Science & Technology Facilities Council

Orderings governed by numerical factorization

Iain Duff[1,2] **and Mario Arioli**[1]

STFC Rutherford Appleton Laboratory[1] and CERFACS[2]

SIAM Applied Linear Algebra Meeting.
16-22 June 2012

We wish to find the least-squares solution to the linear system

$$\mathbf{Ax} = \mathbf{b}$$

where the sparse matrix $\mathbf{A}$ has dimension $m \times n$, $m \geq n$ with rank n, and the solution $\mathbf{x}$ is such that

$$||\mathbf{b} - \mathbf{Ax}||_2$$

is minimized.

We will use the augmented system formulation

$$\begin{bmatrix} I_m & A \\ A^T & 0 \end{bmatrix} \begin{bmatrix} r \\ x \end{bmatrix} = \begin{bmatrix} b \\ 0 \end{bmatrix}.$$

We will look at the selection of an $n \times n$ nonsingular basis matrix, B, and the resulting partition of A as

$$A = \begin{pmatrix} B \\ N \end{pmatrix}.$$

With this partition, the augmented system can be written as:

$$\begin{bmatrix} I_n & 0 & B \\ 0 & I_{m-n} & N \\ B^T & N^T & 0 \end{bmatrix} \begin{bmatrix} r_B \\ r_N \\ x \end{bmatrix} = \begin{bmatrix} b_B \\ b_N \\ 0 \end{bmatrix}.$$

By eliminating the first n-variables, we obtain the reduced system:

$$\begin{bmatrix} I_{m-n} & N \\ N^T & -B^T B \end{bmatrix} \begin{bmatrix} r_N \\ x \end{bmatrix} = \begin{bmatrix} b_N \\ -B^T b_B \end{bmatrix}.$$

This linear system has a matrix that is Symmetric Quasi-Definite (SQD).

By symmetrically preconditioning this system by

$$M = \begin{bmatrix} I_{m-n} & 0 \\ 0 & B^{-T} \end{bmatrix}$$

we have

$$\begin{bmatrix} I_{m-n} & NB^{-1} \\ B^{-T}N^{T} & -I_n \end{bmatrix} \begin{bmatrix} r_N \\ Bx \end{bmatrix} = \begin{bmatrix} b_N \\ -b_B \end{bmatrix}.$$

We will look at three ways of solving these systems using the preconditioning matrix M or $\begin{bmatrix} I_n & 0 \\ 0 & M \end{bmatrix}$.

The first is to use LSQR on the preconditioned augmented system:

$$\begin{bmatrix} I_n & 0 & I \\ 0 & I_{m-n} & NB^{-1} \\ I & B^{-T}N^T & 0 \end{bmatrix} \begin{bmatrix} r_B \\ r_N \\ x \end{bmatrix} = \begin{bmatrix} b_B \\ b_N \\ 0 \end{bmatrix}.$$

which is the same as using LSQR on the augmented matrix

$$\begin{bmatrix} I \\ NB^{-1} \end{bmatrix}$$

.

The other methods we compare with work on the reduced SQD system:

$$\left[\begin{array}{cc} I_{m-n} & NB^{-1} \\ B^{-T}N^T & -I_n \end{array}\right] \left[\begin{array}{c} r_N \\ Bx \end{array}\right] = \left[\begin{array}{c} b_N \\ -b_B \end{array}\right].$$

where we compare MINRES with LSQR (called GLSQR-R by Arioli and Orban see also Saunders BIT 1995).

We note that the conditioning of the matrix NB^{-1} is not related to the conditioning of the original matrix A since, if the QR decomposition of A is $A = QR$ then

$$\begin{bmatrix} B \\ N \end{bmatrix} = \begin{bmatrix} Q_B \\ Q_N \end{bmatrix} R$$

so that

$$NB^{-1} = Q_N Q_B^{-1}$$

which does not involve R that determines the conditioning of A.

