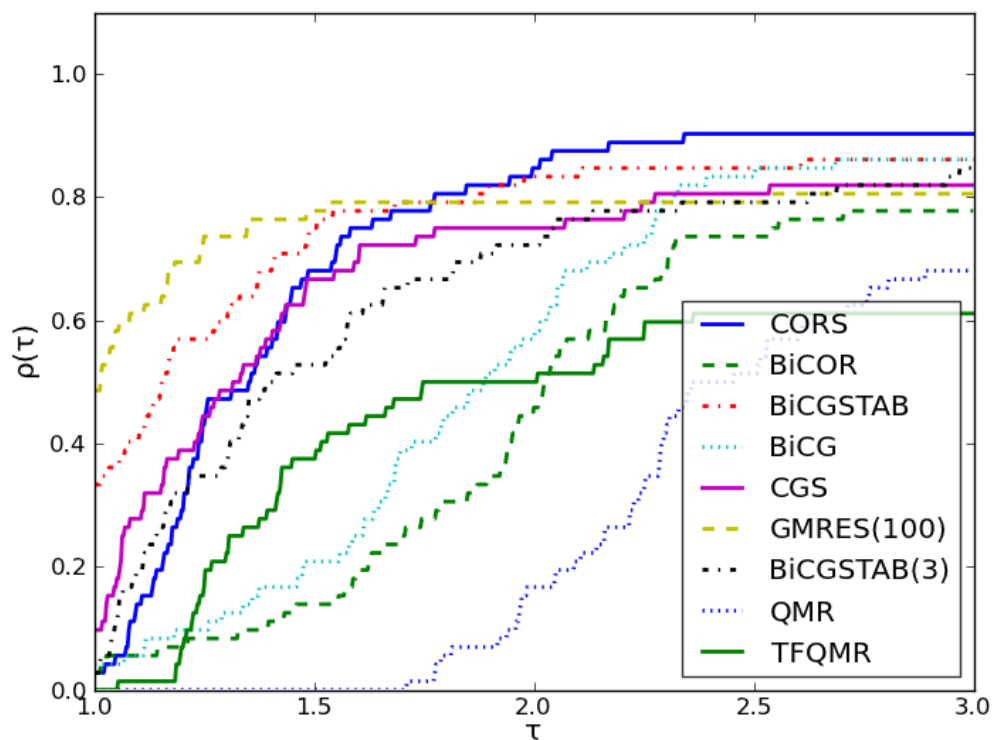




# Performance analysis of the CORS and BiCOR iterative methods for solving nonsymmetric sparse linear systems



Bachelor Thesis in Applied Mathematics

August 2011

Student: S. Baars

Supervisor: Dr. B. Carpentieri



## Abstract

Recently, the iterative methods BiCOR and CORS for solving real nonsymmetric (or complex non-Hermitian), and possibly indefinite sparse linear systems were developed. There is not much known yet about the performance of those methods. We consider iterative methods in general, and go more into detail about CORS and BiCOR. We analyse the performance of BiCOR and CORS by comparing them to seven popular solvers on a large set of publicly available matrices coming from different areas of application. We use different qualities of preconditioners to do this. In our experiments we observe that CORS is a highly competitive solver compared to other popular solvers, like GMRES and BiCGSTAB.

**Keywords:** CORS, BiCOR, GMRES, BiCGSTAB; iterative methods, Krylov subspace methods; performance profiles; preconditioning.

## Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
<b>2</b>	<b>Krylov subspace methods</b>	<b>1</b>
2.1	The Krylov subspace . . . . .	1
2.2	Arnoldi's method . . . . .	2
2.3	Different approaches . . . . .	4
2.4	The GMRES method . . . . .	5
2.5	Preconditioning . . . . .	5
<b>3</b>	<b>The Petrov-Galerkin projection</b>	<b>6</b>
3.1	The basics . . . . .	6
3.2	The two-sided biconjugate A-orthonormalisation method . . . . .	8
3.3	The biconjugate A-orthonormalisation procedure for solving general linear systems	10
<b>4</b>	<b>The BiCOR method</b>	<b>12</b>
<b>5</b>	<b>The CORS method</b>	<b>13</b>
<b>6</b>	<b>Computational aspects</b>	<b>14</b>
6.1	Preconditioning . . . . .	14
6.2	Stopping criteria . . . . .	15
6.3	Implementational aspects . . . . .	16
<b>7</b>	<b>Numerical experiments</b>	<b>17</b>
7.1	Information about the experiments . . . . .	17
7.2	Data analysis . . . . .	18
7.3	Results . . . . .	18
7.3.1	Speed . . . . .	20
7.3.2	Reliability . . . . .	21

<b>8</b>	<b>Conclusion</b>	<b>23</b>
<b>9</b>	<b>Acknowledgments</b>	<b>23</b>
<b>A</b>	<b>Problems</b>	<b>28</b>
A.1	Problem types . . . . .	28
A.1.1	Problems with 2D/3D geometry . . . . .	28
A.1.2	Problems that normally do not have 2D/3D geometry . . . . .	28
A.2	Problem list . . . . .	28
<b>B</b>	<b>Implementation of BiCOR</b>	<b>31</b>
B.1	User documentation . . . . .	31
B.1.1	Argument lists and calling sequence . . . . .	31
B.1.1.1	Initialization of the control parameters . . . . .	31
B.1.1.2	Solving $Ax=b$ . . . . .	31
B.1.2	Control parameters . . . . .	33
B.1.3	Error values . . . . .	34
B.1.4	General information . . . . .	35
B.2	Implementation . . . . .	35
<b>C</b>	<b>Implementation of CORS</b>	<b>39</b>
C.1	User documentation . . . . .	39
C.1.1	Argument lists and calling sequence . . . . .	39
C.1.1.1	Initialization of the control parameters . . . . .	39
C.1.1.2	Solving $Ax=b$ . . . . .	39
C.1.2	Control parameters . . . . .	41
C.1.3	Error values . . . . .	42
C.1.4	General information . . . . .	42
C.2	Implementation . . . . .	42
<b>D</b>	<b>Implementation of the testing application</b>	<b>46</b>
<b>E</b>	<b>Implementation of the data analysis tool</b>	<b>56</b>

# 1 Introduction

Computational simulation of scientific and engineering problems often involves solving large systems of equations of the form

$$Ax = b, \tag{1.1}$$

with  $A \in \mathbb{C}^{m \times n}$ ,  $x \in \mathbb{C}^n$  and  $b \in \mathbb{C}^n$ . The usual way of solving small systems of linear equations of the form (1.1) is by using Gaussian elimination. Gaussian elimination, however, as well as other direct methods, has a cost of  $\mathcal{O}(n^3)$ [33]. This is really expensive if the order  $n$  of the matrix  $A$  is large and also unnecessarily expensive if the matrix is sparse, i.e. it contains many zero entries. If the matrix is sparse, not only the computational cost is expensive, but also the storage cost in the memory. For direct methods it is usually still needed to store  $n^2$  entries in the memory, where one would like to only store the  $\mathcal{O}(n)$  nonzero entries in the matrix. In this case it might be useful to use an iterative method, and a Krylov subspace method in particular.

Developing those methods is a continuously evolving subject of research. Recently, a new family of iterative methods were developed around the two-sided A-orthonormalisation procedure that will be introduced in this thesis. To date, very little is known about the performance of those methods. We test the performance of two of those methods, BiCOR and CORS, and compared those with some other popular iterative methods. We do this mostly using various qualities of preconditioners. To compare BiCOR and CORS to the other iterative methods, we used a FORTRAN implementation, that we also provide here.

As an introduction to the subject, we discuss iterative methods, also called Krylov methods, mostly following Van der Vorst in [40]. Then we show how the BiCOR and CORS method can be derived from the two-sided A-orthonormalisation procedure following [8], and finally, we analyse the results of our experiments.

## 2 Krylov subspace methods

### 2.1 The Krylov subspace

The general idea behind iterative methods is that we want to solve the system  $Ax = b$ , and at each iteration  $i$ , we have an approximate solution  $x_i$ . We can also write this as  $x = x_i + \epsilon_i$  where  $\epsilon_i$  is the error at step  $i$ . Multiplication by  $A$  gives us

$$A\epsilon_i = A(x - x_i) = b - Ax_i.$$

Since we do not have the real solution, we do not know the actual error either. Instead we try to solve the system

$$Mz_i = b - Ax_i$$

for  $z_i$ , with  $M$  an approximation of  $A$  that makes the system easier to solve. If we take  $x_0 = 0$  for instance, the first step would be solving  $Mz_0 = b$ . Since  $M$  is an approximation of  $A$ ,  $z_i$  is an approximation of the error. Thus solving the easier system leads to a better approximation of the solution:  $x_{i+1} = x_i + z_i$ . The basic iteration introduced here, now leads to

$$x_{i+1} = x_i + M^{-1}(b - Ax_i),$$

where  $M$  is called the preconditioner. One uses a preconditioner to speed up convergence. An iterative method converges fast when  $M^{-1}A$  is close to identity. If  $M^{-1}A$  was equal to identity,

we would have convergence in one step. We only write the inverse of  $M$  for notational purposes. In practice,  $M^{-1}$  is usually not calculated. For more information about preconditioners, see section 2.5.

If we now take  $M = I$ , we obtain the well known Richardson iteration [40]

$$x_{i+1} = b + (I - A)x_i = x_i + r_i,$$

with  $r_i = b - Ax_i$  the residual at step  $i$ . We try to find a relation between  $r_{i+1}$  and  $r_i$  by multiplying the above relation by  $-A$  and adding  $b$  to it

$$b - Ax_{i+1} = b - Ax_i - Ar_i$$

so

$$r_{i+1} = (I - A)r_i = (I - A)^{i+1}r_0.$$

It then follows that the approximate solution  $x_{i+1}$  may be written as

$$x_{i+1} = r_0 + r_1 + \dots + r_i = \sum_{k=0}^i (I - A)^k r_0$$

for  $x_0 = 0$ . We can do this without loss of generality, because in case  $x_0$  is nonzero, we could just shift the system by setting  $Ay = b - Ax_0 = \hat{b}$  with  $y_0 = 0$ . We now observe that

$$x_{i+1} \in \text{Span}\{r_0, Ar_0, \dots, A^i r_0\} \equiv \mathcal{K}_{i+1}(A; r_0).$$

The space of dimension  $m$ , spanned by a given vector  $v$ , and increasing powers of  $A$  applied to  $v$  up to the  $(m - 1)$ th power of  $A$  is called the  $m$ -dimensional Krylov subspace generated by  $A$  and  $v$ , and is denoted as  $\mathcal{K}_m(A; v)$  [13, 40].

## 2.2 Arnoldi's method

Assuming the matrix  $A$  has  $n$  eigenvalues  $|\lambda_1| > |\lambda_2| \geq |\lambda_3| \geq \dots \geq |\lambda_n| \geq 0$ , and linearly independent eigenvectors  $\{v_1, v_2, \dots, v_n\}$  with  $Av_i = \lambda_i v_i$ , we may write the solution  $x$  to the system  $Ax = b$  as

$$x = \sum_{j=1}^n \alpha_j v_j.$$

Multiplying both sides by  $A^k$  gives

$$A^k x = \sum_{j=1}^n \alpha_j A^k v_j = \sum_{j=1}^n \alpha_j \lambda_j^k v_j.$$

If we factor out  $\lambda_1^k$  from the right hand side

$$A^k x = \lambda_1^k \sum_{j=1}^n \alpha_j \frac{\lambda_j^k}{\lambda_1^k} v_j$$

we see that, since  $|\lambda_1| > |\lambda_2| \geq |\lambda_3| \geq \dots \geq |\lambda_n| \geq 0$ , this converges to

$$A^k x = \lambda_1^k \alpha_1 v_1$$

as  $k \rightarrow \infty$  and assuming that  $\alpha_1 \neq 0$  [7]. This only holds because  $|\lambda_1|$  has to be strictly greater than  $|\lambda_2|$ . Otherwise we would not be able to factor out the eigenvectors belonging to  $\lambda_2$  to  $\lambda_n$ . This is the idea behind the basic iteration for finding eigenpairs  $(\theta_m, v_m)$  through the Power Method.

Something we observe, is that the obvious basis  $\{r_0, Ar_0, \dots, A^{m-1}r_0\}$  for the Krylov subspace  $\mathcal{K}_m(A; r_0)$  is not very attractive. The vectors  $A^j r_0$  point more and more in the direction of the dominant eigenvector for increasing  $j$ . Hence, the basis vectors will become linearly dependent in finite precision arithmetic. This is why we want to make sure that we have an orthogonal basis for our Krylov subspace [40].

A way to form an orthogonal basis for the Krylov subspace is suggested by Arnoldi [1]. Arnoldi's method often uses modified Gram-Schmidt orthogonalisation in the process of finding this basis. Modified Gram-Schmidt orthogonalisation is used, because normal Gram-Schmidt orthogonalisation can produce linearly dependent vectors due to rounding errors. We first describe this process before describing Arnoldi's method itself.

In the modified Gram-Schmidt process [31], we start with the vector

$$q_1 = \frac{1}{\|x_1\|} x_1,$$

where  $x_1, \dots, x_n$  form an ordinary basis. Now, at the beginning of the  $(k+1)$ -th step, the projections of the vector  $x_{k+1}$  along the vectors  $q_i, \dots, q_k$  are progressively subtracted from  $x_{k+1}$ . We can write the first step of the subtraction within the  $(k+1)$ -th step of the process itself as

$$x_{k+1}^{(1)} = x_{k+1} - q_1^T x_{k+1} q_1.$$

This new vector  $x_{k+1}^{(1)}$  is then projected along  $q_2$  and subtracted again from  $x_{k+1}^{(1)}$ , yielding

$$x_{k+1}^{(2)} = x_{k+1}^{(1)} - q_2^T x_{k+1}^{(1)} q_2.$$

We can continue this process until  $x_{k+1}^{(k)}$  is computed.  $q_{k+1}$  is then given as

$$q_{k+1} = \frac{1}{\|x_{k+1}^{(k)}\|} x_{k+1}^{(k)}.$$

We can now describe the procedure of Arnoldi's method as follows. We start with  $v_1 = r_0 / \|r_0\|_2$ . Then we compute  $Av_1$ , make it orthogonal to  $v_1$ , and normalise the result using the modified Gram Schmidt process described above. This gives us  $v_2$ . In general, we have an orthonormal basis  $v_1, \dots, v_j$  for our Krylov subspace  $\mathcal{K}_j(A; r_0)$ . We expand this basis by calculating  $t = Av_j$  and orthonormalising this vector  $t$  with respect to the basis  $v_1, \dots, v_j$ . This leads to an algorithm as seen in Algorithm 1 to form a basis  $v_1, \dots, v_m$  for  $\mathcal{K}_m(A; r_0)$ . The orthogonalisation can be done in different ways, but the most common way is to use the modified Gram-Schmidt process [33, 40].

We clearly see here that the matrix  $A$  is only accessed through matrix-vector products, which is an advantage compared to direct methods, where the matrix is used directly. This allows us to specify our own matrix-vector product if we store the matrix for example in a sparse format.

Now let  $V_j$  denotes the matrix that has columns  $v_1, \dots, v_j$ . We then see from Arnoldi's method that

$$AV_{m-1} = V_m H_{m,m-1} \tag{2.1}$$

---

**Algorithm 1** *Arnoldi's method using modified Gram-Schmidt orthogonalisation.*

---

```
1:  $v_1 = r_0 / \|r_0\|_2$ 
2: for  $j = 1, 2, \dots, m-1$  do
3:    $t = Av_j$ 
4:   for  $i = 1, \dots, j$  do
5:      $h_{i,j} = v_i^T t$ 
6:      $t = t - h_{i,j} v_i$ 
7:   end for
8:    $h_{j+1,j} = \|t\|_2$ 
9:    $v_{j+1} = t / h_{j+1,j}$ 
10: end for
```

---

where  $H_{m,m-1}$  is an upper Hessenberg matrix, which means that  $h_{i,j} = 0$  for  $i > j + 1$ . The other entries of  $H_{m,m-1}$  are defined by Arnoldi's method. We see that the orthogonalisation becomes increasingly expensive for increasing dimension of the subspace, since every iteration needs one extra inner product and vector update compared to the last iteration to compute a new column of  $H_{m,m-1}$ .

If  $A$  is symmetric however, so is  $H_{m-1,m-1} = V_{m-1}^T A V_{m-1}$ . Therefore, in this case,  $H_{m-1,m-1}$  is tridiagonal. This means that during the orthogonalisation process, most inner products vanish, so the work does not increase. The resulting three-term recurrence relation is known as the Lanczos method [29] that has some well known methods derived from it. The constant amount of work that has to be done during every iteration in the Lanczos method is one matrix vector product, two inner products and two vector updates.

### 2.3 Different approaches

Methods that attempt to generate an approximate solution from the Krylov subspace, like Arnoldi's method, are usually referred to as Krylov subspace methods. There are four classes of Krylov subspace methods that can be distinguished:

- The *Ritz-Galerkin approach*: Require for  $x_k$  that the residual is orthogonal to the current subspace:  $b - Ax_k \perp \mathcal{K}_k(A; r_0)$ .
- The *minimum norm residual approach*: Require for  $x_k$  that the Euclidean norm  $\|b - Ax_k\|_2$  is minimal over  $\mathcal{K}_k(A; r_0)$ .
- The *Petrov-Galerkin approach*: Require for  $x_k$  that the residual  $b - Ax_k$  is orthogonal to some other suitable  $k$ -dimensional subspace.
- The *minimum norm error approach*: Require for  $x_k$  in  $A^T \mathcal{K}_k(A^T; r_0)$  that the Euclidean norm  $\|x_k - x\|_2$  is minimal.

The Ritz-Galerkin approach leads to methods such as the Lanczos method mentioned before and the Conjugate Gradient (CG) method [21]. The minimum norm residual approach leads to methods such as the Generalised Minimum Residual (GMRES) method [36]. The minimum norm error approach leads to some less well known methods that will not be discussed in this paper. And lastly, the Petrov-Galerkin approach leads to methods such as the Biconjugate Gradient (BiCG) method [15], the Quasi-Minimal Residual (QMR) method [18], and the Biconjugate A-Orthogonal Residual (BiCOR) [25] method we discuss later in this thesis. Other methods like the Conjugate Gradients Squared (CGS) [38], Biconjugate Gradient Stabilised (BiCGSTAB) [39] and BiCGSTAB( $\ell$ ) [37] methods and the Conjugate A-Orthogonal Residual Squared (CORS) method [25] we discuss later, are hybrids of the different approaches.



## 2.4 The GMRES method

The GMRES method can be derived using the minimum residual approach. It is an optimal method, in the sense that it minimizes the 2-norm of the residual over the corresponding Krylov space. Starting from Arnoldi's method, in (2.1), we had an orthogonal basis for the Krylov subspace of dimension  $i + 1$ , leading to

$$AV_i = V_{i+1}H_{i+1,i}.$$

We are looking for an  $x_i \in \mathcal{K}_i(A; r_0)$ , such that the residual,  $\|b - Ax_i\|_2$ , is minimal. Since  $x_i \in \mathcal{K}_i(A; r_0)$ , we can also write  $x_i = V_i y$ . The norm of the residual can be rewritten as

$$\|r_i\|_2 = \|b - Ax_i\|_2 = \|b - AV_i y\|_2 = \|\beta V_{i+1} e_1 - V_{i+1} H_{i+1,i} y\|_2,$$

with  $\beta \equiv \|r_0\|_2$ . Now, since the column vectors of  $V_{i+1}$  are orthonormal, we have

$$\|b - Ax_i\|_2 = \|\beta e_1 - H_{i+1,i} y\|_2,$$

which can be solved as a least squares problem. This least squares problem can be solved by making the QR-factorisation of  $H_{i+1,i}$ , and because of the upper Hessenberg structure, this can be done efficiently using Givens matrices. Givens matrices are elementary rotation matrices of the form  $G(i, k, \theta) = I - Y$ , where  $I$  is the identity matrix and  $Y$  is a null matrix except for the elements  $y_{ii} = y_{kk} = 1 - \cos(\theta)$  and  $y_{ik} = -y_{ki} = -\sin(\theta)$ . The Givens rotations remove the subdiagonal elements from the upper Hessenberg matrix  $H_{i+1,i}$ , resulting in an upper triangular matrix  $R_{i,i}$ :

$$H_{i+1,i} = Q_{i+1,i} R_{i,i}$$

where  $Q_{i+1,i}$  is the matrix consisting of the product of successive Givens rotations. Now we can write the least squares problem as the minimisation of

$$\begin{aligned} \|\beta e_1 - H_{i+1,i} y\|_2 &= \|\beta e_1 - Q_{i+1,i} R_{i,i} y\|_2 \\ &= \|Q_{i+1,i}^T \beta e_1 - R_{i,i} y\|_2. \end{aligned}$$

This leads to the minimum norm solution

$$y = R_{i,i}^{-1} Q_{i+1,i}^T \beta e_1,$$

where the approximate solution  $x_i$  is computed as  $x_i = V_i y$ .

To lower the storage and computational costs of the orthogonalisation process, GMRES is usually restarted after  $m$  steps. This method is referred to as GMRES( $m$ ). By restarting GMRES, we lose the optimality property however. The non-restarted version of GMRES is also referred to as full GMRES.

## 2.5 Preconditioning

In general, all iterative methods we mentioned before converge rapidly if the matrix  $A$  of the problem  $Ax = b$  from (1.1) is close to identity. If the matrix is equal to the identity matrix, those methods converge in one step. For most problems however, the matrix  $A$  is far from being close to identity, and therefore one can not be sure that the iterative methods will compute a good approximation of the solution in  $m \ll n$  iterations. Every different method has its own drawbacks. In exact arithmetic, some methods, like full GMRES, lead to the exact solution in

at most  $n$  steps, but that might not be very practical. Other methods, like CG, only work for certain kinds of matrices. In the case of CG, the matrix must be symmetric positive definite. There are also methods, like BiCG, BiCGSTAB and the methods discussed in this thesis, CORS and BiCOR, that might suffer from breakdowns or stagnation. The rate of convergence depends in a very complicated way on the spectral properties (eigenvalue distribution, etc.) of the matrix  $A$  and in real applications, this information is not available [40].

The trick is to use a preconditioner. A preconditioner  $M \approx A$  tries to get the original problem closer to identity, so the spectral properties are better. In general, one can write

$$M_1^{-1}AM_2^{-1}y = M_1^{-1}b$$

where  $M_1$  is the left preconditioner, and  $M_2$  is the right preconditioner with  $x = M_2^{-1}y$ . If we choose, for example  $M_1 = A$  and  $M_2 = I$ , the problem is solved in one step. Note that if we precondition from the right, the residual stays the same as the residual of the original system, because

$$r = b - A\hat{x} = b - AM_2^{-1}y, \quad (2.2)$$

where  $\hat{x}$  is the approximate solution.

Calculating the inverse is usually very expensive, so instead of calculating the exact inverse, one can also try to approximate it. Tools to derive those preconditioners are even more diverse than those used in the derivation of iterative methods [20], and therefore we do not discuss this here.

### 3 The Petrov-Galerkin projection

#### 3.1 The basics

For nonsymmetric matrices, it is desirable to have a three-term recurrence relation similar to the one from the Lanczos method. Due to the work of Faber and Manteuffel, we know that for nonsymmetric matrices, it is not possible to find short-term recurrence relations while keeping the optimality property as for the GMRES method [14]. To reduce memory usage and computational costs, however, it is still very useful to derive non-optimal methods. Let's start with what we already know from Arnoldi's method. This suggested a basis

$$h_{i+1,i}v_{i+1} = Av_i - \sum_{j=1}^i h_{j,i}v_j \quad (3.1)$$

for the Krylov subspace, which could be written as

$$V_{i+1}H_{i+1,i} = AV_i \quad (3.2)$$

in matrix notation. Using  $V_i$  for the projection, we would end up making the new vector orthogonal to Krylov subspace like we would do in the Ritz-Galerkin approach, but that's not what we want here. So suppose that we have  $W_i$  for which  $W_i^T V_i = D_i$  with  $D_i$  a diagonal matrix, and for which  $v_{i+1}$  is orthogonal to  $W_i$ , so  $W_i^T v_{i+1} = 0$ . Then

$$W_i^T AV_i = D_i H_{i,i}. \quad (3.3)$$

Our goal is to find a  $W_i$  such that  $H_{i,i}$  is tridiagonal. In this case we would have a three-term recurrence relation. So from the above statement we see that  $(W_i^T AV_i)^T = V_i^T A^T W_i$  should

also be tridiagonal. This looks very similar to what we have in (3.3), so this suggests we can write a relation similar to (3.1), but now with  $W_i$ .

