ & & \bar{F}_k & \bar{G}_k \end{bmatrix} \in \mathbb{R}^{n \times n}, \quad (13)$$

with a nonzero in the upper right corner of the block,

$$\bar{F}_i = \begin{bmatrix} & f_i \end{bmatrix} \in \mathbb{R}^{s \times s}, \quad (14)$$

while each $\bar{E}_i$ is a dense $s \times s$ matrix, and each $\bar{G}_i \in \mathbb{R}^{s \times s}$ is in upper Hessenberg form. The block-wise grouping of the relaxed Lanczos vectors is given by

$$\tilde{V}_n := [\bar{V}_1 \parallel \dots \parallel \bar{V}_k] \in \mathbb{R}^{N \times n}, \quad (15)$$

$$\tilde{W}_n := [\bar{W}_1 \parallel \dots \parallel \bar{W}_k] \in \mathbb{R}^{N \times n}. \quad (16)$$

When implemented, the s -step Lanczos algorithm iterates $k = n/s$ block iterations yielding both $\tilde{T}_n$ and $\tilde{V}_n$. It then turns $\tilde{T}_n$ and $\tilde{V}_n$ into the matrices T_n and V_n from the classical Lanczos procedure. This back transformation is sketched in the following theorem.

Theorem 1. Let $\tilde{W}_n^T \tilde{V}_n$ be a non-singular matrix and let

$$\tilde{W}_n^T \tilde{V}_n = \tilde{L}_n \tilde{U}_n \quad (17)$$

denote its LU decomposition. Then, $\tilde{T}_n$, $\tilde{V}_n$ and $\tilde{W}_n$ can be transformed into T_n , W_n and V_n , originating from the classical Lanczos method in the absence of breakdowns:

$$T_n = \tilde{U}_n \tilde{T}_n \tilde{U}_n^{-1}, \quad (18)$$

$$V_n = \tilde{V}_n \tilde{U}_n^{-1}, \quad (19)$$

$$W_n^T = \tilde{L}_n^{-1} \tilde{W}_n^T. \quad (20)$$

Proof. See the work of Kim and Chronopoulos (1992). ■

4.2. Normalization scheme. The original s -step Lanczos algorithm (Kim and Chronopoulos, 1992) implemented in double precision floating-point arithmetic can involve a numerical overflow. More precisely, we observed that the floating-point values in $\tilde{W}_n^T \tilde{V}_n$ can grow rapidly. To reduce the possibility of numerical overflow, we introduce the normalization scheme

$$\tilde{\mathbf{v}}_{k+1}^1 := f_{k+1} \mathbf{v}_{k+1}^1, \quad (21)$$

$$\tilde{\mathbf{w}}_{k+1}^1 := f_{k+1} \mathbf{w}_{k+1}^1, \quad (22)$$

where

$$f_{k+1} := \sqrt{|\langle \tilde{\mathbf{w}}_{k+1}^1, \tilde{\mathbf{v}}_{k+1}^1 \rangle|}. \quad (23)$$

Thus, we have

$$\langle \mathbf{w}_{k+1}^1, \mathbf{v}_{k+1}^1 \rangle = \pm 1. \quad (24)$$

This differs from the version of Kim and Chronopoulos (1992), in which any normalization is avoided, corresponding to $f_{k+1} := 1$ in the new scheme. In addition to this normalization scheme, it turns out that it is also convenient to normalize $\tilde{\mathbf{v}}_{k+1}^j$ and $\tilde{\mathbf{w}}_{k+1}^j$, such that

$$\langle \mathbf{w}_{k+1}^j, \mathbf{v}_{k+1}^j \rangle = \pm 1 \quad \text{for } j = 2, \dots, s \quad (25)$$

holds. Therefore, we scale the Lanczos vectors by

$$\mathbf{v}_{k+1}^j = \sigma_{k+1}^j \tilde{\mathbf{v}}_{k+1}^j \quad \text{for } j = 2, \dots, s \quad (26)$$

$$\mathbf{w}_{k+1}^j = \sigma_{k+1}^j \tilde{\mathbf{w}}_{k+1}^j \quad \text{for } j = 2, \dots, s, \quad (27)$$

where the scaling factors are given by

$$\sigma_{k+1}^j := |\langle \tilde{\mathbf{w}}_{k+1}^j, \tilde{\mathbf{v}}_{k+1}^j \rangle|^{-\frac{1}{2}} \quad \text{for } j = 2, \dots, s. \quad (28)$$

This normalization scheme requires a new derivation of the underlying recurrences in Section 5 since these differ from the recurrences of Kim and Chronopoulos (1992). The resulting s -step Lanczos algorithm is depicted in Algorithm 2. Here, each block iteration k generates two blocks of s Lanczos vectors. The first Lanczos vector in each block is computed by a recurrence involving Lanczos vectors from the two previous blocks shown in (30) and (31). The remaining $s - 1$ Lanczos vectors in each block are computed in (32) and (33). The computation of the Lanczos vectors involves the coefficient vectors $\mathbf{g}_k^s$, $\mathbf{e}_k^s$ and $\mathbf{t}_k^2, \mathbf{t}_k^3, \dots, \mathbf{t}_k^s$ that need to be chosen to satisfy the block biorthogonality condition (10). The LU decomposition in Step 16, as well as the back transformation in Step 17, work on the data from the current block iteration. For instance, since by (10) the matrix $\tilde{W}_{sk}^T \tilde{V}_{sk}$ is block diagonal, only an LU decomposition of the $s \times s$ matrix

$$\overline{M}_k := \overline{W}_k^T \overline{V}_k \quad (29)$$

has to be carried out in Step 16.

Finally, consider a block iteration of Algorithm 2 that computes a pair of s Lanczos vectors. We emphasize that each block iteration requires only a single GSP in Step 7.

4.3. Orthogonalizing the s -step Lanczos basis. The coefficient vectors $\mathbf{e}_k^s$, $\mathbf{g}_k^i$, and $\mathbf{t}_k^j$ are chosen to satisfy the biorthogonality (10). This is achieved by solving systems of linear equations with the coefficient matrix (29) that is assumed to be non-singular. The following theorem states the properties of the coefficient vectors.

Theorem 2. *The vectors $\mathbf{e}_k^i$, $\mathbf{g}_k^i$ and $\mathbf{t}_k^j$ with $i = 1, \dots, s$ and $j = 2, \dots, s$ are given by the solutions of the following $s \times s$ systems of linear equations:*

$$\overline{M}_{k-1} \mathbf{e}_k^i = \mathbf{c}_k^i,$$

where

$$\mathbf{c}_k^i := [\langle \mathbf{w}_{k-1}^1, A \mathbf{v}_k^i \rangle, \dots, \langle \mathbf{w}_{k-1}^s, A \mathbf{v}_k^i \rangle]^T, \quad (34)$$

$$\overline{M}_k \mathbf{g}_k^i = \mathbf{d}_k^i,$$

where

$$\mathbf{d}_k^i := [\langle \mathbf{w}_k^1, A \mathbf{v}_k^i \rangle, \dots, \langle \mathbf{w}_k^s, A \mathbf{v}_k^i \rangle]^T, \quad (35)$$

$$\overline{M}_k \mathbf{t}_k^j = \mathbf{b}_k^j,$$

where

$$\mathbf{b}_k^j := [\langle \mathbf{w}_k^1, A^{j-1} \mathbf{v}_{k+1}^1 \rangle, \dots, \langle \mathbf{w}_k^s, A^{j-1} \mathbf{v}_{k+1}^1 \rangle]^T. \quad (36)$$

Algorithm 2. Synchronization-reducing s -step Lanczos.

Input: Non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, starting vectors $\mathbf{v}_1^1, \mathbf{w}_1^1 \in \mathbb{R}^N$ with $\mathbf{v}_1^1 = \mathbf{w}_1^1$ and parameter s .

Output: After $k = n/s$ block iterations, return a tridiagonal matrix $T_n \in \mathbb{R}^{n \times n}$ and Lanczos basis $V_n \in \mathbb{R}^{N \times n}$.

1: Initialize $\overline{V}_0 \leftarrow 0_{N,s}$ and $\overline{W}_0 \leftarrow 0_{N,s}$ and compute

$$\overline{V}_1 \leftarrow [\mathbf{v}_1^1 \parallel A \mathbf{v}_1^1 \parallel \dots \parallel A^{s-1} \mathbf{v}_1^1],$$

$$\overline{W}_1 \leftarrow [\mathbf{w}_1^1 \parallel A^T \mathbf{w}_1^1 \parallel \dots \parallel (A^T)^{s-1} \mathbf{w}_1^1].$$

2: Compute dot products $\langle \mathbf{w}_1^1, A^j \mathbf{v}_1^1 \rangle$ for $j = 0, \dots, 2s - 1$.

3: **for** $k = 1$ **until** Convergence **do**

4: Compute $\overline{M}_k^{i,j}$, $\mathbf{c}_k^s$ and $\mathbf{d}_k^{i,j}$ for $i = 1, \dots, s$ and $j = 1, \dots, s$ using efficient recurrences from Theorem 3.

5: Solve $\overline{M}_{k-1} \mathbf{e}_k^s = \mathbf{c}_k^s$ and $\overline{M}_k \mathbf{g}_k^s = \mathbf{d}_k^s$.