CS decomposition

For the orthonormal matrix Q from our QR factorization of A, there exist matrices $U_B \in \mathrm{R}^{n\times n}$, $U_N^{(m-n)\times(m-n)}$, $V \in \mathrm{R}^{n\times n}$, all orthogonal such that

$$\begin{bmatrix} U_B & 0 \\ 0 & U_N \end{bmatrix} \begin{bmatrix} Q_B \\ Q_N \end{bmatrix} V^T = \begin{bmatrix} I_{n-j} & 0 \\ 0 & C \\ 0 & 0 \\ 0 & S \end{bmatrix} \begin{array}{l} \}\, n-j \\ \}\, j \\ \}\, m-n-j \\ \}\, j \end{array}$$

where $C \in \mathrm{R}^{j\times j}$ and $S \in \mathrm{R}^{j\times j}$ are diagonal matrices $C = diag(c_i)$ and $S = diag(s_i)$ such that

$$c_i^2 + s_i^2 = 1,$$

with $c_i, s_i > 0$ for all i.

Using the CS decomposition, we can compute the singular value decomposition of NB^{-1}:

$$NB^{-1} = Q_N Q_B^{-1} = U_N^T \begin{bmatrix} 0 & 0 \\ 0 & SC^{-1} \end{bmatrix} U_B.$$

Properties of the Spectra

The eigenvalues $\mu_i, \ i = 1, \ldots, m$ of

$$\mathcal{A} = \begin{bmatrix} I_{m-n} & NB^{-1} \\ B^{-T}N^T & -I_n \end{bmatrix}$$

are related to the singular values of the matrix NB^{-1}.

Properties of the Spectra

From the CS-decomposition, we have

$$\begin{bmatrix} U_N & 0 \\ 0 & U_B \end{bmatrix} \begin{bmatrix} I_{m-n} & NB^{-1} \\ B^{-T}N^T & -I_n \end{bmatrix} \begin{bmatrix} U_N^T & 0 \\ 0 & U_B^T \end{bmatrix} =$$

$$\begin{bmatrix} I_{m-n-j} & 0 & 0 & 0 \\ 0 & -I_{n-j} & 0 & 0 \\ 0 & 0 & I_j & SC^{-1} \\ 0 & 0 & SC^{-1} & -I_j \end{bmatrix},$$

Properties of the Spectra

Thus the eigenvalues of $\mathcal{A}$ are

$$
\begin{aligned}
\mu_i &= \frac{1}{c_i} \quad i = 1, \ldots, j \\
\mu_i &= -\frac{1}{c_i} \quad i = j+1, \ldots, 2j
\end{aligned}
$$

with $m - n - j$ eigenvalues of value 1 and $n - j$ of value -1.

Choice of B

We show that the performance of all methods is very dependent on the preconditioner M determined by the choice of B.

Indeed, a major point of this talk is that the choice of the rows that constitute B can have a very significant effect on the power of the preconditioning and on the subsequent solution.

A fascinating illustration of this phenomenon can be seen by looking at the (in)famous Vandermonde matrix, with first column of ones and succeeding columns increasing powers of the data points. If we consider a matrix with 20 columns and 10000 rows, the condition number of this matrix is $\gtrsim 10^{16}$.

If we define our basis matrix B by choosing uniformly spaced data points then the condition number of NB^{-1} is around 6×10^5. However, if the basis matrix is formed using the Chebychev points in the interval then the condition number is around 4!

The following example shows that we can have an apparently good matrix A without a good choice for B.

$$A = Q = \left[\begin{array}{cc} I_{n-1} & 0 \\ 0 & v \end{array} \right]$$

where $v \in \mathbb{R}^{m-n+1}$ and $v_i = 1/\sqrt{m-n+1}$, $\forall i$.

For $m \to \infty$ we have

$$Q^T Q = I_n \quad \text{and} \quad \mathrm{Rank}(Q_B) \to n-1.$$

Rook Pivoting

We will use a sparse rook pivoting algorithm to determine the rows of B (HSL routine MC58). That is, we choose an entry a_{ij} as pivot only if the condition:

$$|a_{ij}| \geq u * \max\{\max_k |a_{kj}|, \max_l |a_{il}|\}$$

holds, where u is a threshold parameter with $0 < u \leq 1$.

HSL routines

MC58 is a HSL package that performs a sparse LU factorization but does not keep the factors and has the option of using rook pivoting. It can be run on rectangular matrices and will identify which rows and columns are in a nonsingular block of order r, the estimated rank of the matrix, where $r \leq \min\{m, n\}$.
MA48 is a HSL package for solving sets of sparse linear equations where the matrix is unsymmetric or rectangular. It also uses an LU factorization.

