Analysis of the BiCG Method

Marissa Renardy

Thesis submitted to the Faculty of the
Virginia Polytechnic Institute and State University
in partial fulfillment of the requirements for the degree of

Master of Science
in
Mathematics

Eric de Sturler, Chair
John F. Rossi
Peter A. Linnell

May 6, 2013
Blacksburg, Virginia

Keywords: Krylov methods, BiCG, GMRES, FOM

Analysis of the BiCG Method

Marissa Renardy

(ABSTRACT)

The Biconjugate Gradient (BiCG) method is an iterative Krylov subspace method that utilizes a 3-term recurrence [7]. BiCG is the basis of several very popular methods, such as BiCGStab [16]. The short recurrence makes BiCG preferable to other Krylov methods because of decreased memory usage and CPU time. However, BiCG does not satisfy any optimality conditions and it has been shown that for up to $\frac{n}{2} - 1$ iterations, a special choice of the left starting vector can cause BiCG to follow *any* 3-term recurrence [8]. Despite this apparent sensitivity, BiCG often converges well in practice. This paper seeks to explain why BiCG converges so well, and what conditions can cause BiCG to behave poorly. We use tools such as the singular value decomposition and eigenvalue decomposition to establish bounds on the residuals of BiCG and make links between BiCG and optimal Krylov methods.

The dedication is left to the reader as an exercise.

Acknowledgments

I would like to thank my committee members for their help and support. In particular, I would like to thank Eric de Sturler for guiding me through this research.

Contents

Chapter 1

Introduction

Krylov subspace methods are a class of iterative methods for solving linear systems that find approximate solutions in successively larger Krylov subspaces. Solutions x_i are found by a projection that is defined by orthogonalization of the residual $b - Ax_i$ against some subspace of the same dimension as the search space. For the purposes of this paper, we are concerned with three such methods: the Generalized Minimal Residual Method (GMRES), Full Orthogonalization Method (FOM), and Biconjugate Gradient Method (BiCG) [14], [12], [7]. The analysis of BiCG will be our main focus. Each of these methods utilizes a different subspace to define the projection, which leads to different convergence behaviors. The convergence behavior of GMRES and FOM are relatively well-understood. GMRES is optimal in the sense that it minimizes the 2-norm of the residual at each step. FOM is not optimal in general, but uses a similar projection space as GMRES. It has been shown that when GMRES is converging well, FOM exhibits similar behavior. In an iteration where GMRES makes no or very little progress, the FOM residual can be undefined or grow significantly [2].

BiCG does not minimize the residual nor the error and hence satisfies no optimality condition. However, BiCG can be preferable because the projection utilized in BiCG results in a 3-term recurrence rather than a full orthogonalization as in GMRES and FOM. Hence, in the cases where BiCG converges, it is often far less expensive computationally. However, the convergence behavior of BiCG is still not well understood. In theory, it has been shown that BiCG can be made to behave almost arbitrarily through the special choice of a left starting vector [8]. In practice, however, with a randomly chosen left starting vector, BiCG has been seen to converge well in most cases. We seek to explain the convergence behavior of BiCG through a theoretical analysis of the method.

Chapter 2

Background

2.1 Iterative Methods

Suppose we are given a large linear system $Ax = b$. If we choose to solve this system directly, for instance by Gaussian elimination, we are theoretically guaranteed to find an exact solution (if it exists) in a finite number of steps. For large systems, however, direct methods are usually too expensive. Iterative methods provide a cheaper alternative. An iterative method continues to approximate the solution until it satisfies a convergence criterion. The main motivation for iterative methods is usually to reduce computer storage and CPU time [17, Chapter 3]. In this paper, we are concerned with a class of iterative methods called Krylov subspace methods.

2.2 Krylov Subspace Methods

The Krylov subspace of dimension i, generated by a matrix A and a vector v, is defined as

$$\mathcal{K}^i(A; v) = \text{span}(v, Av, ..., A^{i-1}v).$$

Krylov subspace methods are iterative methods for solving linear systems $Ax = b$ where $A \in \mathbb{C}^{n\times n}$. At each step i, a Krylov subspace method seeks an approximation x_i in the subspace $x_0 + \mathcal{K}^i(A; r_0)$, where x_0 is an initial guess and r_0 denotes the initial residual $r_0 = b - Ax_0$. In essence, Krylov methods seek to solve an n-dimensional problem through a sequence of lower dimensional problems [15, Lecture 32]. The dimension of the Krylov subspace increases with each iteration until a satisfactory solution is reached.

It is important to note that since the vectors $A^i r_0$ tend toward the direction of the dominant eigenvector, Krylov subspace methods do not use the basis $\{r_0, Ar_0, ..., A^{i-1}r_0\}$. Instead, these methods most commonly use the Arnoldi method with modified Gram-Schmidt or the Lanczos method to build more suitable bases for the Krylov spaces. These methods will be described in the next section.

There are four types of Krylov subspace methods, as presented in [17, Chapter 3]:

1. *The Ritz-Galerkin approach*: Construct the x_i for which $b - Ax_i \perp \mathcal{K}^i(A; r_0)$.

2. *The minimum norm residual approach*: Identify the x_i in $\mathcal{K}^i(A; r_0)$ for which $\|b - Ax_i\|_2$ is minimal.

3. *The Petrov-Galerkin approach*: Find an x_i so that $b - Ax_i$ is orthogonal to some other suitable i-dimensional subspace.

4. *The minimum norm error approach*: Determine x_i in $A^T\mathcal{K}^i(A^T; r_0)$ for which $\|x_i - x\|_2$ is minimal, where x denotes the true solution.

For Hermitian positive definite matrices, the Ritz-Galerkin approach is optimal in the sense that it minimizes the A-norm of the error. This leads to the conjugate gradients (CG) method, which is the most prominent Krylov method [9]. For general matrices, the Ritz-Galerkin approach leads to FOM, the minimum norm residual approach leads to GMRES, and the Petrov-Galerkin approach leads to BiCG. In this paper, we will not discuss the minimum norm error approach.

2.3 The Arnoldi and Lanczos Iterations

Given a matrix A, the Arnoldi iteration is used to compute a unitary matrix V and a Hessenberg matrix H such that $A = VHV^*$ [1]. This is done using modified Gram-Schmidt orthogonalization with the condition

$$Av_i = h_{1,i}v_1 + ... + h_{i,i}v_i + h_{i+1,i}v_{i+1} \tag{2.1}$$

[15, Lecture 33]. The algorithm is shown below.

Algorithm 1 Arnoldi Iteration [15, Lecture 33]

Choose v_1 with $\|v_1\|_2 = 1$.
for $m = 1, 2, ...$ **do**
    $v = Av_m$
    for $j = 1, ..., m$ **do**
        $h_{j,m} = v_j^* v$
        $v = v - h_{j,m}v_j$
    end for
    $h_{m+1,m} = \|v\|$
    $v_{m+1} = v/h_{m+1,m}$
end for

At the ith step of the Arnoldi iteration, we obtain a partial reduction $AV_i = V_{i+1}H_{i+1,i}$ where V_i consists of the first i columns of V and $H_{i+1,i}$ is the $(i+1) \times i$ upper left section of H. This partial reduction is utilized in Krylov methods such as GMRES and FOM, which will be discussed in the next section.

When A is Hermitian, then it follows from $H = V^*AV$ that H is Hermitian. Then since H is also upper Hessenberg, H must be tridiagonal. Thus, (2.1) is replaced by

$$Av_i = h_{i-1,i}v_{i-1} + h_{i,i}v_i + h_{i+1,i}v_{i+1} \tag{2.2}$$

This reduces the Arnoldi iteration to a 3-term recurrence, resulting in the Lanczos iteration [11]. The partial reductions resulting from the Lanczos iteration are utilized in the CG method. Denoting the diagonal elements of H by $\alpha_1, ..., \alpha_n$ and the super-diagonal elements by $\beta_1, ..., \beta_{n-1}$, we obtain Algorithm 2.

Algorithm 2 Lanczos Iteration [15, Lecture 36]

Choose v_1 with $\|v_1\|_2 = 0$. Set $\beta_0 = 0$ and $v_0 = 0$
for $j = 1, 2, ...$ **do**
    $\hat{v} = Av_j$
    $\alpha_j = v_j^* \hat{v}$
    $\hat{v} = \hat{v} - \beta_{j-1} v_{j-1} - \alpha_j v_j$
    $\beta_j = \|\hat{v}\|$
    $v_{j+1} = \hat{v}/\beta_j$
end for

When A is not Hermitian, the Lanczos iteration can be extended to the "nonsymmetric" Lanczos iteration, which builds two biorthogonal sequences instead of one orthogonal sequence [11]. The nonsymmetric Lanczos iteration builds biorthogonal bases for the Krylov subspaces $\mathcal{K}^i(A; v_1)$ and $\mathcal{K}^i(A^*; w_1)$. This is utilized in BiCG. The process for building these bases is shown in Algorithm 3.

Algorithm 3 Nonsymmetric Lanczos Iteration [13, Chapter 7]

Choose v_1, w_1 such that $w_1^* v_1 = 1$.
$\beta_0 = \delta_0 = 0, v_0 = w_0 = 0$
for $j = 1, 2, ...$ **do**
    $\alpha_j = w_j^* A v_j$
    $\hat{v}_{j+1} = Av_j - \alpha_j v_j - \beta_j v_{j-1}$
    $\hat{w}_{j+1} = A^* w_j - \bar{\alpha_j} w_j - \bar{beta_j} w_{j-1}$
    $\delta_{j+1} = |\hat{w}_{j+1}^* \hat{v}_{j+1}|^{1/2}$. If $\delta_{j+1} = 0$, stop.
    $\beta_{j+1} = \hat{w}_{j+1}^* \hat{v}_{j+1}/\delta_{j+1}$
    $w_{j+1} = \hat{w}_{j+1}/\beta_{j+1}$
    $v_{j+1} = \hat{v}_{j+1}/\delta_{j+1}$
end for

2.4 GMRES and FOM

Via the Arnoldi iteration, we can obtain an orthonormal basis $\{v_1, ..., v_{i+1}\}$ for $\mathcal{K}^{i+1}(A; r_0)$ that satisfies $AV_i = V_{i+1} H_{i+1,i}$ where $H_{i+1,i}$ is upper Hessenberg and $V_i = \begin{bmatrix} v_1 & ... & v_i \end{bmatrix}$ is a basis for $K^i(A; r_0)$ [17, Chapter 3].