Let's choose an arbitrary  $w_1 \neq 0$  with  $w_1^T v_1 \neq 0$ . Then we can use (3.1) to generate  $v_2$  and orthogonalise it with respect to  $w_1$ . So then from (3.3) we see that  $h_{1,1} = w_1^T A v_1 / (w_1^T v_1)$ . Since

$$\begin{aligned} v_1^T h_{1,1} w_1 &= v_1^T (w_1^T A v_1) w_1 / (w_1^T v_1) \\ &= (w_1^T A v_1) v_1^T w_1 / (w_1^T v_1) \\ &= w_1^T A v_1 = v_1^T A^T w_1, \end{aligned}$$

we see that  $w_2$  generated from

$$h_{2,1} w_2 = A^T w_1 - h_{1,1} w_1 \quad (3.4)$$

is also orthogonal to  $v_1$ . That is,

$$v_1^T w_2 = \frac{1}{h_{2,1}} (v_1^T A^T w_1 - v_1^T h_{1,1} w_1) = \frac{1}{h_{2,1}} (v_1^T A^T w_1 - v_1^T A^T w_1) = 0.$$

Relation (3.4) indeed looks similar to (3.1).

We can go on with this, and see that we can create bi-orthogonal basis sets  $\{v_j\}$  and  $\{w_j\}$  by making every new  $v_{i+1}$  orthogonal to  $w_1, \dots, w_i$ , and then generating  $w_{i+1}$  using the same coefficients, but with  $A^T$  instead of  $A$ .

Now we have that both  $W_i^T A V_i = D_i H_{i,i}$  and  $V_i^T A^T W_i = D_i H_{i,i}$ . This implies that  $D_i H_{i,i}$  is symmetric, and hence our Hessenberg matrix  $H_{i,i}$  is tridiagonal. This gives us the three-term recurrence relation we wanted with  $v_1, \dots, v_i$  a basis for  $\mathcal{K}_i(A; v_1)$  and  $w_1, \dots, w_i$  a basis for  $\mathcal{K}_i(A^T; w_1)$ . The matrix  $H_{i,i}$  is also commonly denoted as  $T_{i,i}$  due to its tridiagonal form. The three-term recurrence relation we found does not only save a lot of computational power, but also requires less memory. We only have to store the last three vectors of both of the bases.

The two-sided Lanczos method [30] follows from what we derived above. Since we have a tridiagonal matrix, we can write  $v_{i+1}$  and  $w_{i+1}$  at step  $i$  in the construction of the dual basis as

$$\delta_i v_{i+1} = A v_i - \alpha_i v_i - \beta_{i-1} v_{i-1}$$

and

$$\delta_i w_{i+1} = A^T w_i - \alpha_i w_i - \beta_{i-1} w_{i-1}.$$

---

**Algorithm 2** *The two-sided Lanczos method.*

---

- 1: Choose a  $v_1$  and  $w_1$  such that  $w_1^T v_1 = \gamma_1 \neq 0$
  - 2:  $\beta_0 = 0, w_0 = v_0 = 0$
  - 3: **for**  $j = 1, 2, \dots$  **do**
  - 4:    $p = A v_i - \beta_{i-1} v_{i-1}$
  - 5:    $\alpha_i = w_i^T p / \gamma_i$
  - 6:    $p = p - \alpha_i v_i$
  - 7:    $\delta_{i+1} = \|p\|_2$
  - 8:    $v_{i+1} = p / \delta_{i+1}$
  - 9:    $w_{i+1} = (A^T w_i - \beta_{i-1} w_{i-1} - \alpha_i w_i) / \delta_{i+1}$
  - 10:    $\gamma_{i+1} = w_{i+1}^T v_{i+1}$
  - 11:    $\beta_i = \delta_{i+1} \gamma_{i+1} / \gamma_i$
  - 12: **end for**
-

So here we have  $\delta_i = h_{i+1,i}$ ,  $\alpha_i = h_{i,i}$ , and  $\beta_i = h_{i-1,i}$  where available. The full method is given in Algorithm 2. In this algorithm we also use  $\gamma_i = w_i^T v_i$ . The only thing we do here is repeatedly calculating  $v_{i+1}$  and  $w_{i+1}$  using exactly what we derived above.

The method we described above can also be seen as an oblique projection of the residual onto the space orthogonal to the space spanned by  $W$ , which is exactly what our initial condition  $r_k \perp \mathcal{K}_m(A^T, w_1)$  says. With an oblique projection we mean the projection of a vector, onto a space  $\mathcal{K}$  orthogonal to a space  $\mathcal{L}^\perp$ . We also say the projection is along  $\mathcal{L}$  onto  $\mathcal{K}$  [35, 6]. Put in another way, oblique projections are projections that are not orthogonal. Orthogonal projections actually project onto the orthogonal space.




Figure 1: 2D interpretation of the Petrov-Galerkin projection

In the case of the Petrov-Galerkin projection, we are projecting the residual along  $AW$  onto the space orthogonal to the space spanned by  $W$ , as can be seen in Figure 1. In this figure, the purple line resembles the projection. The image was created for a  $2 \times 2$  matrix, using the two-sides Lanczos method to construct the bases  $V$  and  $W$ . We chose  $v_1 = r_0/\|r_0\|_2$  and  $w_1 = r_0$ . After computing  $V$  and  $W$ , we applied the prototype projection method suggested by Saad in [35]. The prototype projection method can be found in Algorithm 3

---

**Algorithm 3** *Prototype projection method*

---

- 1: **for**  $i = 1, 2, \dots$ , until convergence **do**
  - 2:   Select a pair of subspaces  $\mathcal{K}$  and  $\mathcal{L}$
  - 3:   Choose bases  $V = [v_1, \dots, v_i]$  and  $W = [w_1, \dots, w_i]$  for  $\mathcal{K}$  and  $\mathcal{L}$
  - 4:    $r_i = b - Ax_i$
  - 5:    $y = (W^T AV)^{-1} W^T r_i$
  - 6:    $x_{i+1} = x_i + Vy$
  - 7: **end for**
- 

### 3.2 The two-sided biconjugate A-orthonormalisation method

The two-sided biconjugate A-orthonormalisation method [25] is a method similar to the two-sided Lanczos method, and can, like the two-sided Lanczos method, be used for nonsymmetric matrices. Given two vectors  $v_1$  and  $w_1$  for which  $w_1^T A v_1 = 1$ , we define two Lanczos-type vectors  $v_j$  and  $w_j$  very similar to the ones we described in the last section. We again use scalars

$\alpha_j, \beta_j$  and  $\delta_j$ . The two vectors are recursively defined as

$$\delta_{j+1}v_{j+1} = Av_j - \beta_jv_{j-1} - \alpha_jv_j, \quad (3.5)$$

$$\beta_{j+1}w_{j+1} = A^T w_j - \delta_jw_{j-1} - \alpha_jw_j \quad (3.6)$$

where the scalars are chosen as

$$\alpha_j = w_j^T A^2 v_j, \quad \beta_j = w_{j-1}^T A^2 v_j, \quad \delta_j = w_j^T A^2 v_{j-1}.$$

The choice of the scalars assures that the vectors  $v_j$  and  $w_j$  form a biconjugate  $A$ -orthonormal basis. So  $w_i^T Av_j = \delta_{i,j}$ , with  $\delta_{i,j}$  the Kronecker delta. The rest of the procedure can be derived from the two-sided Lanczos algorithm. For the sake of clarity, we show a complete version of the process in Algorithm 4.

---

**Algorithm 4** *The biconjugate  $A$ -orthonormalisation procedure.*

---

```

1: Choose a  $v_1$  and  $w_1$  such that  $w_1^T Av_1 = 1$ 
2:  $\beta_0 = \delta_1 = 0, w_0 = v_0 = 0$ 
3: for  $j = 1, 2, \dots$  do
4:    $\alpha_j = w_j^T A(Av_j)$ 
5:    $\tilde{v}_{j+1} = Av_j - \alpha_jv_j - \beta_jv_{j-1}$ 
6:    $\tilde{w}_{j+1} = A^T w_j - \alpha_jw_j - \delta_jw_{j-1}$ 
7:    $\delta_{j+1} = |\tilde{w}_{j+1}^T A\tilde{v}_{j+1}|^{\frac{1}{2}}$ 
8:    $\beta_{j+1} = \frac{\tilde{w}_{j+1}^T A\tilde{v}_{j+1}}{\delta_{j+1}}$ 
9:    $v_{j+1} = \frac{\tilde{v}_{j+1}}{\delta_{j+1}}$ 
10:   $w_{j+1} = \frac{\tilde{w}_{j+1}}{\beta_{j+1}}$ 
11: end for
```

---

The fact that the basis sets  $\{v_j\}$  and  $\{w_j\}$ , generated from relations (3.5) and (3.6), really form a basis of the Krylov subspaces  $\mathcal{K}_i(A; v_1)$  and  $\mathcal{K}_i(A^T; w_1)$  can be shown in a similar way as we did for the two-sided Lanczos method, and can be found for example in [25]. Additionally, the following relations hold

$$AV_m = V_{m+1}T_{m+1,m}, \quad (3.7)$$

$$A^T W_m = W_{m+1}T_{m,m+1}^T, \quad (3.8)$$

$$W_m^T AV_m = I_m, \quad (3.9)$$

$$W_m^T A^2 V_m = T_m, \quad (3.10)$$

with

$$T_m = \begin{pmatrix} \alpha_1 & \beta_2 & & & \\ \delta_2 & \alpha_2 & \beta_3 & & \\ & \ddots & \ddots & \ddots & \\ & & \delta_{m-1} & \alpha_{m-1} & \beta_m \\ & & & \delta_m & \alpha_m \end{pmatrix}$$

and  $V_m = [v_1, v_2, \dots, v_m]$ ,  $W_m = [w_1, w_2, \dots, w_m]$ .

Due to the three-term recurrence relation like the one in the two-sided Lanczos method, we can overwrite for example  $w_{j-1}$  with  $w_{j+1}$ . After all, we see in Algorithm 4, line 6, that  $w_{j-1}$  is not used after  $\tilde{w}_{j+1}$  has been computed. An advantage of the three-term recurrence relation is that storage is very limited if you compare it to Arnoldi's method.

The method can possible fail if  $\delta_{j+1}$  vanishes while  $\tilde{w}_{j+1}$  and  $A\tilde{v}_{j+1}$  are not the zero vector. One could try to recover from such failures using so-called look-ahead strategies [32] as used in for instance the QMR implementation we use for this thesis.

### 3.3 The biconjugate A-orthonormalisation procedure for solving general linear systems

As one can derive for instance the Biconjugate Gradient method from the two-sided Lanczos method, one can also derive methods for solving  $Ax = b$  from the the Biconjugate A-Orthonormalisation procedure by applying a Petrov-Galerkin projection. We will describe this process in three steps

**Step 1** Run Algorithm 4 for  $m \ll n$  steps and generate  $V_m$ ,  $W_m$  and  $T_m$  as described above.

**Step 2** Compute the approximate solution  $x_m$  that belongs to the Krylov subspace  $x_0 + \mathcal{K}_m(A; v_1)$  by using the Petrov-Galerkin projection to project the residual orthogonally to the space  $A^T \mathcal{K}_m(A^T; w_1)$ , so

$$r_m = b - Ax_m \perp A^T \mathcal{K}_m(A^T; w_1). \quad (3.11)$$

Using matrix notation, we may also write

$$(A^T W_m)^T (b - Ax_m) = 0$$

and since our approximate solution is of the form

$$x_m = x_0 + V_m y_m \quad (3.12)$$

we get, using (3.10)

$$\begin{aligned} (A^T W_m)^T (b - A(x_0 + V_m y_m)) &= (A^T W_m)^T r_0 - (A^T W_m)^T V_m y_m \\ &= W_m^T A r_0 - T_m y_m = 0 \end{aligned}$$

so, if we have  $v_1 = r_0 / \|r_0\|_2$ ,

$$T_m y_m = \|r_0\|_2 e_1 \quad (3.13)$$

with  $e_1$  the first canonical unit vector.

**Step 3** Compute the new residual, and terminate if it meets the stopping criterion. Otherwise, enlarge the Krylov subspace and start again.

By using this method, we not only solve the system  $Ax = b$ , but also implicitly the system  $A^T x' = b'$ . We use the notation used in [8] by denoting vectors belonging to this dual system with primed symbols. We can now say we compute the approximation  $x'_m$  that belongs to the Krylov subspace  $x'_0 + \mathcal{K}_m(A^T; w_1)$ , so we get a relation similar to (3.11)

$$r'_m = b' - A^T x'_m \perp A \mathcal{K}_m(A; v_1).$$

We also get similar relations as above, but now for the dual system

$$(AV_m)^T (b' - A^T x'_m) = 0, \quad (3.14)$$

$$x'_m = x'_0 + W_m y'_m, \quad (3.15)$$

$$T_m^T y'_m = \|r'_0\| e_1. \quad (3.16)$$

We can now update  $x_m$  and  $x'_m$  from  $x_{m-1}$  respectively  $x'_{m-1}$ . Assume the LU factorisation of the tridiagonal matrix  $T_m$  is

$$L_m U_m = T_m.$$

Substituting this expression in (3.12), (3.13) and (3.15), (3.16), we get

$$\begin{aligned} x_m &= x_0 + V_m (L_m U_m)^{-1} (\|r_0\|_2 e_1) \\ &= x_0 + V_m U_m^{-1} L_m^{-1} (\|r_0\|_2 e_1) \\ &= x_0 + P_m z_m \\ x'_m &= x'_0 + W_m (U_m^T L_m^T)^{-1} (\|r'_0\|_2 e_1) \\ &= x'_0 + W_m (L_m^T)^{-1} (U_m^T)^{-1} (\|r'_0\|_2 e_1) \\ &= x'_0 + P'_m z'_m \end{aligned}$$

where we take  $P_m = V_m U_m^{-1}$ ,  $z_m = L_m^{-1} (\|r_0\|_2 e_1)$ ,  $P'_m = W_m (L_m^T)^{-1}$  and  $z'_m = (U_m^T)^{-1} (\|r'_0\|_2 e_1)$ . Because of the structure of  $U_m$ , which only has a nonzero diagonal and superdiagonal, we can easily calculate the elements of  $P_m$ . We see that  $u_{m-1,m} p_{m-1} + u_{m,m} p_m = v_m$ , where  $p_m$  and  $v_m$  are the last column of  $P_m$  and  $V_m$  respectively, and  $u_{i,j}$  is the  $i, j$ th element of  $U_m$ . Now it follows that

$$p_m = \frac{1}{u_{m,m}} (v_m - u_{m-1,m} p_{m-1}). \quad (3.17)$$

In addition, because of the structure of  $L_m$ , which only has a nonzero diagonal (consisting of only ones) and subdiagonal, we find

$$z_m = (z_{m-1}, \zeta_m)^T$$

in which  $\zeta_m = l_{m,m-1} \zeta_{m-1}$ , where  $l_{i,j}$  is the  $i, j$ th element of  $L_m$ . If we now substitute this back in the relation  $x_m = x_0 + P_m z_m$  found above, we get

$$x_m = x_0 + [P_{m-1}, p_m] \begin{bmatrix} z_{m-1} \\ \zeta_m \end{bmatrix} = x_0 + P_{m-1} z_{m-1} + \zeta_m p_m.$$

Noting that  $x_0 + P_{m-1} z_{m-1} = x_{m-1}$ , we finally find

$$x_m = x_{m-1} + \zeta_m p_m. \quad (3.18)$$

If we repeat those steps for the dual system, we find a similar relation

$$x'_m = x'_{m-1} + \zeta'_m p'_m. \quad (3.19)$$

This is the same derivation as used for IOM and DIOM in [35].

We now state two propositions that we will use in the derivation in the next section.

**Proposition 3.1.** *The pairs of the primary and dual direction vectors  $p_i$  and  $p'_j$  form a  $A^2$ -orthonormal set, i.e.  $p_i^T A^2 p_j = \delta_{i,j}$ .*

**Proof.**

$$\begin{aligned} (P'_m)^T A^2 P_m &= (W_m (L_m^T)^{-1})^T A^2 V_m U_m^{-1} \\ &= L_m^{-1} W_m^T A^2 V_m U_m^{-1} \\ &= L_m^{-1} T_m U_m^{-1} \quad (\text{using (3.9)}) \\ &= L_m^{-1} L_m U_m U_m^{-1} \\ &= I \end{aligned}$$

□

**Proposition 3.2.** *The pairs of the primary and dual residual vectors  $r_i$  and  $r'_i$  form a  $A$ -orthonormal set, i.e.  $r_i'^T A r_j = 0$  for  $i \neq j$ .*

**Proof.** Combining (3.7) and (3.12)-(3.13), we get

$$\begin{aligned} r_m &= b - A r_m \\ &= b - A x_0 - A V_m y_m \\ &= r_0 - V_m T_m y_m - \delta_{m+1} v_{m+1} e_m^T y_m \\ &= r_0 - r_0 - \delta_{m+1} v_{m+1} e_m^T y_m \\ &= -\delta_{m+1} v_{m+1} e_m^T y_m. \end{aligned}$$

In a similar way, we get for the dual system, using (3.8) and (3.15)-(3.16),

$$r'_m = -\beta_{m+1} w_{m+1} e_m^T y'_m.$$

Combining the above two relations with (3.9), we now find that  $r_i'^T A r_j = 0$  for  $i \neq j$ .  $\square$

## 4 The BiCOR method

We can now proceed in a similar way as the derivation of the Conjugate Gradient method. For the derivation of CG, see for example [40, 20]. Given an initial guess  $x_0$ , we get the coupled two-term recurrences

$$r_0 = b - A x_0, \quad p_0 = r_0, \quad (4.1)$$

$$x_{j+1} = x_j + \alpha_j p_j, \quad (4.2)$$

$$r_{j+1} = r_j - \alpha_j A p_j, \quad (4.3)$$

$$p_{j+1} = r_{j+1} + \beta_j p_j, \quad \text{for } j = 0, 1, \dots \quad (4.4)$$

where  $r_j = b - A x_j$  is the residual at iteration  $j$  and  $p_j$  is the search direction vector at iteration  $j$  as in (3.17). Here the vectors  $p_j$  are multiples of the vectors  $p_j$  as seen in 3.3. It's important to note that  $\alpha_j$  and  $\beta_j$  are different from the  $\alpha_j$  and  $\beta_j$  used in the previous section (3.2). This is done for consistency with the notation used in the derivation of other methods. The coupled two-term recurrences for the dual system are defined in a similar way:

$$r'_{j+1} = r'_j - \alpha_j A^T p'_j, \quad (4.5)$$

$$p'_{j+1} = r'_{j+1} + \beta_j p'_j, \quad \text{for } j = 0, 1, \dots \quad (4.6)$$

The parameters  $\alpha_j$  and  $\beta_j$  can be determined from the orthogonality relations

$$r_{j+1} \perp \mathcal{L}_m \text{ and } A p_{j+1} \perp \mathcal{L}_m$$

as found in section 3.2. Using propositions 3.1 and 3.2, we find the subspace  $A^T \mathcal{K}_m(A^T; r'_0)$  to be suitable, where  $r'_0 = P(A) r_0$  with  $P(t)$  an arbitrary polynomial in variable  $t$ . A common choice is  $r'_0 = r_0$ , but here we will use  $r'_0 = A r_0$ . If instead of  $A^T \mathcal{K}_m(A^T; r'_0)$  we choose for instance  $\mathcal{K}_m(A; r_0)$ , we get the CG method [21]. Further derivation using (4.1)-(4.6) gives us the following expressions for  $\alpha_j$  and  $\beta_j$ . For the full derivation see for example [28].

$$\alpha_j = \frac{r_j'^T A r_j}{p_j'^T A^2 p_j} \quad (4.7)$$

$$\beta_j = \frac{r_{j+1}'^T A r_{j+1}}{r_j'^T A r_j}. \quad (4.8)$$



Now, combining (4.1)-(4.8), we finally get the Biconjugate Biconjugate  $A$ -Orthogonal Residual method, or simply BiCOR [25, 28, 27, 8, 9]. The complete algorithm can be found in Algorithm 5. In the algorithm, we use the following notations: the dual vectors have a primed symbol, preconditioned vectors have a prefixed  $z$ , and a hat symbol is used for matrix-vector products.

---

**Algorithm 5** *Left preconditioned BiCOR method.*

---

```

1: Compute  $r_0 = b - Ax_0$  for some initial guess  $x_0$ .
2: Choose  $r'_0 = P(A)r_0$  such that  $\langle r'_0, Ar_0 \rangle \neq 0$ , where  $P(t)$  is a polynomial in  $t$ . (For example,  $r'_0 = Ar_0$ ).
3: for  $j = 1, 2, \dots$  do
4:   solve  $Mzr_{j-1} = r_{j-1}$ 
5:   if  $j=1$  then
6:     solve  $M^T zr'_0 = r'_0$ 
7:   end if
8:    $\hat{zr} = Azr_{j-1}$ 
9:    $\rho_{j-1} = \langle zr'_{j-1}, \hat{zr} \rangle$ 
10:  if  $\rho_{j-1} = 0$ , method fails
11:  if  $j = 1$  then
12:     $p_0 = zr_0$ 
13:     $p'_0 = zr'_0$ 
14:     $q_0 = \hat{zr}$ 
15:  else
16:     $\beta_{j-2} = \rho_{j-1} / \rho_{j-2}$ 
17:     $p_{j-1} = zr_{j-1} + \beta_{j-2} p_{j-2}$ 
18:     $p'_{j-1} = zr'_{j-1} + \beta_{j-2} p'_{j-2}$ 
19:     $q_{j-1} = \hat{zr} + \beta_{j-2} q_{j-2}$ 
20:  end if
21:   $\hat{q}'_{j-1} = A^T p'_{j-1}$ 
22:  solve  $M^T \hat{zq}'_{j-1} = \hat{q}'_{j-1}$ 
23:   $\alpha_{j-1} = \rho_{j-1} / \langle \hat{zq}'_{j-1}, q_{j-1} \rangle$ 
24:   $x_j = x_{j-1} + \alpha_{j-1} p_{j-1}$ 
25:   $r_j = r_{j-1} - \alpha_{j-1} q_{j-1}$ 
26:   $zr'_j = zr'_{j-1} - \alpha_{j-1} \hat{zq}'_{j-1}$ 
27:  check convergence; continue if necessary
28: end for

```

---

## 5 The CORS method

Using the same strategy used when deriving the CGS method from the BiCG method, see for example [40], we can derive a transpose-free variant of the BiCOR method, the Conjugate  $A$ -Orthogonal Residual Squared method (CORS) [25, 28, 27, 8, 9].

In the previous section, we could have written the representations of the vectors  $r_j, p_j, r'_j, p'_j$  at step  $j$  as the polynomial representations

$$\begin{aligned}
r_j &= \phi_j(A)r_0, & p_j &= \psi_j(A)r_0, \\
r'_j &= \phi_j(A^T)r'_0, & p'_j &= \psi_j(A^T)r'_0,
\end{aligned}$$

where  $\phi_j$  and  $\psi_j$  are Lanczos-type polynomials of degree less than or equal to  $j$  satisfying  $\psi_j(0) = 1$ . Substituting back in (4.7) and (4.8) gives us

$$\begin{aligned}
\alpha_j &= \frac{r_j^T Ar_j}{p_j^T A^2 p_j} = \frac{r_0^T A \phi_j^2(A) r_0}{r_0^T A^2 \psi_j^2(A) r_0} \\
\beta_j &= \frac{r_{j+1}^T Ar_{j+1}}{r_j^T Ar_j} = \frac{r_0^T A \phi_{j+1}^2(A) r_0}{r_0^T A \phi_j^2(A) r_0}.
\end{aligned}$$

Also note that from (4.3) and (4.4)  $\phi_j$  and  $\psi_j$  can be expressed recursively as

$$\begin{aligned}\phi_{j+1}(t) &= \phi_j(t) - \alpha_j t \psi_j(t), \\ \psi_{j+1}(t) &= \phi_{j+1}(t) + \beta_j \psi_j(t).\end{aligned}$$

Using the strategy mentioned above, we now get the CORS algorithm, as described in Algorithm 6.

