PANM 15, Dolní Maxov, 6.-11. června 2010

SPAM:

Subspace Projected Approximate Matrices

in Numerical Linear Algebra

Jan Brandts and Ricardo da Silva
Korteweg-de Vries Institute for Mathematics
University of Amsterdam, Netherlands

PANM 15, Dolní Maxov, 6.-11. června 2010

SPAM:

Subspace Projected Approximate Matrices

in Numerical Linear Algebra

♡ Definition: Subspace Projected Approximate Matrix

♠ Review: Subspace Methods for Linear Algebra Problems

♣ Investigation: Incorporating SPAM in Subspace Methods

A is a Hermitian $n \times n$ matrix and A_0 an "approximation" of A,

$$\{0\} = \mathcal{U}_0 \subset \mathcal{U}_1 \subset \cdots \subset \mathcal{U}_{n-1} \subset \mathcal{U}_n = \mathcal{C}^n \quad \text{and } \dim(\mathcal{U}_k) = k$$

Π_k is the orthogonal projection on $\mathcal{U}_k^{\perp}$. Then the matrix

$$A_k = A + \Pi_k(A_0 - A)\Pi_k$$

is the k-th SPAM matrix of the pair A, A_0

Π_k is the orthogonal projection on $\mathcal{U}_k^\perp$. The k-th SPAM matrix of the pair A, A_0 is the matrix

$$A_k = A + \Pi_k(A_0 - A)\Pi_k$$

Properties:

consistent definition for $k = 0$; moreover, $A_n = A$

$A_k^* = A_k$, and for all $u \in \mathcal{U}_k$ we have $A_k u = Au$ and $u^* A_k = u^* A$

Example: let $\mathcal{U}_k = \text{span}\{e_1, \ldots, e_k\}$ (first k standard basis vectors)

Then $A_k = A + \Pi_k(A_0 - A)\Pi_k$ is the matrix

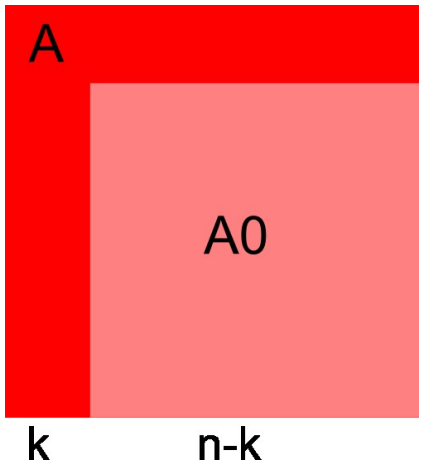

Example: let $\mathcal{U}_k = \text{span}\{e_1, \ldots, e_k\}$ (first k standard basis vectors)

Then $A_k = A + \Pi_k(A_0 - A)\Pi_k$ is the matrix

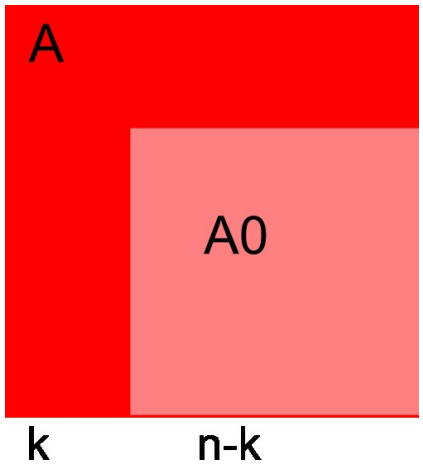

Example: let $\mathcal{U}_k = \text{span}\{e_1, \ldots, e_k\}$ (first k standard basis vectors)

Then $A_k = A + \Pi_k(A_0 - A)\Pi_k$ is the matrix

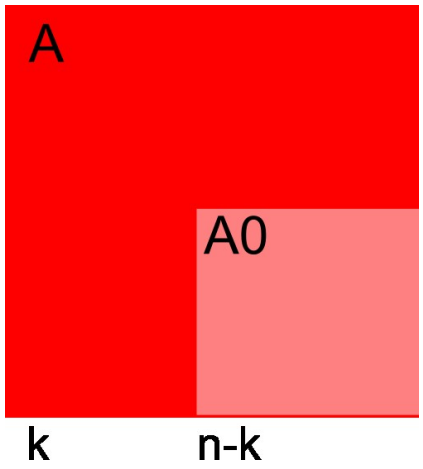

Notice: let the first k columns of U with $U^*U = I$ span $\mathcal{U}_k$.

Then U^*A_kU is a similar combination of U^*AU and U^*A_0U.

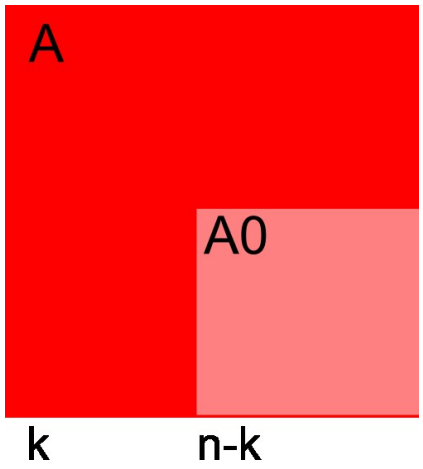

Notice: With respect to bases $\{u_1, \ldots, u_k\}$ for $\mathcal{U}_k$, the matrix A_k is a rank-2 update of A_{k-1} of arrowhead type

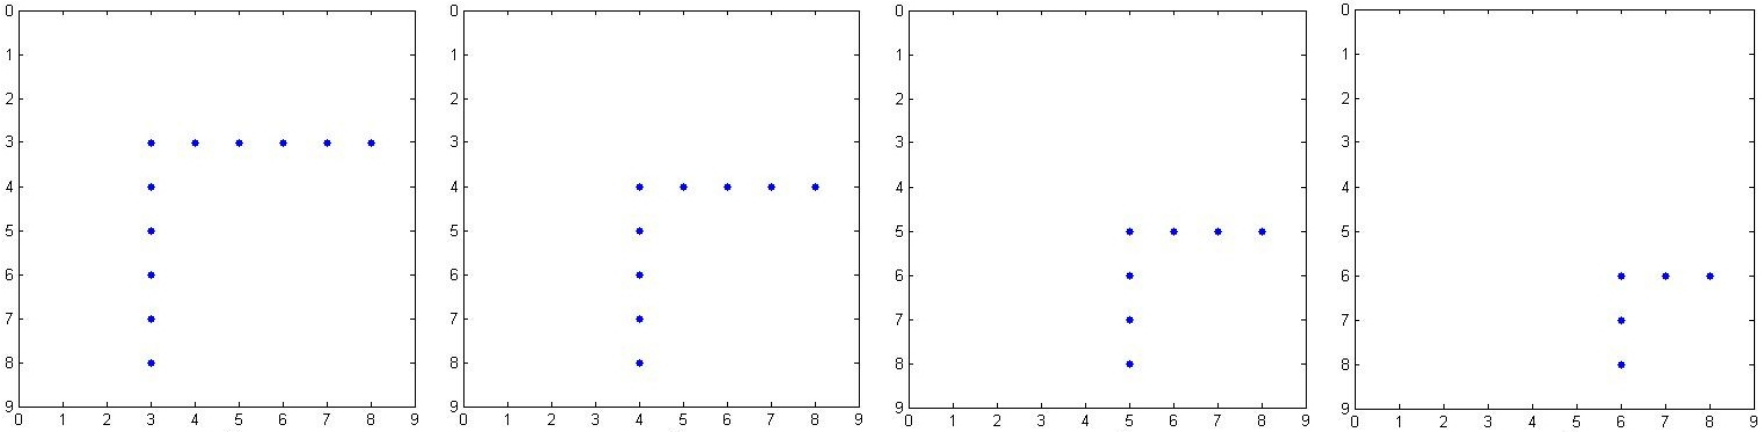

Notice: With respect to bases $\{u_1, \ldots, u_k\}$ for $\mathcal{U}_k$, the matrix A_k is a rank-2 update of A_{k-1} of arrowhead type

