Finite Precision Krylov Methods in a Nutshell

Jens–Peter M. Zemke

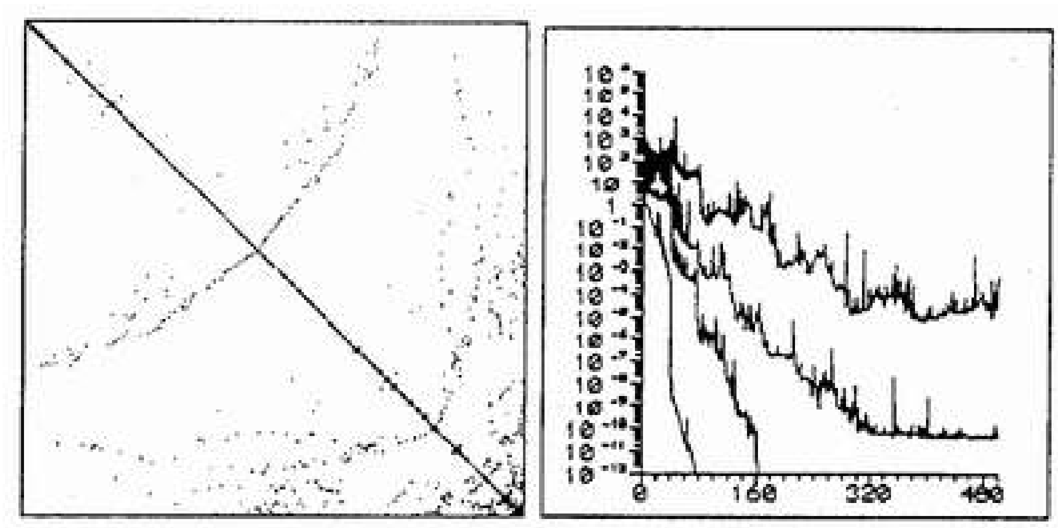

Talk at the Conference

Computational Linear Algebra with Applications

MILOVY 2002,

August 8, 2002, Milovy, Czech Republic

Table Of Contents

The Menagerie of Krylov Methods

- Lanczos based methods (short–term methods)
- Arnoldi based methods (long–term methods)

- eigensolvers: $Av = v\lambda$
- linear system solvers: $Ax = b$
  - (quasi-) orthogonal residual approaches: (Q)OR
  - (quasi-) minimal residual approaches: (Q)MR

Extensions:

- Lanczos based methods:
  - look-ahead
  - product-type (LTPMs)
  - applied to normal equations (CGN)
- Arnoldi based methods:
  - restart (thin/thick, explicit/implicit)
  - truncation (standard/optimal)

A Unified Matrix Description of Krylov Methods

Krylov methods as projection onto 'simpler' matrices:

- o $Q^H Q = I$, $Q^H A Q = H$ Hessenberg (Arnoldi),
- o $\hat{Q}^H Q = I$, $\hat{Q}^H A Q = T$ tridiagonal (Lanczos)

Introduce computed (condensed) matrix $C = T, H$

$$Q^{-1} A Q = C \quad \Rightarrow \quad AQ = QC$$

Iteration implied by unreduced Hessenberg structure:

$$AQ_k = Q_{k+1} \underline{C}_k, \quad Q_k = [q_1, \ldots, q_k], \quad \underline{C}_k \in \mathbb{K}^{(k+1)\times k}$$

Stewart: 'Krylov Decomposition'

Iteration spans Krylov subspace ($q = q_1$):

$$\text{span}\{Q_k\} = \mathcal{K}_k = \text{span}\{q, Aq, \ldots, A^{k-1} q\}$$

Perturbed Krylov Decompositions

A Krylov decomposition analogue holds true in finite precision:

$$\begin{aligned} AQ_k = Q_{k+1}\underline{C}_k - F_k &= Q_k C_k + q_{k+1} c_{k+1,k} e_k^T - F_k \\ &= Q_k C_k + M_k - F_k \end{aligned}$$

We have to investigate the impacts of the method on
- o the structure of the basis Q_k (local orthogonality/duality)
- o the structure of the computed C_k, $\underline{C}_k$
- o the size/structure of the error term $-F_k$

Convergence theory:
- o is usually based on inductively proven properties:
  orthogonality, bi-orthogonality, A-conjugacy, . . .

