Rank One Updates that may Accelerate Restarted GMRES

Jurjen Duintjer Tebbens

Institute of Computer Science

Academy of Sciences of the Czech Republic

German-Polish Workshop for Young Researchers in Applied and Numerical Linear Algebra, Będlewo, February 5, 2006

Outline

1. Motivation
2. The rank one updated system matrix
3. Prescription of the spectrum
4. Prescription of the convergence curve
5. Future work

1. Motivation

We recall a famous result from [Greenbaum, Strakoš - 94], [Greenbaum, Pták, Strakoš - 96], [Arioli, Pták, Strakoš - 98]:

Theorem: *If we have given*

- *A right-hand side* $b \in \mathbb{R}^n$
- *non-increasing values* $\|b\| = f_0 \geq f_1 \geq f_2 \cdots \geq f_{n-1} > f_n = 0$
- *a set of real and possibly complex conjugated values* $\{\theta_1, \ldots, \theta_n\}$

then there exists a class $\mathcal{A}$ *of real* $n \times n$ *matrices such that the residual vectors* r_j *obtained by application of GMRES to*

$$\mathbf{A}x = b, \quad \mathbf{A} \in \mathcal{A},$$

satisfy

$$\|r_j\| = f_j, \quad 0 \leq j \leq n \quad \text{and} \quad \sigma(\mathbf{A}) = \{\theta_1, \ldots, \theta_n\}.$$

The result is often used to show eigenvalues need not say anything about convergence speed.

Question: Can we use it for something constructive ? Given a linear system

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

can we transform $\mathbf{A}$ so that it belongs to a class $\mathcal{A}$ with prescribed convergence curve?

And can we exploit the transformed system to solve the original system ?

The proof shows the class $\mathcal{A}$ is given by matrices of the form

$$\mathbf{A} = \mathbf{WRHW}^T,$$

where $\mathbf{R}$ is upper triangular, $\mathbf{W} = (w_1, \ldots, w_n)$ is orthonormal such that

$$\mathbf{W}^T b = \begin{pmatrix} \pm\sqrt{f_0^2 - f_1^2} \\ \vdots \\ \pm\sqrt{f_{n-2}^2 - f_{n-1}^2} \\ \pm\sqrt{f_{n-1}^2} \end{pmatrix}, \quad \mathbf{H} = \begin{pmatrix} 0 & \ldots & 0 & 1/(b^T w_n) \\ 1 & & 0 & -(b^T w_1)/(b^T w_n) \\ & \ddots & \vdots & \vdots \\ 0 & \ldots & 1 & -(b^T w_{n-1})/(b^T w_n) \end{pmatrix}$$

and $\sigma(\mathbf{RH}) = \{\theta_1, \ldots, \theta_n\}$.

A simple transformation is the following rank one update (other rank one acceleration techniques include [Eirola, Nevanlinna - 1989] and [Bollhöfer, Mehrmann - 1998]):

2. The rank one updated system matrix

Consider the modified system matrix

$$\hat{\mathbf{A}} := \mathbf{A} - by^T,$$

where b is the right hand side and $y \in \mathbb{R}^n$ is a free parameter vector. Then the Sherman-Morrison formula gives

$$\mathbf{A}^{-1}b = (\hat{\mathbf{A}} + by^T)^{-1}b = \hat{\mathbf{A}}^{-1}b - \frac{y^T\hat{\mathbf{A}}^{-1}b}{1 + y^T\hat{\mathbf{A}}^{-1}b}\hat{\mathbf{A}}^{-1}b.$$

Computation of the rightmost expression involves solving the **auxiliary system** defined as

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

Let $\widehat{x}_k$ be a satisfactory approximation to the solution $\widehat{x}$ of the auxiliary system.

Then

$$\mathbf{A}^{-1}b = \widehat{\mathbf{A}}^{-1}b - \frac{y^T\widehat{\mathbf{A}}^{-1}b}{1 + y^T\widehat{\mathbf{A}}^{-1}b}\widehat{\mathbf{A}}^{-1}b = \widehat{x} - \frac{y^T\widehat{x}}{1 + y^T\widehat{x}}\widehat{x}$$

$$= \frac{1}{1 + y^T\widehat{x}}\widehat{x} \approx \frac{1}{1 + y^T\widehat{x}_k}\widehat{x}_k \equiv \bar{x}_k\,,$$

provided $y^T\widehat{x}_k \neq -1$.

