Inexact range-space Krylov solvers for linear systems arising from inverse problems

Serge Gratton, Philippe Toint and Jean Tshimanga

Department of Mathematics, University of Namur, Belgium

(philippe.toint@fundp.ac.be)

Toulouse, February 2010

Motivation: data assimilation for weather forecasting

(Attempt to) predict...

- tomorrow's weather
- the ocean's average temperature next month
- future gravity field
- future currents in the ionosphere
- ...

Data assimilation for weather forecasting (2)

Data: temperature, wind, pressure, . . . everywhere and at all times!

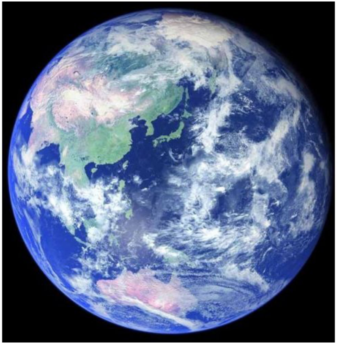

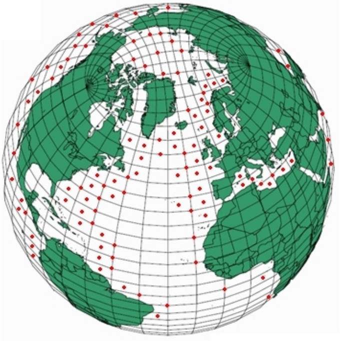

May involve up to 1,000,000,000 variables!

Data assimilation for weather forecasting (3)

The principle:

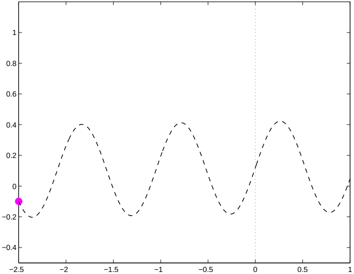

temp. vs. days

- Known situation 2.5 days ago and background prediction

Data assimilation for weather forecasting (3)

The principle:

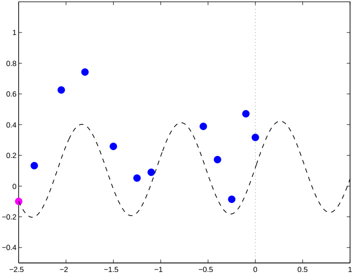

temp. vs. days

- Known situation 2.5 days ago and background prediction
- Record temperature for the past 2.5 days

Data assimilation for weather forecasting (3)

The principle:

Minimize deviation between model and past observations

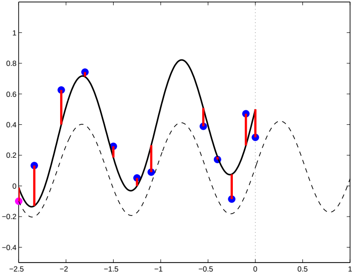

temp. vs. days

- Known situation 2.5 days ago and background prediction
- Record temperature for the past 2.5 days
- Run the model to minimize difference between model and observations

$$\min_{x_0} \frac{1}{2}\|x_0 - x_b\|^2_{B^{-1}} + \frac{1}{2}\sum_{i=0}^{N} \|\mathcal{H}\mathcal{M}(t_i, x_0) - b_i\|^2_{R_i^{-1}}.$$

Data assimilation for weather forecasting (3)

The principle:

Minimize deviation between model and past observations

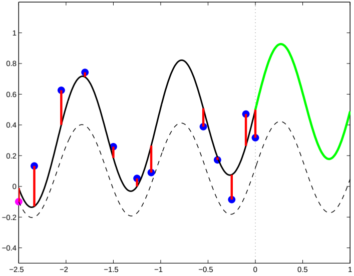

temp. vs. days

- Known situation 2.5 days ago and background prediction
- Record temperature for the past 2.5 days
- Run the model to minimize difference between model and observations
- Predict temperature for the next day

Data assimilation for weather forecasting (4)

Analysis of the ocean's heat content: CERFACS (2009)

