

Algorithm 913: An Elegant IDR(s) Variant that Efficiently Exploits Biorthogonality Properties

MARTIN B. VAN GIJZEN and PETER SONNEVELD, Delft University of Technology

The IDR(s) method that is proposed in Sonneveld and van Gijzen [2008] is a very efficient limited memory method for solving large nonsymmetric systems of linear equations. IDR(s) is based on the induced dimension reduction theorem, that provides a way to construct subsequent residuals that lie in a sequence of shrinking subspaces. The IDR(s) algorithm that is given in Sonneveld and van Gijzen [2008] is a direct translation of the theorem into an algorithm. This translation is not unique. This article derives a new IDR(s) variant, that imposes (one-sided) biorthogonalization conditions on the iteration vectors. The resulting method has lower overhead in vector operations than the original IDR(s) algorithms. In exact arithmetic, both algorithms give the same residual at every $(s + 1)$ -st step, but the intermediate residuals and also the numerical properties differ. We show through numerical experiments that the new variant is more stable and more accurate than the original IDR(s) algorithm, and that it outperforms other state-of-the-art techniques for realistic test problems.

Categories and Subject Descriptors: G.1.3 [Numerical Analysis]: Numerical Linear Algebra—*Linear systems (direct and iterative methods)*; G.4 [Mathematical Software]: *Algorithm design and analysis*

General Terms: Algorithms, Performance, Reliability

Additional Key Words and Phrases: Iterative methods, IDR, IDR(s), Krylov-subspace methods, sparse nonsymmetric linear systems

ACM Reference Format:

van Gijzen, M. B. and Sonneveld, P. 2011. Algorithm 913: An elegant IDR(s) variant that efficiently exploits biorthogonality properties. *ACM Trans. Math. Softw.* 38, 1, Article 5 (November 2011), 19 pages. DOI = 10.1145/2049662.2049667 <http://doi.acm.org/10.1145/2049662.2049667>

1. INTRODUCTION

We consider the linear system

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

with $\mathbf{A} \in \mathbb{C}^{N \times N}$ a large, sparse, and nonsymmetric matrix. In 1980, Sonneveld proposed the iterative method IDR [Wesseling and Sonneveld 1980] for solving such systems. The IDR method generates residuals that are forced to be in subspaces \mathcal{G}_j of decreasing dimension. These nested subspaces are related by $\mathcal{G}_j = (\mathbf{I} - \omega_j \mathbf{A})(\mathcal{S} \cap \mathcal{G}_{j-1})$, where \mathcal{S} is a fixed proper subspace of \mathbb{C}^N , and the ω_j 's are nonzero scalars.

Recently, it was recognized that this IDR approach is quite general and can be used as a framework for deriving iterative methods. This observation has led to the development of IDR(s) [Sonneveld and van Gijzen 2008]. IDR(s) has attracted considerable attention; for example, see Gutknecht [2010], Gutknecht and Zemke [2010], Onoue

Part of this research has been funded by the Dutch BSIK/BRICKS project.

Authors' addresses: M. B. van Gijzen (corresponding author) and P. Sonneveld, Delft Institute of Applied Mathematics, Mekelweg 4, 2628 Delft, The Netherlands; email: m.b.vangijzen@tudelft.nl.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2011 ACM 0098-3500/2011/11-ART5 \$10.00

DOI 10.1145/2049662.2049667 <http://doi.acm.org/10.1145/2049662.2049667>

et al. [2008, 2009], and Simoncini and Szyld [2010]. The examples that are described in Sonneveld and van Gijzen [2008] show that $\text{IDR}(s)$, with $s > 1$ and not too big, outperforms the well-known Bi-CGSTAB method [van der Vorst 1992] for important classes of problems.

Although the working principle of $\text{IDR}(s)$ differs from that of Bi-CGSTAB, it turns out that both methods are mathematically closely related. Specifically, $\text{IDR}(1)$ is mathematically equivalent with Bi-CGSTAB, and $\text{IDR}(s)$ with $s > 1$ is closely related (but not mathematically equivalent) to the Bi-CGSTAB generalisation $\text{ML}(k)\text{BiCGSTAB}$ [Yeung and Chan 1999] of Yeung and Chan. We refer to Sonneveld and van Gijzen [2008] for the details. In Sleijpen et al. [2010], Bi-CGSTAB is considered as an IDR method, and that paper explains how IDR ideas can be incorporated into Bi-CGSTAB.

The prototype $\text{IDR}(s)$ algorithm that is described in Sonneveld and van Gijzen [2008] is only one of many possible variants. One of the possibilities to make alternative IDR methods is a different computation of the *intermediate* residuals. In $\text{IDR}(s)$, the residual is uniquely defined in every $(s+1)$ -st step, [Sonneveld and van Gijzen 2008, Section 5.1]. This step corresponds to the calculation of the first residual in \mathcal{G}_j . In order to advance to \mathcal{G}_{j+1} , s additional residuals in \mathcal{G}_j need to be computed. These intermediate residuals are not uniquely defined and their computation leaves freedom to derive algorithmic variants. In exact arithmetic, the residuals at every $(s+1)$ -st step do not depend on the way the intermediate residuals are computed. The numerical stability and efficiency of the specific IDR algorithm, however, do depend on the computation of the intermediate residuals.

In this article we will derive an elegant, efficient, and in our experience numerically very stable IDR-based method that imposes and exploits as much as possible (one-sided) biorthogonality conditions between the intermediate residuals and the pre-chosen vectors $\mathbf{p}_1, \dots, \mathbf{p}_s$, that define the subspace \mathcal{S} . We will denote this new IDR variant by *IDR(s)-biortho*¹ to distinguish it from *IDR(s)-proto*, the prototype algorithm in Sonneveld and van Gijzen [2008]. *IDR(s)-biortho* uses less vector operations per iteration than *IDR(s)-proto*, and has better stability properties, in particular for large values of s .

This article is organized as follows.

The next section describes a general framework for deriving an IDR-based method. It starts with reviewing the IDR theorem. Then it explains how the theorem can be used to compute the first residual in \mathcal{G}_{j+1} and the corresponding approximation for the solution, given sufficient vectors in \mathcal{G}_j . Furthermore it explains how sufficient intermediate residuals and vectors in \mathcal{G}_{j+1} can be computed in order to advance to the next lower-dimensional subspace, and what freedom there is in generating these intermediate vectors.

Section 3 derives the *IDR(s)-biortho* variant by filling in the freedom in generating the intermediate residuals by imposing biorthogonality conditions between the intermediate residuals and the vectors $\mathbf{p}_1, \dots, \mathbf{p}_s$.

Section 4 presents numerical experiments that compare *IDR(s)-biortho* and *IDR(s)-proto*. It also shows experiments to illustrate the excellent performance of *IDR(s)* in comparison with state-of-the-art methods like Bi-CGSTAB, GMRES [Saad and Schultz 1986], and $\text{BiCGstab}(\ell)$ [Sleijpen and Fokkema 1994].

We make concluding remarks in Section 5.

¹Other authors have used different names for *IDR(s)-biortho*: the method is called *IDR(s)BiO* in Gutknecht [2010], *IDRBiO* in Gutknecht and Zemke [2010], and *Bi-IDR(s)* in Onoue et al. [2009].

2. DESIGNING AN IDR-BASED ALGORITHM

2.1. The IDR Theorem

At the basis of every IDR algorithm is the IDR Theorem [Sonneveld and van Gijzen 2008], which is given next.

THEOREM 2.1. *Let \mathbf{A} be any matrix in $\mathbb{C}^{N \times N}$, let \mathbf{v}_0 be any nonzero vector in \mathbb{C}^N , and let \mathcal{G}_0 be the full Krylov space $\mathcal{K}^N(\mathbf{A}, \mathbf{v}_0)$. Let \mathcal{S} denote any (proper) subspace of \mathbb{C}^N such that \mathcal{S} and \mathcal{G}_0 do not share a nontrivial invariant subspace of \mathbf{A} , and define the sequence \mathcal{G}_j , $j = 1, 2, \dots$ as*

$$\mathcal{G}_j = (\mathbf{I} - \omega_j \mathbf{A})(\mathcal{G}_{j-1} \cap \mathcal{S}),$$

where the ω_j 's are nonzero scalars. Then

- (i) $\mathcal{G}_j \subset \mathcal{G}_{j-1}$ for all $\mathcal{G}_{j-1} \neq \{\mathbf{0}\}$, $j > 0$.
- (ii) $\mathcal{G}_j = \{\mathbf{0}\}$ for some $j \leq N$.

For the proof we refer to Sonneveld and van Gijzen [2008].

