Techniques for Solving Saddle Point Indefinite Linear Systems

Chen Greif, University of British Columbia

PIMS Workshop on Numerical Linear Algebra and Applications
UBC
August 2003

Indefinite linear systems

Recall Alison's short course this morning...

- There exists a vector $\mathbf{x}$ such that $\mathbf{x}^T A\mathbf{x} = 0, \ \mathbf{x} \neq 0.$
- Both positive and negative eigenvalues; no natural energy norm that can be used the way CG is derived; more difficult to handle...
- BUT, there are ways to handle indefinite linear systems in a nicer way than handling general nonsymmetric linear systems.

Not a walk in the park

- ”There is a widely held view that iterative solution of indefinite systems is much less reliable and much less efficient in general.” (*Wathen, Fischer and Silvester* [1997].)
- ”When the original matrix is strongly indefinite, i.e. when it has eigenvalues spread on both sides of the imaginary axis, the usual Krylov subspace methods may fail. The conjugate Gradient approach applied to the normal equations may then become a good alternative.” (*Saad* [1996].)

It is fair to say that the one unifying opinion in 2003 is that preconditioning is crucial. (Again, see Alison’s short course.)

Saddle point linear systems

$$\mathcal{K}u = \begin{pmatrix} A & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} x \\ \lambda \end{pmatrix} = \begin{pmatrix} c \\ d \end{pmatrix}$$

The matrix $\mathbf{A}$:

- square, $(n \times n)$, large, sparse, symmetric.
- in many cases is positive definite, but may be singular.

The matrix $\mathbf{B}$:

- rectangular, $(n \times m)$, $m < n$.
- is assumed to have full column rank.

No lack of associated applications

- Flow through electric networks.
- Navier-Stokes equations.
- Linear elasticity.
- Electromagnetics.
- Constrained least squares.
- Image processing.
- Data interpolation and surface fitting.
- Interior point methods.
- Structural analysis.
- Mixed finite element formulations.

and many more...

A familiar PDE with a constraint

Given data f, find the velocity u and pressure p satisfying

$$\begin{cases} -\nu\Delta u + (u\cdot\nabla)u + \nabla p = f \\ \text{div } u = 0 \end{cases}$$

on a domain $\Omega \subset \mathcal{R}^2$ or $\Omega \subset \mathcal{R}^3$, with some boundary conditions, e.g. $u = g$ on $\partial\Omega$.

Apply a fixed point iteration $\rightarrow$ obtain *Oseen equations*: (w is given)

$$\begin{cases} -\nu\Delta u + (w\cdot\nabla)u + \nabla p = f \\ \text{div } u = 0 \end{cases}$$

Example: Quadratic Programming

Equality-constrained quadratic programs:

Minimize $\frac{1}{2}x^T Ax - x^T c$ subject to $B^T x = d$.

Equivalent Lagrange Multipliers formulation: define

$$\phi(x, y) = x^T Ax - x^T c + \lambda^T (B^T x - d)$$

and compute its stationary points:

$$\nabla\phi = 0.$$

A few observations

- Indefiniteness: The system is indefinite.

$$e_{k+1}^T \mathcal{K} e_{k+1} = 0.$$

- Nonsingularity:
  - Want B to have full column rank, and want no intersection of the null-spaces of A and B^T.
  - If Z is a basis for the null space of B^T, the *reduced Hessian* $Z^T A Z$ is positive definite (and B has full column rank), then $\mathcal{K}$ is nonsingular.

A few observations (cont.)

- Inertia: If A is positive definite, we have n positive eigenvalues and m negative ones:

$$\begin{pmatrix} A & B \\ B^T & 0 \end{pmatrix} = \begin{pmatrix} I & 0 \\ B^T A^{-1} & I \end{pmatrix} \begin{pmatrix} A & 0 \\ 0 & S \end{pmatrix} \begin{pmatrix} I & A^{-1}B \\ 0 & I \end{pmatrix},$$

where $S = -B^T A^{-1} B$ is the *Schur complement*.

In general,

$$i(\mathcal{K}) = (m, m, 0) + i(Z^T A Z).$$

Solution methods

- Direct methods. (Duff, Bunch, Parlett, Kaufman,...)
- Schur complement techniques and preconditioned Krylov solvers. (Golub, Wathen, Elman, Silvester, Ramage, Fischer...)
- Null-space and inertia controlling methods. (Fletcher, Murray, Gill, Saunders, Forsgren, Wright...)
- Multiply second block-row by -1 and apply effective splitting. (Benzi, Golub, Wathen, Dyn, Ferguson...)
- Uzawa and inexact Uzawa. (Arrow, Hurwicz, Uzawa, Bramble, Pasciak, Elman, Golub,...)

Difficulties with direct methods

- Cannot form usual Cholesky due to the indefiniteness.

- Not practical to use Gaussian Elimination with partial pivoting or a variant, because the symmetry and sparsity structures are not taken into account.

Stable factorizations

- There are certain factorizations that are efficient for symmetric matrices, but have to be careful:
  - LDL^T not particularly stable.
    Example (Golub & Van Loan):

$$\begin{pmatrix} \varepsilon & 1 \\ 1 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ \frac{1}{\varepsilon} & 1 \end{pmatrix} \begin{pmatrix} \varepsilon & 0 \\ 0 & -\frac{1}{\varepsilon} \end{pmatrix} \begin{pmatrix} 1 & 0 \\ \frac{1}{\varepsilon} & 1 \end{pmatrix}^T,$$

  - Stable methods (Bunch & Parlett [1971] and others) are based on permutation matrices. The matrix D in LDL^T is block diagonal with either 1×1 or 2×2 block. (In our case, though, the sparsity structure may be less favorable.) See, for example, Golub & Van Loan, Section 4.4, or Nocedal & Wright, Chapter 16.