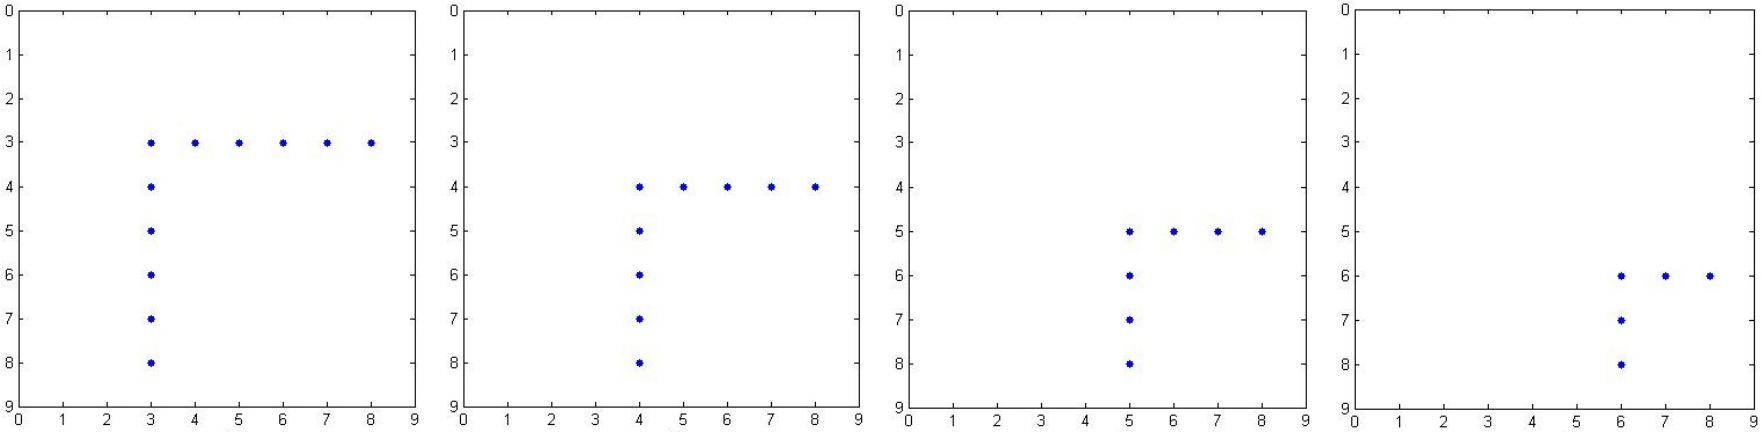

Let u be such that $\mathcal{U}_k = \mathcal{U}_{k-1} \oplus < u >$, with $u \perp \mathcal{U}_{k-1}$ and $\|u\| = 1$

Then A_k is an indefinite Hermitian rank-2 update of A_{k-1},

$$A_k = A_{k-1} + uv^* + vu^* = A_{k-1} + (u|v) \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} (u|v)^*$$

where

$$v = \left(\Pi_{k-1} - \tfrac{1}{2} uu^*\right)(A - A_0)u$$

The computational costs for updating involve one MV with A

Proposition: $U^* A_k U = M = U^* A U$ ("shared Ritz values")

Theorem: If $A - A_0 \geq 0$ then

$$\mu_j \leq \theta_j \leq \lambda_j$$

- μ_j eigenvalue of M (Ritz value);
- θ_j eigenvalue of A_k;
- λ_j eigenvalue of A.

Proved by the Cauchy Interlace Theorem and Weyl's Theorem

Recall: Hermitian eigenvalue problem $Ax = \lambda x$.

Subspace $\mathcal{U}_k$ with orthonormal basis U

Rayleigh-Ritz procedure:

Eigenpairs of $M = U^* A U$ define so-called Ritz pairs: improvingly good approximations to eigenpairs of A

If $\mathcal{U}_k$ is a Krylov subspace, we get the Lanczos method.

Lanczos method

Expand $\mathcal{U}_k$ with the current eigenvalue residual $r = Av_j - \mu_j v_j$

SPAM eigenvalue method (Shephard et al. 2001)

Expand $\mathcal{U}_k$ with the best eigenvector approximation of A_k

- this probably leads to faster convergence in k
- but requires an inner iteration for A_k
- if A_0 is "simple" this iteration is cheap

Approximation from below: constructing A_0 with $A - A_0 \geq 0$

- naturally available in numPDE context with diffusion term

Algebraically: Choose any $H \geq 0$ and set $A_0 = A - H$

Requirements for H:

- $A_0 = A - H$ should be sparser than A and/or of lower rank
- H should be poor approximation of A
- of course these requirements are mutually contradictory