GMRES seeks to compute $x_i \in \mathcal{K}^i(A; r_0)$ such that the residual norm $\|r_i\|_2$ is minimized. Since $x_i \in \mathcal{K}^i(A; r_0)$, we can write $x_i = V_i y$ for some y. Then

$$r_i = r_0 - Ax_i = r_0 - AV_i y = r_0 - V_{i+1} H_{i+1,i} y = V_{i+1}(\|r_0\|_2 e_1 - H_{i+1,i} y). \tag{2.3}$$

So minimizing $\|r_0 - Ax_i\|_2$ is equivalent to minimizing $\|V_{i+1}(\|r_0\|_2 e_1 - H_{i+1,i}y)\|_2$. This is further equivalent to minimizing $\|\|r_0\|_2 e_1 - H_{i+1,i}y\|_2$. GMRES solves this least squares problem for y and then computes $x_i = V_i y$. This is demonstrated in Algorithm 4.

Algorithm 4 Generalized Minimal Residual Method (GMRES) Algorithm [13]

Compute $r_0 = b - Ax_0$, $\beta = \|r_0\|_2$, and $v_1 = r_0/\beta$.
for $m = 1, 2, ...$ until convergence **do**
  Define the $(m+1) \times m$ matrix $H_m = \{h_{i,j}\}_{1 \le i \le m+1, 1 \le j \le m}$. Set $H_m = 0$.
  for $j = 1, 2, ..., m$ **do**
    $w_j = Av_j$.
    for $i = 1, ..., j$ **do**
      $h_{i,j} = v_i^* w_j$
      $w_j = w_j - h_{i,j} v_i$
    end for
    $h_{j+1,j} = \|w_j\|_2$. If $h_{j+1,j} = 0$, perform 13-15 and stop.
    $v_{j+1} = w_j / h_{j+1,j}$
  end for
  Compute y_m to minimize $\|\beta e_1 - H_m y\|_2$.
  $x_m = x_0 + V_m y_m$
  $r_m = r_0 - V_{m+1} H_{m+1,m} y_m$
end for

This results in an orthogonal projection method with projection space $A\mathcal{K}^i(A; r_0)$, as shown in Theorem 2.2. First, we need the following lemma.

Lemma 2.1. *Let $\mathcal{U}$ be an inner product space with inner product $\langle \cdot, \cdot \rangle_\alpha$, let $\mathcal{V}$ be a subspace of $\mathcal{U}$, and let $\hat{x} \in \mathcal{U}$. Then $x \in \mathcal{V}$ minimizes $\|\hat{x} - x\|_\alpha$ if and only if $\langle \hat{x} - x, v \rangle_\alpha = 0$ for all $v \in \mathcal{V}$.*

Proof. See Section 8.9 of [5]. □

Theorem 2.2. *$x_i \in \mathcal{K}^i(A; r_0)$ minimizes $\|r_i\|_2$ if and only if $r_i \perp A\mathcal{K}^i(A; r_0)$.*

Proof. Note that since $x_i \in \mathcal{K}^i(A; r_0)$, $Ax_i \in A\mathcal{K}^i(A; r_0)$. Then letting $\mathcal{V} = A\mathcal{K}^i(A; r_0)$ and $r_0 = \hat{x}$ in the above lemma, we see that $Ax_i \in \mathcal{V}$ minimizes $\|r_i\|_2 = \|r_0 - Ax_i\|_2$ if and only if $r_i = r_0 - Ax_i \perp A\mathcal{K}^i(A; r_0)$. Hence $x_i \in \mathcal{K}^i(A; r_0)$ minimizes $\|r_i\|_2$ if and only if $r_i \perp A\mathcal{K}^i(A; r_0)$. □

Hence, solving the least squares problem $\min \|r_0 - Ax_i\|_2$ in GMRES is equivalent to finding $x_i \in \mathcal{K}^i(A; r_0)$ such that $r_i = r_0 - Ax_i \perp A\mathcal{K}^i(A; r_0)$. This projection space distinguishes GMRES from FOM and BiCG. For our analysis in the following chapters, we will identify GMRES by this projection.

FOM is an oblique projection method based on the Ritz-Galerkin condition $r_0 - Ax_i \perp \mathcal{K}^i(A; r_0)$. It utilizes a similar process as GMRES, but with this slightly different projection space. Letting $v_1 = r_0/\|r_0\|_2$ in the Arnoldi iteration, we get $V_i^* A V_i = H_i$ and $V_i^* r_0 =$

$V_i^*(\|r_0\|_2 v_1) = \|r_0\|_2 e_1$. Then writing $x_i = V_i y_i$,

$$\begin{aligned} V_i^*(r_0 - AV_i y_i) &= V_i^* r_0 - V_i^* A V_i y_i \\ &= \|r_0\|_2 e_1 - H y_i. \end{aligned}$$

Hence $V_i^*(r_0 - Ax_i) = 0$ if and only if $Hy_i = \|r_0\|_2 e_1$, so the solution to the Galerkin condition is given by $x_i = V_i y_i$ where $y_i = H_i^{-1}\|r_0\|_2 e_1$, assuming H_i is nonsingular [13, Chapter 6].

Since FOM does not minimize the residual or the error for general A, FOM is not considered optimal. However, it has been shown that when GMRES converges, FOM converges at nearly the same rate. More precisely, letting $\hat{r}_i$ denote the ith residual in the FOM iteration and r_i the ith residual in the GMRES iteration,

$$\|\hat{r}_i\|_2 = \frac{\|r_i\|_2}{\sqrt{1 - (\|r_i\|_2/\|r_{i-1}\|_2)^2}} \tag{2.4}$$

in exact arithmetic [2]. Hence, the norm of the FOM residual is determined by the convergence of GMRES. In every iteration where GMRES makes reasonable progress, $\|r_i\|_2/\|r_{i-1}\|_2$ is small and FOM is roughly equivalent to GMRES. The FOM residual only becomes large when GMRES (nearly) stagnates.

Algorithm 5 Full Orthogonalization Method (FOM) Algorithm [13]

1: Compute $r_0 = b - Ax_0$, $\beta = \|r_0\|_2$, and $v_1 = r_0/\beta$.
2: **for** $m = 1, 2, ...$ until convergence **do**
3: Define the $m \times m$ matrix $H_m = \{h_{i,j}\}_{i,j=1,...,m}$; Set $H_m = 0$.
4: **for** $j = 1, 2, ..., m$ **do**
5: $w_j = Av_j$.
6: **for** $i = 1, ..., j$ **do**
7: $h_{i,j} = v_i^* w_j$
8: $w_j = w_j - h_{i,j} v_i$
9: **end for**
10: $h_{j+1,j} = \|w_j\|_2$. If $h_{j+1,j} = 0$, perform 13-14 and stop.
11: $v_{j+1} = w_j / h_{j+1,j}$.
12: **end for**
13: $y_m = H_m^{-1}(\beta e_1)$
14: $x_m = x_0 + V_m y_m$.
15: **end for**

At step m, the residuals r_m and $\hat{r}_m$ for GMRES and FOM, respectively, satisfy

$$\begin{array}{lll} \text{GMRES} & r_m = r_0 - Q_m y_m \perp Q_m & \iff \quad y_m = Q_m^* r_0 \\ \text{FOM} & \hat{r}_m = r_0 - Q_m \hat{y}_m \perp \hat{Q}_m & \iff \quad \hat{y}_m = (\hat{Q}_m^* Q_m)^{-1} \hat{Q}_m^* r_0 \end{array} \tag{2.5}$$

where Q_m and $\hat{Q}_m$ are any matrices whose columns form orthonormal bases for $A\mathcal{K}^m(A; r_0)$ and $\mathcal{K}^m(A; r_0)$ respectively. We assume that $(\hat{Q}_m^* Q_m)^{-1}$ exists.

2.5 BiCG

If the matrix A is Hermitian positive-definite, then $H_{i+1,i}$ (as defined above) is tridiagonal, giving us a 3-term recurrence rather than a full orthogonalization for computing new approximations. Using the projection space $\mathcal{K}^i(A; r_0)$, this 3-term recurrence results in the conjugate gradient (CG) method. In the case that A is Hermitian positive-definite, CG and FOM are equivalent. For non-Hermitian matrices, however, we cannot maintain both orthogonality of the columns of V_i and tridiagonality of $H_{i+1,i}$ except in very special cases [6]. In BiCG, we sacrifice the orthogonality of the column vectors v_j to find a tridiagonal matrix $T_{i+1,i}$ such that the Lanczos relations

$$AV_i = V_{i+1}T_{i+1,i} \tag{2.6}$$

hold. BiCG utilizes the nonsymmetric Lanczos iteration to build V_i and a matrix W_i, whose columns form a basis for $\mathcal{K}^i(A^*; \tilde{r}_0)$, where $\tilde{r}_0$ is chosen such that $\tilde{r}_0^* r_0 \neq 0$, such that the columns of V_i and W_i are biorthogonal and analogous Lanczos relations hold for W_i:

$$A^*W_i = W_{i+1}S_{i+1,i} \tag{2.7}$$

where $S_{i+1,i}$ is tridiagonal [15, Lecture 39].

Essentially, BiCG solves the system $Ax = r_0$ by solving both $Ax = r_0$ and $A^*\tilde{x} = \tilde{r}_0$ simultaneously (we may assume $x_0 = 0$) [13, Chapter 7]. At each step, BiCG computes $x_i \in \mathcal{K}^i(A; r_0)$ such that $W_i^*(r_0 - Ax_i) = 0$, i.e. $r_i \perp W_i$. The solution to this equation is given by $x_i = V_i y$ where y satisfies $T_{i,i}y = \|r_0\|_2 e_1$. This r_i becomes the next basis vector for $\mathcal{K}^{i+1}(A; r_0)$, so $V_{i+1} = \begin{bmatrix} V_i & r_i \end{bmatrix}$. Simultaneously, BiCG finds $\tilde{x}_i \in \mathcal{K}^i(A^*; \tilde{r}_0)$ such that $\tilde{r}_i \perp V_i$ to build the next basis vector for $\mathcal{K}^{i+1}(A^*; \tilde{r}_0)$, so $W_{i+1} = \begin{bmatrix} W_i & \tilde{r}_i \end{bmatrix}$. As a result, the jth column of V_i and W_i are r_j and $\tilde{r}_j$, respectively, and the sequences $\{r_j\}$ and $\{\tilde{r}_j\}$ are biorthogonal.