We first use MC58 on the matrix A to find the rows in B but the factorization of B is performed by MA48 using threshold pivoting with a threshold value of 0.01.

This is very much work in progress and so today we show results on rather small matrices, namely those popularized by Paige and Saunders, obtained by Saunders from the DSIR Geophysics Division, Wellington, NZ, and included in the original Harwell-Boeing test set and in the Florida collection of Tim Davis.

id	m	n	nnz	$\kappa(A)$
`well1033`	1033	320	4732	$1.6\ 10^2$
`illc1033`	1033	320	4719	$1.8\ 10^4$
`well1850`	1850	712	8755	$1.1\ 10^2$
`illc1850`	1850	712	8636	$1.4\ 10^4$

Table: Dimensions, number of nonzeros and $\kappa(A) = ||A^\dagger||_2||A||_2$ for Paige-Saunders tests.

Effect of threshold on rook pivoting

We look at the effect of changing the threshold value for `MC58` on these matrices.

u	`MC58`		`MA48`	$\|\|B^{-1}\|\|$	$\|\|NB^{-1}\|\|$
	Time	L/U	L/U		
0.1	0.03	6265	1842	$2.4\ 10^4$	$3.8\ 10^2$
0.5	0.03	7186	1819	$1.4\ 10^4$	$1.0\ 10^2$
1.0	0.32	11710	1835	$1.6\ 10^4$	$6.7\ 10^1$

Table: Data on determination and factorization of basis matrix for WELL1033.

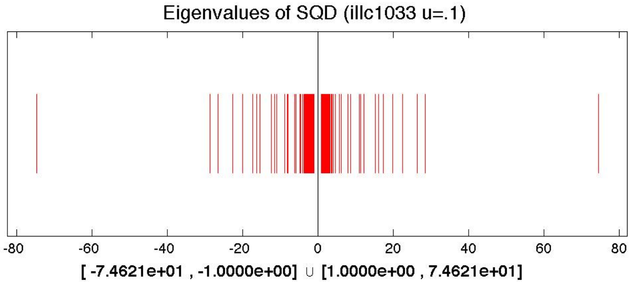
Eigenvalues of SQD (illc1033 u=.1)
-80
-60
-40
-20
0
20
40
60
80
[-7.4621e+01 , -1.0000e+00] ∪ [1.0000e+00 , 7.4621e+01]

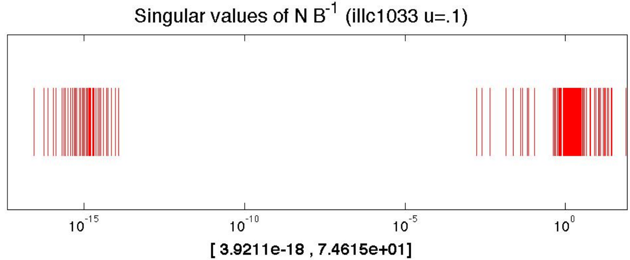
Singular values of N B^{-1} (illc1033 u=.1)
10^{-15}
10^{-10}
10^{-5}
10^{0}
[3.9211e-18 , 7.4615e+01]

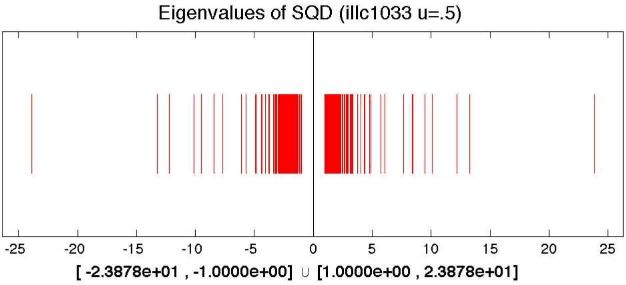
Eigenvalues of SQD (illc1033 u=.5)
-25 -20 -15 -10 -5 0 5 10 15 20 25
[-2.3878e+01 , -1.0000e+00] ∪ [1.0000e+00 , 2.3878e+01]