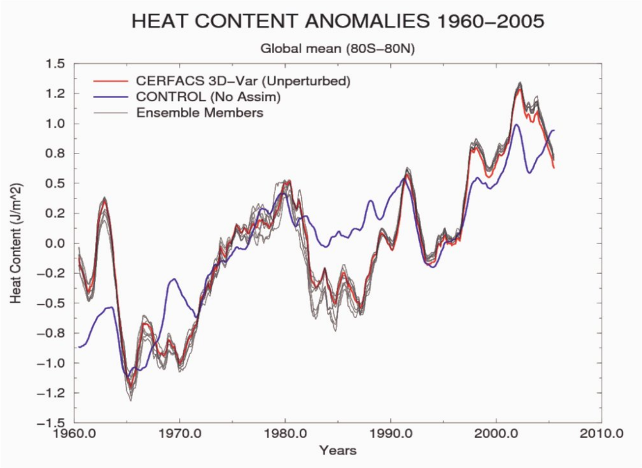

Much better fit!

Data assimilation problem: reformulations (1)

initial formulation:

$$\min_{x_0} \frac{1}{2}\|x_0 - x_b\|^2_{B^{-1}} + \frac{1}{2}\sum_{i=0}^{N} \|\mathcal{H}\mathcal{M}(t_i, x_0) - y_i\|^2_{R_i^{-1}}.$$

linearize, concatenate successive times and define $x_0 = x_s + s$:

$$\min_{x_0} \tfrac{1}{2}(x_s + s - x_b)^T B^{-1}(x_s + s - x_b) + \tfrac{1}{2}(Hs - d)^T R^{-1}(Hs - d)$$

write optimality conditions, using $c = x_b - x_s$:

$$(B^{-1} + H^T R^{-1} H)s = H^T d + B^{-1} c$$

Data assimilation problem: reformulations (2)

precondition using $z = B^{-1/2}s$ and :

$$\Big(I + \underbrace{B^{1/2}H^T R^{-1/2}}_{K^T}\underbrace{R^{-1/2}HB^{1/2}}_{K}\Big)\, z = \underbrace{B^{1/2}H^T R^{-1/2}}_{K^T} R^{-1/2}d + B^{-1/2}c$$

or

precondition using $z = B^{-1}s$:

$$\Big(I + \underbrace{H^T R^{-1}}_{K^T}\underbrace{HB^{-1}}_{L}\Big)\, z = \underbrace{H^T R^{-1}}_{K^T} d + B^{-1}c$$

In practice: use CG with reorthogonalization
(on problems where $n \approx 100,000$)...

The formal problem

Assume we now wish to solve

$$(\gamma I_n + K^T L)s = b$$

where $\gamma \neq 0$

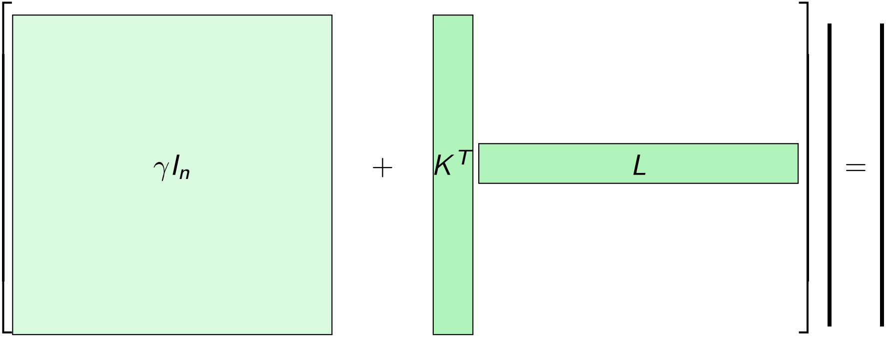

Note: We do not assume full-rank of K or L

The problem's sizes

But

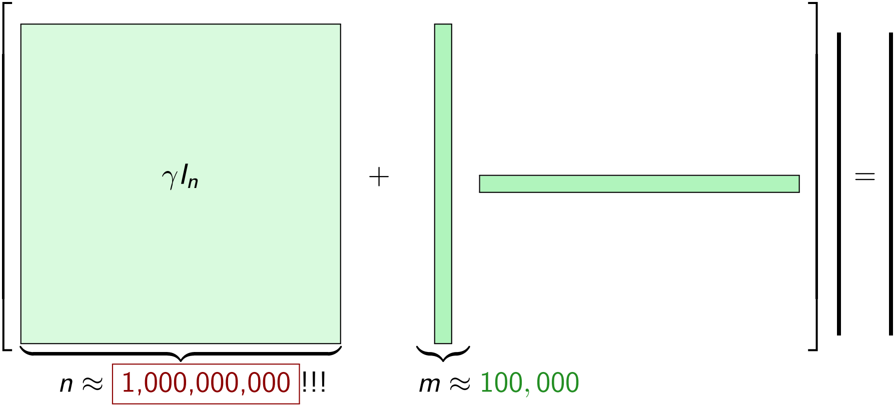

Wish to work in $\mathbb{R}^m$!