6: Orthogonalize $\tilde{\mathbf{v}}_{k+1}^1$ against $\overline{W}_k$ and $\tilde{\mathbf{w}}_{k+1}^1$ against $\overline{V}_k$ by

$$\tilde{\mathbf{v}}_{k+1}^1 \leftarrow A \mathbf{v}_k^s - \overline{V}_{k-1} \mathbf{e}_k^s - \overline{V}_k \mathbf{g}_k^s, \quad (30)$$

$$\tilde{\mathbf{w}}_{k+1}^1 \leftarrow A^T \mathbf{w}_k^s - \overline{W}_{k-1} \mathbf{e}_k^s - \overline{W}_k \mathbf{g}_k^s. \quad (31)$$

7: Compute $2s$ dot products $\langle \tilde{\mathbf{w}}_{k+1}^1, A^j \tilde{\mathbf{v}}_{k+1}^1 \rangle$ for $j = 0, \dots, 2s - 1$ with **global synchronization**.

8: Compute normalization coefficient

$$f_{k+1} \leftarrow \sqrt{|\langle \tilde{\mathbf{w}}_{k+1}^1, \tilde{\mathbf{v}}_{k+1}^1 \rangle|}.$$

9: Carry out normalization

$$\mathbf{v}_{k+1}^1 \leftarrow \tilde{\mathbf{v}}_{k+1}^1 / f_{k+1} \quad \text{and} \quad \mathbf{w}_{k+1}^1 \leftarrow \tilde{\mathbf{w}}_{k+1}^1 / f_{k+1}.$$

10: Normalize, for $j = 0, \dots, 2s - 1$, via

$$\langle \mathbf{w}_{k+1}^1, A^j \mathbf{v}_{k+1}^1 \rangle \leftarrow \langle \tilde{\mathbf{w}}_{k+1}^1, A^j \tilde{\mathbf{v}}_{k+1}^1 \rangle / f_{k+1}^2.$$

11: Compute $\mathbf{b}_k^j$ for $j = 2, \dots, s + 1$ using efficient recurrences from Theorem 3.

12: Solve $\overline{M}_k \mathbf{t}_k^j = \mathbf{b}_k^j$ for $j = 2, \dots, s$.

13: Orthogonalize $[\mathbf{v}_{k+1}^1 \parallel \dots \parallel A^{s-1} \mathbf{v}_{k+1}^1]$ against $\overline{W}_k$ and the block $[A^T \mathbf{w}_{k+1}^1 \parallel \dots \parallel (A^T)^{s-1} \mathbf{w}_{k+1}^1]$ against $\overline{V}_k$ by

$$\tilde{\mathbf{v}}_{k+1}^j \leftarrow A^{j-1} \mathbf{v}_{k+1}^1 - \overline{V}_k \mathbf{t}_k^j, \quad j = 2, \dots, s \quad (32)$$

$$\tilde{\mathbf{w}}_{k+1}^j \leftarrow (A^T)^{j-1} \mathbf{w}_{k+1}^1 - \overline{W}_k \mathbf{t}_k^j, \quad j = 2, \dots, s. \quad (33)$$

14: Compute normalization coefficients σ_{k+1}^j for $j = 2, \dots, s$ using efficient recurrences from Theorem 3.

15: Normalize, for $j = 2, \dots, s$, via

$$\mathbf{v}_{k+1}^j \leftarrow \tilde{\mathbf{v}}_{k+1}^j \sigma_{k+1}^j \quad \text{and} \quad \mathbf{w}_{k+1}^j \leftarrow \tilde{\mathbf{w}}_{k+1}^j \sigma_{k+1}^j.$$

16: Update LU decomposition of $\tilde{L}_{sk} \tilde{U}_{sk} \leftarrow \tilde{W}_{sk}^T \tilde{V}_{sk}$.

17: Update back transformation

$$T_{sk} \leftarrow \tilde{U}_{sk} \tilde{T}_{sk} \tilde{U}_{sk}^{-1}, \quad V_{sk} \leftarrow \tilde{V}_{sk} \tilde{U}_{sk}^{-1},$$

$$W_{sk}^T \leftarrow \tilde{L}_{sk}^{-1} \tilde{W}_{sk}^T.$$

18: **end for**

Proof. Given by Feuerriegel and Bucker (2013a), as well as Kim and Chronopoulos (1992). ■

The solution of these small and dense linear systems is computationally affordable. The next section summarizes the implementation details for setting up the coefficient matrices and the right-hand sides.

5. Deriving efficient recurrence equations

Computing the right-hand side vectors $\mathbf{b}_k^j$, $\mathbf{c}_k^j$, $\mathbf{d}_k^j$ and the coefficient matrix $\overline{\mathbf{M}}_k$ explicitly for all relevant j 's is not only a computationally expensive task, but also involves a considerable number of dot products. In fact, the evaluation of these dot products would destroy the benefits of the s -step approach instantly. Fortunately, there is a remedy to this problem which consists in retrieving these products recursively from the $2s$ dot products $\langle \mathbf{w}_k^1, \mathbf{v}_k^1 \rangle, \langle \mathbf{w}_k^1, A\mathbf{v}_k^1 \rangle, \dots, \langle \mathbf{w}_k^1, A^{2s-1}\mathbf{v}_k^1 \rangle$, as stated in the following theorem. These dot products introduce the only GSP per block iteration.

Throughout this article, the symbols $\mathbf{b}_k^{i,j}$, $\mathbf{c}_k^{i,j}$, $\mathbf{d}_k^{i,j}$, $\mathbf{g}_k^{i,j}$ and $\mathbf{t}_k^{i,j}$ denote the i -th element of the vectors $\mathbf{b}_k^j$, $\mathbf{c}_k^j$, $\mathbf{d}_k^j$, $\mathbf{g}_k^j$ and $\mathbf{t}_k^j$, respectively. The symbol $\overline{\mathbf{M}}_k^{i,j}$ denotes the matrix element of $\overline{\mathbf{M}}_k$ at row i and column j .

Theorem 3. (Efficient recurrences) *The scaling factors (28), as well as the coefficient matrices and the right-hand sides involved in Theorem 2, are given as follows, where we define $\sigma_k^1 := 1$ and use the indices $i = 1, \dots, s$:*

$$\begin{aligned} \sigma_k^j &= |\langle \mathbf{w}_k^1, A^{2j-2}\mathbf{v}_k^1 \rangle - (\mathbf{t}_{k-1}^j)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j|^{-\frac{1}{2}} \\ &\quad \text{for } j = 2, \dots, s, \\ \overline{\mathbf{M}}_k^{i,j} &= \sigma_k^i \sigma_k^j [\langle \mathbf{w}_k^1, A^{i+j-2}\mathbf{v}_k^1 \rangle - (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j] \\ &\quad \text{for } j = 1, \dots, s, \\ \mathbf{b}_k^{i,j} &= \frac{\sigma_k^i}{\sigma_k^s} \left[\langle f_{k+1} \mathbf{w}_{k+1}^1, A^{i+j-s-2}\mathbf{v}_{k+1}^1 \rangle \right. \\ &\quad \left. + \sum_{\substack{\ell=2s \\ +3-i-j}}^s \frac{\sigma_k^\ell \mathbf{g}_k^{\ell,s} \mathbf{b}_k^{i-s+\ell-1,j}}{\sigma_k^{i-s+\ell-1}} \right] \text{ for } j = 2, \dots, s+1, \\ \mathbf{c}_k^{j,s} &= 0 \quad \text{for } j = 1, \dots, s-1, \\ \mathbf{c}_k^{s,s} &= \sigma_k^s \left[\mathbf{b}_{k-1}^{s,s+1} - \left[\mathbf{d}_{k-1}^{s,1}, \dots, \mathbf{d}_{k-1}^{s,s} \right] \mathbf{t}_{k-1}^s \right], \\ \mathbf{d}_k^{i,j} &= \sigma_k^i \sigma_k^j \left[\langle \mathbf{w}_k^1, A^{i+j-1}\mathbf{v}_k^1 \rangle \right. \\ &\quad \left. - (\mathbf{t}_{k-1}^{i+1})^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j - (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^{j+1} \right. \\ &\quad \left. + (\mathbf{t}_{k-1}^i)^T [\mathbf{d}_{k-1}^1 \parallel \dots \parallel \mathbf{d}_{k-1}^s] \mathbf{t}_{k-1}^{j+1} \right] \\ &\quad \text{for } j = 1, \dots, s. \end{aligned}$$

Proof. See Appendix. ■

Table 1 compares the main computational cost for $n = sk$ iterations of the classical Lanczos algorithm,

Algorithm 1, and k block iterations of two s -step variants. The first s -step variant is the one without normalization introduced by Kim and Chronopoulos (1992), while the second is the one proposed in Algorithm 2. Recall that the latter requires only a single GSP per block iteration. This table reports the number of operations, as well as vector storages of size N , neglecting all corresponding costs of vectors of dimension s . Though the s -step variants slightly raise the computational cost, they reduce the number of GSPs by a factor of $\mathcal{O}(s)$.

6. s -Step BiCG and QMR methods

The Lanczos algorithm is now used to design Krylov subspace methods for the solution of linear systems. The aim of this section is to derive new synchronization-reducing s -step variants of BiCG and QMR.

