

RECYCLING BiCGSTAB WITH AN APPLICATION TO PARAMETRIC MODEL ORDER REDUCTION*

KAPIL AHUJA[†], PETER BENNER[‡], ERIC DE STURLER[§], AND LIHONG FENG[¶]

Abstract. Krylov subspace recycling is a process for accelerating the convergence of sequences of linear systems. Based on this technique, the recycling BiCG algorithm has been developed recently. Here, we now generalize and extend this recycling theory to BiCGSTAB. Recycling BiCG focuses on efficiently solving sequences of dual linear systems, while the focus here is on efficiently solving sequences of single linear systems (assuming nonsymmetric matrices for recycling BiCG and recycling BiCGSTAB). As compared with other methods for solving sequences of single linear systems with nonsymmetric matrices (e.g., recycling variants of GMRES), BiCG-based recycling algorithms, like recycling BiCGSTAB, have the advantage that they involve a short-term recurrence and hence do not suffer from storage issues and are also cheaper with respect to the orthogonalizations. We modify the BiCGSTAB algorithm to use a recycle space, which is built from left and right approximate invariant subspaces. Using our algorithm for a parametric model order reduction example gives good results. We show about 40% savings in the number of matrix-vector products and about 35% savings in runtime.

Key words. Krylov subspace recycling, deflation, BiCGSTAB, model reduction, rational Krylov

AMS subject classifications. 65F10, 65N22, 93A15, 93C05

DOI. 10.1137/140972433

1. Introduction. We focus on efficiently solving sequences of linear systems of the following type:

$$(1.1) \quad A^{(\iota)} x^{(\iota, \kappa)} = b^{(\iota, \kappa)},$$

where $A^{(\iota)} \in \mathbb{R}^{n \times n}$ varies with ι ; $b^{(\iota, \kappa)} \in \mathbb{R}^n$ varies with both ι and κ ; the matrices $A^{(\iota)}$ are large, sparse, and nonsymmetric; and the change from one system to the next is small.

Krylov subspace methods are usually used for solving such large and sparse linear systems. For linear systems with nonsymmetric matrices, GMRES [28] is one of the first choices, but it is generally not optimal with respect to the runtime. BiCGSTAB [32] is competitive with GMRES and in many cases performs better than GMRES in time. Also, it does not suffer from storage issues, which is a problem in GMRES.

Krylov subspace recycling is a technique for efficient solution of sequences of linear systems. Here, while solving one system in the sequence, approximate invariant subspaces of the matrix are selected and used to accelerate the convergence of the next

*Received by the editors June 11, 2014; accepted for publication (in revised form) June 1, 2015; published electronically October 29, 2015. This material is based upon work supported by the National Science Foundation (USA) under grants DMS-1025327 and DMS-1217156, and by Council of Scientific and Industrial Research (India) grant 25/(0220)/13/EMR-II.
<http://www.siam.org/journals/sisc/37-5/97243.html>

[†]Discipline of Computer Science and Engineering, Indian Institute of Technology Indore, Indore, India (kahuja@iiti.ac.in).

[‡]Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, D-39106 Germany (benner@mpi-magdeburg.mpg.de).

[§]Department of Mathematics, Virginia Tech, Blacksburg, VA 24061 (sturler@vt.edu).

[¶]Max Planck Institute for Dynamics of Complex Technical Systems, Magdeburg, D-39106 Germany (feng@mpi-magdeburg.mpg.de).

system in the sequence. Since the matrices in the sequence do not change much, this provides substantial reduction in both the number of matrix-vector products and time. See [25] and [5] for more about Krylov subspace recycling.

Here, we have a sequence of linear systems with nonsymmetric matrices, and hence GCRO-DR [25] and GCROT [25] algorithms, which are recycling variants of GMRES, are more suited. However, since there is no optimal method in time for solving linear systems with nonsymmetric matrices, there is no optimal method in time for solving sequences of such linear systems. Like GMRES, its recycling variants may also suffer from storage issues. Hence, we develop a recycling variant of BiCGSTAB based on the work in [4, 3].

We explore the usage of recycling BiCGSTAB for parametric model order reduction (PMOR) [7, 14] that requires a solution of systems of the form (1.1). We show about 40% reduction in the number of matrix-vector products when using recycling as compared with not using recycling in BiCGSTAB. In terms of time, this translates to about 35% savings in runtime.

In related work in this area (specific to BiCGSTAB), [20] discusses a variant of recycling BiCGSTAB (using the terminology of deflation and augmentation instead of Krylov subspace recycling). There are three main differences between that approach and ours. First, we use a different inner product in our derivation as compared with [20]. Second, the focus in [20] is on only using a recycle space, while here we discuss both using and generating a recycle space. Third, we also give numerical experiments demonstrating the usefulness of our approach, while [20] discusses only a theoretical framework. Also, [2, 24, 1] focus on variants of deflated BiCGSTAB for multiple right-hand sides and do not discuss changing matrices.

Since we propose recycling BiCGSTAB as an alternative for GMRES-based recycling solvers, we also compare with GCRO-DR. In the context of PMOR, simplified versions of GCRO-DR have also been proposed [14, 15]. We show that our recycling BiCGSTAB is 10% more efficient in time than GCRO-DR for our test problem.

To simplify notation, we drop the superscripts ι and κ in (1.1). Throughout the paper, $\|\cdot\|$ refers to the two-norm, (\cdot, \cdot) refers to the standard inner product, $*$ indicates the conjugate transpose operation, $\bar{\cdot}$ indicates complex conjugation, and $\underline{\cdot}$ is used to signify a rectangular matrix.

The rest of the paper is divided into six more sections. The bi-Lanczos algorithm [22] and recycling BiCG [5] form the basis of our recycling BiCGSTAB. Hence, we revisit these in sections 2 and 3, respectively. In section 3, we also give a new result related to recycling BiCG. Next, we derive recycling BiCGSTAB in section 4. In section 5, we analyze the subspaces that can be used in recycling BiCGSTAB. Finally, we discuss the application of recycling BiCGSTAB to PMOR in section 6 and give concluding remarks in section 7.

2. The bi-Lanczos algorithm. Consider a primary system $Ax = b$ with x_0 the initial guess and $r_0 = b - Ax_0$ the residual. Also, consider an auxiliary dual system $A^*\tilde{x} = \tilde{b}$ with \tilde{b} a random vector, \tilde{x}_0 the initial guess, and $\tilde{r}_0 = \tilde{b} - A^*\tilde{x}_0$ the residual. This dual system is termed auxiliary because for this work we are not interested in its solution (although the system is real). The bi-Lanczos algorithm remains the same even when the dual system is of interest.

Let the columns of $V_i = [v_1 \ v_2 \ \dots \ v_i]$ define the basis of the primary system Krylov space $\mathcal{K}^i(A, r_0) \equiv \text{span}\{r_0, Ar_0, A^2r_0, \dots, A^{i-1}r_0\}$. Also, let the columns of $\tilde{V}_i = [\tilde{v}_1 \ \tilde{v}_2 \ \dots \ \tilde{v}_i]$ define the basis of the dual system Krylov space $\tilde{\mathcal{K}}^i(A^*, \tilde{r}_0) \equiv \text{span}\{\tilde{r}_0, A^*\tilde{r}_0, A^{2*}\tilde{r}_0, \dots, A^{(i-1)*}\tilde{r}_0\}$.