PANM 15, Dolní Maxov, 6.-11. června 2010

Option: $H = EE^*AEE^*$ where E is a selection of columns of I

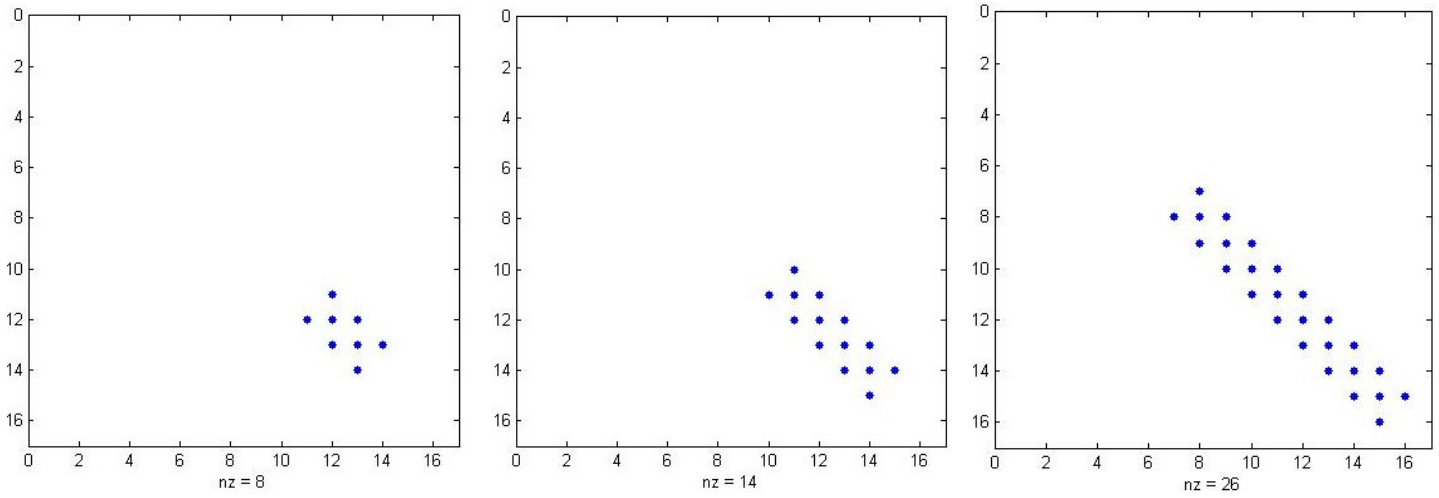

If $A \geq 0$ then also $H \geq 0$; depicted is $A_0 = A - H$

Option: $H = EE^*AEE^*$ where E is a selection of columns of I

Which selection?

- randomly: Ritz-Galerkin projection is not very good
- based on smallest diagonal elements

We will illustrate the latter selection strategy on a reaction-diffusion problem discretized by the finite difference method.

