Some snapshots of numerical linear algebra and optimization

Michael Saunders

MS&E and ICME, Stanford University

Center for Science of Information (CSoI) seminar

Dept of Computer Science, Purdue University

West Lafayette, IN Sep 13, 2017

Abstract

Everyone has something they would like to minimize or maximize. Mathematical models and computer implementations give us the field of numerical optimization. Constraints reflecting physical reality require numerical linear algebra.

We review some of the software and aerospace applications associated with Philip Gill's contributions to numerical optimization. We then review the iterative methods CG, SYMMLQ, and MINRES for solving symmetric $Ax = b$ and show how SYMMLQ provides bounds on the 2-norm of the error for CG iterates.

Partially supported by the
National Institute of General Medical Sciences
of the National Institutes of Health (NIH)
Award U01GM102098

Part I

SIAM Optimization Conference

Vancouver BC, May 2017

Pre-history

Simplex via Cholesky

$B = LQ$, Q not kept, replace one col of B

$$LL^T \leftarrow LL^T + vv^T - ww^T$$
$$\text{or } LL^T - ww^T + vv^T$$

- Gill and Murray (1973) — A numerically stable form of the simplex method
- Saunders (1972) — Large-scale Linear Programming using the Cholesky Factorization

Perhaps a modern version is possible via LL^T or QR:

- Chen, Davis, Hager and Rajamanickam (2008) — Algorithm 887, CHOLMOD
  Supernodal Sparse Cholesky Factorization and Update/Downdate
- Davis (2011) — Algorithm 915, SuiteSparseQR
  Multifrontal multithreaded rank-revealing sparse QR factorization

LU is more sparse than LQ

Early linear programming systems assumed B is close to triangular.

- Helleman and Rarick (1971, 1972) The (partitioned) preassigned pivot procedure P^3, P^4
- Saunders (1976) The complexity of LU updating in the simplex method, MINOS

$$B = \begin{bmatrix} x & & & & & & & \\ & x & & & & & & \\ & x & x & * & & & & * \\ & x & x & x & & & & * \\ & & & x & x & & * & * \\ x & & & & x & x & * & * \\ & & x & x & x & x & x & \\ x & & & & x & x & x & x \end{bmatrix}$$

Markowitz LU is more sparse than P^4

Early 1980s: Rob Burchett, General Electric — Optimal Power Flow problem

Basis matrices were close to symmetric — Not good for P^4

Sparse LU with Markowitz merit function

- Duff and Reid (1977) — Fortran subroutines for sparse unsymmetric linear equations, MA28
- Reid (1982) — A sparsity-exploiting variant of the Bartels-Golub decomposition, LA05, LA15
- Gill, Murray, S, & Wright (1987) — Maintaining LU factors of a general sparse matrix, LUSOL

LUSOL does the linear algebra for MINOS, SQOPT, SNOPT, MILES, PATH, lp_solve

McDonnell Douglas

Huntington Beach, CA

SQP methods

NPSOL, SNOPT

Aerospace Applications of NPSOL and SNOPT

OTIS #1

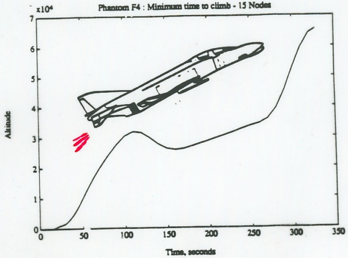

DC-Y single-stage-to-orbit

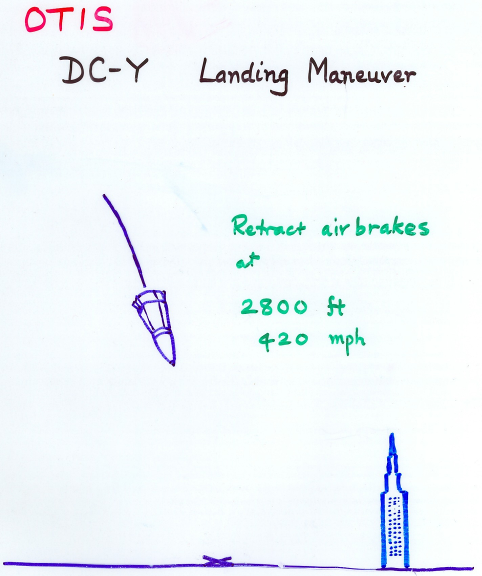
OTIS
DC-Y Landing Maneuver
Retract air brakes at
2800 ft
420 mph

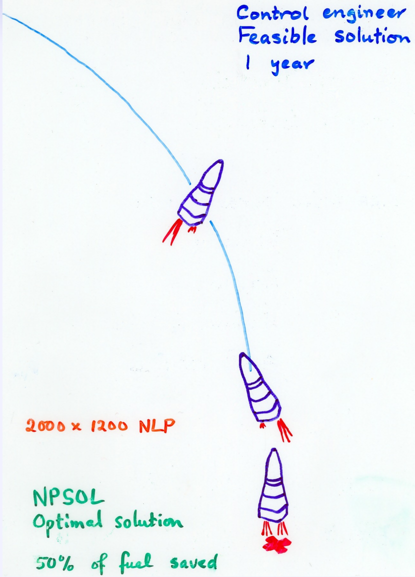
Control engineer
Feasible solution
1 year
2000 × 1200 NLP
NPSOL
Optimal solution
50% of fuel saved

DC-Y landing, 2nd OTIS/NPSOL optimization

- 1st optimization: starting altitude = 2800ft
- 2nd optimization: starting altitude = variable
- New constraint needed: Don't exceed 3g

Optimum starting altitude = 1400ft(!)

Come back Alan Shephard!

DC-X flying model
1/3 scale = 40ft tall

- 1993-95: DC-X made 8 flights
  Flight 8: demonstrated turnaround maneuver; hard landing damaged aeroshell
- 1996: DC-XA made 4 flights
  Flight 3: demonstrated 26-hour turnaround
  Flight 4: landing strut failed to extend; tipped over and exploded
- 1997: McDonnell Douglas merges with Boeing
  Huntington Beach campus becomes part of Boeing
  Philip continued 5-to-8 days for several years (till Rocky Nelson retired)

McDonnell Douglas motivation

The aerospace problems kept getting bigger

SQP needs Hessian H for QP subproblems and null-space operator Z for constraints

NPSOL

- dense quasi-Newton $H = R^T R$
- dense Z from $J^T = QR$

SNOPT

- limited-memory H
- Z from sparse $B = LU$ (reduced-gradient method)
- SQIC can switch from $B = LU$ to block-LU updates of K