We call $\bar{x}_k$ the **back-transformed iterate** for the original system.

We will try to enforce acceleration through special choices of the parameter vector y. The basic form of the algorithms we propose:

1. Apply m initial steps of GMRES to $\mathbf{A}x = b$ **with** $x_0 = 0$ and gain information to construct the parameter vector y.

2. Apply GMRES(m) to $\widehat{\mathbf{A}}\widehat{x} = b$, (**with** $\widehat{x}_0 = 0$ **for the initial cycle**), until a satisfactory approximation $\widehat{x}_k$ is found.

3. Back-transform to obtain the iterate $\bar{x}_k$ for the original system $\bar{x}_k = \frac{1}{1+y^T\widehat{x}_k}\widehat{x}_k$.

The advantage of executing initial cycles with zero initial guesses is that with $x_0 = 0 = \widehat{x}_0$

$$\mathcal{K}_k(\widehat{\mathbf{A}}, b) = \text{span}\{b, (\mathbf{A} - by^T)b, \ldots, (\mathbf{A} - by^T)^{k-1}b\},$$

hence

$$\mathcal{K}_k(\widehat{\mathbf{A}}, b) = \mathcal{K}_k(\mathbf{A}, b)$$

and information gained from the GMRES process applied to $\mathbf{A}x = b$ can be easily translated to a process for $\widehat{\mathbf{A}}\widehat{x} = b$.

In this case we can either prescribe the spectrum or prescribe the convergence curve by the choice of y (but not both).

3. Prescription of the spectrum

Theorem 1: *Let the matrix* $\mathbf{B} \in \mathbb{R}^{l \times l}$ *and the vector* $c \in \mathbb{R}^l$ *be such that the Krylov subspace* $\mathcal{K}_l(\mathbf{B}, c)$ *has full dimension and let* $\{\theta_1, \ldots, \theta_l\}$ *be a set of real and complex conjugate values. Then there exists a vector* $z \in \mathbb{R}^l$ *such that* $\mathbf{B} - cz^T$ *has the eigenvalues* $\theta_1, \ldots, \theta_l$.

A simplified proof for diagonalizable $\mathbf{B}$, $\mathbf{B} = \mathbf{XDX}^{-1}, \mathbf{D} = \text{diag}(d_1, \ldots, d_l)$:

For $z \in \mathbb{R}^l$,

$$\sigma(\mathbf{B} - cz^T) = \sigma\left(\mathbf{X}(\mathbf{D} - \mathbf{X}^{-1}cz^T\mathbf{X})\mathbf{X}^{-1}\right) = \sigma(\mathbf{D} - \mathbf{X}^{-1}cz^T\mathbf{X}).$$

We put $\tilde{c} = \mathbf{X}^{-1}c$ and $\tilde{z} = \mathbf{X}^T z$.

The characteristic polynomial of the rank one updated diagonal matrix $\mathbf{D} - \tilde{c}\tilde{z}^T$ equals

$$\prod_{i=1}^{l}(d_i-\lambda)-(\tilde{c}_1\tilde{z}_1)\prod_{i\neq 1}^{l}(d_i-\lambda)-(\tilde{c}_2\tilde{z}_2)\prod_{i\neq 2}^{l}(d_i-\lambda)\cdots-(\tilde{c}_l\tilde{z}_l)\prod_{i\neq l}^{l}(d_i-\lambda),$$

see e.g. [DT - 2004]. Define the numbers $\alpha_{i,j}$ through

$$\prod_{i=1}^{l}(d_i-\lambda) =: \sum_{i=0}^{l}\alpha_{i,0}\lambda^i, \qquad \prod_{i=1, i\neq j}^{l}(d_i-\lambda) =: \sum_{i=0}^{l-1}\alpha_{i,j}\lambda^i.$$

The coefficients α_i before λ of the whole characteristic polynomial depend upon $(\tilde{c}_1\tilde{z}_1), \ldots, (\tilde{c}_l\tilde{z}_l)$ as

$$\alpha_i = \alpha_{i,0} - \alpha_{i,1}(\tilde{c}_1\tilde{z}_1) - \ldots \alpha_{i,l}(\tilde{c}_l\tilde{z}_l).$$

They can be forced to take any value.