The bi-Lanczos algorithm computes the columns of V_i and \tilde{V}_i such that, in exact arithmetic, $V_i \perp_b \tilde{V}_i$, where \perp_b is referred to as bi-orthogonality; this implies that $\tilde{V}_i^* V_i$ is a diagonal matrix. The columns of V_i and \tilde{V}_i are called Lanczos vectors. There is a degree of freedom in choosing the scaling of the Lanczos vectors [17, 19, 27]. Using the scaling

$$(2.1) \quad \|v_i\| = 1, \quad (v_i, \tilde{v}_i) = 1,$$

we initialize the Lanczos vectors as follows:

$$v_1 = \frac{r_0}{\|r_0\|}, \quad \tilde{v}_1 = \frac{\tilde{r}_0}{(v_1, \tilde{r}_0)}.$$

The $(i + 1)$ st Lanczos vectors are given by

$$(2.2) \quad \begin{aligned} \gamma v_{i+1} &= Av_i - V_i \tau \perp \tilde{V}_i, \\ \tilde{\gamma} \tilde{v}_{i+1} &= A^* \tilde{v}_i - \tilde{V}_i \tilde{\tau} \perp V_i, \end{aligned}$$

where $\gamma, \tilde{\gamma}$ and $\tau, \tilde{\tau}$ are determined by the biorthogonality condition (2.2) and the normalization condition (2.1). The computation of the $(i + 1)$ st Lanczos vectors requires only the i th and the $(i - 1)$ st Lanczos vectors (see [27]). These 3-term recurrences are called the bi-Lanczos relations and are defined as follows:

$$(2.3) \quad \begin{aligned} AV_i &= V_{i+1} \underline{T}_i = V_i T_i + t_{i+1,i} v_{i+1} e_i^T, \\ A^* \tilde{V}_i &= \tilde{V}_{i+1} \underline{\tilde{T}}_i = \tilde{V}_i \tilde{T}_i + \tilde{t}_{i+1,i} \tilde{v}_{i+1} e_i^T, \end{aligned}$$

where T_i, \tilde{T}_i are $i \times i$ tridiagonal matrices, $t_{i+1,i}$ is the last element of the last row of $\underline{T}_i \in \mathbb{C}^{(i+1) \times i}$, and $\tilde{t}_{i+1,i}$ is the last element of the last row of $\underline{\tilde{T}}_i \in \mathbb{C}^{(i+1) \times i}$.

The bi-Lanczos algorithm breaks down when at any step i , $\tilde{v}_i^* v_i = 0$. There exist so-called look-ahead strategies [16, 19] that take multiple Lanczos vectors together in succession and make them block bi-orthogonal to avoid this breakdown.

3. Recycling BiCG revisited. We first introduce a generalization of the bi-Lanczos algorithm [4]. We show that even for a pair of matrices that are not conjugate transposes of each other, one can build bi-orthogonal bases (for the associated two Krylov subspaces) using a short-term recurrence.

Expanding the search space to include a recycle space leads to an augmented bi-orthogonality condition. The augmented bi-Lanczos algorithm, as derived for recycling BiCG [5], computes bi-orthogonal bases for the two Krylov subspaces such that this augmented bi-orthogonality condition is satisfied. Next, we revisit augmented bi-Lanczos [3] and show that it is a special case of generalized bi-Lanczos. Finally, we list the recycling BiCG algorithm from [5].

There are numerous ways of computing good bases for Krylov subspaces $\mathcal{K}^m(B, v_1)$ and $\mathcal{K}^m(\tilde{B}, \tilde{v}_1)$, where B and \tilde{B} are $n \times n$ general matrices, and v_1 and \tilde{v}_1 are any two n dimensional vectors. Let the columns of $V_m = [v_1 \ v_2 \ \dots \ v_m]$ and $\tilde{V}_m = [\tilde{v}_1 \ \tilde{v}_2 \ \dots \ \tilde{v}_m]$ define one such pair of good bases for $\mathcal{K}^m(B, v_1)$ and $\mathcal{K}^m(\tilde{B}, \tilde{v}_1)$, respectively. We compute these bases using the following, in principle, full recurrences:

$$(3.1) \quad \beta_{i+1,i} v_{i+1} = Bv_i - \beta_{ii} v_i - \beta_{i-1,i} v_{i-1} - \dots - \beta_{1i} v_1,$$

$$(3.2) \quad \tilde{\beta}_{i+1,i} \tilde{v}_{i+1} = \tilde{B}\tilde{v}_i - \tilde{\beta}_{ii} \tilde{v}_i - \tilde{\beta}_{i-1,i} \tilde{v}_{i-1} - \dots - \tilde{\beta}_{1i} \tilde{v}_1,$$

where $i \in \{1, 2, 3, \dots, m-1\}$ and $\{\beta_{ij}\}, \{\tilde{\beta}_{ij}\}$ are scalars to be determined. We assume that for $i < m$, $\mathcal{K}^i(B, v_1)$ is not an invariant subspace of B . (Similarly, $\mathcal{K}^i(\tilde{B}, \tilde{v}_1)$ is

not an invariant subspace of \tilde{B} for $i < m$.) We can rewrite (3.1) as follows:

$$Bv_i = \beta_{1i}v_1 + \beta_{2i}v_2 + \cdots + \beta_{i-1,i}v_{i-1} + \beta_{ii}v_i + \beta_{i+1,i}v_{i+1}.$$

Combining these equations, for $i \in \{1, 2, 3, \dots, m-1\}$, into matrix form, we get

$$B[v_1 \ v_2 \ \dots \ v_{m-1}] = [v_1 \ v_2 \ \dots \ v_{m-2} \ v_{m-1} \ v_m] \begin{bmatrix} \beta_{11} & \beta_{12} & \cdots & \beta_{1,m-1} \\ \beta_{21} & \beta_{22} & \cdots & \beta_{2,m-1} \\ 0 & \beta_{32} & \cdots & \beta_{3,m-1} \\ 0 & 0 & \cdots & \beta_{4,m-1} \\ \vdots & \ddots & \ddots & \vdots \\ 0 & 0 & \ddots & \beta_{m-1,m-1} \\ 0 & 0 & \cdots & \beta_{m,m-1} \end{bmatrix},$$

or

$$BV_{m-1} = V_m \underline{H}_{m-1},$$

where \underline{H}_{m-1} is an $m \times (m-1)$ upper Hessenberg matrix. This result also holds for each $i \in \{1, 2, 3, \dots, m-1\}$, i.e.,

$$(3.3) \quad BV_i = V_{i+1} \underline{H}_i.$$

Similarly, using (3.2) and following the steps above, we get the following relation for the dual system:

$$(3.4) \quad \tilde{B}\tilde{V}_i = \tilde{V}_{i+1} \tilde{\underline{H}}_i.$$

The scalars $\{\beta_{ij}\}$ and $\{\tilde{\beta}_{ij}\}$ are determined by a choice of constraints. One option is to enforce that the columns of V_i (and \tilde{V}_i) are orthonormal vectors (as in the Arnoldi algorithm). Another option, as in the bi-Lanczos algorithm, is to enforce¹

$$V_i \perp_b \tilde{V}_i, \quad \|v_i\| = 1, \quad \text{and} \quad (v_i, \tilde{v}_i) = 1,$$

or

$$(3.5) \quad \tilde{V}_i^* V_i = I \quad \text{and} \quad \|v_i\| = 1.$$

If $\tilde{B} = B^*$, then (3.3), (3.4), and (3.5) lead to the bi-Lanczos relations (2.3), which consist of three-term recurrences. Our goal here is to relax the condition $\tilde{B} = B^*$ and still obtain short-term recurrences.

THEOREM 3.1. *Let $B, \tilde{B} \in \mathbb{C}^{n \times n}$, and let the following conditions hold:*

- (a) $B - \tilde{B}^* = \tilde{F}_k \tilde{C}_k^* - C_k F_k^*$, where $C_k, \tilde{C}_k, F_k, \tilde{F}_k \in \mathbb{C}^{n \times k}$,
- (b) $\forall x : Bx \perp \tilde{C}_k, \forall \tilde{x} : \tilde{B}\tilde{x} \perp C_k$,
- (c) $v_1 \perp \tilde{C}_k$, and $\tilde{v}_1 \perp C_k$.

Also, let (3.5) be used as the set of constraints for (3.3) and (3.4). Then, $\beta_{ij} = 0$ and $\tilde{\beta}_{ij} = 0$ for $j > i + 1$, which leads to the following three-term recurrences:

$$\begin{aligned} \beta_{i+1,i}v_{i+1} &= Bv_i - \beta_{ii}v_i - \beta_{i-1,i}v_{i-1}, \\ \tilde{\beta}_{i+1,i}\tilde{v}_{i+1} &= \tilde{B}\tilde{v}_i - \tilde{\beta}_{ii}\tilde{v}_i - \tilde{\beta}_{i-1,i}\tilde{v}_{i-1} \end{aligned}$$

for $i \in \{1, 2, 3, \dots, m-1\}$.

¹In this paper, for ease of exposition, we assume that breakdowns do not happen. Hence, $(\tilde{v}_i, v_i) \neq 0$.