Block-LU updates

Block-LU updates

- Bisschop and Meeraus (1977) Matrix augmentation and partitioning in updating the basis inverse
- Gill, Murray, S, and Wright (1984) Sparse matrix methods in optimization
- Eldersveld and S (1992) A block-LU update for large-scale linear programming, MINOS/SC
- Huynh (2008) A large-scale QP solver based on block-LU updates of the KKT system, QPBLU
- Maes (2010) Maintaining LU factors of a general sparse matrix, QPBLUR
- Wong (2011) Active-set methods for quadratic programming, icQP
- Gill and Wong (2014) Software for large-scale quadratic programming, SQIC
- Gill and Wong (2015) Methods for convex and general quadratic programming, SQIC

$$B_0 = L_0 U_0$$ LUSOL, BG updates

$$B_k \equiv \begin{pmatrix} B_0 & V_k \\ E_k & \end{pmatrix}$$ not implemented

$$K_0 = L_0 U_0$$ LUSOL, MA57, MA97, SuperLU, UMFPACK

$$K_k \equiv \begin{pmatrix} K_0 & V_k \\ V_k^T & D_k \end{pmatrix}$$

Quad Precision

"Carrying somewhat more precision in the arithmetic than twice the precision carried in the data and available for the result will vastly reduce embarrassment due to roundoff-induced anomalies."

"Default evaluation in Quad is the humane option."

— *William Kahan, 2011*

Double MINOS Quad MINOS

Double MINOS	Quad MINOS
`real(8)`	`real(16)`
`eps = 2.22e-16`	`eps = 1.93e-34`
Hardware	Software

We use this humane approach to Quad implementation

2 source codes	2 programs

snopt9 = Double or Quad SQOPT, SNOPT

snPrecision.f90

```
module snModulePrecision
  integer(4), parameter :: ip = 4, rp = 8    ! double
! integer(4), parameter :: ip = 8, rp = 16   ! quad
end module snModulePrecision
```

module sn50lp

```
use snModulePrecision, only :  ip, rp
subroutine s5solveLP ( x, y )
real(rp), intent(inout) ::  x(nb), y(nb)
```

1 source code	2 programs

DQQ procedure
for multiscale LP and NLP

Developed for systems biology models of metabolism

DQQ procedure

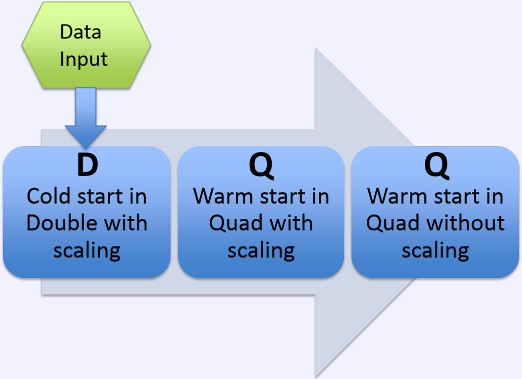

- Ding Ma, Laurence Yang, MS, et al. (2017) Reliable and efficient solution of genome-scale models of Metabolism and macromolecular Expression
  Double and Quad MINOS

Meszaros "problematic" LP test set

	Itns	Times	Final objective	Pinf	Dinf
gen1	369502	205.3	−1.6903658594e−08	−06	−12
	246428	9331.3	1.2935699163e−06	−12	−31
	2394	81.6	1.2953925804e−06	−**45**	−**30**
gen2	44073	60.0	3.2927907828e+00	−04	−11
	1599	359.9	3.2927907840e+00	–	−29
	0	10.4	3.2927907840e+00	–	−**32**
gen4	45369	212.4	1.5793970394e−07	−06	−10
	53849	14812.5	2.8932268196e−06	−12	−30
	37	10.4	2.8933064888e−06	−**54**	−**30**
l30	1229326	876.7	9.5266141574e−01	−10	−09
	275287	7507.1	−7.5190273434e−26	−25	−32
	0	0.2	−4.2586876849e−24	−**24**	−**33**

Pinf, Dinf $= \log_{10}$ Primal/Dual infeasibilities

Systems biology multiscale LP modes

	Itns	Times	Final objective	Pinf	Dinf
TMA_ME	12225	37.1	8.0051076669e−07	−06	−05
	685	61.5	8.7036315385e−07	−24	−30
	0	6.7	8.7036315385e−07	**–**	**−31**
GlcAerWT	62856	9707.3	−2.4489880182e+04	+04	−05
	5580	3995.6	−7.0382449681e+05	−07	−26
	4	60.1	−7.0382449681e+05	**−19**	**−21**
GlcAlift	134693	14552.8	−5.1613878666e+05	−03	−01
	3258	1067.1	−7.0434008750e+05	−09	−26
	2	48.1	−7.0434008750e+05	**−20**	**−22**

Final Pinf/ $\|x^*\|_\infty$ and Dinf/ $\|y^*\|_\infty$ are $\mathbf{O(10^{-30})}$

Quad NLP

Metabolic models and macromolecular expression (ME models)

Laurence Yang, UC San Diego

Quadratic convergence of major iterations (Robinson 1972)

```
     quadMINOS  5.6      (Nov 2014)

Begin tinyME-NLP   cold start NLP with mu = mu0
Itn      304 -- linear constraints satisfied.
Calling funcon.  mu =  0.80000000000000000000000000000000004
nnCon, nnJac, neJac        1073        1755        2681
funcon  sets    2681   out of    2681   constraint gradients.
funobj  sets       1   out of       1   objective  gradients.

Major minor  step     objective  Feasible Optimal   nsb   ncon penalty
    1   304T 0.0E+00  8.00000E-01 6.1E-03 2.1E+03     0      4 1.0E+02
    2   561T 1.0E+00  8.00000E-01 2.6E-14 3.2E-04     0     46 1.0E+02
    3    40T 1.0E+00  8.28869E-01 5.4E-05 3.6E-05     0     87 1.0E+02
    4      7 1.0E+00  8.46923E-01 1.2E-05 2.9E-06     0     96 1.0E+02
    5      0 1.0E+00  8.46948E-01 4.2E-10 2.6E-10     0     97 1.0E+02
    6      0 1.0E+00  8.46948E-01 7.9E-23 1.2E-20     0     98 1.0E+01

EXIT -- optimal solution found                              13.5 secs
```

Quasi-Newton optimization with finite-difference gradients

Design of computer experiments

Selden Crary, indie-physicist, 2015

$n = 11$ points (x_i, y_i) on $[-1, 1]$ square (one twin-point)

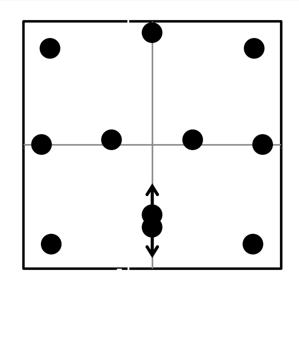

[d,n,p,theta1,theta2]=[2,11,2,0.128,0.069]

Design of computer experiments

Selden Crary, physicist, 2015

IMSPE-optimal designs (integrated mean-squared prediction error)

$$\min \ 1 - \text{trace}(B^{-1}A)$$

A and B: symmetric matrices of order $n+1$

B increasingly ill-conditioned if points approach each other

2D, Gaussian covariance parameters σ, θ_1, θ_2

$$A \propto \int_{-1}^{1}\int_{-1}^{1} \begin{bmatrix} 1 \\ v \end{bmatrix} \begin{bmatrix} 1 & v^T \end{bmatrix} dx\, dy \qquad B = \begin{bmatrix} 0 & \mathbf{1}^T \\ \mathbf{1} & V \end{bmatrix}$$

v_i functions of $\exp(\cdot)$ and $\text{erf}(\cdot)$, $\qquad V_{ij} = \sigma^2 e^{-\theta_1(x_i - x_j)^2 - \theta_2(y_i - y_j)^2}$

Pagoda plot of IMSPE function

Selden Crary, 2015

C^∞ a.e.

