An optimal Q-OR Krylov subspace method for solving linear systems

Gérard MEURANT

October 2016

1. Introduction
2. Q-OR and Q-MR methods
3. Properties of Q-OR methods
4. Construction of "good" bases for Q-OR
5. Avoiding the use of U
6. Properties of the optimal basis
7. The algorithm
8. Numerical experiments

This talk could have been titled:

Yet another Krylov method equivalent to GMRES

Many Krylov methods have been proposed over the years for solving linear systems $Ax = b$

Many of them can be classified as quasi-orthogonal (Q-OR) or quasi-minimum residual (Q-MR)

Q-OR: FOM, BiCG, Hessenberg, . . .

Q-MR: GMRES, QMR, CMRH, . . .

Whatever their definition, these methods share many fundamental properties

See M. Eiermann and O.G. Ernst, *Geometric aspects in the theory of Krylov subspace methods*, Acta Numerica, v 10 n 10 (2001), pp. 251–312

They differ by the basis of the Krylov space that is constructed:

- orthogonal for FOM/GMRES,
- bi-orthogonal for BiCG/QMR,
- based on an LU factorization for Hessenberg/CMRH

What do we know about GMRES?

Let

$$K = \begin{pmatrix} b & Ab & A^2 b & \cdots & A^{n-1} b \end{pmatrix}$$

be the Krylov matrix that we assume of full rank. Then

$$K = VU$$

with V orthogonal (or unitary) and U upper triangular with positive real diagonal entries

The matrix $H = V^* AV$ is upper Hessenberg

We have

$$H = UCU^{-1}$$

where C is the companion matrix for the eigenvalues of A

Let x_k^G (resp. x_k^F) be the iterates for GMRES (resp. FOM) and the residual vectors $r_k^G = b - Ax_k^G$, $r_k^F = b - Ax_k^F$

We assume $x_0 = 0$ and $\|b\| = 1$

We know that

- every decreasing residual norm convergence curve is possible for GMRES

- $|(U^{-1})_{1,k}| = 1/\|r_{k-1}^F\|$
- $\|r_k^G\|^2 = 1/(\mathcal{M}_{k+1}^{-1})_{1,1}$ with $\mathcal{M}_{k+1} = U_{k+1}^* U_{k+1} = K_{k+1}^* K_{k+1}$

- one can construct matrices A with a prescribed spectrum as well as prescribed Ritz values and right-hand sides b such that GMRES (or FOM) yields a prescribed decreasing residual norm convergence curve

- we have two parametrizations of this class of matrices

For all these properties see:

A. Greenbaum and Z. Strakoš, *Matrices that generate the same Krylov residual spaces*, in Recent advances in iterative methods, G.H. Golub, A. Greenbaum and M. Luskin, eds., Springer, (1994), pp. 95–118

A. Greenbaum, V. Pták and Z. Strakoš, *Any nonincreasing convergence curve is possible for GMRES*, SIAM J. Matrix Anal. Appl., v 17 (1996), pp. 465–469

M. Arioli, V. Pták and Z. Strakoš, *Krylov sequences of maximal length and convergence of GMRES*, BIT Numerical Mathematics, v 38 n 4 (1998), pp. 636–643

J. Duintjer Tebbens and G. Meurant, *Any Ritz value behavior is possible for Arnoldi and for GMRES*, SIAM J. Matrix Anal. Appl., v 33 n 3 (2012), pp. 958–978

G. Meurant and J. Duintjer Tebbens, *The role eigenvalues play in forming GMRES residual norms with non-normal matrices*, Numerical Algorithms, v 68, n 1 (2015), pp. 143–165

Q-OR and Q-MR methods

We assume that we have a basis V of the Krylov space (with columns of unit norm) such that $K = VU$ with V nonsingular and U upper triangular and $v_1 = b$

We define $H = UCU^{-1}$. As a consequence $AV = VH$. The iterates are

$$x_k = V_k y^{(k)}$$

where V_k is the matrix of the k first columns of V. The residual r_k is

$$V_k e_1 - AV_k y^{(k)} = V_k(e_1 - H_k y^{(k)}) - h_{k+1,k} y_k^{(k)} v_{k+1} = V_{k+1}(e_1 - \underline{H}_k y^{(k)})$$