Proof. Using (b) and (c), we can show that

$$(3.6) \quad C_k^* \tilde{V}_i = 0 \quad \text{and} \quad \tilde{C}_k^* V_i = 0.$$

We show $C_k^* \tilde{V}_i = 0$ by induction. One can similarly show that $\tilde{C}_k^* V_i = 0$. $C_k^* \tilde{v}_1 = 0$ by (c). Let $C_k^* \tilde{v}_l = 0$ for $l = \{1, 2, \dots, i\}$, and consider the case $l = i + 1$. From (3.2), we know that

$$\tilde{\beta}_{i+1,i} \tilde{v}_{i+1} = \tilde{B} \tilde{v}_i - \tilde{\beta}_{ii} \tilde{v}_i - \tilde{\beta}_{i-1,i} \tilde{v}_{i-1} - \dots - \tilde{\beta}_{1i} \tilde{v}_1.$$

Then, $C_k^* \tilde{v}_{i+1} = 0$ since $C_k^* \tilde{B} \tilde{v}_i = 0$ using (b) and $\tilde{\beta}_{li} C_k^* \tilde{v}_l = 0$ for $l \in \{1, 2, \dots, i\}$ by the induction hypothesis.² This proves (3.6). Multiplying both sides in (3.3) by \tilde{V}_i^* and using (3.5), we get

$$\tilde{V}_i^* B V_i = H_i.$$

Substituting (a) in the above equation leads to

$$\begin{aligned} \tilde{V}_i^* \left(\tilde{B}^* + \tilde{F}_k \tilde{C}_k^* - C_k F_k^* \right) V_i &= H_i \iff \\ \tilde{V}_i^* \tilde{B}^* V_i + \tilde{V}_i^* \tilde{F}_k \tilde{C}_k^* V_i - \tilde{V}_i^* C_k F_k^* V_i &= H_i. \end{aligned}$$

Using (3.6), we get

$$\begin{aligned} \tilde{V}_i^* \tilde{B}^* V_i &= H_i \iff \\ (\tilde{B} \tilde{V}_i)^* V_i &= H_i. \end{aligned}$$

Finally, using (3.4) and (3.5) in the above equation gives

$$\tilde{H}_i^* = H_i.$$

This implies that both H_i and \tilde{H}_i are tridiagonal matrices, and hence $\beta_{ij} = 0$ and $\tilde{\beta}_{ij} = 0$ for $j > i + 1$. \square

We now revisit augmented bi-Lanczos [3] and show that it is a special case of generalized bi-Lanczos. The BiCG algorithm is primarily used where the dual system is not auxiliary. That is, one needs to solve both a primary system and a dual system. The recycling BiCG algorithm (also termed RBiCG) was developed to accelerate the convergence of sequences of such systems.

In RBiCG, we use the matrix U to define the primary system recycle space, and we compute $C = A^{(\iota+1)} U$, where U is derived from an approximate right invariant subspace of $A^{(\iota)}$ and ι denotes the index of the linear system in the sequence of linear systems; see (1.1). Similarly, we use the matrix \tilde{U} to define the dual system recycle space, and we compute $\tilde{C} = A^{(\iota+1)*} \tilde{U}$, where \tilde{U} is derived from an approximate left invariant subspace of $A^{(\iota)}$. U and \tilde{U} are computed such that C and \tilde{C} are bi-orthogonal (see page 35 of [4]). The number of vectors selected for recycling is denoted by k , and hence U, \tilde{U}, C , and $\tilde{C} \in \mathbb{C}^{n \times k}$.

The bi-Lanczos algorithm was modified to compute the columns of V_i and \tilde{V}_i such that

$$[C \ V_i] \perp_b [\tilde{C} \ \tilde{V}_i].$$

²Note that our earlier assumption, that $\mathcal{K}^i(\tilde{B}, \tilde{v}_1)$ is not an invariant subspace of \tilde{B} for $i < m$, shows that $\tilde{\beta}_{i+1,i} \neq 0$.

Using the scaling (2.1), we initialize the Lanczos vectors as

$$v_1 = \frac{(I - CD_c^{-1}\tilde{C}^*)r_0}{\|(I - CD_c^{-1}\tilde{C}^*)r_0\|}, \quad \tilde{v}_1 = \frac{(I - \tilde{C}\mathcal{D}_c^{-1}C^*)\tilde{r}_0}{(v_1, (I - \tilde{C}\mathcal{D}_c^{-1}C^*)\tilde{r}_0)}.$$

Here, $\mathcal{D}_c = \tilde{C}^*C$ is a diagonal matrix (implied by $C \perp_b \tilde{C}$; we also enforce \mathcal{D}_c to have positive, real coefficients). As for the bi-Lanczos algorithm in (2.2), the $(i + 1)$ st Lanczos vectors here are given by

$$\begin{aligned} \gamma v_{i+1} &= Av_i - V_i\tau - C\rho \perp [\tilde{C} \tilde{V}_i], \\ \tilde{\gamma}\tilde{v}_{i+1} &= A^*\tilde{v}_i - \tilde{V}_i\tilde{\tau} - \tilde{C}\tilde{\rho} \perp [C V_i], \end{aligned}$$

where $\gamma, \tilde{\gamma}, \tau, \tilde{\tau}, \rho,$ and $\tilde{\rho}$ are to be determined. The computation of the $(i + 1)$ st Lanczos vector for the primary system now requires the i th and $(i - 1)$ st Lanczos vectors and C (see [3]). This gives a $(3 + k)$ -term recurrence, where k is the number of columns of C . Similarly, we get a $(3 + k)$ -term recurrence for computing the Lanczos vectors for the dual system. We refer to this pair of $(3 + k)$ -term recurrences as the augmented bi-Lanczos relations, and they are given by

$$\begin{aligned} (I - C\hat{C}^*)AV_i &= V_{i+1}\underline{T}_i, \\ (I - \tilde{C}\check{C}^*)A^*\tilde{V}_i &= \tilde{V}_{i+1}\underline{\tilde{T}}_i, \end{aligned}$$

where

$$(3.7) \quad \begin{aligned} \hat{C} &= \begin{bmatrix} \frac{1}{c_1^*e_1}\tilde{c}_1 & \frac{1}{c_2^*e_2}\tilde{c}_2 & \cdots & \frac{1}{c_k^*e_k}\tilde{c}_k \end{bmatrix} = \tilde{C}\mathcal{D}_c^{-1}, \\ \check{C} &= \begin{bmatrix} \frac{1}{\tilde{c}_1^*c_1}c_1 & \frac{1}{\tilde{c}_2^*c_2}c_2 & \cdots & \frac{1}{\tilde{c}_k^*c_k}c_k \end{bmatrix} = C\mathcal{D}_c^{-1}. \end{aligned}$$

THEOREM 3.2. *Let $v_1 = \eta(I - CD_c^{-1}\tilde{C}^*)r_0, \tilde{v}_1 = \tilde{\eta}(I - \tilde{C}\mathcal{D}_c^{-1}C^*)\tilde{r}_0, B = (I - CD_c^{-1}\tilde{C}^*)A,$ and $\tilde{B} = (I - \tilde{C}\mathcal{D}_c^{-1}C^*)A^*,$ where $\eta, \tilde{\eta}$ are scalars and $C, \tilde{C} \in \mathbb{C}^{n \times k}$ s.t. $\mathcal{D}_c = \tilde{C}^*C$ is a diagonal matrix with positive, real coefficients. Also, let (3.5) be used as the set of constraints for (3.3) and (3.4). Then, $\beta_{ij} = 0$ and $\tilde{\beta}_{ij} = 0$ for $j > i + 1,$ which leads to the following short-term recurrences:*

$$\begin{aligned} \beta_{i+1,i}v_{i+1} &= Bv_i - \beta_{ii}v_i - \beta_{i-1,i}v_{i-1}, \\ \tilde{\beta}_{i+1,i}\tilde{v}_{i+1} &= \tilde{B}\tilde{v}_i - \tilde{\beta}_{ii}\tilde{v}_i - \tilde{\beta}_{i-1,i}\tilde{v}_{i-1} \end{aligned}$$

for $i \in \{1, 2, 3, \dots, m - 1\}.$

Proof. We show that conditions (a)–(c) of Theorem 3.1 are satisfied. This demonstrates that augmented bi-Lanczos is a special case of generalized bi-Lanczos. We have $B, \tilde{B} \in \mathbb{C}^{n \times n}$ such that

$$\begin{aligned} B - \tilde{B}^* &= A - CD_c^{-1}\tilde{C}^*A - A + ACD_c^{-1}\tilde{C}^* \\ &= (AC\mathcal{D}_c^{-1})\tilde{C}^* - C(A^*\tilde{C}\mathcal{D}_c^{-1})^*. \end{aligned}$$