This projection allows a more abstract definition of BiCG. At step m, the residual r_m of BiCG satisfies

$$r_m = r_0 - Q_m y_m \perp \tilde{Q}_m \quad \Longleftrightarrow \quad y_m = (\tilde{Q}_m^* Q_m)^{-1} \tilde{Q}_m^* r_0 \tag{2.8}$$

where Q_m is as in (2.5) and $\tilde{Q}_m$ is any matrix whose columns form an orthonormal basis for $\mathcal{K}^m(A^*; \tilde{r}_0)$.

Since the columns of W_i and V_i are biorthogonal, we may write

$$W_i^* V_i = \Delta \tag{2.9}$$

for some diagonal matrix $\Delta = \text{diag}(\delta_i)$. BiCG breaks down if $\delta_i = 0$ since, in such cases, the recurrence coefficient α_i in Algorithm 6 cannot be computed. At such a step, $\tilde{Q}_m^* Q_m$ is singular and the projection in (2.8) is undefined. Such breakdowns can be avoided by using a lookahead strategy or a restart of the algorithm [17, Chapter 7].

In general, BiCG does not minimize the error or the residual. Hence, BiCG is not an optimal method. In fact, it has been shown that a special choice of the left starting vector $\tilde{r}_0$ can cause BiCG to follow any 3-term recurrence for up to $\frac{n}{2} - 1$ iterations [8]. Thus, the convergence behavior of BiCG can be almost arbitrarily controlled. In practice, however, when the left starting vector is chosen randomly, the method tends to converge well regardless of this choice. This paper seeks to explain why BiCG converges so well in practice despite the potential for it to behave poorly.

Algorithm 6 Biconjugate Gradient (BiCG) Algorithm [4]

Compute $r_0 := b - Ax_0$. Choose $\tilde{r}_0$ such that $\tilde{r}_0^* r_0 \neq 0$.
for $j = 0, 1, 2, ...$ until convergence **do**
    $\delta_j = \tilde{r}_j^* r_j$
    $\alpha_j = \tilde{r}_j^* A r_j / \delta_j$
    $\beta_{j-1} = \gamma_{j-1}(\delta_j / \delta_{j-1})$
    $\gamma_j = -\alpha_j - \beta_{j-1}$
    $r_{j+1} = \gamma_j^{-1}(Ar_j - \alpha_j r_j - \beta_{j-1} r_{j-1})$
    $\tilde{r}_{j+1} = \bar{\gamma}_j^{-1}(A^* \tilde{r}_j - \bar{\alpha}_j \tilde{r}_j - \bar{\beta}_{j-1} \tilde{r}_{j-1})$
    $x_{j+1} = -(\alpha_j / \gamma_j) x_j - (\beta_{j-1} / \gamma_j) x_{j-1} - \gamma_j^{-1} r_j$
end for

Chapter 3

BiCG in Eigenspace – Diagonalizable Case

If A is diagonalizable, then we can write $A = X\Lambda X^{-1}$ and $A^* = Y\bar{\Lambda}Y^{-1}$, where the columns of X are unit right eigenvectors of A, the columns of Y are unit left eigenvectors of A, and Λ is the diagonal matrix with the corresponding eigenvalues of A on the diagonal. Furthermore, X and Y can be chosen such that $Y^*X = D$, where D is diagonal with positive real coefficients [4]. Then D induces an inner product given by $\langle x, y\rangle_D = y^*Dx$.

3.1 BiCG in Eigenvector Basis

Letting $\rho_m = X^{-1}r_m$ and $\tilde{\rho}_m = Y^{-1}\tilde{r}_m$, we can express BiCG with respect to this eigenvector basis using the D-inner product:

$$\begin{aligned}
\rho_{j+1} &= X^{-1}r_{j+1} \\
&= X^{-1}\gamma_j^{-1}(Ar_j - \alpha_j r_j - \beta_{j-1}r_{j-1}) \\
&= \gamma_j^{-1}(X^{-1}Ar_j - \alpha_j X^{-1}r_j - \beta_{j-1}X^{-1}r_{j-1}) \\
&= \gamma_j^{-1}(\Lambda\rho_j - \alpha_j\rho_j - \beta_{j-1}\rho_{j-1}) \\
\tilde{\rho}_{j+1} &= Y^{-1}\tilde{r}_{j+1} \\
&= Y^{-1}\bar{\gamma}_j^{-1}(A^*\tilde{r}_j - \bar{\alpha}_j\tilde{r}_j - \bar{\beta}_{j-1}\tilde{r}_{j-1}) \\
&= \bar{\gamma}_j^{-1}(Y^{-1}A^*\tilde{r}_j - \bar{\alpha}_j\tilde{r}_j - \bar{\beta}_{j-1}\tilde{r}_{j-1}) \\
&= \bar{\gamma}_j^{-1}(\bar{\Lambda}\tilde{\rho}_j - \bar{\alpha}_j\tilde{\rho}_j - \bar{\beta}_{j-1}\tilde{\rho}_{j-1})
\end{aligned}$$

where

$$\delta_j = \tilde{r}_j^* r_j = \tilde{\rho}_j^* D\rho_j \tag{3.1}$$

and

$$\alpha_j = \tilde{r}_j^* Ar_j/\delta_j = \tilde{\rho}_j^* D\Lambda\rho_j/(\tilde{\rho}_j^* D\rho_j). \tag{3.2}$$

This gives us the same 3-term recurrences that we saw in Algorithm 6. Hence, solving $Ax = r_0$ with BiCG in the standard basis is equivalent to solving $\Lambda\xi = \rho_0$ in the eigenvector

basis using the D-inner product. This implicit BiCG in the eigenvector basis will help explain some of the properties of BiCG. The residual ρ_m in the eigenvector basis satisfies

$$\rho_m = \rho_0 - Q_m z_m \perp_D \tilde{Q}_m \quad \iff \quad z_m = (\tilde{Q}_m^* D Q_m)^{-1} \tilde{Q}_m^* D \rho_0$$

where the columns of Q_m form an orthonormal basis for $\Lambda\mathcal{K}^m(\Lambda; \rho_0)$ and the columns of $\tilde{Q}_m$ form an orthonormal basis for $\mathcal{K}^m(\bar{\Lambda}; \tilde{\rho}_0)$.

To see the effect of D, it is useful to consider the same problem in the Euclidean inner product with a special starting vector $\tilde{\rho}_0 = DY^{-1}\tilde{r}_0$.

Theorem 3.1. *In the eigenvector basis, solving $\Lambda\xi = \rho_0$ with BiCG in the D-inner product is equivalent to solving $\Lambda\xi = \rho_0$ with BiCG in the Euclidean inner product, but with the left starting vector $\tilde{\rho}_0 = DY^{-1}\tilde{r}_0$* [4].

Proof. Let $\hat{\rho}_0 = Y^{-1}\tilde{r}_0$ be the left starting vector for BiCG in the eigenvector basis and let $\hat{\rho}_m$ denote the mth residual. Note that the $\hat{\rho}_m = p_m(\bar{\Lambda})\hat{\rho}_0$ for some polynomial p_m. Then since D commutes with $\bar{\Lambda}$,

$$\begin{aligned} \hat{\rho}_m^* D &= (p_m(\bar{\Lambda})\hat{\rho}_0)^* D \\ &= \hat{\rho}_0^* p_m(\bar{\Lambda}) D \\ &= \hat{\rho}_0^* D p_m(\bar{\Lambda}) \\ &= (p_m(\bar{\Lambda}) D \hat{\rho}_0)^*. \end{aligned}$$

By replacing $\hat{\rho}_0$ by $\tilde{\rho}_0 = D\hat{\rho}_0$, we get that $\tilde{\rho}_m = p_m(\bar{\Lambda})\tilde{\rho}_0 = p_m(\bar{\Lambda})D\hat{\rho}_0 = Dp_m(\bar{\Lambda})\hat{\rho}_0 = D\hat{\rho}_m$. Then $\delta_m = \tilde{\rho}_m^*\rho_m$ and $\beta_m = \tilde{\rho}_m^*\Lambda\rho_m / \tilde{\rho}_m^*\rho_m$. So the relations (3.1) and (3.2) are now in the Euclidean inner product. Hence, in the eigenvector basis, BiCG with the D-inner product is equivalent to BiCG in the Euclidean inner product with a special starting vector $\tilde{\rho}_0 = DY^{-1}\tilde{r}_0$. □

Thus, a small coefficient in D is equivalent to the damping of the corresponding left eigenvector component in $\tilde{\rho}_0$. This is analyzed further in the following sections.

3.2 How BiCG Approximates FOM

Note that Λ and $\bar{\Lambda}$ have the same eigenvalues (up to complex conjugation), and the right and left eigenvectors are the same and form orthogonal bases. How quickly an eigenvector converges in $\mathcal{K}^m(\Lambda; \rho_0)$ or $\mathcal{K}^m(\bar{\Lambda}; \tilde{\rho}_0)$ depends mainly on the position of the corresponding eigenvalue in the spectrum. Thus the left and right eigenvectors of Λ corresponding to the same eigenvalue should converge in $\mathcal{K}^m(\bar{\Lambda}; \tilde{\rho}_0)$ and $\mathcal{K}^m(\Lambda; \rho_0)$ (respectively) at roughly the same rate (see Figure 3.1). Since the left and right eigenvectors are the same, $\mathcal{K}^m(\bar{\Lambda}; \tilde{\rho}_0)$ and $\mathcal{K}^m(\Lambda; \rho_0)$ approximate the same vectors, and hence these spaces converge to each other. Thus, as long as the components in D are not too small, BiCG in the eigenvector basis approximates a FOM iteration [4]. The effect of small coefficients in D is demonstrated in Figures 3.1 and 3.2. Some consequences of this result are discussed in more detail in Chapter 5.

3.3 Convergence Bounds

Viewing BiCG in the eigenvector bases, convergence properties become more clear. For instance, we see that as the left eigenvectors converge in $\mathcal{K}^m(A^*, \tilde{r}_0)$, the corresponding right eigenvectors are removed from the residual r_m. Similarly, as the right eigenvectors converge in $\mathcal{K}^m(A; r_0)$, the corresponding left eigenvectors are removed from $\tilde{r}_m$ [3]. This is demonstrated in the following two theorems.