The Q-OR method is defined (provided that H_k is nonsingular) by

$$H_k y^{(k)} = e_1$$

This annihilates the first term in the residual

In the Q-MR method $y^{(k)}$ is computed as the solution of the least squares problem

$$\min_y \|e_1 - \underline{H}_k y\|$$

where $\underline{H}_k$ is $(k+1) \times k$. The vector $z_k^M = e_1 - \underline{H}_k y^{(k)}$ is referred as the quasi-residual

The residual vector is $r_k^M = V_{k+1} z_k^M$

Properties of Q-OR and Q-MR methods

We can show by induction that

$$|(U^{-1})_{1,k}| = \frac{1}{\|r_{k-1}^O\|}$$

The inverses of the Q-OR residual norms can be read from the first row of the inverse of U

For *any* Q-OR method we have the same property as for FOM

For these properties and more see:

G. Meurant and J. Duintjer Tebbens, *On the convergence of Q-OR and Q-MR Krylov methods for solving nonsymmetric linear systems*, BIT Numerical Mathematics, v 56 n 1 (2016), pp. 77-97

Construction of “good” bases

We would like to find bases which lead to a “good” convergence of the Q-OR method

The matrix V of the basis is related to the Krylov matrix K by $K = VU$ with U upper triangular

The entries of the first row of U^{-1} are the inverses of the Q-OR residual norms (up to the sign)

Constructing a “good” basis may seem easy since one can think that we can just construct any upper triangular matrix U^{-1} with entries of large modulus on the first row

But, it is not so since the columns of V have to be of unit norm

We can try directly computing U^{-1} from $V = KU^{-1}$

In this way we obtain the vectors v_j straightforwardly, but, again, the columns of V have to be of unit norm

Let $\nu_{i,j}$ be the entries of U^{-1} and

$$v_k = \nu_{1,k} b + \nu_{2,k} A^2 b + \cdots + \nu_{k,k} A^{k-1} b$$

We would like to have $\|v_k\| = 1$ and $|\nu_{1,k}|$ as large as possible

Can we solve this problem?

Let $\tilde{\nu}$ be the vector of the components $\nu_{i,k}, i = 1, \ldots, k$. Then $v_k = K\tilde{\nu}$

We want $\|K_k\tilde{\nu}\| = 1$. This corresponds to

$$\tilde{\nu}^T K_k^T K_k \tilde{\nu} = \tilde{\nu}^T \mathcal{M}_k \tilde{\nu} = 1$$

This is the equation of an (hyper) ellipsoid in $\mathbb{R}^k$ centered at the origin

We have to find a point on the surface of this ellipsoid with a maximum of the absolute value of the first coordinate

The solution is obtained by writing the equation of a tangent hyperplane and asking that it is orthogonal to the first axis

One can show that a solution is $\nu_{1,k} = \sqrt{(\mathcal{M}_k^{-1})_{1,1}}$ and the other components are obtained by solving a linear system of order $k-1$ whose matrix and right-hand side are $\mathcal{M}_{2:k,2:k}$ and $-\nu_{1,k}\mathcal{M}_{2:k,1}$

This yields U^{-1}. If we apply Q-OR with the basis $V = KU^{-1}$ we obtain residual vectors whose norms are

$$\|r_k^O\|^2 = \frac{1}{(\mathcal{M}_{k+1}^{-1})_{1,1}}$$

These values are those that are obtained from GMRES
Therefore, they are the best ones that we can get with the given Krylov subspace. In a sense we have an optimal Q-OR method

Avoiding the use of U

The previous construction is not practical because

1) we do not want to compute $\mathcal{M}_k$ and $\mathcal{M}_k^{-1}$

2) in many cases the matrix U is almost singular and must be (numerically) avoided

Instead we would like to directly construct H column by column. We have

$$H_j = U_j E_j U_j^{-1} + \begin{pmatrix} 0 & \cdots & 0 & \frac{1}{u_{j,j}} U_{1:j,j+1} \end{pmatrix}$$

E_j down-shift matrix

It yields

$$\sum_{j=1}^{k+1} \nu_{1,j} h_{j,k} = 0 \;\Rightarrow\; \nu_{1,k+1} = -\frac{1}{h_{k+1,k}} \sum_{j=1}^{k} \nu_{1,j} h_{j,k}$$

At step k we have already computed $\nu_{1,j}$, $j = 1, \ldots, k$ and we would like to choose $h_{j,k}$, $j = 1, \ldots, k+1$ to maximize the absolute value of $\nu_{1,k+1}$

