

A NOTE ON GMRES PRECONDITIONED BY A PERTURBED LDL^T DECOMPOSITION WITH STATIC PIVOTING*

M. ARIOLI[†], I. S. DUFF[‡], S. GRATTON[§], AND S. PRALET[¶]

Abstract. A strict adherence to threshold pivoting in the direct solution of symmetric indefinite problems can result in substantially more work and storage than forecast by a sparse analysis of the symmetric problem. One way of avoiding this is to use static pivoting where the data structures and pivoting sequence generated by the analysis are respected and pivots that would otherwise be very small are replaced by a user defined quantity. This can give a stable factorization but of a perturbed matrix. The conventional way of solving the sparse linear system is then to use iterative refinement, but there are cases where this fails to converge. In this paper, we discuss the use of more robust iterative methods, namely GMRES and flexible GMRES (FGMRES). We show both theoretically and experimentally that both approaches are more robust than iterative refinement and furthermore that FGMRES is far more robust than GMRES and that, under reasonable hypotheses, FGMRES is backward stable. We also show how restarted variants can be beneficial, although again the GMRES variant is not as robust as FGMRES.

Key words. augmented systems, sparse matrices, multifrontal methods, GMRES, flexible GMRES, static pivoting, roundoff error

AMS subject classifications. 65F05, 65F50, 65F10, 65G50

DOI. 10.1137/060661545

1. Introduction. This paper is concerned with solving the set of linear equations

$$(1.1) \quad Ax = b,$$

where the coefficient matrix $A \in \mathbb{R}^{n \times n}$ is a symmetric indefinite sparse matrix. Our hope is to solve this system using a direct method, but sometimes the cost of doing this is too high in terms of time or memory. We have therefore looked at the possibility of using static pivoting to avoid these problems which are particularly acute if the matrix is highly indefinite as, for example, can happen for saddle-point problems.

We will use a multifrontal approach for our direct method. In this approach, we first determine an order for choosing pivots based on the sparsity structure of A and we then accommodate further pivoting for numerical stability during the subsequent numerical factorization phase. When the matrix is highly indefinite, the resulting pivot sequence used in the numerical factorization can differ substantially from that predicted by the analysis step. A simple way to avoid this problem is to force the elimination through static pivoting.

*Received by the editors May 31, 2006; accepted for publication (in revised form) March 13, 2007; published electronically September 28, 2007.

<http://www.siam.org/journals/sisc/29-5/66154.html>

[†]Rutherford Appleton Laboratory, Chilton, Didcot, Oxon, OX11 0QX, UK (m.arioli@rl.ac.uk). This author's work was supported by EPSRC grant GR/S42170/01.

[‡]Rutherford Appleton Laboratory, Chilton, Didcot, Oxon, OX11 0QX, UK (i.s.duff@rl.ac.uk) and CERFACS, 42 rue G. Coriolis, 31057 Toulouse, France. This author's work was supported by EPSRC grant GR/S42170/01.

[§]CERFACS, 42 rue G. Coriolis, 31057 Toulouse, France (Serge.Gratton@cerfacs.fr).

[¶]SAMTECH s.a. LIEGE Science Park, Rue des Chasseurs-Ardenais, 8, B-4031 Liège, Belgium (stephane.pralet@samcef.com). Part of the work of this author was supported by ENSEEIHT and CNRS.

We assume that the matrix A has been factorized using the HSL MA57 package [7] with the option of using static pivoting [8]. Static pivoting will add a factor τ to a diagonal entry when it is impossible to find a suitable pivot in the fully summed blocks. It is common to choose $\tau \approx \sqrt{\varepsilon} \|A\|$ (ε machine precision).

Therefore, the computed factors \hat{L} and \hat{D} are, in exact arithmetic, the exact factorization of the perturbed problem

$$(1.2) \quad A + E = \hat{L}\hat{D}\hat{L}^T,$$

where E is a diagonal matrix of rank equal to the number of static pivots used during the factorization and $|E| \leq \tau I$. Hereafter, given the matrix B , we denote by $|B|$ the matrix having entries equal to the absolute values of the entries of B . The nonzero diagonal entries in E correspond to the positions at which static pivoting was performed and are all equal to τ . Note that if τ is chosen too small, then the factorization could be very unstable, whereas if it is chosen too large, the factorization will be stable but will not be an accurate factorization of the original matrix (that is, $|E|$ will be large). In section 2, we give a brief description of the static pivoting strategy used in MA57.

Equation (1.2) gives a splitting of A in terms of $M = \hat{L}\hat{D}\hat{L}^T$ and E ,

$$(1.3) \quad A = M - E,$$

and the solution of (1.1) can be expressed as the solution of the equivalent systems

$$(1.4) \quad \begin{cases} (I - EM^{-1})Mx = b, \\ (I - M^{-1}E)x = M^{-1}b. \end{cases}$$

Owing to the symmetry of all the matrices in (1.4), $I - EM^{-1} = (I - M^{-1}E)^T$. Moreover, the first system corresponds to a right preconditioning of (1.1), while the second corresponds to a left preconditioning.

If the spectral radius of the matrix $I - M^{-1}E$ (or $I - EM^{-1}$) is less than one, the system (1.4) can be solved using iterative refinement. This has been used by many authors [5, 8, 13, 18] and is successful over a wide range of matrices although it is sensitive to the value of τ . If, however, the spectral radius is greater than or equal to one (or even close to one), it may be necessary to switch to a more powerful method such as GMRES [17]. Although the matrix is symmetric, we choose GMRES since it gives us much more freedom to work with a wide range of preprocessors and preconditionings.

We have found experimentally that using the factorization (1.2) as a right preconditioning for GMRES works in most cases and is, as expected, much more robust than iterative refinement. However, there are cases where we do not get convergence of a scaled residual to machine precision. Moreover, we have found that restarted GMRES performs better and that using flexible GMRES (FGMRES) [15], even though our preconditioner remains constant, does even better.

We illustrate this through numerical experiment (see section 6) and show theoretically that, under reasonable assumptions, FGMRES preconditioned by our static pivoting factorization is backward stable so that a small scaled residual can be achieved (see section 5). Our analysis also holds for the case of restarted FGMRES that we advocate for controlling the memory requirement while still achieving the desired robustness and accuracy. Indeed, we give theoretical arguments for why restarting often greatly improves the convergence.

In the following, we use the MATLAB notation $[A, B]$ to mean the matrix that consists of the columns of matrix A followed by the columns of matrix B . We denote by ε the machine precision of a finite-precision arithmetic satisfying the IEEE standard, and by $fl(\cdot)$ the result of a sequence of floating-point operations. The computed matrices, vectors, and scalars will be identified by a bar over the symbol. Finally, we will denote by $\|\cdot\|$ the usual Euclidean 2-norm for vectors and the corresponding induced norm for matrices.

2. LDL^T with static pivoting. In the multifrontal context, the factorization can be represented by a tree at each node of which elimination operations are performed on a partially summed frontal matrix

$$(2.1) \quad \begin{bmatrix} F_{11} & F_{12} \\ F_{12}^T & F_{22} \end{bmatrix},$$

and pivots at that stage can be chosen only from within the fully summed block F_{11} . The problem occurs when it is impossible or numerically suicidal to eliminate all of F_{11} , resulting in more work and storage (sometimes dramatically more) than forecast.

Our mixed pivoting approach is based on two phases. In the first phase, we perform numerical pivoting in the block F_{11} of fully summed variables until no remaining variables satisfy the numerical criterion. In the second phase, we eliminate the remaining fully summed variables, adding 1×1 perturbations defined using a threshold $\mu > 0$. Let us denote by a_{ii} the generic 1×1 pivot and by

$$P = \begin{bmatrix} a_{ii} & a_{ij} \\ a_{ji} & a_{jj} \end{bmatrix}$$

the generic 2×2 pivot.