Especially, they can assume the coefficients of the wanted characteristic polynomial

$$\prod_{i=1}^{l} (\lambda - \theta_l)$$

by solving

$$\begin{pmatrix} \alpha_{0,1} & \alpha_{0,2} & \cdots & \alpha_{0,l} \\ \vdots & & & \\ \alpha_{l-1,1} & & & \alpha_{l-1,l} \end{pmatrix} \begin{pmatrix} \tilde{c}_1 \tilde{z}_1 \\ \vdots \\ \tilde{c}_l \tilde{z}_l \end{pmatrix} = \begin{pmatrix} \alpha_{0,0} \\ \vdots \\ \alpha_{l-1,0} \end{pmatrix},$$

provided the system is non-singular. After that, put $c = \mathbf{X}\tilde{c}$ and $z = \mathbf{X}^{-T}\tilde{z}$. $\square$

The proof shows prescription of the eigenvalues of the matrix $\mathbf{A} - by^T$ is too expensive for a large system dimension n.

But: We have to our disposal a small matrix whose eigenvalues (the Ritz values) approximate the eigenvalues of $\mathbf{A}$:

$$\mathbf{H}_k = \mathbf{V}_k^T \mathbf{A} \mathbf{V}_k,$$

the restriction of $\mathbf{A}$ to $\mathcal{K}_k(\mathbf{A}, r_0)$.

Moreover, with $\widehat{x}_0 = 0 = x_0$, $\hat{\mathbf{V}} = \mathbf{V}$ and

$$\hat{\mathbf{H}}_k = \mathbf{V}_k^T(\mathbf{A} - by^T)\mathbf{V}_k = \mathbf{H}_k - \mathbf{V}_k^T b y^T \mathbf{V}_k,$$

and $\widehat{\mathbf{H}}_k$ is again a rank one update of $\mathbf{H}_k$.

Hence with Theorem 1 we can prescribe the eigenvalues of $\widehat{\mathbf{H}}_k$, and presume that the eigenvalues of $\widehat{\mathbf{A}}$ are close.

If the matrix $\mathbf{A}$ is close to normal, convergence can be related to eigenvalues and we may design a parameter vector that eliminates convergence hampering eigenvalues.

Example: A driven cavity matrix from Matrix Market (Cavity01) $\mathbf{A} \in \mathbb{R}^{317\times 317}$, with $\|\mathbf{X}\|\|\mathbf{X}^{-1}\| = \pm 20$. With right hand side $b = (1, \ldots, 1)^T)/\sqrt{317}$, GMRES(10) generates a Hessenberg matrix $\mathbf{H}_{10}$ with

$$\sigma(\mathbf{H}_{10}) = \{0.046, 0.5, \ldots, 7327\}.$$

We can define $y \in \mathbb{R}^{317}$ such that GMRES for

$$(\mathbf{A} - by^T)\widehat{x} = b$$

generates a Hessenberg matrix $\widehat{\mathbf{H}}_{10}$ with eigenvalue 30 instead of 0.046 (the others are not modified). The corresponding auxiliary system converges as

