Krylov methods, an introduction

Paul Van Dooren, CESAME, Univ. Cath. Louvain, Belgium

slides on http://www.auto.ucl.ac.be/∼vdooren/Iterative.html

What we'll talk about ...

- basic ideas of iterative methods
  recursive refinement
- Krylov methods and their variants
  orthogonality vs bi-orthogonality
- some numerical aspects
  error propagation
- some algebraic aspects
  breakdowns
- eigenvalue problems
  projected eigenvalues
- rational approximation
  Padé approximation

Motivation

Every method performs better for some classes of problems ...

- Direct methods
- Jacobi/Gauss-Seidel
- Krylov methods
- Multigrid methods
- Fast multipole methods

but their features can be combined (hybrid, preconditioning)

Advantages of Krylov methods depend on whom to compare with

$\boxed{\textbf{Recurrences and Krylov methods}}$

Solve $Ax = b$ via fixed point of $x_k := b + (I - A)x_{k-1}$

Rewrite this as $x_k := x_{k-1} + r_{k-1}$ using residual $r_{k-1} := b - Ax_{k-1}$

$$\implies \quad x_k = x_0 + r_0 + r_1 + \ldots r_{k-1}$$

From $b - Ax_k = b - Ax_{k-1} - Ar_{k-1}$ we find $r_k := (I - A)r_{k-1}$

$$\implies \quad x_k = x_0 + \left(r_0 + (I - A)r_0 + \ldots (I - A)^{k-1} r_0\right)$$

$$= x_0 + \left[r_0, Ar_0, \ldots, A^{k-1} r_0\right] c$$

A Krylov subspace is a space spanned by

$$\mathcal{K}_k(A, r_0) := Im\left[r_0, Ar_0, \ldots, A^{k-1}r_0\right]$$

We are looking for "good" linear combinations

$$x_k - x_0 = \sum_{j=0}^{k-1} c_j A^j r_0 \in \mathcal{K}_k(A, r_0)$$

There are essentially two different criterions

$$\min \|Ax_k - b\|_2, \quad \Longleftrightarrow \quad \text{make } Ax_k - b \perp \mathcal{K}_k(A, r_0)$$

related to orthogonal recurrence relations (GMRES and FOM)

Two additional classes of methods are related to bi-orthogonal relations (QMR and BI-CG)

Arnoldi process

There always exists an orthogonal matrix $U^T U = I_n$ *such that*

$$U^T A U =: H = \begin{bmatrix} h_{1,1} & h_{1,2} & \dots & h_{1,n} \\ h_{2,1} & h_{2,2} & \ddots & \vdots \\ & \ddots & \ddots & h_{n-1,n} \\ & & h_{n,n-1} & h_{n,n} \end{bmatrix}$$

and the first column u_1 *of* U *can be chosen arbitrarily.*

Equate columns of $AU = UH$. First one: $Au_1 = u_1 h_{1,1} + u_2 h_{2,1}$

$$\Rightarrow \quad h_{1,1} := u_1^T A u_1, \quad \hat{u}_2 := Au_1 - u_1 h_{1,1}, \quad h_{2,1} := \|\hat{u}_2\|_2$$

For the following columns:

$$Au_k = \sum_{j=1}^{k} u_j h_{j,k} + u_{k+1} h_{k+1,k} \quad \Rightarrow \quad h_{j,k} := u_j^T A u_k$$

$$\hat{u}_{k+1} := Au_k - \sum_{j=1}^{k} u_j h_{j,k}, \quad h_{k+1,k} := \|\hat{u}_{k+1}\|_2, \quad u_{k+1} := \hat{u}_{k+1}/h_{k+1,k}$$

In block notation, with $U_k := U(:, 1:k), H_k := H(1:k, 1:k)$:

$$\boxed{A} \cdot \boxed{U_k} = \boxed{U_k} \cdot \boxed{H_k} + \hat{u}_{k+1} e_k^T$$

It is easy to see that $\mathcal{K}_k(H, e_1) = Im \begin{bmatrix} \times & \cdots & \times \\ & \ddots & \vdots \\ & & \times \\ & 0 & \end{bmatrix} = Im \begin{bmatrix} I_k \\ 0 \end{bmatrix}$

Choose U such that $r_0/\|r_0\|_2 = Ue_1,\ A = UHU^T$, then

$$\mathcal{K}_k(A, r_0) = Im\,[U_k]$$

because

$$\left[Ue_1, (UHU^T)Ue_1, \ldots, (UH^{k-1}U^T)Ue_1\right] = U\left[e_1, He_1, \ldots, H^{k-1}e_1\right]$$