Then, our pivoting strategy is defined as follows. We define

$$(2.2) \quad g_1(i) = \frac{\max_{k \neq i} |a_{ki}|}{|a_{ii}|}$$

and

$$(2.3) \quad g_2(i, j) = \left\| \left\| |P^{-1}| \begin{pmatrix} \max_{k \neq i, j} |a_{ki}| \\ \max_{k \neq i, j} |a_{kj}| \end{pmatrix} \right\| \right\|_{\infty}.$$

During the first phase of our algorithm, we perform the usual Duff–Reid algorithm. In this approach, a 1×1 pivot a_{ii} is considered to be stable if and only if

$$(2.4) \quad g_1(i) \leq 1/u,$$

and a 2×2 pivot is considered to be stable if and only if

$$(2.5) \quad g_2(i, j) \leq 1/u.$$

During the second phase we preselect a 1×1 and a 2×2 pivot and decide which will be eliminated. The choice between a 1×1 and a 2×2 pivot is done in three stages. First, if we can eliminate a pivot and ensure a growth factor lower than $1 + 1/\mu$,

where $\mu = \sqrt{\varepsilon}$, i.e., if $g_1(i) \leq 1/\mu$ or $g_2(i, j) \leq 1/\mu$, then we select the one with the lower growth factor. Second, if $g_1(i) \geq 1/\mu$ and $g_2(i, j) \geq 1/\mu$ (we cannot ensure a growth factor lower than $1+1/\mu$), then we compare the quantities $1/|a_{ii}|$ and $\|P^{-1}\|_\infty$ and select the pivot associated with the smallest quantity. This second comparison is guided by the growth factor that would appear if we suppose that the largest off-diagonal entry is bounded by $\max_{ij} |a_{ij}|$. Finally, if no pivot can be chosen, we add some perturbation to make the factorization more stable. More precisely, a 1×1 pivot a_{ii} is selected and perturbed using $\delta = s(a_{ii})(\tau - |a_{ii}|)$, where s is the sign function and $\tau = \mu \max_{ij} |a_{ij}|$.

3. GMRES. The convergence of the GMRES method can be quite problematic for the general case even when we have a favorable distribution of the eigenvalues as in our case of the matrix $I - M^{-1}E$ (or $I - M^{-1}E$) [2].

However, because of the symmetry of A , the spectrum of $I - M^{-1}E$ (or $I - M^{-1}E$) has additional properties that make the behavior more regular. Let us denote by \tilde{I}_m and \tilde{I}_{n-m} the $n \times m$ and $n \times (n - m)$ matrices corresponding to the first m columns and the last $n - m$ columns of the $n \times n$ identity matrix, respectively. We assume that the matrices E , M , and A are symmetrically permuted so that the nonzero entries of E are the first m diagonal entries of the matrix E .

We observe that \tilde{I}_{n-m} is a basis for $\text{Ker}(E)$.

THEOREM 3.1. *Using the previous notation, if E is nonnegative definite, the spectra of $I - M^{-1}E$ and $I - EM^{-1}$ are real.*

Proof. $M^{-1}Ex = \lambda x$ if and only if $Ex = \lambda Mx$. Since M is symmetric and nonsingular and E is positive semidefinite, λ is either zero or real, and (M, E) is an Hermitian definite pencil (see [3, p. 110]). \square

We can introduce the right preconditioned version of GMRES, where we note that, although the vector z_k is computed, it is not stored.

ALGORITHM 3.1.

```

procedure  $[x] = \text{right\_Prec\_gmres}(A, M, b, \text{maxit})$ 
 $x_0 = M^{-1}b, r_0 = b - Ax_0$  and  $\beta = \|r_0\|$ 
 $v_1 = r_0/\beta; k = 0; r = r_0$ 
while  $\|r\| > \varepsilon (\|b\| + \|A\| \|x_k\|)$  and  $k < \text{maxit}$ 
     $k = k + 1;$ 
     $z_k = M^{-1}v_k; w = Az_k;$ 
    for  $i = 1, \dots, k$  do
         $h_{i,k} = v_i^T w;$ 
         $w = w - h_{i,k}v_i;$ 
    end for;
     $h_{k+1,k} = \|w\|;$ 
     $v_{k+1} = w/h_{k+1,k};$ 
     $V_k = [v_1, \dots, v_k];$ 
     $H_k = \{h_{i,j}\}_{1 \leq i \leq j+1; 1 \leq j \leq k};$ 
     $y_k = \arg \min_y \|\beta e_1 - H_k y\|;$ 
    if  $\|\beta e_1 - H_k y_k\| \leq \varepsilon (\|b\| + \|A\| \|x_k\|)$  do
         $x_k = x_0 + M^{-1}V_k y_k$  and  $r = b - Ax_k;$ 
    end if
    end while;
end procedure.

```

4. Flexible GMRES.

4.1. FGMRES versus GMRES. First, we introduce the right-preconditioned version of FGMRES in Algorithm 4.1.

ALGORITHM 4.1.

```

procedure  $[x] = \text{FGMRES}(A, M_i, b, \text{maxit})$ 
   $x_0 = M_0^{-1}b$ ,  $r_0 = b - Ax_0$  and  $\beta = \|r_0\|$ 
   $v_1 = r_0/\beta$ ;  $k = 0$ ;  $r = r_0$ 
  while  $\|r\| > \varepsilon(\|b\| + \|A\| \|x_k\|)$  and  $k < \text{maxit}$ 
     $k = k + 1$ ;
     $z_k = M_k^{-1}v_k$ ;  $w = Az_k$ ;
    for  $i = 1, \dots, k$  do
       $h_{i,k} = v_i^T w$ ;
       $w = w - h_{i,k}v_i$ ;
    end for;
     $h_{k+1,k} = \|w\|$ ;
     $v_{k+1} = w/h_{k+1,k}$ ;
     $Z_k = [z_1, \dots, z_k]$ ;  $V_k = [v_1, \dots, v_k]$ ;
     $H_k = \{h_{i,j}\}_{1 \leq i \leq j+1; 1 \leq j \leq k}$ ;
     $y_k = \arg \min_y \|\beta e_1 - H_k y\|$ ;
    if  $\|\beta e_1 - H_k y_k\| \leq \varepsilon(\|b\| + \|A\| \|x_k\|)$  do
       $x_k = x_0 + Z_k y_k$  and  $r = b - Ax_k$ ;
    end if
  end while;
end procedure.

```

When $M_k = M$ for all k , Algorithms 3.1 and 4.1 differ only in the computation of x_k where in Algorithm 4.1 the vectors z_k are stored in Z_k . Since Algorithm 4.1 needs both V_k and Z_k , it requires additional storage with respect to Algorithm 3.1. In [9, 16], the convergence of FGMRES is analyzed: if each matrix H_k is full rank, the algorithm converges to the solution.

Both algorithms are based on the relations

$$(4.1) \quad C^{(k)} = [r_0, AZ_k] = V_{k+1} \begin{bmatrix} R_k \\ 0 \end{bmatrix}; \quad V_j^T V_j = I_j \quad \forall j.$$

This corresponds to computing the orthogonal factorization of $C^{(k)}$, where each column $C_i^{(k)}$ is computed after the $(i-1)$ th step of the Gram–Schmidt orthogonalization process. Furthermore, we have

$$(4.2) \quad R_k = \begin{bmatrix} \beta e_1 & H_k \end{bmatrix}$$

and

$$(4.3) \quad AZ_k = V_{k+1} H_k.$$

The vector y_k (in both Algorithms 3.1 and 4.1) is computed by a QR algorithm based on Givens rotations using the upper Hessenberg structure of H_k .

5. Roundoff error analysis. In the following, we will denote by $c_p(n, j)$ functions that depend only on the dimension n and the integer j . We will avoid a precise formulation of these dependences, but we assume that each $c_p(n, j)$ grows moderately with n and j . Finally, if $B \in \mathbb{R}^{n \times m}$, $n \geq m$, is a full rank matrix, we denote by $\kappa(B) = \|B\| \|B^+\|$ its spectral condition number where $B^+ = (B^T B)^{-1} B$.

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

1. Error analysis of the Arnoldi–Krylov process. We will analyze in detail the modified Gram–Schmidt approach, but similar results can be achieved using a Householder-based approach as presented in [1, 6, 19].
2. Error analysis of the Givens process used on the upper Hessenberg matrix H_k in order to reduce it to upper triangular form.
3. Error analysis of the computation of x_k in FGMRES and GMRES.
4. Use of the static pivoting properties and of (1.3) in order to have simpler expressions.

The first two stages of the roundoff error analysis are the same for both FGMRES and GMRES, and the last two stages are specific to each algorithm.

5.1. FGMRES roundoff error analysis. In this section, we state the theorem presenting the main result that covers stages 1–3. The proofs are given in the appendices.