What can be said about these properties?

'Standard' error analysis:
- o splits into *forward* and *backward* error analysis.

Does this type of analysis apply to Krylov methods?

All methods fit pictorially into:

$$A\,Q_k - Q_k\,C_k = 0 - F_k.$$

This is a perturbed Krylov decomposition, as subspace equation.

Examination of the methods can be grouped according to

- o methods directly based on the Krylov decomposition
- o methods based on a split Krylov decomposition
- o LTPMs

The matrix C_k plays a crucial role:

- o C_k is Hessenberg or even tridiagonal (basics),
- o C_k may be blocked or banded (block Krylov methods),
- o C_k may have humps, spikes, ... (more sophisticated)

The error analysis and convergence theory splits further up:

- o knowledge on Hessenberg (tridiagonal) matrices
- o knowledge on orthogonality, duality, conjugacy, ...

We start with results on Hessenberg matrices.

A Short Excursion on Matrix Structure

Eigendecomposition of A:

$$AV = VJ_\Lambda$$

Left eigenmatrices:

$$\begin{aligned} \widehat{V} &\equiv V^{-H} & \quad\Rightarrow\quad & \widehat{V}^H A = J_\Lambda \widehat{V}^H \\ \check{V} &\equiv V^{-T} & \quad\Rightarrow\quad & \check{V}^T A = J_\Lambda \check{V}^T \end{aligned}$$

The adjoint of $\lambda I - A$ fulfils

$$\operatorname{adj}(\lambda I - A)(\lambda I - A) = \det(\lambda I - A) I \equiv \chi_A(\lambda) I.$$

Suppose that λ is not contained in the spectrum of A.

We form the resolvent $R(\lambda) = (\lambda I - A)^{-1}$ of λ and obtain

$$\operatorname{adj}(\lambda I - A) = \chi_A(\lambda) R(\lambda) = \chi_A(\lambda)\; V J_{\lambda-\Lambda}^{-1} \widehat{V}^H.$$

One shifted and inverted Jordan block:

$$J_{\lambda-\lambda_i}^{-1} = S_i E_i S_i \equiv S_i \begin{pmatrix} (\lambda-\lambda_i)^{-1} & (\lambda-\lambda_i)^{-2} & \dots & (\lambda-\lambda_i)^{-k} \\ & (\lambda-\lambda_i)^{-1} & & \\ & & \ddots & \vdots \\ & & & (\lambda-\lambda_i)^{-1} \end{pmatrix} S_i,$$

Observation: Terms with negative exponent cancel with factors in the characteristic polynomial $\chi_A(\lambda)$.

The resulting expression is a source of eigenvalue – eigenvector relations.

We express the adjugate with the aid of compound matrices,

$$\text{adj}\, A \equiv S\mathcal{C}_{n-1}(A^T)S.$$

Then we have equality

$$P \equiv \mathcal{C}_{n-1}(\lambda I - A^T) = (SVS)\, G\, (S\widehat{V}^H S) \equiv (SVS)\, \chi_A(\lambda)\, E\, (S\widehat{V}^H S).$$

The compound matrix P is composed of polynomials in λ:

$$p_{ij} = p_{ij}(\lambda; A) \equiv \det L_{ji}, \qquad \text{where} \qquad L \equiv \lambda I - A.$$

G is composed of rational functions in λ:

$$G = \chi_A(\lambda) \cdot (\oplus_i E_i)\,.$$

Since many terms cancel, the elements of G are polynomials.

We divide by maximal factor $(\lambda - \lambda_i)^\ell$ and compute the limes $\lambda \to \lambda_i$.