Galerkin condition (FOM)

Look for $x_k - x_0 \in \mathcal{K}_k$ such that $b - Ax_k \perp \mathcal{K}_k$. Therefore

$$x_k - x_0 = U_k y, \quad U_k^T(b - Ax_k) = 0, \quad \Rightarrow \quad U_k^T A U_k y = U_k^T r_0 = \|r_0\|_2 e_1$$

So we solve using efficient recurrence relations

$$\begin{bmatrix} h_{1,1} & h_{1,2} & \cdots & h_{1,k} \\ h_{2,1} & h_{2,2} & \ddots & \vdots \\ & \ddots & \ddots & h_{k-1,k} \\ & & h_{k,k-1} & h_{k,k} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} = \begin{bmatrix} \|r_0\|_2 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Error bound $\|b - Ax_k\|_2 = |h_{k+1,k}.y_k|$ may grow ...

Minimize residual (GMRES)

Look for $x_k - x_0 \in \mathcal{K}_k$ to minimize $\|b - Ax_k\|_2$. Therefore

$$x_k - x_0 = U_k y, \quad b - Ax_k \in Im[U_{k+1}] \;\Rightarrow\; \|b - Ax_k\|_2 = \|U_{k+1}^T r_0 - U_{k+1}^T A U_k y\|_2$$

So we solve using efficient recurrence relations

$$\begin{bmatrix} h_{1,1} & h_{1,2} & \dots & h_{1,k} \\ h_{2,1} & h_{2,2} & \ddots & \vdots \\ & \ddots & \ddots & h_{k-1,k} \\ & & h_{k,k-1} & h_{k,k} \\ & & & h_{k+1,k} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} = \begin{bmatrix} \|r_0\|_2 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

One can prove that $\|b - Ax_k\|_2 \leq \|b - Ax_{k-1}\|_2$ but it may stall ...

$\boxed{\textbf{GMRES vs FOM}}$

Denote $\rho_k := \|b - Ax_k\|_2$ of FOM and GMRES by ρ_k^F and ρ_k^G, resp.

Then

$$(\rho_k^G)^{-2} = (\rho_k^F)^{-2} + (\rho_{k-1}^G)^{-2} \quad \Rightarrow \quad \rho_k^G \leq \rho_k^F$$

The Arnoldi process "breaks down" when $h_{k+1,k} = 0$ since we need to divide by it.

But then the system is solved since both FOM and GMRES yield the same answer and $\rho_k^G = \rho_k^F = |h_{k+1,k}.y_k| = 0$

$\boxed{\textbf{Stalling}}$

Consider

$$\begin{bmatrix} 0 & \dots & 0 & 1 \\ 1 & 0 & \ddots & 0 \\ & \ddots & \ddots & \vdots \\ & & 1 & 0 \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

Then

$$\rho^G = (1, \dots, 1, 0) \quad \Leftrightarrow \quad \rho^F = (\infty, \dots, \infty, 0)$$

GMRES is always bounded, outperforms FOM but still can stall

Lanczos process

For $A = A^T$ *there exists an orthogonal matrix* $U^T U = I_n$ *such that*

$$U^T A U =: T = \begin{bmatrix} \alpha_1 & \beta_2 & & \\ \beta_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_n \\ & & \beta_n & \alpha_n \end{bmatrix}$$

and the first column u_1 *of* U *can be chosen arbitrarily.*

Same derivation but recurrences are now short :

$$A u_k = u_{k-1}\beta_k + u_k \alpha_k + u_{k+1}\beta_{k+1} \quad \Rightarrow \quad \alpha_k := u_k^T A u_k$$

$$\hat{u}_{k+1} := A u_k - u_{k-1}\beta_k - u_k\alpha_k, \quad \beta_{k+1} := \|\hat{u}_{k+1}\|_2, \quad u_{k+1} := \hat{u}_{k+1}/\beta_{k+1}$$

For $A = A^T \succ 0$ the tri-diagonal matrix T can be factored, yielding two coupled 2-term recurrences instead (Conjugate Gradient).