THEOREM 5.1. *We apply Algorithm 4.1 to solve (1.1), using finite-precision arithmetic conforming to the IEEE standard with relative precision ε . Under the four hypotheses*

$$(5.1) \quad \sigma_{\min}(\bar{H}_k) > c_1(k, 1)\varepsilon \|\bar{H}_k\| + \mathcal{O}(\varepsilon^2) \quad \forall k,$$

$$(5.2) \quad |\bar{s}_k| < 1 - \varepsilon \quad \forall k,$$

where \bar{s}_k are the sines computed during the Givens algorithm applied to \bar{H}_k in order to compute \bar{y}_k , and

$$(5.3) \quad 2.12(n + 1)\varepsilon < 0.01 \quad \text{and} \quad 18.53\varepsilon n^{\frac{3}{2}}\kappa(C^{(k)}) < 0.1 \quad \forall k,$$

there exists \hat{k} , $\hat{k} \leq n$, such that, for all $k \geq \hat{k}$, we have

$$(5.4) \quad \|b - A\bar{x}_k\| \leq c_2(n, k)\varepsilon \left(\|b\| + \|A\| \|\bar{x}_0\| + \|A\| \|\bar{Z}_k\| \|\bar{y}_k\| \right) + \mathcal{O}(\varepsilon^2).$$

Moreover, if $M_i = M$, for all i , and under the hypothesis

$$(5.5) \quad \rho = 1.3 \|\hat{W}_k\| + c_3(k, 1)\varepsilon \|M\| \|\bar{Z}_k\| < 1 \quad \forall k < \hat{k},$$

where $\hat{W}_k = [M\bar{z}_1 - \bar{v}_1, \dots, M\bar{z}_k - \bar{v}_k]$, we have

$$(5.6) \quad \|b - A\bar{x}_k\| \leq c_4(n, k)\gamma\varepsilon \left(\|b\| + \|A\| \|\bar{x}_0\| + \|A\| \|\bar{Z}_k\| \left(\|M(\bar{x}_k - \bar{x}_0)\| + \varepsilon \|\bar{x}_0\| \right) \right) + \mathcal{O}(\varepsilon^2),$$

$$\gamma = \frac{1.3}{1 - \rho}.$$

Proof. See Appendix A. □

Remark 1. We point out that if $\bar{x}_0 = 0$ and $M = I$, FGMRES is numerically equivalent to classical GMRES. Moreover, formula (5.6) implies that GMRES is normwise backward stable, and we obtain the result in [14].

5.1.1. Stage 4: Static pivoting case. We remark that the computation of the matrices \hat{L} and \hat{D} by MA57 with static pivoting in floating-point arithmetic gives the following relations:

$$(5.7) \quad \begin{cases} A + \delta A + \tau \Delta = M, \\ \|\delta A\| \leq c_5(n)\varepsilon \|\hat{L}\|\hat{D}\|\hat{L}^T\|, \\ \|\Delta\| \leq 1. \end{cases}$$

Moreover, we take into account that

$$(5.8) \quad \begin{cases} (M + \delta M)\bar{x}_0 = b, & \text{where} \\ \|\delta M\| \leq c_6(n)\varepsilon \|\hat{L}\|\hat{D}\|\hat{L}^T\|. \end{cases}$$

Relations (5.7) and (5.8) follow by the use of standard techniques similar to those for the roundoff error analysis of Gaussian elimination [12, 13, 20]. We point out that, from (5.7), the perturbation δA must have a norm smaller than τ in order not to dominate the global error that would make the diagonal perturbation ineffective. Therefore, it is reasonable to assume that

$$(5.9) \quad \max(c_5(n), c_6(n))\varepsilon \|\hat{L}\|\hat{D}\|\hat{L}^T\| < \tau.$$

THEOREM 5.2. *Under the hypotheses of Theorem 5.1, hypothesis (5.9), and the hypotheses*

$$(5.10) \quad c_7(n, k)\gamma\varepsilon \|A\| \|\bar{Z}_k\| < 1 \quad \forall k < \hat{k}$$

and

$$(5.11) \quad \max\{\|M^{-1}\|, \|\bar{Z}_k\|\} \leq \frac{\tilde{c}}{\tau},$$

where \tilde{c} is a constant, we have

$$(5.12) \quad \|b - A\bar{x}_k\| \leq 2\mu\varepsilon (\|b\| + \|A\| (\|\bar{x}_0\| + \|\bar{x}_k\|)) + \mathcal{O}(\varepsilon^2),$$

$$\mu = \frac{c_7(n, k)\gamma}{1 - c_7(n, k)\varepsilon\gamma\|A\|\|\bar{Z}_k\|}.$$

Proof. See Appendix B. \square

Remark 2. We point out that hypothesis (5.11), as we will see in section 6, is satisfied for a moderate value of \tilde{c} by our static pivoting strategy on all our test problems. However, the upper bound (5.6) is more general and also gives a good estimate of the error when $\tau \ll \sqrt{\varepsilon}$. In all our experiments, we have $\|\bar{Z}_k\| \|M(\bar{x}_k - \bar{x}_0)\| < 10^4$.

5.2. GMRES roundoff error analysis. For the sake of simplicity, we make the assumption that the solution of the linear system

$$Mq = fl(\bar{V}_k \bar{y}_k)$$

is followed by a few steps of iterative refinement in order to guarantee a good local backward stability. Under this assumption, we have the following theorem.

THEOREM 5.3. *If we apply Algorithm 3.1 to solve (1.1), using finite-precision arithmetic conforming to the IEEE standard with relative precision ε and under the hypotheses (5.1), (5.2), and (5.3) of Theorem 5.1, and if*

$$(5.13) \quad c_8(n, 1)\varepsilon \kappa(M) < 1,$$

then there exists \hat{k} , $\hat{k} \leq n$, such that, for all $k \geq \hat{k}$, we have

$$(5.14) \quad \begin{aligned} \|b - A\bar{x}_k\| \leq & c_9(n, k)\chi\varepsilon \left\{ \|b\| + \|A\| \|\bar{x}_0\| + \|AM^{-1}\| \|M\| \|\bar{x}_k - \bar{x}_0\| \right. \\ & + \left[\|A\| \|\bar{Z}_k\| + \|AM^{-1}\| \|\hat{L}\|\hat{D}\|\hat{L}^T\| \right] \\ & \left. \times \left[\|M(\bar{x}_k - \bar{x}_0)\| + n\varepsilon \|M\| (\|\bar{x}_k - \bar{x}_0\| + \|\bar{x}_0\|) \right] \right\} + \mathcal{O}(\varepsilon^2), \end{aligned}$$

$$\chi = \frac{1.3}{1 - 1.3k\varepsilon}.$$

Proof. See Appendix C. \square

Remark 3. Both (5.6) and (5.14) do not depend on the special initial choice $\bar{x}_0 = M^{-1}b$. They will be true for any choice of \bar{x}_0 . In both FGMRES and GMRES, a restart with a new $\bar{x}_0 = \bar{x}_k$ can improve the situation, reducing the upper bounds significantly.

Remark 4. Finally, we point out that when we replace the values of $c_4(n, k)\gamma$ and $c_9(n, k)\chi$ in (5.6) and (5.14) with $\hat{c} = \max(c_4(n, k)\gamma, c_9(n, k)\chi)$, (5.14) becomes an upper bound for (5.6). The numerical experiments in section 6 support this observation.

5.2.1. Stage 4: Static pivoting case. From (5.7) and (5.9), taking into account that $\tau = \varepsilon^\sigma \|A\|$ with $0 < \sigma \leq 1$, we have

$$(5.15) \quad \|M\| \leq \|A\| + \tau + c_{10}(n)\varepsilon \|\hat{L}\|\hat{D}\|\hat{L}^T\| \leq \|A\| + 2\tau \leq 3\|A\|.$$

From (5.14), we have

$$(5.16) \quad \begin{aligned} \|b - A\bar{x}_k\| \leq & c_{11}(n, k)\chi\varepsilon \left[\|b\| + \|A\| \|\bar{x}_0\| + \|AM^{-1}\| \|A\| \|\bar{x}_k - \bar{x}_0\| \right. \\ & + \|A\| \|\bar{Z}_k\| \|M(\bar{x}_k - \bar{x}_0)\| \\ & \left. + \|AM^{-1}\| \|\hat{L}\|\hat{D}\|\hat{L}^T\| \|A\| \|\bar{x}_k - \bar{x}_0\| \right] + \mathcal{O}(\varepsilon^2). \end{aligned}$$

Moreover, we have

$$AM^{-1} = I - \tau\Delta M^{-1} - c_5(n)\varepsilon \|\hat{L}\|\hat{D}\|\hat{L}^T\|$$

and, from (5.11) and by assuming (5.9) is satisfied,

$$\|AM^{-1}\| \leq 1 + \tilde{c}.$$

Thus we have from (5.16)

$$(5.17) \quad \begin{aligned} \|b - A\bar{x}_k\| \leq & c_{12}(n, k)\chi\varepsilon \left[\|b\| + \|A\| \|\bar{x}_0\| + \|A\| \|\bar{Z}_k\| \|M(\bar{x}_k - \bar{x}_0)\| \right. \\ & \left. + \left[1 + \|\hat{L}\|\hat{D}\|\hat{L}^T\| \right] \|A\| \|\bar{x}_k - \bar{x}_0\| \right] + \mathcal{O}(\varepsilon^2), \end{aligned}$$

$$\chi = \frac{1.3}{1 - 1.3k\varepsilon}.$$

Remark 5. Formula (5.17) suggests that we might expect GMRES to have a final residual larger than the final residual in FGMRES. In particular, we could expect that a factor proportional to $\varepsilon \|A\| \|\hat{L}\| \|\hat{D}\| \|\hat{L}^T\| \|\bar{x}_k - \bar{x}_0\|$ will appear.