Observation: Only few elements of $\lim G$ are non-zero. Choice of eigenvectors based on non-zero positions i, j:

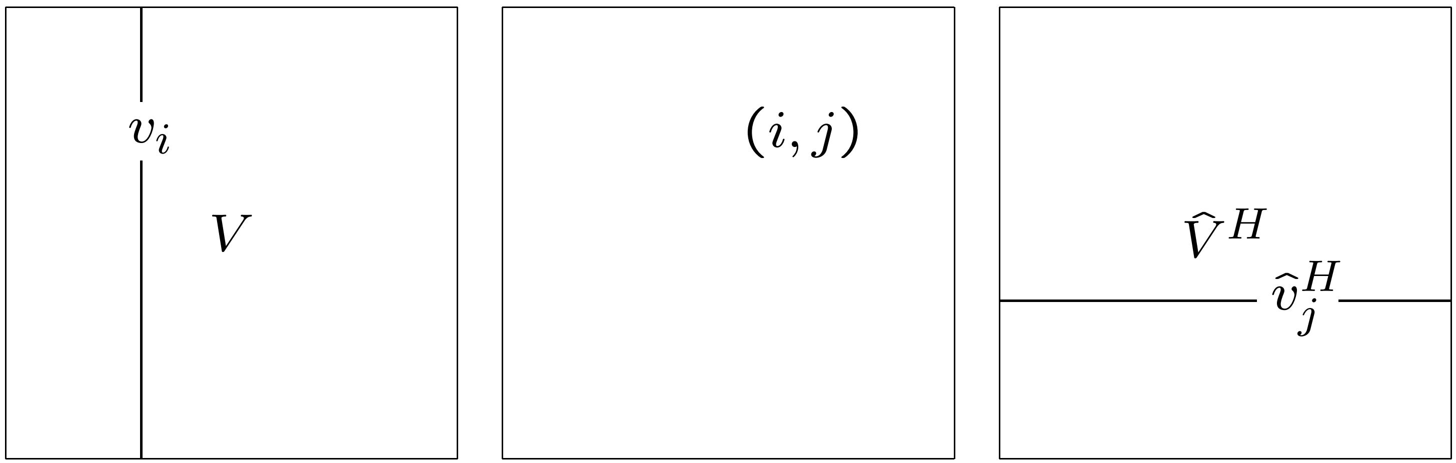

We consider here only the special case of non-derogatory eigenvalues.

We arrive at equations involving the elements of P, $\lim G$ and products of components of left and right eigenvectors.

Theorem: Let $A \in \mathbb{K}^{n\times n}$. Let $\lambda_l = \lambda_{l+1} = \ldots = \lambda_{l+k}$ be a geometrically simple eigenvalue of A. Let $k+1$ be the algebraic multiplicity of λ. Let $\widehat{v}_{l+k}^H$ and v_l be the corresponding left and right eigenvectors with appropriate normalization.

Then

$$v_{jl}\check{v}_{i,l+k} = (-1)^{(j+i+k)} \frac{p_{ji}(\lambda_l; A)}{\prod_{\lambda_s \neq \lambda_l} (\lambda_l - \lambda_s)}$$

holds true.

The sign matrices bear the blame for the minus one, P for the numerator and $\lim G$ for the denominator.

This setting matches every eigenvalue of non-derogatory A.

Unreduced Hessenberg matrices are non-derogatory matrices. This is easily seen by a simple rank argument. In the following let $H = H_m$ be unreduced Hessenberg of size $m \times m$,

$$\mathrm{rank}(H - \theta I) \geq m - 1.$$

Many polynomials can be evaluated in case of Hessenberg matrices:

Theorem: The polynomial p_{ji}, $i \leq j$ has degree $(i-1)+(m-j)$ and can be evaluated as follows:

$$\begin{aligned} p_{ji}(\theta; H) &= \begin{vmatrix} \theta I - H_{1:i-1} & & \star \\ & R_{i+1:j-1} & \\ 0 & & \theta I - H_{j+1:m} \end{vmatrix} \\ &= (-1)^{i+j} \chi_{H_{1:i-1}}(\theta) \prod \mathrm{diag}(H_{i:j}, -1) \chi_{H_{j+1:m}}(\theta). \end{aligned}$$

Denote by $\mathcal{H}(m)$ the set of unreduced Hessenberg matrices of size $m \times m$. The general result on eigenvalue – eigenvector relations can be simplified to read:

Theorem: Let $H \in \mathcal{H}(m)$. Let $i \leq j$. Let θ be an eigenvalue of H with multiplicity $k+1$. Let s be the unique left eigenvector and $\widehat{s}^H$ be the unique right eigenvector to eigenvalue θ.