Defining $F = A^*\tilde{C}\mathcal{D}_c^{-1}$ and $\tilde{F} = ACD_c^{-1},$ we get

$$B - \tilde{B}^* = \tilde{F}\tilde{C}^* - CF^*, \quad \text{where } C, \tilde{C}, F, \tilde{F} \in \mathbb{C}^{n \times k}.$$

Hence, (a) is satisfied. For any \tilde{x} , consider the following:

$$\begin{aligned} C^* \tilde{B} \tilde{x} &= C^*(I - \tilde{C} \mathcal{D}_c^{-1} C^*) A^* \tilde{x} \\ &= (C^* - \mathcal{D}_c \mathcal{D}_c^{-1} C^*) A^* \tilde{x} = 0. \end{aligned}$$

Similarly, for any x , consider the following:

$$\begin{aligned} \tilde{C}^* B x &= \tilde{C}^*(I - C \mathcal{D}_c^{-1} \tilde{C}^*) A x \\ &= (\tilde{C}^* - \mathcal{D}_c \mathcal{D}_c^{-1} \tilde{C}^*) A x = 0. \end{aligned}$$

Hence, (b) is satisfied. Similarly, for v_1 and \tilde{v}_1 chosen in the theorem, $\tilde{C}^* v_1 = 0$ and $C^* \tilde{v}_1 = 0$. Hence, (c) is satisfied. \square

For ease of future derivations, we introduce a slight change of notation. Let x_{-1} and \tilde{x}_{-1} be the initial guesses and $r_{-1} = b - Ax_{-1}$ and $\tilde{r}_{-1} = \tilde{b} - A^* \tilde{x}_{-1}$ the corresponding initial residuals. We define

$$(3.8) \quad \begin{aligned} x_0 &= x_{-1} + U \hat{C}^* r_{-1}, & r_0 &= (I - C \hat{C}^*) r_{-1}, \\ \tilde{x}_0 &= \tilde{x}_{-1} + \tilde{U} \tilde{C}^* \tilde{r}_{-1}, & \tilde{r}_0 &= (I - \tilde{C} \tilde{C}^*) \tilde{r}_{-1} \end{aligned}$$

and follow this convention for $x_0, \tilde{x}_0, r_0,$ and \tilde{r}_0 for the rest of the paper. Algorithm 1 gives the RBiCG algorithm from [5]. Here, we have not given details on how the recycle space is computed in RBiCG. For that, we refer the reader to [5].

ALGORITHM 1. RBiCG [5].

1. Given U (also $C = AU$) and \tilde{U} (also $\tilde{C} = A^* \tilde{U}$) s.t. $C \perp_b \tilde{C}$, compute \check{C} and \hat{C} using (3.7). If U and \tilde{U} are not available, then initialize U, \tilde{U}, \check{C} , and \hat{C} to empty matrices.
 2. Choose x_{-1}, \tilde{x}_{-1} and compute $x_0, \tilde{x}_0, r_0,$ and \tilde{r}_0 using (3.8).
 3. **if** $(r_0, \tilde{r}_0) = 0$, **then** initialize \tilde{x}_{-1} to a random vector.
 4. Set $p_0 = 0, \tilde{p}_0 = 0, \zeta_c = 0, \tilde{\zeta}_c = 0,$ and $\beta_0 = 0$. Choose `tol` and `max_itn`.
 5. **for** $i = 1 \dots \text{max_itn}$ **do**
 - \diamond $p_i = r_{i-1} + \beta_{i-1} p_{i-1};$ $\tilde{p}_i = \tilde{r}_{i-1} + \tilde{\beta}_{i-1} \tilde{p}_{i-1}$
 - \diamond $z_i = A p_i;$ $\tilde{z}_i = A^* \tilde{p}_i$
 - \diamond $\zeta_i = \hat{C}^* z_i;$ $\tilde{\zeta}_i = \tilde{C}^* \tilde{z}_i$
 - \diamond $q_i = z_i - C \zeta_i;$ $\tilde{q}_i = \tilde{z}_i - \tilde{C} \tilde{\zeta}_i$
 - \diamond $\alpha_i = (\tilde{r}_{i-1}, r_{i-1}) / (\tilde{p}_i, q_i);$ $\tilde{\alpha}_i = \tilde{\alpha}_i$
 - \diamond $\zeta_c = \zeta_c + \alpha_i \zeta_i;$ $\tilde{\zeta}_c = \tilde{\zeta}_c + \tilde{\alpha}_i \tilde{\zeta}_i$
 - \diamond $x_i = x_{i-1} + \alpha_i p_i;$ $\tilde{x}_i = \tilde{x}_{i-1} + \tilde{\alpha}_i \tilde{p}_i$
 - \diamond $r_i = r_{i-1} - \alpha_i q_i;$ $\tilde{r}_i = \tilde{r}_{i-1} - \tilde{\alpha}_i \tilde{q}_i$
 - \diamond **if** $\|r_i\| \leq \text{tol}$ and $\|\tilde{r}_i\| \leq \text{tol}$, **then break**
 - \diamond $\beta_i = (\tilde{r}_i, r_i) / (\tilde{r}_{i-1}, r_{i-1})$
 6. **end for**
 7. $x_i = x_i - U \zeta_c;$ $\tilde{x}_i = \tilde{x}_i - \tilde{U} \tilde{\zeta}_c$
-

Like BiCG, breakdowns can happen in RBiCG as well. Besides the breakdown in the underlying augmented bi-Lanczos algorithm ($\tilde{v}_i^* v_i = 0$ at step i , called *serious breakdown*), a breakdown can happen when pivotless LDU decomposition of the tridiagonal matrix (as in the augmented bi-Lanczos relations discussed earlier) does not exist. A LDU decomposition of a matrix A is of the form $A = LDU$, where L and U are unit triangular matrices and D is a diagonal matrix. This is referred to as a breakdown of the *second kind*.

The breakdown in the augmented bi-Lanczos algorithm can be avoided by using look-ahead strategies [16, 19] (as applied for the bi-Lanczos algorithm). The second breakdown can also be avoided in the same way as in BiCG, that is, by performing the LDU decomposition with 2×2 block diagonal elements [6].

4. Recycling BiCGSTAB. In RBiCG [4, 5], the iteration vectors p , \tilde{p} , r , and \tilde{r} are updated using the following recurrences:

$$\begin{aligned} p_i &= r_{i-1} + \beta_{i-1}p_{i-1}, & \tilde{p}_i &= \tilde{r}_{i-1} + \tilde{\beta}_{i-1}\tilde{p}_{i-1}, \\ r_i &= r_{i-1} - \alpha_i B p_i, & \tilde{r}_i &= \tilde{r}_{i-1} - \tilde{\alpha}_i \tilde{B} \tilde{p}_i, \end{aligned}$$

where $B = (I - C\hat{C}^*)A$ and $\tilde{B} = (I - \tilde{C}\tilde{C}^*)A^*$. We first give the polynomial representations of these iteration vectors.

THEOREM 4.1. *Let r_i , p_i , \tilde{r}_i , and \tilde{p}_i be defined as above. Then, for the primary system*

$$r_i = \Theta_i(B)r_0, \quad p_i = \Pi_{i-1}(B)r_0,$$

where $\Theta_i(K)$ and $\Pi_{i-1}(K)$ are i th and $(i-1)$ st degree polynomials, for an arbitrary square matrix K , that satisfy the following recurrences:

$$\begin{aligned} \Theta_i(K) &= \Theta_{i-1}(K) - \alpha_i K \Pi_{i-1}(K), \\ \Pi_{i-1}(K) &= \Theta_{i-1}(K) + \beta_{i-1} \Pi_{i-2}(K). \end{aligned}$$

Similarly, for the dual system

$$\tilde{r}_i = \bar{\Theta}_i(\tilde{B})\tilde{r}_0, \quad \tilde{p}_i = \bar{\Pi}_{i-1}(\tilde{B})\tilde{r}_0,$$

where $\bar{\Theta}_i(K)$ and $\bar{\Pi}_i(K)$ satisfy the following recurrences:

$$\begin{aligned} \bar{\Theta}_i(K) &= \bar{\Theta}_{i-1}(K) - \bar{\alpha}_i K \bar{\Pi}_{i-1}(K), \\ \bar{\Pi}_{i-1}(K) &= \bar{\Theta}_{i-1}(K) + \bar{\beta}_{i-1} \bar{\Pi}_{i-2}(K). \end{aligned}$$