Without loss of generality, we may assume the space \mathcal{S} to be the left null space of some (full rank) $N \times s$ matrix \mathbf{P} .

$$\mathbf{P} = (\mathbf{p}_1 \ \mathbf{p}_2 \ \dots \ \mathbf{p}_s), \quad \mathcal{S} = \mathcal{N}(\mathbf{P}^H)$$

2.2. General Recursions

Let $\mathbf{A}\mathbf{x} = \mathbf{b}$ be an $N \times N$ linear system. A Krylov-type solver produces iterates \mathbf{x}_n for which the residuals $\mathbf{r}_n = \mathbf{b} - \mathbf{A}\mathbf{x}_n$ are in the Krylov spaces $\mathcal{K}^n(\mathbf{A}, \mathbf{r}_0) = \text{span}\{\mathbf{r}_0, \mathbf{A}\mathbf{r}_0, \dots, \mathbf{A}^n \mathbf{r}_0\}$. Here, \mathbf{x}_0 is an initial estimate of the solution.

An IDR-based method can be made by using recursions of the following form

$$\begin{aligned} \mathbf{r}_{n+1} &= \mathbf{r}_n - \alpha \mathbf{A} \mathbf{v}_n - \sum_{i=1}^{\hat{i}} \gamma_i \mathbf{g}_{n-i} \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + \alpha \mathbf{v}_n + \sum_{i=1}^{\hat{i}} \gamma_i \mathbf{u}_{n-i} \end{aligned} \quad (1)$$

in which \mathbf{v}_n is any computable vector in $\mathcal{K}^n(\mathbf{A}, \mathbf{r}_0) \setminus \mathcal{K}^{n-1}(\mathbf{A}, \mathbf{r}_0)$, $\mathbf{g}_{n-i} \in \mathcal{K}^n(\mathbf{A}, \mathbf{r}_0)$, and \mathbf{u}_{n-i} such that

$$\mathbf{g}_{n-i} = \mathbf{A} \mathbf{u}_{n-i}. \quad (2)$$

These recursions are quite general and hold for many Krylov subspace methods.

The IDR theorem can be applied by generating residuals \mathbf{r}_n that are forced to be in the subspaces \mathcal{G}_j , where j is nondecreasing with increasing n . Then, according to Theorem 2.1, $\mathbf{r}_n \in \{\mathbf{0}\}$ for some n .

2.3. A Dimension-Reduction Step: Computing the First Residual in \mathcal{G}_{j+1}

According to Theorem 2.1, the residual \mathbf{r}_{n+1} is in \mathcal{G}_{j+1} if

$$\mathbf{r}_{n+1} = (\mathbf{I} - \omega_{j+1} \mathbf{A}) \mathbf{v}_n \quad \text{with } \mathbf{v}_n \in \mathcal{G}_j \cap \mathcal{S}.$$

If we choose

$$\mathbf{v}_n = \mathbf{r}_n - \sum_{i=1}^{\hat{i}} \gamma_i \mathbf{g}_{n-i} \quad (3)$$

the expression for \mathbf{r}_{n+1} reads

$$\mathbf{r}_{n+1} = \mathbf{r}_n - \omega_{j+1} \mathbf{A} \mathbf{v}_n - \sum_{i=1}^{\hat{i}} \gamma_i \mathbf{g}_{n-i}, \quad (4)$$

which corresponds to (1), with $\alpha = \omega_{j+1}$.

Now suppose that $\mathbf{r}_n, \mathbf{g}_{n-i} \in \mathcal{G}_j$, $i = 1, \dots, \hat{i}$. This implies that $\mathbf{v}_n \in \mathcal{G}_j$. If we choose $\gamma_i, i = 1, \dots, \hat{i}$ such that in addition $\mathbf{v}_n \in \mathcal{S}$, then by Theorem 2.1 we have $\mathbf{r}_{n+1} \in \mathcal{G}_{j+1}$.

If $\mathbf{v}_n \in \mathcal{S} = \mathcal{N}(\mathbf{P}^H)$, it satisfies

$$\mathbf{P}^H \mathbf{v}_n = \mathbf{0}. \quad (5)$$

Combining (3) and (5) yields an $s \times \hat{i}$ linear system for the coefficients γ_i . Except for special circumstances, this system is uniquely solvable if $\hat{i} = s$, which means that we need s vectors $\mathbf{g}_i \in \mathcal{G}_j$ for $\mathbf{r}_{n+1} \in \mathcal{G}_{j+1}$.

Suppose that after n iterations we have exactly s vectors $\mathbf{g}_i \in \mathcal{G}_j$, $i = n-1, \dots, n-s$, and s corresponding vectors \mathbf{u}_i with $\mathbf{g}_i = \mathbf{A} \mathbf{u}_i$ at our disposal. Define the matrices

$$\mathbf{G}_n = (\mathbf{g}_{n-s} \ \mathbf{g}_{n-s+1} \ \dots \ \mathbf{g}_{n-1}), \quad (6)$$

$$\mathbf{U}_n = (\mathbf{u}_{n-s} \ \mathbf{u}_{n-s+1} \ \dots \ \mathbf{u}_{n-1}). \quad (7)$$

Then the computation of the residual $\mathbf{r}_{n+1} \in \mathcal{G}_{n+1}$ can be implemented by the following algorithm.

Calculate: $\mathbf{c} \in \mathbb{C}^s$ from $(\mathbf{P}^H \mathbf{G}_n) \mathbf{c} = \mathbf{P}^H \mathbf{r}_n$,
 $\mathbf{v}_n = \mathbf{r}_n - \mathbf{G}_n \mathbf{c}$,
 $\mathbf{r}_{n+1} = \mathbf{v}_n - \omega_{j+1} \mathbf{A} \mathbf{v}_n$.

According to (4), the new residual satisfies

$$\mathbf{r}_{n+1} = \mathbf{r}_n - \omega_{j+1} \mathbf{A} \mathbf{v}_n - \mathbf{G}_n \mathbf{c}.$$

Multiplying this expression with \mathbf{A}^{-1} yields the corresponding recursion for the iterate.

$$\mathbf{x}_{n+1} = \mathbf{x}_n + \omega_{j+1} \mathbf{v}_n + \mathbf{U}_n \mathbf{c}$$

In the calculation of the first residual in \mathcal{G}_{j+1} , we may choose ω_{j+1} freely, but the same value must be used in the calculations of the subsequent residuals in \mathcal{G}_{j+1} . A natural choice for ω_{j+1} is the value that minimizes the norm of \mathbf{r}_{n+1} , similarly as is done in, amongst others, the Bi-CGSTAB algorithm. Minimizing $\|\mathbf{r}_{n+1}\|_2 = \|\mathbf{v}_n - \omega_{j+1} \mathbf{A} \mathbf{v}_n\|_2$ yields

$$\omega_{j+1} = \frac{(\mathbf{A} \mathbf{v}_n)^H \mathbf{v}_n}{(\mathbf{A} \mathbf{v}_n)^H \mathbf{A} \mathbf{v}_n}.$$

In Sleijpen and van der Vorst [1995], Sleijpen and Van der Vorst propose in the context of Bi-CGSTAB an improvement on this choice. They explain that a small value for the minimal residual parameter ω can have a negative effect on the accuracy of the Bi-CG parameters, and as a consequence on the convergence of Bi-CGSTAB. As a possible cure to this they propose to use not a pure minimal residual step, but to increase the value of ω if the cosine of angle between $\mathbf{A} \mathbf{v}_n$ and \mathbf{v}_n is smaller than a threshold κ . This means that ω is increased if these vectors are too close to being perpendicular. A similar approach can be applied to the IDR(s) algorithm. In the setting of this algorithm the computation of ω_{j+1} according to the strategy of Sleijpen and Van der Vorst [1995] becomes

$$\begin{aligned} \omega_{j+1} &= \frac{(\mathbf{A} \mathbf{v}_n)^H \mathbf{v}_n}{(\mathbf{A} \mathbf{v}_n)^H \mathbf{A} \mathbf{v}_n}, \quad \rho = \frac{(\mathbf{A} \mathbf{v}_n)^H \mathbf{v}_n}{\|\mathbf{A} \mathbf{v}_n\| \|\mathbf{v}_n\|} \\ \text{IF } |\rho| &< \kappa \\ \omega_{j+1} &= \omega_{j+1} \kappa / |\rho| \\ \text{ENDIF} \end{aligned}$$

Sleijpen and Van der Vorst recommend 0.7 as a suitable value for κ . In our experience, this “maintaining the convergence” choice for ω can greatly improve the rate of convergence of IDR(s); see Section 6.4 of Sonneveld and van Gijzen [2008] for an example. In Section 4 we will illustrate through numerical experiments that this choice for ω may also yield a more accurate solution, in particular for larger values of s . Note that the calculation of ω_{j+1} does not require an additional matrix multiplication, since the vector $\mathbf{A}v_n$ can be reused in the update of the residual.