Theorem 3.2. *Let x_k, y_k be the kth columns of X and Y, respectively (so x_k is a right eigenvector and y_k is a left eigenvector). If $y_k \in \mathcal{K}^m(A^*; \tilde{r}_0)$, then the BiCG residual r_m has no component in the direction x_k, i.e. $(\rho_m)_k = 0$. Similarly, if $x_k \in \mathcal{K}^m(A; r_0)$ then $\tilde{r}_m$ has no component in the direction y_k.*

Proof. Suppose $y_k \in \mathcal{K}^m(A^*; \tilde{r}_0)$. Then since $y_k \perp x_j$ for $j \neq k$,

$$0 = y_k^* r_m = y_k^* X \rho_m = y_k^* x_k (\rho_m)_k$$

So $(\rho_m)_k = 0$ and hence $r_m = X\rho_m$ has no component in x_k. If we suppose $x_k \in \mathcal{K}^m(A; r_0)$, then the proof that $(\tilde{\rho}_m)_k = 0$ is analogous. □

We now present a more general result in the case that y_k is almost (but not fully) contained in $\mathcal{K}^m(A^*; \tilde{r}_0)$.

Theorem 3.3. *Suppose $y_k = (1-\epsilon^2)^{1/2} v_1 + \epsilon v_2$ where $v_1 \in \mathcal{K}^m(A^*; \tilde{r}_0)$, $v_2 \perp \mathcal{K}^m(A^*; \tilde{r}_0)$, and $\|v_1\| = \|v_2\| = 1$. Then $|(\rho_m)_k| \leq \frac{\epsilon \|r_m\|}{d_k}$.*

Proof.

$$y_k^* r_m = (1-\epsilon^2)^{1/2} v_1^* r_m + \epsilon v_2^* r_m = \epsilon v_2^* r_m$$

and

$$y_k^* r_m = y_k^* X \rho_m = y_k^* x_k (\rho_m)_k.$$

So $y_k^* x_k (\rho_m)_k = \epsilon v_2^* r_m$, and $(\rho_m)_k = \frac{\epsilon v_2^* r_m}{y_k^* x_k}$. Hence,

$$|(\rho_m)_k| = \frac{\epsilon |v_2^* r_m|}{y_k^* x_k} = \frac{\epsilon}{d_k} |v_2^* r_m| \leq \frac{\epsilon}{d_k} \|v_2\| \|r_m\| = \frac{\epsilon \|r_m\|}{d_k}.$$

□

Hence if d_k is not too small relative to ϵ, then x_k is nearly removed from r_m. However, if some of the eigenvalues are ill-conditioned, then the corresponding coefficients in D will be small. As a result, r_m and $\tilde{r}_m$ may not decrease in the directions of the corresponding eigenvectors. In the next chapter, we will state analogous results for generalized eigenvectors in the case that A is nondiagonalizable.

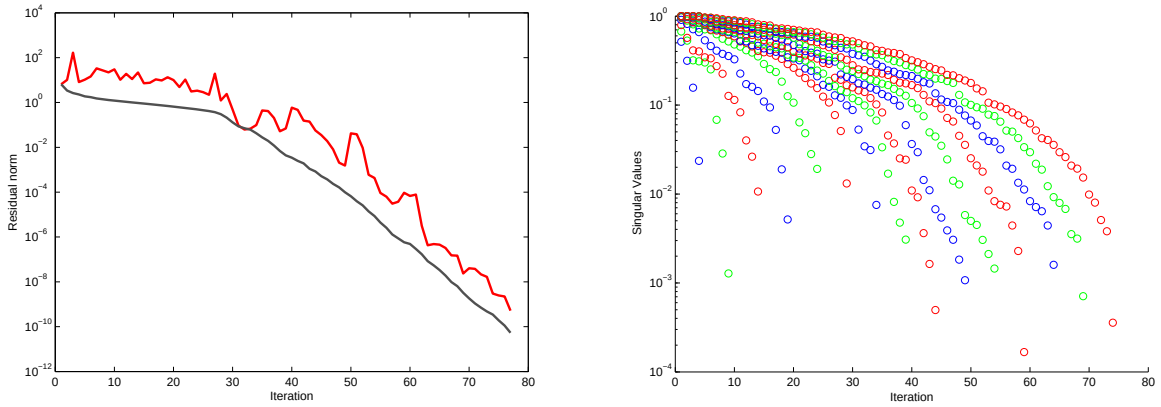

Figure 3.1: BiCG and FOM applied to a convection-diffusion problem with a randomly chosen left starting vector. Left: BiCG residual (red) and FOM residual (grey). Right: Singular values of $\tilde{Q}_m^* \hat{Q}_m$. These are the cosines of the principal angles between projection spaces for FOM and BiCG (discussed further in Chapter 5). Note the singular values are not systematically small, so the principal angles are not close to $\pi/2$ and hence the projection spaces are close.

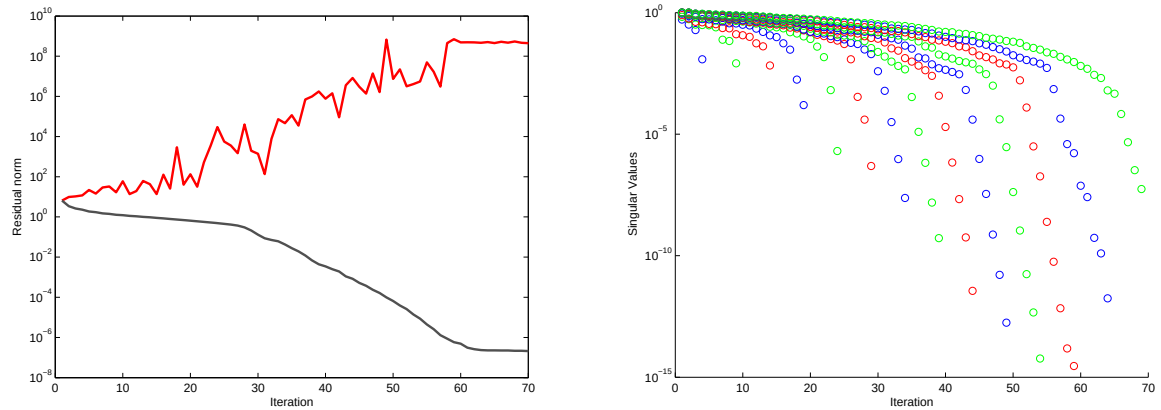

Figure 3.2: BiCG and FOM applied to the same problem as in Figure 3.1, but with the 25 absolute largest eigenvalue components removed from the left starting vector $\tilde{r}_0$. Left: BiCG residual (red) and FOM residual (grey). Right: Singular values of $\tilde{Q}_m^* \hat{Q}_m$. Note that these are much smaller than in the previous figure, hence the principal angles are closer to $\pi/2$.

Chapter 4

BiCG in Eigenspace – Nondiagonalizable Case

In the previous chapter we saw that when A is diagonalizable, there is an equivalent BiCG in the eigenvector basis. We now show that similar results hold when A is not diagonalizable. To do so, we use the Jordan decomposition with bases X and Y of right and left generalized eigenvectors, respectively, such that $Y^*X = D$ where D is diagonal with positive real coefficients. Throughout this section, chains of right and left generalized eigenvectors corresponding to a Jordan block $\mathcal{J}$ are indexed by $x_i = (A - \lambda I)^{m-i}x_m$ and $y_i^* = y_1^*(A - \lambda I)^{i-1}$, respectively, where x_m and y_1^* are generalized eigenvectors of degree m.

4.1 Generalized Eigenvectors

First, we establish some basic properties of chains of generalized eigenvectors.

Lemma 4.1. *Let $\{x_i\}, \{y_i\}$ be chains of right and left generalized eigenvectors, respectively, corresponding to a Jordan block $\mathcal{J}$ of dimension m with eigenvalue λ. Then the following properties hold:*

1. *for $1 \leq i, j < m$, $y_i^*x_j = y_{i+1}^*x_{j+1}$;*
2. *for $1 \leq j < i \leq m$, $y_i^*x_j = 0$*

(i.e. if $Y = \begin{bmatrix} y_1 & \dots & y_m \end{bmatrix}$ and $X = \begin{bmatrix} x_1 & \dots & x_m \end{bmatrix}$, then Y^*X is upper triangular and constant along diagonals).

Proof. Suppose $1 \leq i, j < m$ where m denotes the dimension of the Jordan block. Then by definition, $y_i^*x_j = y_1^*(A - \lambda I)^{i-1}(A - \lambda I)^{m-j}x_m = y_1^*(A - \lambda I)^{m+i-j-1}x_m$ and $y_{i+1}^*x_{j+1} = y_1^*(A - \lambda I)^{i}(A - \lambda I)^{m-(j+1)}x_m = y_1^*(A - \lambda I)^{m+i-j-1}x_m$. So $y_i^*x_j = y_{i+1}^*x_{j+1}$.

Let $1 \leq j < i \leq m$. Then since $i - j \geq 1$, we have $m + i - j - 1 \geq m$. So $(A - \lambda I)^{m+i-j-1}x_m = 0$ and $y_i^*x_j = y_1^*(A - \lambda I)^{m+i-j-1}x_m = 0$. □

Using these properties, we can build biorthogonal bases of right and left generalized eigenvectors corresponding to $\mathcal{J}$.

Theorem 4.2. *Let $\mathcal{J}$ be a Jordan block of dimension m with eigenvalue λ. Then we can choose a basis of left generalized eigenvectors $Y = \begin{bmatrix} y_1 & ... & y_m \end{bmatrix}$ and a basis of right generalized eigenvectors $X = \begin{bmatrix} x_1 & ... & x_m \end{bmatrix}$ corresponding to $\mathcal{J}$ such that the $\{y_i\}$ and $\{x_i\}$ form chains of generalized eigenvectors and $Y^*X = \gamma I$ for some real, positive constant γ.*

Proof. Let $\{x_1, x_2, ..., x_m\}$ and $\{y_1^*, y_2^*, ..., y_m^*\}$ be chains of generalized eigenvectors, where $x_i = (A - \lambda I)^{m-i} x_m$ and $y_j^* = y_1^*(A - \lambda I)^{j-1}$. Write $\tilde{\gamma} = y_1^* x_1$. We want to build a left chain $\{\tilde{y}_1^*, \tilde{y}_2^*, ..., \tilde{y}_m^*\}$ biorthogonal to $\{x_1, x_2, ..., x_m\}$.