Near-normal driven cavity matrix of dimension 317

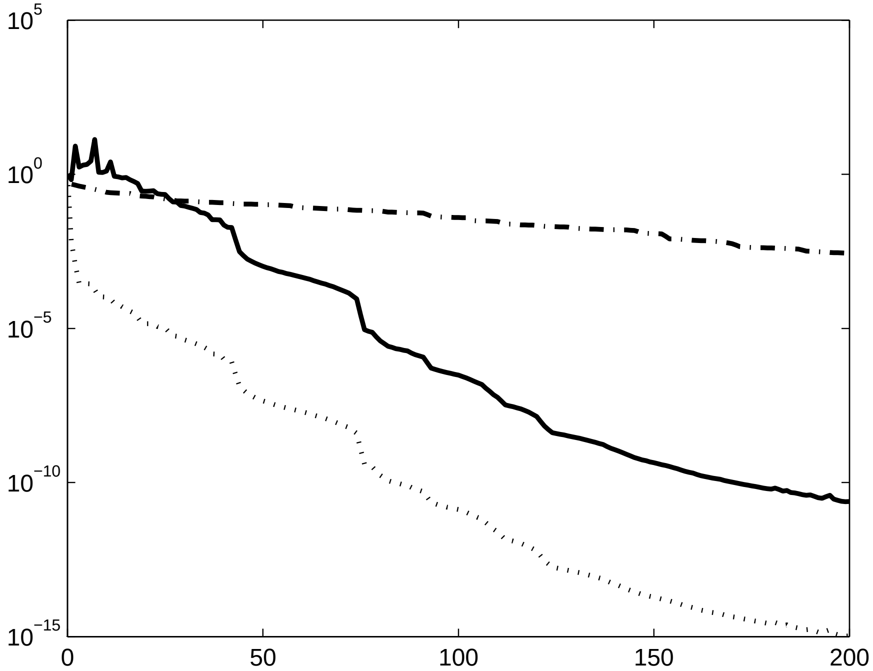

Dash-dotted line: Original system
Dotted line: Rank one updated system
Solid line: Residual norms after back-transformation

We observe quite a large gap, which is due to

$$\|\bar{r}_k\| = \|b - \mathbf{A}\bar{x}_k\| = \frac{\|b - \widehat{\mathbf{A}}\widehat{x}_k\|}{|1 + y^T \widehat{x}_k|} = \frac{\|\widehat{r}_k\|}{|1 + y^T \widehat{x}_k|}.$$

What can we say about the denominator $|1 + y^T \widehat{x}_k|$?

The Sherman-Morrison Theorem states: Under the assumption that $\widehat{\mathbf{A}}$ is non-singular,

$$1 + y^T \widehat{\mathbf{A}}^{-1} b \neq 0 \qquad \Leftrightarrow \qquad \widehat{\mathbf{A}} + b y^T = \mathbf{A} \text{ is non-singular.}$$

Hence unless we have forced a singular matrix $\widehat{\mathbf{A}} = \mathbf{A} - b y^T$, then with convergence we must have

$$1 + y^T \widehat{x}_k \approx 1 + y^T \widehat{\mathbf{A}}^{-1} b \neq 0.$$

More precisely: If $\|\widehat{x} - \widehat{x}_k\| = \varepsilon$, then

$$|1 + y^T \widehat{x}_k| > |1 + y^T \widehat{\mathbf{A}}^{-1} b| - \varepsilon \|y\|.$$

Whatever happens during the GMRES(m) process, at convergence the gap is non-zero.

We also now something about the gap created after the initial restart cycle:

Let $\tilde{\mathbf{H}}_k = \begin{pmatrix} \mathbf{H}_k \\ h_{k+1,k} e_k^T \end{pmatrix}$ and $\widehat{\mathbf{H}}_k = \mathbf{H}_k - \mathbf{V}_k^T b y^T \mathbf{V}_k \equiv \mathbf{H}_k - e_1 z^T$.

Lemma 1: *We introduce the notation*

$$a_1 = z^T(\tilde{\mathbf{H}}_k^T \tilde{\mathbf{H}}_k)^{-1} z, \quad a_2 = z^T(\tilde{\mathbf{H}}_k^T \tilde{\mathbf{H}}_k)^{-1} \tilde{\mathbf{H}}_k^T e_1 \ \text{ and } \ c = e_1^T \tilde{\mathbf{H}}_k (\tilde{\mathbf{H}}_k^T \tilde{\mathbf{H}}_k)^{-1} \tilde{\mathbf{H}}_k^T e_1.$$

Then after k iterations of the auxiliary system, the gap equals

$$1 + y^T \widehat{x}_k = 1 + \frac{a_1(c-1) + a_2(1-a_2)}{a_1(1-c) + (1-a_2)^2}.$$

Apart from the danger of a large gap, **our strategy suffers from too many degrees of freedom**: We can prescribe any spectrum, which one guarantees to overcome stagnation ?

An idea is to **modify only one Ritz value at the time** and define nested rank one updates. Then the Ritz value can be chosen uniquely by a **criterion that makes it as large as possible while keeping the gap small**. We take a closer look at this problem.

Changing only one Ritz value θ to $\widehat{\theta}$ is very easy, it suffices to define

$$z = (\theta - \widehat{\theta}) s_k,$$

where s_k is an appropriately scaled left eigenvector for θ, see e.g. [Zemke - 2006].