"Post", not pole

Need multistarts

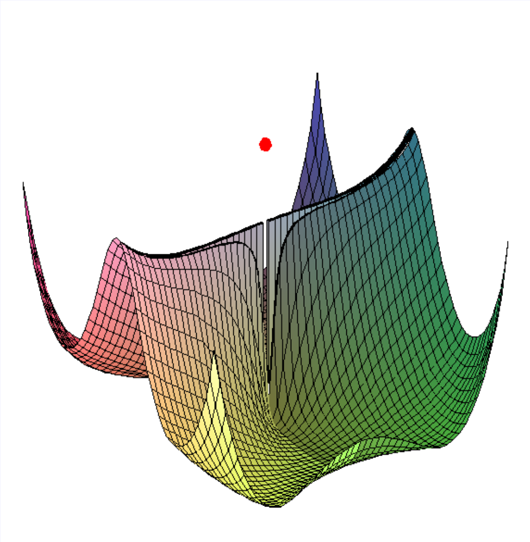

With Maple, Selden has found twin-points, triple-points, . . . and a new class of rational functions (the Nu class)

Smooth sailing through the black-hole/white-hole event horizon

No string theory

No complex numbers in quantum mechanics

. . .

IMSPE, 2D, $n = 11$, $\theta = (0.128, 0.069)$

Quad MINOS unconstrained optimization $\in \mathbb{R}^{22}$ without gradients 6 secs

```
 Itn ph pp      rg      step      objective     nobj  nsb  cond(H)
   1  4  0  3.9E-05  1.0E+03  2.47305090E-05    57   22  4.5E+01
  50  4  1  6.4E-07  6.2E+00  6.01181966E-06  1384   22  2.1E+04
 100  3  1 -7.6E-08  1.1E+00  5.65611811E-06  2726   22  1.4E+05
 150  4  1  8.5E-07  6.0E+00  5.11053080E-06  4102   22  8.2E+03
 200  4  1  2.6E-08  1.1E-01  5.02762155E-06  5464   22  1.0E+07
 239  4  1  1.1E-07  1.5E-06  5.02762154E-06  7478   22  1.0E+10
Search exit  7 -- too many functions.

EXIT -- optimal solution found

No. of iterations                  239   Objective value      5.0276215358E-06
No. of calls to funobj            7538   Calls with mode=2 (f, known g)    244
Calls for forward differencing    4466   Calls for central differencing   1716
Max Primal infeas           0 0.0E+00   Max Dual infeas              2 1.1E-07
```

not great ↑

17 points on $[-1, 1]^2$
(two twin-points)

Quad MINOS design
refined by Selden
via MAPLE

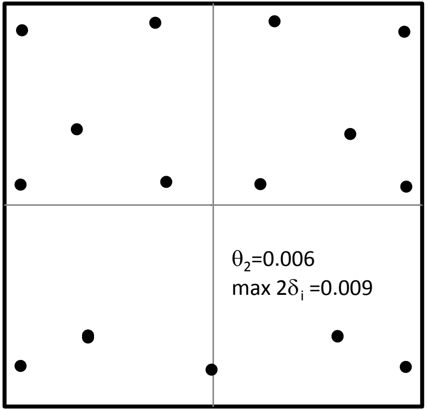

After further downhill searching, both sets of twins are now close, confirming the conjecture of a true double–twin-point design. 20161106

Linear algebra question

A, B real, symmetric, indefinite, ill-conditioned   min $\text{IMSPE} = 1 - \text{trace}(B^{-1}A)$

cf. GEV problem $Ax = \lambda Bx$ $\text{trace}(B^{-1}A) = \sum \lambda_i$

- QZ algorithm ignores symmetry but avoids ill-conditioned B^{-1}
- Will QZ compute real λ_i?

Yuji Nakatsukasa (Oxford) is developing `qdwhgep.m` for $Ax = \lambda Bx$ (real, symmetric)

- Congruence transformations are real $P^T APy = \lambda P^T BPy$
- Eigenvalues can be complex conjugate pairs
- $\text{trace}(B^{-1}A) = \sum \lambda_i$ will be real

PDCO in C++

Ron Estrin, UBC → Stanford

PDCO in C++

Matlab PDCO: regularized convex optimization ($D_1, D_2 \succ 0$, diagonal)

$$\begin{array}{ll} \underset{x,\,r}{\text{minimize}} & \phi(x) + \frac{1}{2}\|D_1 x\|^2 + \frac{1}{2}\|r\|^2 \\ \text{subject to} & Ax + D_2 r = b, \quad \ell \le x \le u, \end{array}$$

- Needed for huge metabolic LP models that are near-block-diagonal
- C++ has `double` and `float128` data-types
- Compiler generates multiple codes
- Switch from double to quad at run-time

1 source code 1 program

spring200

An optimal control problem modeling a spring/mass/damper

spring200

$$\min \; f(y,z,u) = \frac{1}{2}\sum_{t=0}^{T} z_t^2$$

$$y_{t+1} = y_t - 0.01y_t^2 - 0.004z_t + 0.2u_t$$

$$z_{t+1} = z_t + 0.2y_t \qquad t = 0,\dots,T-1$$

$$-1 \le y_t \qquad -0.2 \le u_t \le 0.2$$

$$y_0 = 0 \qquad y_T = 0 \qquad z_0 = 10$$

	Opt tol	Majors	Minors	Superbasics	Objective	Time
double	1e-06	13	576	18	1186.3839	0.05
quad	1e-15	31	1282	113	1186.3820	2.75

quad-MINOS gives an unexpectedly "clean" solution
(many variables exactly zero, including control variables u_t)