Proof. This can be proved by induction, following the derivation in [30, sect. 2, pp. 37–40], but using B instead of A and \tilde{B} instead of A^* . \square

From RBiCG, we know $r_i \perp \tilde{r}_j$ for $j < i$. Using Theorem 4.1, we get that

$$(\bar{\Theta}_j(\tilde{B})\tilde{r}_0, \Theta_i(B)r_0) = 0 \text{ for } j < i.$$

This implies $\Theta_i(B)r_0 \perp \mathcal{K}^i(\tilde{B}, \tilde{r}_0)$, where $\tilde{r}_0, \tilde{B}\tilde{r}_0, \dots, \tilde{B}^{i-1}\tilde{r}_0$ span the subspace $\mathcal{K}^i(\tilde{B}, \tilde{r}_0)$. As observed in [32], the above orthogonality conditions must be satisfied by other bases of $\mathcal{K}^i(\tilde{B}, \tilde{r}_0)$, too. So, other polynomials can be used as well [34], that is,

$$(4.1) \quad (\bar{\Omega}_j(\tilde{B})\tilde{r}_0, \Theta_i(B)r_0) = 0 \text{ for } j < i.$$

Similar to the derivation in [32], we define

$$\bar{\Omega}_i(\tilde{B}) = (I - \bar{\omega}_1 \tilde{B})(I - \bar{\omega}_2 \tilde{B}) \cdots (I - \bar{\omega}_i \tilde{B}),$$

where ω_i is selected to minimize the residual r_i with respect to ω_i . Then, as first proposed in [30], instead of (4.1), we use the following form of inner product:

$$(\tilde{r}_0, \Omega_j(B)\Theta_i(B)r_0) = 0 \text{ for } j < i$$

with

$$\Omega_i(B) = (I - \omega_1 B)(I - \omega_2 B) \cdots (I - \omega_i B).$$

This inner product does not require the transpose of B and hence is appropriate when there is no dual system to solve. Computing the inner product in this fashion, we obtain the recycling BiCGSTAB algorithm (similar to the way BiCGSTAB is obtained from BiCG in [32]). We term our recycling BiCGSTAB as RBiCGSTAB. The algorithm is given in Algorithm 2. Some algorithmic improvements to make the code faster (similar to those discussed in section 6.2 of [23]) are not given here.

ALGORITHM 2. RBiCGSTAB.

1. Given U (also $C = AU$) and \tilde{U} (also $\tilde{C} = A^* \tilde{U}$) s.t. $C \perp_b \tilde{C}$, compute $D_c = \tilde{C}^* C$ and $\hat{C} = \tilde{C} D_c^{-1}$.
2. Choose x_{-1} ; initialize \tilde{r}_{-1} to a random vector; and compute x_0, r_0, \tilde{r}_0 using (3.8).
3. **if** $(r_0, \tilde{r}_0) = 0$, **then** rechoose x_{-1} or reinitialize \tilde{r}_{-1} to avoid this condition.
4. Set scalars β_0 and ω_0 as well as vectors p_0, q_0 , and x_c to zero.
5. Choose **tol** and **max_itn**.
6. **for** $i = 1 \dots \text{max_itn}$ **do**
 - ◇ $p_i = r_{i-1} + \beta_{i-1} p_{i-1} - \beta_{i-1} \omega_{i-1} q_{i-1}$
 - ◇ $q_i = A p_i$
 - ◇ $\zeta_i = \hat{C}^* q_i$
 - ◇ $q_i = q_i - C \zeta_i$
 - ◇ $\alpha_i = \frac{(\tilde{r}_0, r_{i-1})}{(\tilde{r}_0, q_i)}$
 - ◇ $s_i = r_{i-1} - \alpha_i q_i$
 - ◇ $t_i = A s_i$
 - ◇ $\gamma_i = \hat{C}^* t_i$
 - ◇ $t_i = t_i - C \gamma_i$
 - ◇ $\omega_i = \frac{(s_i, t_i)}{(t_i, t_i)}$
 - ◇ $x_i = x_{i-1} + \alpha_i p_i + \omega_i s_i$
 - ◇ $x_c = x_c + \alpha_i \zeta_i + \omega_i \gamma_i$
 - ◇ $r_i = r_{i-1} - \alpha_i q_i - \omega_i t_i$
 - ◇ **if** $\|r_i\| \leq \text{tol}$, **then break**
 - ◇ $\beta_i = \frac{(\tilde{r}_0, r_i)}{(\tilde{r}_0, r_{i-1})} \cdot \frac{\alpha_i}{\omega_i}$
7. **end for**
8. $x_i = x_i - U x_c$

Note. When A does not change for multiple systems, several changes should be made to make this algorithm substantially more efficient. This is done in our implementation for the model reduction problem of section 6.2.

Breakdowns in RBiCG (as discussed in the end of section 3) lead to breakdowns in the RBiCGSTAB algorithm as well. This is similar to how breakdowns in BiCG lead to breakdowns in BiCGSTAB. A breakdown free BiCGSTAB algorithm is proposed in [12], which uses the theory of formal orthogonal polynomials. The same theory can be applied to the RBiCGSTAB algorithm.

The BiCGSTAB algorithm also breaks down when the minimization with respect to ω_i fails. This problem can be avoided by minimizing in two or more dimensions. This led to the development of BiCGSTAB2 [18] and BiCGSTAB(l) [29]. In [4], a recycling BiCGSTAB2 that can be extended to a recycling BiCGSTAB(l) is proposed.

5. Numerical experiments. For BiCG, it has been shown that including a left eigenvector into the search space leads to the removal of the corresponding right




FIG. 1. Coefficients for the PDE.

eigenvector from the right residual (and vice versa) [13]. In our experiments, we demonstrate that recycling a left invariant subspace may improve the convergence rate in the RBiCGSTAB algorithm. These experiments demonstrate how the method works but do not show its competitiveness, which is discussed in the next section (application to model order reduction).

We consider two examples. For the first example, we perform vertex centered finite volume discretization of the PDE

$$-(u_x)_x - (u_y)_y + 10u_x - 10u_y = 0$$

on a 42×42 grid unit square resulting in a 1600×1600 linear system. We use the following boundary conditions: $u_{\text{south}} = 1, u_{\text{west}} = 1, u_{\text{north}} = 0, u_{\text{east}} = 0$. We do not use a preconditioner in this example, the initial guess is a vector of all ones, and the relative convergence tolerance is 10^{-10} .

For the second example, we perform finite difference discretization of the partial differential equation [32]

$$-(Av_x)_x - (Av_y)_y + B(x, y)v_x = F$$

with A as shown in Figure 1, $B(x, y) = 2e^{2(x^2+y^2)}$, and $F = 0$ everywhere except in a small square in the center where $F = 100$ (see Figure 1). The $(0, 1) \times (0, 1)$ domain is discretized on a 129×129 grid resulting in a 16129×16129 linear system. We use the following boundary conditions:

$$\begin{aligned} v(0, y) &= v(1, y) = v(x, 0) = 1, \\ v(x, 1) &= 0. \end{aligned}$$

We use an ILUTP [27] preconditioner with a drop tolerance of 0.1 (split-preconditioned). The initial guess is 0.5 times a vector of all ones, and the relative convergence tolerance is 10^{-8} .

For each example, we do three experiments. First, we solve the system without recycling. Second, we use the right invariant subspace (corresponding to the smallest magnitude eigenvalues) as the recycle space. This is implemented by setting $\tilde{U} = U$.

Finally, we use both the left and right invariant subspaces (again, corresponding to the smallest magnitude eigenvalues) as the recycle space.

For the first example's second set of experiments, we use five exact right eigenvectors computed using the MATLAB function `eigs`. For the first example's third set of experiments, we use five exact left eigenvectors and five exact right eigenvectors (for a total of ten), again computed using the MATLAB function `eigs`.

For the second example's second set of experiments, we use twenty approximate right eigenvectors obtained by solving the problem twice with RBiCG. We solved the problem twice so that we can get better approximation of the invariant subspace. For the second example's third set of experiments, we use twenty approximate left eigenvectors and twenty approximate right eigenvectors (for a total of forty), again obtained by solving the problem twice with RBiCG.