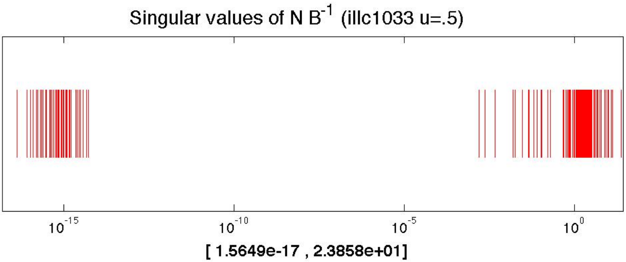
Singular values of N B^{-1} (illc1033 u=.5)
10^{-15} 10^{-10} 10^{-5} 10^{0}
[1.5649e-17 , 2.3858e+01]

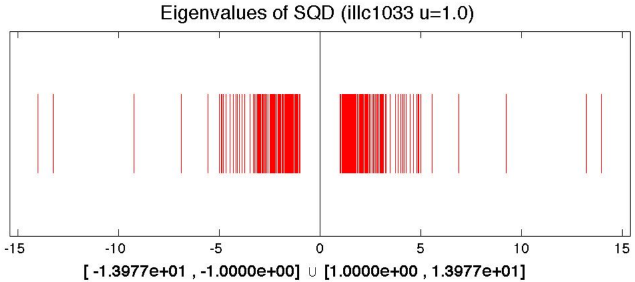
Eigenvalues of SQD (illc1033 u=1.0)
-15
-10
-5
0
5
10
15
[-1.3977e+01 , -1.0000e+00] ∪ [1.0000e+00 , 1.3977e+01]

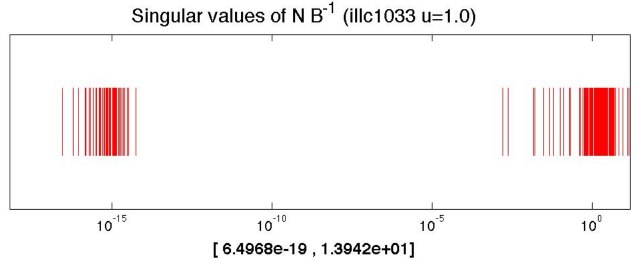
Singular values of N B^-1 (illc1033 u=1.0)
10^-15
10^-10
10^-5
10^0
[6.4968e-19 , 1.3942e+01]

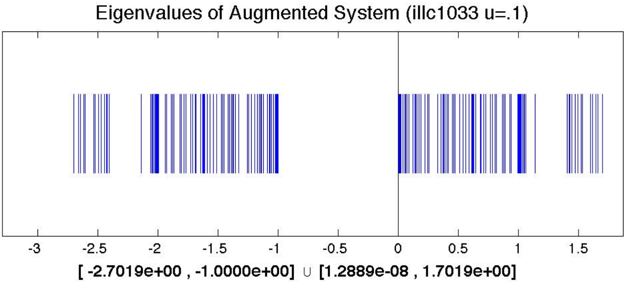
Eigenvalues of Augmented System (illc1033 u=.1)
-3 -2.5 -2 -1.5 -1 -0.5 0 0.5 1 1.5
[-2.7019e+00 , -1.0000e+00] ∪ [1.2889e-08 , 1.7019e+00]

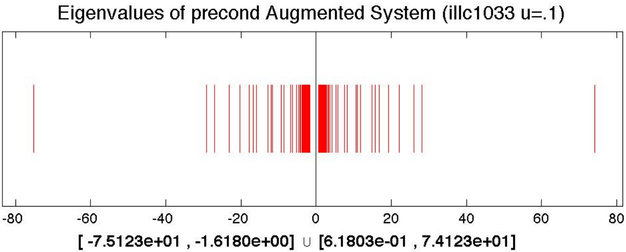
Eigenvalues of precond Augmented System (illc1033 u=.1)
-80 -60 -40 -20 0 20 40 60 80
[-7.5123e+01 , -1.6180e+00] ∪ [6.1803e-01 , 7.4123e+01]