double-MINOS

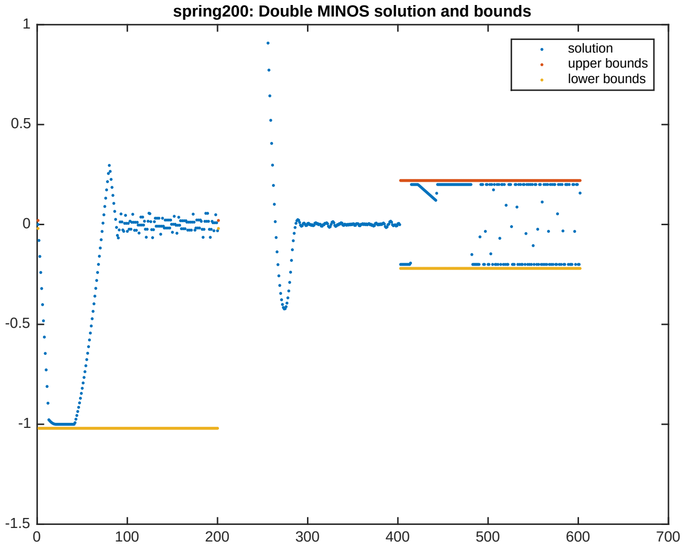

quad-MINOS

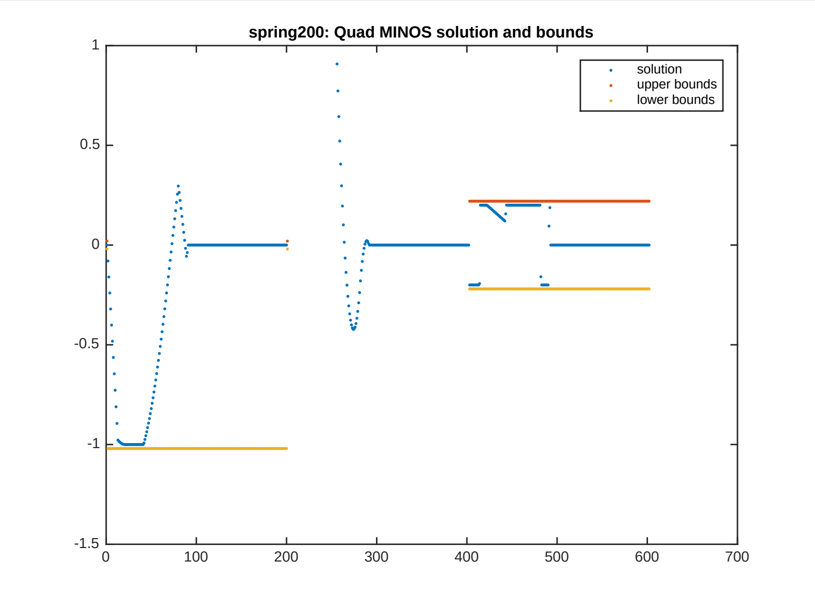

spring200: Quad MINOS solution and bounds
solution
upper bounds
lower bounds
1
0.5
0
-0.5
-1
-1.5
0
100
200
300
400
500
600
700

Part II

Householder Conference

Blacksberg VA, June 2017

Joint work with Ron Estrin and Dominique Orban

Lanczos for symmetric $Ax = b$

Error bounds for SYMMLQ and hence CG

Assume exact arithmetic

Check experimentally

Previous work

Error estimates for CG

Golub and Strakoš (1994)
Golub and Meurant (MMQ 1994, 1997)
Meurant (1997, 2005)
Brezinski (1999)
Frommer, Kahl, Lippert, and Rittich (2013)

Finite-precision analyses

- Strakoš and Tichý (2002)
  On error estimation in the CG method and why it works in finite precision computations
  ETNA 13
- Meurant (2006), The Lanczos and CG Algorithms: From Theory to Finite Precision Computations
  SIAM
- Greif, Paige, Titley-Peloquin, and Varah (2016)
  Numerical equivalences among Krylov subspace algorithms for skew-symmetric matrices
  SIMAX 37
- Paige (2017), Accuracy of the Lanczos process for the eigenproblem and solution of equations
  SIMAX soon (hot off the press!)

The Lanczos process for A, b

For $k = 1, 2, \ldots, \ell$

Lanczos generates $V_k = \begin{bmatrix} v_1 & v_2 & \ldots & v_k \end{bmatrix}$ and $\{\alpha_k, \beta_k > 0\}$ such that

$$\beta_1 v_1 = b \qquad AV_k = V_{k+1}\underline{T}_k \qquad \|v_k\| = 1 \qquad \beta_{\ell+1} = 0$$

$$\underline{T}_k = \left[\begin{array}{cccc} \alpha_1 & \beta_2 & & \\ \beta_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_k \\ & & \beta_k & \alpha_k \\ \hline & & & \beta_{k+1} \end{array}\right] = \begin{bmatrix} T_k \\ \cdots \; \times \end{bmatrix}$$

SYMMLQ, CG, MINRES for $Ax = b$

- $x_k = V_k y_k$
- $r_k = b - Ax_k = V_{k+1}(\beta_1 e_1 - \underline{T}_k y_k)$ — r_k small if $\underline{T}_k y_k \approx \beta_1 e_1$
- 3 ways to make r_k small — 3 subproblems for choosing y_k

$$\begin{bmatrix} \alpha_1 & \beta_2 & & \\ \beta_2 & \alpha_2 & \ddots & \\ & \ddots & \ddots & \beta_k \\ \hline & & \beta_k & \alpha_k \\ \hline & & & \beta_{k+1} \end{bmatrix} y_k \approx \begin{bmatrix} \beta_1 \\ 0 \\ \vdots \\ \hline 0 \\ \hline 0 \end{bmatrix}$$

SYMMLQ: $\min \|y_k\|$ st $\underline{T}_{k-1}^T y_k = \beta_1 e_1$

CG: $T_k y_k = \beta_1 e_1$

MINRES: $\underline{T}_k y_k \approx \beta_1 e_1$

SYMMLQ

$$\min \|y_k\| \;\; \textbf{st} \;\; \underline{T}_{k-1}^T y_k = \beta_1 e_1 \qquad \textbf{(then } x_k^L = V_k y_k\textbf{)}$$

$$\textbf{Needs } \underline{T}_{k-1}^T = \begin{bmatrix} L_{k-1} & 0 \end{bmatrix} Q_k$$

$$x_k^L = W_{k-1} z_{k-1} = x_{k-1}^L + \zeta_{k-1} w_{k-1}$$

moves in theoretically orthogonal directions

SYMMLQ recursions

$$\underline{T}_{k-1}^T Q_k^T = \begin{bmatrix} L_{k-1} & 0 \end{bmatrix} \qquad L_{k-1} z_{k-1} = \beta_1 e_1$$

$$T_k Q_k^T = \bar{L}_k = \begin{bmatrix} & L_{k-1} & & \\ 0 & \epsilon_k & \delta_k & \bar{\gamma}_k \end{bmatrix} \qquad \bar{L}_k \bar{z}_k = \beta_1 e_1$$

$$\bar{W}_k = V_k Q_k^T = \begin{bmatrix} W_{k-1} & \bar{w}_k \end{bmatrix} \qquad \bar{z}_k = \begin{bmatrix} z_{k-1} \\ \bar{\zeta}_k \end{bmatrix}$$