Of course, other than the previously discussed strategies for the computation of ω_{j+1} are possible; see, for example, Simoncini and Szyld [2010] for a new approach based on precomputed Ritz values.

The preceding framework explains how a residual in \mathcal{G}_{j+1} can be computed given $\mathbf{r}_n, \mathbf{g}_{n-i} \in \mathcal{G}_j, i = 1, \dots, s$. Next we will discuss a technique for computing these vectors.

2.4. Computing Additional Vectors in \mathcal{G}_{j+1}

The procedure that is outlined in the previous section can be used directly to compute a new residual $\mathbf{r}_{n+2} \in \mathcal{G}_{j+1}$: since $\mathbf{g}_i \in \mathcal{G}_j, i = n-1, \dots, n-s$ and $\mathbf{r}_{n+1} \in \mathcal{G}_{j+1} \subset \mathcal{G}_j$, the computations

$$\begin{aligned} \text{Calculate: } \mathbf{c} &\in \mathbb{C}^s \text{ from } (\mathbf{P}^H \mathbf{G}_n) \mathbf{c} = \mathbf{P}^H \mathbf{r}_{n+1}, \\ \mathbf{v}_{n+1} &= \mathbf{r}_{n+1} - \mathbf{G}_n \mathbf{c}, \\ \mathbf{r}_{n+2} &= \mathbf{v}_{n+1} - \omega_{j+1} \mathbf{A} \mathbf{v}_{n+1} \end{aligned}$$

yield a residual that satisfies $\mathbf{r}_{n+2} \in \mathcal{G}_{j+1}$.

Furthermore, we observe that the *residual difference vector* $(\mathbf{r}_{n+2} - \mathbf{r}_{n+1})$ is in the space \mathcal{G}_{j+1} . Since $\mathbf{A}^{-1}(\mathbf{r}_{n+2} - \mathbf{r}_{n+1}) = -(\mathbf{x}_{n+2} - \mathbf{x}_{n+1})$ we have found a suitable pair of vectors $\mathbf{g}_{n+1}, \mathbf{u}_{n+1}$.

$$\mathbf{g}_{n+1} = -(\mathbf{r}_{n+2} - \mathbf{r}_{n+1}), \quad \mathbf{u}_{n+1} = \mathbf{x}_{n+2} - \mathbf{x}_{n+1}$$

In a practical algorithm, the computation of \mathbf{g}_{n+1} and of \mathbf{u}_{n+1} precedes the computation of \mathbf{r}_{n+2} and of \mathbf{x}_{n+2} . First, the update vector for the iterate can be computed by

$$\mathbf{u}_{n+1} = \omega_{j+1} \mathbf{v}_{n+1} + \mathbf{U}_n \mathbf{c},$$

followed by the computation of \mathbf{g}_{n+1} by

$$\mathbf{g}_{n+1} = \mathbf{A} \mathbf{u}_{n+1} \tag{8}$$

to preserve in finite precision arithmetic as much as possible the relation between \mathbf{u}_{n+1} and \mathbf{g}_{n+1} . The iterate and residual are then updated through

$$\mathbf{x}_{n+2} = \mathbf{x}_{n+1} + \mathbf{u}_{n+1}, \quad \mathbf{r}_{n+2} = \mathbf{r}_{n+1} - \mathbf{g}_{n+1}. \tag{9}$$

The vector \mathbf{g}_{n+1} is in the space \mathcal{G}_{j+1} , and hence also in \mathcal{G}_j . This means that we can use this vector in the calculation of new vectors in \mathcal{G}_{j+1} , and discard an old vector, for example, \mathbf{g}_{n-s} . This can be done by defining the matrices \mathbf{G}_{n+2} and \mathbf{U}_{n+2} as

$$\mathbf{G}_{n+2} = (\mathbf{g}_{n+1} \ \mathbf{g}_{n-s+1} \ \dots \ \mathbf{g}_{n-1}), \tag{10}$$

$$\mathbf{U}_{n+2} = (\mathbf{u}_{n+1} \ \mathbf{u}_{n-s+1} \ \dots \ \mathbf{u}_{n-1}). \tag{11}$$

The advantage of this procedure is that it saves vector space: storage for exactly s \mathbf{g} -vectors and s \mathbf{u} -vectors is needed.

We can repeat the preceding procedure s times to compute $\mathbf{r}_{n+s+1}, \mathbf{g}_{n+k} \in \mathcal{G}_{j+1}, k = 1, \dots, s$, and the corresponding vectors $\mathbf{x}_{n+s+1}, \mathbf{u}_{n+k}, k = 1, \dots, s$, which are the vectors that are needed to compute a residual in \mathcal{G}_{j+2} .

The aforesaid relations define (apart from initialization of the vectors) a complete IDR-method. In fact, the algorithm that is outlined before is almost the same as the

IDR(s) method from Sonneveld and van Gijzen [2008]. The only difference is that the original IDR(s) method also computes $\mathbf{g}_n = -(\mathbf{r}_{n+1} - \mathbf{r}_n)$, which vector is then included in \mathbf{G}_{n+k} , $k = 1, \dots, s$. Leaving this vector out simplifies the IDR(s)-biortho algorithm that we present in the next section.

In the previous algorithm, vectors in \mathcal{G}_{j+1} are generated by direct application of the IDR theorem. The computations of the first residual in \mathcal{G}_{j+1} are almost the same as the computation of the following s residuals in \mathcal{G}_{j+1} . However, in computing the intermediate residuals, there is more freedom that can be exploited. In the preceding algorithm, a residual is updated by

$$\mathbf{r}_{n+k+1} = \mathbf{r}_{n+k} - \mathbf{g}_{n+k}.$$

Here, \mathbf{r}_{n+k+1} , \mathbf{r}_{n+k} , and \mathbf{g}_{n+k} are in \mathcal{G}_{j+1} . But in order to compute a new residual in \mathcal{G}_{j+1} we could also have used a more general linear combination of vectors in \mathcal{G}_{j+1} .

$$\mathbf{r}_{n+k+1} = \mathbf{r}_{n+k} - \sum_{i=1}^k \beta_i \mathbf{g}_{n+i}$$

Clearly, the vector \mathbf{r}_{n+k+1} computed in this way is also in \mathcal{G}_{j+1} . We can choose the parameters β_i to give the intermediate residuals a desirable property, like minimum norm. In the algorithm that we present in the next section we will use the parameters β_i such that the intermediate residual \mathbf{r}_{n+k+1} is orthogonal to $\mathbf{p}_1, \dots, \mathbf{p}_k$.

The same freedom that we have for computing a new residual in \mathcal{G}_{j+1} , we have for computing the vectors \mathbf{g}_{n+k} : linear combinations of vectors in \mathcal{G}_{j+1} are still in \mathcal{G}_{j+1} . Let

$$\bar{\mathbf{g}} = -(\mathbf{r}_{n+k+1} - \mathbf{r}_{n+k}).$$

Then the vector

$$\mathbf{g}_{n+k} = \bar{\mathbf{g}} - \sum_{i=1}^{k-1} \alpha_i \mathbf{g}_{n+i}$$

is also in \mathcal{G}_{j+1} , and can be used in the subsequent computations. Again, the parameters α_i can be chosen such that the vector \mathbf{g}_{n+k} gets some favorable properties. In the algorithm that we present in the next section we will chose the parameters α_i such that the vector \mathbf{g}_{n+k} is made orthogonal to $\mathbf{p}_1 \dots \mathbf{p}_{k-1}$.

Apart from the initialization of the variables, we have now given a complete framework for an IDR-based solution algorithm. To initialize the recursions, values for \mathbf{x}_s , \mathbf{r}_s , \mathbf{U}_0 , and \mathbf{G}_0 have to be computed. This can be done by any Krylov method. Figure 1 presents a framework, including an (implicit) initialization in the first s steps, for an IDR-based algorithm. The freedom that is left open is in the choice of the parameters α_i and β_i , and ω . In the algorithm we have omitted the indices for the iteration number. Vectors on the left are overwritten by vectors on the right.