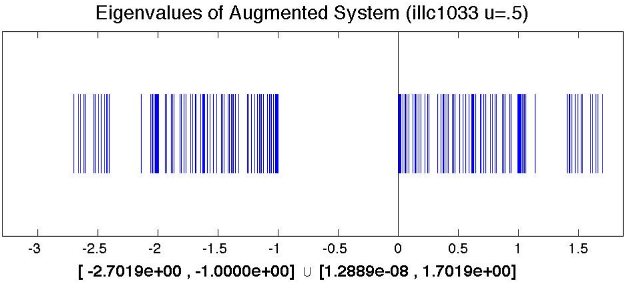
Eigenvalues of Augmented System (illc1033 u=.5)
-3
-2.5
-2
-1.5
-1
-0.5
0
0.5
1
1.5
[-2.7019e+00 , -1.0000e+00] ∪ [1.2889e-08 , 1.7019e+00]

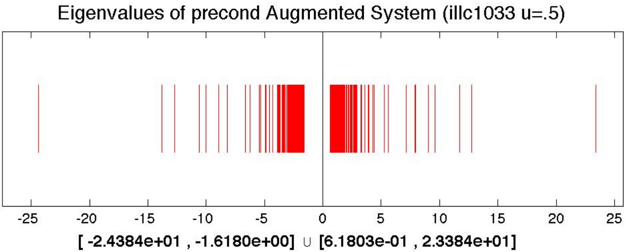
Eigenvalues of precond Augmented System (illc1033 u=.5)
-25
-20
-15
-10
-5
0
5
10
15
20
25
[-2.4384e+01 , -1.6180e+00] ∪ [6.1803e-01 , 2.3384e+01]

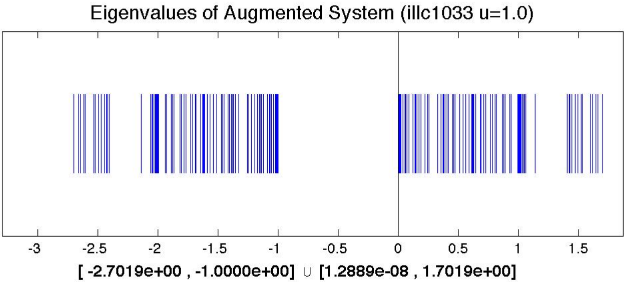
Eigenvalues of Augmented System (illc1033 u=1.0)
-3
-2.5
-2
-1.5
-1
-0.5
0
0.5
1
1.5
[-2.7019e+00 , -1.0000e+00] ∪ [1.2889e-08 , 1.7019e+00]

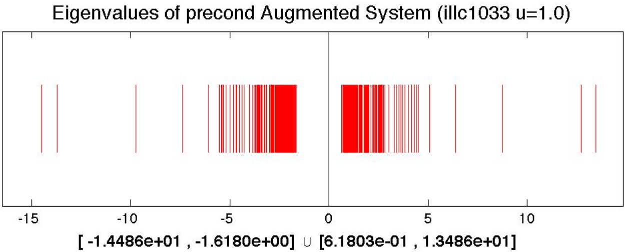
Eigenvalues of precond Augmented System (illc1033 u=1.0)
-15
-10
-5
0
5
10
[-1.4486e+01 , -1.6180e+00] ∪ [6.1803e-01 , 1.3486e+01]

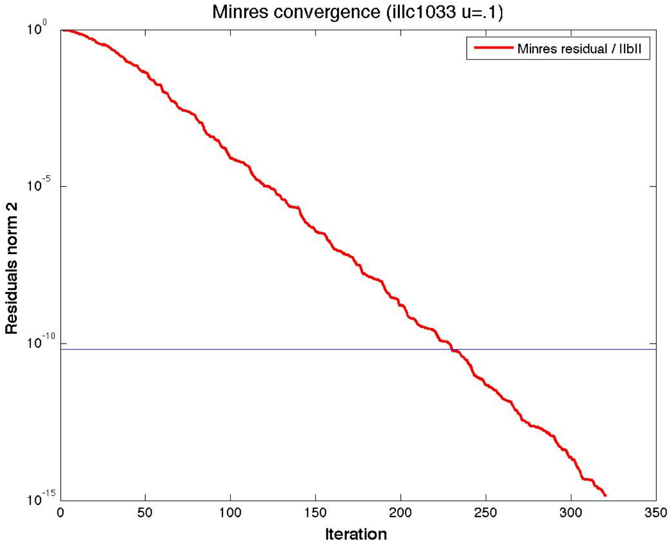
Minres convergence (illc1033 u=.1)
Minres residual / llbll
Residuals norm 2
Iteration
10^{0}
10^{-5}
10^{-10}
10^{-15}
0
50
100
150
200
250
300
350