---

**Algorithm 6** *Left preconditioned CORS method.*

---

```

1: Compute  $r_0 = b - Ax_0$  for some initial guess  $x_0$ .
2: Choose  $r'_0 = P(A)r_0$  such that  $\langle r'_0, Ar_0 \rangle \neq 0$ , where  $P(t)$  is a polynomial in  $t$ . (For example,  $r'_0 = Ar_0$ ).
3: for  $j = 1, 2, \dots$  do
4:   solve  $Mzr_{j-1} = r_{j-1}$ 
5:    $\hat{z}r = Azr_{j-1}$ 
6:    $\rho_{j-1} = \langle r'_0, \hat{z}r \rangle$ 
7:   if  $\rho_{j-1} = 0$ , method fails
8:   if  $j = 1$  then
9:      $e_0 = r_0$ 
10:    solve  $Mze_0 = e_0$ 
11:     $d_0 = \hat{z}r$ 
12:     $q_0 = \hat{z}r$ 
13:  else
14:     $\beta_{j-2} = \rho_{j-1} / \rho_{j-2}$ 
15:     $e_{j-1} = r_{j-1} + \beta_{j-2} h_{j-2}$ 
16:     $ze_{j-1} = zr_{j-1} + \beta_{j-2} f_{j-2}$ 
17:     $d_{j-1} = \hat{z}r + \beta_{j-2} g_{j-2}$ 
18:     $q_{j-1} = d_{j-1} + \beta_{j-2} (g_{j-2} + \beta_{j-2} q_{j-2})$ 
19:  end if
20:  solve  $Mzq = q_{j-1}$ 
21:   $\hat{z}q = Azq$ 
22:   $\alpha_{j-1} = \rho_{j-1} / \langle r'_0, \hat{z}q \rangle$ 
23:   $h_{j-1} = e_{j-1} - \alpha_{j-1} q_{j-1}$ 
24:   $f_{j-1} = ze_{j-1} - \alpha_{j-1} zq$ 
25:   $g_{j-1} = d_{j-1} - \alpha_{j-1} \hat{z}q$ 
26:   $x_j = x_{j-1} + \alpha_{j-1} (2ze_{j-1} - \alpha_{j-1} zq)$ 
27:   $r_j = r_{j-1} - \alpha_{j-1} (2d_{j-1} - \alpha_{j-1} \hat{z}q)$ 
28:  check convergence; continue if necessary
29: end for

```

---

From our experiments, we find that CORS is highly competitive to all other popular algorithms (see section 7). However, like the CGS method, it is based on squaring the residual, which might result in a substantial buildup of rounding errors and worse approximate solutions, or possibly even overflow. This also means that CORS might in general need more time to complete a calculation than other methods, if they both succeed.

## 6 Computational aspects

### 6.1 Preconditioning

In our experiments, we use preconditioners constructed by ILUPACK [5]. The algorithms of ILUPACK compute an incomplete LU-factorisation  $A = LDU + E$ . Here  $L$  is a lower triangular matrix with unit diagonal,  $D$  is a diagonal matrix and  $U$  is an upper triangular matrix with unit diagonal.  $LDU$  is an approximation of the standard LU-factorisation that can be used as a preconditioner. Furthermore  $\|E\| < \tau$  where  $\tau$  is the drop tolerance. The matrices  $L$ ,  $D$  and

$U$  are easily implicitly computed. In the case of ILUPACK, the diagonal matrix is not a real diagonal matrix. ILUPACK computes

$$\tilde{P}^T A \tilde{Q} = \begin{pmatrix} B & F \\ E & C \end{pmatrix} \approx \begin{pmatrix} L_B & 0 \\ L_E & I \end{pmatrix} \begin{pmatrix} D_B & 0 \\ 0 & S_C \end{pmatrix} \begin{pmatrix} U_B & U_F \\ 0 & I \end{pmatrix}$$

and then uses the inverse

$$(\tilde{P}^T A \tilde{Q})^{-1} = \begin{pmatrix} B & F \\ E & C \end{pmatrix}^{-1} \approx \begin{pmatrix} \tilde{B}^{-1} & 0 \\ 0 & 0 \end{pmatrix} + \begin{pmatrix} \tilde{B}^{-1} F \\ I \end{pmatrix} S_C^{-1} (-E \tilde{B}^{-1} \quad I)$$

where  $\tilde{B} = L_B D_B U_B$  [4].

## 6.2 Stopping criteria

An iterative method will never provide an exact solution with a zero residual,  $r = b - Ax = 0$ , unless of course  $b$  is equal to zero. For this reason, we have to choose a good stopping criterion that we can use in all the different solvers we use for testing. A really small relative error  $\|x - y\|/\|x\|$  with respect to the approximate solution  $y$  is usually enough, but this can not always be achieved. Also, since we do not have the actual solution, we can not explicitly calculate the relative error. Therefore, our stopping criterion will be based on the backward error analysis introduced by Wilkinson [41].

A calculated solution  $\hat{x}$  of a system  $Ax = b$  can be seen as the (exact) solution of the perturbed problem

$$(A + \delta A)\hat{x} = (b + \delta b).$$

The so called backward error measures the distance between the data of the original system and the perturbed system. The uncertainties in the data can either be due to measurements, or due to accumulation or propagation of roundoff errors [22]. If the backward error is not larger than those uncertainties, we may assume that the approximation is accurate. Componentwise perturbations and normwise perturbations can be used to calculate backward error. These lead to explicit formulas to calculate the backward error. It is generally accepted that for iterative methods, the use of normwise perturbations is appropriate [17]. We use this strategy to stop our solvers.

At iteration  $j$  of an iterative method, we compute an approximation  $x_j$  of the actual solution  $x = A^{-1}b$ . We can see  $x_j$  as the solution of the perturbed problem  $(A + \delta A)x_j = (b + \delta b)$ . We introduce

$$\begin{aligned} \eta_j &= \min \{ \epsilon > 0 : (A + \delta A)x_j = (b + \delta b), \|\delta A\|_2 \leq \epsilon\alpha, \|\delta b\|_2 \leq \epsilon\beta \} \\ &= \frac{\|b - Ax_j\|_2}{\alpha\|x_j\|_2 + \beta} \end{aligned}$$

as the normwise backward error [22]. When the machine precision has been reached by our method, the method does not converge any further, so at best, the backward error is as small as the machine precision. In the testing application, we stop when  $\eta \leq 10^{-10}$ .

Common choices for  $\alpha$  and  $\beta$  are, respectively,  $\|A\|_2$  and  $\|b\|_2$ . In this case,  $\eta_j$  is called the normwise relative backward error [22]. For the sake of simplicity, however, we have chosen to use  $\alpha = 0$  and  $\beta = \|b\|_2$  in the testing application.

Value	Meaning
-1	An error occurred
1	Convergence has been achieved or the user may check for convergence
2	The user must perform a matrix-vector product
3	The user must perform the preconditioning operation

Table 2: Return values of `IACT` for our CORS implementation

### 6.3 Implementational aspects

Implementations of iterative methods basically require vector updates, scalar products and matrix-vector products. The first two are standard routines that are implemented in the BLAS library, but for the matrix-vector products, the user might want to provide their own implementations. This is mainly because matrices can be stored in various ways. Sometimes matrices are not even stored explicitly, but only as a subroutines. The same holds for preconditioners. ILUPACK for example does not provide an explicit matrix to use for preconditioning.

We could just let the user implement matrix-vector products in the code themselves, but that is not very user-friendly. For this reason, we allow the user to specify their own matrix-vector products and preconditioning operations, using reverse communication [12]. Reverse communication is commonly used in FORTRAN implementations of iterative methods, for example in the Harwell Subroutine Library (HSL) [23]. Here we explain how it works.

In the call to the iterative method, several variables are provided. One of those is the reverse communication variable. Once you call the function for the first time, it has to have a certain value, so the method knows it's the first time you call it. In our case the variable is `IACT` and this default value is zero. Other values of the reverse communication variable tell you to perform for instance a matrix-vector product, a preconditioning operation, or they tell you that an error occurred or convergence has been achieved. See for an example of the values of `IACT` Table 2.

One of the other variables is an array of vectors the method will use during the process. Once the user is told to perform for example a matrix-vector product, other variables are used to tell the user which vectors in the array to use. In our case those variables are `LOCY` and `LOCZ`. Meaning the location of  $y$  and  $z$  coming from the assignment  $y = Az$ . So one reads from the `LOCZ`-th vector and writes to the `LOCY`-th vector. Once the user performed the operation he is supposed to perform, the same subroutine is called again with the same argument. This process is repeated until convergence is observed by either the user or the algorithm itself, depending on whether the user wants to check or not.

The reverse communication method is overall very fast, because no memory has to be allocated during any of the operations. The user only has to perform a certain operation. It's also very user friendly, because the user can use any implementation of a matrix-vector product, preconditioning operation, or convergence check.

A different way to implement this in FORTRAN would be allowing the user to pass a function or subroutine to the subroutine that is then called by the subroutine itself, but this limits the user to passing only subroutines or functions that require a set amount of variables, whereas the user probably wants to pass more variables. A way to do this in object oriented languages, like C++, is by overloading operators.

## 7 Numerical experiments

### 7.1 Information about the experiments

In our experiments, we consider a collection of various matrices available from the University of Florida Sparse Matrix Collection from Tim Davis [10]. The matrices we used are a reasonable representation of all nonsymmetric and real matrices available in the collection, covering every field of research in the collection. To analyse the performance of BiCOR and CORS, we compared them to the popular methods BiCG, BiCGSTAB, CGS, GMRES, BiCGSTAB( $\ell$ ), QMR and TFQMR. For GMRES we used a value of restart equal to 100. This reduced the memory needed to run the solver on the largest problems. We chose the value  $\ell = 3$  for BiCGSTAB( $\ell$ ), because this yielded the best results.

We implemented the BICOR and CORS methods in FORTRAN 77 by ourselves. For BiCG, BiCGSTAB, CGS and GMRES we used implementations from the Harwell Subroutine Library (HSL) [23]. The implementation of BiCGSTAB( $\ell$ ) was obtained from Van der Vorst's website, [16] and the implementations of QMR and TFQMR came from QMRPACK [19].

The tests were run on a PC equipped with an AMD Athlon™ 7850 Dual-Core Processor running at 2,8 GHz and 4GB 800 MHz DDR2-RAM. Our code was compiled with the GNU FORTRAN compiler (gfortran) version 4.5.2 that came with Ubuntu 11.04. The implementations of all the solvers we tested were in FORTRAN 77, the testing application calling those implementations was in FORTRAN 2003 to allow us to keep running the application without having to restart for every other preconditioner or matrix. This needed memory allocation and making sure the results were saved on the local disk required flushing the file after every solver completed.

To read the matrices from our hard-drive, we downloaded the matrices in MatrixMarket format [3]. To store our matrices in the main memory, we used the compressed sparse row format. We used SPARSKIT [34] to convert the matrices from the coordinate format used in the MatrixMarket script to the compressed sparse row format. The matrix-vector product and transpose matrix-vector product were performed by AMUX and ATMUX in SPARSKIT. The preconditioning operation was performed by ILUPACK. Those libraries, as well as the implementations of all the different solvers, needed a BLAS implementation, for which we used the ACML library optimised for AMD processors.

The data we gathered from running the different solvers included the amount of time it took to solve a problem, the amount of matrix-vector products, and, if a solver did not complete, the reason why. We ran every solver six times if it completed the first time to be able to get rid of any flaws caused by processes running in the background. To minimise this effect, no applications were started other than the default startup applications, excluding Ubuntu One, and including Dropbox to make sure results were not lost, a terminal, and nautilus. During the process, the CPU was monitored to make sure nothing interfered with the testing application. As a result, the testing application ran at 99%-100% of one core essentially all of the time.

To make a better selection of the matrices to test, we excluded matrices that completed faster than 0.05 seconds, because the results we got from the CPU.TIME routine were only accurate to up to 2 decimals. Additionally, we excluded those matrices of which the problems took too long to solve, i.e. took more than 20000 matrix-vector operations, for none of the iterative solvers using the best preconditioner. The reason we stopped at 20000 matrix-vector products is that for the bigger problems it would take five days or more to get up to  $n$  iterations for any solver. This would mean that we would be done maybe two years from now. The reason we excluded the problems that failed to complete for every solver was so we did not have to

rerun them for worse preconditioners.

We also checked if a result of a given solver on a given problem was much different from the average of the other runs with the same solver on the same problem. If it was more than 10% off, it was excluded from the results. This rarely ever happened.

## 7.2 Data analysis

To analyse the data we gathered from running the testing application, we make use of performance profiles of the computation time and the amount of matrix-vector products as suggested by Dolan and Moré in [11]. In the performance profiles, we can see what solver is most likely to solve a certain problem after a certain amount of time or with a certain amount of matrix-vector products compared to the other solvers. The best solver for every matrix gets a value of one associated with them, and the other solvers get a value greater than one, that is the ratio between this solver and the best solver. So the performance ratio of a solver  $s$  on a problem  $p$  is given by

$$r_{p,s} = \frac{t_{p,s}}{\min_s t_{p,s}}$$

and the cumulative distribution function for the performance ratio is given by

$$\rho_s(\tau) = \frac{1}{n_p} \text{size}\{p : r_{p,s} \leq \tau\}$$


where  $n_p$  is the total amount of problems we tested. So at  $\tau = \tau_1$ , a certain solver  $s$  has a probability  $\rho_s(\tau_1)$  of solving a problem at a ratio  $\tau_1$  worse than the fastest solver.  $\rho(1)$ , is of particular interest, because we can see there how many times a solver was the best.

If a solver does not solve a problem, the ratio  $r_M$  is assigned. This ratio should be higher than the highest ratio found for any solver on any problem that did not fail. In this way, the solver will still have a value assigned for the certain problem where it failed, but we simply will not plot for  $\tau \geq r_M$ . So we will see the solver that solved the most problems overall on top when we look at the far right of the plot.


## 7.3 Results

We ran our preliminary tests with three different preconditioners constructed with drop tolerances of 0.1, 1.0 and 10.0 on a total of more than 100 matrices, but ended up with only 72 matrices that satisfied our criteria. We solved the linear system using preconditioning from the right. This means that we solved the system  $AM_2^{-1}y = b$  with our solution  $x = M_2^{-1}y$ . From (2.2), we see that we did not have to adjust our stopping criterion to work on the preconditioned system. If we would have used preconditioning from the left, in a real implementation one would have had to adjust every solver to have a stopping criterion based on the preconditioned system. We could, however, in reverse communication just replace the matrix-vector products with a preconditioning operation and a matrix-vector product, and the preconditioning operation with a vector copy.

The time it took to complete the experiments with 9 solvers on 72 matrices was over 120 hours. We analyse the results in the next sections.

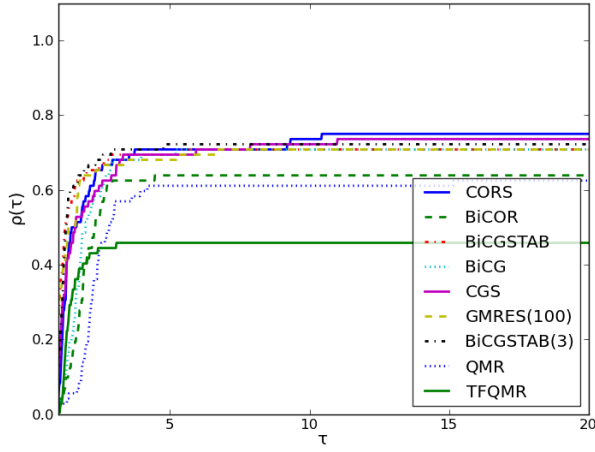


(a) Ratio of CPU time




(b) Ratio of matrix-vector products

Figure 2: Performance profiles with a tolerance of 0.1




(a) Ratio of CPU time




(b) Ratio of matrix-vector products

Figure 3: Performance profiles with a tolerance of 1.0



(a) Ratio of CPU time



(b) Ratio of matrix-vector products

Figure 4: Performance profiles with a tolerance of 10.0

### 7.3.1 Speed

As one can see in Figures 2-4, our tests revealed that GMRES and BiCGSTAB were in general the fastest solvers. In terms of matrix-vector products, GMRES was of course the fastest, because it uses only one matrix-vector product per iteration, where the other methods use two. In terms of time however, GMRES became better compared to the other solvers for better preconditioners. For the preconditioner constructed with a tolerance of 0.1, GMRES had 35 wins where BiCGSTAB had 24, with a tolerance of 1.0, they both had 19 wins, and with a tolerance of 10, BiCGSTAB had 19 wins and GMRES 15. Because GMRES got worse for sparser preconditioners. This also meant that the other solvers got relatively more wins.

Here it must be noted that we used a restart value of 100 for GMRES, but for really large problems, where memory use is an issue, we would not be able to use such a high value for the restart. In such a case, it would be more fair to have a value of restart that makes the memory use of GMRES similar to that of the other methods. We tried this, but this gave such bad results (worse than TFQMR), that we decided to use a value of 100.

If we do not only look at the winners, but at a slightly bigger region of interest, say  $\tau \leq 2$ , we see that CORS, BiCGSTAB, GMRES and CGS are the most competitive solvers. For the best preconditioner, the one with a drop tolerance of 0.1, CORS is even on top after  $\tau \approx 1.7$  (see Figure 5).

We also see that CORS is considerably faster in terms of matrix-vector products (see Figures 2-4). It performs a bit worse in terms of time is mainly due to the amount of scalar times a vector plus a different vector operations, or simply Scalar A X Plus Y (SAXPY) operations. Those are the most expensive operations done in the algorithms themselves. CORS uses 12 of such per iteration where BiCOR for example only uses 6.

We now classify the solvers according to the information we gathered about the different problems, see Appendix A.2. In this case, we excluded solvers like QMR and TFQMR from this analysis, because they performed very badly. We also only name the solvers where we saw something notable. First, we see that the bigger the problem, the better CORS and BiCGSTAB perform and the worse GMRES performs. This is mainly of interest, because one tends to use iterative methods only for bigger problems. For smaller problems one could as well just use direct methods.




Figure 5: Performance profile with a tolerance of 0.1 on a smaller region: Ratio of CPU time



Solver	breakdown	iterations	NAN
CORS	1	2	0
BiCOR	11	3	0
BiCGSTAB	7	0	0
BiCG	5	2	0
CGS	0	7	0
GMRES(100)	0	12	0
BiCGSTAB(3)	0	8	2
QMR	0	13	0
TFQMR	0	25	0

Table 4: Failures with a preconditioner tolerance of 0.1

We also investigated if the percentage of pattern symmetry and value symmetry mattered. Here we see that GMRES performs better for the highly nonsymmetric problems, and CORS and BiCGSTAB for the more symmetric problems. We also noted that for diagonal dominant matrices, BiCGSTAB was better and GMRES worse.

Now we get to the kind of problems. Here we see that GMRES performed better for economic, and semiconductor device problems. BiCGSTAB performed better for circuit simulation, computational fluid dynamics, and semiconductor device problems. We also saw that CORS performed considerably well on circuit simulation problems, and that BiCOR and BiCG were very efficient for the electromagnetics problems.

### 7.3.2 Reliability

If we look at a region of large values of  $\tau$  in the performance profiles in Figure 2, we find that CORS ends up solving 4 more problems than the two next best solvers, BiCG and CGS. For this preconditioner, CORS only fails to solve three problems. CORS also ends up being on top for the preconditioner constructed with a tolerance of 1.0. For the preconditioner constructed with a tolerance of 10, we see that BiCGSTAB( $\ell$ ) ends up on top. CORS is the second best solver, and its performance comes very close to BiCGSTAB( $\ell$ ) if we look at values of  $\tau$  greater than 20, but this is not shown in the performance profiles. If we look at the other preconditioner tolerances, we see that BiCGSTAB( $\ell$ ) does not even come close to the performance of CORS.

We’re now interested in what happens when CORS fails, since it’s the solver that has the least amount of failures. We find that for the preconditioner constructed with a tolerance of 0.1, CORS exceeds the maximum amount of iterations twice, and breaks down once. A breakdown in the implementations of CORS, BiCOR, BiCGSTAB, CGS and BiCG means that  $|\rho_{j-1}| < un$  and  $|\rho_{j-1}| < u\|r_{j-1}\|_2\|r'_{j-1}\|_2$  where  $n$  is the size of the problem,  $u$  is the machine precision and the other variables as in Algorithm 5. This is as it is adopted in the HSL. The GMRES method, being an optimal method, can not break down, so for GMRES, we do not see any breakdowns. BiCGSTAB( $\ell$ ) returned quite a lot of not-a-number answers, which might be due to the breakdown implementation which is different from that of the HSL.

So let’s look at the only case where CORS broke down for the best tolerance we tested. This was on the `torso1` matrix. For this matrix, we find that not only CORS, but also BiCGSTAB broke down, BiCOR and BiCG converged, and the others exceeded the maximum amount of allowed iterations. If we look at the convergence history of the 2-norm of the residual of CORS, BiCOR, BiCGSTAB, BiCG, CGS and GMRES, we see that for BiCOR and BiCG the residual

started reducing at a nice rate after about 200 iterations. For GMRES, the residual stayed constant after about 80 iterations, and for CGS the residual heavily fluctuated somewhere above  $10^5$ . The residual of BICGSTAB suddenly increased after about 300 iterations and the method broke down after doing a few more steps. CORS at first showed about the same behaviour as CGS, then fluctuated less heavily, but did still not converge, and after that, CORS broke down.

The two matrices where CORS took too many iterations were the `cryg10000` and `invextr1_new` matrices. For the first one, CORS did not seem to converge at all, but for the last one, CORS converged at steady rate, but unfortunately not fast enough to complete within the set maximum amount of iterations, as can be seen in Figure 6(a). What we can also see in this figure is the relatively wild behaviour of CGS due to the squaring of the residual, and the behaviour of GMRES, which usually converges steadily, but in this case not at all. What we can conclude, is that CORS is able to solve most problems, and therefore is the most robust method for this preconditioner.

In Figure 6(b), we again see that here CGS is a lot wilder than CORS, which seems to be the usual behaviour. This is probably the reason why CORS is more stable than CGS, while they are based on the same ideas. In this figure we also see the breakdown of BiCGSTAB. The method first behaves like other solvers, but then sees a sudden increase in the residual after which the residual stays the same and then the method breaks down. This is the standard behaviour we observed for this method and other methods.

The default drop tolerance used by ILUPACK is 0.01 instead of the 0.1 we used as lowest tolerance. We did this mainly to make sure the solvers did not complete too quickly, and so we could find out which solver was most robust. The construction of the preconditioner for the bigger problems only took a small amount of time compared to solving the problem itself, so it would be reasonable to use a better preconditioner. So we ran some more tests. We tried several better preconditioners on the matrices that failed for all methods for the preconditioner with a drop tolerance of 0.1. In those tests, CORS again turned out to be the most reliable method in every test. Some other solvers were able to compete with CORS in some tests, but CORS was the only one that turned out to be the most reliable in every single test.

If we look at BiCOR we see that a lot of breakdowns occur. We checked that in all cases the value of  $\rho$  was indeed smaller than the machine precision. Two new methods, BiCOR Stabilised (BiCORSTAB) [28] and Composite Step BiCOR (CSBiCOR) [26], have already been developed to prevent those failures. We however did not have a chance to test those methods.

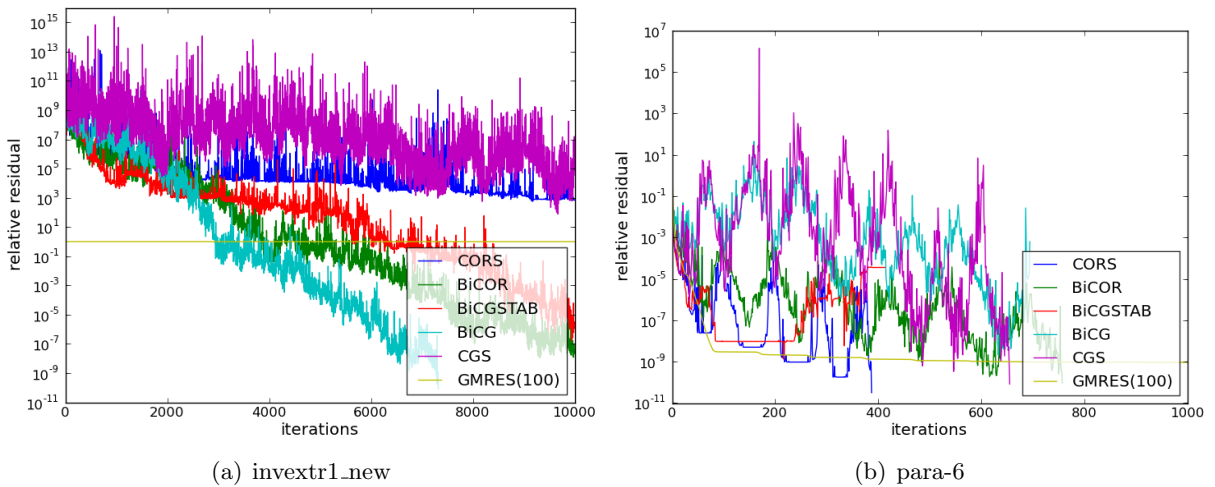


Figure 6: Convergence history of the relative residual on two problems

The last question we tried to answer was how the distribution of the eigenvalues of the preconditioned matrix affected convergence. If the preconditioner is good, then  $AM^{-1}$  is close to identity, so we may expect that most of the eigenvalues are close to one. We calculated the eigenvalues of the 20 smallest matrices, but it was not possible to draw any conclusions.