6. Numerical experiments.

6.1. Test problems. We have run FGMRES and GMRES on several test problems. For the sake of simplicity, we present only the results obtained on particularly tough problems where iterative refinement fails for several different static pivoting strategies and values of τ [8]. We decided to omit the other results either because they are very similar to our results or because iterative refinement produced a scaled residual that is at rounding error level.

The dimension, the number of nonzero entries, and the origin of the three test problems that we use to illustrate the theory presented in section 5 are shown in Table 6.1.

TABLE 6.1
Test problems.

	n	nnz	$\ A\ _\infty$	$\ A\ $	Description
CONT201	80595	239596	8.2	8.1	KKT matrix Convex QP (M2)
CONT300	180895	562496	8.2	8.1	KKT matrix Convex QP (M2)
TUMA1	22967	76199	7.8	4.9	Mixed-hybrid finite-element




FIG. 6.1. $\|L\| \|D\| \|L^T\|$ versus $1/\tau$.

6.2. Numerical results. The smallest example, TUMA1, presents some interesting features and illustrates the behavior of test examples, where iterative refinement converges for the range of τ we are interested in. In Figure 6.1, we summarize the

TABLE 6.2
TUMA1 results.

τ	$\frac{\ b - A\bar{x}_k\ }{\ b\ + \ A\ _\infty \ \bar{x}_k\ }$		$\ M(\bar{x}_k - \bar{x}_0)\ $			$\ L\ D\ L^T\ _\infty$
	GMRES (#It)	FGMRES (#It)	$\ Z_k\ $	FGMRES	GMRES	
1.0e-03	1.0e-14 (26)	7.2e-17 (9)	1.2e+02	3.5e-03	3.5e-03	4.4e+04
1.0e-04	1.8e-16 (6)	3.1e-17 (6)	4.7e+01	4.4e-04	4.4e-04	1.8e+05
1.0e-05	1.3e-16 (5)	1.9e-17 (5)	4.4e+01	4.5e-05	4.5e-05	1.8e+06
1.0e-06	1.3e-16 (4)	1.9e-17 (4)	4.4e+01	4.5e-06	4.5e-06	1.8e+07
1.0e-07	1.2e-16 (3)	2.0e-17 (3)	4.3e+01	4.5e-07	4.5e-07	1.8e+08
1.0e-08	1.3e-16 (3)	1.8e-17 (3)	4.3e+01	4.5e-08	4.5e-08	1.8e+09
1.0e-09	2.8e-15 (31)	1.8e-17 (3)	2.6e+01	4.0e-08	4.0e-08	1.8e+10
1.0e-10	4.2e-13 (31)	1.8e-17 (3)	8.8e+00	4.0e-07	4.0e-07	1.8e+11
1.0e-11	1.0e-10 (31)	6.2e-17 (3)	6.8e+00	4.0e-06	4.0e-06	1.8e+12
1.0e-12	1.0e-08 (31)	2.2e-17 (4)	3.2e+01	4.3e-05	4.3e-05	1.8e+13
1.0e-13	2.4e-07 (31)	1.9e-17 (6)	1.3e+02	3.9e-04	3.9e-04	1.8e+14
1.0e-14	8.6e-06 (31)	2.1e-17 (10)	1.8e+02	4.3e-03	4.3e-03	1.8e+15

TABLE 6.3
CONT201 results.

τ	$\frac{\ b - A\bar{x}_k\ }{\ b\ + \ A\ _\infty \ \bar{x}_k\ }$		$\ M(\bar{x}_k - \bar{x}_0)\ $			$\ L\ D\ L^T\ _\infty$
	GMRES (#It)	FGMRES (#It)	$\ Z_k\ $	FGMRES	GMRES	
1.0e-03	1.8e-05 (31)	9.8e-06 (31)	*	1.5e-04	7.1e-04	8.3e+07
1.0e-04	2.0e-07 (31)	2.0e-07 (31)	*	1.9e-05	1.5e-05	1.8e+08
1.0e-05	1.8e-12 (31)	1.1e-16 (30)	4.1e+05	1.3e-05	5.9e-06	4.4e+09
1.0e-06	1.1e-11 (31)	2.1e-16 (15)	2.7e+06	7.8e-07	7.8e-07	1.8e+10
1.0e-07	4.8e-11 (31)	1.8e-16 (9)	1.4e+08	8.7e-08	8.7e-08	1.9e+12
1.0e-08	2.7e-10 (31)	5.8e-17 (6)	2.1e+07	1.3e-06	1.3e-06	1.8e+13
1.0e-09	1.8e-09 (31)	4.5e-17 (5)	1.1e+07	1.3e-06	1.3e-06	1.5e+13
1.0e-10	3.2e-09 (31)	7.2e-17 (4)	3.4e+05	9.2e-06	9.2e-06	1.5e+14
1.0e-11	2.1e-09 (31)	4.5e-17 (4)	1.9e+03	2.8e-04	2.8e-04	2.6e+15
1.0e-12	4.5e-07 (31)	3.8e-17 (5)	2.0e+02	9.5e-04	9.5e-04	1.6e+16
1.0e-13	1.3e-04 (31)	2.6e-16 (8)	1.6e+02	1.1e-02	1.1e-02	4.1e+17
1.0e-14	2.3e-01 (31)	2.5e-14 (12)	4.3e+02	1.0e-02	1.9e-02	9.2e+18

behavior of $\|L\|D\|L^T\|$ versus $1/\tau$ for all our test examples. We point out that $\|L\|D\|L^T\|$ can grow much more than $1/\tau$ because of the threshold pivoting strategy, and this growth is the reason for the poor behavior of GMRES.

In Tables 6.2, 6.3, and 6.4, we present the parameters and the scaled residual

$$(6.1) \quad \frac{\|b - A\bar{x}_k\|}{\|b\| + \|A\|_\infty \|\bar{x}_k\|}$$

for right-preconditioned GMRES (GMRES column) and flexible GMRES (FGMRES column), for the test problems TUMA1, CONT201, and CONT300, respectively. We also show the values of $\|Z_k\|$ and $\|M(\bar{x}_k - \bar{x}_0)\|$ at convergence, as they are crucial quantities for the residual bounds in Theorems 5.1 and 5.3.

TABLE 6.4
CONT300 results.

τ	$\frac{\ b - A\bar{x}_k\ }{\ b\ + \ A\ \ \bar{x}_k\ }$		$\ Z_k\ $	$\ M(\bar{x}_k - \bar{x}_0)\ $		$\ L\ D\ L^T\ _\infty$
	GMRES (#It)	FGMRES (#It)		FGMRES	GMRES	
1.0e-03	3.6e-05 (31)	2.5e-05 (31)	*	1.3e-04	8.7e-04	2.5e+08
1.0e-04	5.5e-07 (31)	5.5e-07 (31)	*	2.8e-05	6.5e-05	4.3e+09
1.0e-05	8.7e-09 (31)	8.7e-09 (31)	*	6.1e-06	3.7e-06	1.4e+11
1.0e-06	6.9e-11 (31)	1.4e-16 (23)	3.0e+06	9.8e-07	5.7e-07	6.2e+11
1.0e-07	2.1e-10 (31)	8.2e-17 (12)	7.6e+06	2.3e-07	2.3e-07	2.0e+12
1.0e-08	1.4e-08 (31)	1.2e-16 (8)	7.5e+07	1.8e-06	1.8e-06	4.1e+13
1.0e-09	1.6e-05 (31)	8.8e-17 (8)	3.7e+07	2.8e-04	2.8e-04	3.7e+15
1.0e-10	6.8e-07 (31)	4.1e-17 (6)	3.8e+05	3.6e-04	3.6e-04	9.6e+15
1.0e-11	1.6e-06 (31)	8.7e-17 (5)	1.4e+03	5.3e-03	5.3e-03	1.0e+17
1.0e-12	1.1e-06 (31)	2.7e-16 (5)	1.5e+02	1.0e-02	1.0e-02	1.9e+17
1.0e-13	3.4e-03 (31)	9.2e-16 (7)	1.3e+02	1.9e-01	1.9e-01	1.3e+19
1.0e-14	1.4e-01 (31)	1.8e-14 (12)	2.1e+02	4.7e-02	4.7e-02	6.6e+19




FIG. 6.2. GMRES versus FGMRES on CONT201 test example: $\tau = 10^{-6}, 10^{-8}, 10^{-10}$.