Then

$$(-1)^k \, \check{s}(i) s(j) = \left[\frac{\chi_{H_{1:i-1}} \chi_{H_{j+1:m}}}{\chi_{H_{1:m}}^{(k+1)}} (\theta) \right] \prod_{l=i}^{j-1} h_{l+1,l} \qquad (1)$$

holds true.

Remark: We ignored the implicit scaling in the eigenvectors imposed by the choice of eigenvector-matrices, i.e. by $\check{S}^T S = I$.

Error Analysis Revisited

For simplicity we assume that the perturbed Krylov decomposition

$$M_k = AQ_k - Q_kC_k + F_k$$

is diagonalisable, i.e. that A and C_k are diagonalisable. Let $y_j \equiv Q_k s_j$.

Theorem: The recurrence of the basis vectors in eigenparts is given by

$$\hat{v}_i^H q_{k+1} = \frac{\left(\lambda_i - \theta_j\right) \hat{v}_i^H y_j + \hat{v}_i^H F_k s_j}{c_{k+1,k} s_{kj}} \quad \forall\ i, j(, k).$$

This *local error amplification formula* consists of four ingredients:

- o the left eigenpart of q_{k+1}: $\hat{v}_i^H q_{k+1}$,
- o a measure of convergence: $(\lambda_i - \theta_j)\hat{v}_i^H y_j$,
- o an error term: $\hat{v}_i^H F_k s_j$,
- o an amplification factor: $c_{k+1,k} s_{kj}$.

The formula depends on the *Ritz pair* of the actual step. Using the eigenvector basis we can get rid of the *Ritz vector*:

$$I = SS^{-1} = S\check{S}^T \quad \Rightarrow \quad e_l = S\check{S}^T e_l \equiv \sum_{j=1}^{k} \check{s}_{lj} s_j.$$

Theorem: The recurrence between vectors q_l and q_{k+1} is given by

$$\left[\sum_{j=1}^{k} \frac{c_{k+1,k}\check{s}_{lj}s_{kj}}{\lambda_i - \theta_j}\right] \hat{v}_i^H q_{k+1} = \hat{v}_i^H q_l + \hat{v}_i^H F_k \left[\sum_{j=1}^{k} \left(\frac{\check{s}_{lj}}{\lambda_i - \theta_j}\right) s_j\right].$$

For $l = 1$ we obtain a formula that reveals how the errors affect the recurrence from the beginning:

$$\left[\sum_{j=1}^{k} \frac{c_{k+1,k}\check{s}_{1j}s_{kj}}{\lambda_i - \theta_j}\right] \hat{v}_i^H q_{k+1} = \hat{v}_i^H q_1 + \hat{v}_i^H F_k \left[\sum_{j=1}^{k} \left(\frac{\check{s}_{1j}}{\lambda_i - \theta_j}\right) s_j\right].$$

Interpretation: The size of the deviation depends on the *size* of the *first component* of the *left* eigenvector $\widehat{s}_j$ of C_k and the *shape and size* of the *right* eigenvector s_j.