Illustration: Lanczos for 1d-reaction diffusion FD

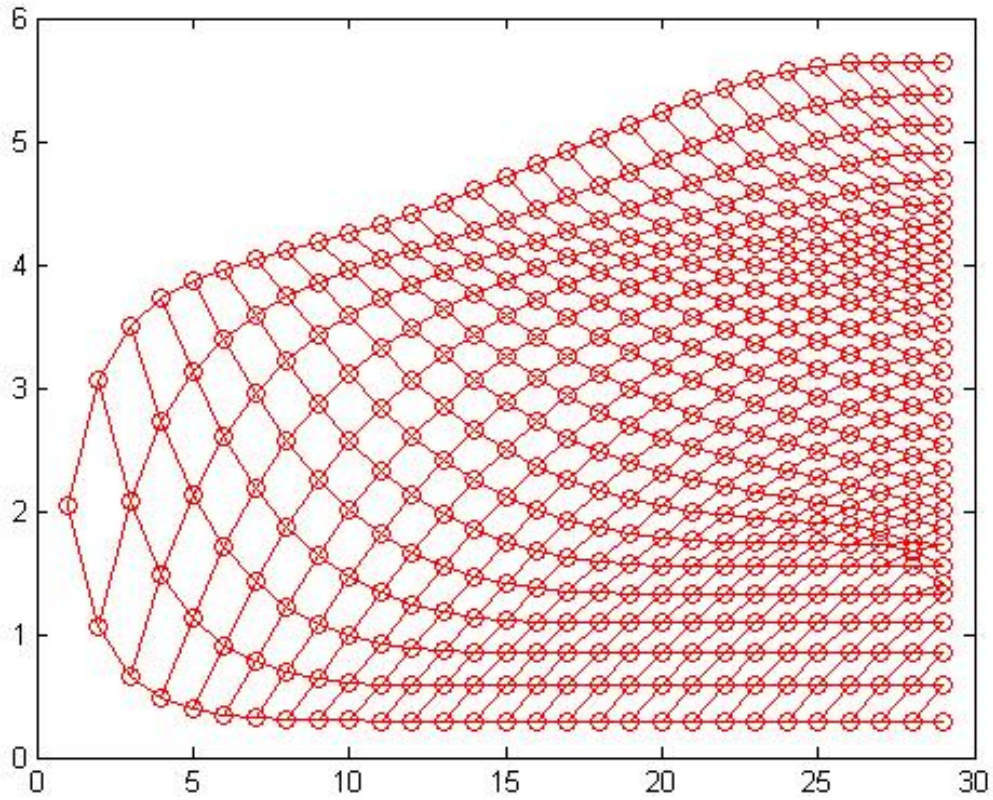

Illustration: SPAM with rank-2 approximation A_0

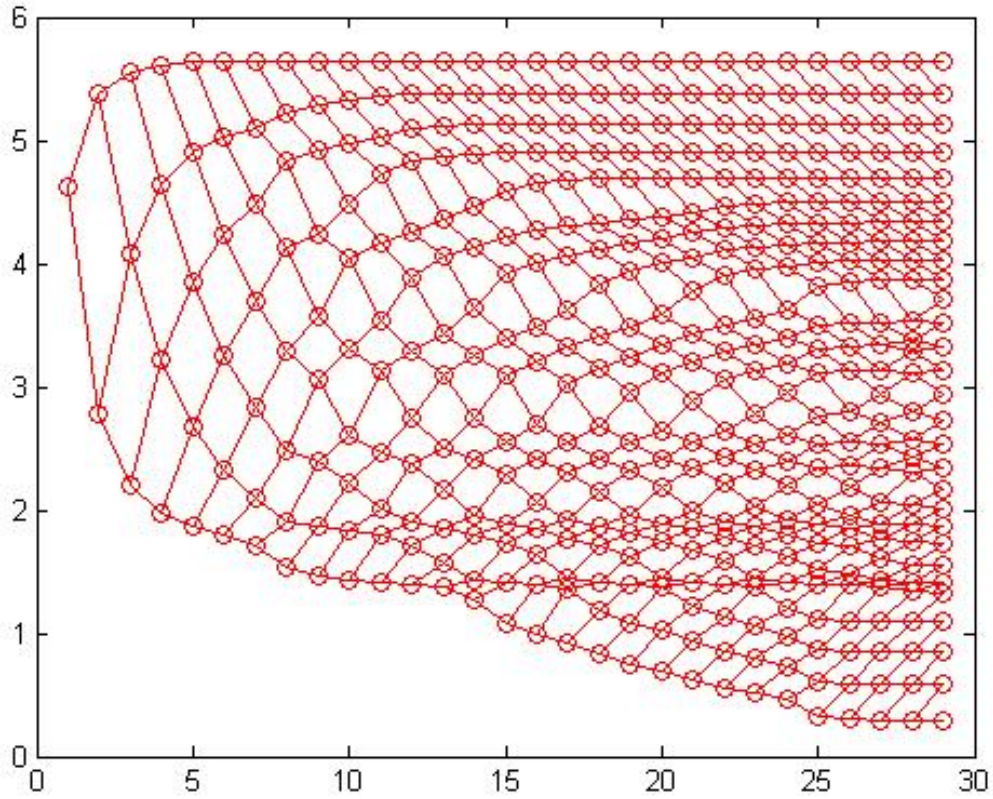

Illustration: SPAM with rank-4 approximation A_0

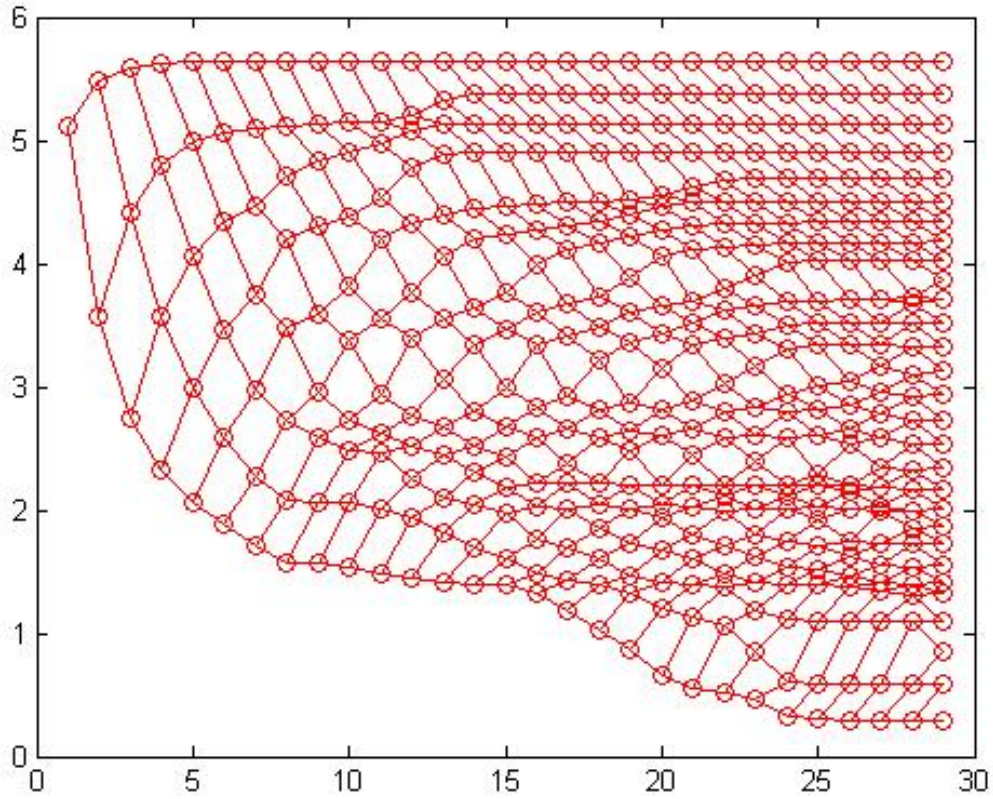

Illustration: SPAM with rank-5 approximation A_0

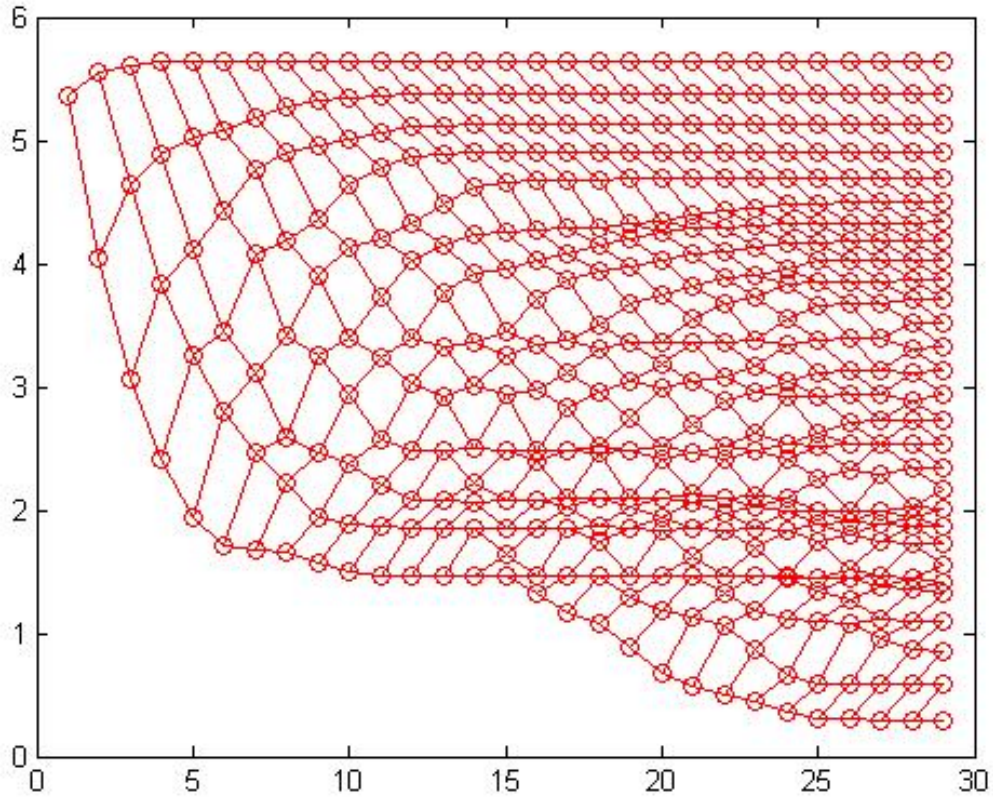

Illustration: SPAM with rank-6 approximation A_0