The standard GMRES for unsymmetric systems $Ax = b$

Based on the sequence of nested Krylov spaces:

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

Main idea:

At iteration k,

- build an orthonormal basis of $\mathcal{K}_k(A,b)$
- "solve" the problem in $\mathcal{K}_k(A,b)$ using this basis
- check for convergence?

+ get the solution in $\mathbf{R}^n$

"solve" may be:

- minimize the residual of the restricted problem $\Rightarrow$ GMRES
- solve a (small) system of linear equations $\Rightarrow$ FOM

GMRES for $Ax = b$ (2)

How to do that?

1. using $\mathcal{K}_{k-1}(A,b) \subset \mathcal{K}_k(A,b)$, incrementally build the basis of the span of

$$V_k = [v_1, v_2, \ldots, v_{k-1}, v_k] \quad \text{with} \quad V_k^T V_k = I_k$$

by

- computing Av_{k-1} (to create a new dimension)
- projecting this vector on $\mathcal{K}_{k-1}(A,b)^{\perp}$ and normalizing the result

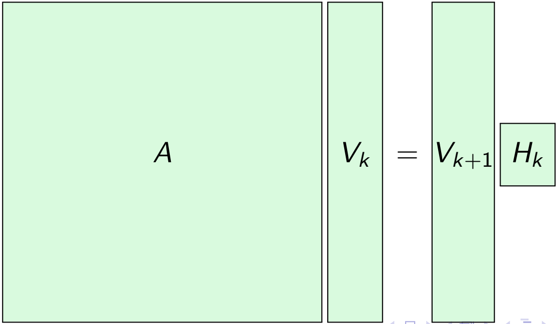

GMRES for $Ax = b$ (3)

How to do that?

2. Reduce the problem to $\mathcal{K}_k(A,b)$ (i.e. $x_k \in \mathcal{K}_k(A,b)$)

$$\| \underbrace{AV_k y_k - b}_{\text{size n}} \| \;=\; \| V_{k+1} H_k y_k - \beta V_{k+1} e_1 \| \;=\; \| \underbrace{H_k y_k - \beta e_1}_{\text{size } k} \|$$

Then solve

$$\min_y \| H_k y - \beta e_1 \| \;\rightarrow\; y_k \qquad \text{or} \qquad \text{solve}_y \;\; H_k^{\square} y = \beta e_1 \;\rightarrow\; y_k$$

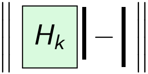

or

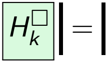

(minimum residual) (Galerkin)

(negligeable cost. . .)

GMRES for $Ax = b$ (4)

How to do that?

3. Test convergence: terminate if

$$\|H_k y_k - \beta e_1\| \leq \epsilon_A \quad \text{or} \quad \frac{\|H_k y_k - \beta e_1\|}{\|H_k\|\,\|y_k\| + \beta} \leq \epsilon_R$$

4. Reconstruct solution in $\mathbf{R}^n$:

$$x_k = V_k y_y$$

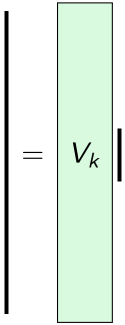

GMRES, FOM, MINRES and CG for $Ax = b$

- $\{\|r_k\|\}$ decreases monotonically, where $r_k = AV_k y_k - b$

(GMRES)

$$f_k = y_k^T V_k^T A V_k y_k - b^T V_k y_k \text{ decreases monotonically}$$

(FOM)

- Can be extended to exploit symmetry $\Rightarrow$ MINRES, CG

(in exact arithmetic)

- Performs well in practice, but high storage cost (V_k).