3. AN EFFICIENT IDR(S) VARIANT THAT EXPLOITS BIORTHOGONALITY PROPERTIES

3.1. General Idea

In this section we will fill in the freedom that we have left in the framework IDR algorithm. As in the previous section we assume that \mathbf{r}_{n+1} is the first residuals in \mathcal{G}_{j+1} . We fill in the freedom by constructing vectors that satisfy the following biorthogonality conditions:

$$\mathbf{g}_{n+k} \perp \mathbf{p}_i, \quad i = 1, \dots, k-1, \quad k = 2, \dots, s, \quad (12)$$

and

$$\mathbf{r}_{n+k+1} \perp \mathbf{p}_i, \quad i = 1, \dots, k, \quad k = 1, \dots, s. \quad (13)$$

```

Require:  $A \in \mathbb{C}^{N \times N}$ ;  $x, b \in \mathbb{C}^N$ ;  $P \in \mathbb{C}^{N \times s}$ ;  $TOL \in (0, 1)$ ;
Ensure:  $x$  such that  $\|b - Ax\| \leq TOL \cdot \|b\|$ 
  {Initialization.}
  Calculate  $r = b - Ax$ ;
   $G = O \in \mathbb{C}^{N \times s}$ ;  $U = O \in \mathbb{C}^{N \times s}$ ;
   $M = I \in \mathbb{C}^{s \times s}$ ;  $\omega = 1$ ;

  {Loop over  $\mathcal{G}_j$  spaces}
  while  $\|r\| > TOL$  do
    {Compute  $s$  independent vectors  $g_k$  in  $\mathcal{G}_j$  space}
    for  $k = 1$  to  $s$  do
      Compute  $f = P^H r$ ;
      {Note:  $M = P^H G$ }
      Solve  $c$  from  $Mc = f$ ;
       $v = r - Gc$ ;
       $u_k = Uc + \omega v$ ;
       $g_k = Au_k$ ;
      {Linear combinations of vectors  $\in \mathcal{G}_j$  are still in  $\mathcal{G}_j$ :}
      Select  $\alpha_i$  and  $\beta_i$ ,  $i = 1, \dots, k-1$ ;
       $g_k = g_k - \sum_{i=1}^{k-1} \alpha_i g_i$ ;  $u_k = u_k - \sum_{i=1}^{k-1} \alpha_i u_i$ ;
       $r = r - \sum_{i=1}^k \beta_i g_i$ ;  $x = x + \sum_{i=1}^k \beta_i u_i$ ;
      {Update of the  $k$ -th column of  $M$ :}
       $M_{:,k} = P^H g_k$ ;
      {Overwrite  $k$ -th column of  $G$  by  $g_k$ , and of  $U$  by  $u_k$ }
       $G_{:,k} = g_k$ ;  $U_{:,k} = u_k$ ;
    end for
    {Entering  $\mathcal{G}_{j+1}$ }
     $f = P^H r$ ;
    Solve  $c$  from  $Mc = f$ ;
     $v = r - Gc$ ;
     $t = Av$ ;
    Select  $\omega$ ;
     $x = x + Uc + \omega v$ ;
    {Alternative computation:  $r = v - \omega t$ };
     $r = r - Gc - \omega t$ ;
  end while

```

Fig. 1. A framework for an IDR-based algorithm.

As we will see, these relations lead to important simplifications in the algorithm.

3.2. A Dimension-Reduction Step: Computing the First Residual in \mathcal{G}_{j+1}

The biorthogonality condition for the intermediate residuals (13) implies that the first intermediate residual is orthogonal to \mathbf{p}_1 , the second to \mathbf{p}_1 and to \mathbf{p}_2 , etc. Hence, the last intermediate residual before making a dimension reduction step, that is, \mathbf{r}_n is orthogonal to $\mathbf{p}_1, \dots, \mathbf{p}_s$. Consequently,

$$\mathbf{r}_n \in \mathcal{G}_j \cap \mathcal{S}.$$

Now, by Theorem 2.1

$$\mathbf{r}_{n+1} = (\mathbf{I} - \omega_{j+1} \mathbf{A}) \mathbf{r}_n \in \mathcal{G}_{j+1}.$$

With the standard choice for ω_{j+1} , the dimension reduction step simplifies to a standard minimal residual step.

3.3. Computing Additional Vectors in \mathcal{G}_{j+1}

In order to calculate a vector $\mathbf{v}_{n+k} \in \mathcal{G}_j \cap \mathcal{S}$, a system of the form $(\mathbf{P}^H \mathbf{G}_{n+k})\mathbf{c} = \mathbf{P}^H \mathbf{r}_{n+k}$ has to be solved. Using the conditions (12) and (13) this system gets a simple form. Let

$$\mu_{i,k} = \mathbf{p}_i^H \mathbf{g}_{n+k}, \quad i = 1, \dots, s.$$

Then, because of (12), $\mu_{ik} = 0$ for $i < k$. Furthermore, let

$$\phi_i = \mathbf{p}_i^H \mathbf{r}_{n+k}, \quad i = 1, \dots, s.$$

Then, because of (13), $\phi_i = 0$ for $i < k$. Consequently, the system $(\mathbf{P}^H \mathbf{G}_{n+k})\mathbf{c} = \mathbf{P}^H \mathbf{r}_{n+k}$ has the following structure.

$$\begin{pmatrix} \mu_{1,1} & 0 & \dots & \dots & 0 \\ \mu_{2,1} & \mu_{2,2} & \ddots & & \vdots \\ \vdots & \vdots & \ddots & \ddots & \vdots \\ \vdots & \vdots & & \ddots & 0 \\ \mu_{s,1} & \mu_{s,2} & \dots & \dots & \mu_{s,s} \end{pmatrix} \begin{pmatrix} \gamma_1 \\ \vdots \\ \vdots \\ \vdots \\ \gamma_s \end{pmatrix} = \begin{pmatrix} 0 \\ \vdots \\ \phi_k \\ \vdots \\ \phi_s \end{pmatrix}$$

Clearly, $\gamma_1, \dots, \gamma_{k-1}$ are zero, and the update for \mathbf{v}_{n+k} becomes

$$\mathbf{v}_{n+k} = \mathbf{r}_{n+k} - \sum_{i=k}^s \gamma_i \mathbf{g}_{n+i-s-1}.$$

Next, we compute “temporary” vectors \mathbf{u}_{n+k} and \mathbf{g}_{n+k} by

$$\mathbf{u}_{n+k} = \omega_{j+1} \mathbf{v}_{n+k} + \sum_{i=k}^s \gamma_i \mathbf{u}_{n+i-s-1}; \quad \mathbf{g}_{n+k} = \mathbf{A} \mathbf{u}_{n+k}.$$

Next, the vector $\mathbf{g}_{n+k} \in \mathcal{G}_{j+1}$ is made orthogonal to $\mathbf{p}_1, \dots, \mathbf{p}_{k-1}$ by the following procedure, that includes the corresponding updates to compute \mathbf{u}_{n+k} .

For $i = 1$ to $k-1$

$$\begin{aligned} \alpha &= \frac{\mathbf{p}_i^H \mathbf{g}_{n+k}}{\mu_{i,i}}; \\ \mathbf{g}_{n+k} &= \mathbf{g}_{n+k} - \alpha \mathbf{g}_{n+i}; \\ \mathbf{u}_{n+k} &= \mathbf{u}_{n+k} - \alpha \mathbf{u}_{n+i}. \end{aligned}$$

End for

The preceding algorithm is similar to the modified Gram-Schmidt procedure for orthogonalizing a vector with respect to a set of orthogonal vectors. Alternatively, we could have used the classical Gram-Schmidt-like procedure as we have used for the computation of a vector \mathbf{v}_{n+k} that is orthogonal to $\mathbf{p}_1, \dots, \mathbf{p}_s$.

The next step in the algorithm is to compute the next intermediate residual \mathbf{r}_{n+k+1} that is orthogonal to $\mathbf{p}_1, \dots, \mathbf{p}_k$. It is easy to check that such a residual can be computed by

$$\mathbf{r}_{n+k+1} = \mathbf{r}_{n+k} - \frac{\phi_k}{\mu_{k,k}} \mathbf{g}_{n+k}. \quad (14)$$

The corresponding approximate solution then becomes

$$\mathbf{x}_{n+k+1} = \mathbf{x}_{n+k} + \frac{\phi_k}{\mu_{k,k}} \mathbf{u}_{n+k}.$$

The outline of the algorithm suggests that we have to compute the inner products as $\phi_i = \mathbf{p}_i^H \mathbf{r}_{n+k+1}$, $i = k, \dots, s$. From (14), however, it follows that

$$\mathbf{p}_i^H \mathbf{r}_{n+k+1} = \mathbf{p}_i^H \mathbf{r}_{n+k} - \frac{\phi_k}{\mu_{k,k}} \mathbf{p}_i \mathbf{g}_{n+k} \quad i = k+1, \dots, s,$$

and hence that

$$\phi_i = \phi_i - \frac{\phi_k \mu_{i,k}}{\mu_{k,k}}, \quad i = k+1, \dots, s.$$

So given $\mu_{i,k}$, $i = k+1, \dots, s$, the new value for ϕ_i can be computed via a scalar update.