SYMMLQ recursions

$$\underline{T}_{k-1}^T Q_k^T = \begin{bmatrix} L_{k-1} & 0 \end{bmatrix} \qquad L_{k-1} z_{k-1} = \beta_1 e_1$$

$$T_k Q_k^T = \bar{L}_k = \begin{bmatrix} L_{k-1} & & & \\ 0 & \epsilon_k & \delta_k & \bar{\gamma}_k \end{bmatrix} \qquad \bar{L}_k \bar{z}_k = \beta_1 e_1$$

$$\bar{W}_k = V_k Q_k^T = \begin{bmatrix} W_{k-1} & \bar{w}_k \end{bmatrix} \qquad \bar{z}_k = \begin{bmatrix} z_{k-1} \\ \bar{\zeta}_k \end{bmatrix}$$

$$x_k^L = W_{k-1} z_{k-1} = x_{k-1}^L + \zeta_{k-1} w_{k-1}$$

$$x_k^C = \bar{W}_k \bar{z}_k \qquad = x_k^L + \bar{\zeta}_k \bar{w}_k$$

W_{k-1}, $\bar{W}_k$ theoretically have orthonormal columns

SYMMLQ error bound

$$x_k^L = W_{k-1} z_{k-1}, \qquad x_k^C = \bar{W}_k \bar{z}_k$$

W_{k-1}, $\bar{W}_k$ have theoretically orthonormal columns

$$\|x_k^L\|^2 = \|z_{k-1}\|^2 = \sum_1^{k-1} \zeta_j^2$$

$$\|x_*\|^2 = \|z_\ell\|^2 \quad = \sum_1^{\ell} \zeta_j^2$$

$$\|x_* - x_k^L\|^2 = \|x_*\|^2 - \|x_k^L\|^2$$

To bound the SYMMLQ error we need a bound on $\|x_*\|^2 = b^T A^{-2} b$

Bounding $\|x_*\|^2 = b^TA^{-2}b$

Needs Golub and Meurant

Golub and Meurant (1994, 1997)

Estimate bilinear forms $u^T f(A) v$ using Gaussian-quadrature theory

Theorem

For SPD A and suitable f, fix $\lambda_{\mathrm{est}} \in (0, \lambda_{\min}(A))$ and choose ω_k such that $\widetilde{T}_k = \begin{bmatrix} T_{k-1} & \beta_k e_{k-1} \\ \beta_k e_{k-1}^T & \omega_k \end{bmatrix}$, $\lambda_{\min}(\widetilde{T}_k) = \lambda_{\mathrm{est}}$. Then $b^T f(A) b \leq \|b\|^2 \, e_1^T f(\widetilde{T}_k) e_1$.

$$f(\xi) = \xi^{-2} \text{ gives } \|x_*\|^2 = b^T A^{-2} b \leq \|b\|^2 \, e_1^T \widetilde{T}_k^{-2} e_1$$

Golub and Meurant (1994, 1997)

Estimate bilinear forms $u^T f(A)v$ using Gaussian-quadrature theory

Theorem

For SPD A and suitable f, fix $\lambda_{\text{est}} \in (0, \lambda_{\min}(A))$ and choose ω_k such that
$\widetilde{T}_k = \begin{bmatrix} T_{k-1} & \beta_k e_{k-1} \\ \beta_k e_{k-1}^T & \omega_k \end{bmatrix}$, $\lambda_{\min}(\widetilde{T}_k) = \lambda_{\text{est}}$. Then $b^T f(A) b \leq \|b\|^2 \, e_1^T f(\widetilde{T}_k) e_1$.

$$f(\xi) = \xi^{-2} \text{ gives } \|x_*\|^2 = b^T A^{-2} b \leq \|b\|^2 \, e_1^T \, \widetilde{T}_k^{-2} e_1$$

Theorem

$\omega_k = \lambda_{\text{est}} + \eta$, where η is last entry solution of $(T_{k-1} - \lambda_{\text{est}} I) u_{k-1} = \beta_k^2 e_{k-1}$.

QR on $(T_{k-1} - \lambda_{\text{est}} I)$ gives η, ω_k LQ on $\widetilde{T}_k$ gives $\|b\|^2 \, e_1^T \, \widetilde{T}_k^{-2} e_1$

Computing $\beta_1^2 e_1^T \widetilde{T}_k^{-2} e_1$

$\widetilde{T}_k = \widetilde{L}_k \widetilde{Q}_k$ is almost the same as $T_k = \bar{L}_k Q_k$.

- Solve $\widetilde{L}_k \widetilde{z}_k = \beta_1 e_1$ to get $\widetilde{z}_k = \begin{bmatrix} z_{k-1} \\ \widetilde{\zeta}_k \end{bmatrix}$
- $\|x_*\|^2 \leq \beta_1^2 e_1^T \widetilde{T}_k^{-2} e_1 = \|\beta_1 \widetilde{L}_k^{-1} e_1\|^2 = \|\widetilde{z}_k\|^2$
- We already solve $L_{k-1} z_{k-1} = \beta_1 e_1$ and have $\|x_k^L\|^2 = \|z_{k-1}\|^2$

Hence

$$\begin{aligned} \|x_* - x_k^L\|^2 &= \|x_*\|^2 - \|x_k^L\|^2 \\ &\leq \|\widetilde{z}_k\|^2 - \|z_{k-1}\|^2 = \widetilde{\zeta}_k^2 \end{aligned}$$

and we can bound the SYMMLQ error in $O(1)$ work per iteration:

$$\|x_* - x_k^L\| \leq \epsilon_k^L \equiv |\widetilde{\zeta}_k|$$

CG error $\leq$ SYMMLQ error

Theorem (Estrin, Orban, and S. 2017)

For positive-semidefinite consistent $Ax = b$,

$$\|x_k^L\| \leq \|x_k^C\|$$
$$\|x_* - x_k^C\| \leq \|x_* - x_k^L\|$$

Immediate consequence: $\|x_* - x_k^C\| \leq \|x_* - x_k^L\| \leq \epsilon_k^L$

Better bound: $\|x_* - x_k^C\| \leq \epsilon_k^C := \sqrt{(\epsilon_k^L)^2 - \bar{\zeta}_k^2}$ $\quad (x_k^C = x_k^L + \bar{\zeta}_k \bar{w}_k)$

Golub and Strakoš (1994)

- Store $x_k^C, \ldots, x_{k+d}^C$ for moderate values of d (sliding window approach)
- Take $x_{k+d} \approx x_*$ for some purposes

Lemma (Estrin, Orban, and S. 2017)

$$\theta_k := x_*^T x_k^C - \|x_k^C\|^2 \geq 0$$

$$\|x_* - x_k^C\| \leq \sqrt{(\epsilon_k^C)^2 - 2\theta_k} \qquad \text{(not computable)}$$

$$\theta_k^{(d)} := (x_{k+d})^T x_k^C - \|x_k^C\|^2 \leq \theta_k$$

Better bound: $$\|x_* - x_k^C\| \leq \sqrt{(\epsilon_k^C)^2 - 2\theta_k^{(d)}}$$

Summary so far

- For SPD $Ax = b$, we derived upper bounds on the SYMMLQ and CG errors, assuming exact arithmetic. The results hold if $Ax = b$ is semidefinite and consistent.
- Numerical experiments show the bounds hold until convergence, but rigorous finite-precision analysis is desirable.
- If A is indefinite, the SYMMLQ upper bound becomes an estimate. Could obtain a bound by treating $b^T A^{-2} b$ as a quadratic form in A^2, but this is expensive (2 applications of A per iteration).