6.1. Solving linear systems using the Lanczos basis.

The Lanczos basis $V_n \in \mathbb{R}^{N \times n}$ generated by any Lanczos algorithm can be used to iteratively solve non-symmetric systems of linear equations,

$$A\mathbf{x} = \mathbf{b},$$

where $\mathbf{x}, \mathbf{b} \in \mathbb{R}^N$. Given an initial guess $\mathbf{x}_0$ to the exact solution $\mathbf{x}$, the current approximation is given by

$$\mathbf{x}_n = \mathbf{x}_0 + V_n \mathbf{z}_n, \quad (37)$$

where the coefficient vector $\mathbf{z}_n \in \mathbb{R}^n$ needs to be determined. For reasons of numerical stability, we will also use another basis,

$$P_n = [\mathbf{p}_1 \parallel \dots \parallel \mathbf{p}_n] \in \mathbb{R}^{N \times n},$$

defined by

$$V_n = P_n U_n.$$

Here, U_n is the $n \times n$ unit upper triangular factor from the LU decomposition of the $(n+1) \times n$ tridiagonal matrix,

$$\underline{T}_n := \begin{bmatrix} T_n & \\ 0 & \dots & 0 & \gamma_{n+1} \end{bmatrix},$$

which is obtained from appending another row at the bottom of the tridiagonal matrix T_n defined in (6). This LU decomposition,

$$\begin{aligned} L_n U_n &= \begin{bmatrix} \tau_1 & & & & \\ \omega_2 & \tau_2 & & & \\ & \ddots & \ddots & & \\ & & \omega_n & \tau_n & \\ & & & \omega_{n+1} & \end{bmatrix} \begin{bmatrix} 1 & \mu_2 & & & \\ & 1 & \ddots & & \\ & & \ddots & \mu_n & \\ & & & 1 & \\ & & & & 1 \end{bmatrix}, \end{aligned} \quad (38)$$

Table 1. Comparison of cost for $n = sk$ iterations of the classical Lanczos algorithms and k block iterations of the synchronization-reducing s -step Lanczos variants.

Operation/Storage	Classical (Algorithm 1)	Unnormalized s -step (Kim and Chronopoulos, 1992)	Normalized s -step (Algorithm 2)
Dot products (of length N)	$2 sk$	$2 sk$	$2 sk$
Vector updates (of length N)	$6 sk$	$2 s(s-1)k + 4 sk$	$2 s(s-1)k + 8 sk - k$
Matrix-vector products	$2 sk$	$2 sk + k$	$2 sk + k$
Synchronization points	$\mathcal{O}(sk)$	k	k
Back transformations	—	$2k$	$2k$
Vector storage (of length N)	4	$4 s$	$4 s$

is computed for $n \geq 2$ via

$$\omega_n = \gamma_n, \quad (39)$$

$$\mu_n = \beta_n / \tau_{n-1}, \quad (40)$$

$$\tau_n = \alpha_n - \omega_n \mu_n \quad \text{with } \tau_1 := \alpha_1. \quad (41)$$

Recall from Section 3 that this LU decomposition is the basic building block of the Lanczos variant based on coupled two-term recurrences. We will use this variant in the remainder of the paper because it has a better reputation in terms of numerical stability.

From the structure of U_n according to (38), we find

$$\mathbf{p}_n = \mathbf{v}_n - \mu_n \mathbf{p}_{n-1}. \quad (42)$$

Furthermore, the relation of this formulation based on coupled-two term recurrences to the corresponding one based on three-term recurrences is given by

$$AP_n = V_{n+1} L_n. \quad (43)$$

Now, we replace the vector $\mathbf{z}_n$ in (37) by the n -dimensional coefficient vector

$$\mathbf{y}_n := U_n \mathbf{z}_n$$

and arrive at

$$\mathbf{x}_n = \mathbf{x}_0 + P_n \mathbf{y}_n. \quad (44)$$

The residual vector defined by

$$\mathbf{r}_n := \mathbf{b} - A \mathbf{x}_n$$

is then, using (43), transformed into

$$\begin{aligned} \mathbf{r}_n &= \mathbf{r}_0 - AP_n \mathbf{y}_n = \omega_1 \mathbf{v}_1 - V_{n+1} L_n \mathbf{y}_n \\ &= V_{n+1} \left([\omega_1, 0, \dots, 0]^T - L_n \mathbf{y}_n \right), \end{aligned} \quad (45)$$

where $\mathbf{v}_1 := \mathbf{r}_0 / \omega_1$ with a scaling factor $\omega_1 := \|\mathbf{r}_0\|_2$.

6.2. Synchronization-reducing s -step biconjugate gradient algorithm. The idea of BiCG (Fletcher, 1976) is to determine the coefficient vector $\mathbf{y}_n$ by zeroing out the first n entries of the vector in parentheses in (45). More

precisely, if the symbol $\bullet$ denotes an arbitrary value in the $(n+1)$ -th component of that vector, we require that this vector satisfy

$$[\omega_1, 0, \dots, 0]^T - L_n \mathbf{y}_n = [0, \dots, 0, \bullet]^T. \quad (46)$$

Let $\mathbf{y}_n = [\kappa_1, \dots, \kappa_n]^T$ denote the entries of the coefficient vector. Inserting this ansatz into (45) gives

$$\begin{aligned} \mathbf{r}_n &= V_{n+1} [0, \dots, 0, -\omega_{n+1} \kappa_n]^T \\ &= -\omega_{n+1} \kappa_n \mathbf{v}_{n+1}. \end{aligned}$$

Assuming that (46) is fulfilled in the first $n-1$ components, we find for the $(n-1)$ -th component of (46) that

$$0 = -\omega_n \kappa_{n-1} - \tau_n \kappa_n.$$

Hence

$$\kappa_n = -\frac{\omega_n \kappa_{n-1}}{\tau_n} \quad \text{with } \kappa_0 := -1. \quad (47)$$

The process of fixing $\mathbf{y}_n$ is easily updated in each iteration step because $\mathbf{y}_{n-1}$ coincides with the first $n-1$ components of $\mathbf{y}_n$; more precisely,

$$\mathbf{y}_n = \begin{bmatrix} \mathbf{y}_{n-1} \\ \kappa_n \end{bmatrix}.$$

Thus, from (44), the n -th approximation is obtained via

$$\mathbf{x}_n = \mathbf{x}_0 + P_{n-1} \mathbf{y}_{n-1} + \kappa_n \mathbf{p}_n = \mathbf{x}_{n-1} + \kappa_n \mathbf{p}_n. \quad (48)$$

We now let Algorithm 2 generate the Lanczos basis V_n and put together (39)–(41), (42), (47), and (48). The pseudocode of the resulting synchronization-reducing s -step BiCG is presented in Algorithm 3. The k -th block iteration of the s -step BiCG algorithm computes s Lanczos vectors and updates the LU decomposition of $\underline{T}_n$ from (38) with $n = sk$. Throughout the s -step BiCG algorithm, the call to the s -step Lanczos method in Step 3 introduces the only GSP per block iteration.

6.3. Synchronization-reducing s -step quasi-minimal residual algorithm. The idea of QMR (Freund and Nachtigal, 1991; 1994) is to find the coefficient vector

Algorithm 3. Synchronization-reducing s -step BiCG.

Input: Non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, right-hand side $\mathbf{b} \in \mathbb{R}^N$, initial guess $\mathbf{x}_0 \in \mathbb{R}^N$ and parameter s .

Output: After $k = n/s$ block iterations, the algorithm returns an approximation $\mathbf{x}_n$ to the solution $A^{-1}\mathbf{b}$.

- 1: Initialize vectors $\mathbf{r}_0 \leftarrow \mathbf{b} - A\mathbf{x}_0$, $\mathbf{v}_1 \leftarrow \mathbf{r}_0 / \|\mathbf{r}_0\|_2$, $\mathbf{p}_1 \leftarrow \mathbf{v}_1$ and set scalars $\kappa_0 \leftarrow -1$, $\omega_1 \leftarrow \|\mathbf{r}_0\|_2$, $\mu_1 \leftarrow 0$, and $\tau_1 \leftarrow \alpha_1$ where α_1 is taken from the Lanczos process initialized with $\mathbf{w}_1^1 \leftarrow \mathbf{v}_1^1$.
- 2: **for** $k = 1$ **until** Convergence **do**
- 3: Compute next s Lanczos vectors $\mathbf{v}_{s(k-1)+1}, \dots, \mathbf{v}_{sk}$ as well as next entries in $\underline{T}_{sk}$, i.e., $\alpha_{s(k-1)+1}, \dots, \alpha_{sk}$, $\beta_{s(k-1)+1}, \dots, \beta_{sk}$, $\gamma_{s(k-1)+1}, \dots, \gamma_{sk}$.
- 4: **for** $n = s(k-1) + 1$ **to** sk **do**
- 5: Update LU decomposition of $\underline{T}_{sk}$ via $\omega_n \leftarrow \gamma_n$, $\mu_n \leftarrow \beta_n / \tau_{n-1}$, $\tau_n \leftarrow \alpha_n - \omega_n \mu_n$.
- 6: Compute vector $\mathbf{p}_n \leftarrow \mathbf{v}_n - \mu_n \mathbf{p}_{n-1}$.
- 7: Set $\kappa_n \leftarrow -\omega_n \kappa_{n-1} / \tau_n$ and compute approximation $\mathbf{x}_n \leftarrow \mathbf{x}_{n-1} + \kappa_n \mathbf{p}_n$.
- 8: **end for**
- 9: **end for**

$\mathbf{y}_n$ by minimizing the Euclidean norm of the vector in parentheses in (45). To solve these least-squares problems, we follow Sauren and Bückner (1998) as stated in the subsequent Theorem 4.