In fact, with a multiple of an exact left eigenvector y of $\mathbf{A}$ for λ_1,

$$y^T(\mathbf{A}-by^T) = (\lambda_1-y^Tb)y^T, \quad y^T\widehat{\mathbf{A}}^{-1} = y^T(\mathbf{A}-by^T)^{-1} = \frac{1}{\lambda_1 - y^Tb}y^T.$$

Let $\mathbf{A} = \mathbf{X}\Lambda\mathbf{X}^{-1}$ be the spectral decomposition of $\mathbf{A}$ where the elements of Λ appear in increasing order. Then $y = \alpha e_1^T\mathbf{X}^{-1}$ and

$$(\mathbf{A}-by^T)\mathbf{X} = \mathbf{X}\Lambda-\alpha be_1^T = \mathbf{X}(\Lambda-\alpha\mathbf{X}^{-1}be_1^T) = \mathbf{X}\left(\begin{pmatrix}\lambda_1 & & \\ & \ddots & \\ & & \lambda_n\end{pmatrix} - \alpha(\mathbf{X}^{-1}b, 0, \dots)\right)$$

hence $\widehat{\lambda}_1 = \lambda_1 - \alpha e_1^T\mathbf{X}^{-1}b$ and

$$1 + y^T\widehat{\mathbf{A}}^{-1}b = 1 + \frac{y^Tb}{\lambda_1 - y^Tb} = \frac{\lambda_1}{\lambda_1 - y^Tb} = \frac{\lambda_1}{\lambda_1 - \alpha e_1^T\mathbf{X}^{-1}b} = \frac{\lambda_1}{\widehat{\lambda}_1}\,!$$

We observed (but did not prove yet) that in case of modifying only Ritz values, with only left Ritz vectors, then

$$1 + y^T \widehat{\mathbf{A}}^{-1} b \qquad \sim \qquad \frac{\theta_1}{\widehat{\theta}_1},$$

too.

The modified Ritz value should be chosen as small as possible, but ensure non-stagnation of the auxiliary system. Example:

Near-normal driven cavity matrix of dimension 317

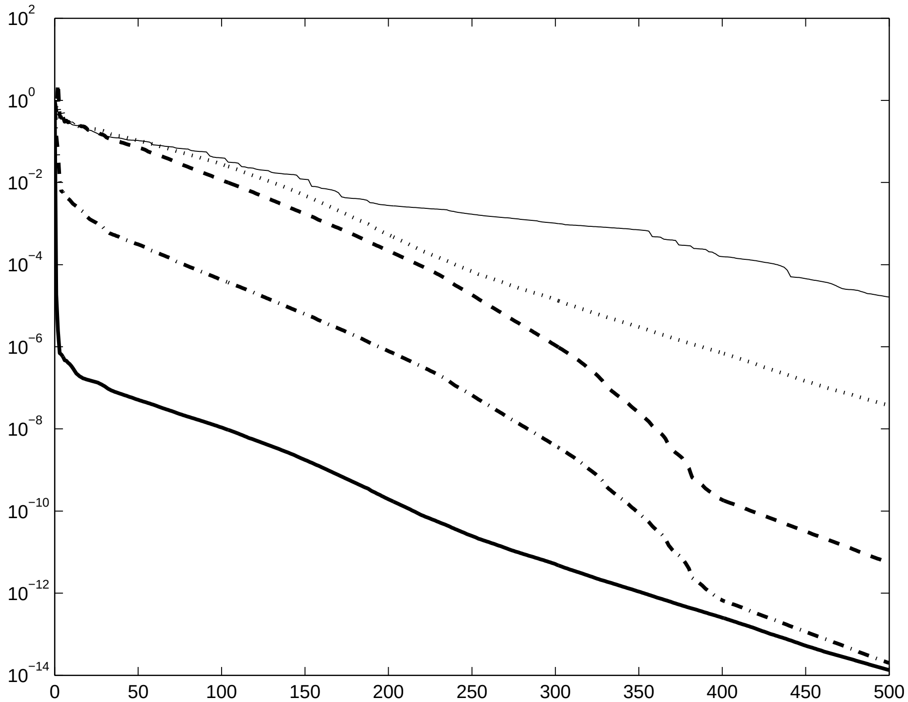

Thin line: Original system
Dotted line: Back-transformed system, $\widehat{\theta} = 10^6 \cdot \theta$
Dashed line: Back-transformed system, $\widehat{\theta} = 10^3 \cdot \theta$
Dash-dotted line: Rank one updated system, $\widehat{\theta} = 10^3 \cdot \theta$
Solid line: Rank one updated system, $\widehat{\theta} = 10^6 \cdot \theta$

4. Prescription of the convergence curve

For significantly non-normal matrices, it is in general not clear what properties cause restarted GMRES to stagnate, hence here we will prescribe convergence curves instead of the spectrum.