Figure 7: Typical distribution of eigenvalues, in this case of the `powersim` matrix.

## 8 Conclusion

When we started comparing CORS and BiCOR to other iterative methods, we had no idea whether they were competitive or not. In our experiments we found that BiCOR broke down many times, and therefore is not very attractive for solving realistic applications. We also found that the BiCGSTAB method and GMRES method with sufficiently large restart, are the most popular methods in use today for a reason: they turned out to be the fastest methods. In terms of stability however, CORS proves to be the best. It might not be as fast, mostly due to the larger amount of SAXPY operations, but reliability comes with a cost. We also see this when we compare for example  $\text{BiCGSTAB}(\ell)$  to normal BiCGSTAB, which is generally a lot faster than  $\text{BiCGSTAB}(\ell)$ .

The most interesting case when using iterative methods is a large problem with a good preconditioner. Bigger problems are more vulnerable to the performance of the methods, simply because they take longer to solve. Also, because ILUPACK constructs the preconditioners quite fast, with an amount of nonzeros of the order of the problem itself, one most likely wants to use a better preconditioner. We found CORS to excel in both cases: it was better for better preconditioners, and also faster for bigger problems in comparison with other solvers.

We conclude that the CORS method turns out to be a valuable addition to the long list of iterative methods already available.

## 9 Acknowledgments

I would really like to thank Bruno Carpentieri for all his support and help during every step of the process and for developing BiCOR and CORS in the first place to make this research possible. My thanks also go to Matthias Bollhöfer (Institute of “Computational Mathematics”,

Technische Universität Braunschweig, Germany) for his support in the usage of the ILUPACK software. And finally, I would like to thank Sjoerd Meesters, Paulus Meessen and Jeroen Lanting for thoroughly reading the manuscript and giving advice on how to improve it, and Diederik Perdok, who read the manuscript with no knowledge of numerical mathematics in general, but still gave really good comments.

## References

- [1] W. E. Arnoldi. The principle of minimized iterations in the solution of the matrix eigenvalue problem. *Quarterly of Applied Mathematics*, 9:17–29, 1951.
- [2] R. Barrett, M. Berry, T. Chan, J. Demmel, J. Donato, J. Dongarra, V. Eijkhout, R. Pozo, C. Romine, and H. van der Vorst. *Templates for the Solution of Linear Systems: Building Blocks for Iterative Methods*. SIAM, Philadelphia, 1993. Obtainable from [research.att.com/netlib/linalg](http://research.att.com/netlib/linalg) using ftp.
- [3] R. F. Boisvert, R. Pozo, K. Remington, R. F. Barrett, and J. J. Dongarra. Matrix market: A web resource for test matrix collections. In *The Quality of Numerical Software: Assessment and Enhancement*, pages 125–137. Chapman & Hall, 1997.
- [4] M. Bollhöfer and Y. Saad. Multilevel preconditioners constructed from inverse-based ILUs. 27(5):1627–1650, 2006.
- [5] M. Bollhöfer, Y. Saad, and O. Schenk. ILUPACK — preconditioning software package, June 2011. <http://ilupack.tu-bs.de/>. Release 2.4.
- [6] C. Brezinski and L. Wuytack. *Projection methods for systems of equations*. North Holland, 1997.
- [7] R. L. Burden and J. D. Faires. *Numerical Analysis*. Thompson, 8 edition, 2005.
- [8] B. Carpentieri, Y.-F. Jing, and T.-Z. Huang. The BiCOR and CORS algorithms for solving nonsymmetric linear systems. *SIAM J. Scientific Computing*, 2011. In press.
- [9] B. Carpentieri, Y.-F. Jing, T.-Z. Huang, W.-C. Pi, and X.-Q. Sheng. A novel family of iterative solvers for Method of Moments discretizations of Maxwells equations. In L. Gürel, editor, *CEM’11 Computational Electromagnetics*, pages 85–90. Bilkent University, Computational Electromagnetics Research Center, August 2011.
- [10] T. A. Davis. University of florida sparse matrix collection. Technical report, 1994.
- [11] E. D. Dolan and J. J. More. Benchmarking optimization software with performance profiles. *Math. Programming, Ser. A*, 91:201–212, 2002.
- [12] J. Dongarra, V. Eijkhout, and A. Kalhan. Reverse communication interface for linear algebra templates for iterative methods. Technical Report UT-CS-95-291, May 1995.
- [13] J. J. Dongarra, I. S. Duff, D. C. Sorensen, and H. A. van der Vorst. *Numerical linear algebra for high-performance computers*, volume 7 of *Software, Environments, and Tools*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 1998.
- [14] V. Faber and T. Manteuffel. Necessary and sufficient conditions for the existence of a conjugate gradient method. *SIAM Journal on Numerical Analysis*, 21(2):352–362, Apr. 1984.

- [15] R. Fletcher. *Conjugate gradient methods for indefinite systems*, volume 506 of *Lecture Notes Math.*, pages 73–89. Springer-Verlag, Berlin, 1976.
- [16] D. R. Fokkema. Bicgstab(ell), full version. <http://www.staff.science.uu.nl/~vorst102/software.html>.
- [17] V. Frayssé, L. Giraud, S. Gratton, and J. Langou. A set of GMRES routines for real and complex arithmetics on high performance computers. *ACM Trans. Math. Softw.*, 31(2):228–238, 2005.
- [18] R. W. Freund and N. M. Nachtigal. QMR: A quasi-minimal residual method for non-Hermitian linear systems. *Numerische Mathematik*, 60:315–340, 1991.
- [19] R. W. Freund and N. M. Nachtigal. QMRPACK: A package of QMR algorithms. *ACM Trans. Math. Softw.*, 22(1):46–77, 1996.
- [20] A. Greenbaum. *Iterative Methods for Solving Linear Systems*. SIAM, Philadelphia, 1997.
- [21] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49(6):409–436, 1952.
- [22] N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2002.
- [23] HSL(2011). A collection of fortran codes for large scale scientific computation. <http://www.hsl.rl.ac.uk/>.
- [24] I. C. F. Ipsen and C. D. Meyer. The idea behind Krylov methods. *The American Mathematical Monthly*, 105(10):889–899, 1998.
- [25] Y.-F. Jing, B. Carpentieri, and T.-Z. Huang. Experiments with Lanczos biconjugate A-orthonormalization methods for MoM discretizations of Maxwell’s equations. *Progress In Electromagnetics Research, PIER 99*, pages 427–451, 2009.
- [26] Y.-F. Jing, T.-Z. Huang, B. Carpentieri, and Y. Duan. Investigating the composite step biconjugate A-orthogonal residual method for non-hermitian linear systems in Electromagnetics. In L. Gürel, editor, *CEM’11 Computational Electromagnetics*, pages 80–84. Bilkent University, Computational Electromagnetics Research Center, August 2011.
- [27] Y.-F. Jing, T.-Z. Huang, Y. Duan, and B. Carpentieri. A comparative study of iterative solutions to linear systems arising in quantum mechanics. *Journal of Computational Physics*, 229:8511–8520, November 2010.
- [28] Y.-F. Jing, T.-Z. Huang, Y. Zhang, L. Li, G.-H. Cheng, Z.-G. Ren, Y. Duan, T. Sogabe, and B. Carpentieri. Lanczos-type variants of the COCR method for complex nonsymmetric linear systems. *Journal of Computational Physics*, 228(17):6376–6394, 2009.
- [29] C. Lanczos. An iteration method for the solution of the eigenvalue problem of linear differential and integral operators. *J. Res. Nat. Bur. Standards*, 45:255–282, 1950.
- [30] C. Lanczos. Solution of systems of linear equations by minimized iterations. *J. Res. Natl. Bur. Stand.*, 49:33–53, 1952.
- [31] S. J. Leon. *Linear algebra with applications*. Prentice-Hall, pub-PH:adr, seventh edition, 2006.

- [32] B. N. Parlett, D. R. Taylor, and Z. A. Liu. A look-ahead Lanczos algorithm for unsymmetric matrices. *Math. Comp.*, 44:105–124, 1985.
- [33] A. Quarteroni, R. Sacco, and F. Saleri. *Numerical mathematics*. Springer, New York, 2000.
- [34] Y. Saad. Sparskit: a basic tool kit for sparse matrix computations - version 2, 1994.
- [35] Y. Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.
- [36] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM J. Scientific and Statistical Computing*, 7:856–869, 1986.
- [37] G. L. G. Sleijpen and D. R. Fokkema. BiCGstab(L) for linear equations involving unsymmetric matrices with complex spectrum. *Elect. Trans. Numer. Anal.*, 1:11–32, 1993.
- [38] P. Sonneveld. CGS, a fast Lanczos-type solver for nonsymmetric linear systems. *SIAM J. Scientific and Statistical Computing*, 10:36–52, 1989.
- [39] H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems. *SIAM J. Scientific and Statistical Computing*, 13:631–644, 1992.
- [40] H. A. van der Vorst. *Iterative Krylov methods for large linear systems*, volume 13 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, UK, 2003.
- [41] J. H. Wilkinson. *Rounding Errors in Algebraic Processes*. Notes on Applied Science No. 32, Her Majesty’s Stationery Office, London, 1963. Also published by Prentice-Hall, Englewood Cliffs, NJ, USA. Reprinted by Dover, New York, 1994.

# Appendices

In the appendices, one can find a description of all problems used, the implementation and documentation of the CORS and BiCOR methods, and the two main programs used in the analysis of the results. In addition to this, two Python modules were written, 5 more Python scripts, and 4 more Fortran applications. Since the appendix would be a lot longer if those were also included, and does not really add any valuable content to the thesis, those were left out.

## **A Problems**

### **A.1 Problem types**

List taken from <http://www.cise.ufl.edu/research/sparse/matrices/kind.html>

#### **A.1.1 Problems with 2D/3D geometry**

- 2D/3D problem
- acoustics problem
- computational fluid dynamics problem
- computer graphics/vision problem
- electromagnetics problem
- materials problem
- model reduction problem
- robotics problem
- semiconductor device problem
- structural problem
- thermal problem

#### **A.1.2 Problems that normally do not have 2D/3D geometry**

- chemical process simulation problem
- circuit simulation problem
- counter-example problem
- economic problem
- frequency-domain circuit simulation problem
- least squares problem
- linear programming problem
- optimization problem
- power network problem
- statistical/mathematical problem
- theoretical/quantum chemistry problem
- combinatorial problem
- graph problems

### **A.2 Problem list**



matrix name	number of rows	nonzeros	nonzero pattern symmetry	numeric value symmetry	row diagonal dominance	column diagonal dominance	kind
torso1	116,158	8,516,500	42%	0%	0.00%	0.15%	2D/3D problem
shermanACb	18,510	145,149	15%	3%	38.10%	37.36%	2D/3D problem
av41092	41,092	1,683,902	0%	0%	0.00%	0.00%	2D/3D problem
Baumann	112,211	748,331	100%	0%	16.58%	97.05%	2D/3D problem
heart3	2,339	680,341	100%	0%	0.00%	0.00%	2D/3D problem
chem_master1	40,401	201,201	100%	0%	1.98%	100.00%	2D/3D problem
e40r0100	17,281	553,562	31%	0%	0.14%	0.14%	2D/3D problem
Zd_Jac3	22,835	1,915,726	0%	0%	0.00%	0.00%	chemical process simulation problem
std1_Jac2_db	21,982	498,771	33%	0%	66.96%	54.41%	chemical process simulation problem
memplus	17,758	99,147	100%	50%	75.96%	87.85%	circuit simulation problem
ASIC_320k	321,821	1,931,828	100%	36%	71.29%	71.96%	circuit simulation problem
hcircuit	105,676	513,072	100%	20%	83.19%	84.66%	circuit simulation problem
scircuit	170,998	958,936	100%	80%	97.64%	97.39%	circuit simulation problem
ASIC_680k	682,862	2,638,997	100%	0%	85.46%	94.34%	circuit simulation problem
circuit_3	12,127	48,137	77%	30%	61.21%	63.71%	circuit simulation problem
transient	178,866	961,368	100%	24%	90.18%	90.74%	circuit simulation problem
trans4	116,835	749,800	85%	30%	57.54%	52.20%	circuit simulation problem sequence
lung2	109,460	492,564	57%	0%	49.61%	49.61%	computational fluid dynamics problem
airfoil_2d	14,214	259,688	98%	0%	4.48%	18.01%	computational fluid dynamics problem
atmosmodl	1,489,752	10,319,760	100%	67%	100.00%	100.00%	computational fluid dynamics problem
Ill.Stokes	20,896	191,368	99%	33%	17.67%	17.58%	computational fluid dynamics problem
atmosmodd	1,270,432	8,814,880	100%	67%	100.00%	100.00%	computational fluid dynamics problem
goodwin	7,320	324,772	64%	0%	4.09%	0.17%	computational fluid dynamics problem
poisson3Db	85,623	2,374,949	100%	0%	2.04%	10.50%	computational fluid dynamics problem
invextr1_new	30,412	1,793,881	97%	72%	2.98%	5.29%	computational fluid dynamics problem
GT01R	7,980	430,909	88%	0%	3.72%	4.11%	computational fluid dynamics problem
raefsky1	3,242	293,409	100%	9%	25.87%	25.87%	computational fluid dynamics problem sequence
cage11	39,082	559,722	100%	18%	97.40%	92.67%	directed weighted graph
language	399,130	1,216,334	6%	0%	81.25%	78.62%	directed weighted graph
psmigr_2	3,140	540,022	48%	0%	0.00%	0.00%	economic problem
g7jac160	47,430	564,952	3%	0%	25.94%	28.87%	economic problem
g7jac060	17,730	183,325	4%	0%	25.87%	29.20%	economic problem
mark3jac040	18,289	106,803	7%	1%	23.17%	27.14%	economic problem
jan99jac040	13,694	72,734	0%	0%	29.89%	65.81%	economic problem
mark3jac080sc	36,609	214,643	7%	1%	17.44%	29.55%	economic problem

g7jac140	41,490	488,633	3%	0%	25.95%	28.91%	economic problem
fp	7,548	834,222	76%	0%	3.11%	3.11%	electromagnetics problem
dw8192	8,192	41,746	96%	92%	10.64%	10.64%	electromagnetics problem
utm5940	5,940	83,842	53%	0%	12.82%	15.57%	electromagnetics problem
tmt_unsym	917,825	4,584,801	100%	0%	49.95%	49.96%	electromagnetics problem
viscoplastic2	32,769	381,326	57%	0%	43.94%	8.07%	materials problem
cryg10000	10,000	49,699	100%	0%	54.72%	1.97%	materials problem
inlet	11,730	328,323	61%	0%	3.30%	1.50%	model reduction problem
flowmeter5	9,669	67,391	100%	6%	73.72%	73.62%	model reduction problem
chipcool1	20,082	281,150	100%	9%	4.17%	4.95%	model reduction problem
crashbasis	160,000	1,750,416	55%	0%	49.73%	95.75%	optimization problem
hvdcl	24,842	158,426	98%	10%	5.41%	4.58%	power network problem
powersim	15,838	64,424	59%	53%	36.49%	49.09%	power network problem
TSOPF_RS_b39_c19	38,098	684,206	6%	0%	0.49%	0.49%	power network problem
nmos3	18,588	237,130	100%	17%	18.15%	19.91%	semiconductor device problem
matrix_9	103,430	1,205,518	100%	17%	10.99%	34.87%	semiconductor device problem
matrix-new_3	125,329	893,984	99%	28%	57.78%	62.60%	semiconductor device problem
igbt3	10,938	130,500	100%	17%	25.60%	3.74%	semiconductor device problem
2D_27628_bjtcai	27,628	206,670	100%	22%	43.23%	49.79%	semiconductor device problem
ohne2	181,343	6,869,939	100%	9%	23.86%	2.09%	semiconductor device problem
3D_51448_3D	51,448	537,038	99%	19%	34.34%	44.88%	semiconductor device problem
3D_28984_Tetra	28,984	285,092	99%	36%	47.07%	49.30%	semiconductor device problem
2D_54019_highK	54,019	486,129	100%	19%	31.67%	43.35%	semiconductor device problem
ibm_matrix_2	51,448	537,038	99%	19%	35.72%	44.57%	semiconductor device problem
wang3	26,064	177,168	100%	98%	84.12%	85.09%	semiconductor device problem
sme3Db	29,067	2,081,063	100%	44%	0.00%	0.00%	structural problem
sme3Da	12,504	874,887	100%	44%	0.00%	0.00%	structural problem
t2d_q4	9,801	87,025	100%	69%	74.97%	75.16%	structural problem sequence
venkat50	62,424	1,717,777	100%	6%	0.00%	0.00%	subsequent computational fluid dynamics problem
barrier2-10	115,625	2,158,759	100%	20%	27.31%	4.96%	subsequent semiconductor device problem
para-6	155,924	2,094,873	100%	37%	27.60%	4.02%	subsequent semiconductor device problem
barrier2-4	113,076	2,129,496	100%	19%	25.31%	5.07%	subsequent semiconductor device problem
para-9	155,924	2,094,873	100%	18%	27.71%	4.02%	subsequent semiconductor device problem
epb1	14,734	95,053	73%	0%	52.65%	60.66%	thermal problem
thermomech_dK	204,316	2,846,228	100%	67%	0.00%	0.00%	thermal problem
ted_A	10,605	424,587	57%	11%	0.00%	0.00%	thermal problem
FEM_3D_thermal1	17,880	430,740	100%	95%	0.00%	0.00%	thermal problem

## B Implementation of BiCOR

### B.1 User documentation

The implementation of the algorithm is based on the methods used in the Harwell Subroutine Library (HSL) [23]. Since the implementation is quite similar, the documentation is also quite similar to the documentation for the various subroutines in the HSL.

#### B.1.1 Argument lists and calling sequence

##### B.1.1.1 Initialization of the control parameters

The following subroutines have to be called before using the algorithm with BICORA(D). For single precision we have

```
CALL BICORI(ICNTL, CNTL, ISAVE, RSAVE)
```

and for double precision we have

```
CALL BICORID(ICNTL, CNTL, ISAVE, RSAVE)
```

where

- ICNTL is an INTEGER array of length 8 that does not have to be set by the user. On return, it contains the default values as described in section B.1.2.
- CNTL is a REAL (DOUBLE PRECISION in the D version) array of length 5 that does not have to be set by the user. On return, it contains the default values and described in section B.1.2.
- ISAVE is an INTEGER array of length 17 that must not be altered by the user.
- RSAVE is a REAL (DOUBLE PRECISION in the D version) array of length 9 that must not be altered by the user.

##### B.1.1.2 Solving $Ax=b$

Here we will actually solve  $Ax = b$ . For single precision we have

```
CALL BICORA(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL, INFO,  
+          ISAVE, RSAVE)
```

and for double precision we have

```
CALL BICORAD(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL, INFO,  
+          ISAVE, RSAVE)
```

where

IACT

is an INTEGER that indicate the action the user has to preform on every return of the BICORA/AD routines. Prior to the first call to BICORA/AD, IACT should be set to 0. Possible values are as follows:

- 1 An error occurred, and the user must terminate the computation. The reason for the error is in INFO(1). See section B.1.3 for more information.
- 1 If ICNTL(4)=0 (the default value), convergence has been achieved, and the user should terminate the computation. If ICNTL(4) is nonzero, the user may test for convergence. If convergence has not been achieved, BICORA/AD should be called again, without changes to its arguments.
- 2 The user must preform the matrix-vector product

$$y = Az$$

and recall BICORA/AD. The vectors  $y$  and  $z$  are in columns LOCY and LOCZ of array W respectively. The user should not change  $z$ .

- 3 The user must preform the preconditioning operation

$$y = Mz$$

where  $M$  is the preconditioner. The vectors  $y$  and  $z$  are in columns LOCY and LOCZ of array W respectively. The user should not change  $z$ . Preconditioning is only used when ICNTL(3) is nonzero.

- 4 The user must preform the transpose matrix-vector product

$$y = A^T z$$

and recall BICORA/AD. The vectors  $y$  and  $z$  are in columns LOCY and LOCZ of array W respectively. The user should not change  $z$ .

- 5 The user must preform the transpose preconditioning operation

$$y = M^T z$$

where  $M$  is the preconditioner. The vectors  $y$  and  $z$  are in columns LOCY and LOCZ of array W respectively. The user should not change  $z$ . Preconditioning is only used when ICNTL(3) is nonzero.

N

is an INTEGER variable that must be set by the user to the order of the matrix  $A$ . The variable must be preserved between calls to BICORA/AD. This argument is not altered by the routine.

W

is a REAL (DOUBLE PRECISION in the D version) two-dimensional array with dimensions (LWD, 8). Prior to the first call to BICORA/AD, the first column must hold the right hand side  $b$ . If ICNTL(5) is nonzero, the second column must contain the initial estimate of the solution  $x$ . On exit, the first column holds the residual  $r = b - Ax$  and the second column holds the estimate of the solution  $x$ . Other than the vector contained in the LOCYth column of W, W should remain unchanged between calls to BICORA/AD.

LWD	is an INTEGER variable that must be set by the user to the first dimension of $w$ . This argument is not altered by the routine. It should be greater than $N$ .
LOCY, LOCZ	are INTEGER variables that need not be set by the user. On return with $IAC T > 1$ , they indicate which columns of $w$ should be used to preform the operations as specified under $IAC T$ (see above). These arguments must not be altered by the user between calls to BICORA/AD.
RESID	is a REAL (DOUBLE PRECISION in the D version) variable that need not be set by the user. On return with $IAC T=1$ , it contains the 2-norm of the residual vector $\ b - A\hat{x}\ _2$ , where $\hat{x}$ is the current estimate of the solution.
ICNTL	is an INTEGER array of length 8 that has to be set by the user. The default values are set by a call to BICORA/AD as described in section B.1.1.1. Details of the control parameters are given in section B.1.2. This argument is not altered by the routine.
CNTL	is a REAL (DOUBLE PRECISION in the D version) array of length 5 that does not have to be set by the user. that has to be set by the user. The default values are set by a call to BICORA/AD as described in section B.1.1.1. Details of the control parameters are given in section B.1.2. This argument is not altered by the routine.
INFO	is an INTEGER array of length 4 that need not be set by the user. It is used to store information about the subroutine. On return of BICORA/AD, $INFO(1)$ tells if the subroutine was successful (value 0) or if an error occurred (non-zero values). More information about this is in section B.1.3. $INFO(2)$ holds the amount of iterations preformed by the subroutine. $INFO(3)$ and $INFO(4)$ are unused.
ISAVE	is an INTEGER array of length 17 that must not be altered by the user.
RSAVE	is a REAL (DOUBLE PRECISION in the D version) array of length 9 that must not be altered by the user.

### B.1.2 Control parameters

ICNTL and CNTL contain the control parameters of BICORA/AD. ICNTL controls the actions BICORA/AD takes, and CNTL controls the tolerances used by BICORA/AD. The default values are set by BICORI/ID.

ICNTL(1)	is the stream number for error messages and has default value 6. Printing of error messages is suppressed if $ICNTL(1) \leq 0$ .
ICNTL(2)	is the stream number for warning messages and has default value 6. Printing of warning messages is suppressed if $ICNTL(1) \leq 0$ .
ICNTL(3)	controls whether the user wishes to use preconditioning. It has default value 0 and in this case no preconditioning is used. If $ICNTL(3)$ is non-zero, the user will be expected to preform preconditioning when $IAC T = 3$ .
ICNTL(4)	controls whether the convergence test offered by BICORA/AD is used. It has default value 0 and in this case the computed solution $\hat{x}$ is accepted if the 2-norm of the residual ( $\ b - A\hat{x}\ _2$ ) is less or equal to $\max(CNTL(1) * (CNTL(2) + \ x\ _2 * CNTL(3)), CNTL(4))$ . If the user does not want to use this test, $ICNTL(4)$ should be non-zero. In this case, the user will be expected to test for convergence when $IAC T = 1$ .