But $h_{k+1,k}$ has to be chosen to obtain a vector v_{k+1} of unit norm

Let

$$\tilde{v} = Av_k - \sum_{j=1}^{k} h_{j,k} v_j$$

the next basis vector is $v_{k+1} = \tilde{v}/h_{k+1,k}$ with $h_{k+1,k} = \|\tilde{v}\|$

$$|\nu_{1,k+1}| = \frac{|\nu^T y|}{\|d - By\|}$$

with

$$d = Av_k, \quad B = V_k = \begin{pmatrix} v_1 & \cdots & v_k \end{pmatrix}, \quad y = \begin{pmatrix} h_{1,k} & \cdots & h_{k,k} \end{pmatrix}^T$$

$$\nu = \begin{pmatrix} \nu_{1,1} & \cdots & \nu_{1,k} \end{pmatrix}$$

We need to minimize $1/|\nu_{1,k+1}|^2$

We would like to solve

$$\gamma_{opt} = \min_{y \in \mathbb{R}^k, \nu^T y \neq 0} \frac{\|d - By\|^2}{(\nu^T y)^2}$$

The minimum is given by

$$\gamma_{opt} = \frac{\alpha}{\alpha \nu^T (B^T B)^{-1} \nu + \omega^2}$$

with $\alpha = d^T d - d^T B (B^T B)^{-1} B^T d$ and $\omega = d^T B (B^T B)^{-1} \nu$

Moreover, if $\omega \neq 0$, a solution y_{opt} of the minimization problem is given by

$$\begin{aligned} y_{opt} &= (B^T B)^{-1} B^T d + \frac{\alpha}{\omega} (B^T B)^{-1} \nu \\ &= s + \frac{\alpha}{\omega} p \end{aligned}$$

This is obtained by finding the largest possible value of γ such that

$$\frac{1}{|\nu_{1,k+1}|^2} = \frac{\|b - By\|^2}{(\nu^T y)^2} \geq \gamma$$

which can be written in matrix form as

$$\begin{pmatrix} y^T & 1 \end{pmatrix} \begin{pmatrix} B^T B - \gamma \nu \nu^T & -B^T b \\ -b^T B & b^T b \end{pmatrix} \begin{pmatrix} y \\ 1 \end{pmatrix} \geq 0$$

In our case we have to solve

$$(V_k^T V_k) s = V_k^T A v_k, \quad (V_k^T V_k) p = \nu$$

Properties of the optimal basis

$$\tilde{v} = (I - V_k(V_k^T V_k)^{-1} V_k^T) A v_k - \frac{\alpha}{\omega} V_k (V_k^T V_k)^{-1} \nu$$

$$V_{k+1}^T v_{k+1} = \frac{1}{\nu_{1,k+1}} \begin{pmatrix} \nu_{1,1} \\ \vdots \\ \nu_{1,k} \\ \nu_{1,k+1} \end{pmatrix}$$

$$V_k^T V_k = \begin{pmatrix} 1 & \frac{1}{\nu_{1,2}} & \frac{1}{\nu_{1,3}} & \cdots & \frac{1}{\nu_{1,k}} \\ \frac{1}{\nu_{1,2}} & 1 & \frac{\nu_{1,2}}{\nu_{1,3}} & \cdots & \frac{\nu_{1,2}}{\nu_{1,k}} \\ \frac{1}{\nu_{1,3}} & \frac{\nu_{1,2}}{\nu_{1,3}} & 1 & \cdots & \frac{\nu_{1,3}}{\nu_{1,k}} \\ \vdots & \vdots & & \ddots & \vdots \\ \frac{1}{\nu_{1,k}} & \frac{\nu_{1,2}}{\nu_{1,k}} & \cdots & & 1 \end{pmatrix}$$

When the method converges, the basis is more and more orthogonal

The inverse of $V_k^T V_k$ is tridiagonal and the matrix $V_k^T A V_k$ is upper triangular