In all test examples, the values of $\|\bar{Z}_k\|$ and $\|M(\bar{x}_k - \bar{x}_0)\|$ for FGMRES show that

$$\|\bar{Z}_k\| \|M(\bar{x}_k - \bar{x}_0)\| < 10^4$$

when $\tau = 10^{-9}$ (see Table 6.4). However, for the remaining values of τ the latter product is less than 2×10^2 .

In Figures 6.2 and 6.3, we show the convergence histories of the scaled residual (6.1) for FGMRES and the right-preconditioned version of GMRES for the CONT201 and CONT300 test problems.




FIG. 6.3. *GMRES versus FGMRES on CONT300 test example: $\tau = 10^{-6}, 10^{-8}, 10^{-10}$.*




FIG. 6.4. *Restarted GMRES versus FGMRES on CONT201 test example: $\tau = 10^{-8}$.*

First, we point out that right-preconditioned GMRES is not backward stable. The FGMRES dependence on τ shows that the best trade-off is obtained for the recommended choice of $\tau = 10^{-8}$, where both iterative refinement and right-preconditioned GMRES are not able to obtain a scaled residual close to ε .

Both FGMRES and right-preconditioned GMRES can benefit from a restarting process. In particular, the restarted right-preconditioned GMRES converges with a scaled residual at machine precision for a small value of the restart parameter. In Figures 6.4 and 6.5, we compare the restarted right-preconditioned GMRES with




FIG. 6.5. Restarted GMRES versus FGMRES on CONT300 test example: $\tau = 10^{-8}$.

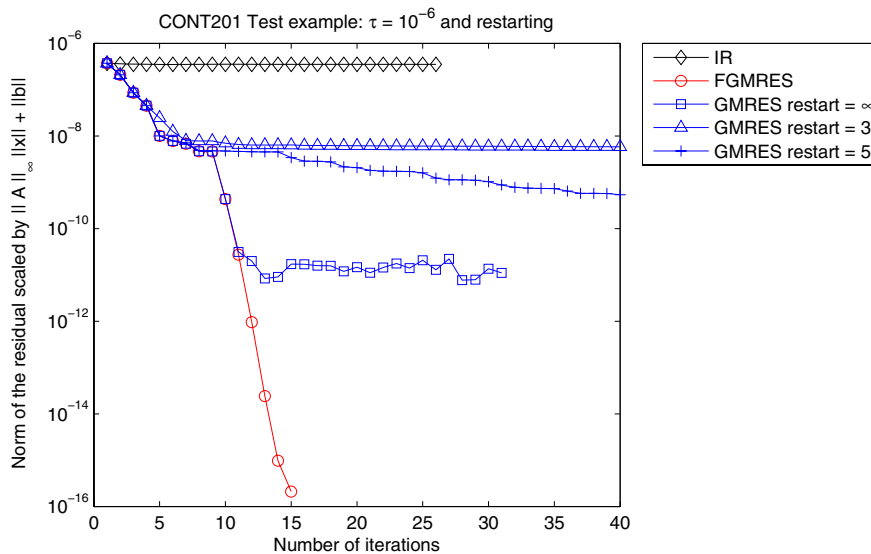


FIG. 6.6. Restarted GMRES on CONT201 test example: $\tau = 10^{-6}$.

FGMRES for $\tau = 10^{-8}$ for the CONT201 and CONT300 test examples, respectively. In the same figures, we show the behavior of classical iterative refinement which is unable to obtain a scaled residual close to rounding. Furthermore, we show in Figure 6.6 that, even with restarting, the preconditioned GMRES approach is not robust.

7. Conclusions. We have shown by experiment and analysis that FGMRES is a powerful method for obtaining the solution of sets of sparse linear equations when

a direct method using static pivoting has been used to factorize the matrix. This factorization is then used as a preconditioner for an iterative method. In particular, we show that GMRES converges when Richardson’s method (iterative refinement) does not, that restarted GMRES performs better, and that FGMRES succeeds when these other methods fail. We also see that the convergence of FGMRES is much less sensitive to the choice of the static pivoting parameter.

Our analysis gives sufficient conditions for convergence, and we observe that FGMRES often converges even when our hypotheses are not satisfied. The main reason for this is that the bound $\|L\|D\|L^T\|$ on the norm of the solution of the preconditioned system is rather crude. Indeed, taking into account that $\|M\| \approx \|A\|$, we see that there is severe cancellation in the product and that small entries in D are usually balanced by large entries in L .

However, we see clearly that the analysis showing that stronger hypotheses are required for the convergence of preconditioned GMRES is reflected in the poorer convergence of this method, even when restarting is used.

A great benefit of our analysis and experiments is that the FGMRES iterative method can be used with confidence to solve systems preconditioned with a static pivoting factorization and that the method is far less sensitive to the choice of static pivoting parameter than either GMRES or iterative refinement.

Appendix A: Proof of Theorem 5.1.

Stage 1. By standard techniques [13], the computed matrix-by-vector products satisfy the relations

$$(A.1) \quad fl(r_0) = r_0 + f_1, \quad \|f_1\| \leq c_{13}(n, 1)\varepsilon (\|b\| + \|A\| \|\bar{x}_0\|) + \mathcal{O}(\varepsilon^2),$$

$$(A.2) \quad fl(A\bar{Z}_k) = A\bar{Z}_k + F_k, \quad \|F_k\| \leq c_{14}(n, k)\varepsilon \|A\| \|\bar{Z}_k\| + \mathcal{O}(\varepsilon^2),$$

where $\bar{x}_0 = fl(x_0)$.

Following [4, 10], the Gram–Schmidt orthogonalization process applied to $fl(C^{(k)})$ computes an upper triangular matrix \bar{R}_k for which there exists an orthonormal matrix \hat{V}_{k+1} that satisfies the relations

$$(A.3) \quad \begin{cases} [fl(r_0); fl(A\bar{Z})] + [f_2; E_k] = \hat{V}_{k+1}\bar{R}_k, & \hat{V}_{k+1}^T \hat{V}_{k+1} = I_{k+1}, \\ \|f_2\| \leq c_{15}(n, 1)\varepsilon \|r_0\| + \mathcal{O}(\varepsilon^2), & \|E_k\| \leq c_{16}(n, k)\varepsilon \|A\bar{Z}_k\| + \mathcal{O}(\varepsilon^2) \end{cases}$$

under the hypothesis (5.3).

By combining (A.1), (A.2), and (A.3), we have

$$(A.4) \quad \begin{cases} [r_0; A\bar{Z}_k] + [f_1 + f_2; F_k + E_k] = \hat{V}_{k+1}\bar{R}_k, \\ \|f_1 + f_2\| \leq c_{17}(n, 1)\varepsilon (\|b\| + \|A\| \|\bar{x}_0\|) + \mathcal{O}(\varepsilon^2), \quad \text{and} \\ \|F_k + E_k\| \leq c_{18}(n, k)\varepsilon (\|A\bar{Z}_k\| + \|A\| \|\bar{Z}_k\|) + \mathcal{O}(\varepsilon^2). \end{cases}$$

Stage 2. The second part of FGMRES computes the vector $\bar{y}_k = fl(y_k)$ by Givens or Householder algorithms [11, 12]. The vector \bar{y}_k satisfies the relations

$$(A.5) \quad \begin{cases} \bar{y}_k = \arg \min_y \|\bar{\beta}e_1 + g^{[k]} - (\bar{H}_k + \Delta\bar{H}_k)y\|, \\ \|\Delta\bar{H}_k\| \leq c_{19}(k, 1)\varepsilon \|\bar{H}_k\| + \mathcal{O}(\varepsilon^2) \text{ and } \|g^{[k]}\| \leq c_{20}(k, 1)\varepsilon \bar{\beta} + \mathcal{O}(\varepsilon^2), \end{cases}$$

where $\bar{\beta} = (\bar{R}_k)_{11} = \|r_0 + f_1 + f_2\|$ and the columns of \bar{H}_k are columns $2, \dots, k + 1$ of

\bar{R}_k . Note that $\bar{\beta}$ is independent of k because at step k we compute only the k th column of \bar{R}_k , leaving the earlier columns unchanged. The computation of \bar{y}_k is performed in two stages, and contributions to the matrix $\Delta\bar{H}_k$ come from two sources.