3.4. A Preconditioned Version of IDR(s)-Biortho

Preconditioning can simply be implemented by applying the unpreconditioned algorithm to an explicitly preconditioned problem. In the case of right preconditioning with preconditioner \mathbf{B} this means that IDR(s) is applied to the system

$$\mathbf{A}\mathbf{B}^{-1}\mathbf{y} = \mathbf{b}$$

and every matrix-vector multiplication becomes a multiplication with $\mathbf{A}\mathbf{B}^{-1}$, that is, an application of the preconditioner followed by a multiplication with \mathbf{A} . In order to get the solution of the problem $\mathbf{A}\mathbf{x} = \mathbf{b}$, we have to scale back the solution

$$\mathbf{x} = \mathbf{B}^{-1}\mathbf{y},$$

so one extra preconditioning step has to be performed after the iterative process has terminated.

As is shown in Onoue et al. [2008] for IDR(s)-proto, it is possible to avoid this extra operation and to make preconditioning implicit. The same ideas that are described in Onoue et al. [2008] can also be applied to IDR(s)-biortho.

The idea is to slightly rearrange the update of the solution. For example, the dimension-reduction step for the explicitly right-preconditioned problem reads

$$\begin{aligned} \mathbf{t} &= \mathbf{A}\mathbf{B}^{-1}\mathbf{r}_n, \\ \mathbf{r}_{n+1} &= \mathbf{r}_n - \omega_{j+1}\mathbf{t}, \\ \mathbf{y}_{n+1} &= \mathbf{y}_n + \omega_{j+1}\mathbf{r}_n. \end{aligned}$$

Multiplying the recursion for \mathbf{y} with \mathbf{B}^{-1} gives the recursion for \mathbf{x} :

$$\begin{aligned} \mathbf{t} &= \mathbf{A}\mathbf{B}^{-1}\mathbf{r}_n, \\ \mathbf{r}_{n+1} &= \mathbf{r}_n - \omega_{j+1}\mathbf{t}, \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + \omega_{j+1}\mathbf{B}^{-1}\mathbf{r}_n, \end{aligned}$$

which can be implemented as

$$\begin{aligned} \mathbf{v} &= \mathbf{B}^{-1}\mathbf{r}_n, \\ \mathbf{t} &= \mathbf{A}\mathbf{v}, \\ \mathbf{r}_{n+1} &= \mathbf{r}_n - \omega_{j+1}\mathbf{t}, \\ \mathbf{x}_{n+1} &= \mathbf{x}_n + \omega_{j+1}\mathbf{v}. \end{aligned}$$

The same technique can be used to make the preconditioning operation in the intermediate steps implicit.

```

Require:  $A \in \mathbb{C}^{N \times N}$ ;  $x, b \in \mathbb{C}^N$ ;  $P \in \mathbb{C}^{N \times s}$ ;  $TOL \in (0, 1)$ ;
Ensure:  $x_n$  such that  $\|b - Ax\| \leq TOL \cdot \|b\|$ ;
  {Initialization}
  Calculate  $r = b - Ax$ ;
   $g_i = u_i = 0 \in \mathbb{C}^N, i = 1, \dots, s$ ;  $M = I \in \mathbb{C}^{s \times s}$ ;  $\omega = 1$ ;

  {Loop over  $\mathcal{G}_j$  spaces, for  $j = 0, 1, 2, 3, \dots$ }
  while  $\|r\| > TOL$  do
     $f = P^H r, (\phi_1, \dots, \phi_s)^T = f$ ;
    for  $k = 1$  to  $s$  do
      Solve  $c$  from  $Mc = f, (\gamma_1, \dots, \gamma_s)^T = c$ ;
       $v = r - \sum_{i=k}^s \gamma_i g_i$ ;
      {Preconditioning operation}
       $v = B^{-1}v$ ;
       $u_k = \omega v + \sum_{i=k}^s \gamma_i u_i$ ;
       $g_k = Au_k$ ;
      for  $i = 1$  to  $k - 1$  do
         $\alpha = \frac{p_i^H g_k}{\mu_{i,i}}$ ;
         $g_k = g_k - \alpha g_i$ ;
         $u_k = u_k - \alpha u_i$ ;
      end for
       $\mu_{i,k} = p_i^H g_k, M_{i,k} = \mu_{i,k}, i = k \dots s$ ;
       $\beta = \frac{\phi_k}{\mu_{k,k}}$ ;
       $r = r - \beta g_k$ ;
       $x = x + \beta u_k$ ;
      if  $k + 1 \leq s$  then
         $\phi_i = 0, i = 1, \dots, k$ ;
         $\phi_i = \phi_i - \beta \mu_{i,k}, i = k + 1, \dots, s$ ;
         $f = (\phi_1, \dots, \phi_s)^T$ ;
      end if
    end for
    {Entering  $\mathcal{G}_{j+1}$ }
    {Preconditioning}
     $v = B^{-1}r$ ;
     $t = Av$ ;
    {Calculation of  $\omega$  using “maintaining the convergence” strategy}
     $\omega = t^H r / t^H t$ ;  $\rho = (t^H r) / (\|t\| \|r\|)$ ;
    if  $|\rho| < \kappa$  then
       $\omega = \omega \kappa / |\rho|$ ;
    end if
     $r = r - \omega t$ ;
     $x = x + \omega v$ ;
  end while

```

Fig. 2. Preconditioned IDR(s)-biortho, the preconditioner is denoted by B .

3.5. The IDR(s)-Biortho Algorithm

Figure 2 presents the complete IDR(s)-biortho algorithm, including preconditioning and the computation of ω_{j+1} according to the “maintaining the convergence” strategy. In the algorithm we have omitted the indices for the iteration number.

Table I. Memory Requirements and Vector Operations per Preconditioned Matrix-Vector Product

Method	DOT	AXPY	Memory Requirements (vectors)
IDR(1)	2	3	7
IDR(2)	$2\frac{2}{3}$	$4\frac{2}{3}$	10
IDR(4)	$4\frac{2}{3}$	$8\frac{2}{3}$	16
IDR(8)	$8\frac{2}{3}$	$16\frac{2}{3}$	28
GMRES	$\frac{n+1}{2}$	$\frac{n+1}{2}$	$n+3$
Bi-CG	1	$2\frac{1}{2}$	7
QMR	1	4	13
CGS	1	3	8
Bi-CGSTAB	2	3	7
BiCGstab(2)	$2\frac{1}{4}$	$3\frac{3}{4}$	9

The algorithm is quite efficient in terms of vector operations and even more efficient than IDR(s)-proto, despite the additional orthogonalization operations. The operation count for the main operations to perform the dimension reduction step yields: one preconditioned matrix-vector product, two vector updates, and two inner products. For the intermediate steps we get: s preconditioned matrix-vector products, $2s+2$ vector updates, and $s+1$ inner products. Hence, for a full cycle of $s+1$ IDR(s) steps we get: $s+1$ preconditioned matrix-vector products, s^2+s+2 inner products, and $2s^2+2s+2$ vector updates. IDR(s)-biortho requires slightly less vector updates than IDR(s)-proto, and the same number of inner products and preconditioned matrix-vector multiplications. The original IDR(s) method requires $2s^2 + \frac{7}{2}s + \frac{5}{2}$ vector-updates.

Table I gives an overview of the number of vector operations per preconditioned matrix-vector multiplication for some values of s for IDR(s)-biortho, and for comparison also for the Krylov methods that we will use in the numerical experiments. This table also gives the memory requirements (excluding storage of the system matrix and of the preconditioner, but including storage for the right-hand-side vector and the solution).

4. NUMERICAL EXPERIMENTS

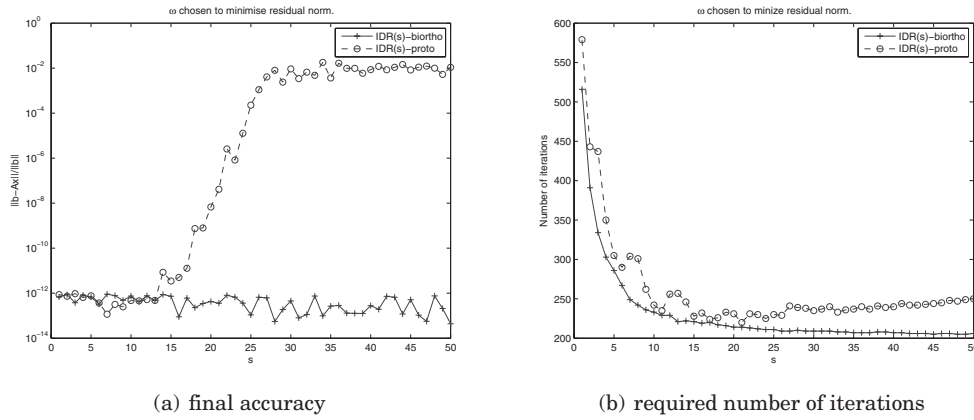
In this section we present numerical examples to compare the numerical behavior of IDR(s)-proto and IDR(s)-biortho. For another evaluation of the performance of the two IDR(s) variants, based on selected test matrices from the Florida Sparse Matrix Collection [Davis and Hu 2010], we refer to Onoue et al. [2009]. We also present experiments to evaluate the performance of the biortho variant of IDR(s) in comparison with a number of state-of-the-art Krylov methods. Other performance comparisons can be found in Sonneveld and van Gijzen [2008], van Gijzen and Sonneveld [2008], and in Onoue et al. [2008]; the latter reference uses a test set from the University of Florida Sparse Matrix Collection. In these three references the IDR(s)-proto variant is used.