Next step: Application of the eigenvector – eigenvalue relation (1), (set $k = 1, i = 1, m = k, j = k$):

$$(-1)^k \,\check{s}(i) s(j) = \left[\frac{\chi_{C_{1:i-1}}\chi_{C_{j+1:m}}}{\chi^{(k+1)}_{C_{1:m}}}(\theta)\right] \prod_{l=i}^{j-1} c_{l+1,l}.$$

Theorem: The recurrence between basis vectors q_1 and q_{k+1} can be described by

$$\left[\sum_{j=1}^{k} \frac{\prod_{p=1}^{k} c_{p+1,p}}{\prod_{s\neq j}\left(\theta_s - \theta_j\right)\left(\lambda_i - \theta_j\right)}\right] \widehat{v}_i^H q_{k+1} = \widehat{v}_i^H q_1 + \widehat{v}_i^H F_k \left[\sum_{j=1}^{k}\left(\frac{\check{s}_{1j}}{\lambda_i - \theta_j}\right) s_j\right]$$

A result from polynomial interpolation theory (Lagrange):

$$\sum_{j=1}^{k} \frac{1}{\prod_{l\neq j}\left(\theta_j-\theta_l\right)\left(\lambda_i-\theta_j\right)} = \frac{1}{\chi_{C_k}(\lambda_i)} \sum_{j=1}^{k} \frac{\prod_{l\neq j}(\lambda_i-\theta_l)}{\prod_{l\neq j}\left(\theta_j-\theta_l\right)} = \frac{1}{\chi_{C_k}(\lambda_i)}$$

The following theorem holds true:

Theorem: The recurrence between basis vectors q_1 and q_{k+1} can be described by

$$\widehat{v}_i^H q_{k+1} = \frac{\chi_{C_k}(\lambda_i)}{\prod_{p=1}^{k} c_{p+1,p}} \left(\widehat{v}_i^H q_1 + \widehat{v}_i^H F_k \left[\sum_{j=1}^{k} \left(\frac{\check{s}_{1j}}{\lambda_i-\theta_j}\right) s_j\right]\right).$$

Similarly we can get rid of the eigenvectors s_j in the error term:

$$e_l^T \left[\sum_{j=1}^{k} \left(\frac{\breve{s}_{1j}}{\lambda_i - \theta_j} \right) s_j \right] = \sum_{j=1}^{k} \left(\frac{\breve{s}_{1j} s_{lj}}{\lambda_i - \theta_j} \right) = \frac{\prod_{p=1}^{l} c_{p+1,p} \chi_{C_{l+1:k}}(\lambda_i)}{\chi_{C_k}(\lambda_i)}$$

This results in the following theorem:

Theorem: The recurrence between basis vectors q_1 and q_{k+1} can be described by

$$\begin{aligned} \hat{v}_i^H q_{k+1} &= \frac{\chi_{C_k}(\lambda_i)}{\prod_{p=1}^{k} c_{p+1,p}} \left(\hat{v}_i^H q_1 + \hat{v}_i^H \sum_{l=1}^{k} \frac{\prod_{p=1}^{l} c_{p+1,p} \chi_{C_{l+1:k}}(\lambda_i)}{\chi_{C_k}(\lambda_i)} f_l \right) \\ &= \frac{\chi_{C_k}(\lambda_i)}{\prod_{p=1}^{k} c_{p+1,p}} \hat{v}_i^H q_1 + \sum_{l=1}^{k} \left(\frac{\chi_{C_{l+1:k}}(\lambda_i)}{\prod_{p=l+1}^{k} c_{p+1,p}} \hat{v}_i^H f_l \right). \end{aligned}$$

Multiplication by the right eigenvectors v_i and summation gives the familiar result:

Theorem: The recurrence of the basis vectors of a finite precision Krylov method can be described by

$$q_{k+1} = \frac{\chi_{C_k}(A)}{\prod_{p=1}^{k} c_{p+1,p}}\, q_1 + \sum_{l=1}^{k} \left(\frac{\chi_{C_{l+1:k}}(A)}{\prod_{p=l+1}^{k} c_{p+1,p}}\, f_l \right).$$

This result holds true even for non-diagonalisable matrices A, C_k.

The method can be interpreted as an *additive mixture* of several instances of the same method with several starting vectors.

A *severe deviation* occurs when one of the characteristic polynomials $\chi_{C_{l+1:k}}(A)$ becomes large compared to $\chi_{C_k}(A)$.

Open Questions

- Can Krylov methods be forward or backward stable?

- If so, which can?