We know the entries of $(V_k^T V_k)^{-1}$
Let α_j be the diagonal entries and β_j the subdiagonal entries

$$\alpha_1 = \frac{\nu_{1,2}^2}{\nu_{1,2}^2 - 1}, \quad \beta_1 = -\frac{\nu_{1,2}}{\nu_{1,2}^2 - 1}$$

$$\alpha_i = \frac{\nu_{1,i-1}^2}{\nu_{1,i}^2 - \nu_{1,i-1}^2} + \frac{\nu_{1,i+1}^2}{\nu_{1,i+1}^2 - \nu_{1,i}^2}, \quad \beta_i = -\frac{\nu_{1,i}\nu_{1,i+1}}{\nu_{1,i+1}^2 - \nu_{1,i}^2}$$

and

$$\alpha_k = \frac{\nu_{1,k}^2}{\nu_{1,k}^2 - \nu_{1,k-1}^2}$$

This result could help us obtaining lower bounds on the singular values of V_k

$$p = (V_k^T V_k)^{-1}\nu = \begin{pmatrix} 0 \\ \vdots \\ 0 \\ \nu_{1,k} \end{pmatrix}$$

We will use this relation to simplify the construction of the basis vectors

The relation giving $V_k^T V_k$ cannot be used numerically because it will lead to a discrepancy between the computed vectors v_j and the computed $V_k^T V_k$

We can simplify the formulas for the new vector

$$\omega = p^T V_k^T A v_k = \nu_{1,k} v_k^T A v_k$$

Let $y_{opt} = s + \frac{\alpha}{\omega} p$

$$\begin{aligned} \tilde{v} &= A v_k - V_k y_{opt} \\ &= A v_k - V_k s - \frac{\alpha}{\omega} V_k p \\ &= A v_k - V_k s - \frac{\alpha}{\omega} \nu_{1,k} v_k \\ &= A v_k - V_k s - \frac{\alpha}{v_k^T A v_k} v_k \end{aligned}$$

and

$$h_{1:k,k} = s + \beta e_k, \quad \beta = \frac{\alpha}{v_k^T A v_k}$$

The Q-OR optimal algorithm

We compute incrementally the inverses of the Cholesky factors of $V_k^T V_k$

Let $v_k^A = Av_k$

1- $v_k^V = V_{k-1}^T v_k, \;\; v_k^{tA} = V_k^T v_k^A$

2- $\ell_k = \tilde{L}_{k-1} v_k^V, \; y_k^T = \ell_k^T \tilde{L}_{k-1}$

3- if $\ell_k^T \ell_k < 1, \; \ell_{k,k} = \sqrt{1 - \ell_k^T \ell_k}$, else $(p_k^v)^T = y_k^T V_{k-1}^T$, $\ell_{k,k} = \|v_k - p_k^v\|$ end

4-

$$\tilde{L}_k = \begin{pmatrix} \tilde{L}_{k-1} & 0 \\ -\frac{1}{\ell_{k,k}} y_k^T & \frac{1}{\ell_{k,k}} \end{pmatrix}$$

5- $\ell_A = \tilde{L}_k v_k^{tA}, \quad s = \tilde{L}_k^T \ell_A$

6- $\alpha = (v_k^A)^T v_k^A - \ell_A^T \ell_A, \;\; \beta = \frac{\alpha}{(v_k^{tA})_k}$

7-

$$h_{1:k,k} = \begin{pmatrix} h_{1,k} \\ \vdots \\ h_{k,k} \end{pmatrix} = s + \beta e_k$$

8-

$$\tilde{v} = v_k^A - V_k\, h_{1:k,k},\ \ h_{k+1,k} = \|\tilde{v}\|,\ \ \nu_{1,k+1} = -\frac{1}{h_{k+1,k}} \nu^T h_{1:k,k}$$

$$\nu = \begin{pmatrix} \nu_{1,1} & \cdots & \nu_{1,k+1} \end{pmatrix}^T$$

9- $v_{k+1} = \frac{1}{h_{k+1,k}} \tilde{v}$ and $v_{k+1}^A = A v_{k+1}$

10- if needed, solve $H_k y^{(k)} = \|b\| e_1$ using Givens rotations, $x_k = V_k y^{(k)}$

In this algorithm almost everything is expressed in terms of matrix-vector products

Numerical experiments

fs 680 1 of order 680 scaled by the inverse of its diagonal
It has 2184 non zero entries. The norm of A is 3.8168 and its condition number is 8.6944 10^3