The standard GMRES algorithm

$s =$ **GMRES(** K**,** L**,** b **)**

1. Define $\beta_1 = \|b\|$ and $v_1 = b/\beta_1$.
2. For $k = 1, \ldots, m$,
   1. $w_k = K^T L v_k$
   2. for $i = 1, \ldots, k$,
      1. $H_{i,k} = v_i^T w_k$
      2. $w_k \leftarrow w_k - H_{i,k} v_i$
   3. $H_{k,k} \leftarrow H_{k,k} + \gamma$,
   4. $\beta_{k+1} = H_{k+1,k} = \|w_k\|$,
   5. $v_{k+1} = w_k / \beta_{k+1}$,
   6. $y_k = \arg\min_y \|Hy - \beta_1 e_1\|$,
   7. if $\|Hy_k - \beta_1 e_1\| < \epsilon$, break.
3. Return $s = V_k y_k$.

Range-space GMRES: the main idea

Return to the case of interest where

$$A = \gamma I_n + K^T L \quad \text{and} \quad b = K^T d.$$

Observe that

$$\text{span}_{i=0,\dots,k-1}\left[\left(\gamma I_n + K^T L\right)^i b\right] = \text{span}_{i=0,\dots,k-1}\left[\left(K^T L\right)^i b\right]$$

$$\begin{aligned}
\mathcal{K}_k(\gamma I_n + K^T L, b) &= \text{span}(b, K^T L b, \dots, (K^T L)^{k-1} b) \\
&= \text{span}(K^T d, K^T L K^T d, \dots, (K^T L)^{k-1} K^T d) \\
&= K^T \text{span}(d, L K^T d, \dots, (L K^T)^{k-1} d)
\end{aligned}$$

$$\mathcal{K}_k(\gamma I_n + K^T L, b) = K^T \mathcal{K}_k(L K^T, d)$$

(Gratton, Tshimanga for CG)

The range-space GMRES (1)

Main objectives

- all vectors now of size m! Factor K^T in the algorithm ($v = K^T \hat{v}$)
- good variational properties maintained
- need to compute norms in $\mathbb{R}^n$:

$$\|v\|^2 = \|K^T \hat{v}\|^2 = \hat{v}^T \underbrace{K K^T \hat{v}}_{\hat{z}} = \hat{v}^T \hat{z}$$

- store $\hat{V}_k$ and $\hat{Z}_k$ (but of size m)
- additional product by K to compute $\|v_k\|\ldots$

No free lunch. . . for the unsymmetric case

The range-space GMRES (2)

$s =$ **RSGMR0(** K**,** L**,** d **)**

1. Define $p_1 = K^T d, \quad \hat{z}_1 = K p_1,$
2. Set $\beta_1 = \sqrt{d^T \hat{z}_1}, \quad \hat{v}_1 = d/\beta_1 \quad \hat{z}_1 \leftarrow \hat{z}_1/\beta_1$ and $p_1 \leftarrow p_0/\beta_1.$
3. For $k = 1, \ldots, m,$
   1. $\hat{w}_k = L p_k$
   2. for $i = 1, \ldots, k,$
      1. $H_{i,k} = \hat{z}_i^T \hat{w}_k$
      2. $\hat{w}_k \leftarrow \hat{w}_k - H_{i,k} \hat{v}_i$
   3. $H_{k,k} \leftarrow H_{k,k} + \gamma,$
   4. $p_{k+1} = K^T \hat{w}_k, \quad \hat{z}_{k+1} = K p_k, \quad \beta_{k+1} = H_{k+1,k} = \sqrt{\hat{z}_{k+1}^T \hat{w}_k},$
   5. $\hat{v}_{k+1} \leftarrow \hat{w}_k/\beta_{k+1}, \quad \hat{z}_{k+1} \leftarrow \hat{z}_k/\beta_{k+1},, \quad p_{k+1} \leftarrow p_k/H_{k+1,k},$
   6. $y_k = \arg\min_y \|Hy - \beta_1 e_1\|,$
   7. if $\|Hy_k - \beta_1 e_1\| < \epsilon$, break.
4. Return $s = K^T \hat{V}_k y_k.$

The range-space GMRES (3)

If $b \notin \text{range}(K^T)\ldots$

- change K (and L)!

$$\overline{K} = \begin{bmatrix} K \\ b^T \end{bmatrix} \quad \text{and} \quad \overline{L} = \begin{bmatrix} L \\ 0^T \end{bmatrix}$$

and

$$\overline{K}^T \overline{L} = K^T L \quad \text{with} \quad \overline{K}^T e_{m+1} = b$$

- vectors of size $m+1$.