Let $\tilde{y}_1^* = \sum_{i=1}^m \alpha_i y_i^*$ with $\alpha_1 \neq 0$. Then $\tilde{y}_1$ is a left generalized eigenvector of degree m and we can define the α_i recursively such that $\tilde{y}_1^*$ generates a chain biorthogonal to $\{x_1, x_2, ..., x_m\}$:

$$\begin{aligned} \tilde{y}_1^* x_1 &= \sum_{i=1}^m \alpha_i y_i^* x_1 \\ &= \alpha_1 y_1^* x_1 = \alpha_1 \tilde{\gamma} \neq 0 \\ \tilde{y}_1^* x_2 &= \sum_{i=1}^m \alpha_i y_i^* x_2 \\ &= \alpha_1 y_1^* x_2 + \alpha_2 y_2^* x_2 \\ &= \alpha_1 y_1^* x_2 + \alpha_2 \tilde{\gamma}. \end{aligned}$$

So we can choose $\alpha_2 = -\frac{\alpha_1}{\tilde{\gamma}} y_1^* x_2$ to make $\tilde{y}_1 \perp x_2$.

$$\begin{aligned} \tilde{y}_1^* x_3 &= \alpha_1 y_1^* x_3 + \alpha_2 y_2^* x_3 + \alpha_3 y_3^* x_3 \\ &= \alpha_1 y_1^* x_3 + \alpha_2 y_2^* x_3 + \alpha_3 \tilde{\gamma}. \end{aligned}$$

So we can choose $\alpha_3 = -\frac{1}{\tilde{\gamma}}(\alpha_1 y_1^* x_3 + \alpha_2 y_2^* x_3)$ to make $\tilde{y}_1 \perp x_3$. Similarly, for any $k > 1$ we get

$$\begin{aligned} \tilde{y}_1^* x_k &= \sum_{i=1}^k \alpha_i y_i^* x_k \\ &= \sum_{i=1}^{k-1} \alpha_i y_i^* x_k + \alpha_k \tilde{\gamma}. \end{aligned}$$

So we can choose $\alpha_k = -\frac{1}{\gamma}(\sum_{i=1}^{k-1} \alpha_i y_i^* x_k)$ to make $\tilde{y}_1 \perp x_k$.

Then $\tilde{y}_1 \perp x_2, ..., x_m$. Consider the chain $\{\tilde{y}_1^*, \tilde{y}_2^*, ..., \tilde{y}_m^*\}$ generated by $\tilde{y}_1$, and let $Y = \begin{bmatrix} \tilde{y}_1 & ... & \tilde{y}_m \end{bmatrix}$ and $X = \begin{bmatrix} x_1 & ... & x_m \end{bmatrix}$. Since $\{\tilde{y}_1^*, ..., \tilde{y}_m^*\}$ is a chain, Y^*X is upper triangular and constant along diagonals. Since $\tilde{y}_1^* \perp x_2, ..., x_m$, the first row of Y^*X is zero except for the first entry. So Y^*X is diagonal. So the chains $\{\tilde{y}_1^*, ..., \tilde{y}_m^*\}$ and $\{x_1, ..., x_m\}$ are biorthogonal. In particular, $Y^*X = \gamma I$ where $\gamma = \tilde{y}_1^* x_1 = \alpha_1 \tilde{\gamma}$. We can choose γ to be real and positive by scaling α_1. □

We need to generalize this result to the case of multiple Jordan blocks. To do so, it is sufficient to show that the right generalized eigenvectors of one Jordan block can be chosen to be orthogonal to the left generalized eigenvectors of any other Jordan block.

Lemma 4.3. *Let $\mathcal{J}_1$ and $\mathcal{J}_2$ be distinct Jordan blocks with corresponding eigenvalues λ_1 and λ_2 respectively (λ_1 and λ_2 need not be distinct). Let y be a left generalized eigenvector of A corresponding to $\mathcal{J}_1$ and x a right generalized eigenvector of A corresponding to $\mathcal{J}_2$. Then $y \perp x$.*

Proof. We may assume that, in the case of repeated eigenvalues, we have already chosen the left and right eigenvectors to be biorthogonal. If there are no repeated eigenvalues, then biorthogonality of the left and right eigenvectors is automatically satisfied. So the invariant subspaces corresponding to $\mathcal{J}_1$ and $\mathcal{J}_2$ are already determined.

Let y be a left generalized eigenvector corresponding to $\mathcal{J}_1$ and x a right generalized eigenvector corresponding to $\mathcal{J}_2$. Let X be a matrix whose columns form an orthonormal basis for the right invariant subspace associated with $\mathcal{J}_2$, so $x \in R(X)$. Then we can find Z such that $\begin{bmatrix} X & Z \end{bmatrix}$ is unitary and as a result

$$\begin{bmatrix} X & Z \end{bmatrix}^* A \begin{bmatrix} X & Z \end{bmatrix} = \begin{bmatrix} L & H \\ 0 & M \end{bmatrix},$$

where L and M are upper triangular. Note that the eigenvalues of M are the eigenvalues of A (except for λ_2), with the same multiplicities, so we may assume M contains all Jordan blocks of A other than $\mathcal{J}_2$. Further, if v is a left generalized eigenvector of M then Zv is a left generalized eigenvector of A corresponding to some Jordan block other than $\mathcal{J}_2$. Conversely, if u is a left generalized eigenvector of A not corresponding to $\mathcal{J}_2$, then $u = Zv$ for some left generalized eigenvector v of M. So $y = Zv$ for some v, so $y \in R(Z) \perp R(X)$. Then since $x \in R(X)$, we have $y \perp x$. □

We can now prove the general result:

Theorem 4.4. *Let A be any square matrix. Then there exist bases $X = \begin{bmatrix} x_1 & ... & x_n \end{bmatrix}$ and $Y = \begin{bmatrix} y_1 & ... & y_n \end{bmatrix}$ such that $A = X\mathcal{J}X^{-1}$ and $A^* = Y\mathcal{J}^*Y^{-1}$ are Jordan decompositions and $Y^*X = D$ where D is diagonal. Further, D is constant on the blocks corresponding to Jordan blocks of $\mathcal{J}$.*

Proof. Suppose $\mathcal{J}$ consists of k Jordan blocks $\mathcal{J}_1, ..., \mathcal{J}_k$. For each block $\mathcal{J}_i$, construct bases X_i and Y_i of right and left generalized eigenvectors as in Theorem 4.2. Then $Y_i^*X_i = d_iI$ for some real constant d_i and by the previous lemma, $Y_i^*X_j = 0$ for $i \neq j$. Let $Y = \begin{bmatrix} Y_1 & ... & Y_k \end{bmatrix}$ and $X = \begin{bmatrix} X_1 & ... & X_k \end{bmatrix}$. Then $Y^*X = D$ where D is diagonal and constant on the blocks corresponding to distinct Jordan blocks. □

In particular, we can scale the y_i and x_i so that D is positive definite, so D induces an inner product.

4.2 BiCG in Generalized Eigenvector Basis

Let $A = X\mathcal{J}X^{-1}$ and $A^* = Y\mathcal{J}^*Y^{-1}$ be Jordan decompositions satisfying the conclusions of Theorem 4.4. Let $r_i = X\rho_i$, $\tilde{r}_i = Y\tilde{\rho}_i$, and $D = Y^*X$. We may assume D is real, diagonal,

and positive definite. Note that the block constant structure of D means that D and $\mathcal{J}$ commute. Then

$$\begin{aligned} r_m &= \gamma_{m-1}^{-1}(Ar_{m-1} - \alpha_{m-1}r_{m-1} - \beta_{m-2}r_{m-2}) \\ &= \gamma_{m-1}^{-1}(X\mathcal{J}\rho_{m-1} - \alpha_{m-1}X\rho_{m-1} - \beta_{m-2}X\rho_{m-2}) \\ &= X\gamma_{m-1}^{-1}(\mathcal{J}\rho_{m-1} - \alpha_{m-1}\rho_{m-1} - \beta_{m-2}\rho_{m-2}). \end{aligned}$$

So $\rho_m = X^{-1}r_m = \gamma_{m-1}^{-1}(\mathcal{J}\rho_{m-1} - \alpha_{m-1}\rho_{m-1} - \beta_{m-2}\rho_{m-2})$. Similarly,

$$\begin{aligned} \tilde{r}_m &= \bar{\gamma}_{m-1}^{-1}(A^*\tilde{r}_{m-1} - \bar{\alpha}_{m-1}\tilde{r}_{m-1} - \bar{\beta}_{m-2}\tilde{r}_{m-2}) \\ &= \bar{\gamma}_{m-1}^{-1}(Y\mathcal{J}^*\tilde{\rho}_{m-1} - \bar{\alpha}_{m-1}Y\tilde{\rho}_{m-1} - \bar{\beta}_{m-2}Y\tilde{\rho}_{m-2}) \\ &= Y\bar{\gamma}_{m-1}^{-1}(\mathcal{J}^*\tilde{\rho}_{m-1} - \bar{\alpha}_{m-1}\tilde{\rho}_{m-1} - \bar{\beta}_{m-2}\tilde{\rho}_{m-2}). \end{aligned}$$

So $\tilde{\rho}_m = Y^{-1}\tilde{r}_m = \bar{\gamma}_{m-1}^{-1}(\mathcal{J}^*\tilde{\rho}_{m-1} - \bar{\alpha}_{m-1}\tilde{\rho}_{m-1} - \bar{\beta}_{m-2}\tilde{\rho}_{m-2})$ where

$$\begin{aligned} \alpha_{m-1} &= \tilde{r}_{m-1}^* Ar_{m-1}/\delta_{m-1} \\ &= (\tilde{\rho}_{m-1}Y^*X\mathcal{J}X^{-1}X\rho_{m-1})/(\tilde{\rho}_{m-1}^*D\rho_{m-1}) \\ &= (\tilde{\rho}_{m-1}D\mathcal{J}\rho_{m-1})/(\tilde{\rho}_{m-1}^*D\rho_{m-1}) \\ \delta_{m-1} &= \tilde{r}_{m-1}^* r_{m-1} = \tilde{\rho}_{m-1}Y^*X\rho_{m-1} \\ &= \tilde{\rho}_{m-1}^*D\rho_{m-1}. \end{aligned}$$

Thus the standard BiCG is equivalent to solving $\mathcal{J}\zeta = \rho_0$ and $\mathcal{J}^*\tilde{\zeta} = \tilde{\rho}_0$ in the bases of generalized eigenvectors with the D-inner product. By the same argument as in Theorem 3.1, this is also equivalent to solving $\mathcal{J}\zeta = \rho_0$ and $\mathcal{J}^*\tilde{\zeta} = \tilde{\rho}_0$ in the bases of generalized eigenvectors using the standard inner product with starting vector $\tilde{\rho}_0 = DY^{-1}\tilde{r}_0$.

4.3 Convergence Bounds

We can now state results analogous to those in Section 3.3.