ICNTL(5)	controls whether the user wishes to supply an initial estimate of the solution $x$ . It has default value 0 and in this case the initial estimate is set to the zero vector. If the user wishes to supply an initial estimate, ICNTL(5) should be non-zero. In this case, the initial estimate should be put in the second column of $W$ prior to the first call to BICORA/AD.
ICNTL(6)	determines the maximum number of iterations allowed. It has default value -1, and in this case the maximum number of iterations is equal to the order of the matrix $A$ ( $N$ ). If the user wishes to use a different maximum number of iterations, ICNTL(6) should be set to this number. In case of a negative number, the default will be used.
ICNTL(7), ICNTL(8)	have default value 0 and are unused by BICORA/AD
CNTL(1)	is one of the two convergence tolerances, as described under ICNTL(4). CNTL(1) has default value $\sqrt{u}$ , where $u$ is the relative machine precision. If ICNTL(4) is non-zero, this will not be used. See section B.2 for more information.
CNTL(2)	is the first variable used in the normwise backward error, and has a default value $\ b\ _2$ . The default value is set in BICORA/AD, not in BICORI/ID. See section B.2 for more information.
CNTL(3)	is the second variable used in the normwise backward error, and has a default value of zero. If this is left zero, the norm of $x$ will also not be calculated. See section B.2 for more information.
CNTL(4)	is one of the two convergence tolerances, as described under ICNTL(4). CNTL(2) has default value 0. If ICNTL(4) is non-zero, this will not be used. See section B.2 for more information.
CNTL(5)	is the breakdown tolerance. It has default value $u$ , where $u$ is the relative machine precision. If $\rho$ is close enough to zero according to this tolerance, the method has broken down. See section B.2 for more information.
CNTL(4), CNTL(5)	have default value 0 and are unused by BICORA/AD

### B.1.3 Error values

Upon the return of BICORA/AD, negative values for INFO(1) indicate an error and positive values indicate a warning. If everything went well, the value should be zero. Error messages are written to ICNTL(1) and warnings to ICNTL(2). Possible non-zero values for INFO(1) are:

- 1 The value of  $N$  is out of range ( $< 1$ ). There is an immediate return without any input parameters changed.
- 2 The value of  $LWD$  is out of range ( $< N$ ). There is an immediate return without any input parameters changed.
- 3 The algorithm has broken down.
- 4 The maximum amount of iterations determined by ICNTL(6) if it is not the default or  $N$  if ICNTL(6) is the default has been exceeded.
- 1 The convergence tolerance specified by the user in CNTL(1) lies outside the interval  $(u, 1.0)$  where  $u$  is the machine precision. CNTL(1) is reset to the default value  $\sqrt{u}$ .

#### **B.1.4 General information**

**Files needed to run the algorithm:**

`bicor.f`, `ddeps.f`

**Routines called:**

**BLAS** `SNRM2/DNRM2`, `SCOPY/DCOPY`, `SAXPY/DAXPY`, `SSCAL/DSCAL`, `SDOT/DDOT`

**HSL** `FD15A/AD`

**Restriction:**

$LWD \geq N \geq 1$

#### **B.2 Implementation**

```

1      SUBROUTINE BICORID(ICNTL,CNTL,ISAVE,RSAVE)
2 C Variables passed to the subroutine
3      IMPLICIT NONE
4      DOUBLE PRECISION CNTL(5)
5      INTEGER ICNTL(8)
6      INTEGER ISAVE(17)
7      DOUBLE PRECISION RSAVE(9)
8 C Local variables
9      INTEGER I
10     DOUBLE PRECISION ZERO
11     PARAMETER (ZERO=0.0D+0)
12     DOUBLE PRECISION FD15AD
13     EXTERNAL FD15AD
14     INTRINSIC SQRT
15     ICNTL(1) = 6
16     ICNTL(2) = 6
17     ICNTL(3) = 0
18     ICNTL(4) = 0
19     ICNTL(5) = 0
20     ICNTL(6) = -1
21     ICNTL(7) = 0
22     ICNTL(8) = 0
23     CNTL(1) = SQRT(FD15AD('E'))
24     CNTL(2) = ZERO
25     CNTL(3) = ZERO
26     CNTL(4) = ZERO
27     CNTL(5) = FD15AD('E')
28     DO 10 I = 1, 15
29       ISAVE(I) = 0
30 10 CONTINUE
31     DO 20 I = 1, 9
32       RSAVE(I) = 0.0
33 20 CONTINUE
34     RETURN
35     END
36     SUBROUTINE BICORAD(IACT,N,W,LDW,LOCY,LOCZ,RESID,ICNTL,CNTL,INFO,
37 + ISAVE,RSAVE)
38 C Variables passed to the subroutine
39     IMPLICIT NONE
40     DOUBLE PRECISION RESID
41     INTEGER IACT,LDW,LOCY,LOCZ,N
42     DOUBLE PRECISION CNTL(5),W(LDW,9)
43     INTEGER ICNTL(8),INFO(4)
44     INTEGER ISAVE(17)
45     DOUBLE PRECISION RSAVE(9)
46 C Local variables
47     DOUBLE PRECISION ONE, ZERO
48     PARAMETER (ONE=1.0D+0,ZERO=0.0D+0)
49     DOUBLE PRECISION BNRM2, RNRM2, RTNRM2, ALPHA, BETA, RHO, RHO1,
50 + XNRM2
51     INTEGER B, R, X, RPRM, ZR, ZRPRM, ZRHAT, P, PPRM, Q, QPRM,
52 + QPRMHAT, ZQPRMHAT, ITMAX, IPOS, I
53     DOUBLE PRECISION DDOT, DNRM2, FD15AD
54     EXTERNAL DDOT, DNRM2, FD15AD
55     INTRINSIC ABS,MAX,SQRT
56     EXTERNAL DAXPY,DCOPY,DSCAL
57 C Code
58 C Load all the local variables as they were on the last run
59     IPOS = ISAVE(1)
60     ITMAX = ISAVE(2)

```

```

61     B = ISAVE(3)
62     X = ISAVE(4)
63     R = ISAVE(5)
64     RPRM = ISAVE(6)
65     ZR = ISAVE(7)
66     ZRPRM = ISAVE(8)
67     ZRHAT = ISAVE(9)
68     P = ISAVE(10)
69     PPRM = ISAVE(11)
70     Q = ISAVE(12)
71     QPRM = ISAVE(13)
72     QPRMHAT = ISAVE(14)
73     ZQPRMHAT = ISAVE(15)
74     BNRM2 = RSAVE(1)
75     ALPHA = RSAVE(2)
76     BETA = RSAVE(3)
77     RHO = RSAVE(4)
78     RHO1 = RSAVE(5)
79     XNRM2 = RSAVE(6)
80     IF (IACT.EQ.0) GO TO 10
81     IF (IACT.LT.0) GO TO 1000
82     IF (IACT.EQ.1 .AND. ICNTL(4).EQ.0) GO TO 1000
83     IF (IACT.EQ.1 .AND. BNRM2.EQ.ZERO) GO TO 1000
84     IF (IPOS.EQ.1) GO TO 40
85     IF (IPOS.EQ.2) GO TO 60
86     IF (IPOS.EQ.3) GO TO 70
87     IF (IPOS.EQ.4) GO TO 80
88     IF (IPOS.EQ.5) GO TO 90
89     IF (IPOS.EQ.6) GO TO 100
90     IF (IPOS.EQ.7) GO TO 110
91     IF (IPOS.EQ.8) GO TO 120
92 10 CONTINUE
93     INFO(1) = 0
94 C No negative order possible
95     IF (N.LE.0) THEN
96       INFO(1) = -1
97 C W can't be larger than the order
98     ELSE IF (LDW.LT.MAX(1,N)) THEN
99       INFO(1) = -2
100    END IF
101 C Something went wrong, return an error
102    IF (INFO(1).LT.0) THEN
103      IACT = -1
104      IF (ICNTL(1).GT.0) WRITE (ICNTL(1),FMT=9000) INFO(1)
105      GO TO 1000
106    END IF
107    B = 1
108    X = 2
109    R = 1
110    RPRM = 3
111    ZR = 4
112    ZRPRM = 5
113    ZRHAT = 6
114    P = 7
115    PPRM = 8
116    Q = 9
117    QPRM = 3
118    QPRMHAT = 6
119    ZQPRMHAT = 10
120    INFO(2) = 0

```



```

121 C Max amount of iterations is N
122     ITMAX = N
123 C or ICNTL(6) if specified
124     IF (ICNTL(6).GT.0) ITMAX = ICNTL(6)
125 C If the 2 norm of b is zero, that means that b is zero, so the solution
126 C is zero, the residual is zero, everything is zero
127     BNRM2 = DNRM2(N,W(1,B),1)
128     IF (BNRM2.EQ.ZERO) THEN
129         IACT = 1
130         DO 20 I = 1,N
131             W(I,X) = ZERO
132             W(I,B) = ZERO
133     20 CONTINUE
134     RESID = ZERO
135     GO TO 1000
136 END IF
137 C In this case, the user may test for convergence when IACT = 1 is
138 C returned.
139     IF (ICNTL(4).EQ.0) THEN
140         IF (CNTL(1).LT.FD15AD('E') .OR. CNTL(1).GT.ONE) THEN
141             INFO(1) = 1
142             IF (ICNTL(2).GT.0) THEN
143                 WRITE (ICNTL(2),FMT=9010) INFO(1)
144                 WRITE (ICNTL(2),FMT=9020)
145             END IF
146             CNTL(1) = SQRT(FD15AD('E'))
147         END IF
148         IF (CNTL(2).EQ.ZERO) THEN
149             CNTL(2) = BNRM2
150         END IF
151     END IF
152 C Initial estimate for x is the 0 vector
153     IF (ICNTL(5).EQ.0) THEN
154         DO 30 I = 1,N
155             W(I,X) = ZERO
156     30 CONTINUE
157     GO TO 50
158 ELSE
159 C or if ICNTL(5) is not 0, you need to have specified W(1,X)
160     IF (DNRM2(N,W(1,X),1).EQ.ZERO) GO TO 50
161     IPOS = 1
162     IACT = 2
163     LOCY = P
164     LOCZ = X
165     GO TO 1000
166 END IF
167 C We have x and b, so r = -p (is Ax) + r (is b), so b-Ax
168 C We don't need b any more
169     40 CALL DAXPY(N,-ONE,W(1,P),1,W(1,R),1)
170     50 CONTINUE
171 C Set r prime as Ar
172     IPOS = 2
173     IACT = 2
174     LOCY = RPRM
175     LOCZ = R
176     GO TO 1000
177     60 CONTINUE
178 C Calculate zr prime on the first run
179     IF (ICNTL(3).NE.0) THEN
180         IPOS = 3

```

```

181     IACT = 5
182     LOCY = ZRPRM
183     LOCZ = RPRM
184     GO TO 1000
185 ELSE
186     CALL DCOPY(N,W(1,RPRM),1,W(1,ZRPRM),1)
187 END IF
188 70 CONTINUE
189     INFO(2) = INFO(2) + 1
190 C Check maximum number of iterations has not been exceeded.
191     IF (INFO(2).GT.ITMAX) THEN
192         INFO(1) = -4
193         IACT = -1
194         IF (ICNTL(1).GT.0) THEN
195             WRITE (ICNTL(1),FMT=9000) INFO(1)
196             WRITE (ICNTL(1),FMT=9030) ITMAX
197         END IF
198         GO TO 1000
199     END IF
200 C Perform the preconditioning operation
201     IF (ICNTL(3).NE.0) THEN
202         IPOS = 4
203         IACT = 3
204         LOCY = ZR
205         LOCZ = R
206         GO TO 1000
207     ELSE
208         CALL DCOPY(N,W(1,R),1,W(1,ZR),1)
209     END IF
210     80 CONTINUE
211 C Calculate zr hat
212     IPOS = 5
213     IACT = 2
214     LOCY = ZRHAT
215     LOCZ = ZR
216     GO TO 1000
217     90 CONTINUE
218 C See if the algorithm broke down. Otherwise, we can use rho in the
219 C remaining part of the algorithm
220     RHO = DDOT(N, W(1,ZRPRM), 1, W(1,ZRHAT), 1)
221     IF (ABS(RHO).LT.CNTL(5)*N) THEN
222         RNRM2 = DNRM2(N,W(1,R),1)
223         RTNRM2 = DNRM2(N,W(1,RPRM),1)
224         IF (ABS(RHO).LT.CNTL(5)*RNRM2*RTNRM2) THEN
225             INFO(1) = -3
226             IACT = -1
227             IF (ICNTL(1).GT.0) WRITE (ICNTL(1),FMT=9000) INFO(1)
228             GO TO 1000
229         END IF
230     END IF
231     IF (INFO(2).GT.1) THEN
232         BETA = RHO/RHO1
233         CALL DSCAL(N,BETA,W(1,P),1)
234         CALL DAXPY(N,ONE,W(1,ZR),1,W(1,P),1)
235         CALL DSCAL(N,BETA,W(1,PPRM),1)
236         CALL DAXPY(N,ONE,W(1,ZRPRM),1,W(1,PPRM),1)
237         CALL DSCAL(N,BETA,W(1,Q),1)
238         CALL DAXPY(N,ONE,W(1,ZRHAT),1,W(1,Q),1)
239     ELSE
240         CALL DCOPY(N,W(1,ZR),1,W(1,P),1)

```

```

241      CALL DCCOPY(N,W(1,ZRPRM),1,W(1,PPRM),1)
242      CALL DCCOPY(N,W(1,ZRHAT),1,W(1,Q),1)
243      END IF
244      IPOS = 6
245      IACT = 4
246      LOCY = QPRMHAT
247      LOCZ = PPRM
248      GO TO 1000
249 100 CONTINUE
250 C Perform preconditioning
251      IF (ICNTL(3).NE.0) THEN
252          IPOS = 7
253          IACT = 5
254          LOCY = ZQPRMHAT
255          LOCZ = QPRMHAT
256          GO TO 1000
257      ELSE
258          CALL DCCOPY(N,W(1,QPRMHAT),1,W(1,ZQPRMHAT),1)
259      END IF
260 110 CONTINUE
261      ALPHA = RHO/DDOT(N, W(1, ZQPRMHAT), 1, W(1,Q), 1)
262      CALL DAXPY(N,ALPHA,W(1,P),1,W(1,X),1)
263      CALL DAXPY(N,-ALPHA,W(1,Q),1,W(1,R),1)
264      CALL DAXPY(N,-ALPHA,W(1,ZQPRMHAT),1,W(1,ZRPRM),1)
265      RESID = DNRM2(N,W(1,R),1)
266      IPOS = 8
267 C The user can check the error if ICNTL(4) is non-zero at IACT.EQ.1
268      IF (ICNTL(4).NE.0) THEN
269          IACT = 1
270          GO TO 1000
271      ELSE
272          IF (CNTL(3).NE.ZERO) THEN
273              XNRM2 = DNRM2(N,W(1,X),1)
274          END IF
275          IF (RESID.LE.MAX(CNTL(1)*(CNTL(2)+XNRM2*CNTL(3)),CNTL(4))) THEN
276              IACT = 1
277              GO TO 1000
278          END IF
279      END IF
280 120 CONTINUE
281      RHO1 = RHO
282      GO TO 70
283 1000 CONTINUE
284 C Save all the local variables to use on the next run
285      ISAVE(1) = IPOS
286      ISAVE(2) = ITMAX
287      ISAVE(3) = B
288      ISAVE(4) = X
289      ISAVE(5) = R
290      ISAVE(6) = RPRM
291      ISAVE(7) = ZR
292      ISAVE(8) = ZRPRM
293      ISAVE(9) = ZRHAT
294      ISAVE(10) = P
295      ISAVE(11) = PPRM
296      ISAVE(12) = Q
297      ISAVE(13) = QPRM
298      ISAVE(14) = QPRMHAT
299      ISAVE(15) = ZQPRMHAT
300      RSAVE(1) = BNRM2

```

```

301      RSAVE(2) = ALPHA
302      RSAVE(3) = BETA
303      RSAVE(4) = RHO
304      RSAVE(5) = RHO1
305      RSAVE(6) = XNRM2
306      RETURN
307 9000 FORMAT (' Error message from BICOR. INFO(1) = ',I4)
308 9010 FORMAT (' Warning message from BICOR. INFO(1) = ',I4)
309 9020 FORMAT (' Convergence tolerance out of range.')
310 9030 FORMAT (' Number of iterations required exceeds the maximum of ',
311             + I8,' allowed by ICNTL(6)')
312      END

```

## C Implementation of CORS

### C.1 User documentation

The implementation of the algorithm is based on the methods used in the Harwell Subroutine Library (HSL) [23]. Since the implementation is quite similar, the documentation is also quite similar to the documentation for the various subroutines in the HSL.

#### C.1.1 Argument lists and calling sequence

##### C.1.1.1 Initialization of the control parameters

The following subroutines have to be called before using the algorithm with `CORSA(D)`. For single precision we have

```
CALL CORSI(ICNTL, CNTL, ISAVE, RSAVE)
```

and for double precision we have

```
CALL CORSID(ICNTL, CNTL, ISAVE, RSAVE)
```

where

- ICNTL is an INTEGER array of length 8 that does not have to be set by the user. On return, it contains the default values as described in section C.1.2.
- CNTL is a REAL (DOUBLE PRECISION in the D version) array of length 5 that does not have to be set by the user. On return, it contains the default values and described in section C.1.2.
- ISAVE is an INTEGER array of length 19 that must not be altered by the user.
- RSAVE is a REAL (DOUBLE PRECISION in the D version) array of length 9 that must not be altered by the user.

##### C.1.1.2 Solving $Ax=b$

Here we will actually solve  $Ax = b$ . For single precision we have

```
CALL CORSA(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL, INFO,  
+          ISAVE, RSAVE)
```

and for double precision we have

```
CALL CORSAD(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL, INFO,  
+          ISAVE, RSAVE)
```

where

IACT is an INTEGER that indicate the action the user has to preform on every return of the CORSA/AD routines. Prior to the first call to CORSA/AD, IACT should be set to 0. Possible values are as follows:

- 1 An error occurred, and the user must terminate the computation. The reason for the error is in INFO(1). See section C.1.3 for more information.
- 1 If ICNTL(4)=0 (the default value), convergence has been achieved, and the user should terminate the computation. If ICNTL(4) is nonzero, the user may test for convergence. If convergence has not been achieved, CORSA/AD should be called again, without changes to its arguments.
- 2 The user must preform the matrix-vector product

$$y = Az$$

and recall CORSA/AD. The vectors  $y$  and  $z$  are in columns LOCY and LOCZ of array W respectively. The user should not change  $z$ .

- 3 The user must preform the preconditioning operation

$$y = Mz$$

where  $M$  is the preconditioner. The vectors  $y$  and  $z$  are in columns LOCY and LOCZ of array W respectively. The user should not change  $z$ . Preconditioning is only used when ICNTL(3) is nonzero.

N is an INTEGER variable that must be set by the user to the order of the matrix  $A$ . The variable must be preserved between calls to CORSA/AD. This argument is not altered by the routine.

W is a REAL (DOUBLE PRECISION in the D version) two-dimensional array with dimensions (LWD, 13). Prior to the first call to CORSA/AD, the first column must hold the right hand side  $b$ . If ICNTL(5) is nonzero, the second column must contain the initial estimate of the solution  $x$ . On exit, the first column holds the residual  $r = b - A\hat{x}$  and the second column holds the estimate of the solution  $x$ . Other than the vector contained in the LOCYth column of W, W should remain unchanged between calls to CORSA/AD.

LWD is an INTEGER variable that must be set by the user to the first dimension of W. This argument is not altered by the routine. It should be greater than N.

LOCY, LOCZ are INTEGER variables that need not be set by the user. On return with IACT > 1, they indicate which columns of W should be used to preform the operations as specified under IACT (see above). These arguments must not be altered by the user between calls to CORSA/AD.

RESID is a REAL (DOUBLE PRECISION in the D version) variable that need not be set by the user. On return with IACT=1, it contains the 2-norm of the residual vector  $\|b - A\hat{x}\|_2$ , where  $\hat{x}$  is the current estimate of the solution.

ICNTL is an INTEGER array of length 8 that has to be set by the user. The default values are set by a call to CORSA/AD as described in section C.1.1.1. Details of the control parameters are given in section C.1.2. This argument is not altered by the routine.

CNTL	is a REAL (DOUBLE PRECISION in the D version) array of length 5 that does not have to be set by the user. that has to be set by the user. The default values are set by a call to CORSA/AD as described in section C.1.1.1. Details of the control parameters are given in section C.1.2. This argument is not altered by the routine.
INFO	is an INTEGER array of length 4 that need not be set by the user. It is used to store information about the subroutine. On return of CORSA/AD, INFO(1) tells if the subroutine was successful (value 0) or if an error occurred (non-zero values). More information about this is in section C.1.3. INFO(2) holds the amount of iterations preformed by the subroutine. INFO(3) and INFO(4) are unused.
ISAVE	is an INTEGER array of length 17 that must not be altered by the user.
RSAVE	is a REAL (DOUBLE PRECISION in the D version) array of length 9 that must not be altered by the user.

### C.1.2 Control parameters

ICNTL and CNTL contain the control parameters of CORSA/AD. ICNTL controls the actions CORSA/AD takes, and CNTL controls the tolerances used by CORSA/AD. The default values are set by CORSI/ID.

ICNTL(1)	is the stream number for error messages and has default value 6. Printing of error messages is suppressed if $\text{ICNTL}(1) \leq 0$ .
ICNTL(2)	is the stream number for warning messages and has default value 6. Printing of warning messages is suppressed if $\text{ICNTL}(1) \leq 0$ .
ICNTL(3)	controls whether the user wishes to use preconditioning. It has default value 0 and in this case no preconditioning is used. If ICNTL(3) is non-zero, the user will be expected to preform preconditioning when $\text{IACT} = 3$ .
ICNTL(4)	controls whether the convergence test offered by CORSA/AD is used. It has default value 0 and in this case the computed solution $\hat{x}$ is accepted if the 2-norm of the residual ( $\ b - A\hat{x}\ _2$ ) is less or equal to $\max(\text{CNTL}(1) * (\text{CNTL}(2) + \ x\ _2 * \text{CNTL}(3)), \text{CNTL}(4))$ . If the user does not want to use this test, ICNTL(4) should be non-zero. In this case, the user will be expected to test for convergence when $\text{IACT} = 1$ .
ICNTL(5)	controls whether the user wishes to supply an initial estimate of the solution $x$ . It has default value 0 and in this case the initial estimate is set to the zero vector. If the user wishes to supply an initial estimate, ICNTL(5) should be non-zero. In this case, the initial estimate should be put in the second column of $w$ prior to the first call to CORSA/AD.
ICNTL(6)	determines the maximum number of iterations allowed. It has default value -1, and in this case the maximum number of iterations is equal to the order of the matrix $A$ ( $N$ ). If the user wishes to use a different maximum number of iterations, ICNTL(6) should be set to this number. In case of a negative number, the default will be used.
ICNTL(7), ICNTL(8)	have default value 0 and are unused by CORSA/AD
CNTL(1)	is one of the two convergence tolerances, as described under ICNTL(4). CNTL(1) has default value $\sqrt{u}$ , where $u$ is the relative machine precision. If ICNTL(4) is non-zero, this will not be used. See section C.2 for more information.

- CNTL(2) is the first variable used in the normwise backward error, and has a default value  $\|b\|_2$ . The default value is set in CORSA/AD, not in CORSI/ID. See section C.2 for more information.
- CNTL(3) is the second variable used in the normwise backward error, and has a default value of zero. If this is left zero, the norm of  $x$  will also not be calculated. See section C.2 for more information.
- CNTL(4) is one of the two convergence tolerances, as described under ICNTL(4). CNTL(2) has default value 0. If ICNTL(4) is non-zero, this will not be used. See section C.2 for more information.
- CNTL(5) is the breakdown tolerance. It has default value  $u$ , where  $u$  is the relative machine precision. If  $\rho$  is close enough to zero according to this tolerance, the method has broken down. See section C.2 for more information.

### C.1.3 Error values

Upon the return of CORSA/AD, negative values for INFO(1) indicate an error and positive values indicate a warning. If everything went well, the value should be zero. Error messages are written to ICNTL(1) and warnings to ICNTL(2). Possible non-zero values for INFO(1) are:

- 1 The value of N is out of range ( $< 1$ ). There is an immediate return without any input parameters changed.
- 2 The value of LWD is out of range ( $< N$ ). There is an immediate return without any input parameters changed.
- 3 The algorithm has broken down.
- 4 The maximum amount of iterations determined by ICNTL(6) if it is not the default or N if ICNTL(6) is the default has been exceeded.
- 1 The convergence tolerance specified by the user in CNTL(1) lies outside the interval  $(u, 1.0)$  where  $u$  is the machine precision. CNTL(1) is reset to the default value  $\sqrt{u}$ .