- Are there any sets of matrices A for which Krylov methods are stable?

- Does the stability depend on the starting vector?

- Are there any a priori results on
  - the behaviour to be expected and
  - the rate of convergence?

$A \in \mathbb{R}^{100\times 100}$ normal, eigenvalues equidistant in $[0, 1]$.

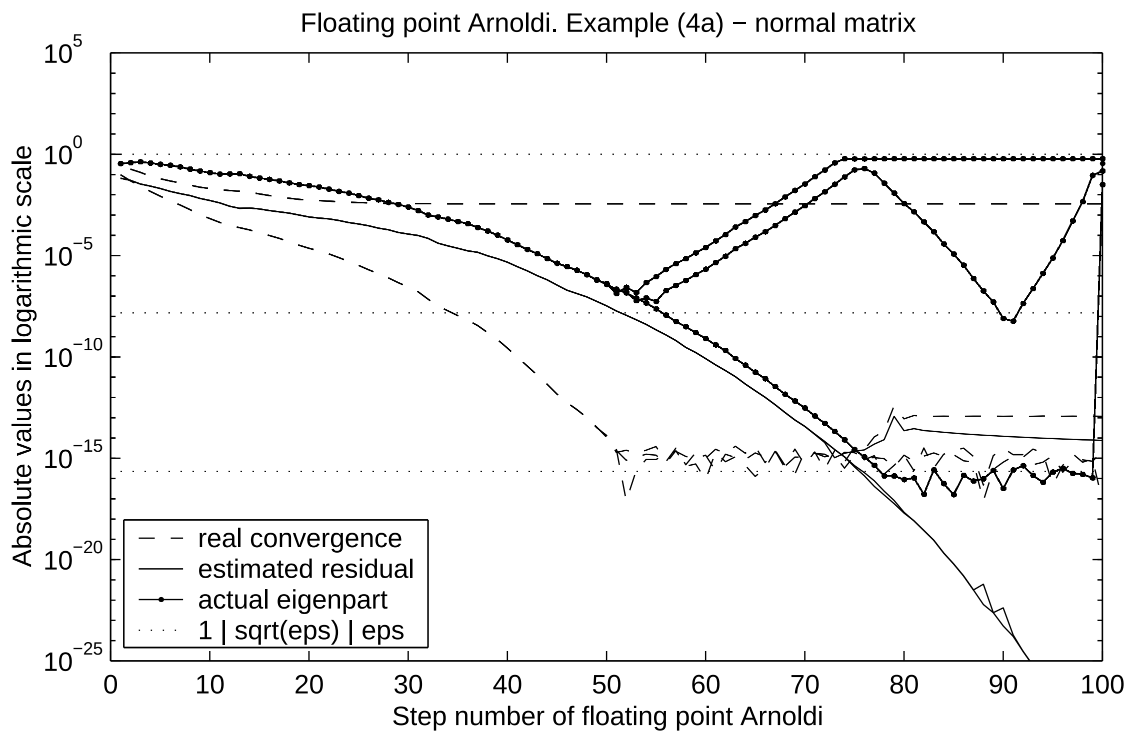

Behaviour of CGS-Arnoldi, MGS-Arnoldi, DO-Arnoldi, convergence to eigenvalue of largest modulus.

$A \in \mathbb{R}^{100\times 100}$ non-normal, eigenvalues equidistant in $[0, 1]$.