Theorem 4.5. *Let x_k, y_k be as in Theorem 4.4. If $y_k \in \text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$, then x_k is removed from r_m. Similarly, if $x_k \in \text{span}(r_0, ..., r_{m-1})$, then y_k is removed from $\tilde{r}_m$.*

Proof. Suppose $y_k \in \text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$. Since $y_k \perp x_j$ for $j \neq k$,

$$0 = y_k^* r_m = y_k^* X\rho_m = y_k^* x_k(\rho_m)_k$$

Then since $y_k^* x_k = d_k \neq 0$, we must have $(\rho_m)_k = 0$. So $r_m = X\rho_m$ has no component in x_k. If we suppose $x_k \in \text{span}(r_0, ..., r_{m-1})$, then the proof that y_k is removed from $\tilde{r}_m$ is analogous. □

A more general result holds when y_k is almost (but not fully) contained in $\text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$.

Theorem 4.6. *Suppose $y_k = (1-\epsilon^2)^{1/2}v_1 + \epsilon v_2$ where $v_1 \in \text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$, $v_2 \perp \text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$, and $\|v_1\| = \|v_2\| = \|y_k\|$. Then $|(\rho_m)_k \leq \frac{\epsilon\|r_m\|}{d_k}\|y_k\|$.*

Proof. Since $r_m \perp \text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$, $v_1^* r_m = 0$. Then

$$y_k^* r_m = (1-\epsilon^2)^{1/2} v_1^* r_m + \epsilon v_2^* r_m = \epsilon v_2^* r_m.$$

Furthermore, since $y_k \perp x_j$ for $j \neq k$,

$$y_k^* r_m = y_k^* X \rho_m = y_k^* x_k (\rho_m)_k.$$

Thus $y_k^* x_k (\rho_m)_k = \epsilon v_2^* r_m$, and $(\rho_m)_k = \frac{\epsilon v_2^* r_m}{y_k^* x_k}$. Then

$$|(\rho_m)_k| = \frac{\epsilon |v_2^* r_m|}{y_k^* x_k} = \frac{\epsilon}{d_k} |v_2^* r_m| \leq \frac{\epsilon}{d_k} \|v_2\| \|r_m\| = \frac{\epsilon \|r_m\|}{d_k} \|y_k\|.$$

□

So if ϵ is small and d_k is not too close to zero, then x_k is nearly removed from r_m.

We now generalize these results to invariant subspaces. Let $\hat{Y}_k = [y_i]_{i \in B}$, $B \subseteq \{1, 2, ..., n\}$, be an $n \times k$ matrix whose columns form a basis for a left-invariant subspace of A. Let X be the matrix of right generalized eigenvectors of A.

Theorem 4.7. *If* $\text{range}(\hat{Y}_k) \subseteq \text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$*, then* $X_k = [x_i]_{i \in B}$ *is removed from* r_m*.*

Proof. Let $\hat{D} = \text{diag}(d_i)_{i \in B}$ and $\hat{\rho} = [(\rho_m)_i]_{i \in B}$. Then since $y_i \perp x_j$ for $i \neq j$,

$$0 = \hat{Y}_k^* r_m = \hat{Y}_k^* X \rho_m = \hat{D} \hat{\rho}.$$

Then $(\rho_m)_i = 0$ for all $i \in B$, and the corresponding subspace X_k is removed from r_m. □

We now suppose $\hat{Y}_k$ is not fully contained in $\text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$, but that $\hat{Y}_k$ is contained up to some components of small length σ_i.

Theorem 4.8. *Let* $Y_k = \hat{Y}_k Z$ *be an orthonormal matrix whose columns span the same subspace as* $\hat{Y}_k$*. Let* V *be a matrix whose columns form an orthonormal basis for* $\text{span}(\tilde{r}_0, ..., \tilde{r}_{m-1})$*, and let* $V^* Y_k = \Phi \Omega \Psi^*$ *be an SVD. Define* $\sigma_1 \geq 0$ *such that* $\omega_i = (1 - \sigma_i^2)^{1/2}$*. Then* $\|Z^* \hat{D} \hat{\rho}\| \leq \sigma \|r_m\|$ *where* $\hat{D} = \text{diag}(d_i)_{i \in B}$*,* $\hat{\rho} = [(\rho_m)_i]_{i \in B}$*, and* $\sigma = \max_i(\sigma_i)$*.*

Proof. Rewriting $V^* Y_k = \Phi \Omega \Psi^*$, we get $\Phi^* V^* Y_k \Psi = (V\Phi)^* (Y_k \Psi) = \Omega$. Then

$$Y_k \Psi = V \Phi \Omega + W \Sigma.$$

where W is $n \times k$ with orthonormal columns, $\Sigma = \text{diag}(\sigma_i)$, and $W^* V = 0$. Write $\tilde{Y}_k = Y_k \Psi, V_1 = V\Phi$ and $V_2 = W$. Then

$$\begin{aligned} \tilde{Y}_k^* r_m &= (V_1 \Omega + V_2 \Sigma)^* r_m \\ &= \Omega^* V_1^* r_m + \Sigma^* V_2^* r_m \\ &= \Sigma^* V_2^* r_m \end{aligned}$$

and

$$\begin{aligned}\tilde{Y}_k^* r_m &= (\hat{Y}_k Z\Psi)^* r_m \\ &= \Psi^* Z^* \hat{Y}_k^* r_m \\ &= \Psi^* Z^* \hat{Y}_k^* X \rho_m \\ &= \Psi^* Z^* \hat{D}\hat{\rho}\end{aligned}$$

So $\Sigma^* V_2^* r_m = \Psi^* Z^* \hat{D}\hat{\rho}$. Equivalently, $Z^* \hat{D}\hat{\rho} = \Psi\Sigma^* V_2^* r_m$. Then

$$\|Z^* \hat{D}\hat{\rho}\| \leq \|\Sigma\|\|r_m\| = \sigma\|r_m\|$$

where $\sigma = \max_i(\sigma_i)$, i.e. $(1-\sigma^2)^{1/2}$ is the smallest singular value of $V^* Y_k$. □

So if σ is small, then so is $\|Z^* \hat{D}\hat{\rho}\|$. If $Z^* \hat{D}$ is well-conditioned, then $\|\hat{\rho}\|$ will be small. Hence X_k will be nearly removed from r_m.

Chapter 5

Angles Between Krylov Spaces

We have seen that if the projection spaces are close, then methods will approximate one another. For instance, in Section 3.2 we explained that BiCG in the eigenvector basis approximates a FOM iteration if the coefficients in $D = Y^*X$ are not too small, as demonstrated in Figure 3.1. We now quantify this effect by analyzing the residuals of BiCG, GMRES, and FOM based on the principal angles between the corresponding projection subspaces.

As established in Chapter 2, the residuals at the m^{th} iteration of GMRES, BiCG, and FOM satisfy the following:

$$\begin{array}{lll} \text{GMRES} & r = r_0 - Q_m y \perp Q_m & y = Q_m^* r_0. \\ \text{BiCG} & \tilde{r} = r_0 - Q_m \tilde{y} \perp \tilde{Q}_m & \tilde{y} = (\tilde{Q}_m^* Q_m)^{-1} \tilde{Q}_m^* r_0. \\ \text{FOM} & \hat{r} = r_0 - Q_m \hat{y} \perp \hat{Q}_m & \hat{y} = (\hat{Q}_m^* Q_m)^{-1} \hat{Q}_m^* r_0. \end{array}$$

where Q_m, $\tilde{Q}_m$, and $\hat{Q}_m$ are any convenient matrices whose columns form orthonormal bases for $A\mathcal{K}^m(A; r_0)$, $\mathcal{K}^m(A^*; r_0)$, and $\mathcal{K}^m(A; r_0)$, respectively [4]. We assume that $(\tilde{Q}_m^* Q_m)^{-1}$ and $(\hat{Q}_m^* Q_m)^{-1}$ exist. This chapter is concerned with establishing relationships based on the principal angles between $Q_m, \tilde{Q}_m$, and $\hat{Q}_m$. Some of these relations are stated in [4] without proof, and the proofs are presented here.

5.1 Relations in the Euclidean Inner Product

Since $\mathcal{K}^m(A; r_0) \cap A\mathcal{K}^m(A; r_0) = A\mathcal{K}^{m-1}(A; r_0)$, we can construct Q_m and $\hat{Q}_m$ such that $Q_{m-1} = \hat{Q}_{m-1}$ is an orthonormal basis for $A\mathcal{K}^{m-1}(A; r_0)$ so that

$$\hat{q}_m = \frac{r_0 - \sum_{i=1}^{m-1} \langle r_0, \hat{q}_i \rangle \hat{q}_i}{\| r_0 - \sum_{i=1}^{m-1} \langle r_0, \hat{q}_i \rangle \hat{q}_i \|} \text{ and } q_m = \frac{A^m r_0 - \sum_{i=1}^{m-1} \langle A^m r_0, q_i \rangle q_i}{\| A^m r_0 - \sum_{i=1}^{m-1} \langle A^m r_0, q_i \rangle q_i \|} \alpha$$

with $|\alpha| = 1$ and $q_m^* \hat{q}_m = \omega_m \in \mathbb{R} > 0$. Here, q_i and $\hat{q}_i$ denote the ith columns of Q_m and $\hat{Q}_m$ respectively.

Theorem 5.1. *$r - \tilde{r} = Q_m U \Sigma^{-1} (I - \Sigma^2)^{1/2} C^* r = \sum_{i=1}^{m} Q_m u_i (c_i^* r \tan(\phi_i))$ where $Q_m^* \tilde{Q}_m = U \Sigma V^*$ is a singular value decomposition, $\tilde{Q}_m V = Q_m U \Sigma + C(I - \Sigma^2)^{1/2}$, and $\phi_i = \arccos(\sigma_i)$ are the principal angles between Q_m and $\tilde{Q}_m$ [4].*

Proof.

$$
\begin{aligned}
r - \tilde{r} &= Q_m\tilde{y} - Q_m y \\
&= Q_m(\tilde{Q}_m^* Q_m)^{-1}\tilde{Q}_m^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m((Q_m^*\tilde{Q}_m)^*)^{-1}\tilde{Q}_m^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m((U\Sigma V^*)^*)^{-1}\tilde{Q}_m^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m(V\Sigma U^*)^{-1}\tilde{Q}_m^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m U\Sigma^{-1}V^*\tilde{Q}_m^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m U\Sigma^{-1}(\tilde{Q}_m V)^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m U\Sigma^{-1}(Q_m U\Sigma + C(I-\Sigma^2)^{1/2})^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m U\Sigma^{-1}(\Sigma U^* Q_m^* + (I-\Sigma^2)^{1/2}C^*) r_0 - Q_m Q_m^* r_0 \\
&= Q_m Q_m^* r_0 + Q_m U\Sigma^{-1}(I-\Sigma^2)^{1/2}C^* r_0 - Q_m Q_m^* r_0 \\
&= Q_m U\Sigma^{-1}(I-\Sigma^2)^{1/2}C^* r_0 \\
&= Q_m U\Sigma^{-1}(I-\Sigma^2)^{1/2}C^* r \text{ (since } C^* Q_m = 0 \text{ and } r = r_0 - Q_m y) \\
&= \sum_{i=1}^{m} Q_m u_i \left(c_i^* r \frac{(1-\sigma^2)^{1/2}}{\sigma} \right) \\
&= \sum_{i=1}^{m} Q_m u_i (c_i^* r \tan(\phi_i)). \qquad (5.1)
\end{aligned}
$$

□

Hence, if the principal angles between Q_m and $\tilde{Q}_m$ are relatively small (i.e. the singular values σ_i are not close to zero), the BiCG residual is not far from the GMRES residual. If we look only at the principal angles, however, we get a pessimistic bound. This is demonstrated in the following figure.