Theorem 4. The unique solution of

$$\mathbf{y}_n := \arg \min_{\mathbf{y} \in \mathbb{R}^n} \left\| [\omega_1, 0, \dots, 0]^T - L_n \mathbf{y} \right\|_2 \quad (49)$$

is given by the recurrences

$$\mathbf{y}_n = \begin{bmatrix} \mathbf{y}_{n-1} \\ 0 \end{bmatrix} + \mathbf{g}_n, \quad (50)$$

$$\mathbf{g}_n = \theta_n \begin{bmatrix} \mathbf{g}_{n-1} \\ 0 \end{bmatrix} + [0, \dots, 0, \xi_n]^T, \quad n \geq 2, \quad (51)$$

with $\mathbf{y}_1 = \mathbf{g}_1 = [\xi_1]$, where

$$\theta_n = \frac{|\tau_n|^2 (1 - \sigma_n)}{\sigma_n |\tau_n|^2 + |\omega_{n+1}|^2}, \quad (52)$$

$$\xi_n = -\frac{\omega_n \bar{\tau}_n \xi_{n-1}}{\sigma_n |\tau_n|^2 + |\omega_{n+1}|^2}, \quad \xi_0 = -1, \quad (53)$$

$$\sigma_{n+1} = \frac{\sigma_n |\tau_n|^2}{\sigma_n |\tau_n|^2 + |\omega_{n+1}|^2}, \quad \sigma_1 = 1. \quad (54)$$

The current approximation is given by

$$\mathbf{x}_n = \mathbf{x}_{n-1} + \mathbf{d}_n, \quad (55)$$

where $\mathbf{d}_n = \theta_n \mathbf{d}_{n-1} + \xi_n \mathbf{p}_n$.

Again, we let Algorithm 2 generate the Lanczos basis V_n and put together the equations from Theorem 4. The resulting s -step QMR algorithm is given by the pseudo-code in Algorithm 4. Once more, the invocation

Algorithm 4. Synchronization-reducing s -step QMR.

Input: Non-symmetric matrix $A \in \mathbb{R}^{N \times N}$, right-hand side $\mathbf{b} \in \mathbb{R}^N$, initial guess $\mathbf{x}_0 \in \mathbb{R}^N$ and parameter s .

Output: After $k = n/s$ block iterations, the algorithm returns an approximation $\mathbf{x}_n$ to the solution $A^{-1}\mathbf{b}$.

- 1: Initialize vectors $\mathbf{r}_0 \leftarrow \mathbf{b} - A\mathbf{x}_0$, $\mathbf{v}_1 \leftarrow \mathbf{r}_0 / \|\mathbf{r}_0\|_2$, $\mathbf{p}_1 \leftarrow \mathbf{v}_1$, $\mathbf{d}_0 \leftarrow \mathbf{0}_N$ and set scalars $\xi_0 \leftarrow -1$, $\sigma_1 \leftarrow 1$, $\omega_1 \leftarrow \|\mathbf{r}_0\|_2$, $\mu_1 \leftarrow 0$, and $\tau_1 \leftarrow \alpha_1$, where α_1 is taken from the Lanczos process initialized with $\mathbf{w}_1^1 \leftarrow \mathbf{v}_1^1$.
- 2: **for** $k = 1$ **until** Convergence **do**
- 3: Compute next s Lanczos vectors $\mathbf{v}_{s(k-1)+1}, \dots, \mathbf{v}_{sk}$, as well as next entries in $\underline{T}_{sk}$, i.e., $\alpha_{s(k-1)+1}, \dots, \alpha_{sk}$, $\beta_{s(k-1)+1}, \dots, \beta_{sk}$, $\gamma_{s(k-1)+1}, \dots, \gamma_{sk}$.
- 4: **for** $n = s(k-1) + 1$ **to** sk **do**
- 5: Update LU decomposition of $\underline{T}_{sk}$ via $\omega_n \leftarrow \gamma_n$, $\mu_n \leftarrow \beta_n / \tau_{n-1}$, $\tau_n \leftarrow \alpha_n - \omega_n \mu_n$.
- 6: Compute vector $\mathbf{p}_n \leftarrow \mathbf{v}_n - \mu_n \mathbf{p}_{n-1}$.
- 7: $\theta_n \leftarrow \frac{|\tau_n|^2 (1 - \sigma_n)}{\sigma_n |\tau_n|^2 + |\omega_{n+1}|^2}$,
 $\xi_n \leftarrow -\frac{\omega_n \bar{\tau}_n \xi_{n-1}}{\sigma_n |\tau_n|^2 + |\omega_{n+1}|^2}$,
 $\sigma_{n+1} \leftarrow \frac{\sigma_n |\tau_n|^2}{\sigma_n |\tau_n|^2 + |\omega_{n+1}|^2}$.
- 8: Compute $\mathbf{d}_n \leftarrow \theta_n \mathbf{d}_{n-1} + \xi_n \mathbf{p}_n$ and approximation $\mathbf{x}_n \leftarrow \mathbf{x}_{n-1} + \mathbf{d}_n$.
- 9: **end for**
- 10: **end for**

of the s -step Lanczos subroutine in Step 3 represents the only GSP per block iteration.

7. Numerical experiments

This section compares classical and synchronization-reducing s -step variants of Lanczos, BiCG and QMR in terms of numerical accuracy and parallel performance, using an example by Freund and Nachtigal (1991). We consider the differential equation

$$-\Delta u + 40(xu_x + yu_y + zu_z) - 250u = f \quad (56)$$

on $\Omega = (0, 1) \times (0, 1) \times (0, 1)$ with $u = 0$ on the boundary. Using first-order centered differences and $\sqrt[3]{N}$ discretization points in each direction, we arrive at a linear system of order N . The right-hand side $f(x, y, z)$ is determined such that the vector of all ones is the exact solution. The initial guess is set to $\mathbf{x}_0 := \mathbf{0}_N$.

To this end, we carried out a parallel implementation of the Lanczos algorithm using PETSC (Balay *et al.*, 1997) and included it as an additional eigenvalue solver

Fig. 1. Barrier time versus number of processes on the *Nehalem*-based cluster.

inside SLEPC (Hernandez *et al.*, 2005). Both s -step BiCG and QMR are implemented in parallel using PETSC 3.0 and compared to their classical variants implemented in PETSC. We use a block Jacobi preconditioner and choose $N = 2^{24}$ since PETSC recommends at least 10 000 to 20 000 unknowns per process. All computations are performed on a *Nehalem*-based cluster¹ whose barrier times measured via PETSC are given in Figure 1. A synchronization with a reduction operation accounts for an average of 5.97×10^{-4} s with 256 processes and 7.26×10^{-4} s with 512 processes.

7.1. Synchronization-reducing s -step Lanczos algorithm. Figure 2 compares the Lanczos implementations. All algorithms are started with

$$v_1^1 = w_1^1 = [1, \dots, 1]^T. \quad (57)$$

The diagram shows the convergence history of the error of the biorthogonality property of A . The deviation of the unnormalized 2-step variant exceeds 10^0 from iteration 6 onwards and the algorithm breaks down at iteration 45 with a numerical overflow. The normalized 2-step variant competes well until iteration 35 with a deviation below 10^0 . However, when increasing s to $s = 5$, the divergence of the normalized variant already starts around iteration 33 with a deviation above 10^0 . The tendency for growing numerical instabilities when s is increased is well known (Carson and Demmel, 2014; Kim, 2010).

Previous research (Gustafsson *et al.*, 2012b; Kim and Chronopoulos, 1992; Kim and Kim, 2005) has already demonstrated that, compared with the classical algorithm, execution time is reduced by the unnormalized s -step algorithm. As a result, we excluded the unnormalized s -step method from the analysis and only remark that

¹Each node of this cluster at RWTH Aachen University, Germany, consists of 2 sockets, each equipped with Intel Xeon X5570 quadcore processors running at 2.93 GHz. Each core has a separate L1 and L2 cache, while 4 cores share an L3 cache of size 8 MB. So each node of this cluster is made up of 8 cores called processes hereafter. The nodes are connected by a quad data rate InfiniBand network.

the parallel performance characteristics of the normalized and unnormalized variants are almost identical. The focus is instead upon the parallel performance of the normalized s -step Lanczos algorithm. We compared the new normalized s -step variant to the classical one implemented in SLEPC. When spurious eigenvalues are present, the Lanczos process can converge to wrong values. This behavior appears at the same time as when the Lanczos vectors start to lose biorthogonality. As a possible remedy, one can apply so-called *reorthogonalization*, which is supported by our SLEPC implementation; see the work of Feuerriegel and Bücker (2013a) for more details.

Figure 3 compares the relative (left) and absolute (right) time savings per iteration when using the normalized s -step variant instead of the classical algorithm. Here, reorthogonalization is not used. The s -step algorithm performs slower than the classical algorithm for $p = 8$ and $p = 16$ processes. However, in this case it performs faster for both $p = 32$ and $p = 64$ processes. With 64 processes, for instance, the corresponding time saving per iteration accounts for 0.0042 s or 16.39%, respectively.

7.2. Numerical accuracy of synchronization-reducing linear solvers. In exact arithmetic, the s -step variants and their classical counterparts are mathematically equivalent. However, their numerical behavior can differ in finite-precision arithmetic. The relative residual norm for BiCG is plotted against the iteration number in Figure 4. The corresponding convergence behavior for QMR is shown in Figure 5. The findings are as follows: (i) the relative residual norms for classical and s -step algorithms are similar, (ii) an increasing iteration index n augments the numerical discrepancy between s -step and classical solvers, (iii) numerical instabilities grow with an increase in step size s —as the s -step Lanczos basis becomes more unstable, it can even result in the possibility of a breakdown.

Fig. 2. Convergence history of the biorthogonality property for the classical algorithm, as well as for the unnormalized and the normalized s -step Lanczos variant.

Fig. 3. Relative (left) and absolute (right) time savings per iteration, when using the normalized s -step variant instead of the classical algorithm. We choose $N = 4096$, $s = 2$ and vary the number of processes p .