### C.1.4 General information

**Files needed to run the algorithm:**

`cors.f`, `ddeps.f`

**Routines called:**

**BLAS** SNRM2/DNRM2, SCOPY/DCOPY, SAXPY/DAXPY, SSCAL/DSCAL, SDOT/DDOT

**HSL** FD15A/AD

**Restriction:**

$$\text{LWD} \geq \text{N} \geq 1$$

## C.2 Implementation

```

1      SUBROUTINE CORSID(ICNTL,CNTL,ISAVE,RSAVE)
2 C Variables passed to the subroutine
3      IMPLICIT NONE
4      DOUBLE PRECISION CNTL(5)
5      INTEGER ICNTL(8)
6      INTEGER ISAVE(19)
7      DOUBLE PRECISION RSAVE(9)
8 C Local variables
9      INTEGER I
10     DOUBLE PRECISION ZERO
11     PARAMETER (ZERO=0.0D+0)
12     DOUBLE PRECISION FD15AD
13     EXTERNAL FD15AD
14     INTRINSIC SQRT
15     ICNTL(1) = 6
16     ICNTL(2) = 6
17     ICNTL(3) = 0
18     ICNTL(4) = 0
19     ICNTL(5) = 0
20     ICNTL(6) = -1
21     ICNTL(7) = 0
22     ICNTL(8) = 0
23     CNTL(1) = SQRT(FD15AD('E'))
24     CNTL(2) = ZERO
25     CNTL(3) = ZERO
26     CNTL(4) = ZERO
27     CNTL(5) = FD15AD('E')
28     DO 10 I = 1, 19
29       ISAVE(I) = 0
30 10 CONTINUE
31     DO 20 I = 1, 9
32       RSAVE(I) = 0.0
33 20 CONTINUE
34     RETURN
35     END
36     SUBROUTINE CORSAD(IACT,N,W,LDW,LOCY,LOCZ,RESID,ICNTL,CNTL,INFO,
37 + ISAVE,RSAVE)
38 C Variables passed to the subroutine
39     IMPLICIT NONE
40     DOUBLE PRECISION RESID
41     INTEGER IACT,LDW,LOCY,LOCZ,N
42     DOUBLE PRECISION CNTL(5),W(LDW,13)
43     INTEGER ICNTL(8),INFO(4)
44     INTEGER ISAVE(19)
45     DOUBLE PRECISION RSAVE(9)
46 C Local variables
47     DOUBLE PRECISION TWO, ONE, ZERO
48     PARAMETER (TWO=2.0D+0,ONE=1.0D+0,ZERO=0.0D+0)
49     DOUBLE PRECISION BNRM2, RNRM2, RTNRM2, ALPHA, BETA, RHO, RHO1,
50 + XNRM2
51     INTEGER B, R, X, RPRM, ZR, ZRHAT, E, ZE, D, Q, ZQ, ZQHAT, H, F, G,
52 + ITMAX, IPOS, I
53     DOUBLE PRECISION DDOT, DNRM2, FD15AD
54     EXTERNAL DDOT, DNRM2, FD15AD
55     INTRINSIC ABS,MAX,SQRT
56     EXTERNAL DAXPY,DCOPY,DSCAL
57 C Code
58 C Load all the local variables as they were on the last run
59     IPOS = ISAVE(1)
60     ITMAX = ISAVE(2)
61     B = ISAVE(3)
62     X = ISAVE(4)
63     R = ISAVE(5)
64     RPRM = ISAVE(6)
65     ZR = ISAVE(7)
66     ZRHAT = ISAVE(8)
67     E = ISAVE(9)
68     ZE = ISAVE(10)
69     D = ISAVE(11)
70     Q = ISAVE(12)
71     ZQ = ISAVE(13)
72     ZQHAT = ISAVE(14)
73     H = ISAVE(15)
74     F = ISAVE(16)
75     G = ISAVE(17)
76     BNRM2 = RSAVE(1)
77     ALPHA = RSAVE(2)
78     BETA = RSAVE(3)
79     RHO = RSAVE(4)
80     RHO1 = RSAVE(5)
81     XNRM2 = RSAVE(6)
82     IF (IACT.EQ.0) GO TO 10
83     IF (IACT.LT.0) GO TO 1000
84     IF (IACT.EQ.1 .AND. ICNTL(4).EQ.0) GO TO 1000
85     IF (IACT.EQ.1 .AND. BNRM2.EQ.ZERO) GO TO 1000
86     IF (IPOS.EQ.1) GO TO 40
87     IF (IPOS.EQ.2) GO TO 60
88     IF (IPOS.EQ.3) GO TO 70
89     IF (IPOS.EQ.4) GO TO 80
90     IF (IPOS.EQ.5) GO TO 90
91     IF (IPOS.EQ.6) GO TO 100
92     IF (IPOS.EQ.7) GO TO 110
93     IF (IPOS.EQ.8) GO TO 120
94 10 CONTINUE
95     INFO(1) = 0
96 C No negative order possible
97     IF (N.LE.0) THEN
98       INFO(1) = -1
99 C W can't be larger than the order
100    ELSE IF (LDW.LT.MAX(1,N)) THEN
101      INFO(1) = -2
102    END IF
103 C Something went wrong, return an error
104    IF (INFO(1).LT.0) THEN
105      IACT = -1
106      IF (ICNTL(1).GT.0) WRITE (ICNTL(1),FMT=9000) INFO(1)
107      GO TO 1000
108    END IF
109    B = 1
110    X = 2
111    R = 1
112    RPRM = 3
113    ZR = 4
114    ZRHAT = 5
115    E = 6
116    ZE = 7
117    D = 8
118    Q = 9
119    ZQ = 5
120    ZQHAT = 10

```

```

121      H = 11
122      F = 12
123      G = 13
124      INFO(2) = 0
125 C Max amount of iterations is N
126      ITMAX = N
127 C or ICNTL(6) if specified
128      IF (ICNTL(6).GT.0) ITMAX = ICNTL(6)
129 C If the 2 norm of b is zero, that means that b is zero, so the solution
130 C is zero, the residual is zero, everything is zero
131      BNRM2 = DNRM2(N,W(1,B),1)
132      IF (BNRM2.EQ.ZERO) THEN
133          IACT = 1
134          DO 20 I = 1,N
135              W(I,X) = ZERO
136              W(I,B) = ZERO
137      20    CONTINUE
138      RESID = ZERO
139      GO TO 1000
140      END IF
141 C In this case, the user may test for convergence when IACT = 1 is
142 C returned.
143      IF (ICNTL(4).EQ.0) THEN
144          IF (CNTL(1).LT.FD15AD('E') .OR. CNTL(1).GT.ONE) THEN
145              INFO(1) = 1
146              IF (ICNTL(2).GT.0) THEN
147                  WRITE (ICNTL(2),FMT=9010) INFO(1)
148                  WRITE (ICNTL(2),FMT=9020)
149              END IF
150              CNTL(1) = SQRT(FD15AD('E'))
151          END IF
152          IF (CNTL(2).EQ.ZERO) THEN
153              CNTL(2) = BNRM2
154          END IF
155      END IF
156 C Initial estimate for x is the 0 vector
157      IF (ICNTL(5).EQ.0) THEN
158          DO 30 I = 1,N
159              W(I,X) = ZERO
160      30    CONTINUE
161      GO TO 50
162      ELSE
163 C or if ICNTL(5) is not 0, you need to have specified W(1,X)
164      IF (DNRM2(N,W(1,X),1).EQ.ZERO) GO TO 50
165      IPOS = 1
166      IACT = 2
167      LOCY = Q
168      LOCZ = X
169      GO TO 1000
170      END IF
171 C We have x and b, so r = -q (is Ax) + r (is b), so b-Ax
172 C We don't need b any more
173      40  CALL DAXPY(N,-ONE,W(1,Q),1,W(1,R),1)
174      50  CONTINUE
175 C Set r prime as Ar
176      IPOS = 2
177      IACT = 2
178      LOCY = RPRM
179      LOCZ = R
180      GO TO 1000

```

```

181      60  CONTINUE
182      INFO(2) = INFO(2) + 1
183 C Check maximum number of iterations has not been exceeded.
184      IF (INFO(2).GT.ITMAX) THEN
185          INFO(1) = -4
186          IACT = -1
187          IF (ICNTL(1).GT.0) THEN
188              WRITE (ICNTL(1),FMT=9000) INFO(1)
189              WRITE (ICNTL(1),FMT=9030) ITMAX
190          END IF
191          GO TO 1000
192      END IF
193 C Perform the preconditioning operation
194      IF (ICNTL(3).NE.0) THEN
195          IPOS = 3
196          IACT = 3
197          LOCY = ZR
198          LOCZ = R
199          GO TO 1000
200      ELSE
201          CALL DCOPY(N,W(1,R),1,W(1,ZR),1)
202      END IF
203      70  CONTINUE
204 C Calculate zr hat
205      IPOS = 4
206      IACT = 2
207      LOCY = ZRHAT
208      LOCZ = ZR
209      GO TO 1000
210      80  CONTINUE
211 C See if the algorithm broke down. Otherwise, we can use rho in the
212 C remaining part of the algorithm
213      RHO = DDOT(N, W(1,RPM), 1, W(1,ZRHAT), 1)
214      IF (ABS(RHO).LT.CNTL(5)*N) THEN
215          RNRM2 = DNRM2(N,W(1,R),1)
216          RTNRM2 = DNRM2(N,W(1,RPM),1)
217          IF (ABS(RHO).LT.CNTL(5)*RNRM2*RTNRM2) THEN
218              INFO(1) = -3
219              IACT = -1
220              IF (ICNTL(1).GT.0) WRITE (ICNTL(1),FMT=9000) INFO(1)
221              GO TO 1000
222          END IF
223      END IF
224      IF (INFO(2).GT.1) THEN
225          BETA = RHO/RHO1
226 C e = r + beta*h
227          CALL DCOPY(N,W(1,R),1,W(1,E),1)
228          CALL DAXPY(N,BETA,W(1,H),1,W(1,E),1)
229 C ze = zr + beta*f
230          CALL DCOPY(N,W(1,ZR),1,W(1,ZE),1)
231          CALL DAXPY(N,BETA,W(1,F),1,W(1,ZE),1)
232 C d = zrhat + beta*g
233          CALL DCOPY(N,W(1,ZRHAT),1,W(1,D),1)
234          CALL DAXPY(N,BETA,W(1,G),1,W(1,D),1)
235 C q = d+beta*(g+beta*q)
236          CALL DSCAL(N,BETA,W(1,Q),1)
237          CALL DAXPY(N,ONE,W(1,G),1,W(1,Q),1)
238          CALL DSCAL(N,BETA,W(1,Q),1)
239          CALL DAXPY(N,ONE,W(1,D),1,W(1,Q),1)
240      ELSE

```



```

241      CALL DCPY(N,W(1,R),1,W(1,E),1)
242      CALL DCPY(N,W(1,ZRHAT),1,W(1,D),1)
243      CALL DCPY(N,W(1,ZRHAT),1,W(1,Q),1)
244 C Calculate ze on the first run by a preconditioning operation
245      IF (ICNTL(3).NE.0) THEN
246          IPOS = 5
247          IACT = 3
248          LOCY = ZE
249          LOCZ = E
250          GO TO 1000
251      ELSE
252          CALL DCPY(N,W(1,E),1,W(1,ZE),1)
253      END IF
254      END IF
255      90 CONTINUE
256 C Perform preconditioning
257      IF (ICNTL(3).NE.0) THEN
258          IPOS = 6
259          IACT = 3
260          LOCY = ZQ
261          LOCZ = Q
262          GO TO 1000
263      ELSE
264          CALL DCPY(N,W(1,Q),1,W(1,ZQ),1)
265      END IF
266      100 CONTINUE
267          IPOS = 7
268          IACT = 2
269          LOCY = ZQHAT
270          LOCZ = ZQ
271          GO TO 1000
272      110 CONTINUE
273          ALPHA = RHO/DDOT(N, W(1, RPRM), 1, W(1,ZQHAT), 1)
274 C h=e-alpha*q
275          CALL DCPY(N,W(1,E),1,W(1,H),1)
276          CALL DAXPY(N,-ALPHA,W(1,Q),1,W(1,E),1)
277 C f=ze-alpha*zq
278          CALL DSCAL(N,-ALPHA,W(1,ZQ),1)
279          CALL DCPY(N,W(1,ZQ),1,W(1,F),1)
280          CALL DAXPY(N,ONE,W(1,ZE),1,W(1,F),1)
281 C g=d-alpha*zqhat
282          CALL DSCAL(N,-ALPHA,W(1,ZQHAT),1)
283          CALL DCPY(N,W(1,ZQHAT),1,W(1,G),1)
284          CALL DAXPY(N,ONE,W(1,D),1,W(1,G),1)
285 C x=x+alpha*(2*ze-alpha*zq), -alpha*zq is already stored in zq
286          CALL DAXPY(N,TWO,W(1,ZE),1,W(1,ZQ),1)
287          CALL DAXPY(N,ALPHA,W(1,ZQ),1,W(1,X),1)
288 C r=r-alpha*(2*d-alpha*zqhat), -alpha*zqhat is already stored in zqhat
289          CALL DAXPY(N,TWO,W(1,D),1,W(1,ZQHAT),1)
290          CALL DAXPY(N,-ALPHA,W(1,ZQHAT),1,W(1,R),1)
291          RESID = DNRM2(N,W(1,R),1)
292          IPOS = 8
293 C The user can check the error if ICNTL(4) is non-zero at IACT.EQ.1
294          IF (ICNTL(4).NE.0) THEN
295              IACT = 1
296              GO TO 1000
297          ELSE
298              IF (CNTL(3).NE.ZERO) THEN
299                  XNRM2 = DNRM2(N,W(1,X),1)
300              END IF

```

```

301      IF (RESID.LE.MAX(CNTL(1)*(CNTL(2)+XNRM2*CNTL(3)),CNTL(4))) THEN
302          IACT = 1
303          GO TO 1000
304      END IF
305      END IF
306      120 CONTINUE
307          RHO1 = RHO
308          GO TO 60
309      1000 CONTINUE
310 C Save all the local variables to use on the next run
311          ISAVE(1) = IPOS
312          ISAVE(2) = ITMAX
313          ISAVE(3) = B
314          ISAVE(4) = X
315          ISAVE(5) = R
316          ISAVE(6) = RPRM
317          ISAVE(7) = ZR
318          ISAVE(8) = ZRHAT
319          ISAVE(9) = E
320          ISAVE(10) = ZE
321          ISAVE(11) = D
322          ISAVE(12) = Q
323          ISAVE(13) = ZQ
324          ISAVE(14) = ZQHAT
325          ISAVE(15) = H
326          ISAVE(16) = F
327          ISAVE(17) = G
328          RSAVE(1) = BNRM2
329          RSAVE(2) = ALPHA
330          RSAVE(3) = BETA
331          RSAVE(4) = RHO
332          RSAVE(5) = RHO1
333          RSAVE(6) = XNRM2
334      RETURN
335      9000 FORMAT (' Error message from CORS. INFO(1) = ',I4)
336      9010 FORMAT (' Warning message from CORS. INFO(1) = ',I4)
337      9020 FORMAT (' Convergence tolerance out of range.')
338      9030 FORMAT (' Number of iterations required exceeds the maximum of ',
339          +      I8,' allowed by ICNTL(6)')
340      END

```

## D Implementation of the testing application

```

1      PROGRAM TEST
2
3 C !!!!!!!!!!!!!!!!!!! WARNING !!!!!!!!!!!!!!!!!!!!!
4 C IF ILUPACK RETURNS AN ERROR, THERE WILL APPEAR A
5 C *** glibc detected *** double free or corruption
6 C ERROR. TO PREVENT THIS, RUN THE PROGRAM WITH
7 C MALLOC_CHECK_=0
8
9
10 C Solve the linear system A x = b
11
12 C .. Parameters ..
13     IMPLICIT NONE
14     INTEGER N, LDW
15 C ..
16 C .. Local Scalars ..
17     DOUBLE PRECISION RESID
18     INTEGER I, IACT, ROWS, COLS, NNZ, RNNZ
19     INTEGER LOCY, LOCZ
20     CHARACTER REP*10
21     CHARACTER FIELD*7
22     CHARACTER SYMM*19
23     DOUBLE PRECISION BNRM2
24 C ..
25 C .. Local Arrays ..
26     INTEGER LOCY2(2), LOCZ2(2)
27     DOUBLE PRECISION CNTL(5)
28     DOUBLE PRECISION, DIMENSION(:, :), ALLOCATABLE :: W
29     INTEGER ICNTL(8), INFO(4)
30     INTEGER ISAVE(19)
31     DOUBLE PRECISION RSAVE(9)
32     COMPLEX, DIMENSION(:), ALLOCATABLE :: CVAL
33     DOUBLE PRECISION, DIMENSION(:), ALLOCATABLE :: RVAL
34     INTEGER, DIMENSION(:), ALLOCATABLE :: IVAL
35     INTEGER, DIMENSION(:), ALLOCATABLE :: INDX
36     INTEGER, DIMENSION(:), ALLOCATABLE :: JNDX
37     DOUBLE PRECISION, DIMENSION(:), ALLOCATABLE :: RVALO
38     INTEGER, DIMENSION(:), ALLOCATABLE :: INDXO
39     INTEGER, DIMENSION(:), ALLOCATABLE :: JNDXO
40     DOUBLE PRECISION, DIMENSION(:), ALLOCATABLE :: SOLUTION
41     DOUBLE PRECISION, DIMENSION(:), ALLOCATABLE :: RHS
42     DOUBLE PRECISION, DIMENSION(:), ALLOCATABLE :: TEMP
43 C ..
44 C .. Local Variables needed for ILUPACK ..
45     DOUBLE PRECISION, DIMENSION(:), ALLOCATABLE :: ILUA
46     INTEGER, DIMENSION(:), ALLOCATABLE :: ILUIA
47     INTEGER, DIMENSION(:), ALLOCATABLE :: ILUJA
48     INTEGER, DIMENSION(:), ALLOCATABLE :: ILUIND
49 C ..
50 C .. ILUPACK external parameters
51     INTEGER ILUMATCHING, ILUMAXIT, ILULFIL, ILULFILS, ILUNRESTART,
52 +         ILUIERR, ILUMIXEDPRECISION
53     CHARACTER ILUORDERING*20
54     DOUBLE PRECISION ILUDROPTOL, ILUDROPTOLS, ILUCONDEST, ILURESTOL,
55 +         ILUELBOU
56     INTEGER*8 ILUPARAM, ILUPREC
57     INTEGER DGNLAMGFACTOR, DGNLAMGNNZ
58     EXTERNAL DGNLAMGINIT, DGNLAMGSOL, DGNLAMGFACTOR, DGNLAMGTSOL,
59 +         DGNLAMGNNZ, DGNLAMGDELETE
60 C ..

```

```

61 C .. Local Variables needed for BICGSTAB1 ..
62     INTEGER L, LDRW, MXMV, LDWB
63     DOUBLE PRECISION TOL
64     DOUBLE PRECISION, DIMENSION(:, :), ALLOCATABLE :: WORK
65     DOUBLE PRECISION, DIMENSION(:, :), ALLOCATABLE :: RWORK
66     INTEGER, DIMENSION(:), ALLOCATABLE :: IWORK
67     INTEGER, DIMENSION(:, :), ALLOCATABLE :: IWORK2, IWORK3
68     EXTERNAL BISTBL, PRECSOLVE, MV
69 C ..
70 C .. Variables for GMRES
71     INTEGER M
72     LOGICAL LSAVE(4)
73     DOUBLE PRECISION, DIMENSION(:, :), ALLOCATABLE :: H
74     INTEGER LDH
75 C ..
76 C .. Variables for QMR
77     INTEGER MAXPQ, MAXVW, MVEC
78 C ..
79 C .. Local Variables needed for testing ..
80     REAL TARRAY(2)
81     CHARACTER MNAME*50, TOL_STRING*5
82     CHARACTER MFNAME*70
83     CHARACTER FNAME*70
84     INTEGER MVOPP, MAXMVP
85     INTEGER FSTAT
86     INTEGER NPRT, PREC, CALLS, RUNS
87     PARAMETER (NPRT=20)
88     DOUBLE PRECISION PREC_TOL, PREC_UPPER_TOL
89 C ..
90 C .. External Subroutines ..
91     EXTERNAL MMREAD, MMINFO
92     EXTERNAL CORSAD, CORSID
93     EXTERNAL BICORAD, BICORID
94     EXTERNAL MI26AD, MI26ID
95     EXTERNAL MI25AD, MI25ID
96     EXTERNAL MI24AD, MI24ID
97     EXTERNAL MI23AD, MI23ID
98     EXTERNAL AMUX, ATMUX
99     DOUBLE PRECISION FD15AD, DNRM2
100     EXTERNAL DCOPY, DNRM2, FD15AD
101     EXTERNAL CLEARALL
102     DOUBLE PRECISION ETA
103
104 C Some parameters
105     PARAMETER (ETA = 1.0D-10)
106     PARAMETER (MAXMVP = 20000)
107
108 C Open the list with matrices
109     OPEN(11, FILE='matrices.txt', STATUS='OLD')
110 5 CONTINUE
111 C ..
112 C .. Load our matrix and convert to the sparse matrix format ..
113     READ(11, *, END=2000) MNAME
114     WRITE(6, 4010) TRIM(MNAME)
115     MFNAME = '/home/sven/matrices/' // TRIM(MNAME) // '.mtx'
116     OPEN(1, FILE=MFNAME, STATUS='OLD')
117     WRITE(6, 4020)
118     CALL MMINFO(1, REP, FIELD, SYMM, ROWS, COLS, NNZ)
119     IF (ROWS.NE.COLS) THEN
120         WRITE(6, 9990)

```

```

121      GO TO 5
122      END IF
123
124      N = ROWS
125      LDW = N
126
127      ALLOCATE (TEMP(N))
128      ALLOCATE (RVAL(NNZ))
129      ALLOCATE (CVAL(NNZ))
130      ALLOCATE (IVAL(NNZ))
131      ALLOCATE (INDX(NNZ))
132      ALLOCATE (JNDX(NNZ))
133      ALLOCATE (RVALO(NNZ))
134      ALLOCATE (JNDXO(NNZ))
135      ALLOCATE (INDXO(N+1))
136
137      CALL MMREAD(1, REP, FIELD, SYMM, ROWS, COLS, NNZ, NNZ,
138 +             INDX, JNDX, IVAL, RVAL, CVAL)
139      CLOSE(1)
140 C .. column indices and row pointers
141      WRITE(6, 4030)
142      CALL COOCSR(ROWS, NNZ, RVAL, INDX, JNDX, RVALO, JNDXO, INDXO)
143
144      DEALLOCATE (RVAL)
145      DEALLOCATE (CVAL)
146      DEALLOCATE (IVAL)
147      DEALLOCATE (INDX)
148      DEALLOCATE (JNDX)
149
150 C ..
151 C Preconditioning is required
152      PREC = 0
153      ICN1(3) = PREC
154 C Set right hand side, b
155      WRITE(6, 4035)
156
157      ALLOCATE (SOLUTION(N))
158      ALLOCATE (RHS(N))
159
160      OPEN(13, FILE='/home/sven/matrices/' // TRIM(MNAME)
161 +      // '_rhs1.mtx', STATUS='OLD', IOSTAT=FSTAT)
162      IF (FSTAT.NE.0) THEN
163          OPEN(13, FILE='/home/sven/matrices/' // TRIM(MNAME)
164 +      // '_b.mtx', STATUS='OLD', IOSTAT=FSTAT)
165          IF (FSTAT.NE.0) THEN
166              WRITE(6, 4036)
167              SOLUTION = 1.0D+0
168              CALL AMUX(N, SOLUTION, RHS, RVALO, JNDXO, INDXO)
169          ELSE
170              CALL MMREAD(13, REP, FIELD, SYMM, ROWS, COLS, RNNZ, N,
171 +              SOLUTION, SOLUTION, SOLUTION, RHS, SOLUTION)
172          END IF
173      ELSE
174          CALL MMREAD(13, REP, FIELD, SYMM, ROWS, COLS, RNNZ, N,
175 +              SOLUTION, SOLUTION, SOLUTION, RHS, SOLUTION)
176      END IF
177      CLOSE(13)
178      BNRM2 = DNRM2(N, RHS, 1)
179 C Preconditioner switch
180      PREC_TOL = 0.01