The range-space GMRES (4)

$s =$ **RSGMR(** K**,** L**,** b **)**

1. Define $\beta_1 = \|b\|$, $p_1 = b$, $u = Kb$, $\hat{z}_1 = u/\beta_1$, and $\hat{v}_1 = e_{m+1}/\beta_1$.
2. For $k = 1, \ldots, m+1$,
   1. $\hat{w}_k^T = [(Lp_k)^T \; 0]$, $\hat{w}_k \leftarrow \hat{w}_k/\beta_k$,
   2. for $i = 1, \ldots, k$,
      1. $H_{i,k} = [\hat{z}_i^T \; 0] \, \hat{w}_k$
      2. $\hat{w}_k \leftarrow \hat{w}_k - H_{i,k}\hat{v}_i$
   3. $H_{k,k} \leftarrow H_{k,k} + \gamma$,
   4. $p_{k+1} = [K^T \; b] \, \hat{w}_k$, $\hat{z}_{k+1} = Kp_{k+1}$, $\zeta_{k+1} = [u^T \; \beta_1^2] \, \hat{w}_k$,
   5. $\beta_{k+1} = H_{k+1,k} = \sqrt{[\hat{z}_{k+1}^T \; \zeta_{k+1}] \, \hat{w}_k}$,
   6. $\hat{v}_{k+1} \leftarrow \hat{w}_k/\beta_{k+1}$, $\hat{z}_{k+1} \leftarrow \hat{z}_k/\beta_{k+1}$,
   7. $y_k = \arg\min_y \|Hy - \beta_1 e_1\|$,
   8. if $\|Hy_k - \beta_1 e_1\| < \epsilon$, break.
3. Return $s = [K^T \; b]\hat{V}_k y_k$.

Full- vs range-space Krylov methods

At iteration k:

	GMRES	RSGMR
storage internal flops products by	$n(k+1)+k(k+3)/2$ $4nk+3n+[sol]$ K^T, L	$n+(2m+1)k+k(k+3)/2$ $4mk+7m+[sol]$ K^T, K, L
	FOM (sym)	RSFOM (sym)
storage internal flops products by	$n(k+1)+k(k+3)/2$ $4nk+3n+[sol]$ K^T, K	$(2m+1)k+k(k+3)/2$ $4mk+6m+[sol]$ K^T, K

Can we reduce cost further?

Inexact products: the context

Possible answer: inexact matrix-vector products

(Simoncini and Szyld, van den Eshof and Sleipen, Giraud, Gratton and Langou, . . .)

Motivations:

- stability wrt roundoff errors
  (remember iterates of RSGMR belong to range(K^T)!)
- allow cheap products (truncated B^{-1}, R^{-1}, simplified models,. . .)

Two error models for the result of $p \approx Av$:

1. Backward:

$$p = (A + E)v \quad \text{with} \quad \|E\| \leq \tau \|A\|$$

2. Forward:

$$p = Av + e \quad \text{with} \quad \|e\| \leq \tau \|Av\|.$$

Inexact products: results for the backward error model

Define

$$q_k = H_k y_k - \beta e_1, \quad G = \max\left[\, \|K\|, \|L\| \,\right] \quad \omega_k = \max_{1,\ldots,k} \|\hat{v}_i\|$$

$$\kappa(K) = \text{condition number of } K$$

(. . . after some analysis. . .)

Assume the backward error model. Then

$$\begin{aligned} \|r_k\| &\leq \sqrt{2(k+1)}\,\|q_k\| \\ &\qquad + \|K\|\omega_k\left[\tau_* \gamma \sqrt{k}\|y_k\| + 4\,G^2 \textstyle\sum_{i=1}^{k} |[y_k]_i|\,\tau_i\right] \\ &\leq \sqrt{2(k+1)}\left[\|q_k\| + \tau_{\max}\kappa(K)\,(\gamma + 4\,G^2)\|y_k\|\right]. \end{aligned}$$

Inexact products: results for the forward error model

Assume the forward error model. Then

$$\begin{aligned} \|r_k\| &\le \sqrt{2(k+1)}\,\|q_k\| + \sqrt{2}\left[\tau_*\gamma\sqrt{k}\|y_k\| + 4\,G\,\|K\|\sum_{i=1}^{k}|[y_k]_i|\,\tau_i\right] \\ &\le \sqrt{2(k+1)}\left[\|q_k\| + \tau_{\max}\,(\gamma + 4\,G\,\|K\|)\|y_k\|\right] \end{aligned}$$

Note in both sets of bounds:

- first of these bounds allow for variable accuracy requirements
- special role of τ_*

CG with inexact products

Is CG a reasonable framework for inexact products?