Minimize residual (MINRES)

Look for $x_k - x_0 \in \mathcal{K}_k$ to minimize $\|b - Ax_k\|_2$ using efficient recurrence relations

$$\begin{bmatrix} \alpha_1 & \beta_2 & & \\ \beta_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_k \\ & & \beta_k & \alpha_k \\ & & & \alpha_{k+1} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} = \begin{bmatrix} \|r_0\|_2 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

One shows that $\|b - Ax_k\|_2$ decreases linearly with approximate factor $(\sqrt{\kappa} - 1)/(\sqrt{\kappa} + 1)$, where $\kappa := \|A\|\|A^{-1}\|$

The complexity of the different methods up to step k is

	Arnoldi	Lanczos	Conj G
Ax	k	k	k
orthog.	$2k^2n$	$9kn$	$10kn$
storage	kn	$3n$	$4n$

This clearly shows the need for other approaches for unsymmetric A if k gets large.

- Partial orthogonalization (IOM, ...)
- Restarts (FOM(m), GMRES(m), ...)
- Consider $A^TAx = A^Tb$ or $\begin{bmatrix} I & A \\ A^T & 0 \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}$

Unsymmetric Lanczos process

There exist invertible matrices V, W *such that* $W^T V = I_n$ *and*

$$W^T A V =: T = \begin{bmatrix} \alpha_1 & \beta_2 & & \\ \gamma_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_n \\ & & \gamma_n & \alpha_n \end{bmatrix}$$

for almost all first columns $v_1,\ w_1$ *of* $V,\ W$.

Derivation now involves columns of $AV = VT$ and $A^T W = W T^T$:

$$A v_k = v_{k-1}\beta_k + v_k \alpha_k + v_{k+1}\gamma_{k+1}$$

$$A^T w_k = w_{k-1}\gamma_k + w_k \alpha_k + w_{k+1}\beta_{k+1}$$

Without breakdowns we have with $T_k := T(1:k, 1:k)$, $V_k := V(:, 1:k), W_k := W(:, 1:k)$, in block notation:

$$\boxed{A} \cdot \boxed{V_k} = \boxed{V_k} \cdot \boxed{T_k} + \beta_{k+1} v_{k+1} e_k^T$$

$$\boxed{A^T} \cdot \boxed{W_k} = \boxed{W_k} \cdot \boxed{T_k^T} + \gamma_{k+1} w_{k+1} e_k^T$$

$$W_k^T A V_k = T_k, \quad W_k^T V_k = I_k \quad (\gamma_{k+1} \beta_{k+1} w_{k+1}^T v_{k+1} \neq 0)$$

Coupled Krylov subspaces

Since $\mathcal{K}_k(T, e_1) = Im \begin{bmatrix} I_k \\ 0 \end{bmatrix}$ and $\mathcal{K}_k(T^T, e_1) = Im \begin{bmatrix} I_k \\ 0 \end{bmatrix}$

we have that

$$\mathcal{K}_k(A, v_1) = Im\,[V_k]\,, \quad \mathcal{K}_k(A^T, w_1) = Im\,[W_k]$$

because

$$\left[Ve_1, (VTW^T)Ve_1, \ldots, (VT^{k-1}W^T)Ve_1\right] = V\left[e_1, Te_1, \ldots, T^{k-1}e_1\right]$$

$$\left[We_1, (WT^TV^T)We_1, \ldots, (W{T^T}^{k-1}V^T)We_1\right] = W\left[e_1, T^Te_1, \ldots, {T^T}^{k-1}e_1\right]$$

$\boxed{\textbf{Galerkin condition}}$

Look for $x_k - x_0 \in \mathcal{K}_k(A, r_0)$ such that $b - Ax_k \perp \mathcal{K}_k(A^T, w_1)$.
Therefore

$$x_k - x_0 = V_k y, \quad W_k^T(b - Ax_k) = 0, \quad \Rightarrow \quad W_k^T A V_k y = W_k^T r_0 = \|r_0\|_2 e_1$$

So we solve recursively

$$\begin{bmatrix} \alpha_1 & \beta_2 & & \\ \gamma_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_k \\ & & \gamma_k & \alpha_k \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} = \begin{bmatrix} \|r_0\|_2 \\ 0 \\ \vdots \\ 0 \end{bmatrix}$$

But now $\|W^T(b - Ax_k)\|_2 = |\gamma_{k+1}.y_k|$ ...

Minimize quasi residual (QMR)

Look for $x_k - x_0 \in \mathcal{K}_k(A, r_0)$ to minimize $\|W^T(b - Ax_k)\|_2$.