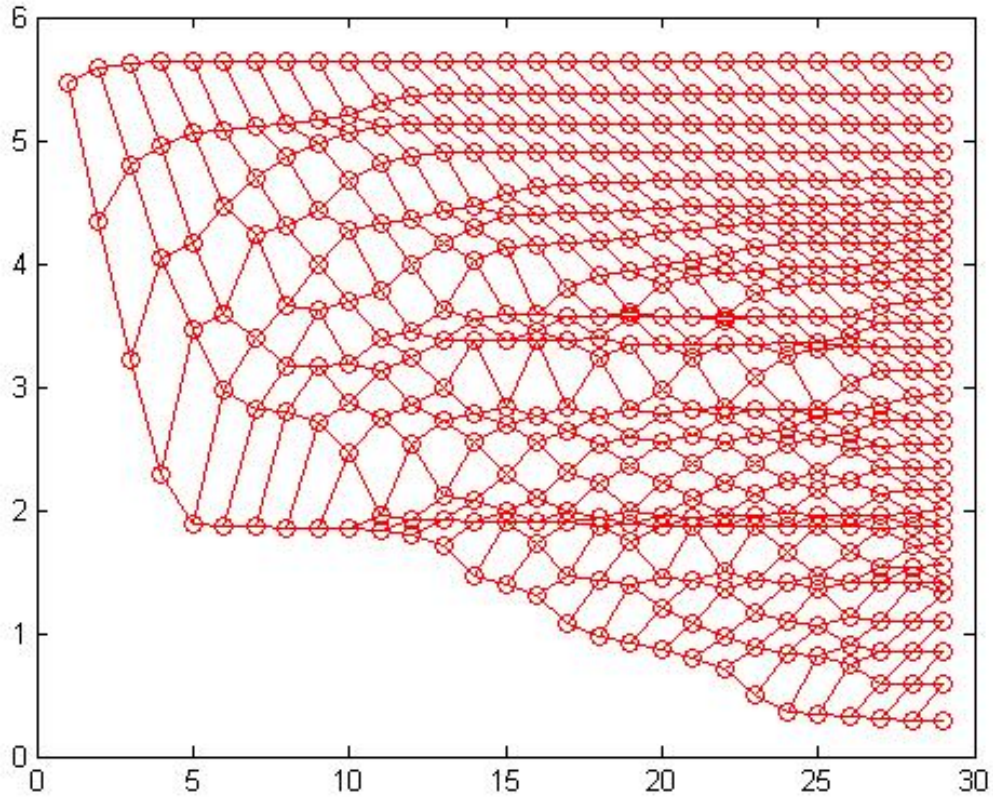

Illustration: SPAM with rank-7 approximation A_0

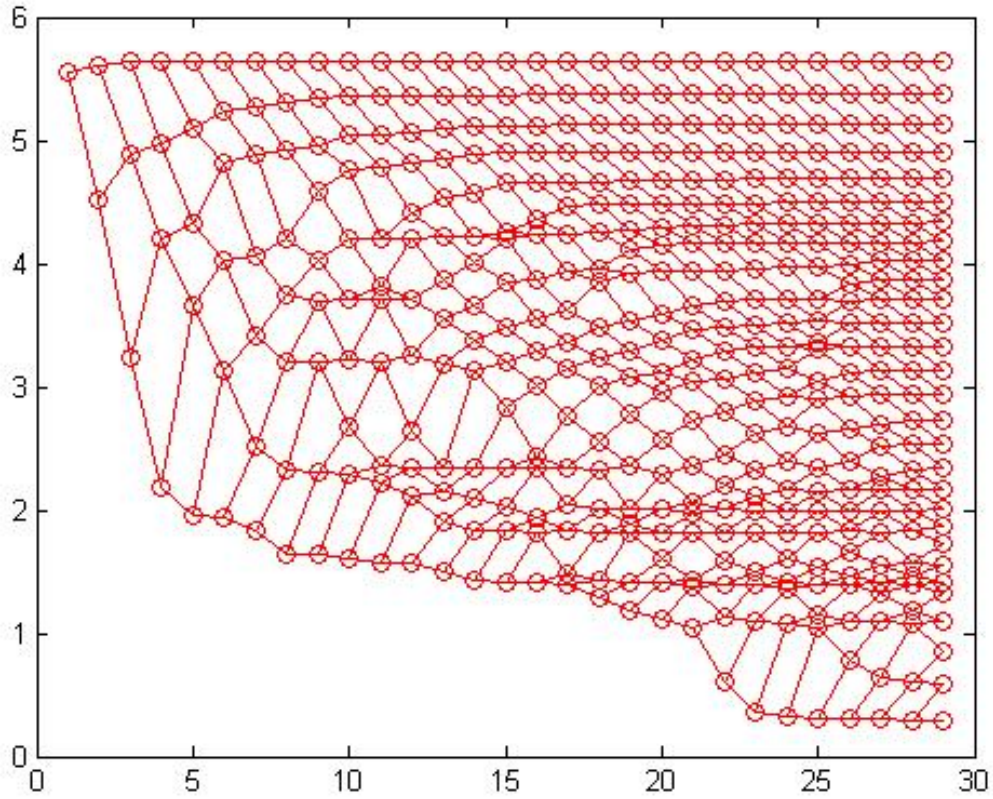

Illustration: SPAM with rank-8 approximation A_0

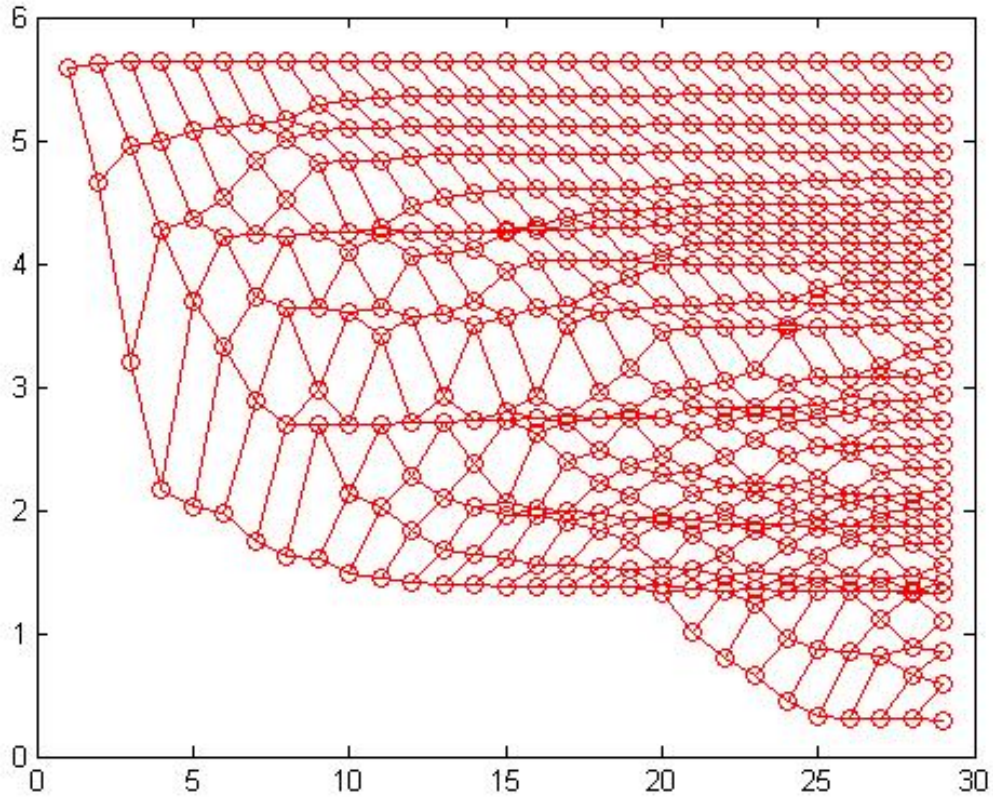

Illustration: SPAM with rank-9 approximation A_0

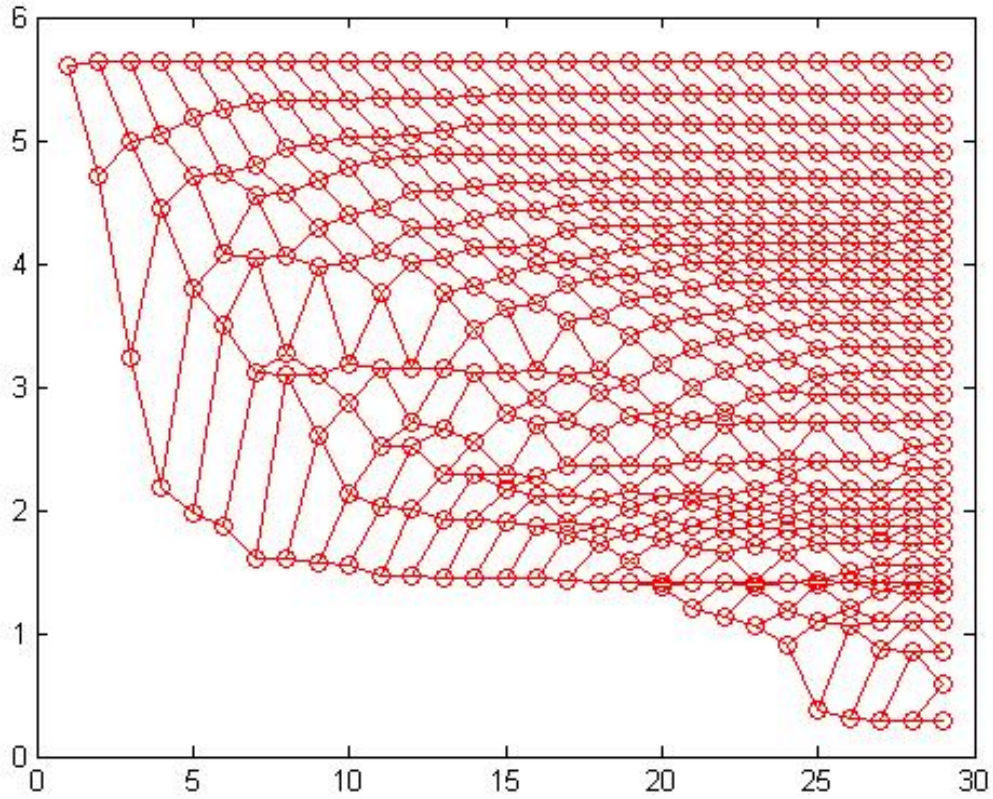

Illustration: SPAM with rank-10 approximation A_0