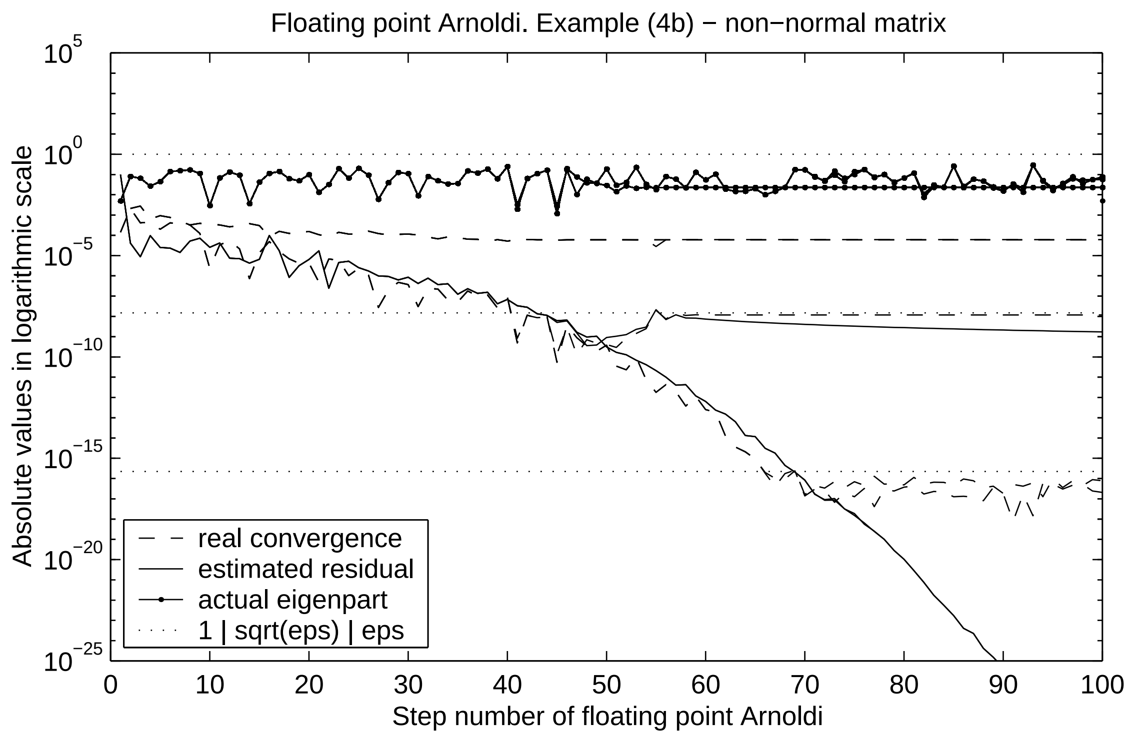

Behaviour of CGS-Arnoldi, MGS-Arnoldi, DO-Arnoldi, convergence to eigenvalue of largest modulus.

$A = A^T \in \mathbb{R}^{100\times 100}$, random entries in $[0, 1]$. Perron root well separated.