First, a sequence of Givens (or Householder) rotations $G^{(i)}$ is computed in order to reduce the matrix \bar{H} to the upper triangular form U_k . The floating-point computation of the matrices $G^{(i)}$ gives

$$fl(G^{(i)}) = \bar{G}^{(i)} = \begin{bmatrix} I_{i-1} & & & \\ & \bar{c}_i & -\bar{s}_i & \\ & \bar{s}_i & \bar{c}_i & \\ & & & I_{n-i-1} \end{bmatrix}, \quad i = 1, \dots, k,$$

$$\bar{c}_i = fl\left(\frac{(\bar{H}_k)_{i,i}}{\sqrt{(\bar{H}_k)_{i,i}^2 + (\bar{H}_k)_{i+1,i}^2}}\right) \quad \text{and} \quad \bar{s}_i = fl\left(\frac{(\bar{H}_k)_{i+1,i}}{\sqrt{(\bar{H}_k)_{i,i}^2 + (\bar{H}_k)_{i+1,i}^2}}\right).$$

The $\bar{G}^{(i)}$ matrices are also applied to the vector $\bar{\beta}e_1$, and, from the error analysis presented in [20], the floating-point arithmetic will produce an exact orthogonal matrix $G^{[k]}$ such that

$$fl(\bar{G}^{(k)} \dots \bar{G}^{(1)} \bar{\beta}e_1) = G^{[k]}(\bar{\beta}e_1 + g^{[k]}),$$

$$fl(\bar{G}^{(k)} \dots \bar{G}^{(1)} \bar{H}_k) = G^{[k]}(\bar{H}_k + \Delta\bar{H}_k^{(1)}),$$

$$\|\Delta\bar{H}_k^{(1)}\| \leq c_{21}(k, 1)\varepsilon \|\bar{H}_k\| + \mathcal{O}(\varepsilon^2).$$

Second, the \bar{y}_k vector is computed by solving the upper triangular system. The standard backward substitution algorithm will introduce an additional perturbation $\Delta\bar{H}_k^{(2)}$ of \bar{H}_k but will leave the perturbation $g^{[k]}$ untouched. The perturbation $\Delta\bar{H}_k^{(2)}$ will also have the same upper Hessenberg structure of \bar{H}_k and

$$\|\Delta\bar{H}_k^{(2)}\| \leq c_{22}(k, 1)\varepsilon \|\bar{H}_k\| + \mathcal{O}(\varepsilon^2).$$

This follows from the upper triangular structure of the perturbation to U_k induced by the backward substitution algorithm and from the structure and orthogonality of $G^{[k]}$. Finally, we point out that in the relations (A.5) we have (see (5.1))

$$\begin{cases} \Delta\bar{H}_k = \Delta\bar{H}_k^{(1)} + \Delta\bar{H}_k^{(2)} & \text{and} \\ c_1(k, 1) = c_{21}(k, 1) + c_{22}(k, 1), \end{cases}$$

where $c_1(k, 1)$ is the constant of (5.1). Moreover, because of the special structure of $\bar{\beta}e_1$ and the orthogonality of $G^{[k]}$ we have

$$g_j^{[k]} = 0, \quad j = k + 1, \dots, n,$$

and, denoting by $\bar{h}^{[k]} = fl(\bar{G}^{(k)} \dots \bar{G}^{(1)} \bar{\beta}e_1)$, we have

$$(A.6) \quad \begin{aligned} \alpha_k &= \|G^{[k]}(\bar{\beta}e_1 + g^{[k]} - (\bar{H}_k + \Delta\bar{H}_k)\bar{y}_k)\| \\ &= \|(G^{[k]}(\bar{\beta}e_1 + g^{[k]}))_{k+1}\| = |\bar{h}_{k+1}^{[k]}|. \end{aligned}$$

A direct analysis of $\bar{h}^{[k]}$ shows that

$$\begin{cases} \bar{h}_1^{[1]} &= \bar{\beta}\bar{c}_1(1 + \mu_1), & |\mu_1| \leq \varepsilon, \\ \bar{h}_2^{[1]} &= \bar{\beta}\bar{s}_1(1 + \rho_1), & |\rho_1| \leq \varepsilon, \\ \bar{h}_j^{[k]} &= \bar{h}_j^{[k-1]}, & j = 1, \dots, k-1, \quad k \geq 2, \\ \bar{h}_k^{[k]} &= \bar{h}_k^{[k-1]}\bar{c}_k(1 + \mu_k), & |\mu_k| \leq \varepsilon, \\ \bar{h}_{k+1}^{[k]} &= \bar{h}_k^{[k-1]}\bar{s}_k(1 + \rho_k), & |\rho_k| \leq \varepsilon. \end{cases}$$

Therefore, from using (A.3), (A.4), (A.5), (A.6), and the orthogonality of

$$\tilde{V}_{k+1} = \hat{V}_{k+1}G^{[k]T},$$

we have

$$\begin{aligned} \alpha_k &= \|\tilde{V}_{k+1}G^{[k]}(\bar{\beta}e_1 + g^{[k]} - (\bar{H}_k + \Delta\bar{H}_k)\bar{y}_k)\| \\ &= \|r_0 + f_1 + f_2 + \hat{V}_{k+1}g^{[k]} - A(\bar{Z}_k + A^{-1}(F_k + E_k + \hat{V}_{k+1}\Delta\bar{H}_k))\bar{y}_k\| \\ &= \|r_0 + \delta r_0 - A(\bar{Z}_k + \hat{Z}_k)\bar{y}_k\|, \end{aligned}$$

where $\delta r_0 = f_1 + f_2 + f_3$ and $\hat{Z}_k = A^{-1}(F_k + E_k + \hat{V}_k\Delta\bar{H}_k)$. Then

$$(A.7) \quad \alpha_k = \alpha_{k-1}|\bar{s}_k|(1 + \rho_k), \quad |\rho_k| \leq \varepsilon.$$

Thus, under the hypotheses (5.2), (5.1), and (5.3) and from (A.5), the matrices $(\bar{Z}_k + \hat{Z}_k)$ have full rank for all k ; i.e.,

$$\bar{H}_k + \Delta\bar{H}_k$$

is full rank for all k , and the values of α_k converge monotonically to zero for a finite value of $k = \hat{k}$. In the worst case this will happen for $\hat{k} = n$.

Stage 3. The last part of FGMRES is the computation of $\bar{x}_k = fl(x_0 + \bar{Z}\bar{y}_k)$. The value \bar{x}_k satisfies the relations

$$(A.8) \quad \begin{cases} \bar{x}_k = \bar{x}_0 + \bar{Z}_k\bar{y}_k + \delta x_k, \\ \|\delta x_k\| \leq c_3(k, 1)\varepsilon\|\bar{Z}_k\|\|\bar{y}_k\| + \varepsilon\|\bar{x}_0\| + \mathcal{O}(\varepsilon^2). \end{cases}$$

Therefore, we have

$$(A.9) \quad \alpha_k = \|r_0 + \delta r_0 + A\delta x_k - A\delta x_k - A\bar{Z}_k\bar{y}_k - A\hat{Z}_k\bar{y}_k\|.$$

From (A.8), we have

$$(A.10) \quad \begin{cases} \alpha_k = \|b - A\bar{x}_k + w\|, \\ w = \delta r_0 + A\delta x_k - A\hat{Z}_k\bar{y}_k, \end{cases}$$

and then

$$(A.11) \quad \|b - A\bar{x}_k\| \leq \|w\| + \alpha_k.$$

From (A.4), (A.5), (A.8), (A.10), and (A.11), we have

$$(A.12) \quad \begin{cases} \|\delta r_0\| \leq c_{24}(n, 1)\varepsilon (\|b\| + \|A\| \|\bar{x}_0\|) + \mathcal{O}(\varepsilon^2), \\ \|A\delta x_k\| \leq c_{25}(n, 1)\varepsilon \|A\| \left[\|\bar{Z}_k\| \|\bar{y}_k\| + \|\bar{x}_0\| \right] + \mathcal{O}(\varepsilon^2) \end{cases}$$

and, finally,

$$(A.13) \quad \|w\| \leq c_{26}(n, k)\varepsilon \left(\|b\| + \|A\| \|\bar{x}_0\| + \|\bar{H}_k\| \|\bar{y}_k\| + \|A\| \|\bar{Z}_k\| \|\bar{y}_k\| \right) + \mathcal{O}(\varepsilon^2).$$

Therefore, under hypothesis (5.1) and from (A.11), we have

$$(A.14) \quad \|b - A\bar{x}_k\| \leq \alpha_k + c_2(n, k)\varepsilon \left(\|b\| + \|A\| \|\bar{x}_0\| + \|\bar{H}_k\| \|\bar{y}_k\| + \|A\| \|\bar{Z}_k\| \|\bar{y}_k\| \right) + \mathcal{O}(\varepsilon^2)$$

and, then, taking into account that, from (A.4),

$$(A.15) \quad \|\bar{H}_k\| \leq \|A\bar{Z}_k\| + \mathcal{O}(\varepsilon) \quad \forall k < \hat{k},$$

we have (5.4).