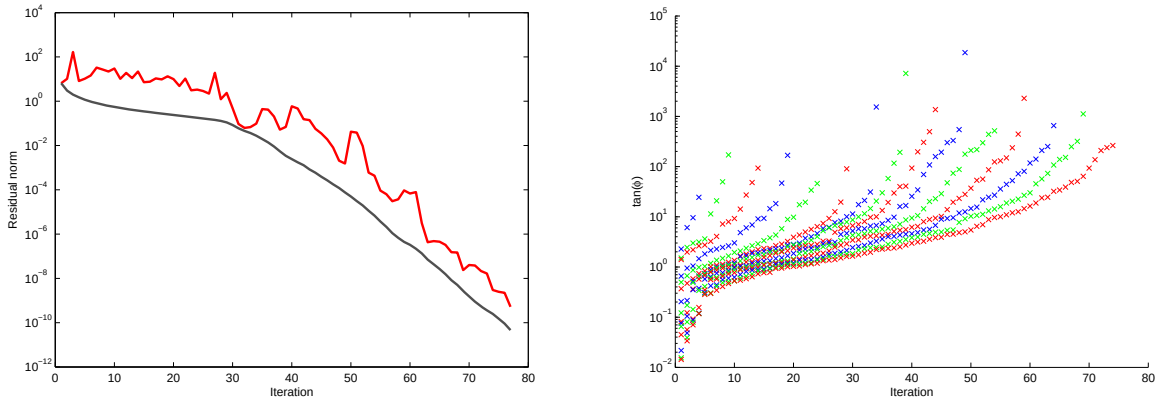

Figure 5.1: Left: BiCG residual (red) and GMRES residual (grey) for a convection-diffusion problem. Right: Tangents of principal angles between Q_m and $\tilde{Q}_m$.

The following theorem relates the residuals of BiCG and FOM.

Theorem 5.2. $\tilde{r}-\hat{r} = \tan(\theta_m)q_m(g_m^* r) - \left(\sum_{i=1}^{m} Q_m u_i(c_i^* r \tan(\phi_i))\right)$ *where* g_m *is a unit vector and* $\theta_m = \arccos(\omega_m)$ *where* $\omega_m = q_m^* \hat{q}_m$.

Proof.

$$Q_m^* \hat{Q}_m = \begin{bmatrix} I_{(m-1)\times(m-1)} & 0 \\ 0 & \omega_m \end{bmatrix} = \Omega = I\Omega I^*.$$
$$\hat{Q}_m = Q_m\Omega + G(I-\Omega^2)^{1/2}.$$

where $G = \begin{bmatrix} g_1 & g_2 & \dots & g_m \end{bmatrix}$ has orthonormal columns. By Theorem 5.1, $r-\tilde{r} = Q_m U\Sigma^{-1}(I-\Sigma^2)^{1/2}C^* r = \sum_{i=1}^{m} Q_m u_i(c_i^* r\tan(\phi_i))$, and

$$\begin{aligned} r-\hat{r} &= Q_m\Omega^{-1}(I-\Omega^2)^{1/2}G^* r \\ &= Q_m \begin{bmatrix} 0_{(m-1)\times(m-1)} & 0_{(m-1)\times 1} \\ 0_{1\times(m-1)} & \frac{(1-\omega_m^2)^{1/2}}{\omega_m} \end{bmatrix} \begin{bmatrix} g_1^* r \\ \vdots \\ g_m^* r \end{bmatrix} \\ &= q_m \frac{(1-\omega_m^2)^{1/2}}{\omega_m} g_m^* r. \end{aligned} \tag{5.2}$$

As a result, $\|\hat{r}-r\| = \left(\frac{1-\omega_m^2}{\omega_m^2}\|G^* r\|^2\right)^{1/2}$ and so

$$\begin{aligned} \|\hat{r}\| &= \left(\|r\|_2^2 + \frac{1-\omega^2}{\omega^2}\|G^* r\|^2\right)^{1/2} \\ &= \left(\|r\|^2 + \frac{1-\omega_m^2}{\omega_m^2}\|r\|^2\right)^{1/2} \\ &= \left(\frac{1}{\omega_m^2}\|r\|^2\right)^{1/2} \\ &= \frac{1}{\omega_m}\|r\|. \end{aligned}$$

This leads to (2.4), but with a different derivation than presented in [8]. (5.1) and (5.2) are combined to obtain

$$\begin{aligned} \tilde{r}-\hat{r} &= (r-\hat{r}) - (r-\tilde{r}) \\ &= q_m \frac{(1-\omega_m^2)^{1/2}}{\omega_m} g_m^* r - Q_m U\Sigma^{-1}(I-\Sigma^2)^{1/2}C^* r \\ &= \tan(\theta_m)q_m(g_m^* r) - \left(\sum_{i=1}^{m} Q_m u_i(c_i^* r\tan(\phi_i))\right) \end{aligned}$$

where $\theta_m = \arccos(\omega_m)$ is the angle between q_m and $\hat{q}_m$. □

Referring back to Figures 3.1 and 3.2, we see that the BiCG and FOM residuals are close when the singular values of $\tilde{Q}_m^* \hat{Q}_m$ are bounded away from zero, i.e. the principal angles are bounded away from $\pi/2$. In Figure 3.1, the singular values are mostly near 1, with the smallest singular value only as small as 10^{-4}. In this case, the residuals are close. In Figure 3.2, the singular values systematically get much smaller, some as small as 10^{-15}. Hence, the principal angles are close to $\pi/2$ and the projection spaces are far from each other. In this case, the BiCG residual does not converge, while the FOM residual does.

Theorem 5.3. *If we construct Q_m and $\hat{Q}_m$ such that $Q_{m-1} = \hat{Q}_{m-1}$, then the first $m-1$ singular values of $Q_m^* \tilde{Q}_m$ are bounded below by the smallest singular value of $\hat{Q}_m^* \tilde{Q}_m$.*

Proof. Since $Q_{m-1} = \hat{Q}_{m-1}$, we have that $Q_{m-1}^* \tilde{Q}_m = \hat{Q}_{m-1}^* \tilde{Q}_m$. The singular values μ_i of $Q_{m-1}^* \tilde{Q}_m$ and σ_i of $Q_m^* \tilde{Q}_m$ satisfy $\sigma_1 \geq \mu_1 \geq \sigma_2 \geq ... \geq \sigma_{m-1} \geq \mu_{m-1} \geq \sigma_m \geq 0$. The singular values γ_i of $\hat{Q}_m^* \tilde{Q}_m$ satisfy $\gamma_1 \geq \mu_1 \geq ... \geq \mu_{m-1} \geq \gamma_m$ [10, Section 7.3]. Then σ_k is bounded below by γ_{k+1} for all $1 \leq k \leq m-1$. In particular, if the γ_i are bounded away from zero, then so are the σ_i except possibly for σ_m. □

Consequently, if all but a few of the γ_i or μ_i converge to 1, then we have the same for the σ_i. This implies that if BiCG well-approximates FOM, then BiCG must also approximate GMRES up to a possible large error in one direction.

5.2 Relations in the D-Inner Product

In Chapters 3 and 4, we discussed an implicit BiCG in the D-inner product where D was a diagonal matrix with positive real coefficients. Thus it is helpful to analyze the effect of this inner product on the above results. In this section, we generalize the results from Section 5.1 to the case of the D-inner product.

For any diagonal matrix D with positive real coefficients, $\langle x, y \rangle_D = y^* D x$ defines an inner product. Redefining orthogonality and normalization with this D-inner product, we can compute Q_m and $\tilde{Q}_m$ such that $Q_m^* D Q_m = I$ and $\tilde{Q}_m^* D \tilde{Q}_m = I$. If we define GMRES and BiCG with respect to the D-inner product, the residuals satisfy

$$\begin{array}{lll} \text{GMRES} & r = r_0 - Q_m y \perp_D Q_m & y = Q_m^* D r_0. \\ \text{BiCG} & \tilde{r} = r_0 - Q_m \tilde{y} \perp_D \tilde{Q}_m & \tilde{y} = (\tilde{Q}_m^* D Q_m)^{-1} \tilde{Q}_m^* D r_0. \end{array}$$

Let $Q_m^* D \tilde{Q}_m = W \Delta Z^*$ be a singular value decomposition. So $W^* W = I$, $Z^* Z = I$, and the diagonal coefficients δ_i of Δ are nonnegative real numbers.