Therefore $x_k - x_0 = V_k y, \quad b - Ax_k \in Im[V_{k+1}]$,

$$\Rightarrow \|W^T(b - Ax_k)\|_2 = \|W_{k+1}^T r_0 - W_{k+1}^T A V_k y\|_2$$

So we solve using efficient recurrence relations

$$\begin{bmatrix} \alpha_1 & \beta_2 & & \\ \gamma_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_k \\ & & \gamma_k & \alpha_k \\ & & & \gamma_{k+1} \end{bmatrix} \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_k \end{bmatrix} = \begin{bmatrix} \|r_0\|_2 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

One can only prove that $\|W^T(b - Ax_k)\|_2 \le \|W^T(b - Ax_{k-1})\|_2$

Variants

- Avoid transposes via inner products $(x, A^T y) = (Ax, y)$
- Factorize $T = L.U$ to get coupled 2-term recurrences rather than 3-term recurrences
- Re-orthogonalize to compensate for loss of bi-orthogonality
- Apply look-ahead to avoid breakdowns
- Restart rather than look-ahead or re-orthogonalization
- Block versions for all of the above

Loss of orthogonality

The orthogonalization process $\hat{u}_{k+1} := Au_k - \sum_{j=1}^{k} u_j h_{j,k}$ yields under finite precision

$$\boxed{A} \cdot \boxed{U_k} = \boxed{U_k} \cdot \boxed{H_k} + \hat{u}_{k+1} e_k^T + \boxed{E_k}, \qquad \|E_k\| \approx \epsilon \|A\|$$

but the error $F_k := U_k^T U_k - I_k$ can grow arbitrarily (even with MGS) unless one orthogonalizes once again : $\|F_k\| \approx \epsilon$

The same comment holds for Lanczos (higher relative cost !)

For unsymmetric Lanczos bounds grow with $\|W_k\| \|V_k\|$

$\boxed{\textbf{Breakdowns}}$

Arnoldi “stops” when $h_{k+1,k} = 0$, implying $Ax_k = b$ (solved)

Lanczos “stops” when $\gamma_{k+1} = 0$, implying $Ax_k = b$ (solved)

when $\beta_{k+1} = 0$, choose $v_{k+1} \perp V_k$, e.g. $v_{k+1} = w_{k+1}$

when $w_{k+1}^T v_{k+1} = 0$, no bi-orthogonality, serious breakdown

Breakdown is “cured” by going to the block version (lookahead)

$$\det(W_{k+1}^T V_{k+1}) \neq 0$$

Recap solving $Ax = b$

Choose x_0, $\Rightarrow\ r_0 := b - Ax_0$ and look for $x_k - x_0 \in \mathcal{K}_k(A, r_0)$

Four different methods

	$Ax_k - b \perp \mathcal{K}_k(A, r_0)$	$\min \|Ax_k - b\|_2$
$O(k^2 n)$	FOM	GMRES
$O(kn)$	CG ($A = A^T \succ 0$)	MINRES ($A = A^T$)
	$Ax_k - b \perp \mathcal{K}_k(A^T, w_1)$	$\min \|W^T(Ax_k - b)\|_2$
$O(kn)$	BICG	QMR

Notice that with full reorthogonalization all methods are $O(k^2 n)$

Many variants try to cure orthogonality and erratic convergence

All methods improve with preconditioning (application-dependent)

$\boxed{\textbf{Eigenvalue problems}}$

Bauer-Fike Theorem applied to

$$AU_k - U_k H_k = h_{k+1,k} u_{k+1} e_k^T$$

yields (for $X^{-1}AX = \Lambda$)

$$\exists i : |\lambda_j(H_k) - \lambda_i(A)| \leq |h_{k+1,k}| \kappa(X)$$

i.e. each eigenvalue of H_k approximates "well" some eigenvalue of A

- For $A = A^T,\ \kappa(X) = 1$
- Breakdown $h_{k+1,k} = 0$ is good since $\lambda_j(H_k) = \lambda_i(A)$
- Improved bounds exist if we know something about $\Lambda(A)$

Bauer-Fike Theorem applied to

$$AV_k - V_k T_k = \hat{v}_{k+1} e_k^T$$

$$AW_k - W_k T_k^T = \hat{w}_{k+1} e_k^T$$

yields (for $X^{-1}AX = \Lambda$)

$$\exists i : |\lambda(T_k) - \lambda_i(A)| \leq \|V_k^\dagger\|_2 \|\hat{v}_{k+1}\|_2 \kappa(X), \|W_k^\dagger\|_2 \|\hat{w}_{k+1}\|_2 \kappa(X)$$

i.e. each eigenvalue of T_k approximates "well" some eigenvalue of A

- Breakdowns $\hat{v}_{k+1} = 0,\ \hat{w}_{k+1} = 0$ are good since $\lambda_j(T_k) = \lambda_i(A)$
- Breakdown $\hat{w}_{k+1}^T \hat{v}_{k+1} = 0$ does not help
- Result still holds when bi-orthogonality gets lost