The results are shown in Figures 2 (a) and (b). For the first example, using the right invariant subspace or using both the left and right invariant subspaces works equally well. However, for the second example, we see that using only the right invariant subspace leads to convergence that is worse than BiCGSTAB without recycling, and much worse than RBiCGSTAB using both the left and right invariant subspaces. This shows that recycling a left invariant subspace may improve the convergence rate in the RBiCGSTAB algorithm.


Next, we analyze why the first example does not need left invariant subspace and the second example does, by considering the cosines of the principal angles between the left and right invariant subspaces associated with the ten smallest magnitude eigenvalues. Table 1 lists these angles.

From the table, we see that the principal angles between the left and right invariant subspaces for the second example are substantially larger than those for the first example (since the cosines of the principal angles is lesser for the second example as compared with the first example). Since for a normal matrix, left and right invariant subspaces are identical, we conclude that the second example is "more" nonnormal than the first.

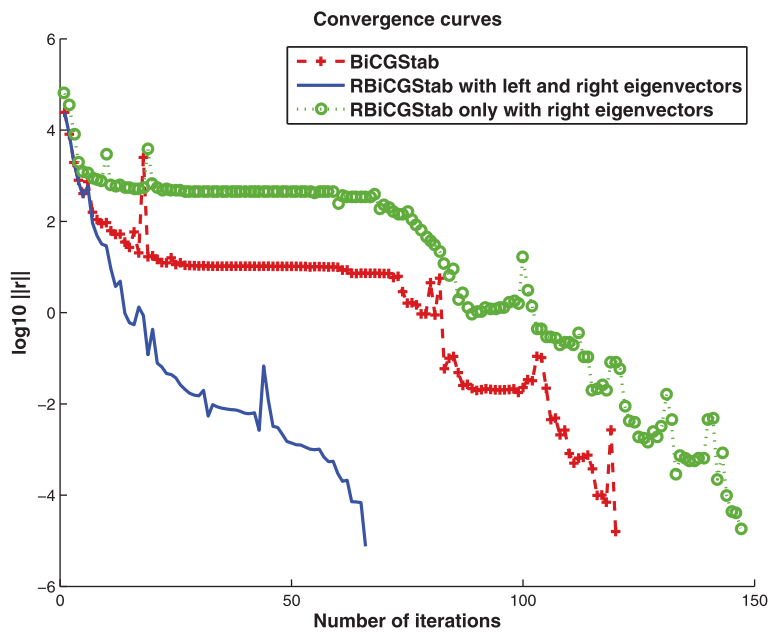
6. Applications. We first discuss several techniques for parametric model order reduction and the one we are using (in section 6.1) followed by the description of how recycling BiCGSTAB is applied for sequences of linear systems arising in parametric model order reduction (in section 6.2).

6.1. Parametric model order reduction. Numerical simulation is an essential tool for solving science and engineering problems. However, simulating large-scale models leads to overwhelming demands on computational resources. This is the main motivation for model reduction. The goal is to produce a surrogate model of much smaller dimension that provides a high-fidelity approximation of the input-output behavior of the original model. Often the models have design parameters associated with them, e.g., boundary conditions, geometry, material properties etc. Changes in these design parameters require the generation of new reduced models, which makes the model reduction process very cumbersome. One practical application where such a challenge arises is micro-electro-mechanical systems (MEMS) design [15, 9]. The goal of PMOR [7, 14] is to generate a reduced model such that parametric dependence, as in the original model, is preserved (or retained).

We focus on physical processes that are modeled as parameterized partial differential equations (PDEs). For PMOR, the PDE is first semidiscretized using classical techniques (e.g., finite differences, finite elements, etc.), and then model reduction is



(a) Example 1: Left eigenvectors not needed, but recycling effective.



(b) Example 2: Left eigenvectors needed for recycling to be effective.

FIG. 2. Convergence curves for two examples using RBiCGSTAB. The second example demonstrates that recycling a left invariant subspace may improve the convergence rate in the RBiCGSTAB algorithm.

applied to the resulting parameterized state-space model [7]:

$$(6.1) \quad G: \begin{cases} E(p) \dot{x}(t) = A(p) x(t) + B(p) u(t), \\ y(t) = C(p) x(t) \end{cases}$$

TABLE 1

For each example, we give the cosines of the principal angles between the exact left and right invariant subspaces of dimension 10, associated with the smallest magnitude eigenvalues. For the first example, we compute invariant subspaces of the matrix obtained after discretization. While for the second example, we compute invariant subspaces of the preconditioned matrix since we split-precondition the linear system obtained after discretization.

Example 1	Example 2
0.9998	0.1039
0.9837	0.0302
0.9207	0.0195
0.7617	0.0106
0.4352	0.0089
0.3987	0.0049
0.2273	0.0043
0.0963	0.0027
0.0235	0.0018
0.0064	0.0012

or

$$(6.2) \quad G(s, p) = C(p) (sE(p) - A(p))^{-1} B(p).$$

Sometimes, the physical process is directly available in the form of a parametrized state-space model. Equations (6.1)–(6.2) represent a multiple input multiple output (MIMO) linear dynamical system, where p is the parameter vector; $u(t) = [u_1(t), \dots, u_m(t)]^T: \mathbb{R} \rightarrow \mathbb{R}^m$ is the input; $y(t): \mathbb{R} \rightarrow \mathbb{R}^l$ is the output; $x(t): \mathbb{R} \rightarrow \mathbb{R}^n$ is the state vector; $E(p), A(p) \in \mathbb{R}^{n \times n}$, $B(p) \in \mathbb{R}^{n \times m}$, and $C(p) \in \mathbb{R}^{l \times n}$ are the system matrices; and s is the frequency domain variable corresponding to t in the time domain.

Above, (6.1) denotes the dynamical system, and (6.2) gives the transfer function of the system obtained after Laplace transformation. By a common abuse of notation, we denote both with G . The dimension of the underlying state-space, n , is called the dimension or order of G . Unless explicitly stated, for the rest of this paper all dynamical systems are assumed to be of the above type.

There are various ways of performing PMOR [11, 7, 26, 10]. This includes moment matching, local \mathcal{H}_2 -optimality, and reduced basis approaches. For this work, we focus on moment matching based PMOR because of its flexibility (few limits on the system properties) and low computational cost in many industrial applications.

Moment matching based PMOR algorithms [10, 15] require solution of sequences of linear systems of the type (1.1), which is a key bottleneck when using these algorithms for reducing larger models. Specifically, the systems have the form as follows:

$$\begin{aligned} A^{(1)} x^{(1,1)} &= b^{(1,1)}, \\ A^{(1)} x^{(1,2)} &= b^{(1,2)}, \\ A^{(1)} x^{(1,3)} &= b^{(1,3)}, \\ &\vdots \\ A^{(2)} x^{(2,1)} &= b^{(2,1)}, \\ A^{(2)} x^{(2,2)} &= b^{(2,2)}, \\ A^{(2)} x^{(2,3)} &= b^{(2,3)}, \\ &\vdots \end{aligned}$$

In the example used here, each matrix has the form

$$A^{(i)} = A_0 + \sigma_i A_1 + \delta_i A_2 \quad \text{with } i \in \mathbb{N}.$$

Please note that $\sigma_i A_1$ and $\delta_i A_2$ are small perturbations to A_0 . As discussed in the introduction, our goal here is to use Krylov subspace recycling (recycling BiCGSTAB specifically) to efficiently solve such systems.

6.2. Application to PMOR. Our test dynamical system comes from a silicon nitride membrane model [8]. Such a membrane can be part of many devices, e.g., a gas sensor chip, a microthruster, an optical filter, etc. We use the moment matching based PMOR algorithm, described in [10, 15], to compute a reduced model. This leads to a sequence of linear systems of the form (1.1) and size 60,020.

Whenever the matrix changes in the sequence, we call RBiCG to perform the linear solve. This helps to approximate both left and right invariant subspaces, which are not easily available from the RBiCGSTAB iterations. (Sometimes a left invariant subspace is available from a right invariant subspace [2, 24].) The primary system right-hand side comes from the PDE. We take a vector of all ones as the dual system right-hand side. We call RBiCGSTAB for all remaining systems with the same matrix. This corresponds to linear systems where only the right-hand sides change. This is an effective strategy because it has been shown that the recycle space can be useful for multiple consecutive systems [25, 21, 24, 1]. Moreover, using RBiCG for all systems will be expensive since an unnecessary dual system would be solved at each step in the sequence. It needs to be emphasized here that there is a need to implement this efficiently (as discussed in Algorithm 2).