7.3. Parallel performance of synchronization-reducing linear solvers. Figure 6 compares the ratio of the total time that is spent on various linear algebra operations. Due to barrier operations arising from GSPs, one would expect a gradual increase in the proportional influence of inner products. For the classical BiCG, the percentage of the total time resulting from inner products grows from around 12% for 64 processes to almost 47% for 512 processes. For the 2-step BiCG, however, there is only a moderate rise from approximately 8% to 13% when varying the processes from 64 to 512. As almost half of the total computation time using 512 processes is spent on the evaluation of inner products, the scalability to a large number of processes is limited for the classical BiCG. The 2-step BiCG variant scales significantly better.

Figure 7 compares the same linear algebra operations in terms of their absolute time consumption. The overall runtime of the classical BiCG is affected by inner products to a large extent. In contrast, the 2-step BiCG computes all inner products in less time. With 512 processes, the classical BiCG needs 0.76 s for computing all inner products, whereas the 2-step BiCG needs only 0.17 s.

In Fig. 8, the average speedup for a single iteration

is depicted. With an increasing number of processes, the speedup of both s -step solvers, BiCG and QMR, ascends more linearly when compared to the classical variants that begin to flatten out. Measurements for step sizes $s > 2$ are unavailable for more than 256 processes due to breakdowns. As a trend, however, we conclude that, given a network with a relatively time-consuming barrier operation, the scalability of the new s -step variants is significantly improved in comparison with the classical algorithms.

8. Concluding remarks

The purpose of s -step methods is to reduce the number of global synchronization points on distributed-memory computers by a factor of $\mathcal{O}(s)$. Rather than carrying out s separate iterations of a traditional method, these methods rely on using a single block iteration that is equivalent in increasing the dimension of the Krylov subspace, i.e., by restructuring the original algorithms in such a way that multiple inner products are grouped for joint execution. We derive a new s -step Lanczos algorithm with normalization of the underlying Krylov

Fig. 4. Comparison of the classical BiCG method and s -step variants across different step sizes without (top) and with block Jacobi (bottom) preconditioning using two processes.

Fig. 5. Comparison of the classical QMR method and s -step variants across different step sizes without (top) and with block Jacobi (bottom) preconditioning using two processes.

basis using efficient recurrences. Numerical experiments indicate that this new normalized s -step variant—like the previous unnormalized s -step one—is more scalable than the traditional Lanczos algorithm. In addition, this new normalized variant exhibits improved numerical accuracy compared to a previous unnormalized variant. Consequently, this s -step Lanczos algorithm shows a possible path to advance parallel scalability on current large-scale and future extreme-scale supercomputers.

However, there is still room for further improvement. Most notably, the numerical stability tends to decrease with an increase in s . Future work is necessary to investigate promising remedies, such as the use of a basis that is different from the monomial basis (Carson *et al.*, 2014). Another viable alternative is to introduce residual replacement strategies (Carson and Demmel, 2014; van der Vorst and Ye, 2000). The techniques

addressed by Gustafsson *et al.* (2012a) might also further improve the numerical stability. In addition, look-ahead techniques (Freund and Nachtigal, 1991; 1994; Freund and Hochbruck, 1991; 1992) are advantageous in preventing breakdowns.

Acknowledgment

Parts of this research were conducted while the authors were in residence at the Institute for Scientific Computing, the Center for Computational Engineering Science, and the Aachen Institute for Advanced Study in Computational Engineering Science at RWTH Aachen University, Germany. Financial support from Deutsche Forschungsgemeinschaft (German Research Foundation) through the grant GSC 111 is gratefully acknowledged.

Fig. 6. Proportional time spent in linear algebra operations measured across 200 iterations including initialization for the classical BiCG (left) and the s -step BiCG method (right).

Fig. 7. Total time spent in linear algebra operations measured across 200 iterations including initialization for the classical BiCG (left) and the s -step BiCG method (right).

References

- Balay, S., Gropp, W.D., McInnes, L.C. and Smith, B.F. (1997). Efficient management of parallelism in object oriented numerical software libraries, in E. Arge *et al.* (Eds.), *Modern Software Tools in Scientific Computing*, Birkhäuser Press, Boston, MA, pp. 163–202.
- Bücker, H.M. (2002). Iteratively solving large sparse linear systems on parallel computers, in J. Grotendorst, D. Marx and A. Muramatsu (Eds.), *Quantum Simulations of Complex Many-Body Systems: From Theory to Algorithms*, NIC Series, Vol. 10, John Von Neumann Institute for Computing, Jülich, pp. 521–548.
- Bücker, H.M. and Sauren, M. (1996). A parallel version of the quasi-minimal residual method based on coupled two-term recurrences, in J. Waśniewski *et al.* (Eds.), *Applied Parallel Computing*, Lecture Notes in Computer Science, Vol. 1184, Springer, Berlin, pp. 157–165.
- Bücker, H.M. and Sauren, M. (1997). A variant of the biconjugate gradient method suitable for massively parallel computing, in G. Bilardi *et al.* (Eds.), *Solving Irregularly Structured Problems in Parallel*, Lecture Notes in Computer Science, Vol. 1253, Springer, Berlin, pp. 72–79.
- Bücker, H.M. and Sauren, M. (1999). Reducing global synchronization in the biconjugate gradient method, in T. Yang (Ed.), *Parallel Numerical Computations with Applications*, Kluwer Academic Publishers, Norwell, MA, pp. 63–76.
- Cappello, F., Geist, A., Gropp, B., Kale, L., Kramer, B. and Snir, M. (2009). Toward exascale resilience, *International Journal of High Performance Computing Applications* **23**(4): 374–388.
- Carson, E. and Demmel, J. (2014). A residual replacement strategy for improving the maximum attainable accuracy of s -step Krylov subspace methods, *SIAM Journal on Matrix Analysis and Applications* **35**(1): 22–43.
- Carson, E., Knight, N. and Demmel, J. (2014). Avoiding communication in nonsymmetric Lanczos-based Krylov subspace methods, *SIAM Journal on Scientific Computing* **35**(5): S42–S61.
- Chronopoulos, A.T. (1986). A class of parallel iterative methods implemented on multiprocessors, *Technical report UIUCDCS-R-86-1267*, Department of Computer Science, University of Illinois, Urbana, IL.
- Chronopoulos, A.T. and Gear, C.W. (1989). s -Step iterative methods for symmetric linear systems, *Journal of Computational and Applied Mathematics* **25**(2): 153–168.
- Chronopoulos, A.T. and Swanson, C.D. (1996). Parallel iterative s -step methods for unsymmetric linear systems, *Parallel Computing* **22**(5): 623–641.
- Curfman McInnes, L., Smith, B., Zhang, H. and Mills, R.T. (2014). Hierarchical Krylov and nested Krylov methods for extreme-scale computing, *Parallel Computing* **40**(1): 17–31.
- Davis, N.E., Robey, R.W., Ferenbaugh, C.R., Nicholaeff, D. and Trujillo, D.P. (2012). Paradigmatic shifts for exascale supercomputing, *Journal of Supercomputing* **62**(2): 1023–1044.

Fig. 8. Average speedup $S = T_{\text{par}}(1)/T_{\text{par}}(p)$ for a single iteration of BiCG (top) and QMR (bottom) averaged across 200 iterations excluding initialization time.

- Demmel, J., Heath, M. and van der Vorst, H. (1993). Parallel numerical linear algebra, *Acta Numerica* **2**(1): 111–197.
- Duff, I.S. (2012). European exascale software initiative: Numerical libraries, solvers and algorithms, in D. Hutchison *et al.* (Eds.), *Euro-Par 2011: Parallel Processing Workshops*, Lecture Notes in Computer Science, Vol. 7155, Springer, Berlin, pp. 295–304.
- Duff, I.S. and van der Vorst, H.A. (1999). Developments and trends in the parallel solution of linear systems, *Parallel Computing* **25**(13–14): 1931–1970.
- Feuerriegel, S. and Bücker, H.M. (2013a). A normalization scheme for the non-symmetric s -step Lanczos algorithm, in J. Kolodziej *et al.* (Eds.), *ICA3PP, Part II*, Lecture Notes in Computer Science, Vol. 8286, Springer, Berlin, pp. 30–39.
- Feuerriegel, S. and Bücker, H.M. (2013b). Synchronization-reducing variants of the biconjugate gradient and the quasi-minimal residual methods, in J. Kolodziej *et al.* (Eds.), *ICA3PP, Part I*, Lecture Notes in Computer Science, Vol. 8285, Springer, Berlin, pp. 226–235.
- Fischer, B. and Freund, R. (1994). An inner product-free conjugate gradient-like algorithm for Hermitian positive definite systems, in J. Brown *et al.* (Eds.), *Cornelius Lanczos International Centenary Conference*, SIAM, Philadelphia, PA, pp. 288–290.
- Fletcher, R. (1976). Conjugate gradient methods for indefinite systems, in G. Watson (Ed.), *Numerical Analysis*, Lecture Notes in Computer Science, Vol. 506, Springer, Berlin, pp. 73–89.
- Freund, R. and Nachtigal, N. (1991). QMR: A quasi-minimal residual method for non-Hermitian linear systems, *Numerische Mathematik* **60**(1): 315–339.
- Freund, R.W. and Hochbruck, M. (1991). A biconjugate gradient type algorithm on massively parallel architectures, in R. Vichnevetsky and J.J.H. Miller (Eds.), *IMACS'91*, Criterion Press, Dublin, pp. 720–721.
- Freund, R.W. and Hochbruck, M. (1992). A biconjugate gradient-type algorithm for the iterative solution of non-Hermitian linear systems on massively parallel architectures, in C. Brezinski and U. Kulisch (Eds.), *IMACS'91*, North Holland, Amsterdam, pp. 169–178.
- Freund, R.W. and Nachtigal, N.M. (1994). An implementation of the QMR method based on coupled two-term recurrences, *SIAM Journal on Scientific Computing* **15**(2): 313.
- Ghysels, P., Ashby, T.J., Meerbergen, K. and Vanroose, W. (2013). Hiding global communication latency in the GMRES algorithm on massively parallel machines, *SIAM Journal on Scientific Computing* **35**(1): C48–C71.
- Ghysels, P. and Vanroose, W. (2014). Hiding global synchronization latency in the preconditioned conjugate gradient algorithm, *Parallel Computing* **40**(7): 224–238.
- Gustafsson, M., Demmel, J. and Holmgren, S. (2012a). Numerical evaluation of the communication-avoiding Lanczos algorithm, *Technical Report 2012-001*, Department of Information Technology, Uppsala University, Uppsala.
- Gustafsson, M., Kormann, K. and Holmgren, S. (2012b). Communication-efficient algorithms for numerical quantum dynamics, in K. Jónasson (Ed.), *Applied Parallel and Scientific Computing*, Lecture Notes in Computer Science, Vol. 7134, Springer, Berlin, pp. 368–378.
- Gutknecht, M.H. (1997). Lanczos-type solvers for nonsymmetric linear systems of equations, *Acta Numerica* **6**(1): 271–397.
- Hernandez, V., Roman, J.E. and Vidal, V. (2005). SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems, *ACM Transactions on Mathematical Software* **31**(3): 351–362.
- Hoemmen, M.F. (2010). *Communication-avoiding Krylov Subspace Methods*, Ph.D. thesis, EECS Department, University of California, Berkeley, CA.