```

```

181      PREC_UPPER_TOL = 5.0
182 7 CONTINUE
183      RUNS = 0
184      PREC_TOL = PREC_TOL * 10.0
185      WRITE(TOL_STRING, '(F5.2)') PREC_TOL
186      TOL_STRING = REPEAT('0', 5 - LEN_TRIM(ADJUSTL(TOL_STRING)))
187 +      // ADJUSTL(TOL_STRING)
188      FNAME = 'output/' // TRIM(MNAME) // '_' // TOL_STRING // '.txt'
189      OPEN(10, FILE=FNAME, STATUS='NEW', IOSTAT=FSTAT)
190      IF (FSTAT.NE.0) THEN
191          WRITE(6, FMT=4037)
192          RUNS = 100
193          GO TO 1010
194      END IF
195      WRITE(10, FMT=3010) TRIM(MNAME), N, NNZ
196
197      ALLOCATE (W(N, 13))
198
199 C Initialize data for the preconditioner
200 C
201      WRITE(6, FMT=4040)
202
203      ALLOCATE (ILUJA(NNZ))
204      ALLOCATE (ILUIA(N+1))
205      ALLOCATE (ILUA(NNZ))
206      ALLOCATE (ILUIND(N))
207
208      ILUJA = JNDXO
209      ILUIA = INDXO
210      CALL DCOPY(NNZ, RVALO, 1, ILUA, 1)
211
212      CALL DGNLAMGINIT(N, ILUIA, ILUJA, ILUA, ILUMATCHING,
213 +      ILUORDERING, ILUDROPTOL, ILUDROPTOLS,
214 +      ILUCONDEST, ILURESTOL, ILUMAXIT, ILUELBOV,
215 +      ILULFIL, ILULFILS, ILUNRESTART,
216 +      ILUMIXEDPRECISION, ILUIND)
217
218 c threshold for ILU, default: 1e-2
219      ILUDROPTOL=PREC_TOL
220 c
221 c threshold for the approximate Schur complements, default: 0.1*droptol
222      ILUDROPTOLS=0.1*ILUDROPTOL
223
224      ILUMIXEDPRECISION = 0
225      CALL CPU_TIME(TARRAY(1))
226
227      ILUIERR=DGNLAMGFACTOR(ILUPARAM, ILUPREC, N, ILUIA, ILUJA,
228 +      ILUA, ILUMATCHING, ILUORDERING, ILUDROPTOL, ILUDROPTOLS,
229 +      ILUCONDEST, ILURESTOL, ILUMAXIT, ILUELBOV, ILULFIL,
230 +      ILULFILS, ILUNRESTART, ILUMIXEDPRECISION, ILUIND)
231
232      CALL CPU_TIME(TARRAY(2))
233
234      IF (ILUIERR.EQ.-1) THEN
235          WRITE (6, '(A)') 'Error. input matrix may be wrong.'
236      ELSEIF (ILUIERR.EQ.-2) THEN
237          WRITE (6, '(A)')
238 +      'matrix L overflow, increase elbow and retry'
239      ELSEIF (ILUIERR.EQ.-3) THEN
240          WRITE (6, '(A)')

```

```

241 + 'matrix U overflow, increase elbow and retry'
242 ELSEIF (ILUIERR.EQ.-4) THEN
243   WRITE (6,'(A)') 'Illegal value for lfil'
244 ELSEIF (ILUIERR.EQ.-5) THEN
245   WRITE (6,'(A)') 'zero row encountered'
246 ELSEIF (ILUIERR.EQ.-6) THEN
247   WRITE (6,'(A)') 'zero column encountered'
248 ELSEIF (ILUIERR.EQ.-7) THEN
249   WRITE (6,'(A)') 'buffers are too small'
250 ELSEIF (ILUIERR.NE.0) THEN
251   WRITE (6,'(A,I3)')
252 + 'zero pivot encountered at step number', ILUIERR
253 ENDIF
254 IF (ILUIERR.NE.0) THEN
255   WRITE (10, FMT=3032) ILUIERR
256   GO TO 1000
257 END IF
258
259 WRITE (10, FMT=3020) TARRAY(2)-TARRAY(1), N,
260 + DGNLAMGNZ (ILUPARAM, ILUPREC)
261 CALL FLUSH(10)
262 WRITE (6, FMT=4050)
263 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
264 CCCCCCCCCCCCCCCCCCCCCCCCCC   CORS   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
265 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
266   WRITE (10, FMT=3030) 'CORS'
267
268 CALLS = 0
269 10 CONTINUE
270 CALLS = CALLS + 1
271 RUNS = RUNS + 1
272 C Clear everything
273
274 CALL CORSID(ICNTL, CNTL, ISAVE, RSAVE)
275 CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
276 + INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
277
278 C Perform an iteration of the CORS method
279
280 CALL CPU_TIME(TARRAY(1))
281 20 CONTINUE
282 CALL CORSAD(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
283 + INFO, ISAVE, RSAVE)
284
285 IF (MVOPP.GE.MAXMVP) THEN
286   CALL CPU_TIME(TARRAY(2))
287   WRITE (6, FMT=9020) -88
288   CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, -88)
289   GO TO 30
290 END IF
291
292 IF (IACT.LT.0) THEN
293   CALL CPU_TIME(TARRAY(2))
294   WRITE (6, FMT=9020) INFO(1)
295   CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
296   GO TO 30
297 END IF
298
299 IF (IACT.EQ.2) THEN
300 C Perform the matrix-vector product

```

```

301   CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,LOCZ), TEMP, N)
302   CALL AMUX(N, TEMP, W(1, LOCY), RVALO, JNDXO, INDXO)
303   MVOPP = MVOPP + 1
304   GO TO 20
305 END IF
306
307 IF (IACT.EQ.3) THEN
308 C Perform the preconditioning operation
309   CALL DCOPI(N, W(1,LOCZ), 1, W(1,LOCY), 1)
310   GO TO 20
311 END IF
312
313 CALL CPU_TIME(TARRAY(2))
314
315 CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
316 IF (ISNAN(TEMP(1))) THEN
317   INFO(1) = -99
318   WRITE (6, FMT=9020) INFO(1)
319   CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
320   GO TO 30
321 END IF
322
323 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
324
325 C Solution found
326   WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
327   IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
328
329 IF (CALLS.LE.5) THEN
330   GO TO 10
331 END IF
332
333 30 CONTINUE
334
335 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
336 CCCCCCCCCCCCCCCCCCCCCCCC   BICOR   CCCCCCCCCCCCCCCCCCCCCCCCCCCC
337 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
338   WRITE (10,3030) 'BiCOR'
339
340 CALLS = 0
341 110 CONTINUE
342 CALLS = CALLS + 1
343 RUNS = RUNS + 1
344 C Clear everything
345
346 CALL BICORID(ICNTL, CNTL, ISAVE, RSAVE)
347 CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
348 + INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
349
350 C Perform an iteration of the BICOR method
351
352 CALL CPU_TIME(TARRAY(1))
353 120 CONTINUE
354 CALL BICORAD(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
355 + INFO, ISAVE, RSAVE)
356
357 IF (MVOPP.GE.MAXMVP) THEN
358   CALL CPU_TIME(TARRAY(2))
359   WRITE (6, FMT=9020) -88
360   CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, -88)

```

```

361      GO TO 130
362      END IF
363
364      IF (IACT.LT.0) THEN
365          CALL CPU_TIME(TARRAY(2))
366          WRITE (6, FMT=9020) INFO(1)
367          CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
368          GO TO 130
369      END IF
370
371      IF (IACT.EQ.2) THEN
372 C Perform the matrix-vector product
373          CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,LOCZ), TEMP, N)
374          CALL AMUX(N, TEMP, W(1, LOCY), RVALO, JNDXO, INDXO)
375          MVOPP = MVOPP + 1
376          GO TO 120
377      END IF
378
379      IF (IACT.EQ.3) THEN
380 C Perform the preconditioning operation
381          CALL DCOPY(N, W(1,LOCZ), 1, W(1,LOCY), 1)
382          GO TO 120
383      END IF
384
385      IF (IACT.EQ.4) THEN
386 C Perform the matrix-vector product
387          CALL ATMUX(N, W(1, LOCZ), TEMP, RVALO, JNDXO, INDXO)
388          CALL DGNLAMGTSOL(ILUPARAM, ILUPREC, TEMP, W(1,LOCY), N)
389          MVOPP = MVOPP + 1
390          GO TO 120
391      END IF
392
393      IF (IACT.EQ.5) THEN
394 C Perform the preconditioning operation
395          CALL DCOPY(N, W(1,LOCZ), 1, W(1,LOCY), 1)
396          GO TO 120
397      END IF
398
399      CALL CPU_TIME(TARRAY(2))
400
401      CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
402      IF (ISNAN(TEMP(1))) THEN
403          INFO(1) = -99
404          WRITE (6, FMT=9020) INFO(1)
405          CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
406          GO TO 130
407      END IF
408
409      CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
410
411 C Solution found
412      WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
413      IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
414
415      IF (CALLS.LE.5) THEN
416          GO TO 110
417      END IF
418
419      130 CONTINUE
420

```

```

422 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
      BICG-STAB          CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
423 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
      WRITE(10,3030) 'BiCG-stab'
424
425
426     CALLS = 0
427   210 CONTINUE
428       CALLS = CALLS + 1
429       RUNS = RUNS + 1
430 C Clear everything
431
432     CALL MI26ID(ICNTL, CNTL, ISAVE, RSAVE)
433     CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
434 +              INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
435
436 C Perform an iteration of the BICG-STAB method
437
438     CALL CPU_TIME(TARRAY(1))
439   220 CONTINUE
440     CALL MI26AD(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
441 +              INFO, ISAVE, RSAVE)
442
443     IF (MVOPP.GE.MAXMVP) THEN
444         CALL CPU_TIME(TARRAY(2))
445         WRITE (6, FMT=9020) -88
446         CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, -88)
447         GO TO 230
448     END IF
449
450     IF (IACT.LT.0) THEN
451         CALL CPU_TIME(TARRAY(2))
452         WRITE (6, FMT=9020) INFO(1)
453         CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
454         GO TO 230
455     END IF
456
457     IF (IACT.EQ.2) THEN
458 C Perform the matrix-vector product
459         CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,LOCZ), TEMP, N)
460         CALL AMUX(N, TEMP, W(1, LOCY), RVALO, JNDXO, IND XO)
461         MVOPP = MVOPP + 1
462         GO TO 220
463     END IF
464
465     IF (IACT.EQ.3) THEN
466 C Perform the preconditioning operation
467         CALL DCOPY(N, W(1,LOCZ), 1, W(1,LOCY), 1)
468         GO TO 220
469     END IF
470
471     CALL CPU_TIME(TARRAY(2))
472
473     CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
474     IF (ISNAN(TEMP(1))) THEN
475         INFO(1) = -99
476         WRITE (6, FMT=9020) INFO(1)
477         CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
478         GO TO 230
479     END IF
480
```

```

481 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
482 C Solution found
483 WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
484 IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
485
486 IF (CALLS.LE.5) THEN
487 GO TO 210
488 END IF
489
490 230 CONTINUE
491
492 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
493 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC BICG CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
494 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
495 WRITE(10,3030) 'BiCG'
496
497 CALLS = 0
498 310 CONTINUE
499 CALLS = CALLS + 1
500 RUNS = RUNS + 1
501 C Clear everything
502 CALL MI25ID(ICNTL, CNTL, ISAVE, RSAVE)
503 CALL CLEARALL(IACT, N, W, LDW, LOCY2, LOCZ2, RESID, ICNTL, CNTL,
504 + INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
505 C Perform an iteration of the method
506
507 CALL CPU_TIME(TARRAY(1))
508 320 CONTINUE
509 CALL MI25AD(IACT, N, W, LDW, LOCY2, LOCZ2, RESID, ICNTL, CNTL,
510 + INFO, ISAVE, RSAVE)
511
512 IF (MVOPP.GE.MAXMVP) THEN
513 CALL CPU_TIME(TARRAY(2))
514 WRITE (6, FMT=9020) -88
515 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, -88)
516 GO TO 330
517 END IF
518
519 IF (IACT.LT.0) THEN
520 CALL CPU_TIME(TARRAY(2))
521 WRITE (6, FMT=9020) INFO(1)
522 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
523 GO TO 330
524 END IF
525
526 IF (IACT.EQ.2) THEN
527 C Perform the matrix-vector products
528 CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,LOCZ2(1)),
529 + TEMP, N)
530 CALL AMUX(N, TEMP, W(1, LOCY2(1)), RVALO, JNDXO,
531 + IND XO)
532 CALL ATMUX(N, W(1,LOCZ2(2)), TEMP, RVALO, JNDXO,
533 + IND XO)
534 CALL DGNLAMGTSOL(ILUPARAM, ILUPREC, TEMP,
535 + W(1, LOCY2(2)), N)
536 MVOPP = MVOPP + 2
537 GO TO 320
538

```

```

541 IF
542
543 IF (IACT.EQ.3) THEN
544 C Perform the preconditioning operations
545 CALL DCOPY(N, W(1,LOCZ2(1)), 1, W(1,LOCY2(1)), 1)
546 CALL DCOPY(N, W(1,LOCZ2(2)), 1, W(1,LOCY2(2)), 1)
547 GO TO 320
548 END IF
549
550 CALL CPU_TIME(TARRAY(2))
551
552 CALL DGNLAMGSQL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
553 IF (ISNAN(TEMP(1))) THEN
554 INFO(1) = -99
555 WRITE (6, FMT=9020) INFO(1)
556 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
557 GO TO 330
558 END IF
559
560 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
561
562 C Solution found
563 WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
564 IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
565
566 IF (CALLS.LE.5) THEN
567 GO TO 310
568 END IF
569
570 330 CONTINUE
571
572 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
573 CCCCCCCCCCCCCCCCCCCCCCCCCCGS CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
574 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
575 WRITE(10,3030) 'CGS'
576
577 CALLS = 0
578 410 CONTINUE
579 CALLS = CALLS + 1
580 RUNS = RUNS + 1
581 C Clear everything
582
583 CALL MI23ID(ICNTL, CNTL, ISAVE, RSAVE)
584 CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
585 + INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
586
587 C Perform an iteration of the BICG-STAB method
588
589 CALL CPU_TIME(TARRAY(1))
590 420 CONTINUE
591 CALL MI23AD(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
592 + INFO, ISAVE, RSAVE)
593
594 IF (MVOPP.GE.MAXMVP) THEN
595 CALL CPU_TIME(TARRAY(2))
596 WRITE (6, FMT=9020) -88
597 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, -88)
598 GO TO 430
599 END IF
600
```

```

602      IF (IACT.T.0) THEN
603          CALL CPU_TIME(TARRAY(2))
604          WRITE (6, FMT=9020) INFO(1)
605          CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
606          GO TO 430
607      END IF
608
609 C Perform the matrix-vector product
610     CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,LOCZ), TEMP, N)
611     CALL AMUX(N, TEMP, W(1, LOCY), RVALO, JNDXO, IND XO)
612     MVOPP = MVOPP + 1
613     GO TO 420
614 END IF
615
616 IF (IACT.EQ.3) THEN
617 C Perform the preconditioning operation
618     CALL DCOPY(N, W(1,LOCZ), 1, W(1,LOCY), 1)
619     GO TO 420
620 END IF
621
622 CALL CPU_TIME(TARRAY(2))
623
624 CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
625 IF (ISNAN(TEMP(1))) THEN
626     INFO(1) = -99
627     WRITE (6, FMT=9020) INFO(1)
628     CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
629     GO TO 430
630 END IF
631
632 CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
633
634 C Solution found
635     WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
636     IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
637
638 IF (CALLS.LE.5) THEN
639     GO TO 410
640 END IF
641
642 430 CONTINUE
643
644 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
645 CCCCCCCCCCCCCCCCCCCCCCCCCcccccc GMRES CCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
646 CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
647     WRITE(10,3030) 'GMRES'
648
649     CALLS = 0
650     M = 100
651     LDH = M+1
652     ALLOCATE(H(LDH,M+2))
653     DEALLOCATE(W)
654     ALLOCATE(W(LDW,M+7))
655 510 CONTINUE
656     CALLS = CALLS + 1
657     RUNS = RUNS + 1
658 C Clear everything
659
660 CALL MI24ID(ICNTL, CNTL, ISAVE, RSAVE, LSAVE)
```

```

661      CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
662      +          INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
663
664 C   Perform an iteration of the method
665
666      CALL CPU_TIME(TARRAY(1))
667 520 CONTINUE
668      CALL MI24AD(IACT, N, M, W, LDW, LOCY, LOCZ, H, LDH, RESID,
669      +          ICNTL, CNTL, INFO, ISAVE, RSAVE, LSAVE)
670
671      IF (MVOPP.GE.MAXMVP) THEN
672          CALL CPU_TIME(TARRAY(2))
673          WRITE (6, FMT=9020) -88
674          CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, -88)
675          GO TO 530
676      END IF
677
678      IF (IACT.LT.0) THEN
679          CALL CPU_TIME(TARRAY(2))
680          WRITE (6, FMT=9020) INFO(1)
681          CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
682          GO TO 530
683      END IF
684
685      IF (IACT.EQ.2) THEN
686 C   Perform the matrix-vector products
687          CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,LOCZ), TEMP, N)
688          CALL AMUX(N, TEMP, W(1, LOCY), RVALO, JNDXO, IND XO)
689          MVOPP = MVOPP + 1
690          GO TO 520
691      END IF
692
693      IF (IACT.EQ.3) THEN
694 C   Perform the preconditioning operations
695          CALL DCOPY(N, W(1,LOCZ), 1, W(1,LOCY), 1)
696          GO TO 520
697      END IF
698
699      CALL CPU_TIME(TARRAY(2))
700
701      CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
702      IF (ISNAN(TEMP(1))) THEN
703          INFO(1) = -99
704          WRITE (6, FMT=9020) INFO(1)
705          CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
706          GO TO 530
707      END IF
708
709      CALL SAVERESULTS(TARRAY, INFO(2), MVOPP, INFO(1))
710
711 C   Solution found
712      WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
713      IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
714
715      IF (CALLS.LE.5) THEN
716          GO TO 510
717      END IF
718
719 530 CONTINUE
720

```



```

721      DEALLOCATE (H)
722
723      CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
724      CCCCCCCCCCCCCCCCCCCCCCCCCC     BICGSTAB1     CCCCCCCCCCCCCCCCCCCCCCCCCC
725      CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
726      L = 1
727      600 CONTINUE
728      L = L + 1
729      WRITE(10,3031) L
730      CALLS = 0
731      ALLOCATE(WORK(N, 3+2*(L+1)))
732      LDRW = (L+1)*(3+2*(L+1))
733      LDWB = N*(3+2*(L+1))
734      ALLOCATE(RWORK(L+1, 3+2*(L+1)))
735      ALLOCATE(IWORK(L+1))
736      610 CONTINUE
737      CALLS = CALLS + 1
738      RUNS = RUNS + 1
739
740      MXMV = MIN(2 * N, MAXMVP)
741      TOL = ETA
742      C Clear everything
743      CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
744      +          INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
745
746      C Perform an iteration of the method
747
748      CALL CPU_TIME(TARRAY(1))
749      CALL BISTBL(L, N, W(1,2), W(1,1), MV, PRECSOLVE, TOL,
750      +          MXMV, WORK, LDWB, RWORK, LDRW, IWORK, INFO, NNZ,
751      +          RVALO, JNDXO, INDXO, ILUPARAM, ILUPREC, TEMP, BNRM2)
752
753      IF (INFO(1).NE.0) THEN
754      CALL CPU_TIME(TARRAY(2))
755      WRITE (6, FMT=9020) INFO(1)
756      CALL SAVERESULTS(TARRAY, L * INFO(2), MXMV, INFO(1))
757      GO TO 630
758      END IF
759
760      CALL CPU_TIME(TARRAY(2))
761
762      CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,2), TEMP, N)
763      IF (ISNAN(TEMP(1))) THEN
764      INFO(1) = -99
765      WRITE (6, FMT=9020) INFO(1)
766      CALL SAVERESULTS(TARRAY, L * INFO(2), MXMV, INFO(1))
767      GO TO 630
768      END IF
769
770      CALL SAVERESULTS(TARRAY, L * INFO(2), MXMV, INFO(1))
771
772      C Solution found
773      WRITE (6, FMT=9000) INFO(2), (TEMP(I), I=1, NPRT)
774      IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
775
776      IF (CALLS.LE.5) THEN
777      GO TO 610
778      END IF
779
780      630 CONTINUE

```

```

781
782      DEALLOCATE(WORK)
783      DEALLOCATE(RWORK)
784      DEALLOCATE(IWORK)
785
786      IF (L.LT.3) THEN
787      GO TO 600
788      END IF
789
790      CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
791      CCCCCCCCCCCCCCCCCCCCCCCCCC     QMR     CCCCCCCCCCCCCCCCCCCCCCCCCC
792      CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
793      WRITE(10,3030) 'QMR'
794
795      CALLS = 0
796      MAXVW = 1
797      MAXPQ = 1
798      MVEC = MAXPQ + MAXVW
799      M = MAXPQ + MAXVW + 2
800      L = N
801
802      ALLOCATE(WORK(M, 8*M+18))
803      ALLOCATE(IWORK2(6, L+2))
804      ALLOCATE(IWORK3(M, 13))
805      DEALLOCATE(W)
806      ALLOCATE(W(LDW, 5*MVEC+3))
807
808      710 CONTINUE
809      CALLS = CALLS + 1
810      RUNS = RUNS + 1
811      C Clear everything
812      CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
813      +          INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
814
815      CALL DCOPY(N, RHS, 1, W(1,2), 1)
816      RESID = 1.0
817      TOL = ETA
818      MAXVW = 1
819      MAXPQ = 1
820      MVEC = MAXPQ + MAXVW
821      M = MAXPQ + MAXVW + 2
822      L = N
823
824      C Perform an iteration of the method
825
826      CALL CPU_TIME(TARRAY(1))
827      720 CONTINUE
828      CALL DUCPL(LDW, N, L, MAXPQ, MAXVW, M, MVEC, RESID, WORK, IWORK2,
829      +          IWORK3, W, TOL, INFO)
830
831      IF (MVOPP.GE.MAXMVP) THEN
832      CALL CPU_TIME(TARRAY(2))
833      WRITE (6, FMT=9020) -88
834      CALL SAVERESULTS(TARRAY, L, MVOPP, -88)
835      INFO(2) = -1
836      CALL DUCPL(LDW, N, L, MAXPQ, MAXVW, M, MVEC, RESID, WORK,
837      +          IWORK2, IWORK3, W, TOL, INFO)
838      GO TO 730
839      END IF
840

```

```

841     IF (INFO(2).EQ.1) THEN
842 C   Perform the matrix-vector products
843     CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,INFO(3)), TEMP, N)
844     CALL AMUX(N, TEMP, W(1, INFO(4)), RVALO, JNDXO, IND XO)
845     MVOPP = MVOPP + 1
846     GO TO 720
847   END IF
848
849   IF (INFO(2).EQ.2) THEN
850 C   Perform the transpose matrix-vector products
851     CALL ATMUX(N, W(1,INFO(3)), TEMP, RVALO, JNDXO, IND XO)
852     CALL DGNLAMGSOL(ILUPARAM, ILUPREC, TEMP, W(1, INFO(4)), N)
853     MVOPP = MVOPP + 1
854     GO TO 720
855   END IF
856
857   IF (INFO(1).NE.0) THEN
858     CALL CPU_TIME(TARRAY(2))
859     WRITE (6, FMT=9020) INFO(1)
860     CALL SAVERESULTS(TARRAY, L, MVOPP, INFO(1))
861     GO TO 730
862   END IF
863
864   CALL CPU_TIME(TARRAY(2))
865
866   CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,1), TEMP, N)
867   IF (ISNAN(TEMP(1))) THEN
868     INFO(1) = -99
869     WRITE (6, FMT=9020) INFO(1)
870     CALL SAVERESULTS(TARRAY, L, MVOPP, INFO(1))
871     GO TO 730
872   END IF
873
874   CALL SAVERESULTS(TARRAY, L, MVOPP, INFO(1))
875
876 C   Solution found
877   WRITE (6, FMT=9000) L, (TEMP(I), I=1, NPRT)
878   IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
879
880   IF (CALLS.LE.5) THEN
881     GO TO 710
882   END IF
883
884 730 CONTINUE
885
886   DEALLOCATE(WORK)
887   DEALLOCATE(IWORK2)
888   DEALLOCATE(IWORK3)
889
890   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
891   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCC TFQMR CCCCCCCCCCCCCCCCCCCCCCCCCC
892   CCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCCC
893   WRITE(10,3030) 'TFQMR'
894
895   CALLS = 0
896
897   DEALLOCATE(W)
898   ALLOCATE(W(LDW, 13))
899
900 810 CONTINUE