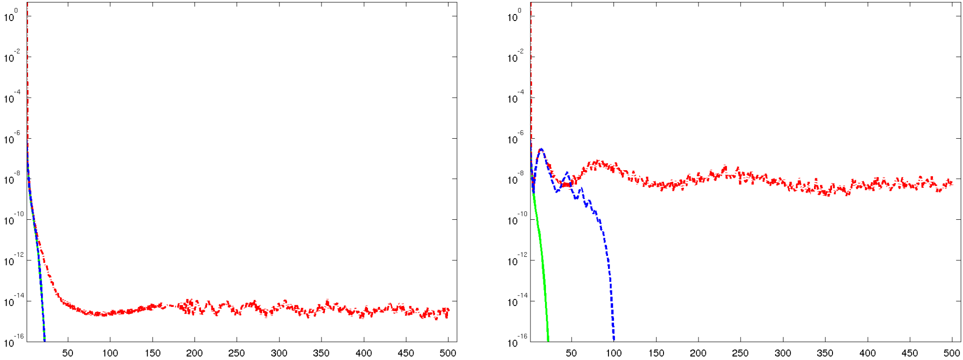

Comparing $\|r_k\|/(\|A\|\,\|s_*\|)$ for FOM, CG with reorthog and CG for exact(left) and inexact (right) products ($\tau = 10^{-9}$, $\kappa \approx 10^6$)

RSGMR and the error models (2)

Is the error model important? $(\epsilon = 10^{-5},\ \kappa \approx 10^{2})$

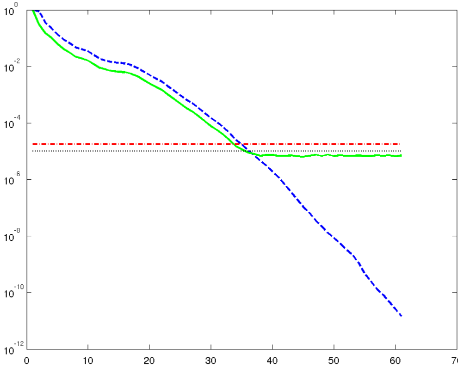

Backward error model

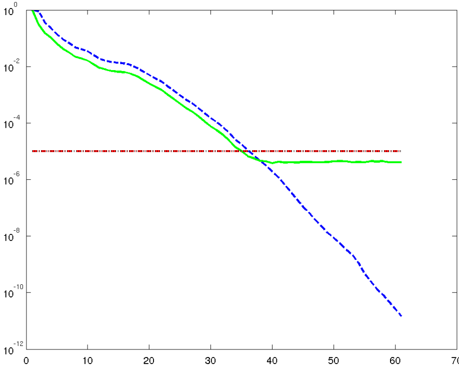

Forward error model

(normalized $\|r_k\|$, normalized $\|q_k\|$, accuracy threshold τ)

RSGMR and the error models (2)

Yes, it can definitely make a difference $(\epsilon = 10^{-5},\ \kappa \approx 10^{9})$