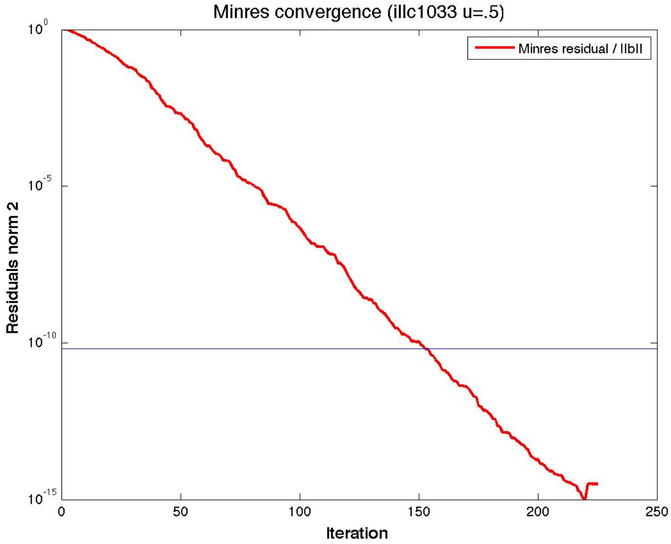
Minres convergence (illc1033 u=.5)
Minres residual / ||b||
Residuals norm 2
Iteration
10^{0}
10^{-5}
10^{-10}
10^{-15}
0
50
100
150
200
250

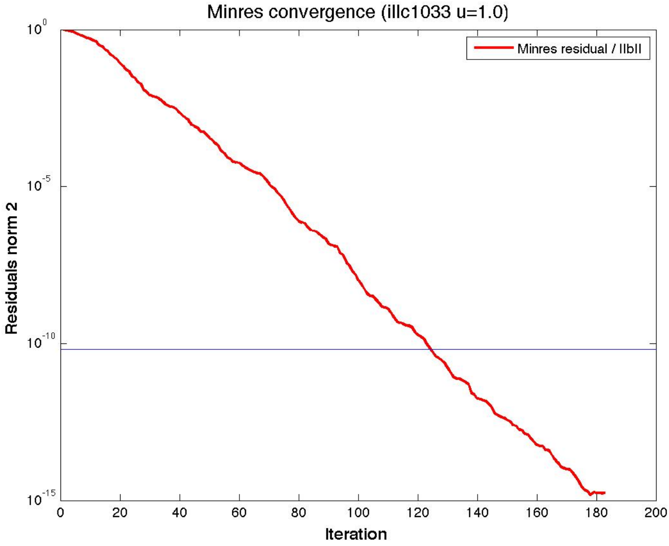
Minres convergence (illc1033 u=1.0)
Minres residual / ||b||
Residuals norm 2
Iteration
10^{0}
10^{-5}
10^{-10}
10^{-15}
0
20
40
60
80
100
120
140
160
180
200

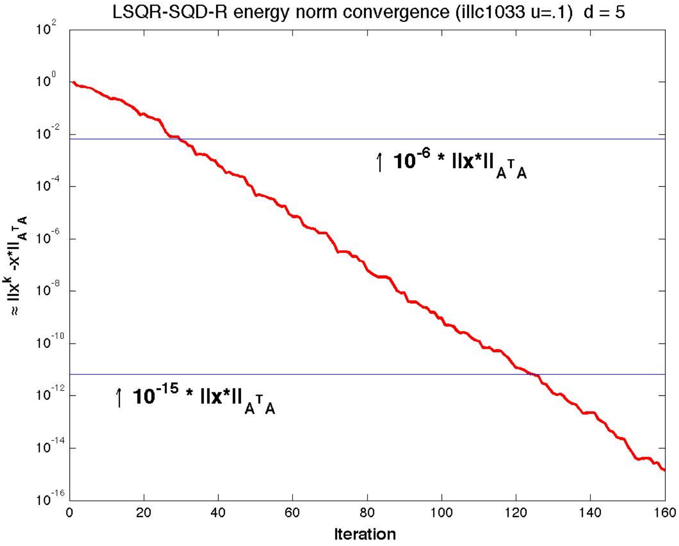
LSQR-SQD-R energy norm convergence (illc1033 u=.1) d = 5
$\approx \|x^k - x^*\|_{A^TA}$
Iteration
$\uparrow 10^{-6} * \|x^*\|_{A^TA}$
$\uparrow 10^{-15} * \|x^*\|_{A^TA}$
10^{2}
10^{0}
10^{-2}
10^{-4}
10^{-6}
10^{-8}
10^{-10}
10^{-12}
10^{-14}
10^{-16}
0
20
40
60
80
100
120
140
160