```

```

901   CALLS = CALLS + 1
902   RUNS = RUNS + 1
903 C   Clear everything
904
905   CALL CLEARALL(IACT, N, W, LDW, LOCY, LOCZ, RESID, ICNTL, CNTL,
906   +           INFO, ISAVE, RSAVE, RHS, PREC, TEMP, ETA)
907
908   CALL DCOPI(N, RHS, 1, W(1,2), 1)
909   L = N
910   TOL = ETA
911
912 C   Perform an iteration of the method
913
914   CALL CPU_TIME(TARRAY(1))
915 820 CONTINUE
916   CALL DUTFX (LDW, N, L, W, TOL, INFO)
917
918   IF (MVOPP.GE.MAXMVP) THEN
919     CALL CPU_TIME(TARRAY(2))
920     WRITE (6, FMT=9020) -88
921     CALL SAVERESULTS(TARRAY, L, MVOPP, -88)
922     INFO(2) = -1
923     CALL DUTFX (LDW, N, L, W, TOL, INFO)
924     GO TO 830
925   END IF
926
927   IF (INFO(2).EQ.1) THEN
928 C   Perform the matrix-vector products
929     CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,INFO(3)), TEMP, N)
930     CALL AMUX(N, TEMP, W(1, INFO(4)), RVALO, JNDXO, IND XO)
931     MVOPP = MVOPP + 1
932     GO TO 820
933   END IF
934
935   IF (INFO(1).NE.0) THEN
936     CALL CPU_TIME(TARRAY(2))
937     WRITE (6, FMT=9020) INFO(1)
938     CALL SAVERESULTS(TARRAY, L, MVOPP, INFO(1))
939     GO TO 830
940   END IF
941
942   CALL CPU_TIME(TARRAY(2))
943
944   CALL DGNLAMGSOL(ILUPARAM, ILUPREC, W(1,1), TEMP, N)
945   IF (ISNAN(TEMP(1))) THEN
946     INFO(1) = -99
947     WRITE (6, FMT=9020) INFO(1)
948     CALL SAVERESULTS(TARRAY, L, MVOPP, INFO(1))
949     GO TO 830
950   END IF
951
952   CALL SAVERESULTS(TARRAY, L, MVOPP, INFO(1))
953
954 C   Solution found
955   WRITE (6, FMT=9000) L, (TEMP(I), I=1, NPRT)
956   IF (INFO(1).GT.0) WRITE (6, FMT=9010) INFO(1)
957
958   IF (CALLS.LE.5) THEN
959     GO TO 810
960   END IF

```

```

961
962 830 CONTINUE
963
964 1000 CONTINUE
965   CLOSE(10)
966 C Deallocate everything
967 C Deallocate data for the preconditioner
968   IF (ILUIERR.EQ.0) THEN
969     CALL DGNLAMGDELETE(ILUPARAM, ILUPREC)
970   END IF
971
972   DEALLOCATE(ILUIA)
973   DEALLOCATE(ILUJA)
974   DEALLOCATE(ILUA)
975   DEALLOCATE(ILUIND)
976
977   DEALLOCATE(W)
978
979 1010 CONTINUE
980
981   IF (PREC_TOL.LT.PREC_UPPER_TOL) THEN
982     IF (RUNS.GT.10) THEN
983       GO TO 7
984     END IF
985   END IF
986
987   DEALLOCATE(TEMP)
988   DEALLOCATE(SOLUTION)
989   DEALLOCATE(RHS)
990   DEALLOCATE(RVALO)
991   DEALLOCATE(JNDXO)
992   DEALLOCATE(INDXO)
993
994   GO TO 5
995 2000 CONTINUE
996   CLOSE(11)
997   STOP
998 3010 FORMAT ( '% Matrix:      name,          N,          NNZ', / A18,
999 +          ', ', I8, ', ', ' ', I10 )
1000 3020 FORMAT ( '% Prec:      time,          N,          NNZ', / F18.3,
1001 +          ', ', I8,
1002 +          ', ', ' ', I10,
1003 +          / '% Results:    time,          iter,    m-v prod, error')
1004 3030 FORMAT ( '% ' A)
1005 3031 FORMAT ( '% BICGSTAB(', I1, ')')
1006 3032 FORMAT ( '% ' I4)
1007 4010 FORMAT ( ' New matrix: ', A)
1008 4020 FORMAT ( ' Reading matrix ' )
1009 4030 FORMAT ( ' Converting to CSR ' )
1010 4035 FORMAT ( ' Forming RHS ' )
1011 4036 FORMAT ( ' No RHS file found ' )
1012 4037 FORMAT ( ' Already done ' )
1013 4040 FORMAT ( ' Building preconditioner ' )
1014 4050 FORMAT ( ' Starting solvers ' )
1015 9000 FORMAT ( / ' Solution found', / I6, ' iterations required ',
1016 +          '// Solution = ', / (1P, 5D10.2))
1017 9010 FORMAT ( ' Warning: INFO( 1 ) = ', I3, ' on exit ' )
1018 9020 FORMAT ( ' Error return: INFO( 1 ) = ', I3, ' on exit ' )
1019 9990 FORMAT ( ' MATRIX NOT SQUARE ' )
1020   END PROGRAM TEST

```

```

1021
1022 C Subroutines used by BiCGSTAB(ell)
1023   SUBROUTINE MV(N, X, Y, NNZ, RVALO, JNDXO,
1024 +     INDXO, ILUPARAM, ILUPREC, TEMP)
1025     INTEGER N
1026     DOUBLE PRECISION X(N), Y(N)
1027     DOUBLE PRECISION RVALO(NNZ), TEMP(N)
1028     INTEGER INDXO(N+1), JNDXO(NNZ)
1029     INTEGER*8 ILUPARAM, ILUPREC
1030     EXTERNAL AMUX, DGNLAMGSOL
1031     CALL DGNLAMGSOL(ILUPARAM, ILUPREC, X, TEMP, N)
1032     CALL AMUX(N, TEMP, Y, RVALO, JNDXO, INDXO)
1033   END SUBROUTINE MV
1034
1035   SUBROUTINE PRECSOLVE(N, X)
1036     INTEGER N
1037     DOUBLE PRECISION X(N)
1038 C Right preconditioning. Do nothing
1039   END SUBROUTINE PRECSOLVE
1040
1041 C Subroutine to save all results
1042   SUBROUTINE SAVERESULTS(TARRAY, ITER, MVOPP, ERROR)
1043     REAL TARRAY(2)
1044     INTEGER ITER, MVOPP, ERROR
1045     WRITE(10,100) TARRAY(2)-TARRAY(1), ITER, MVOPP, ERROR
1046     CALL FLUSH(10)
1047     MVOPP = 0
1048     RETURN
1049 100 FORMAT (F18.3, ', ', I8, ', ', ' ', I10,
1050 +          ', ', ' ', I3)
1051   END SUBROUTINE SAVERESULTS

```

## E Implementation of the data analysis tool

```

1  #!/usr/bin/env python
2  # -*- coding: UTF-8 -*-
3
4  import os
5  import sys
6  import scipy
7  import matplotlib.pyplot as pyplot
8  from optparse import OptionParser
9  import shutil
10
11 import TableToLatex
12 import MatrixFileFetcher
13
14 class Parser(object):
15     ''' The parser class that parses the matrix files of a certain
16         preconditioner tolerance'''
17
18     TIMETOL = 0.05
19     OFFTOL = 0.1
20     UPPERTOL = 150.0
21     MATRIXMODE = 4
22     COLUMNS = 4
23     MAX_MVPS = 20000
24
25     def __init__(self, prec_tol, prec_tols):
26         ''' constructor '''
27         self.mode = 0
28         self.solver = 0
29         self.options = self.set_options()
30         self.solver_dict = ['CORS', 'BiCOR', 'BiCGSTAB', 'BiCG',
31                             'CGS', 'GMRES(100)', 'BiCGSTAB(2)', 'BiCGSTAB(3)', 'QMR', 'TFQMR']
32         self.solver_used_amount = len(self.solver_dict)
33         self.exclude = sorted(self.options.exclude, reverse=True)
34         for i in self.exclude:
35             self.solver_dict.pop(i)
36         self.solver_amount = len(self.solver_dict)
37         self.prec_tol = prec_tol
38         self.prec_tols = prec_tols
39         self.prefix = './test'
40         self.matrix_file = ''
41         self.matrix_parser = MatrixParser(Parser.COLUMNS)
42         self.matrix_name = ''
43         self.matrix_number = 0
44         self.data = SolverData(self.solver_amount)
45         self.current = scipy.zeros([self.solver_used_amount,
46                                     Parser.COLUMNS + 1])
47         self.plotting_pattern = ['b-', 'g--', 'r-', 'c:', 'm-', 'y--',
48                                 'k-', 'b:', 'g-', 'r--']
49
50         self.skip = False
51         self.rerun = ''
52         self.time = 0
53
54         self.latex_list = []
55         self.current_latex_list = []
56         self.first = True
57
58     def run(self):
59         ''' run the parser '''
60         filename_list = os.listdir(self.prefix)
61         filename_list = sorted(filename_list, key=str.lower)

```

```

61
62     try:
63         matrix_file_list, dummy = MatrixFileFetcher.fetch()
64         matrix_file_list = matrix_file_list \
65             if (self.options.n <= 0 or self.options.n > len(matrix_file_list)) \
66             else matrix_file_list[0:self.options.n]
67     except IOError:
68         matrix_file_list = []
69     for filename in filename_list:
70         matrix, sep, prec_tol = filename.rpartition('_')
71         try:
72             prec_tol = float(prec_tol.rpartition('.')[0])
73         except ValueError:
74             print prec_tol
75             self.quit_on_error()
76         if prec_tol != self.prec_tol or (matrix_file_list and matrix not in
77             matrix_file_list):
78             if matrix_file_list and matrix not in matrix_file_list:
79                 print 'Matrix not in matrix file list: ', matrix
80             continue
81         self.current = scipy.zeros([self.solver_used_amount,
82                                     Parser.COLUMNS + 1])
83
84         self.mode = 0
85         self.solver = 0
86         self.skip = False
87         self.matrix_file = self.prefix+'/'+filename
88         matrix_file = open(self.matrix_file)
89         lines = matrix_file.readlines()
90         matrix_file.close()
91         for line in lines:
92             if line.startswith('%'):
93                 # line is a comment
94                 self.mode += 1
95             if self.mode > Parser.MATRIXMODE:
96                 # read the data of the previous solver into self.current
97                 self.read()
98                 self.solver += 1
99                 self.matrix_parser.clear()
100             elif self.mode == 1:
101                 # line with the matrix name
102                 spline = self.split(line)
103                 self.matrix_name = spline[0]
104                 self.current_latex_list = [self.matrix_name]
105                 if self.options.verbose:
106                     print 'Read: {}: {}'.format(self.matrix_number + 1,
107                                                 self.matrix_name)
108             elif self.mode >= Parser.MATRIXMODE:
109                 # line contains solver data
110                 spline = self.split(line)
111                 self.matrix_parser.add(spline)
112                 if self.skip:
113                     break
114             if self.skip:
115                 continue
116             self.read()
117             self.write_time_to_file()
118             self.add()
119             self.draw_plot()
120             self.draw_latex_tables()
121             self.draw_table()

```

```

120
121 def split(self, line):
122     ''' separate the data on each line '''
123     spline = line.split(',')
124     for i in range(len(spline)):
125         spline[i] = spline[i].strip()
126     return spline
127
128 def read(self):
129     ''' read the data for one parser '''
130     if self.solver in self.exclude:
131         return
132
133     data = self.matrix_parser.read()
134
135     times = data[:, 0]
136     self.time += sum(times)
137     length = len(times)
138     if length == 0:
139         print 'Length is 0: {}'.format(self.matrix_name)
140         self.skip = True
141         return
142
143     average = float(sum(times)) / length
144
145     if self.matrix_parser.fail:
146         excl = len([item for item in self.exclude if item < self.solver])
147         err = int(data[0, 3])
148         # 1 is breakdown, 2 is iterations, 3 is NAN
149         error_type = 2 if (err < -3 and err > -99) \
150             or ((self.solver == 6 or self.solver == 7) and err == 1) \
151             or ((self.solver == 8 or self.solver == 9) and err == 4) \
152             else (3 if err == -99 else 1)
153         self.data.solver[self.solver - excl].error(error_type)
154         self.current[self.solver, 3] = data[0, 3]
155         self.write_to_current(average, data, error_type)
156         return
157
158     # check the time values for validity
159     for value in times:
160         if value < Parser.TIMETOL:
161             print 'Time too short: {}, {}, {}'.format(value,
162                 average, self.matrix_name)
163             if not self.skip:
164                 self.skip = True
165                 self._move_matrix('small')
166                 self._move_results('small')
167         if value > (1.0 + Parser.OFFTOL) * average:
168             print 'Result not accurate enough: {}, {}, {}'.format(value, average, self.matrix_name)
169             length -= 1
170             average = float(average * (length + 1)) / length
171
172     self.write_to_current(average, data)
173
174 def write_to_current(self, average, data, error=None):
175     ''' write the data to self.current '''
176     self.current[self.solver, 0] = average
177     self.current[self.solver, 1] = data[0, 1]
178     self.current[self.solver, 2] = data[0, 2]

```

```

180     self.current[self.solver, 3] = data[0, 3]
181     if error is None:
182         self.current_latex_list.append('')
183     elif self.solver not in self.exclude:
184         self.current_latex_list.append(
185             '\\colorbox{yellow}{iterations}' if error == 2 else \
186             '\\colorbox{magenta}{NAN}' if error == 3 else \
187             '\\colorbox{red}{breakdown}')
188
189 def add(self):
190     ''' add the data for the current matrix to self.data usning the
191         ratio method '''
192     if self.options.restriction is not None and not \
193         eval(str(self.current[int(self.options.restriction[0]), 3]) \
194             + self.options.restriction[1]):
195         return
196     for i in range(3):
197         self.ratio(i)
198     if not self.skip:
199         self.matrix_number += 1
200         self.latex_list.append(self.current_latex_list)
201         print 'Added: {}: {}'.format(self.matrix_number,
202             self.matrix_name)
203
204 def ratio(self, column):
205     ''' calculate the ratio and add it to self.data '''
206     invalid_runs = sorted([item for item in range(self.solver_used_amount) \
207         if self.current[item, 3] != 0 or item in self.exclude],
208         reverse = True)
209     stat = self.current[:, column]
210     valid_runs = list(stat)
211     for i in invalid_runs:
212         valid_runs.pop(i)
213     if valid_runs == []:
214         if column == 1:
215             print 'No valid runs at all: {}'.format(self.matrix_name)
216             if self.prec_tol == self.prec_tols[0]:
217                 self._move_matrix('bad')
218                 self._move_results('bad')
219                 self.skip = True
220             winner = 1
221         else:
222             winner = min(valid_runs)
223         j = 0
224         latex_list = {}
225         # calculate the ratios of all used solvers and add them where needed
226         for i in range(self.solver_used_amount):
227             if i in self.exclude:
228                 continue
229             if self.current[i, 3] != 0:
230                 ratio = Parser.UPPERTOL
231                 if column == 0 and self.options.verbose:
232                     print 'Failed method: {}, {}'.format(self.matrix_name,
233                         i)
234             elif column == 2 and (self.current[i, 3] == -88 or \
235                 (self.current[i, 3] == 1 and (i == 6 or i == 7))) and \
236                 self.current[i, column] < 20 * winner and \
237                 self.current[i, column] >= Parser.MAX_MVPS:
238                 if self.options.verbose:
239                     print 'Max runs not enough: {}, {}, {}, {}'.format(\

```

```

240         self.matrix_name, i, winner,
241         self.current[i, column])
242     self.rerun += '{!s}\n{:d}\n{:05.2f}\n{:d}\n'.format(\
243         self.matrix_name, i, self.prec_tol, int(winner))
244
245     else:
246         ratio = float(stat[i]) / winner
247         if ratio >= Parser.UPPERTOL:
248             print 'Ratio too high: {}, {}, {}, {}, {}'.format(stat[i],
249                 winner, self.matrix_name, i, self.prec_tol)
250             print stat
251             self.quit_on_error()
252         if column == 2 and stat[i] > Parser.MAX_MVPS:
253             print 'Solver converged with more than {} mvps: {}, \
254                 {}'.format(Parser.MAX_MVPS, stat[i], self.matrix_name)
255         self.data.solver[j].append(ratio, column)
256         latex_list[j] = ratio
257         j += 1
258     if column == 0:
259         sorted_latex_list = sorted(latex_list, key=latex_list.get)
260         for i, item in enumerate(sorted_latex_list):
261             if latex_list[item] != Parser.UPPERTOL:
262                 self.current_latex_list[item+1] = ('\\colorbox{green!'+\
263                     str(int(100/len(latex_list)*(len(latex_list)-i)))+'\
264                     'black}{'+str(i+1)+'}')
265                 #print self.current_latex_list[item]
266
267     def draw_latex_tables(self):
268         ''' draw a latex table of all the errors '''
269         tex_parser = TableToLatex.TableToLatex()
270         header_list = list(self.solver_dict)
271         header_list.insert(0, 'matrix name')
272         tex_parser.set_header(header_list,
273             '\p{100px}|' + '\p{65px}|'*len(header_list)-1))
274         tex_parser.add_package('xcolor')
275         tex_parser.set_table(self.latex_list)
276         tex_parser.make('table' + str(self.prec_tol))
277         tex_parser = TableToLatex.TableToLatex()
278         header_list = ['Solver', 'breakdown', 'iterations', 'NAN']
279         tex_parser.set_header(header_list)
280         table = []
281         for i in range(self.solver_amount):
282             table.append([self.solver_dict[i],
283                 str(self.data.solver[i].errors),
284                 str(self.data.solver[i].iter_errors),
285                 str(self.data.solver[i].nan_errors)])
286         tex_parser.use_separator(False)
287         tex_parser.set_table(table)
288         tex_parser.set_caption('Failures with a preconditioner tolerance of '\
289             + str(self.prec_tol), str(self.prec_tol))
290         tex_parser.make('failure_table' + str(self.prec_tol))
291
292     def draw_table(self):
293         ''' draw a table containing some useful numbers '''
294         print 'Errors'
295         for i in range(self.solver_amount):
296             print '{:12}: {}'.format(self.solver_dict[i],
297                 self.data.solver[i].errors)
298         print 'More than max iterations'
299         for i in range(self.solver_amount):
300             print '{:12}: {}'.format(self.solver_dict[i],
301                 self.data.solver[i].iter_errors)
302
303         print 'Winner in terms of time'
304         for i in range(self.solver_amount):
305             print '{:12}: {}'.format(self.solver_dict[i],
306                 self.data.solver[i].timer.count(1.0))
307
308         print 'Winner in terms of matrix-vector products'
309         for i in range(self.solver_amount):
310             print '{:12}: {}'.format(self.solver_dict[i],
311                 self.data.solver[i].mvps.count(1.0))
312
313     def draw_plot(self):
314         ''' draw the plots '''
315         time_plot = self.make_figure(1)
316         mvp_plot = self.make_figure(2)
317         iter_plot = self.make_figure(3)
318         time_plot.savefig(str(self.prec_tol) + '_time.png', bbox_inches='tight')
319         mvp_plot.savefig(str(self.prec_tol) + '_mvp.png', bbox_inches='tight')
320         iter_plot.savefig(str(self.prec_tol) + '_iter.png', bbox_inches='tight')
321         #pyplot.show()
322
323     def make_figure(self, number):
324         ''' make a figure '''
325         figure = pyplot.figure()
326         #pyplot.subplots_adjust(left=0.1, right=0.1, top=0.1, bottom=0.1)
327         plot = figure.add_subplot(111)
328         for i in range(self.solver_amount):
329             if number == 1:
330                 data = self.data.solver[i].timer
331             elif number == 2:
332                 data = self.data.solver[i].mvps
333             elif number == 3:
334                 data = self.data.solver[i].iterations
335             self._plot(data, i, plot)
336         leg = plot.legend(self.solver_dict, loc=4)
337         frame = leg.get_frame()
338         frame.set_alpha(0.8)
339         plot.set_xlabel(u' ', size='large')
340         plot.set_ylabel(u' ( )', size='large')
341         return figure
342
343     def _plot(self, data, index, plot):
344         ''' specify what the plot looks like '''
345         xlim = self.options.xlim if self.options.xlim > 1 else 10
346         x = scipy.arange(1, xlim, xlim/1000.0)
347         y = []
348         for xi in x:
349             y.append(1.0 / self.matrix_number * \
350                 len([item for item in data if item <= xi]))
351         plot.plot(x, y, self.plotting_pattern[index], linewidth=2)
352         plot.set_xlim([1, xlim])
353         plot.set_ylim([0, 1.1])
354         plot.hold(True)
355
356     def write_time_to_file(self):
357         if self.first:
358             time_file = open('time.txt', 'w')
359             self.first = False
360         else:
361             time_file = open('time.txt', 'a')
362             time_file.write(self.matrix_name+' ('+str(round(sum(self.current[:,0])))+')')

```

```

        \n')
    time_file.close()
360
361
362 def _move_matrix(self, dest):
363     ''' move a bad matrix '''
364     for pref in ['', '_b', '_rhs1', '_x']:
365         name = self.matrix_name + pref + '.mtx'
366         try:
367             os.rename('/home/sven/matrices/' + name,
368                     '/home/sven/matrices-' + dest + '/' + name)
369         except OSError, err:
370             if pref == '':
371                 print 'Moving matrix failed: {}, {}'.format(name, err)
372
373 def _move_results(self, dest):
374     ''' move bad test results '''
375     if dest == 'small':
376         pref = './test-small/'
377     else:
378         pref = '/home/sven/matrices-bad/'
379     for tol in self.prec_tols:
380         try:
381             name = self.matrix_name + '_%05.2f.txt' % tol
382             os.rename(self.prefix + '/' + name, pref + name)
383         except OSError, (errno, err):
384             if errno == 18:
385                 try:
386                     shutil.move(self.prefix + '/' + name, pref + name)
387                 except IOError:
388                     if dest == 'small' or tol == self.prec_tols[0]:
389                         print 'Moving results failed: {}, {}'.format(name, err)
390                     else:
391                         pass
392                     elif dest == 'small' or tol == self.prec_tols[0]:
393                         print 'Moving results failed: {}, {}'.format(name, err)
394                     else:
395                         pass
396
397     return
398
399 def quit_on_error(self):
400     ''' error '''
401     print 'Aborting: an error occurred'
402     sys.exit(1)
403
404 def set_options(self):
405     ''' set all options '''
406     option_parser = OptionParser()
407     option_parser.add_option('-x', '--xlim', dest='xlim', default=10,
408                             type='float',
409                             help='the maximum ratio on the x axis of the plots')
410     option_parser.add_option('-l', dest='l', default=3,
411                             type='int',
412                             help='the value of l to use in BiCGSTAB(l), 0 means both')
413     option_parser.add_option('-e', '--exclude', dest='exclude',
414                             type='int', action='append',
415                             help='solver to exclude')
416     option_parser.add_option('-r', '--restriction', dest='restriction',
417                             type='string', nargs=2,
418                             help='only use the matrices where solver has an error \

```

```

        conform to the supplied test, i.e. 5 \<-4\ \
        means solver 5 has an error value smaller than -4')
420
421 option_parser.add_option('-v', '--verbose', dest='verbose',
422                         action='store_true', default=False,
423                         help='print more output (failed methods)')
424 option_parser.add_option('-n', '--filelim', dest='n', default=0,
425                         type='int',
426                         help='the amount of matrices to parse')
427 (options, args) = option_parser.parse_args()
428 # add bicgstab(1) to the exclude list
429 if options.exclude is None:
430     options.exclude = []
431 if options.l == 2:
432     options.exclude.append(options.l + 5)
433 elif options.l == 3:
434     options.exclude.append(options.l + 3)
435 return options
436
437 class MatrixParser(object):
438     def __init__(self, columns):
439         self.runs = 6
440         self.columns = columns
441         self.clear()
442
443     def clear(self):
444         self.length = 0
445         self.data = scipy.zeros([self.runs, self.columns])
446         self.fail = False
447
448     def add(self, line):
449         for i, value in enumerate(line):
450             self.data[self.length, i] = \
451                 float(value) if '.' in value else int(value)
452         if int(line[3]) != 0:
453             self.fail = True
454         self.length += 1
455
456     def read(self):
457         return self.data[0:self.length, :]
458
459 class SolverData(object):
460     ''' class that contains the data of all solvers using multiple
461         Solver objects'''
462     def __init__(self, amount):
463         self.solver = []
464         self.amount = amount
465         for i in range(self.amount):
466             self.solver.append(Solver())
467
468 class Solver(object):
469     ''' class that contain the data of one solver '''
470     def __init__(self):
471         self.timer = []
472         self.iterations = []
473         self.mvps = []
474         self.errors = 0
475         self.iter_errors = 0
476         self.nan_errors = 0
477
478     def error(self, error_type):

```



```

479         if error_type == 1:
480             self.errors += 1
481         elif error_type == 2:
482             self.iter_errors += 1
483         elif error_type == 3:
484             self.nan_errors += 1
485
486     def append(self, ratio, number):
487         if number == 0:
488             self.timer.append(ratio)
489         elif number == 1:
490             self.iterations.append(ratio)
491         elif number == 2:
492             self.mvps.append(ratio)
493
494     def main():
495         ''' main method '''
496         def write_to_file(text):
497             '''write the results to a file'''
498             f = open('rerun.txt', 'w')
499             f.write(text)
500             f.close()
501
502         tolerances = [0.1, 1.0, 10.0]
503         rerun = ''
504         time = 0
505         for i, tol in enumerate(tolerances):
506             parser = Parser(tol, tolerances)
507             parser.first = (i == 0)
508             parser.run()
509             rerun += parser.rerun
510             time += parser.time
511             parser = None
512         print 'Total time: {}'.format(time)
513         write_to_file(rerun)
514
515     if __name__ == "__main__":
516         main()

```