Numerical results

SYMMLQ error for UTEP/Dubcova1 $n = 16129$ SPD $\kappa(A) = 10^3$

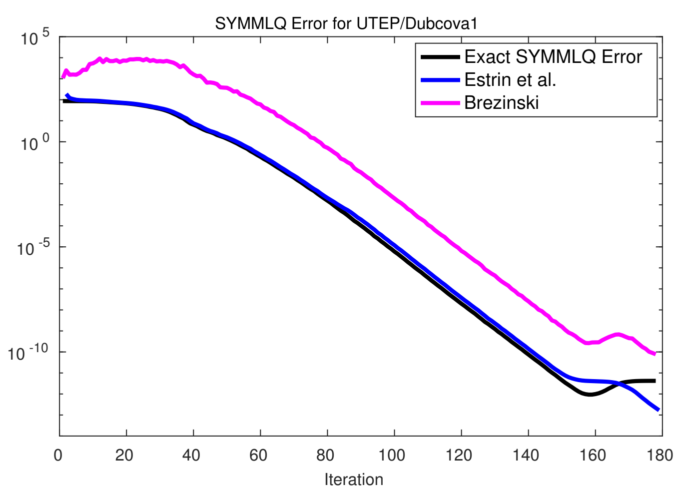

CG error for UTEP/Dubcova1 $n = 16129$ SPD $\kappa(A) = 10^3$

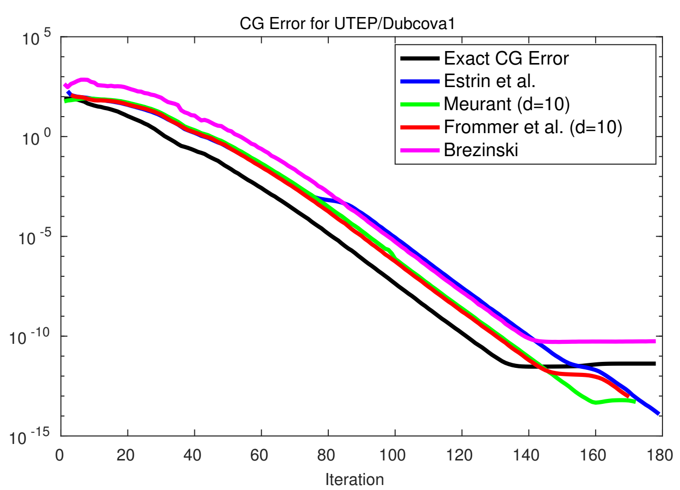

CG error for Nasa/nasa4704 $n = 4704$ SPD $\kappa(A) = 10^7$ $d = 10$

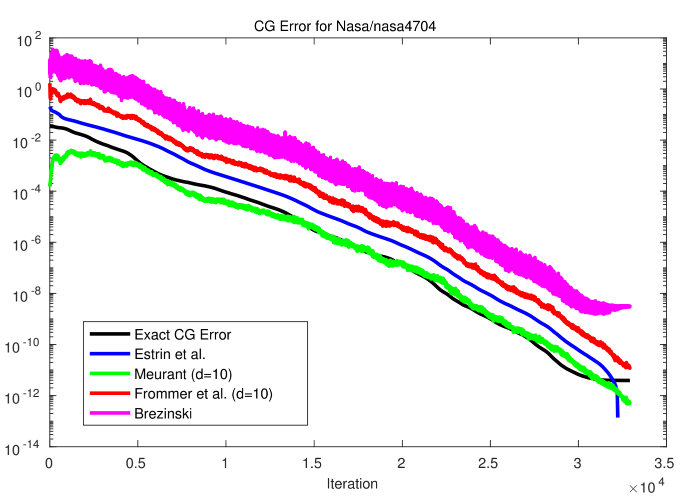

CG error for Nasa/nasa4704 $n = 4704$ SPD $\kappa(A) = 10^7$ $d = 100$

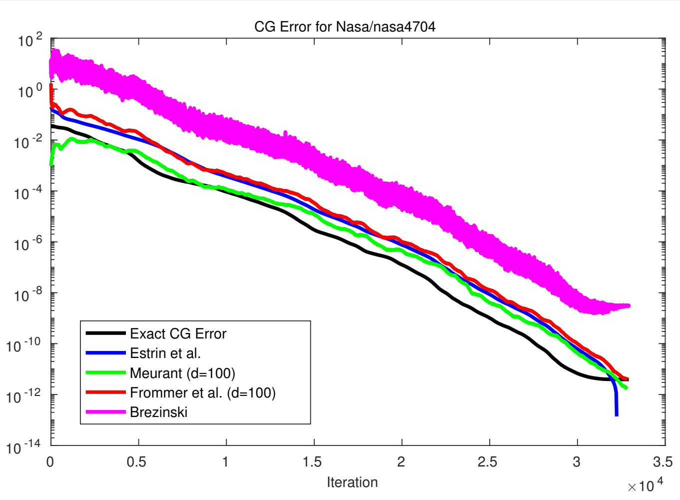

SYMMLQ error for HB/bcsstk28 $n = 4410$ SPD $\kappa(A) = 10^8$

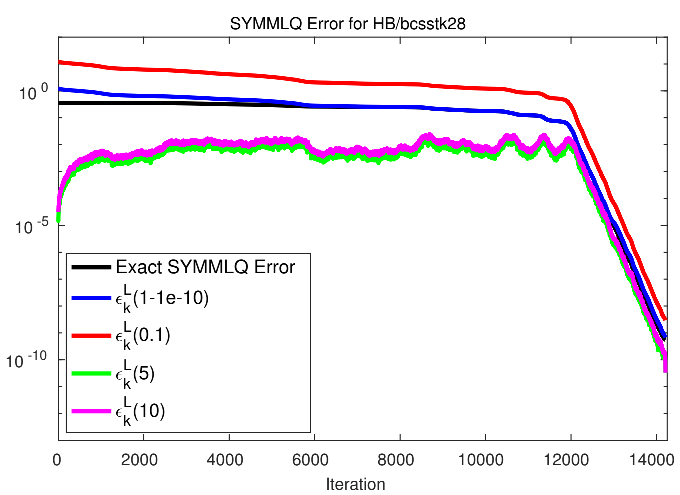

SYMMLQ error for GHS_psdef/wathen100 $n = 30401$ SPD $\kappa(A) = 10^3$