Eigenvalue convergence

Note that Krylov spaces are related to the power method

$$\mathcal{K}_k(A, r) := Im\left[r, Ar, \ldots, A^{k-1}r\right]$$

and that under exact arithmetic

$$\mathcal{K}_k((A - cI), r) = \mathcal{K}_k(A, r)$$

For Arnoldi, $\Lambda(H_k)$ should converge to “outer” spectrum of $A - cI$

The same holds for Lanczos and $\Lambda(T_k)$ (real spectrum if $A = A^T$)

In practice one often converges to “outer” eigenvalues
but the full story is more complex ...

Convergence strongly influenced by rational transformation or implicit shift technique

Rational transformation $\hat{A} := (bI - cA)(dI - A)^{-1}$ transforms spectrum also (inverse iteration $\hat{A} := (dI - A)^{-1}$ is a special case)

$$\mathcal{K}_k(\hat{A}, r) := Im\left[r, \hat{A}r, \ldots, \hat{A}^{k-1}r\right]$$

tends to eigenspace with spectrum of A closest to d

- approximate inverses needed for multiplication with $\hat{A}$
- loss of orthogonality indicates convergence of eigenvalues
- symmetric Lanczos extensively studied

Implicit shifted ARPACK

One implicit shift QR step applied to H_k :

$$H_k - \mu I = Q.R, \quad \hat{H}_k = R.Q + \mu I$$

(and a deflation) deletes the undesired eigenvalue μ from $\hat{H}_k$ and yields a new starting vector $\hat{r} := (A - \mu I)r$ such that

$$\mathcal{K}_{k-1}(A, \hat{r}) := Im\left[\hat{r}, A\hat{r}, \ldots, A^{k-2}\hat{r}\right]$$