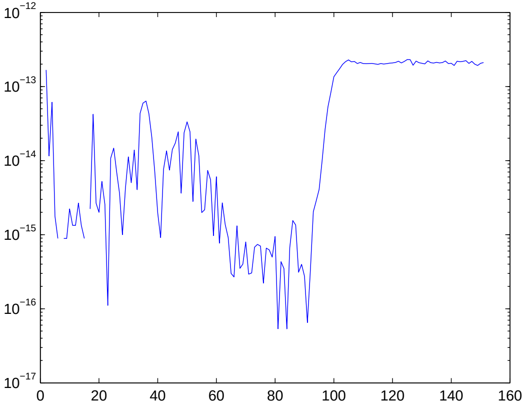

Difference of the true residual norms of GMRES-MGS and Q-OR optimal, fs 680 1c, $n = 680$

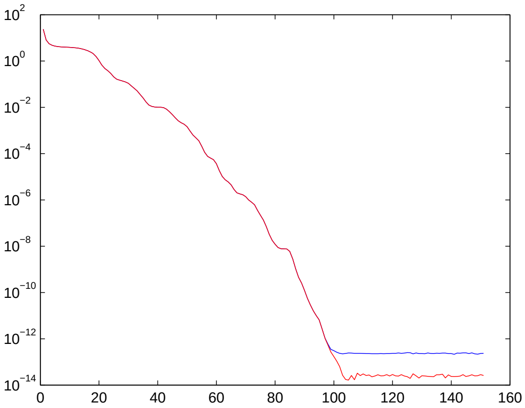

True residual norms of GMRES-MGS (blue) and Q-OR optimal (red), fs 680 1c, $n = 680$

True residual norms for $k = 150$ (maximum attainable accuracy)

- GMRES-CGS 6.8377 10^{-11}
- GMRES-CGS with reorthogonalization 2.79327 10^{-14}
- GMRES-CGS with double reorthogonalization 1.75040 10^{-14}
- GMRES-MGS 2.36046 10^{-13}
- GMRES-MGS with reorthogonalization 2.51184 10^{-14}
- GMRES-MGS with double reorthogonalization 1.59114 10^{-14}
- GMRES-Householder 1.51153 10^{-13}
- QOR opt 2.59770 10^{-14}

SUPG scheme (Streamwise upwind Galerkin) for a convection-diffusion equation in a square with a mesh size of $1/41$
The diffusion coefficient is $\nu = 0.01$
This matrix is of order 1600 and has 13924 non zero entries. Its norm is $4.8716\ 10^{-2}$ and the condition number is 40.518

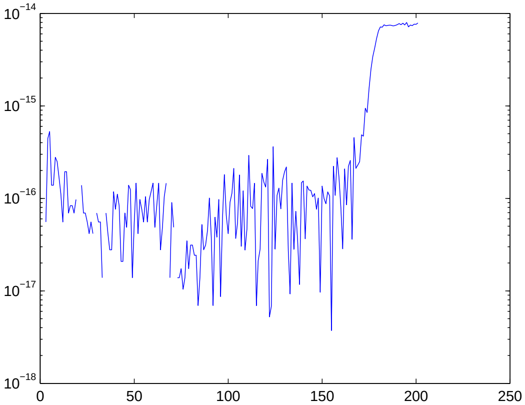

Difference of the true residual norms of GMRES-MGS and Q-OR optimal, supg 1600, $n = 1600$

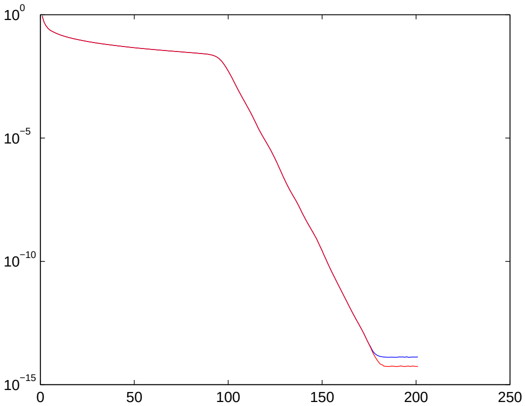

True residual norms of GMRES-MGS (blue) and Q-OR optimal (red), supg 1600, $n = 1600$