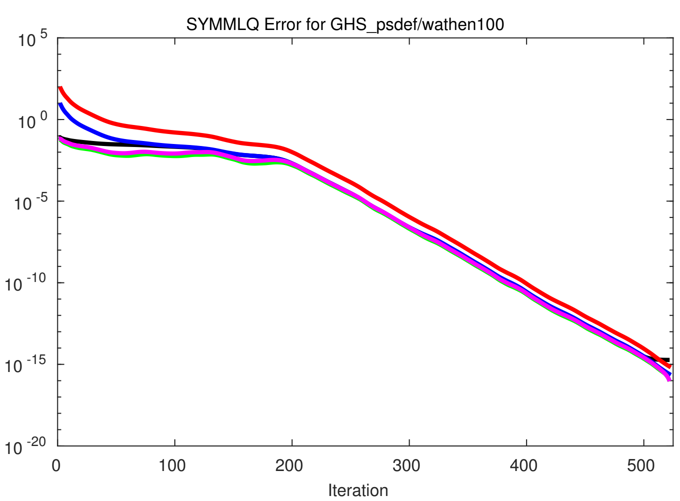

CG error for HB/bcsstk28 $n = 4410$ SPD $\kappa(A) = 10^8$

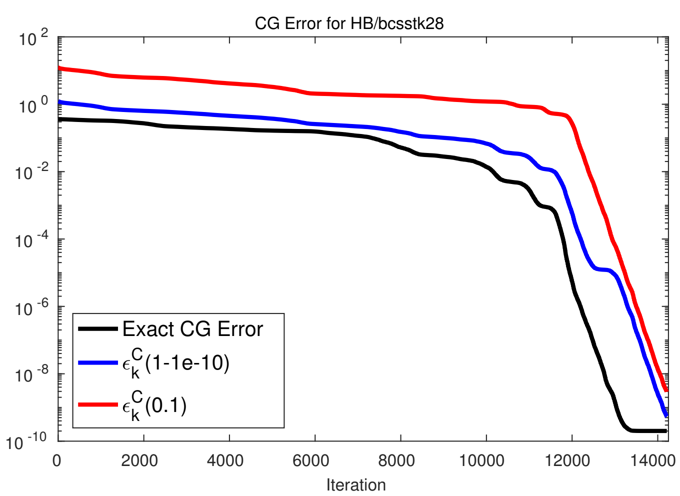

CG error for GHS_psdef/wathen100 $n = 30401$ SPD $\kappa(A) = 10^3$

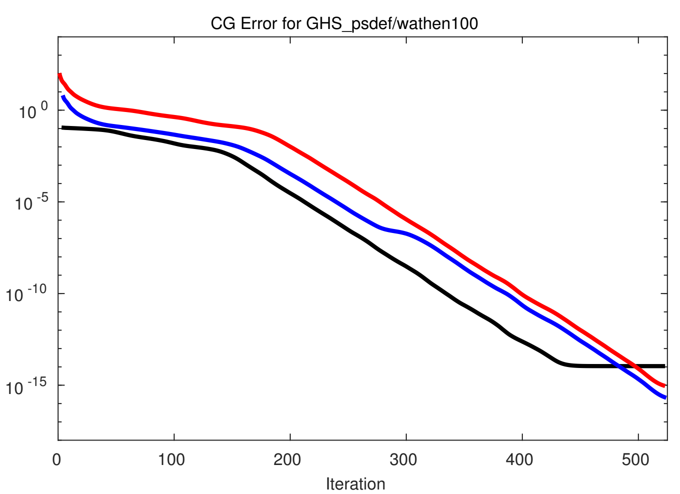

SYMMLQ error for PARSEC/Na5 $n = 5822$ indef $\kappa(A) = 10^3$

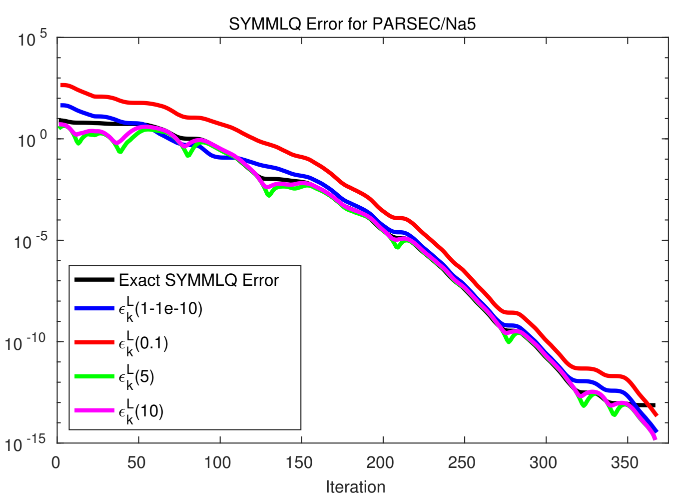

SYMMLQ error for PARSEC/SiNa $n = 5743$ indef $\kappa(A) = 10^2$

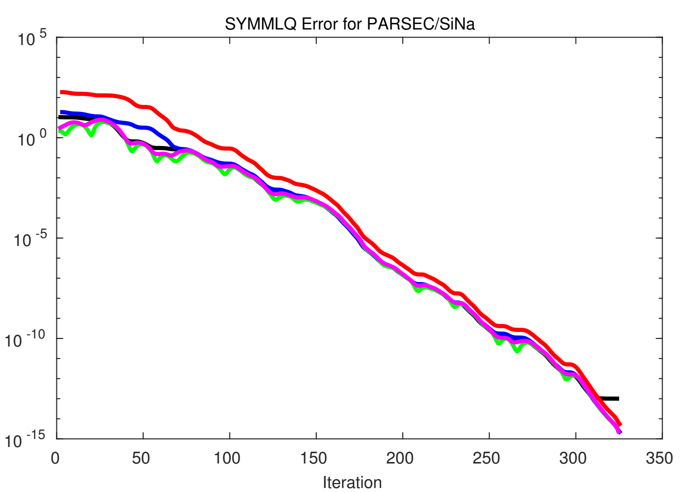

Reminder: CG vs MINRES

on SPD $Ax = b$

CG vs MINRES

- D. Titley-Peloquin (2010), Backward Perturbation Analysis of Least Squares Problems, PhD thesis, McGill University