From the analysis of [4, 10], the Gram–Schmidt orthogonalization process applied to $fl(C^{(k)})$ computes matrices $\bar{V}_k = fl(V_k)$ that lose orthogonality quickly. Nevertheless, under the mild hypothesis (5.3) on the matrix $C^{(k)}$ the spectral condition number of each \bar{V}_k is close to one and

$$\begin{aligned} 1 &\leq \|\bar{V}_k\| \leq k, \\ \kappa(\bar{V}_k) &\leq 1.3 \end{aligned}$$

so that

$$\|\bar{V}_k^+\| \leq 1.3/\|\bar{V}_k\| \leq 1.3,$$

where $\bar{V}_k^+ = (\bar{V}^T \bar{V}_k)^{-1} \bar{V}_k^T$ (see [10, Thm. 3.1] and [14]).

Taking into account the previous relations, it is possible to express \bar{y}_k as a function of $\bar{x}_k - x_0$. For each column \bar{z}_i of \bar{Z}_k a matrix \mathcal{G}_i exists and satisfies

$$(M + \mathcal{G}_i)\bar{z}_i = \bar{v}_i, \quad \|\mathcal{G}_i\| \leq c_{27}(n, 1)\varepsilon \|\hat{L}\| |\hat{D}| \|\hat{L}^T\| + \mathcal{O}(\varepsilon^2),$$

where \bar{v}_i is the i th column of \bar{V}_k . Then, we have the relation

$$(A.16) \quad M\bar{Z}_k = \bar{V}_k + \hat{W}_k + \mathcal{O}(\varepsilon^2).$$

From the relations (A.8) and (A.16), it follows that

$$\begin{aligned} \bar{x}_k - \bar{x}_0 - \delta x_k &= \bar{Z}_k \bar{y}_k, \\ \bar{V}_k^+ M(\bar{x}_k - \bar{x}_0 - \delta x_k) &= \bar{V}_k^+ M \bar{Z}_k \bar{y}_k, \\ \bar{V}_k^+ (M(\bar{x}_k - \bar{x}_0) - M\delta x_k) &= (\bar{V}_k^+ \bar{V}_k \bar{y}_k + \bar{V}_k^+ \hat{W}_k \bar{y}_k) + \mathcal{O}(\varepsilon^2), \\ \bar{V}_k^+ (M(\bar{x}_k - \bar{x}_0) - M\delta x_k) - \bar{V}_k^+ \hat{W}_k \bar{y}_k + \mathcal{O}(\varepsilon^2) &= \bar{y}_k. \end{aligned}$$

Then, we have

$$\|\bar{y}_k\| \leq 1.3\|M(\bar{x}_k - \bar{x}_0)\| + 1.3\|\hat{W}_k\|\|\bar{y}_k\| + \|M\|\|\delta x_k\| + \mathcal{O}(\varepsilon^2)$$

and

$$\begin{aligned} \|\bar{y}_k\| &\leq 1.3\|M(\bar{x}_k - \bar{x}_0)\| + 1.3\|\hat{W}_k\|\|\bar{y}_k\| \\ &\quad + c_3(k, 1)\varepsilon\|M\|\|\bar{Z}_k\|\|\bar{y}_k\| + \varepsilon\|\bar{x}_0\| + \mathcal{O}(\varepsilon^2). \end{aligned}$$

From hypothesis (5.5), we can state the inequality

$$(A.17) \quad \|\bar{y}_k\| \leq \gamma\|M(\bar{x}_k - \bar{x}_0)\| + \gamma\varepsilon\|\bar{x}_0\| + \mathcal{O}(\varepsilon^2), \quad \text{where } \gamma = \frac{1.3}{1 - \rho}.$$

Finally, by substituting (A.17) into (A.14) and taking into account (A.15), we have the following upper bound for the residual:

$$\begin{aligned} \|b - A\bar{x}_k\| &\leq \alpha_k + c_4(n, k)\gamma\varepsilon \left[\|b\| + \|A\|\|\bar{x}_0\| \right. \\ &\quad \left. + \left(\|A\bar{Z}_k\| + \|A\|\|\bar{Z}_k\| \right) \left(\|M(\bar{x}_k - \bar{x}_0)\| + \varepsilon\|\bar{x}_0\| \right) \right] + \mathcal{O}(\varepsilon^2). \end{aligned}$$

Furthermore, hypothesis (5.1) implies that there exists a value \hat{k} (possibly equal to n in the worst case) for which $\alpha_k = 0$. Then, we can conclude that for $k \geq \hat{k}$ we have (5.6).

Appendix B: Proof of Theorem 5.2. We can bound $\|M(\bar{x}_k - \bar{x}_0)\|$ using (5.7) and (5.8). Under hypothesis (5.9), we have

$$(B.1) \quad \begin{aligned} \|M(\bar{x}_k - \bar{x}_0)\| &\leq \|A\bar{x}_k - b\| + \tau\|\bar{x}_k\| + \|\delta A\|\|\bar{x}_k\| + \|\delta M\|\|\bar{x}_0\| \\ &\leq \|A\bar{x}_k - b\| + 2\tau(\|\bar{x}_k\| + \|\bar{x}_0\|). \end{aligned}$$

From the relations (B.1) and (5.6), we have

$$\begin{aligned} \|b - A\bar{x}_k\| &\leq c_4(n, k)\gamma\varepsilon (\|b\| + \|A\|\|\bar{x}_0\| \\ &\quad + \|A\|\|\bar{Z}_k\| (\|A\bar{x}_k - b\| + \tau(\|\bar{x}_k\| + \|\bar{x}_0\|))) + \mathcal{O}(\varepsilon^2). \end{aligned}$$

Under the hypotheses (5.9) and (5.10), we have

$$(B.2) \quad \begin{cases} \|b - A\bar{x}_k\| \leq 2\mu\varepsilon \left(\|b\| + \|A\|\|\bar{x}_0\| \right. \\ \quad \left. + \tau\|A\|\|\bar{Z}_k\| (\|\bar{x}_k\| + \|\bar{x}_0\|) \right) + \mathcal{O}(\varepsilon^2), \\ \mu = \frac{c_7(n, k)\gamma}{1 - c_7(n, k)\varepsilon\gamma\|A\|\|\bar{Z}_k\|}. \end{cases}$$

Under this final hypothesis (5.11), we obtain (5.12), i.e., the normwise backward stability for FGMRES using as preconditioner the M computed by MA57 with static pivoting.

Appendix C: Proof of Theorem 5.3. The first two stages of GMRES roundoff error analysis are the same as the first two stages in the proof of Theorem 5.1. Under the assumption of performing a few steps of iterative refinement in solving the final linear system

$$Mq = fl(\bar{V}_k \bar{y}_k),$$

we have the following relations:

$$(C.1) \quad \begin{cases} (M + \delta M)(I + \Gamma)^{-1} (\bar{x}_k - (I + \Gamma) \bar{x}_0) = \bar{V}_k \bar{y}_k + \tilde{f}, \\ |\Gamma| \leq \varepsilon I, \\ \|\tilde{f}\| \leq k\varepsilon \|\bar{V}_k\| \|\bar{y}_k\| + \mathcal{O}(\varepsilon^2) \leq 1.3k\varepsilon \|\bar{y}_k\| + \mathcal{O}(\varepsilon^2), \\ \|\delta M\| \leq n\varepsilon \|M\| + \mathcal{O}(\varepsilon^2). \end{cases}$$

Taking into account that

$$(M + \delta M)(I + \Gamma)^{-1} = M + (\delta M - M\Gamma)(I + \Gamma)^{-1} = M + \widetilde{\delta M}$$

with

$$\|\widetilde{\delta M}\| \leq (n + 1) \frac{\varepsilon}{1 - \varepsilon} \|M\| = c_{28}(n, 1)\varepsilon \|M\|,$$

we have, from (C.1), that

$$(C.2) \quad (M + \widetilde{\delta M})(\bar{x}_k - \bar{x}_0) = \bar{V}_k \bar{y}_k + \tilde{f} + (M + \widetilde{\delta M})\Gamma \bar{x}_0.$$