- Kandalla, K., Yang, U., Keasler, J., Kolev, T., Moody, A., Subramoni, H., Tomko, K., Vienne, J., De Supinski, B. and Panda, D. (2012). Designing non-blocking allreduce with collective offload on InfiniBand clusters: A case study with conjugate gradient solvers, *Proceedings of the 2012 IEEE 26th International Parallel Distributed Processing Symposium (IPDPS), Shanghai, China*, pp. 1156–1167.
- Kim, S.K. (2010). Efficient biorthogonal Lanczos algorithm on message passing parallel computer, in C.H. Hsu and V. Malyskin (Eds.), *MTPP 2010, Lecture Notes in Computer Science*, Vol. 6083, Springer, Berlin, pp. 293–299.
- Kim, S.K. and Chronopoulos, A. (1991). A class of Lanczos-like algorithms implemented on parallel computers, *Parallel Computing* **17**(6–7): 763–778.
- Kim, S.K. and Chronopoulos, A.T. (1992). An efficient nonsymmetric Lanczos method on parallel vector computers, *Journal of Computational and Applied Mathematics* **42**(3): 357–374.
- Kim, S.K. and Kim, T.H. (2005). A study on the efficient parallel block Lanczos method, in J. Zhang, J.-H. He and Y. Fu (Eds.), *Computational and Information Science, Lecture Notes in Computer Science*, Vol. 3314, Springer, Berlin/Heidelberg, pp. 231–237.
- Lanczos, C. (1950). An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, *Journal of Research of the National Bureau of Standards* **45**(4): 255–282.
- Meurant, G. (1986). The conjugate gradient method on supercomputers, *Supercomputer* **13**: 9–17.
- Mohiyuddin, M., Hoemmen, M., Demmel, J. and Yelick, K. (2009). Minimizing communication in sparse matrix solvers, *Conference on High Performance Computing Networking, Storage and Analysis, Portland, OR, USA*, pp. 36:1–36:12.
- Saad, Y. (1989). Krylov subspace methods on supercomputers, *SIAM Journal on Scientific and Statistical Computing* **10**(6): 1200–1232.
- Sauren, M. and Bucker, H.M. (1998). On deriving the quasi-minimal residual method, *SIAM Review* **40**(4): 922–926.
- Shalf, J., Dosanjh, S. and Morrison, J. (2011). Exascale computing technology challenges, in D. Hutchison *et al.* (Eds.), *High Performance Computing for Computational Science, VECPAR 2010, Lecture Notes in Computer Science*, Vol. 6449, Springer, Berlin, pp. 1–25.
- van der Vorst, H. (1990). Iterative methods for the solution of large systems of equations on supercomputers, *Advances in Water Resources* **13**(3): 137–146.
- van der Vorst, H.A. and Ye, Q. (2000). Residual replacement strategies for Krylov subspace iterative methods for the convergence of true residuals, *SIAM Journal on Scientific Computing* **22**(3): 835–852.
- Van Rosendale, J. (1983). Minimizing inner product data dependencies in conjugate gradient iteration, *NASA Contractor Report NASA-CR-172178*, NASA Langley Research Center, Hampton, VA.

Zhu, S.-X., Gu, T.-X. and Liu, X.-P. (2014). Minimizing synchronizations in sparse iterative solvers for distributed supercomputers, *Computers & Mathematics with Applications* **67**(1): 199–209.

Zuo, X., Gu, T.-X. and Mo, Z. (2010). An improved GPBi-CG algorithm suitable for distributed parallel computing, *Applied Mathematics and Computation* **215**(12): 4101–4109.

Stefan Feuerriegel graduated from the Aachen Institute for Advanced Study in Computational Engineering Science (AICES) at RWTH Aachen University in 2011. Since then, he has been a Ph.D. student at the Chair of Information Systems Research of the University of Freiburg with a focus on big data analytics and text mining.

H. Martin Bucker is currently a full professor in the Department of Mathematics and Computer Science, Friedrich Schiller University Jena, Germany. Prior to this, he was a researcher at Forschungszentrum Jülich, a visiting scientist at Argonne National Laboratory, and a senior researcher at RWTH Aachen University. His research focuses on parallel algorithms, high-performance computing, combinatorial scientific computing, and automatic differentiation.

Appendix

Deriving efficient recurrence equations

Recurrence for the normalization factor. Recall from Section 4.2 that the normalization scheme guarantees $\langle \mathbf{w}_k^j, \mathbf{v}_k^j \rangle = \pm 1$ by simply scaling the vectors $\tilde{\mathbf{v}}_k^j$ and $\tilde{\mathbf{w}}_k^j$ by scalars σ_k^j . We formally define $\sigma_k^1 := 1$. For all other $j = 2, \dots, s$, this normalization factor can be rewritten as

$$\sigma_k^j = |\langle \tilde{\mathbf{w}}_k^j, \tilde{\mathbf{v}}_k^j \rangle|^{-\frac{1}{2}} \stackrel{(29)}{=} \left| \frac{\sigma_k^j \sigma_k^j}{\overline{M}_k^{j,j}} \right|^{\frac{1}{2}}$$

$$\stackrel{(A5)}{=} \left| \langle \mathbf{w}_k^1, A^{2j-2} \mathbf{v}_k^1 \rangle - \left(\mathbf{t}_{k-1}^j \right)^T \overline{M}_{k-1} \mathbf{t}_{k-1}^j \right|^{-\frac{1}{2}}.$$

Defining auxiliary variables $\Omega_k^{i,j}$, $\Xi_k^{i,j}$ and $\Psi_k^{i,j}$. To simplify the subsequent derivation of the recurrence relations, we introduce three auxiliary variables. The first variable is given by

$$\begin{aligned} \Omega_k^{i,j} &= \langle \mathbf{w}_k^1, A^{i-1} \overline{V}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &= \langle (A^T)^{i-1} \mathbf{w}_k^1, \overline{V}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &\stackrel{(32)}{=} \langle \tilde{\mathbf{w}}_k^i + \overline{W}_{k-1} \mathbf{t}_{k-1}^i, \overline{V}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &\stackrel{(10)}{=} \langle \overline{W}_{k-1} \mathbf{t}_{k-1}^i, \overline{V}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &\stackrel{(29)}{=} (\mathbf{t}_{k-1}^i)^T \overline{M}_{k-1} \mathbf{t}_{k-1}^j. \end{aligned}$$

In addition, we let the second variable be

$$\begin{aligned}\Xi_k^{i,j} &= \langle \overline{W}_{k-1} \mathbf{t}_{k-1}^i, A^{j-1} \mathbf{v}_k^1 \rangle \\ &\stackrel{(33)}{=} \langle \overline{W}_{k-1} \mathbf{t}_{k-1}^i, \tilde{\mathbf{v}}_k^j + \overline{V}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &\stackrel{(10)}{=} (\mathbf{t}_{k-1}^i)^T \overline{M}_{k-1} \mathbf{t}_{k-1}^j.\end{aligned}$$

Finally, the third auxiliary variable is given by

$$\begin{aligned}\Psi_k^{i,j} &= (\mathbf{g}_k^s)^T \langle \overline{W}_k, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &= \sum_{\ell=1}^s \mathbf{g}_k^{\ell,s} \langle \mathbf{w}_k^\ell, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(33)}{=} \sum_{\ell=1}^s \sigma_k^\ell \mathbf{g}_k^{\ell,s} \\ &\quad \langle (A^T)^{\ell-1} \mathbf{w}_k^1 - \overline{W}_{k-1} \mathbf{t}_{k-1}^\ell, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(32)}{=} \sum_{\ell=1}^s \sigma_k^\ell \mathbf{g}_k^{\ell,s} \left[\langle \mathbf{w}_k^1, A^{i+j-s-3+\ell} \mathbf{v}_{k+1}^1 \rangle \right. \\ &\quad \left. - \underbrace{\langle \overline{W}_{k-1} \mathbf{t}_{k-1}^\ell, \tilde{\mathbf{v}}_{k+1}^{i+j-s-1} + \overline{V}_k \mathbf{t}_k^{i+j-s-1} \rangle}_{=0} \right] \\ &\stackrel{(A3)}{=} \sum_{\ell=1}^s \frac{\sigma_k^\ell \mathbf{g}_k^{\ell,s} \mathbf{b}_k^{i-s+\ell-1,j}}{\sigma_k^{i-s+\ell-1}} \\ &= \sum_{\substack{\ell=2s \\ +3-i-j}}^s \frac{\sigma_k^\ell \mathbf{g}_k^{\ell,s} \mathbf{b}_k^{i-s+\ell-1,j}}{\sigma_k^{i-s+\ell-1}}.\end{aligned}\tag{A1}$$