Theorem 2: *Let* $\|b\| = f_0 \geq f_1 \geq f_2 \cdots \geq f_k > 0, k < n,$ *be a non-increasing sequence of real values. If* $\mathcal{K}_k(\mathbf{A}, b)$ *has full dimension, then there exists at least one* $y \in \mathbb{R}^n$ *such that the residual vectors* $\widehat{r}_j$ *obtained by application of the GMRES method to the auxiliary system with* $\widehat{\mathbf{A}} = \mathbf{A} - by^T$ *and initial guess* $\widehat{x}_0 = 0$ *satisfy*

$$\|\widehat{r}_j\| = f_j, \quad 0 \leq j \leq k.$$

The proof [DT - 2004] is based on the observation that the convergence speed of GMRES depends on the distance from b to the subspaces

$$\hat{\mathbf{A}}\mathcal{K}_j(\hat{\mathbf{A}}, b) = (\mathbf{A} - by^T)\mathcal{K}_j(\mathbf{A}, b),\ j = 1, \ldots, k,$$

[Greenbaum, Strakoš - 94], which we can modify with the parameter vector y. We can

1. construct a vector $y \in \mathbb{R}^n$ that prescribes the first k residual norms generated by GMRES applied to $(\mathbf{A} - by^T)\hat{x} = b$ with initial guess $\hat{x}_0 = 0$;

2. apply GMRES(m) to $(\mathbf{A} - by^T)\hat{x} = b$ and presume that the first k residual norms of every restart decrease similarly as the prescribed norms of the initial cycle.

Example: Linear system from a convection-diffusion problem

$$-e^{-xy} \triangle u + (10 + ye^{-xy})u_x + (10 + xe^{-xy})u_y - 60u = 1$$

on the unit square with Dirichlet boundary condition $u = 0$ on $\partial([0,1])^2$. Finite difference approximation on a 35×35 grid yields the stiffness matrix $\mathbf{A} \in \mathbb{R}^{1225 \times 1225}$. Then $\mathbf{A} = \mathbf{XDX}^{-1}$, where $\|\mathbf{X}\|\|\mathbf{X}^{-1}\| = \pm 225$.

The right-hand side is $b = (1/35, \ldots, 1/35)^T$ and we choose $x_0 = 0$. We prescribe the first 9 residual norms of GMRES(30):

$\|\hat{r}_1\| = 0.9$, $\|\hat{r}_2\| = 0.8$, $\|\hat{r}_3\| = 0.7$, $\|\hat{r}_4\| = 0.6$, $\|\hat{r}_5\| = 0.5$, $\|\hat{r}_6\| = 0.4$, $\|\hat{r}_7\| = 0.3$, $\|\hat{r}_8\| = 0.2$, $\|\hat{r}_9\| = 0.1$.

1225 × 1225 non-normal convection-diffusion system

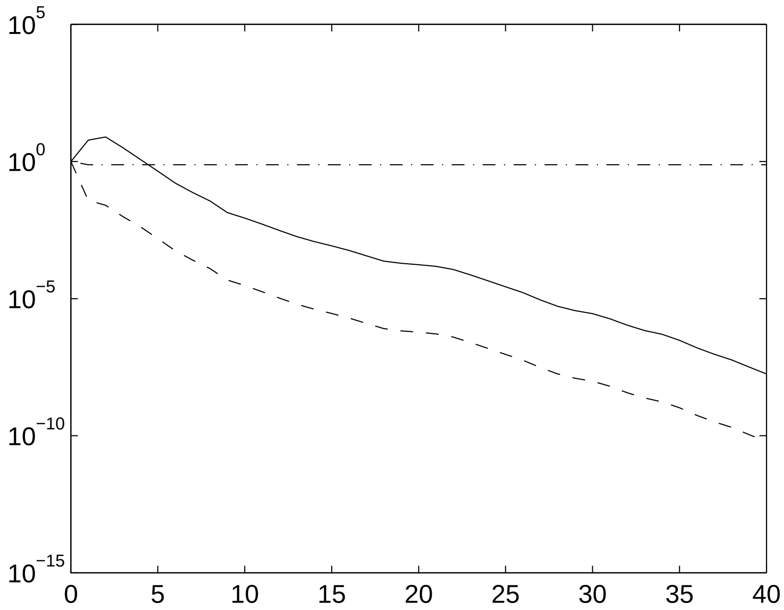

Dash-dotted line: Original system
Dashed line: Rank one updated system
Solid curve: Residual norms after back-transformation