Therefore, we have from (5.13) and (C.2) that (A.8) can be replaced by

$$(C.3) \quad \bar{x}_k = \bar{x}_0 + (M + \widetilde{\delta M})^{-1}(\bar{V}_k \bar{y}_k + \tilde{f}) + \Gamma \bar{x}_0$$

$$(C.4) \quad = \bar{x}_0 + M^{-1} \bar{V}_k \bar{y}_k + \delta \bar{x}_k$$

with

$$\begin{aligned} \delta \bar{x}_k &= \Gamma \bar{x}_0 + M^{-1} \tilde{f} - M^{-1} \widetilde{\delta M} M^{-1} \bar{V}_k \bar{y}_k + \mathcal{O}(\varepsilon^2) \\ &= \Gamma \bar{x}_0 + M^{-1} \tilde{f} - M^{-1} \widetilde{\delta M} (\bar{x}_k - \bar{x}_0 - \delta \bar{x}_k) + \mathcal{O}(\varepsilon^2) \end{aligned}$$

and

$$(C.5) \quad \begin{aligned} \delta \bar{x}_k &= (I - M^{-1} \widetilde{\delta M})^{-1} (\Gamma \bar{x}_0 + M^{-1} \tilde{f} - M^{-1} \widetilde{\delta M} (\bar{x}_k - \bar{x}_0)) + \mathcal{O}(\varepsilon^2) \\ &= \Gamma \bar{x}_0 + M^{-1} \tilde{f} - M^{-1} \widetilde{\delta M} (\bar{x}_k - \bar{x}_0) + \mathcal{O}(\varepsilon^2). \end{aligned}$$

Under hypotheses (5.3) and (5.13), we obtain from (C.3), (C.4), (C.5), and (A.16)

$$(C.6) \quad \begin{cases} \alpha_k = \|b - A\bar{x}_k + w_g\|, & \text{where} \\ w_g = \delta r_0 + A\delta x_k - A\hat{Z}_k \bar{y}_k + A(M^{-1} \bar{V}_k \bar{y}_k - \bar{Z}_k \bar{y}_k) \\ = \delta r_0 + A\delta x_k - A\hat{Z}_k \bar{y}_k + AM^{-1} \hat{W}_k \bar{y}_k + \mathcal{O}(\varepsilon^2), \end{cases}$$

and then

$$(C.7) \quad \|b - A\bar{x}_k\| \leq \alpha_k + \|w_g\|.$$

From (A.4), (A.5), (A.7), (C.3), (C.5), (C.6), and (C.7), we have

$$(C.8) \quad \begin{cases} \|\delta r_0\| \leq c_{29}(n, 1)\varepsilon (\|b\| + \|A\| \|\bar{x}_0\|) + \mathcal{O}(\varepsilon^2) \\ \|A\delta x_k\| \leq c_{30}(n, 1)\varepsilon \left[\|AM^{-1}\| \left(\|M\| \|\bar{x}_k - \bar{x}_0\| + \|\bar{y}_k\| \right) \right. \\ \left. + \|A\| \|\bar{x}_0\| \right] + \mathcal{O}(\varepsilon^2) \end{cases}$$

and, finally,

$$(C.9) \quad \|w_g\| \leq c_{31}(n, k)\varepsilon \left[\|b\| + \|A\| \|\bar{x}_0\| + \|\bar{H}_k\| \|\bar{y}_k\| + \|A\| \|\bar{Z}_k\| \|\bar{y}_k\| \right. \\ \left. + \|AM^{-1}\| \left(\|M\| \|\bar{x}_k - \bar{x}_0\| + \|\bar{y}_k\| + \|\hat{L}\| |\hat{D}| |\hat{L}^T| \|\bar{y}_k\| \right) \right] + \mathcal{O}(\varepsilon^2).$$

From (C.2), we can compute an upper bound for $\|\bar{y}_k\|$ by multiplying the equation by \bar{V}_k^+ :

$$(C.10) \quad \bar{y}_k = \bar{V}_k^+ (M + \delta\bar{M})(\bar{x}_k - \bar{x}_0) - \bar{V}_k^+ (\tilde{f} + (M + \delta\bar{M}) \Gamma \bar{x}_0)$$

$$(C.11) \quad = \bar{V}_k^+ M(\bar{x}_k - \bar{x}_0) - \bar{V}_k^+ (\tilde{f} - \delta\bar{M}(\bar{x}_k - \bar{x}_0) + (M + \delta\bar{M}) \Gamma \bar{x}_0),$$

and then

$$(C.12) \quad \|\bar{y}_k\| \leq 1.3\|M(\bar{x}_k - \bar{x}_0)\| + 1.3k\varepsilon \|\bar{y}_k\| \\ + 1.3n\varepsilon \|M\| (\|\bar{x}_k - \bar{x}_0\| + \|\bar{x}_0\|) + \mathcal{O}(\varepsilon^2).$$

From hypothesis (5.3), we can assume that $1.3k\varepsilon < 1$, and thus we have

$$(C.13) \quad \begin{cases} \|\bar{y}_k\| \leq \chi \|M(\bar{x}_k - \bar{x}_0)\| + \chi n\varepsilon \|M\| (\|\bar{x}_k - \bar{x}_0\| + \|\bar{x}_0\|) + \mathcal{O}(\varepsilon^2), \\ \chi = \frac{1.3}{1 - 1.3k\varepsilon}. \end{cases}$$

Furthermore, as in the proof of Theorem 5.1, hypothesis (5.1) implies that there exists a value \hat{k} (possibly equal to n in the worst case) for which $\alpha_k = 0$. Then, under hypotheses (5.1), (5.3) and from (C.13), (C.7), and (C.9), we can conclude that for $k \geq \hat{k}$ we have (5.14).

Acknowledgment. We would like to thank the anonymous referees for their helpful comments that have hopefully made the paper easier to read.

REFERENCES

- [1] M. ARIOLI AND C. FASSINO, *Roundoff error analysis of algorithms based on Krylov subspace methods*, BIT, 36 (1996), pp. 189–206.
- [2] M. ARIOLI, V. PTÁK, AND Z. STRAKOŠ, *Krylov sequences of maximal length and convergence of GMRES*, BIT, 38 (1998), pp. 1–9.
- [3] Z. BAI, J. DEMMEL, J. DONGARRA, A. RUHE, AND H. VAN DER VORST, *Templates for the Solution of Algebraic Eigenvalue Problems: A Practical Guide*, SIAM, Philadelphia, 2000.

- [4] Å. BJÖRCK AND C. C. PAIGE, *Loss and recapture of orthogonality in the modified Gram–Schmidt algorithm*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 176–190.
- [5] J. DEMMEL, Y. HIDA, W. KAHAN, X. S. LI, S. MUKHERJEEK, AND E. J. RIEDY, *Error bounds from extra precise iterative refinement*, ACM Trans. Math. Software, 32 (2006), pp. 325–351.
- [6] J. DRKOSOVA, A. GEENBAUM, M. ROZLOŽNÍK, AND Z. STRAKOŠ, *Numerical stability of GMRES*, BIT, 35 (1995), pp. 309–330.
- [7] I. S. DUFF, *MA57—A code for the solution of sparse symmetric indefinite systems*, ACM Trans. Math. Software, 30 (2004), pp. 118–144.
- [8] I. S. DUFF AND S. PRALET, *Towards a stable static pivoting strategy for the sequential and parallel solution of sparse symmetric indefinite systems*, SIAM J. Matrix Anal. Appl., to appear.
- [9] L. GIRAUD, S. GRATTON, AND J. LANGOU, *A Note on Relaxed and Flexible GMRES*, Technical report TR/PA/04/41, CERFACS, Toulouse, France, 2004.
- [10] L. GIRAUD AND J. LANGOU, *When modified Gram–Schmidt generates a well-conditioned set of vectors*, IMA J. Numer. Anal., 22 (2002), pp. 521–528.
- [11] G. H. GOLUB, *Numerical methods for solving linear least squares problems*, Numer. Math., 7 (1965), pp. 206–216.
- [12] G. H. GOLUB AND C. F. V. LOAN, *Matrix Computations*, 2nd ed., The Johns Hopkins University Press, Baltimore, MD, 1989.
- [13] N. J. HIGHAM, *Accuracy and Stability of Numerical Algorithms*, 2nd ed., SIAM, Philadelphia, 2002.
- [14] C. C. PAIGE, M. ROZLOŽNÍK, AND Z. STRAKOŠ, *Modified Gram–Schmidt (MGS), least squares, and backward stability of MGS-GMRES*, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 264–284.
- [15] Y. SAAD, *A flexible inner-outer preconditioned GMRES algorithm*, SIAM J. Sci. Comput., 14 (1993), pp. 461–469.
- [16] Y. SAAD, *Iterative Methods for Sparse Linear Systems*, 2nd ed., SIAM, Philadelphia, 2003.
- [17] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems*, SIAM J. Sci. Stat. Comput., 7 (1986), pp. 856–869.
- [18] R. D. SKEEL, *Iterative refinement implies numerical stability for Gaussian elimination*, Math. Comp., 35 (1980), pp. 817–832.
- [19] H. F. WALKER, *Implementation of the GMRES method using Householder transformations*, SIAM J. Sci. Stat. Comput., 9 (1988), pp. 152–163.
- [20] J. H. WILKINSON, *The Algebraic Eigenvalue Problem*, Oxford University Press, New York, 1965.