Here, we entered in the last step the non-zero pattern of $\mathbf{b}_k^{i,j}$ from Eqn. (A2) in order to adjust the summation. In order to access only non-zero entries, the sum of both indices in $\mathbf{b}_k^{i-s+\ell-1,j}$ must be greater than or equal to $s+2$, i.e.,

$$(i-s+\ell-1) + j \geq s+2 \quad \Leftrightarrow \quad \ell \geq 2s+3-i-j.$$

Recurrence for $\mathbf{b}_k^{i,j}$. The following recurrence holds:

$$\begin{aligned}\mathbf{b}_k^{i,j} &= \langle \mathbf{w}_k^i, A^{j-1} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(33)}{=} \sigma_k^i \langle (A^T)^{i-1} \mathbf{w}_k^1, A^{j-1} \mathbf{v}_{k+1}^1 \rangle \\ &\quad - \sigma_k^i \langle \overline{W}_{k-1} \mathbf{t}_{k-1}^i, A^{j-1} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(32)}{=} \sigma_k^i \langle \mathbf{w}_k^1, A^{i+j-2} \mathbf{v}_{k+1}^1 \rangle \\ &\quad - \sigma_k^i \underbrace{\langle \overline{W}_{k-1} \mathbf{t}_{k-1}^i, \tilde{\mathbf{v}}_{k+1}^j + \overline{V}_k \mathbf{t}_k^j \rangle}_{=0} \\ &= \sigma_k^i \langle \mathbf{w}_k^1, A^{i+j-2} \mathbf{v}_{k+1}^1 \rangle\end{aligned}$$

$$\begin{aligned}&= \sigma_k^i \langle (A^T)^{s-1} \mathbf{w}_k^1, A^{i+j-s-1} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(33)}{=} \frac{\sigma_k^i}{\sigma_k^s} \langle \mathbf{w}_k^s + \overline{W}_{k-1} \mathbf{t}_{k-1}^s, A^{i+j-s-1} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(32)}{=} \frac{\sigma_k^i}{\sigma_k^s} \langle \mathbf{w}_k^s, A^{i+j-s-1} \mathbf{v}_{k+1}^1 \rangle \\ &\quad + \frac{\sigma_k^i}{\sigma_k^s} \underbrace{\langle \overline{W}_{k-1} \mathbf{t}_{k-1}^s, \tilde{\mathbf{v}}_{k+1}^{i+j-s} + \overline{V}_k \mathbf{t}_k^{i+j-s} \rangle}_{=0} \\ &= \frac{\sigma_k^i}{\sigma_k^s} \langle A^T \mathbf{w}_k^s, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\stackrel{(30)}{=} \frac{\sigma_k^i}{\sigma_k^s} \langle \tilde{\mathbf{w}}_{k+1}^1 + \overline{W}_{k-1} \mathbf{e}_k^s \\ &\quad + \overline{W}_k \mathbf{g}_k^s, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &= \frac{\sigma_k^i}{\sigma_k^s} \langle \tilde{\mathbf{w}}_{k+1}^1, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\quad + \frac{\sigma_k^i}{\sigma_k^s} \langle \overline{W}_{k-1} \mathbf{e}_k^s, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\quad + \frac{\sigma_k^i}{\sigma_k^s} (\mathbf{g}_k^s)^T \langle \overline{W}_k, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\quad \text{cf. Step 13 in Algorithm 2} \\ &\stackrel{(32)}{=} \frac{\sigma_k^i}{\sigma_k^s} \langle \tilde{\mathbf{w}}_{k+1}^1, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \\ &\quad + \frac{\sigma_k^i}{\sigma_k^s} \underbrace{\langle \overline{W}_{k-1} \mathbf{e}_k^s, \tilde{\mathbf{v}}_{k+1}^{i+j-s-1} + \overline{V}_k \mathbf{t}_k^{i+j-s-1} \rangle}_{=0} \\ &\quad + \frac{\sigma_k^i}{\sigma_k^s} \Psi_k^{i,j} \\ &\stackrel{(A1)}{=} \frac{\sigma_k^i}{\sigma_k^s} \left[\langle \mathbf{f}_{k+1} \mathbf{w}_{k+1}^1, A^{i+j-s-2} \mathbf{v}_{k+1}^1 \rangle \right. \\ &\quad \left. + \sum_{\substack{\ell=2s \\ +3-i-j}}^s \frac{\sigma_k^\ell \mathbf{g}_k^{\ell,s} \mathbf{b}_k^{i-s+\ell-1,j}}{\sigma_k^{i-s+\ell-1}} \right]\end{aligned}$$

and then the matrix $\mathbf{b}_k$ features the following structure:

$$\mathbf{b}_k = \begin{bmatrix} 0 & & & 0 \\ & \ddots & 0 & \beta_1 \\ & 0 & \ddots & \vdots \\ 0 & \beta_1 & \cdots & \beta_{s-1} \end{bmatrix}.\tag{A2}$$

First, the anti-diagonal shape origins from the above equation:

$$\mathbf{b}_k^{i,j} = \sigma_k^i \langle \mathbf{w}_k^1, A^{i+j-2} \mathbf{v}_{k+1}^1 \rangle,\tag{A3}$$

with $i+j = \text{const}$, as the sum $i+j$ determines the values on the anti-diagonals in Eqn. (A2). Second, the non-zero

pattern comes from

$$\begin{aligned}
 b_k^{i,j} &= \langle w_k^i, A^{j-1} v_{k+1}^1 \rangle = \langle (A^T)^{j-1} w_k^i, v_{k+1}^1 \rangle \\
 &= \sigma_k^i \langle (A^T)^{j-1} \tilde{w}_k^i, v_{k+1}^1 \rangle \\
 &\stackrel{(33)}{=} \sigma_k^i \langle (A^T)^{j-1} \left[(A^T)^{i-1} w_k^1 - \overline{W}_{k-1} t_{k-1}^i \right], v_{k+1}^1 \rangle \\
 &= \sigma_k^i \langle (A^T)^{j+i-2} w_k^1 - (A^T)^{j-1} \overline{W}_{k-1} t_{k-1}^i, v_{k+1}^1 \rangle \\
 &\stackrel{(33)}{=} \sigma_k^i \langle \tilde{w}_k^{i+j-1} + \overline{W}_{k-1} t_{k-1}^{i+j-1} \\
 &\quad - (A^T)^{j-1} \overline{W}_{k-1} t_{k-1}^i, v_{k+1}^1 \rangle \\
 &= \sigma_k^i \left[\langle \tilde{w}_k^{i+j-1}, v_{k+1}^1 \rangle + \underbrace{\langle \overline{W}_{k-1} t_{k-1}^{i+j-1}, v_{k+1}^1 \rangle}_{=0} \right. \\
 &\quad \left. - \langle (A^T)^{j-1} \overline{W}_{k-1} t_{k-1}^i, v_{k+1}^1 \rangle \right] \\
 &= \sigma_k^i \left[\langle \tilde{w}_k^{i+j-1}, v_{k+1}^1 \rangle - \langle \overline{W}_{k-1} t_{k-1}^i, A^{j-1} v_{k+1}^1 \rangle \right] \\
 &\stackrel{(32)}{=} \sigma_k^i \left[\langle \tilde{w}_k^{i+j-1}, v_{k+1}^1 \rangle \right. \\
 &\quad \left. - \underbrace{\langle \overline{W}_{k-1} t_{k-1}^i, \tilde{v}_{k+1}^j + \overline{V}_k t_k^j \rangle}_{=0} \right] \\
 &= \sigma_k^i \cdot 0 = 0 \quad \text{for } i+j-1 \leq s.
 \end{aligned}$$

Recurrence for $c_k^{i,j}$. The matrix c_k , whose entries are the elements $c_k^{i,j}$, features the following non-zero pattern:

$$c_k = \begin{bmatrix} \vdots & & \vdots \\ 0 & \dots & 0 \\ c_k^{s,1} & \dots & c_k^{s,s} \end{bmatrix}. \quad (\text{A4})$$