While solving a linear system with RBiCG, Lanczos vectors are generated at each iterative step. These Lanczos vectors are used to build the recycle space (U , \tilde{U} , C , and \tilde{C}). We have the flexibility in deciding when to build the recycle space. One option is to wait for RBiCG to converge, save all the Lanczos vectors, and then build the recycle space. The problem with this approach is that this requires large amounts of memory and may be computationally expensive as well. Instead, we divide the RBiCG iteration in cycles of a certain number of iterations (to be chosen, e.g., 50 iterations). At the end of each cycle, we use the stored Lanczos vectors from that cycle to build or improve the recycle space and then discard these Lanczos vectors, except for the last few that are needed to continue the RBiCG iteration. Note that we do not restart the RBiCG iteration. It is kept on hold until we finish updating the recycle space. At the end of the first cycle, we may build a new recycle space or update a recycle space constructed while solving a previous linear system. For more details on computing the recycle space, please see section 4.1 in [5]. In the following discussion, we denote the length of a cycle by s .

For this experiment, we take $s = 25$ and $k = 20$ (the number of vectors selected for recycling as defined earlier in section 3). These values are chosen based on experience with other recycling algorithms [25]. The linear systems are split-preconditioned with an incomplete LU preconditioner with threshold and pivoting (ILUTP) [27]. The drop tolerance is taken as 10^{-4} . For RBiCG, we take a vector of all zeros as the initial guess for both the primary system and the dual system. For RBiCGSTAB, we take a vector of all zeros as the initial guess as well. (We are only solving the primary system in RBiCGSTAB.) Using the solution from the previous linear system as an initial guess leads to poor convergence for BiCGSTAB with and without recycling. In general, a better initial guess may be based on knowledge of the system. The relative convergence tolerance for the iterative solves is taken as 10^{-8} .

The number of matrix-vector products required to solve systems 1 through 63 are given in Figure 3, and the corresponding timing data is given in Figure 4. In both the figures, the peaks in the recycling BiCGSTAB plot correspond to when the matrix changes and RBiCG is called (three times, at the 1st, 22nd, and 43rd linear system). For all other steps, when only the right-hand side changes, RBiCGSTAB is called.

When recomputing the recycle space for the 22nd system, we use the recycle space generated while solving the 1st system, and improve it. This is evident in Figures 3 and 4, where both the number of matrix-vector products and time for solving the 22nd linear system are less than those of the 1st system. While solving the 22nd linear system, if we do not use the recycle space generated after solving the 1st system, then the matrix-vector product count is 78, which is larger than 66 matrix-vector products required when we do use the recycle space from the 1st system (as reported in Figure 3).

The same process happens when recomputing the recycle space for the 43rd system. Here, the recycle space from the 22nd linear system is improved. Again, the Figures 3 and 4 show the decrease in the matrix-vector product count and time for the 43rd system as compared with the 22nd system. While solving the 43rd linear system, if we do not use the recycle space generated after solving the 22nd system, then the matrix-vector product count is 72, which is larger than 58 matrix-vector products required when we do use the recycle space from the 22nd system (as reported in Figure 3).

First, we compare our results with BiCGSTAB. The RBiCG and RBiCGSTAB combination requires about 40% fewer matrix-vector products in total. Also, computing the reduced model with recycling takes about 35% less time than without recycling. This demonstrates the effectiveness of recycling Krylov subspaces for PMOR.

Second, we compare our results with GCRO-DR. Although the RBiCG and RBiCGSTAB combination does not beat GCRO-DR in the number of matrix vector products, it is 10% more efficient than GCRO-DR in time and reduces storage requirements. We often see the same when comparing BiCGSTAB with GMRES for some examples. That is, although BiCGSTAB is more expensive than GMRES in the number of matrix vector products, it is cheaper than GMRES in time, the reason being that BiCGSTAB is cheaper than GMRES with respect to the orthogonalizations.

7. Conclusions. For solving linear systems with nonsymmetric matrices, BiCGSTAB is one of the best available algorithms. As compared with GMRES, which is the most commonly used algorithm for such linear systems, it has the advantage of a short term recurrence and has cheaper iterations, which often makes it faster in time.

For solving sequences of linear systems with nonsymmetric matrices, it is advantageous to use Krylov subspace recycling for the BiCGSTAB algorithm, and hence we propose the RBiCGSTAB algorithm. We have demonstrated the usefulness of RBiCGSTAB for a parametric model order reduction example.

Here, we have used RBiCG to solve a linear system and generate the recycle space. We have then used RBiCGSTAB, using this generated space, to solve the subsequent systems until the matrix changes. RBiCG in this approach can be replaced with other solvers like GCRO-DR [25] as well.

In the future, we plan to test RBiCGSTAB for other application areas (e.g., acoustics problems). We also plan to extend the recycling framework of RBiCGSTAB to BiCGSTAB(l) [29] and IDR [33, 31]. In section 4, we saw that BiCGSTAB (and

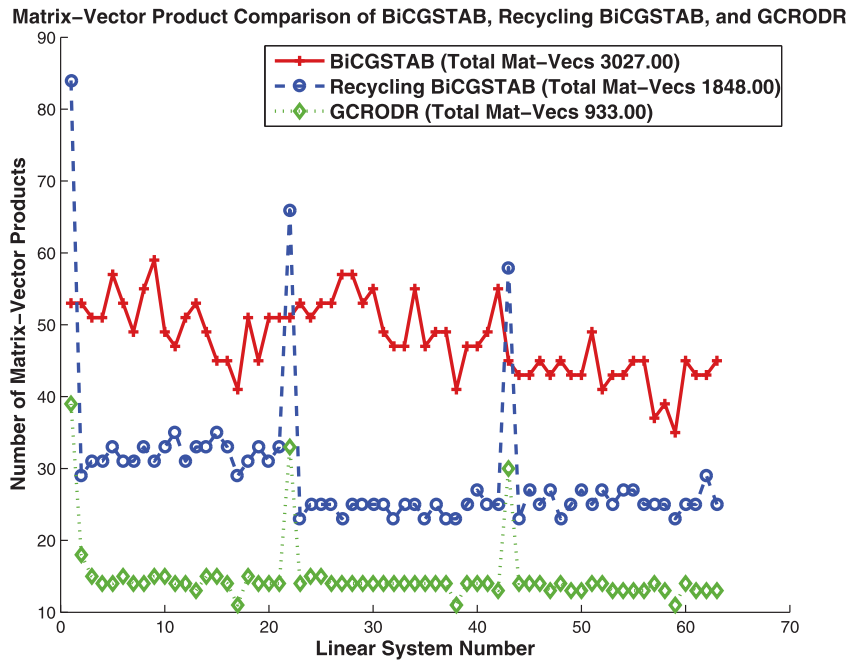


FIG. 3. Comparison of matrix-vector product count when using BiCGSTAB, RBiCGSTAB, and GCRODR as the linear solvers for PMOR.

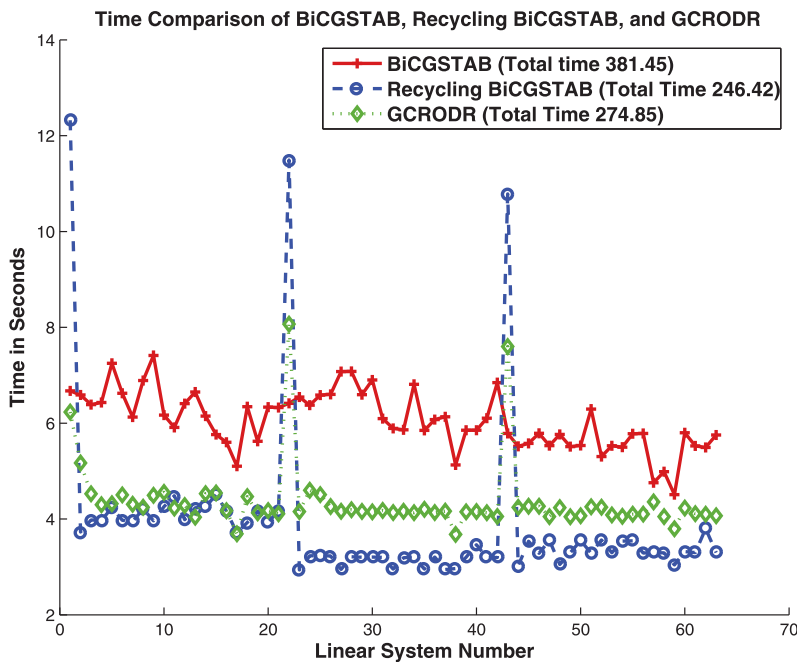


FIG. 4. Comparison of time (in seconds) when using BiCGSTAB, RBiCGSTAB, and GCRODR as the linear solvers for PMOR.