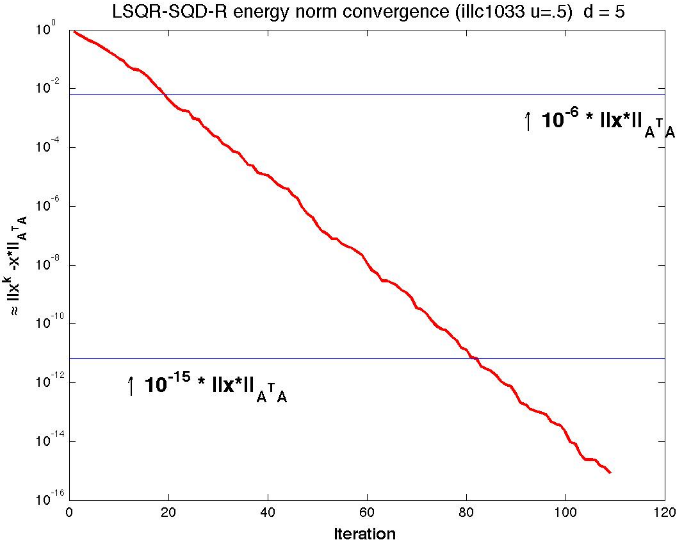
LSQR-SQD-R energy norm convergence (illc1033 u=.5) d = 5
$\approx \|x^k - x^*\|_{A^TA}$
Iteration
$\uparrow 10^{-6} * \|x^*\|_{A^TA}$
$\uparrow 10^{-15} * \|x^*\|_{A^TA}$
10^{0}
10^{-2}
10^{-4}
10^{-6}
10^{-8}
10^{-10}
10^{-12}
10^{-14}
10^{-16}
0
20
40
60
80
100
120

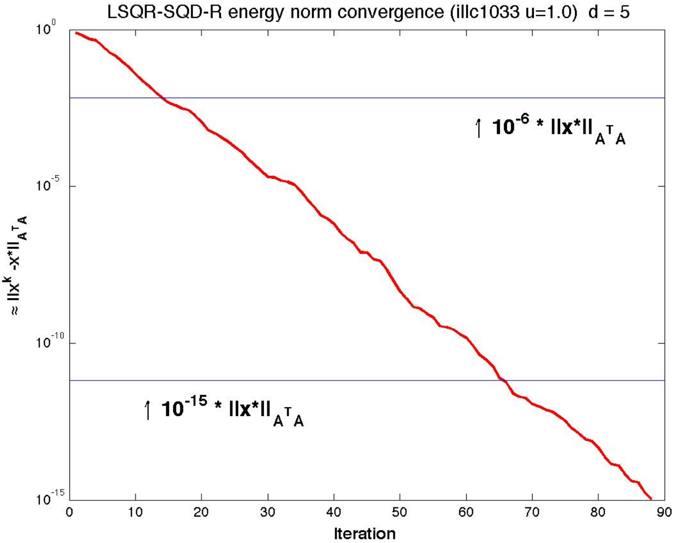
LSQR-SQD-R energy norm convergence (illc1033 u=1.0) d = 5
$\approx \|x^k - x^*\|_{A^TA}$
Iteration
$\uparrow 10^{-6} * \|x^*\|_{A^TA}$
$\uparrow 10^{-15} * \|x^*\|_{A^TA}$
10^0
10^{-5}
10^{-10}
10^{-15}
0
10
20
30
40
50
60
70
80
90