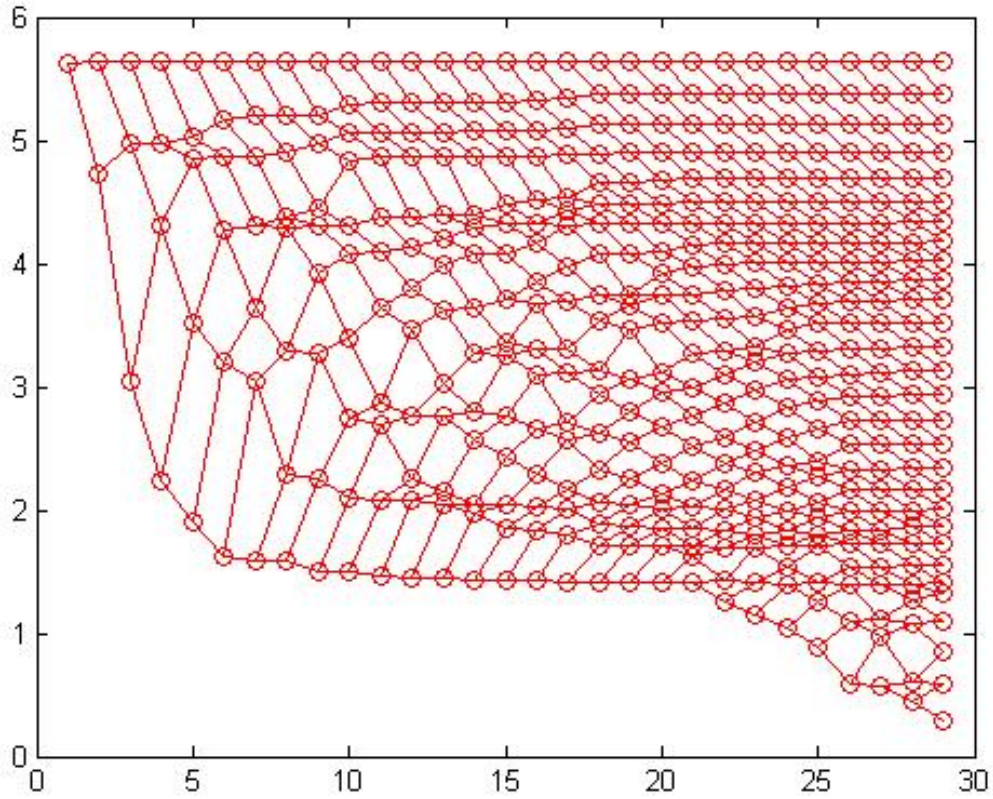

PANM 15, Dolní Maxov, 6.-11. června 2010

Question: SPAM seems to converge more quickly than Lanczos

Lanczos is, however, started with a random vector

SPAM "starts" with an eigenvector of A_0

It would be more fair to start Lanczos with the same vector

Illustration: Lanczos and SPAM with same startvector

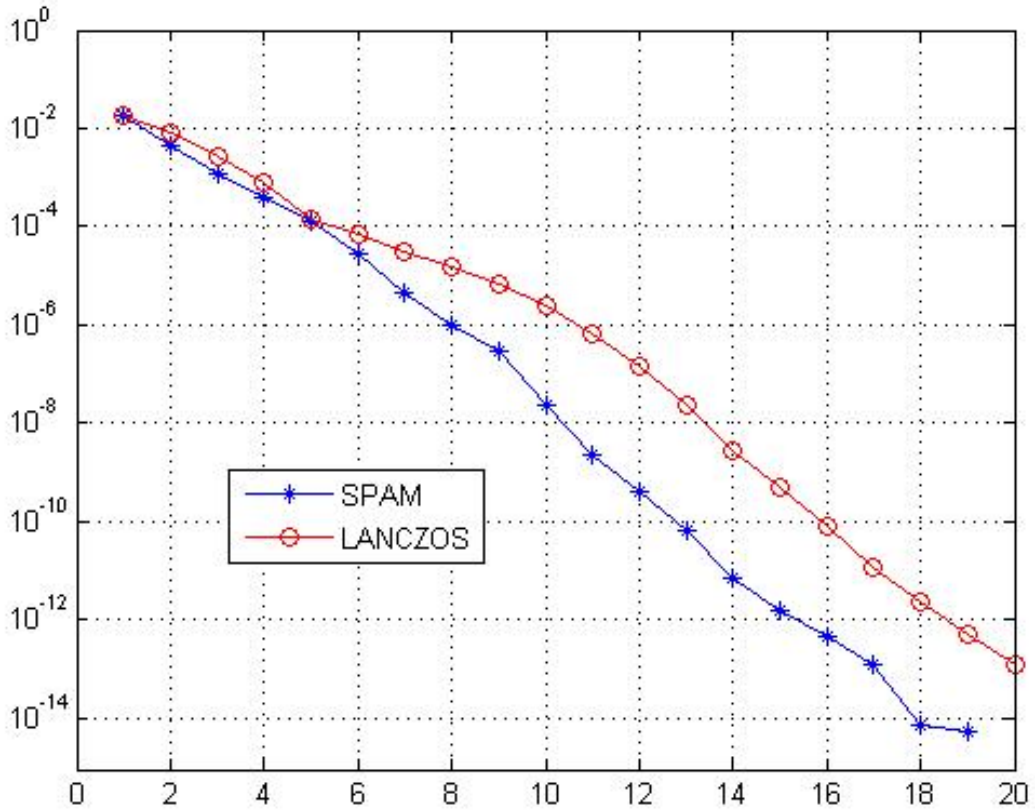