The experiments that are presented in this section have been performed on a standard desktop computer with an Intel Core 2 duo processor and 4Gb of RAM using MatLab 7.5.

In all our experiments we take for $\mathbf{p}_1, \dots, \mathbf{p}_s$ the orthogonalization of s normally distributed random vectors, with mean 0 and standard deviation 1.

4.1. Mathematical Equivalence of IDR(s)-Proto and IDR(s)-Biortho

The first numerical example validates that IDR(s)-proto and IDR(s)-biortho (in exact arithmetic) yield the same residual at every $(s+1)$ -st iteration. This property is shown to be true under mild conditions in Sonneveld and van Gijzen [2008, Section 5.1]. To investigate this numerically we consider the the ADD20 matrix and corresponding right-hand-side vector from the MATRIX MARKET collection. The parameter ω_j is

Fig. 4. Accuracy and number of iterations for the two IDR(s) variants with “minimal residual” choice for ω .Table II. Sherman 2: Iterations and Final Accuracy for Increasing s

Method	Iterations	$\ b - Ax\ /\ b\ $	Method	Iterations	$\ b - Ax\ /\ b\ $
IDR(20)-proto	2160	$3.5 \cdot 10^{-2}$	IDR(20)-biortho	2160	$1.5 \cdot 10^{-3}$
IDR(40)-proto	1501	$1.5 \cdot 10^{-1}$	IDR(40)-biortho	969	$8.2 \cdot 10^{-5}$
IDR(60)-proto	1791	$6.6 \cdot 10^{-2}$	IDR(60)-biortho	945	$9.8 \cdot 10^{-5}$
IDR(80)-proto	1684	$5.1 \cdot 10^{-1}$	IDR(80)-biortho	883	$9.8 \cdot 10^{-5}$
IDR(100)-proto	2145	$9.3 \cdot 10^{-1}$	IDR(100)-biortho	536	$9.8 \cdot 10^{-5}$
IDR(120)-proto	1480	$5.7 \cdot 10^{-1}$	IDR(120)-biortho	473	$8.2 \cdot 10^{-5}$
IDR(140)-proto	2028	$1.7 \cdot 10^{-1}$	IDR(140)-biortho	166	$7.6 \cdot 10^{-5}$

in the new variant clearly improve the accuracy of the algorithm: the required accuracy for IDR(s)-biortho is always achieved.

Figure 4(b) shows the number of iterations to achieve the required tolerance. We have repeated the same experiments with the “maintaining the convergence” choice for ω_j . The most striking difference with the experiments with the “minimum residual choice” for ω_j is that the IDR(s)-proto algorithm remains accurate for much larger s . Only for s larger than 27 the accuracy is lower than expected, whereas for the minimum residual choice for ω_j the accuracy is lower than the expected level for s larger than 12. Again the final accuracy for the IDR(s)-biortho algorithm always satisfies the required tolerance.

This figure shows that for s small, say up to $s = 10$, the number of iterations drops significantly with increasing s . However, for larger values of s no gain can be obtained. This observation, that we have often made for reasonably well-conditioned problems, has recently been explained in Sonneveld [2010].

4.3. A Problem that Requires a Large Value for s

For most problems, it is for efficiency reasons advisable to take a small value for s , for example, $s = 4$. Problems with an ill-conditioned system matrix, however, may require a large value for s . The following example, the matrix SHERMAN2 from the MATRIX-MARKET collection with corresponding right-hand-side vector, is such a problem. We solve this problem with the two IDR(s) variants, with values of s ranging from 20 to 140. The required tolerance is $\|r_i\|/\|b\| < 10^{-4}$. The norm of the recursively updated residual is used in the termination criterion. Table II gives for the two IDR(s) variants the number of iterations and the accuracy that is reached (the norm of the true residual divided by the norm of the right-hand-side vector b) for the different values of s . The

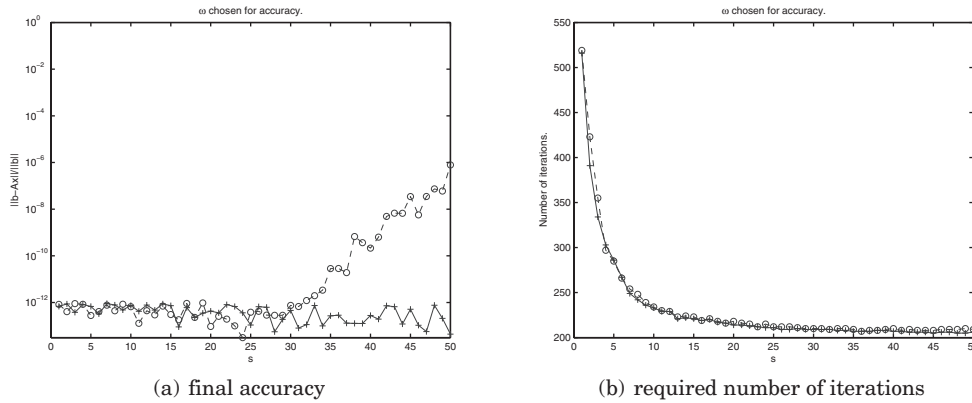


Fig. 5. Accuracy and number of iterations for the two IDR(s) variants with “maintaining the convergence” choice for ω .

maximum number of iterations is set to 2160, which is two times the problem size of 1080.

The table shows that IDR(s)-proto never achieves the required accuracy. IDR(s)-biortho, on the other hand, satisfies the accuracy for all the tested values for s , except for $s = 20$. We remark that the MatLab built-in methods Bi-CGSTAB, QMR, Bi-CG, and CGS do not converge for this problem. Full GMRES converges after 119 iterations.

4.4. A Problem for which IDR(s)-Proto is Inaccurate for Small s

The next example that we consider shows a rare case where IDR(s)-proto does not converge for small s , whereas IDR(s)-biortho converges as expected. The example is an extreme case where IDR(s)-proto stagnates at a certain level, after which it gradually starts to diverge. We have observed this phenomenon more often if very high accuracy (close to machine precision) is required. In this example, however, stagnation of IDR(1)-proto and IDR(2)-proto occurs at a much higher level.

The example is the finite difference discretization of the following convection-diffusion-reaction equation with homogeneous Dirichlet boundary conditions on the unit cube.

$$-\epsilon \Delta u + \vec{\beta} \cdot \nabla u - ru = F$$

The right-hand-side vector F is defined by the solution $u(x, y, z) = x(1-x)y(1-y)z(1-z)$. The problem is discretized using central differences with a grid size $h = 0.025$. The resulting linear system consists of approximately 60,000 equations. We take the following values for the parameters: $\epsilon = 1$ (diffusion), $\vec{\beta} = (0/\sqrt{5} \ 250/\sqrt{5} \ 500/\sqrt{5})^T$ (convection), and $r = 400$ (reaction). The resulting matrix is highly nonsymmetric and indefinite, properties that make the resulting system difficult to solve with an iterative solver.

The ω_j 's are again computed using the “maintaining the convergence” strategy (Figure 5). The required tolerance is $\|r_i\|/\|b\| < 10^{-10}$. The maximum number of matrix-vector multiplications is set to 650.

Figure 6(a) shows the convergence for IDR(s)-proto, for $s = 1, 2, 4$, and 8 , together with the convergence for GMRES (optimal). For comparison we have also included the convergence of Bi-CGSTAB (MatLab implementation). We remark that IDR(1) and the MatLab implementation of Bi-CGSTAB are not mathematically equivalent due to the different choice for p_1 (random in IDR(1) and $p_1 = r_0$ in Bi-CGSTAB), and because of the different strategy for computing ω_j (“maintaining the convergence” in IDR(1)

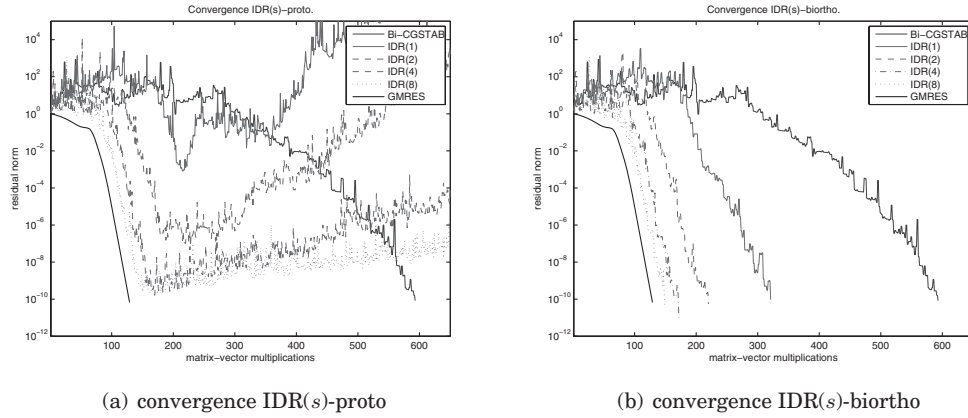


Fig. 6. Accuracy and number of iterations for the two IDR(s) variants, minimal residual choice for ω .