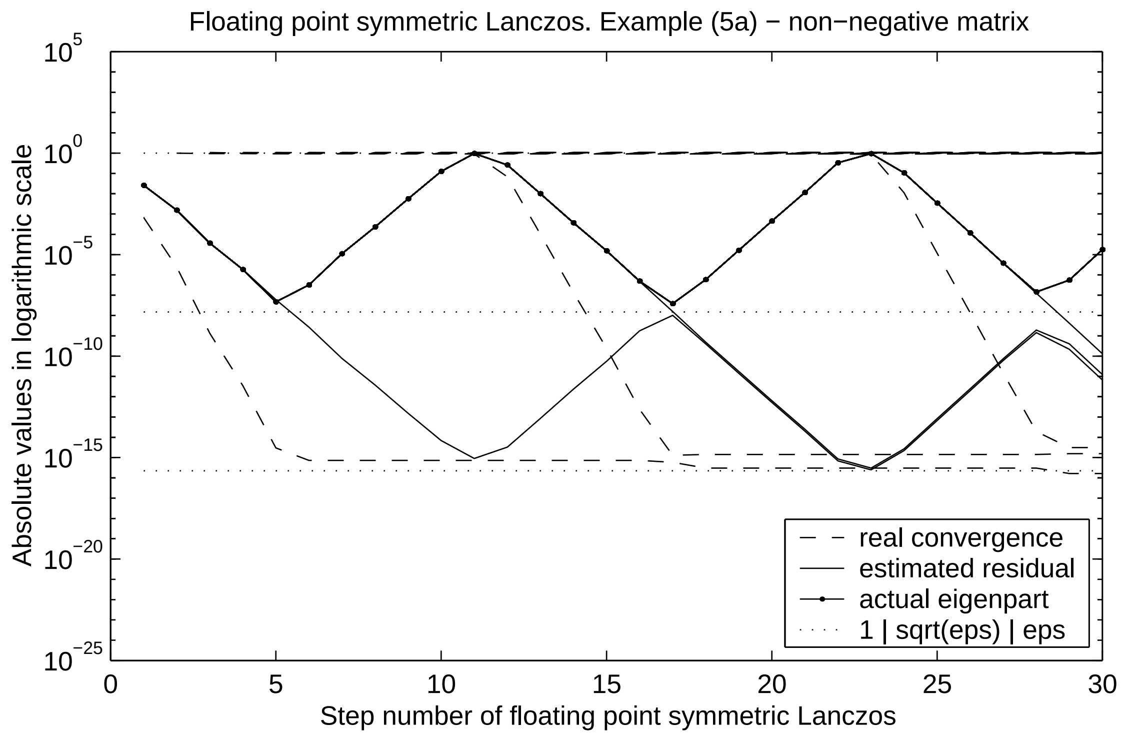

Behaviour of symmetric Lanczos, convergence to eigenvalue of largest modulus.

$A = A^T \in \mathbb{R}^{100\times 100}$, random entries in $[0,1]$. Perron root well separated.

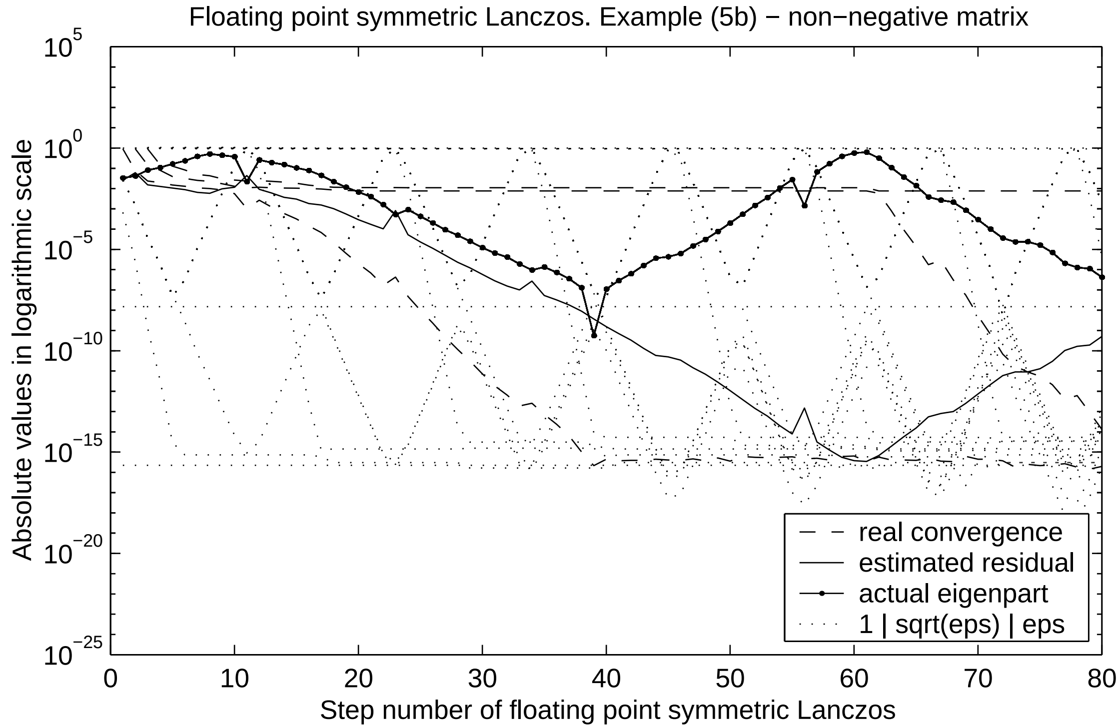

Behaviour of symmetric Lanczos, convergence to eigenvalue of largest and second largest modulus.