Backward errors for x_k

$$\min_{\xi, E, f} \xi \qquad \text{st} \qquad (A + E)x_k = b + f, \qquad \frac{\|E\|}{\|A\|} \leq \alpha\xi, \qquad \frac{\|f\|}{\|b\|} \leq \beta\xi$$

CG vs MINRES

- D. Titley-Peloquin (2010), Backward Perturbation Analysis of Least Squares Problems, PhD thesis, McGill University

Backward errors for x_k

$$\min_{\xi, E, f} \xi \quad \text{st} \quad (A + E)x_k = b + f, \quad \frac{\|E\|}{\|A\|} \leq \alpha\xi, \quad \frac{\|f\|}{\|b\|} \leq \beta\xi$$

- D. C.-L. Fong and S. (2012), CG versus MINRES: An empirical comparison, SQU Journal for Science

Theorem

MINRES backward errors $\|E_k\| \propto \|r_k\| / \|x_k\|$ and $\|f_k\| \propto \|r_k\|$ decrease monotonically

CG vs MINRES, $n = 19779$, backward errors $\|r_k\| \,/\, \|x_k\|$

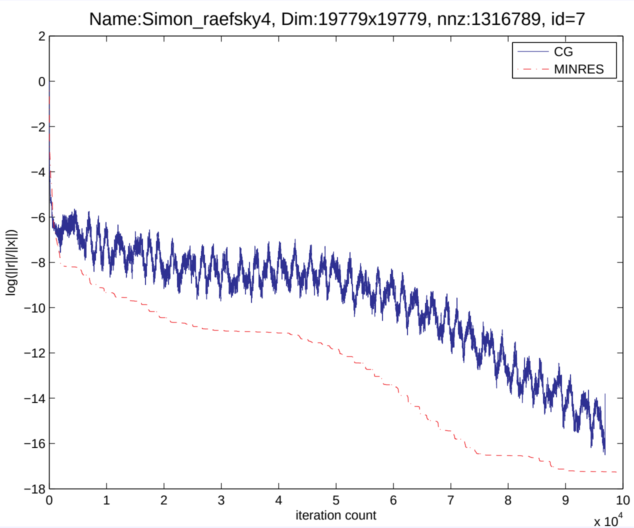

CG vs MINRES, $n = 245874$, backward errors $\|r_k\| \,/\, \|x_k\|$

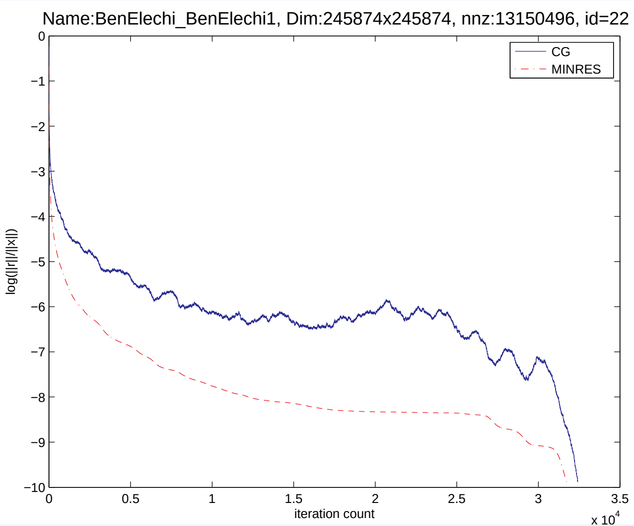

Conclusions

Conclusions

- Derived a **cheap** estimate of errors $\|x_* - x_k\|$ for SYMMLQ and CG.
- When A is SPD, the estimates are upper bounds (assuming exact arithmetic, but empirically in practice until convergence).
- Requires underestimate of smallest nonzero eigenvalue.
  - Common to Gauss-Radau quadrature based methods.
  - Depending on application (e.g. some PDEs) may be reasonable to obtain.
  - Easy for damped least-squares $(A^T A + \lambda^2 I)x = A^T b$. Hence good for LSLQ and LSQR.
- When A is indefinite, the error bound for SYMMLQ seems a good estimate.
- Extend to LSLQ for least-squares problems.

References

- R. Estrin, D. Orban, and S.
  Euclidean-norm error bounds for SYMMLQ and CG
  SIMAX (revised Aug 2017)

References

- R. Estrin, D. Orban, and S.
  Euclidean-norm error bounds for SYMMLQ and CG
  SIMAX (revised Aug 2017)

- R. Estrin, D. Orban, and S.
  LSLQ: An iterative method for linear least-squares with an error minimization property
  SIMAX (in revision Sep 2017)

References

- R. Estrin, D. Orban, and S.
  Euclidean-norm error bounds for SYMMLQ and CG
  SIMAX (revised Aug 2017)

- R. Estrin, D. Orban, and S.
  LSLQ: An iterative method for linear least-squares with an error minimization property
  SIMAX (in revision Sep 2017)
    - Seismic inverse problem, PDE-constrained optimization

References

- R. Estrin, D. Orban, and S.
  Euclidean-norm error bounds for SYMMLQ and CG
  SIMAX (revised Aug 2017)

- R. Estrin, D. Orban, and S.
  LSLQ: An iterative method for linear least-squares with an error minimization property
  SIMAX (in revision Sep 2017)

  - Seismic inverse problem, PDE-constrained optimization
  - Error in gradient of penalty function is bounded by error in x

References

- R. Estrin, D. Orban, and S.
  Euclidean-norm error bounds for SYMMLQ and CG
  SIMAX (revised Aug 2017)
- R. Estrin, D. Orban, and S.
  LSLQ: An iterative method for linear least-squares with an error minimization property
  SIMAX (in revision Sep 2017)

  - Seismic inverse problem, PDE-constrained optimization
  - Error in gradient of penalty function is bounded by error in x
  - Monotonic error in LSLQ iterates $\Rightarrow$ monotonic decrease in error in gradient

Special thanks

- Chris Paige

Special thanks

- Chris Paige
- Gene Golub

Special thanks

- Chris Paige
- Gene Golub
- Gene Golub, Gérard Meurant, Zdeněk Strakoš

Special thanks

- Chris Paige
- Gene Golub
- Gene Golub, Gérard Meurant, Zdeněk Strakoš
- Ron Estrin, Dominique Orban

Special thanks

- Chris Paige
- Gene Golub
- Gene Golub, Gérard Meurant, Zdeněk Strakoš
- Ron Estrin, Dominique Orban
- David Gleich

Special thanks

- Chris Paige
- Gene Golub
- Gene Golub, Gérard Meurant, Zdeněk Strakoš
- Ron Estrin, Dominique Orban
- David Gleich
- Yuja Wang, youtube

Special thanks

- Chris Paige
- Gene Golub
- Gene Golub, Gérard Meurant, Zdeněk Strakoš
- Ron Estrin, Dominique Orban
- David Gleich
- Yuja Wang, youtube
- Late-night talk shows (come back Jay Leno!)