We prove the above structure by induction. If $c_0 := 0_{s,s}$, then it is sufficient to focus on the induction step from $k-1$ to k . Assuming the structure holds for $c_k^{i,j}$, we derive

$$\begin{aligned}
 c_k^{i,j} &= \langle w_{k-1}^i, A v_k^j \rangle = \langle A^T w_{k-1}^i, v_k^j \rangle \\
 &= \sigma_{k-1}^i \langle A^T \tilde{w}_{k-1}^i, v_k^j \rangle \\
 &\stackrel{(33)}{=} \sigma_{k-1}^i \langle A^T \left[(A^T)^{i-1} w_{k-1}^1 - \overline{W}_{k-2} t_{k-2}^i \right], v_k^j \rangle \\
 &= \sigma_{k-1}^i \langle (A^T)^i w_{k-1}^1 - A^T \overline{W}_{k-2} t_{k-2}^i, v_k^j \rangle \\
 &\stackrel{(33)}{=} \sigma_{k-1}^i \left[\langle \tilde{w}_{k-1}^{i+1} + \overline{W}_{k-2} t_{k-2}^{i+1}, v_k^j \rangle \right. \\
 &\quad \left. - \langle A^T \overline{W}_{k-2} t_{k-2}^i, v_k^j \rangle \right]
 \end{aligned}$$

$$\begin{aligned}
 &= \sigma_{k-1}^i \left[\langle \tilde{w}_{k-1}^{i+1}, v_k^j \rangle + \underbrace{\langle \overline{W}_{k-2} t_{k-2}^{i+1}, v_k^j \rangle}_{=0} \right. \\
 &\quad \left. - \langle \overline{W}_{k-2} t_{k-2}^i, A v_k^j \rangle \right] \\
 &\stackrel{(32)}{=} \sigma_{k-1}^i \left[\langle \tilde{w}_{k-1}^{i+1}, v_k^j \rangle \right. \\
 &\quad \left. - \sigma_k^j \langle \overline{W}_{k-2} t_{k-2}^i, A [A^{j-1} v_k^1 - \overline{V}_{k-1} t_{k-1}^j] \rangle \right] \\
 &\stackrel{(32)}{=} \sigma_{k-1}^i \left[\langle \tilde{w}_{k-1}^{i+1}, v_k^j \rangle - \sigma_k^j \langle \overline{W}_{k-2} t_{k-2}^i, \right. \\
 &\quad \left. v_k^{j+1} + \overline{V}_{k-1} t_{k-1}^{j+1} - A \overline{V}_{k-1} t_{k-1}^j \rangle \right] \\
 &= \sigma_{k-1}^i \langle \tilde{w}_{k-1}^{i+1}, v_k^j \rangle \\
 &\quad - \sigma_{k-1}^i \sigma_k^j \underbrace{\langle \overline{W}_{k-2} t_{k-2}^i, v_k^{j+1} + \overline{V}_{k-1} t_{k-1}^{j+1} \rangle}_{=0} \\
 &\quad + \sigma_{k-1}^i \sigma_k^j \langle \overline{W}_{k-2} t_{k-2}^i, A \overline{V}_{k-1} t_{k-1}^j \rangle \\
 &\stackrel{(34)}{=} \sigma_{k-1}^i \left[\langle \tilde{w}_{k-1}^{i+1}, v_k^j \rangle + \sigma_k^j (t_{k-2}^i)^T c_{k-1} t_{k-1}^j \right] \\
 &= \sigma_{k-1}^i \cdot 0 = 0
 \end{aligned}$$

for $i < s$ by induction hypotheses $c_{k-1}^{i,j} = 0$ for all $i < s$. Furthermore,

$$\begin{aligned}
 c_k^{s,j} &= \langle w_{k-1}^s, A v_k^j \rangle \\
 &\stackrel{(32)}{=} \sigma_k^j \langle w_{k-1}^s, A [A^{j-1} v_k^1 - \overline{V}_{k-1} t_{k-1}^j] \rangle \\
 &\stackrel{(35),(36)}{=} \sigma_k^j \left[b_{k-1}^{s,j+1} - d_{k-1}^{s,\bullet} t_{k-1}^j \right]
 \end{aligned}$$

holds where $d_{k-1}^{s,\bullet}$ denotes the s -th row of d_{k-1} . When computing $c_k^{s,j}$ recursively, the additional value of $b_{k-1}^{s,s+1}$ is unavoidably required, but which we also determine in the above recursive fashion.

Recurrence for $d_k^{i,j}$. Next, we find

$$\begin{aligned}
 d_k^{i,j} &= \langle w_k^i, A v_k^j \rangle \\
 &= \sigma_k^i \sigma_k^j \left[\langle (A^T)^{i-1} w_k^1 \right. \\
 &\quad \left. - \overline{W}_{k-1} t_{k-1}^i, A [A^{j-1} v_k^1 - \overline{V}_{k-1} t_{k-1}^j] \rangle \right] \\
 &= \sigma_k^i \sigma_k^j \left[\langle (A^T)^{i-1} w_k^1 \right. \\
 &\quad \left. - \overline{W}_{k-1} t_{k-1}^i, A^j v_k^1 - A \overline{V}_{k-1} t_{k-1}^j \rangle \right] \\
 &\stackrel{(35)}{=} \sigma_k^i \sigma_k^j \left[\langle w_k^1, A^{i+j-1} v_k^1 \rangle \right. \\
 &\quad \left. - \Omega_k^{i+1,j} - \Xi_k^{i,j+1} + (t_{k-1}^i)^T d_{k-1} t_{k-1}^j \right].
 \end{aligned}$$

We notice that the given expressions

$$\begin{aligned}\Omega_k^{i,j} &= (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j, \\ \Xi_k^{i,j} &= (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j\end{aligned}$$

cannot be evaluated for both $\Omega_k^{s+1,j}$ and $\Xi_k^{i,s+1}$ as the variable $\mathbf{t}_{k-1}^{s+1}$ is undefined. Therefore, we need to derive alternative recursions for the special cases $\Omega_k^{s+1,j}$ and $\Xi_k^{i,s+1}$ as follows:

$$\begin{aligned}\Omega_k^{s+1,j} &:= \langle \mathbf{w}_k^1, A^s \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &= \langle (A^T)^{s-1} \mathbf{w}_k^1, A \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &\stackrel{(33)}{=} \langle \frac{\mathbf{w}_k^s}{\sigma_k^s} + \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^s, A \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &= \frac{1}{\sigma_k^s} (\overline{\mathbf{V}}_{k-1}^T A^T \mathbf{w}_k^s)^T \mathbf{t}_{k-1}^j \\ &\quad + \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^s, A \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &= \frac{1}{\sigma_k^s} (\overline{\mathbf{W}}_{k-1}^T A \mathbf{v}_k^s)^T \mathbf{t}_{k-1}^j \\ &\quad + \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^s, A \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle\end{aligned}$$

after Kim and Chronopoulos (1992, p. 366). Then, using (34) and (35), we arrive at

$$\Omega_k^{s+1,j} = \frac{(\mathbf{c}_k^{\bullet,s})^T \mathbf{t}_{k-1}^j}{\sigma_k^s} + (\mathbf{t}_{k-1}^s)^T \mathbf{d}_{k-1} \mathbf{t}_{k-1}^j,$$

where $\mathbf{c}_k^{\bullet,s}$ is the s -th column of $\mathbf{c}_k$. Using (A4), we finally find

$$\Omega_k^{s+1,j} = \frac{\mathbf{c}_k^{s,s} \mathbf{t}_{k-1}^j}{\sigma_k^s} + (\mathbf{t}_{k-1}^s)^T \mathbf{d}_{k-1} \mathbf{t}_{k-1}^j.$$

Similarly, we derive

$$\begin{aligned}\Xi_k^{i,s+1} &:= \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, A^s \mathbf{v}_k^1 \rangle \\ &\stackrel{(32)}{=} \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, A \left[\frac{\mathbf{v}_k^s}{\sigma_k^s} + \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^s \right] \rangle \\ &= \frac{\langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, A \mathbf{v}_k^s \rangle}{\sigma_k^s} \\ &\quad + \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, A \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^s \rangle \\ &\stackrel{(34),(35)}{=} \frac{(\mathbf{t}_{k-1}^i)^T \mathbf{c}_k^{s,s}}{\sigma_k^s} + (\mathbf{t}_{k-1}^i)^T \mathbf{d}_{k-1} \mathbf{t}_{k-1}^s \\ &\stackrel{(A4)}{=} \frac{\mathbf{t}_{k-1}^{s,i} \mathbf{c}_k^{s,s}}{\sigma_k^s} + (\mathbf{t}_{k-1}^i)^T \mathbf{d}_{k-1} \mathbf{t}_{k-1}^s.\end{aligned}$$

Recurrence for $\overline{\mathbf{M}}_k^{i,j}$. Finally, we have

$$\begin{aligned}\overline{\mathbf{M}}_k^{i,j} &= \langle \mathbf{w}_k^i, \mathbf{v}_k^j \rangle \\ &= \sigma_k^i \sigma_k^j \langle (A^T)^{i-1} \mathbf{w}_k^1 \\ &\quad - \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, A^{j-1} \mathbf{v}_k^1 - \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &= \sigma_k^i \sigma_k^j \left[\langle \mathbf{w}_k^1, A^{i+j-2} \mathbf{v}_k^1 \rangle \right. \\ &\quad - \langle \mathbf{w}_k^1, A^{i-1} \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \\ &\quad - \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, A^{j-1} \mathbf{v}_k^1 \rangle \\ &\quad \left. + \langle \overline{\mathbf{W}}_{k-1} \mathbf{t}_{k-1}^i, \overline{\mathbf{V}}_{k-1} \mathbf{t}_{k-1}^j \rangle \right] \\ &\stackrel{(29)}{=} \sigma_k^i \sigma_k^j \left[\langle \mathbf{w}_k^1, A^{i+j-2} \mathbf{v}_k^1 \rangle - \Omega_k^{i,j} \right. \\ &\quad \left. - \Xi_k^{i,j} + (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j \right] \\ &= \sigma_k^i \sigma_k^j \left[\langle \mathbf{w}_k^1, A^{i+j-2} \mathbf{v}_k^1 \rangle \right. \\ &\quad \left. - (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j \right],\end{aligned}\tag{A5}$$

where the last equation follows from

$$\Omega_k^{i,j} = \Xi_k^{i,j} = (\mathbf{t}_{k-1}^i)^T \overline{\mathbf{M}}_{k-1} \mathbf{t}_{k-1}^j.\tag{A6}$$

Received: 15 April 2014
Revised: 17 October 2014