Lemma 5.4. *Assume $m < \frac{n}{2}$. Then $\tilde{Q}_m Z = Q_m W \Delta + J(I - \Delta^2)^{1/2}$ where $J^* D Q_m = 0$ and $J^* D J = I$, i.e. J is orthonormal in the D-inner product.*

Proof. Set $C = \tilde{Q}_m Z - Q_m W\Delta$. Then for all $1 \le k \le m$,

$$\begin{aligned}
\|c_k\|_D = \|Ce_k\|_D &= \|(\tilde{Q}_m Z - Q_m W\Delta)e_k\|_D \\
&= \|\tilde{Q}_m z_k - Q_m w_k \delta_k\|_D \\
&= ((\tilde{Q}_m z_k - Q_m w_k\delta_k)^* D(\tilde{Q}_m z_k - Q_m w_k\delta_k))^{1/2} \\
&= ((z_k^*\tilde{Q}_m^* - \delta_k w_k^* Q_m^*) D(\tilde{Q}_m z_k - Q_m w_k\delta_k))^{1/2} \\
&= (z_k^*\tilde{Q}_m^* D\tilde{Q}_m z_k - z_k^*\tilde{Q}_m^* DQ_m w_k\delta_k - \delta_k w_k^* Q_m^* D\tilde{Q}_m z_k \\
&\quad + \delta_k w_k^* Q_m^* DQ_m w_k\delta_k)^{1/2} \\
&= (z_k^* z_k - z_k^* Z\Delta W^* w_k\delta_k - \delta_k w_k^* W\Delta Z^* z_k + \delta_k w_k^* w_k\delta_k)^{1/2} \\
&= (1 - e_k^*\Delta e_k\delta_k - \delta_k e_k^*\Delta e_k + \delta_k^2)^{1/2} \\
&= (1 - \delta_k^2 - \delta_k^2 + \delta_k^2)^{1/2} \\
&= (1 - \delta_k^2)^{1/2}.
\end{aligned}$$

Then $c_k = j_k(1-\delta^2)^{1/2}$ for some unit vector j_k. Note that if $\delta_k = 1$, then $c_k = 0$ and we can choose j_k freely such that j_k is a unit vector. Hence, $C = J(I-\Delta^2)^{1/2}$ for some J with unit columns. Thus,

$$\begin{aligned}
\tilde{Q}_m Z - Q_m W\Delta &= J(I-\Delta^2)^{1/2} \\
\tilde{Q}_m Z &= Q_m W\Delta + J(I-\Delta^2)^{1/2}
\end{aligned}$$

and

$$\begin{aligned}
(J(I-\Delta^2)^{1/2})^* DQ_m &= (Z^*\tilde{Q}_m^* - \Delta W^* Q_m^*)DQ_m \\
&= Z^*\tilde{Q}_m^* DQ_m - \Delta W^* Q_m^* DQ_m \\
&= Z^* Z\Delta W^* - \Delta W^* \\
&= \Delta W^* - \Delta W^* \\
&= 0.
\end{aligned}$$

Therefore, $0 = (J(I-\Delta^2)^{1/2})^* DQ_m = (I-\Delta^2)^{1/2} J^* DQ_m$. For ease of discussion, we will first assume that $\delta_k < 1$ for all k. Then $(I-\Delta^2)^{1/2}$ is nonsingular, and therefore $J^* DQ_m = 0$ and J is D-orthogonal to Q_m. For the next step, we need the following result: Since $Z^* Z = I$,

$$(\tilde{Q}_m Z)^* D(\tilde{Q}_m Z) = Z^*\tilde{Q}_m^* D\tilde{Q}_m Z = Z^* Z = I.$$

This yields

$$\begin{aligned}
I &= (\tilde{Q}_m Z)^* D(\tilde{Q}_m Z) \\
&= (Q_m W\Delta + J(I-\Delta^2)^{1/2})^* D(Q_m W\Delta + J(I-\Delta^2)^{1/2}) \\
&= \Delta W^* Q_m^* DQ_m W\Delta + \Delta W^* Q_m^* DJ(I-\Delta^2)^{1/2} + \\
&\qquad (I-\Delta^2)^{1/2} J^* DQ_m W\Delta + (I-\Delta^2)^{1/2} J^* DJ(I-\Delta^2)^{1/2} \\
&= \Delta^2 + \Delta W^*(J^* DQ_m)(I-\Delta^2)^{1/2} + (I-\Delta^2)^{1/2}(J^* DQ_m)W\Delta \\
&\quad + (I-\Delta^2)^{1/2} J^* DJ(I-\Delta^2)^{1/2} \\
&= \Delta^2 + (I-\Delta^2)^{1/2} J^* DJ(I-\Delta^2)^{1/2} \\
\iff I - \Delta^2 &= (I-\Delta^2)^{1/2} J^* DJ(I-\Delta^2)^{1/2} \\
\iff I &= J^* DJ.
\end{aligned}$$

Hence, J is a D-orthonormal matrix.

Suppose $\delta_k = 1$ for some k. Let j_k denote the kth column of J. Since $m < \frac{n}{2}$, there are less than $\frac{n}{2}$ other columns vectors in J and less than $\frac{n}{2}$ column vectors in Q_m. Since we are working in a vector space of dimension n, there exists a unit vector D-orthogonal to all other column vectors in J as well as to Q_m. Then since j_k can be chosen freely, we can choose j_k to be this vector. The same reasoning can be used for any number of $\delta_i = 1$. Then we can construct J such that $J^*DQ_m = 0$ and $J^*DJ = I$. □

We can now establish a relationship between the residuals of GMRES and BiCG.

Theorem 5.5. *If $m < \frac{n}{2}$, then the residuals at the mth iteration satisfy*

$$r - \tilde{r} = Q_m W \Delta^{-1}(I - \Delta^2)^{1/2} J^* D r = \sum_{i=1}^{m} Q_m w_i (j_i^* D r \tan(\psi_i))$$

where $\psi_i = \arccos(\delta_i)$.

Proof.

$$\begin{aligned}
r - \tilde{r} &= Q_m \tilde{y} - Q_m y \\
&= Q_m(\tilde{Q}_m^* D Q_m)^{-1} \tilde{Q}_m^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m((Q_m^* D \tilde{Q}_m)^*)^{-1} \tilde{Q}_m^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m((W \Delta Z^*)^*)^{-1} \tilde{Q}_m^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m(Z \Delta W^*)^{-1} \tilde{Q}_m^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m W \Delta^{-1} Z^* \tilde{Q}_m^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m W \Delta^{-1} (\tilde{Q}_m Z)^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m W \Delta^{-1} (Q_m W \Delta + J(I - \Delta^2)^{1/2})^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m W \Delta^{-1} (\Delta W^* Q_m^* + (I - \Delta^2)^{1/2} J^*) D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m Q_m^* D r_0 + Q_m W \Delta^{-1} (I - \Delta^2)^{1/2} J^* D r_0 - Q_m Q_m^* D r_0 \\
&= Q_m W \Delta^{-1} (I - \Delta^2)^{1/2} J^* D r_0.
\end{aligned}$$

Since $J^*DQ_m = 0$, we have

$$\begin{aligned}
J^* D r &= J^* D (r_0 - Q_m y) \\
&= J^* D r_0 - J^* D Q_m y \\
&= J^* D r_0.
\end{aligned}$$

So

$$\begin{aligned}
r - \tilde{r} &= Q_m W \Delta^{-1} (I - \Delta^2)^{1/2} J^* D r_0 \\
&= Q_m W \Delta^{-1} (I - \Delta^2)^{1/2} J^* D r \\
&= \sum_{i=1}^{m} Q_m w_i (j_i^* D r \tan(\psi_i)).
\end{aligned}$$

where $\psi_i = \arccos(\delta_i)$. □

So BiCG will approximate GMRES in the D-inner product if the δ_i are close to 1, provided $m < \frac{n}{2}$.

Chapter 6

Conclusions

In this thesis we seek to explain the convergence behavior of BiCG through a theoretical analysis of the method. We find that the BiCG residual will be close to the GMRES residual if the angles between the respective projection subspaces are relatively small, as demonstrated in Figure 5.1. However, the BiCG residual may not converge if some of the canonical angles are systematically close to $\pi/2$, i.e. the projection spaces are far from one another. These results hold in both the Euclidean and the D-inner product.

Furthermore, there is an equivalent BiCG in the (generalized) eigenvector basis using the D-inner product, or using the Euclidean inner product with a special left starting vector. Viewing BiCG in this basis makes several mathematical properties more clear. We see that small components in $D = Y^*X$ are equivalent to damping the corresponding left eigenvector components in the left starting vector. Such small components can prevent BiCG from converging because it prevents the damping of the corresponding right eigenvector components in the residual. Moreover, we see that as left eigenvectors converge in the left Krylov space, the corresponding right eigenvectors are removed from the residual. More generally, as left invariant subspaces converge in the left Krylov space, components in the corresponding right subspaces are removed from the residual. When a left eigenvector/invariant subspace is *almost* contained in the left Krylov space, then the residual component in the direction of the corresponding right eigenvector/subspace is bounded. This bound depends mainly on the coefficients in D.

For future work, we would like to further explore the convergence of invariant subspaces discussed in Section 4.3 and explore the properties of Z^*D in Theorem 4.8.

Bibliography

[1] W. E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quarterly of Applied Mathematics*, 9(17):17–29, 1951.

[2] Jane Cullum and Anne Greenbaum. Relations between Galerkin and norm-minimizing iterative methods for solving linear systems. *SIAM Journal on Matrix Analysis and Applications*, 17(2):223–247, 1996.

[3] Eric de Sturler. Bicg explained. In *Proceedings of the Householder International Symposium in Numerical Algebra*, pages 254–257, June 1999.

[4] Eric de Sturler. The convergence behavior of biconjugate gradients. Computational Mathematics and Applications Seminar, University of Oxford, November 2010.

[5] Harry Dym. *Linear Algebra in Action*, volume 78 of *Graduate Studes in Mathematics*. American Mathematical Society, 2007.

[6] Vance Faber and Thomas Manteuffel. Necessary and sufficient conditions for the existence of a conjugate gradient method. *SIAM Journal on Numerical Analysis*, 21(2):352–362, 1984.

[7] R. Fletcher. Conjugate gradient methods for indefinite systems. In G. A. Watson, editor, *Proceedings of the Dundee Conference on Numerical Analysis, 1975*, volume 506, pages 73–89. Springer-Verlag, 1976.

[8] Anne Greenbaum. On the role of the left starting vector in the two-sided lanczos algorithm and nonsymmetric linear system solvers. In D. F. Griffiths, D. J. Higham, and G. A. Watson, editors, *Proceedings of the 17th Dundee Biennial Conference on Numerical Analysis*, pages 124–132. Addison Wesley Longman Inc., June 1997.

[9] Magnus R. Hestenes and Eduard Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49(6):409–436, 1952.

[10] Roger A. Horn and Charles R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

[11] Cornelius Lanczos. 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, 1950.

[12] Yousef Saad. Krylov subspace methods for solving large unsymmetric linear systems. *Mathematics of Computation*, 37(155):105–126, 1981.

[13] Yousef Saad. *Iterative Methods for Sparse Linear Systems*. PWS Publishing, 1996.

[14] Yousef Saad and Martin H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7:856–869, 1986.

[15] Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra*. SIAM, 1997.

[16] Henk A. van der Vorst. Bi-CGSTAB: A fast and smoothly convergent variant of bi-cg for the solution of nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 13:631–644, 1992.

[17] Henk A. van der Vorst. *Iterative Krylov Methods for Large Linear Systems*. Cambridge University Press, 2003.