True residual norms for $k = 200$

- GMRES-CGS $1.54043\ 10^{-13}$
- GMRES-CGS with reorthogonalization $7.05585\ 10^{-15}$
- GMRES-CGS with double reorthogonalization $7.23790\ 10^{-15}$
- GMRES-MGS $1.33776\ 10^{-14}$
- GMRES-MGS with reorthogonalization $6.70649\ 10^{-15}$
- GMRES-MGS with double reorthogonalization $6.70339\ 10^{-15}$
- GMRES-Householder $2.03961\ 10^{-14}$
- QOR opt $5.50626\ 10^{-15}$

A smaller matrix for the same problem, $n = 100$

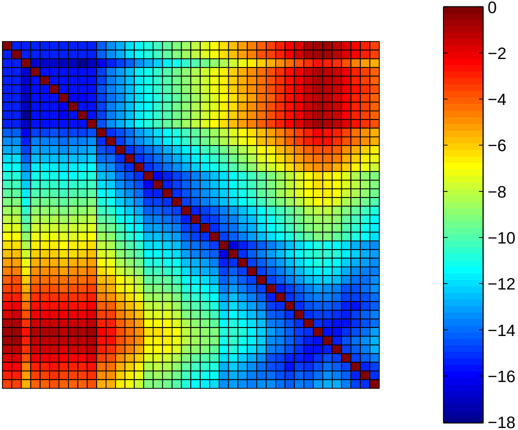

Supg 100, $\log_{10}$ of $|V^T V|$, GMRES-MGS

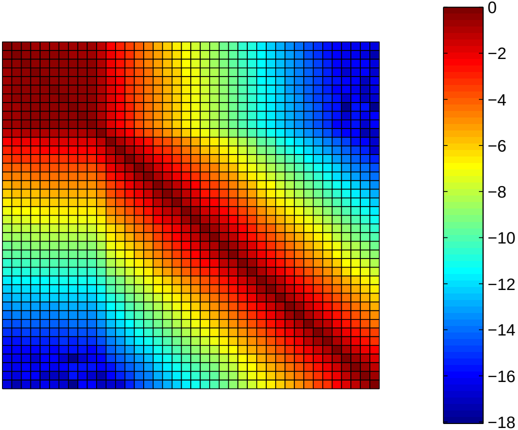

Supg 100, $\log_{10}$ of $|V^T V|$, QOR opt

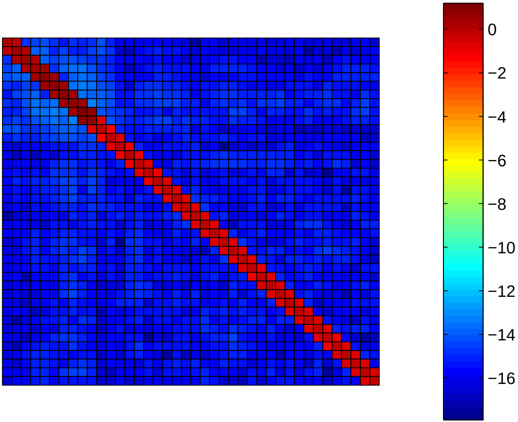

Supg 100, $\log_{10}$ of $|(V^T V)^{-1}|$, QOR opt

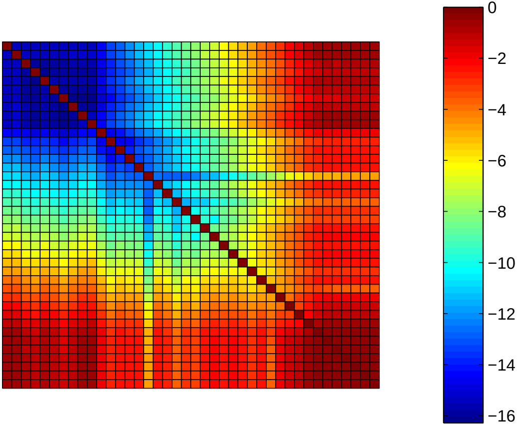

Supg 100, $\log_{10}$ of $|V^T V|$, GMRES-CGS

Conclusion

Using the properties of the Q-OR methods we were able to construct a non-orthogonal basis for which Q-OR gives the same residual norms as GMRES

The algorithm is slightly more expensive than GMRES

But, it is more parallel than GMRES-MGS and most of the operations are matrix-vector products

In many cases the maximum attainable accuracy is better than with GMRES-MGS

However, (at least theoretically), the algorithm is not breakdown-free

It remains to study its stability in finite precision arithmetic and to see how to use it on parallel computers