Downloaded 11/17/15 to 128.173.38.51. Redistribution subject to SIAM license or copyright; see http://www.siam.org/journals/ojsa.php

RBiCGSTAB) performs one-dimensional minimization of the residual. This minimization can be done in higher dimensions as well (say, l), leading to BiCGSTAB(l).³ Like BiCGSTAB, the induced dimension reduction method involves a short term recurrence and has been shown to perform better than BiCGSTAB in many cases [31].

Acknowledgment. We thank the anonymous reviewers for their careful and helpful suggestions, which greatly helped us to improve this paper.

REFERENCES

- [1] A. M. ABDEL-REHIM, A. STATHOPOULOS, AND K. ORGINOS, *Extending the eigCG algorithm to nonsymmetric Lanczos for linear systems with multiple right-hand sides*, Numer. Linear Algebra Appl., 21 (2014), pp. 473–493.
- [2] A. M. ABDEL-REHIM, W. WILCOX, AND R. B. MORGAN, *Deflated BiCGStab for linear equations in QCD problems*, in Proceedings of Science, 2007, pp. 026/1–026/7.
- [3] K. AHUJA, *Recycling Bi-Lanczos Algorithms: BiCG, CGS, and BiCGSTAB*, Master's thesis, Department of Mathematics, Virginia Tech, 2009; also available online from <http://scholar.lib.vt.edu/theses/available/etd-08252009-161256/>.
- [4] K. AHUJA, *Recycling Krylov Subspaces and Preconditioners*, Ph.D. thesis, Department of Mathematics, Virginia Tech, 2011; also available online from <http://scholar.lib.vt.edu/theses/available/etd-11112011-010340/>.
- [5] K. AHUJA, E. DE STURLER, S. GUGERCIN, AND E. CHANG, *Recycling BiCG with an application to model reduction*, SIAM J. Sci. Comput., 34 (2012), pp. A1925–A1949.
- [6] R. E. BANK AND T. F. CHAN, *An analysis of the composite step biconjugate gradient method*, Numer. Math., 66 (1993), pp. 295–319.
- [7] U. BAUR, C. BEATTIE, P. BENNER, AND S. GUGERCIN, *Interpolatory projection methods for parameterized model reduction*, SIAM J. Sci. Comput., 33 (2011), pp. 2489–2518.
- [8] T. BECHTOLD, D. HOHLFELD, E. B. RUDNYI, AND M. GÜNTHER, *Efficient extraction of thin-film thermal parameters from numerical models via parametric model order reduction*, J. Micromechanics Microengineering, 20 (2010), 045030.
- [9] T. BECHTOLD, G. SCHRAG, AND L. FENG, eds., *System-Level Modeling of MEMS*, Advanced Micro & Nanosystems, Wiley-VCH, Weinheim, Germany, 2013.
- [10] P. BENNER AND L. FENG, *A robust algorithm for parametric model order reduction based on implicit moment matching*, in Reduced Order Methods for Modeling and Computational Reduction, A. Quarteroni and G. Rozza, eds., MS&A. Model. Simul. Appl. 9, Springer, Berlin, 2014, pp. 159–186.
- [11] P. BENNER, S. GUGERCIN, AND K. WILLCOX, *A survey of model reduction methods for parametric systems*, Technical report MPIMD/13-14, Max Planck Institute, Magdeburg, 2013.
- [12] Z.-H. CAO, *Avoiding breakdown in variants of the BI-CGSTAB algorithm*, Linear Algebra Appl., 263 (1997), pp. 113–132.
- [13] E. DE STURLER, *BiCG explained*, in Householder Symposium XIV, Proceedings of the Householder International Symposium in Numerical Algebra, Chateau Whistler, Whistler, BC, Canada, 1999, pp. 254–257.
- [14] L. FENG, P. BENNER, AND J. KORVINK, *Parametric model order reduction accelerated by subspace recycling*, in Proceedings of 48th IEEE Conference on Decision & Control and 28th Chinese Control Conference, 2009, pp. 4328–4333.
- [15] L. FENG, P. BENNER, AND J. G. KORVINK, *Subspace recycling accelerates the parametric macro-modeling of MEMS*, Int. J. Numer. Methods Eng., 94 (2013), pp. 84–110.
- [16] R. W. FREUND, M. H. GUTKNECHT, AND N. M. NACHTIGAL, *An implementation of the lookahead Lanczos algorithm for non-Hermitian matrices*, SIAM J. Sci. Comput., 14 (1993), pp. 137–158.
- [17] A. GREENBAUM, *Iterative Methods for Solving Linear Systems*, SIAM, Philadelphia, 1997.
- [18] M. H. GUTKNECHT, *Variants of BICGSTAB for matrices with complex spectrum*, SIAM J. Sci. Stat. Comput., 14 (1993), pp. 1020–1033.
- [19] M. H. GUTKNECHT, *Lanczos-type solvers for nonsymmetric linear systems of equations*, Acta Numer., 6 (1997), pp. 271–397.
- [20] M. H. GUTKNECHT, *Deflated and Augmented Krylov Subspace Methods: A Framework for Deflated BiCG and Related Solvers*, <http://www.sam.math.ethz.ch/~mhg/> (2014).

³One-dimensional minimization was first extended to two dimensions in [18].

- [21] M. E. KILMER AND E. DE STURLER, *Recycling subspace information for diffuse optical tomography*, SIAM J. Sci. Comput., 27 (2006), pp. 2140–2166.
- [22] C. LANCZOS, *Solution of systems of linear equations by minimized iterations*, J. Res. Natl. Bureau Standards, 49 (1952), pp. 33–53.
- [23] L. A. MOTTA MELLO, E. DE STURLER, G. H. PAULINO, AND E. C. NELLI SILVA, *Recycling Krylov subspaces for efficient large-scale electrical impedance tomography*, Comput. Methods Appl. Mech. Engrg., 199 (2010), pp. 3101–3110.
- [24] R. B. MORGAN AND D. A. NICELY, *Restarting the nonsymmetric Lanczos algorithm for eigenvalues and linear equations including multiple right-hand sides*, SIAM J. Sci. Comput., 33 (2011), pp. 3037–3056.
- [25] M. L. PARKS, E. DE STURLER, G. MACKEY, D. D. JOHNSON, AND S. MAITI, *Recycling Krylov subspaces for sequences of linear systems*, SIAM J. Sci. Comput., 28 (2006), pp. 1651–1674.
- [26] A. T. PATERA AND G. ROZZ, *Reduced basis approximation and a posteriori error estimation for parameterized partial differential equations*, Version 1.0, to appear in MIT Papalardo Graduate Monographs in Mechanical Engineering; also available online from <http://augustine.mit.edu> (2006).
- [27] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003.
- [28] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [29] G. L. G. SLEIJPEN AND D. R. FOKKEMA, *BiCGstab(1) for linear equations involving unsymmetric matrices with complex spectrum*, Electron. Trans. Numer. Anal., 1 (1993), pp. 11–32.
- [30] P. SONNEVELD, *CGS, a fast Lanczos-type solver for nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 10 (1989), pp. 36–52.
- [31] P. SONNEVELD AND M. B. VAN GIJZEN, *IDR(s): A family of simple and fast algorithms for solving large nonsymmetric systems of linear equations*, SIAM J. Sci. Comput., 31 (2008), pp. 1035–1062.
- [32] H. A. VAN DER VORST, *Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 13 (1992), pp. 631–644.
- [33] P. WESSELING AND P. SONNEVELD, *Numerical experiments with a multiple grid and a preconditioned Lanczos type method*, in Approximation Methods for Navier–Stokes Problems, R. Rautmann, ed., Lecture Notes in Math. 771, Springer, Berlin 1980, pp. 543–562.
- [34] S.-L. ZHANG, *GPBi-CG: Generalized product-type methods based on Bi-CG for solving nonsymmetric linear systems*, SIAM J. Sci. Comput., 18 (1997), pp. 537–551.