5. Future work

- Near-normal case: The succeeding of our approach depends essentially on the quality of the used approximate eigenvalues. This holds in fact for many other GMRES accelerating techniques too. But our technique keeps computations in the small projected space. It should be compared with Morgan's augmentation technique, which is among spectrum modifying techniques optimal in some sense, see [Eiermann, Ernst - 2000].
- Non-normal case: More heuristical idea, less clear why the one stagnation can be overcome but the next can not. The influence of the choice of prescribed residual norms on the quality of the back-transformation has to be better understood.

We can extend the strategy to updates of higher rank. For example, with $\widehat{\mathbf{A}} = \mathbf{A} - (b, v_2, \ldots, v_k)\mathbf{Y}^T$ we can prescribe all elements of the Hessenberg matrix.

Note our techniques hold for other residual minimizing methods as well. For example, for symmetric indefinite systems it may be applied to MINRES. Here, prescription of spectrum should be even more effective. We would give up symmetry for the sake of definiteness.

Example: The shifted Laplacian $\mathbf{L} - 0.1\,\mathbf{I}$ of dimension 400. The smallest eigenvalues are -0.088 and -0.072. The Hessenberg matrix $\mathbf{H}_{15}$ has 2 negative eigenvalues -0.088 and -0.0452. We define an auxiliary Hessenberg matrix $\widehat{\mathbf{H}}_{15}$ where these smallest eigenvalues take the value 1. In the resulting auxiliary matrix -0.088 has vanished.

The shifted Laplacian $\mathbf{L} - 0.1\,\mathbf{I}$ of dimension 400

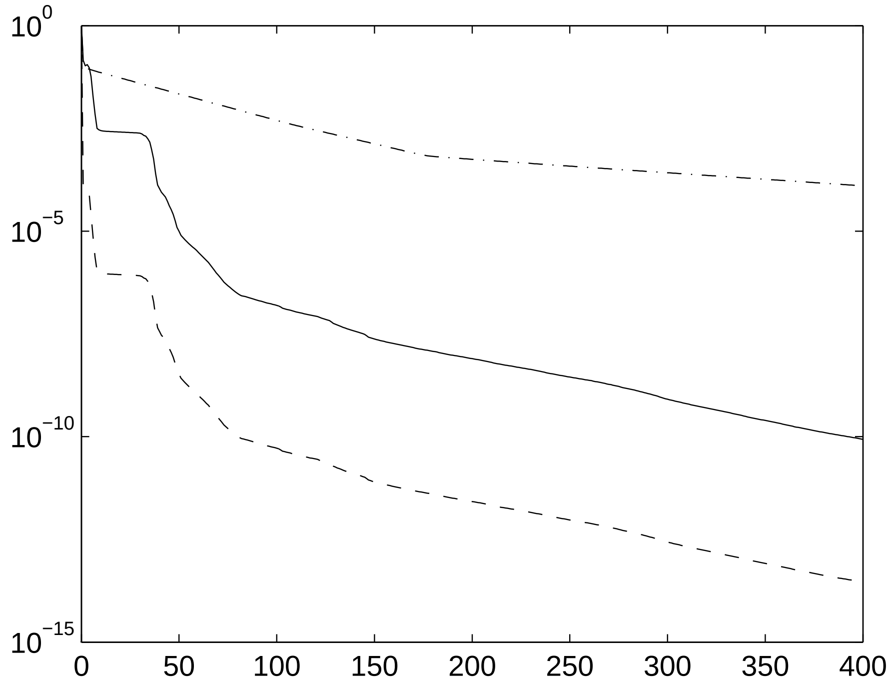

Dash-dotted line: Original system
Dashed line: Rank one updated system
Solid curve: Residual norms after back-transformation

Thank you for your attention

This work was supported by the Program Information Society under project 1ET400300415.