Illustration: Lanczos and SPAM with same startvector

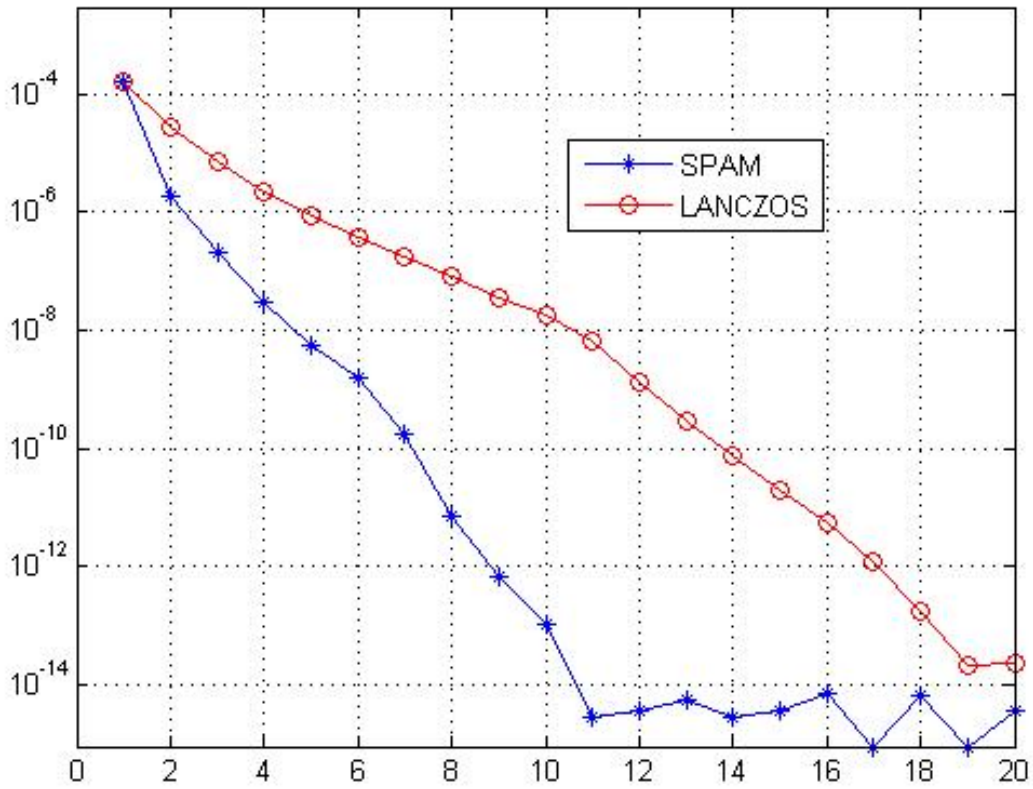

PANM 15, Dolní Maxov, 6.-11. června 2010

SPAM without outer iteration

In SPAM, the outer iteration ensures the interlace property and the monotonicity of the eigenvalue approximations

But, it is not strictly necessary

In the following experiments we compare the maximum eigenvalue of A_k with the maximum eigenvalue of SPAM