Table III. Convection-Diffusion-Reaction Problem: Iterations and Final Accuracy for Increasing s

Method	MATVECS	$\ \mathbf{b} - \mathbf{Ax}\ /\ \mathbf{b}\ $	Method	MATVECS	$\ \mathbf{b} - \mathbf{Ax}\ /\ \mathbf{b}\ $
IDR(1)-proto	650	$1.4 \cdot 10^{+6}$	IDR(1)-biortho	321	$9.3 \cdot 10^{-11}$
IDR(2)-proto	650	$1.7 \cdot 10^{+3}$	IDR(2)-biortho	220	$5.4 \cdot 10^{-11}$
IDR(4)-proto	650	$1.0 \cdot 10^{-4}$	IDR(4)-biortho	172	$8.6 \cdot 10^{-12}$
IDR(8)-proto	650	$1.8 \cdot 10^{-7}$	IDR(8)-biortho	149	$5.7 \cdot 10^{-11}$
GMRES	129	$6.6 \cdot 10^{-11}$	Bi-CGSTAB	593	$5.1 \cdot 10^{-11}$

and “minimum residual” in Bi-CGSTAB). These choices are standard for Bi-CGSTAB. As can be seen, the initial convergence of all IDR(s) variants is satisfactory, for larger s the convergence is close to the optimal GMRES convergence, and the convergence is significantly faster than for Bi-CGSTAB. However, the convergence stagnates for all IDR(s) variants above the required, quite strict tolerance. Once stagnation occurs, IDR(s)-proto gradually starts to diverge. For the higher values of s the required tolerance is almost achieved, but for small s the stagnation occurs at a level that is much higher than to be expected.

Figure 6(b) shows the convergence for IDR(s)-biortho. Now, no stagnation and subsequent divergence occurs. Also, the difference in rate of convergence with Bi-CGSTAB is quite striking. The convergence of IDR(4) is close to the optimal GMRES convergence, while the convergence of Bi-CGSTAB lags considerably behind. Such behavior is typical for difficult (nonsymmetric indefinite) problems.

Table III shows for all methods the different the number of matrix-vector multiplications methods and the final accuracy. The table shows that IDR(s)-biortho is accurate: the required tolerance is always achieved. In terms of matrix-vector multiplications, IDR(s)-biortho is much faster than Bi-CGSTAB for all values of s , and almost as fast as GMRES for $s = 4$ and $s = 8$. Note that we use full GMRES, which is optimal with respect to the number of matrix-vector multiplications, at the expense of a large overhead in vector operations and high storage requirements. The resulting CPU-times are also much lower for IDR(s)-biortho: 1.2s for IDR(4) against 10.6s for GMRES and 6.0s for Bi-CGSTAB. For completeness we also give the time for MatLab’s sparse direct solve (“\”), which takes 10.2s for this example.

4.5. A Performance Comparison for a Navier-Stokes Problem Using IFISS

IFISS is a MatLab open-source package associated with the book Elman et al. [2005] by Elman, Silvester and Wathen.² IFISS can be used to model a range of incompressible fluid flow problems and provides an ideal testing environment for iterative solvers and preconditioners.

We compare the performance of IDR(*s*)-biortho with Bi-CGstab(*ℓ*), and with full GMRES. We have performed our experiments with version 3.0 of IFISS, which has implementations of all three methods. We also report on results for Bi-CGSTAB, which were actually obtained for the mathematically equivalent method BiCGstab(1).

The test problem that we consider describes flow over a step (see Elman et al. [2005, Example 7.1.2]), which is modeled as a Navier-Stokes problem with zero forcing term. The steady-state Navier-Stokes equations are given by

$$\begin{aligned} -\eta \nabla^2 \mathbf{u} + \mathbf{u} \cdot \nabla \mathbf{u} + \nabla p &= 0, \\ \nabla \cdot \mathbf{u} &= 0, \end{aligned}$$

where $\eta > 0$ is a given constant called the kinematic viscosity. The domain is L-shaped. The x -coordinate ranges from -1 to 5 . The y -coordinate ranges from 0 to 1 for x between -1 and 0 , and between -1 and 1 elsewhere: there is a step in the domain at $x = 0$. A Poiseuille flow profile is imposed on the inflow boundary $x = -1$, $0 \leq y \leq 1$ and a zero velocity condition on the walls. The Neumann condition

$$\eta \frac{\partial u_x}{\partial x} - p = 0 \quad \frac{\partial u_y}{\partial x} = 0$$

is applied at the outflow boundary $x = 5$, $-1 \leq y \leq 1$. The problem is discretized with biquadratic Q_2 elements for the velocities and bilinear Q_1 elements for the pressures. The resulting nonlinear system can be solved with Newton's method, which implies that in every iteration a linear system has to be solved to compute the Newton updates. This system has the following form.

$$\begin{pmatrix} \mathbf{F} & \mathbf{B}^T \\ \mathbf{B} & \mathbf{O} \end{pmatrix} \begin{pmatrix} \Delta \mathbf{u} \\ \Delta p \end{pmatrix} = \begin{pmatrix} \mathbf{f} \\ \mathbf{g} \end{pmatrix}$$

Here, the submatrix \mathbf{F} is nonsymmetric, and becomes increasingly more nonsymmetric if η is decreased.

As a test problem we consider the linear system after one Newton iteration. A block-triangular preconditioner of the form

$$\begin{pmatrix} \mathbf{F} & \mathbf{B}^T \\ \mathbf{O} & -\mathbf{M}_s \end{pmatrix}$$

is applied to speed-up the convergence of the iterative methods. Here, \mathbf{M}_s is an approximation to the Schur complement $\mathbf{S} = \mathbf{B}\mathbf{F}^{-1}\mathbf{B}^T$. The specific preconditioner we have selected for our experiments is the modified pressure-correction preconditioner [Elman et al. 2005]. Each application of this preconditioner requires three solves of subsystems: one solve with \mathbf{F} and two solves with the approximate Schur complement \mathbf{M}_s . These solves are approximately done with one sweep of the AMG solver of IFISS.

In the numerical experiments, we have taken a mesh size $h = 2^{-6}$, which yields a system of 102,659 equations. We have performed experiments with increasing Reynolds numbers Re . The Reynolds number is related to the kinematic viscosity by $Re = 2/\eta$. All systems are solved to a tolerance (= reduction of the residuals norm) of 10^{-6} .

²The open source-code is described in Elman et al. [2007], and can be downloaded from <http://www.manchester.ac.uk/ifiss> and <http://www.cs.umd/elman/ifiss.html>.

Table IV. Navier-Stokes Test Problem: Matrix-Vector Multiplications and Computing Times

Re	GMRES	Bi-CGSTAB	BiCGstab(2)	IDR(4)	IDR(8)
100	59 (21.8s)	106 (36.8s)	92 (32.0s)	72 (24.9s)	71 (25.1s)
200	76 (31.9s)	142 (57.2s)	144 (56.2s)	97 (37.8s)	92 (37.1s)
400	114 (52.2s)	280 (115.9s)	272 (113.6s)	156 (65.1s)	143 (60.6s)
800	181 (100.0s)	486 (284.6s)	652 (317.7s)	284 (138.6s)	231 (114.3s)

Table V. Atmospheric Model Problems: Matrix-Vector Multiplications and Computing Times

Iterative\Matrix Method\Size	ATMOSMODD $N = 1270432$	ATMOSMODJ $N = 1270432$	ATMOSMODL $N = 1489752$	ATMOSMODM $N = 1489752$
GMRES(50)	690 (980s)	1069 (1515s)	64 (95s)	38 (50s)
Bi-CGSTAB	479 (145s)	439 (131s)	89 (35s)	78 (28s)
CGS	No convergence	No convergence	No convergence	No convergence
Bi-CG	448 (79s)	430 (75s)	112 (24s)	76 (16s)
QMR	446 (94s)	436 (91s)	112 (28s)	76 (19s)
IDR(1)	520 (45s)	550 (48s)	106 (11s)	78 (9s)
IDR(2)	351 (40s)	311 (36s)	83 (12s)	59 (8s)
IDR(4)	268 (48s)	267 (48s)	69 (15s)	49 (11s)
IDR(8)	247 (80s)	248 (80s)	63 (25s)	43 (18s)

Table IV gives the numbers of matrix-vector multiplications, and in between brackets the computing times. In comparison with Bi-CGSTAB and BiCGstab(2), IDR(4), and especially IDR(8) are considerably faster, in particular for large Reynolds numbers. The difference in solution time for the higher Reynolds numbers is a factor of two to three.