yields $\hat{H}_{k-1}$ as projected matrix

Choose shifts μ_i recursively

$$\hat{r} := [\Pi(A - \mu_i)]r = p(A)r$$

such that $p(s)$ filters the complex plane from undesired eigenvalues

Filter after a number of implicit steps

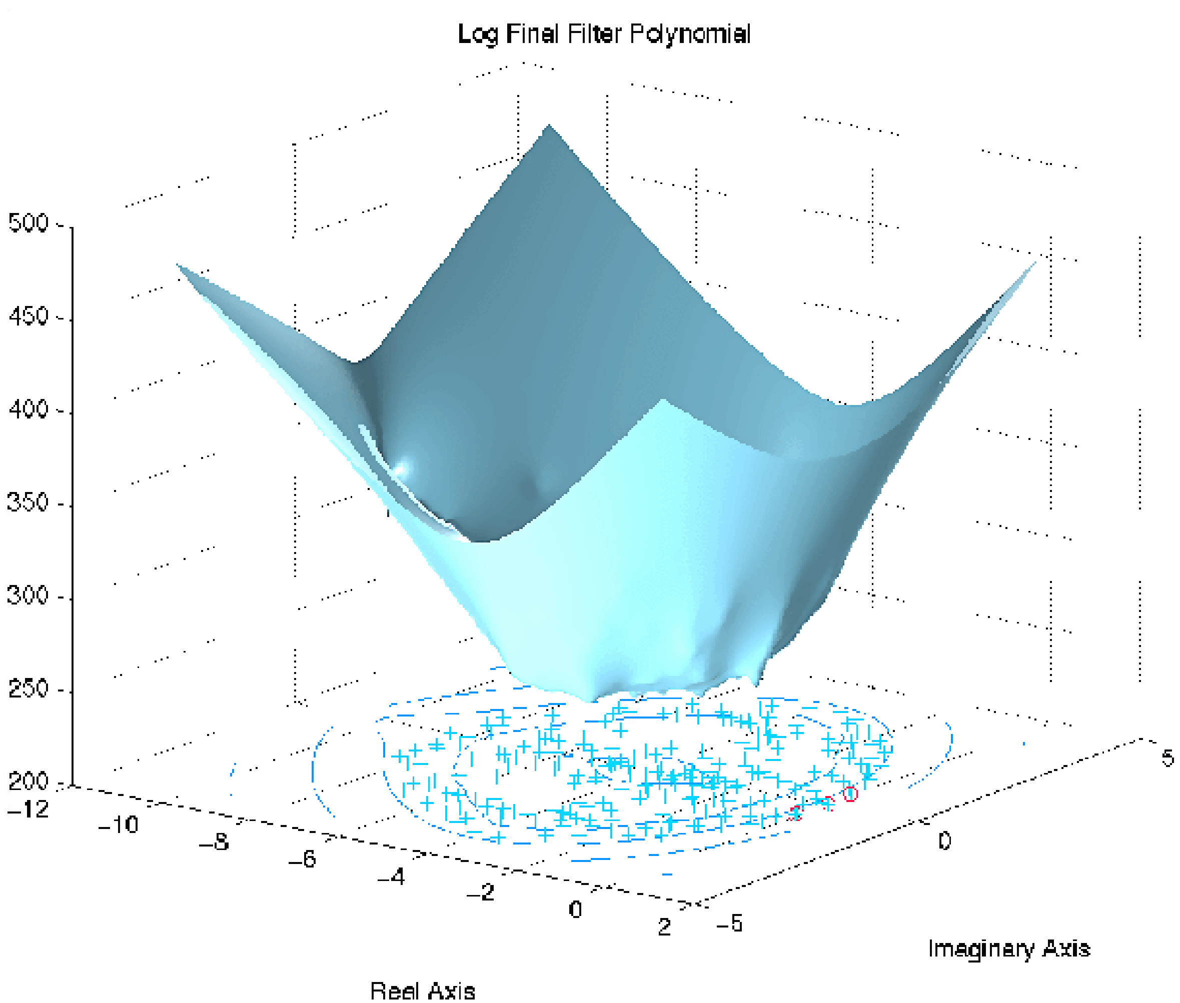

Implicit shifts ideas

- allow to keep k small (deflations)
- behaves numerically better than explicit calculation of $\hat{r}$
- extends to the Lanczos and unsymmetric Lanczos algorithm
- extends to the block Arnoldi and block Lanczos algorithm
- extends to the generalized eigenvalue problem

$$\det(\lambda B - A) = 0 \;\Leftrightarrow\; \det(\lambda I - B^{-1}A) = 0$$

implicit QR steps become implicit QZ steps

Rational approximation

Approximate $R(s)$ of degree N by $\hat{R}(s)$ of degree $n << N$. Assume

$$R(s) := c(\lambda I_N - A)^{-1}b, \quad \hat{R}(s) := \hat{c}(\lambda I_n - \hat{A})^{-1}\hat{b}$$

To interpolate at a point μ up to order $2k$ (Padé approximation):

$$R(s) - \hat{R}(s) = O(s-\mu)^{2k}$$

one chooses

$$\hat{A} = W_k A V_k, \quad \hat{b} = W_k b, \quad \hat{c} = c V_k,$$

where

$$\tilde{A} := (A - \mu I)^{-1}, \quad Im[V_k] = \mathcal{K}_k(\tilde{A}, \tilde{A}b), \quad Im[W_k] = \mathcal{K}_k(\tilde{A}^T, \tilde{A}^T c^T)$$

Also valid for multipoint, multi-input, multi-output and Arnoldi

We did NOT cover

- least squares variants
- implementation aspects
- preconditioning
- singular value problems
- block versions
- multiple right hand sides
- restarts ...

Conclusion

Krylov subspace algorithms are used to

- solve large scale systems of equations
- find eigenvalues, generalized eigenvalues and singular values of large scale matrix problems
- approximate the exponential or high degree rational functions by lower degree ones

There are numerous variants and sophisticated algorithms

It is still a very active area of research

References

[1] R. Lehoucq, D. Sorensen, C. Yang, ARPACK user's guide, *SIAM*, 1997.

[2] B. Parlett, The symmetric eigenvalue problem, *Prentice Hall*, 1980.

[3] Y. Saad, Iterative methods for sparse linear systems, *SIAM*, 2nd Ed., 2003.

[4] H. van der Vorst, Iterative Krylov Methods for Large Linear Systems, *Cambridge Monographs on Applied and Computational Mathematics*, 2003.

[5] P. Van Dooren, Gramian based model reduction of large-scale dynamical systems, in *Numerical Analysis*, Chapman and Hall, CRC Press, London, 2000.