Illustration: SPAM without louter iteration

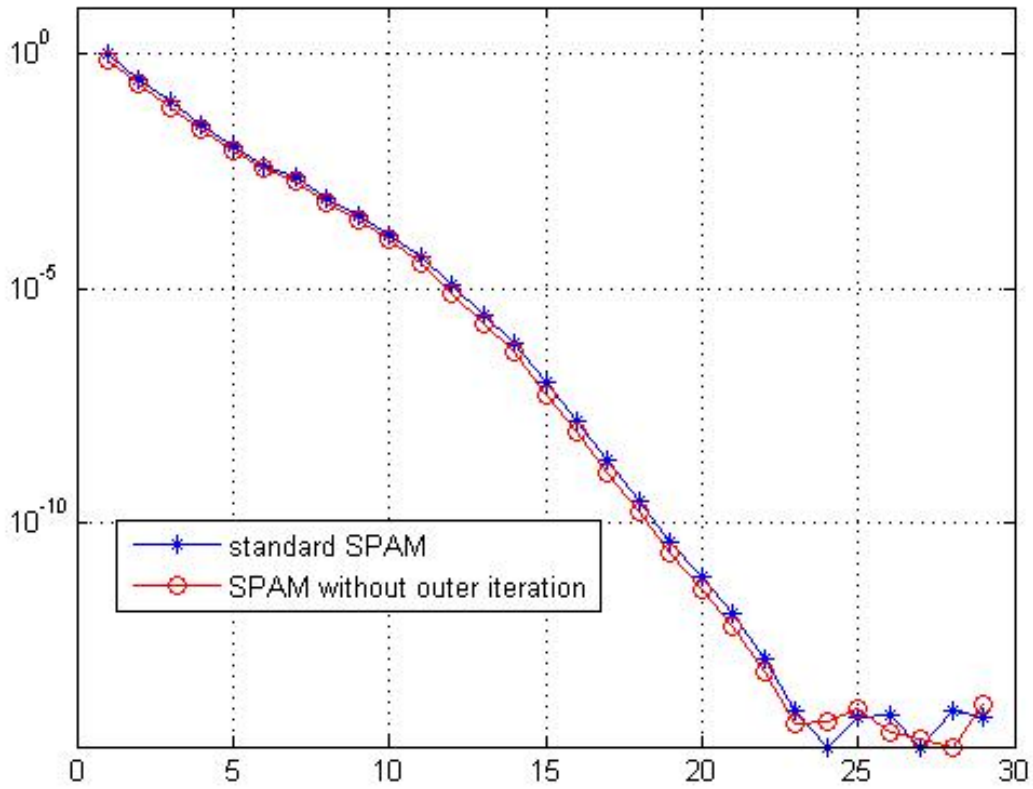

Illustration: SPAM without louter iteration

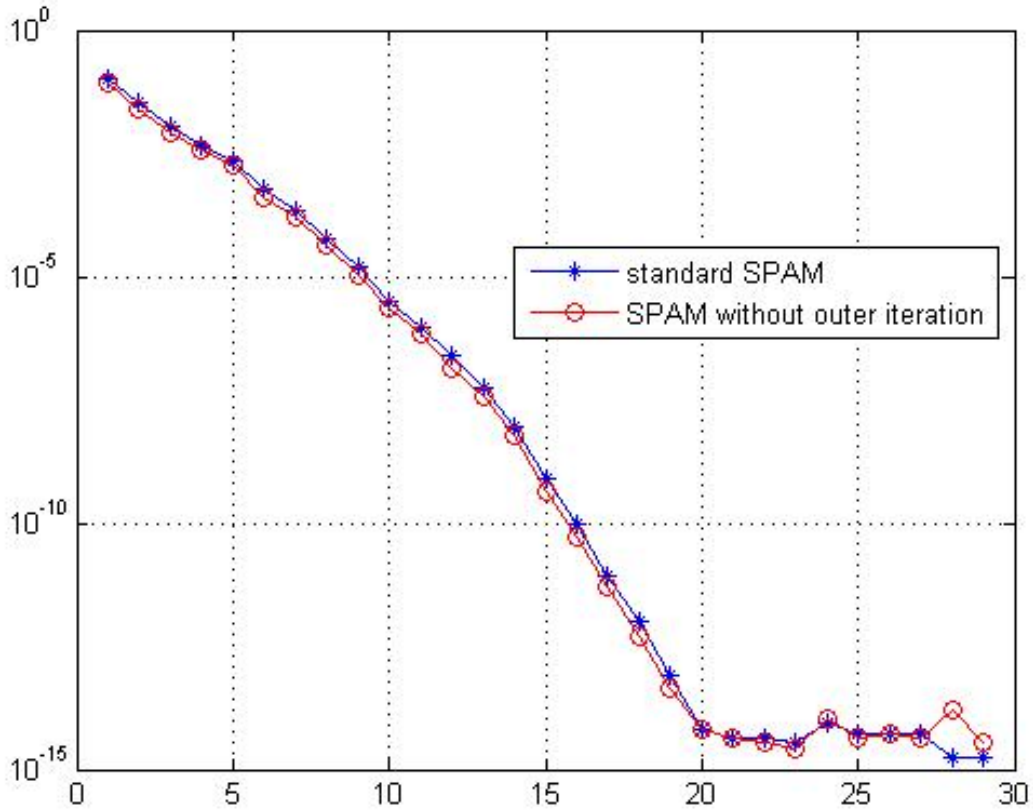

Illustration: SPAM without louter iteration

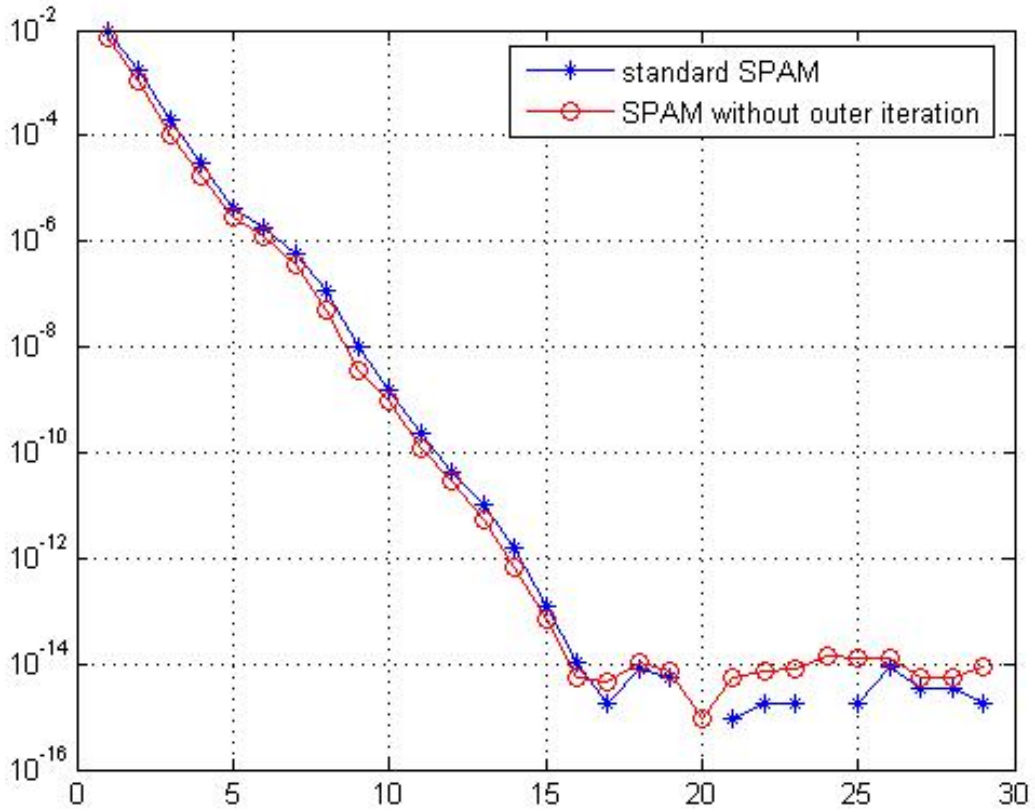

Other approximations A_0 (not necessarily with $A - A_0 \geq 0$)

In the original paper by Shephard et. al:

- A_0 is chosen k-diagonal with lower bandwidth then A
- tensor product aproximations

Their only goal is to have a cheaper MV, they do not care much about the approximation properties or the rank of A_k

The inner iteration

In the inner iteration we need to find an eigenvector of A_k

- a good start vector is available (eigenvector of A_{k-1})
- in fact, a good initial search space $\mathcal{U}_k$ is available

Shephard et al. propose to use SPAM recursively

This requires a range of approximations $A_0, \hat{A}_0, \hat{\hat{A}}_0, \ldots$

Jacobi-Davidson

JD converges quadratically if the start vector is close enough to the eigenvector

Often, Lanczos is used to get close enough to this eigenvector.

SPAM seems a good alternative:

- the rank of A_k is low for small k
- moreover, A_k can be used in JD's inner iteration

PANM 15, Dolní Maxov, 6.-11. června 2010

Preconditioned classical methods for $Ax = b$

Initial guess x_0, residual $r_0 = b - Ax_0$

repeat until $\|r_j\|$ small enough:

- solve $A_0 u_{j+1} = r_j$ and set $c_{j+1} = Au_{j+1}$
- $x_{j+1} = x_j + u_{j+1}$ with residual $r_{j+1} = r_j - c_{j+1}$

end loop

Here A_0 is a fixed preconditioner, for example, A's diagonal

repeat until $\|r_j\|$ small enough:

- solve $A_j u_{j+1} = r_j$ and set $c_{j+1} = A u_{j+1}$
- $x_{j+1} = x_j + u_{j+1}$ with residual $r_{j+1} = r_j - c_{j+1}$
- $A_{j+1} = A_j + uv^* + vu^*$

end loop

Logical choice is $u = u_{j+1}$, orthonormalized to all previous u_j,

because the action of A on u_{j+1} has already been stored in c_{j+1}

repeat until $\|r_j\|$ small enough:

- solve $A_j u_{j+1} = r_j$ and set $c_{j+1} = A u_{j+1}$
- $x_{j+1} = x_j + u_{j+1}$ with residual $r_{j+1} = r_j - c_{j+1}$
- $\widehat{u}_{j+1}$ is orthonormal to all previous $\widehat{u}_j$
- $\widehat{c}_{j+1} = A\widehat{u}_{j+1}$ (without performing the MV!)
- $A_{j+1} = A_j + \widehat{u}_{j+1} v^* + v \widehat{u}_{j+1}^*$

end loop

Solving $A_j u_{j+1} = r_j$ is cheap using the Woodbury formula:

- A_j is a rank $2j$ update of A_0;
- A_j is a rank-2 update of A_{j-1}

These updates are of the form $uv^* + vu^*$ and thus involve the solution of $A_0 w = u$ and $A_0 z = v$ only.

This would lead to a two-term recursion for the approximations.

Conclusions:

The SPAM matrices contain eigenvalue approximations better than the Rayleigh-Ritz approximations.

It is unclear if the additional computational effort to compute them is worthwhile; in some cases it is, in other cases it is not.

In the linear system context the SPAM matrix is a preconditioner that for some generic choices leads to standard Krylov methods ("updating the preconditioner")