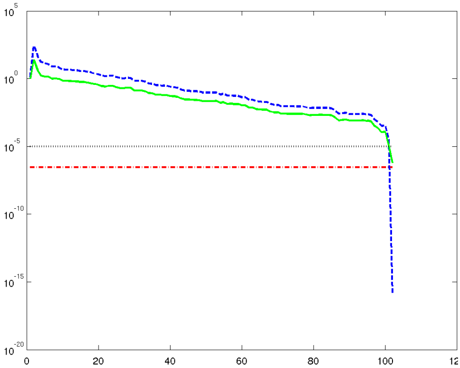

Backward error model

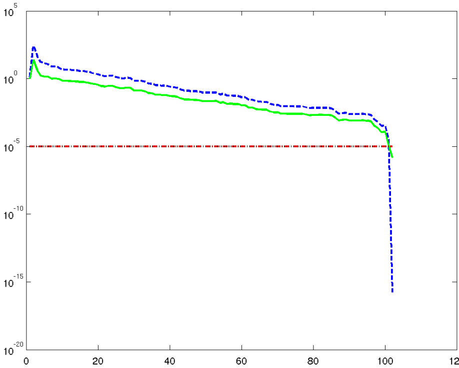

Forward error model

(normalized $\|r_k\|$, normalized $\|q_k\|$, accuracy threshold τ)

Fixed vs variable accuracy thresholds (1)

Can we use variable accuracy thresholds efficiently? $(\epsilon = 10^{-5},\ \kappa \approx 10^2)$

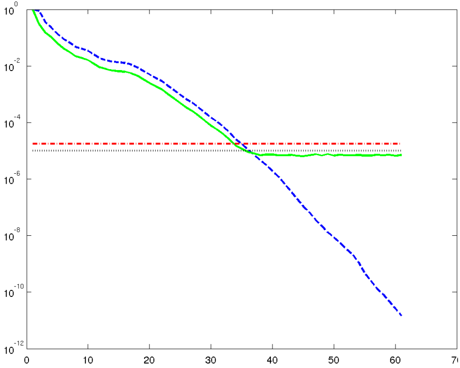

Fixed τ

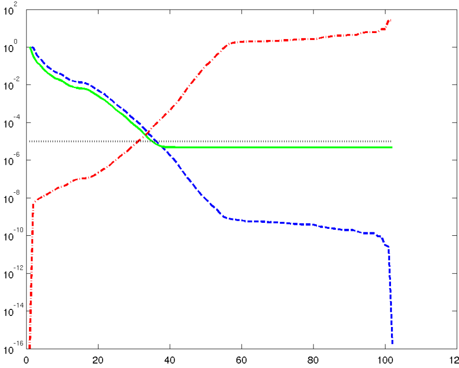

$\tau \approx 1/\|q_k\|$

(normalized $\|r_k\|$, normalized $\|q_k\|$, accuracy threshold τ)

Fixed vs variable accuracy thresholds (2)

Maybe. . . , not obvious. $(\epsilon = 10^{-5},\ \kappa \approx 10^{9})$

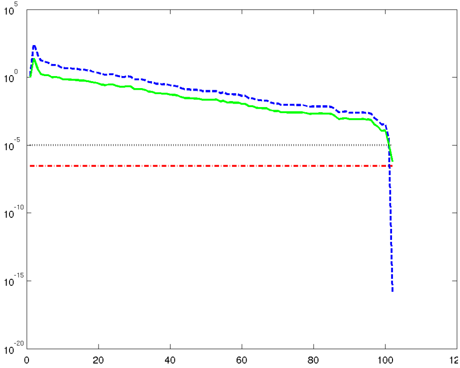

Fixed τ

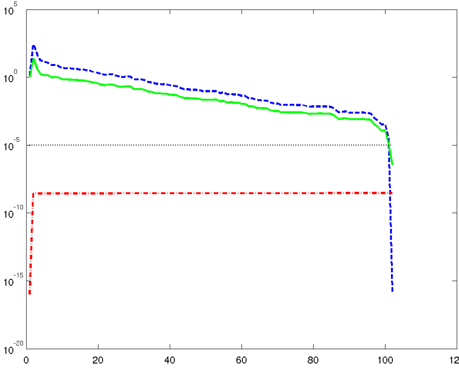

$\tau \approx 1/\|q_k\|$

(normalized $\|r_k\|$, normalized $\|q_k\|$, accuracy threshold τ)

Conclusions

- Range space methods may be designed to gain from low rank
- Further gains may be obtained from inexact products
- Formal bounds on the residual norms are available in this context
- Forward error modelling gives more flexibility than backward
- Many open questions . . . but very interesting
- Opens further doors for algorithm design:
  - efficiently spending one's "inaccuracy budget"
  - short recurrence methods
  - inexact full-space methods using forward error(?)
  - . . .
- True application: a real challenge
  (but we are working on it!)

Many thanks for your attention!