MINRES vs GLSQR

It is known that Krylov space methods based on

$$\mathcal{A} = \begin{bmatrix} I_{m-n} & NB^{-1} \\ B^{-T}N^{T} & -I_n \end{bmatrix}$$

will display some form of stagnation at every other step because of the symmetry of the spectrum of $\mathcal{A}$ (see Theorem 6.9.9 in Fischer and Freund, Golub, and Nachtigal 1991). In particular, the convergence rate of MINRES (Paige-Saunders 1975,Saunders 1995) will require in exact arithmetic double the number of iterations required by GLSQR.

Normal equations for SQD

We observe that the normal equations

$$\mathcal{A}^2 \begin{bmatrix} r_N \\ Bx \end{bmatrix} = \mathcal{A} \begin{bmatrix} b_N \\ -b_B \end{bmatrix}$$

have a very interesting structure because

Normal equations for SQD

We observe that the normal equations

$$\mathcal{A}^2 \begin{bmatrix} r_N \\ Bx \end{bmatrix} = \mathcal{A} \begin{bmatrix} b_N \\ -b_B \end{bmatrix}$$

have a very interesting structure because

$$\mathcal{A}^2 = \begin{bmatrix} I_{m-n} + NB^{-1}B^{-T}N^T & 0 \\ 0 & I_n + B^{-T}N^TNB^{-1} \end{bmatrix} = \mathcal{D}$$

Prob.	Aug. System	SQD	
	LSQR	MINRES	GLSQR-2
illc1033	3.9e-13 (95)	6.7e-13 (182)	1.0e-13 (94)
illc1850	4.5e-14 (142)	6.2e-14 (270)	1.5e-14 (138)
well1033	1.2e-14 (105)	7.3e-15 (204)	4.8e-15 (105)
well1850	3.8e-15 (148)	1.5e-14 (285)	1.3e-15 (146)

Table: Forward errors $||x_{QR} - x^{(k)}||_2/||x_{QR}||_2(u = 1\ ,\ d = 5)$. In parentheses the corresponding number of iterations.

GLSQR is slightly more accurate than LSQR.

Conclusions

- Identifying and using a basis of the overdetermined coefficient matrix for the solution of linear least squares problems yields an effective preconditioner for both LSQR and MINRES.
- The choice of the basis matrix is important.
- Rook pivoting is an effective way of defining a good basis.
- For more information on solution of SQD systems using LSQR, see talk by Mario Arioli at Sparse Days at CERFACS, next week in Toulouse.

THANK YOU FOR YOUR ATTENTION

Effect of threshold on rook pivoting

u	MC58		MA48	$\|\|B^{-1}\|\|$	$\|\|NB^{-1}\|\|$
	Time	L/U	L/U		
0.1	0.02	5542	1821	$6.3\ 10^5$	$1.5\ 10^2$
0.5	0.05	7317	1803	$4.5\ 10^5$	$1.0\ 10^2$
1.0	0.10	10092	1866	$2.9\ 10^5$	$5.1\ 10^1$

Table: Data on determination and factorization of basis matrix for ILL1033.

Effect of threshold on rook pivoting

u	MC58		MA48	$\|\|B^{-1}\|\|$	$\|\|NB^{-1}\|\|$
	Time	L/U	L/U		
0.1	0.05	9996	4215	$3.0\ 10^5$	$4.3\ 10^3$
0.5	0.54	13927	4361	$2.6\ 10^4$	$3.2\ 10^2$
1.0	1.63	42381	4586	$3.0\ 10^4$	$2.5\ 10^2$

Table: Data on determination and factorization of basis matrix for WELL1850.

Effect of threshold on rook pivoting

u	MC58		MA48	$\|\|B^{-1}\|\|$	$\|\|NB^{-1}\|\|$
	Time	L/U	L/U		
0.1	0.05	9965	4248	$1.8\ 10^6$	$5.8\ 10^3$
0.5	0.32	13941	4341	$3.7\ 10^4$	$1.7\ 10^2$
1.0	1.39	24804	4621	$9.6\ 10^4$	$2.9\ 10^2$

Table: Data on determination and factorization of basis matrix for ILL1850.