The preconditioner described earlier is quite effective in reducing the number of iterations, but is expensive to apply. This makes the preconditioned matrix-vector multiplication expensive. As a result, the time per iteration is basically determined by the preconditioned matrix-vector multiplication, and overhead for vector operations is small compared to the cost of the preconditioned matrix-vector multiplications. This situation is particularly advantageous for GMRES, since this method gives an optimal reduction of the residual norm for a given number of iterations (= preconditioned matrix-vector multiplications).

The computing times for GMRES are slightly lower than for IDR(s). However, this comes at the price of a much higher memory consumption: for every iteration an extra vector of dimension N has to be stored, whereas the storage costs for IDR(s) are fixed. For comparison: for the test problem with $Re = 800$, GMRES requires the storage of 184 vectors, IDR(4) of 16 vectors, and IDR(8) of 28.

4.6. A Performance Comparison for Large Atmospheric Test Problems Using Built-In MatLab Routines

The final set of test problems is part of the University of Florida Sparse Matrix Collection [Davis and Hu 2010]. The four test problems come from numerical weather prediction and (three-dimensional) atmospheric modeling. The systems are large, in excess of a million of equations, and have a block tridiagonal structure with seven nonzero diagonals. The matrices are mildly nonsymmetric, and strictly diagonally dominant. As a result of these properties, the systems are well suited for iterative solution methods.

In the experiments we do not apply a preconditioner. This is realistic for this type of applications, where the matrix is normally not explicitly available. We compare the performance of IDR(s) with that of the built-in MatLab routines Bi-CGSTAB, CGS [Sonneveld 1989], Bi-CG [Fletcher 1976], QMR [Freund and Nachtigal 1991], and restarted GMRES, with restarts every 50th iteration. The systems are solved to a tolerance (= reduction of the residuals norm) of 10^{-8} . Table V gives the number of matrix-vector products that each of the methods requires, and in between brackets the

required CPU times. The results show that IDR(2) and IDR(4) are both considerably faster than all the other methods, both in terms of matrix-vector multiplications and in terms of CPU-time. The preceding problems cannot be solved on our computer with the MatLab direct solution method due to insufficient memory.

5. CONCLUDING REMARKS

We have presented a new variant of IDR(s), called IDR(s)-biortho, that is slightly less expensive in vector overhead and according to our experiments more accurate and more stable than the original IDR(s), in particular for large s .

For most problems it is not necessary to choose the parameter s large. In our experience $s = 4$ is a good default value. For this value there is normally little difference in numerical behavior between IDR(s)-biortho and IDR(s)-proto. However, some problems require larger values of s , such as the ill-conditioned problem SHERMAN2 that we have presented in the numerical examples. In particular for such cases we consider the new IDR(s) variant an important improvement. In rare cases there is also a difference in numerical behavior between IDR(s)-proto and IDR(s)-biortho for small values of s . We have presented such an example for which IDR(s)-proto failed to converge up to the required accuracy for small values of s , whereas IDR(s)-biortho reached the required accuracy for all the tested values of s .

We have also compared the performance of IDR(s)-biortho with state-of-the-art Krylov methods like GMRES, Bi-CGSTAB, and BiCGstab(ℓ). These experiments confirm that IDR(s) is quite competitive and outperforms the other methods for important problem classes.

ACKNOWLEDGMENTS

Professor Fujino of the Kyushu University is acknowledged for pointing us to the Toeplitz test problem. We thank the referees who made many detailed and constructive comments to an earlier version of this manuscript.

REFERENCES

- DAVIS, T. A. AND HU, Y. F. 2011. The University of Florida sparse matrix collection. *ACM Trans. Math. Softw.* To appear.
- ELMAN, H., SILVESTER, D., AND WATHEN, A. 2005. Finite elements and fast iterative solvers with application in incompressible fluid dynamics. In *Oxford Science Publications*, Oxford University Press.
- ELMAN, H., RAMMAGE, A., AND SILVESTER, D. 2007. Algorithm 866: IFISS, a Matlab toolbox for modelling incompressible flow. *ACM Trans. Math. Softw.* 33, 2, Article 14.
- FLETCHER, R. 1976. Conjugate gradient methods for indefinite systems. In *Lecture Notes in Mathematics*, vol. 506. Springer, 73–89.
- FREUND, R. W. AND NACHTIGAL, N. M. 1991. QMR: A quasi-minimal residual method for non-Hermitian linear systems. *Numer. Math.* 60, 315–339.
- GUTKNECHT, M. H. 1993. Variants of BICGSTAB for matrices with complex spectrum. *SIAM J. Sci. Comput.* 14, 5, 1020–1033.
- GUTKNECHT, M. H. 2010. IDR explained. *Electr. Trans. Numer. Anal.* 36, 126–148.
- GUTKNECHT, M. H. AND ZEMKE, J. P. M. 2010. Eigenvalue computations based on IDR. Res. rep. no. 2010–13, SAM, ETH Zürich, Switzerland.
- ONOUE, Y., FUJINO, S., AND NAKASHIMA, N. 2008. A difference between easy and profound preconditionings of IDR(s) method. *Trans. Japan Soc. Comput. Engin. Sci.* (Paper 20080023 (in Japanese).)
- ONOUE, Y., FUJINO, S., AND NAKASHIMA, N. 2009. An overview of a family of new iterative methods based on IDR theorem and its estimation. In *Proceedings of the International Multi-Conference of Engineers and Computer Scientists IMECS'09*. 2129–2134.
- SAAD, Y. AND SCHULTZ, M. H. 1986. GMRES: A generalized minimum residual algorithm for solving nonsymmetric linear systems *SIAM J. Sci. Statist. Comput.* 7, 856–869.

- SIMONCINI, V. AND SZYLD, D. B. 2010. Interpreting IDR as a Petrov-Galerkin method. *SIAM J. Sci. Comput.* **32**, 1898–1912.
- SLEIJPEN, G. L. G. AND FOKKEMA, D. R. 1994. BiCGstab(ℓ) for linear equations involving matrices with complex spectrum. *ETNA* **1**, 11–32.
- SLEIJPEN, G. L. G., SONNEVELD, P., AND VAN GLJZEN, M. B. 2010. Bi-CGSTAB as an induced dimension reduction method. *Appl. Numer. Math.* **60**, 1100–1114.
- SLEIJPEN, G. L. G. AND VAN DER VORST, H. A. 1995. Maintaining convergence properties of BiCGstab methods in finite precision arithmetic. *Numer. Algor.* **10**, 203–223.
- SONNEVELD, P. 1989. CGS: A fast Lanczos-type solver for nonsymmetric linear systems. *SIAM J. Sci. Statist. Comput.* **10**, 36–52.
- SONNEVELD, P. 2010. On the convergence behaviour of IDR(s). Tech. rep. 10-08, Department of Applied Mathematical Analysis, Delft University of Technology, Delft, The Netherlands.
- SONNEVELD, P. AND VAN GLJZEN, M. B. 2008. IDR(s): a family of simple and fast algorithms for solving large nonsymmetric linear systems. *SIAM J. Sci. Comput.* **31**, 2, 1035–1062.
- VAN DER VORST, H. A. 1992. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Sci. Comput.* **13**, 631–644.
- VAN GLJZEN, M. B. AND SONNEVELD, P. 2008. The IDR(s) method for solving nonsymmetric systems: A performance study for CFD problems. In *High Performance Algorithms for Computational Science and Their Applications*, Research Institute for Mathematical Sciences, Kyoto University, Japan.
- WESSELING, P. AND SONNEVELD, P. 1980. Numerical experiments with a multiple grid- and a preconditioned Lanczos type method. In *Lecture Notes in Mathematics*, vol. 771. Springer, 543–562.
- YEUNG, M. AND CHAN, T. F. 1999. ML(k)BiCGSTAB: A BiCGSTAB variant based on multiple Lanczos starting vectors. *SIAM J. Sci. Comput.* **21**, 1263–1290.
- ZHANG, S. L. 1997. GPBi-CG: Generalized product-type methods based on Bi-CG for solving nonsymmetric linear systems. *SIAM J. Sci. Comput.* **18**, 2, 537–551.

Received September 2010; accepted January 2011