Difficulties with iterative methods

- As mentioned before, there are Krylov subspace solvers for indefinite symmetric matrices; however when applied straightforwardly to augmented systems no advantage of structure is fully taken.

- Without preconditioning, Krylov subspace solvers perform poorly.

- Sometimes it is worthwhile doing 'unthinkable' things... Like destroying the symmetry! (E.g. Permute rows in above 2×2 example; or multiply second block row by -1: Benzi, Golub, Bai.)

Solution approaches

What Hamlet **REALLY** asked:

> To *eliminate*, or not to *eliminate*, **THIS** is the question!

- Solve for λ first, then compute x.
- Solve for x first, then compute λ.
- Solve for each in an alternating fashion.
- Solve simultaneously for both x and λ.

Schur complement techniques

- **Solve for λ first, then compute x.**
  Performing block Gaussian elimination we get

$$B^T A^{-1} B \lambda = B^T A^{-1} c - d.$$

  - The cost will mainly depend on the cost of inverting A (and on how large n and m are).
  - Recovering x is cheap.
  - The approach is termed a *range-space* technique, or a *displacement* method.

Comments

- Unfortunately in most cases computing the Schur complement is prohibitively expensive due to the cost of inverting A.
- The above will not work at all if A is singular, and will work poorly if A is ill-conditioned.
- However let's cheer up:
  - cases where A is easy to invert *do* exist (structural analysis, electric networks - see Strang [1986,1988]).
  - Cases where A^{-1} is pretty much known explicitly (Quasi-Newton method).
  - Cases where m is so small that the number of backsolves for forming the Schur Complement is nice and small.

Elimination of constraints

- **Solve for x first, then compute λ.**
  - Direct elimination: reduce B to upper trapezoidal form via the QR factorization, or Gaussian Elimination with complete pivoting. As a result, obtain a reduced *unconstrained* least square problem.
  - Null Space method:
    * x is expressed as $x = Yx_Y + Zx_Z$, where $B^T Z = 0$ and $[Y|Z]$ is nonsingular.
    * By algebraic manipulations we obtain $Z^T AZx_z = -Z^T AYx_Y + Z^T c$.
    * Advantage: we do not rely on nonsingularity of A. Disadvantage: It might be expensive to compute the null-space matrix Z and to form $Z^T AZ$.

See Gill, Murray & Wright [1981], for an early discussion.

Direct attack of system

- Methods based on forming the SVD, or the generalized SVD associated with the matrices A and B, in order to simplify the original linear system. (Golub, Gander & Von Mat [1989], others.

- The method of weighting and augmented Lagrangian methods. Give a large weight to certain constraints (though it may cause ill-conditioning because of scaling) or insertion of penalty terms. (Powell & Reid [1969], Hestenes [1969], Powell [1969], Polyak [1970], Miele et. al. [1972], Lawson & Hanson [1974], Berksekas [1975], Van Loan [1985], and references therein.)

- Generate two sequences for the approximations of the unknowns and the Lagrange multipliers.

Preconditioning using the Schur complement

- Preconditioning the augmented system by a block diagonal matrix, typically associated with the Schur complement. (Elman et. al. , Saad et. al. , and others.)

- Effective and elegant approximations to the Schur Complement (in the context of Navier-Stokes equations) by Elman et. al.: Associate Schur complement to the mass matrix.

Uzawa algorithm (Arrow, Hurwicz, Uzawa [1958])

Construct a sequence of approximations to x and y, as follows:
For $k = 0, 1, \ldots$

Solve $Ax_{k+1} = c - By_k$

Compute $y_{k+1} = y_k + \alpha B^T x_{k+1}$

1. Optimal value of parameter α is $\frac{2}{\lambda_{min}(S)+\lambda_{max}(S)}$, where $S = B^T A^{-1} B$ is the Schur complement.

2. There is no need to find the *exact* solution of the 'inner' system. (Bank, Welfert & Yserentant [1990], Elman & Golub [1994], Bramble, Pasciak & Vassilev [1997])

$\mathcal{K}$ may have an Ill-Conditioned or Singular $(1,1)$ Block.

- Domain decomposition methods applied to linear elasticity and structural mechanics problems. (*Farhat & Roux* , *Klawonn & Widlund*.)
- Construction of smooth surfaces from aggregated data; computation of thin plate splines.
(*Dyn & Ferguson*, *Sibson & Stone*.)
- Electromagnetics and magnetostatics problems. (*Perugia, Simoncini & Arioli*.)
- Geophysical inverse problems.
(*Haber & Ascher*.)

We concentrate on:

[Joint work with Gene Golub]

1. ”Preprocessing”: eliminating the singularity of the (1,1) Block.
   (a) ...by using the Augmented Lagrangian method.
   (b) ...by reducing the system size. (Will not describe.)

2. Positive definite block preconditioners. (Preserve inertia, maintain symmetry, but certainly not clear if are superior to indefinite preconditioners.)

Augmented Lagrangian Approach

$$\begin{cases} Ax + B\lambda &= c \\ B^T x &= d \end{cases}$$

Since $BWB^T x = BWd$, can transform system into

$$\begin{cases} (A + BWB^T)x + B\lambda &= c + BWd \\ B^T x &= d \end{cases}$$

Connection to optimization: Fletcher, Hestenes, Powell, others...

Application to BVPs: Fortin, Glowinski, others...

Connection to Optimization

Constrained minimization problems:

minimize $f(x)$ subject to $c_i(x) = 0$

- A popular class of methods: introduce a penalty parameter:

$$Q(x;\mu) = f(x) + \frac{1}{2\mu}\sum c_i^2(x).$$

- Driving μ to zero penalizes constraint violations.

- Introducing explicit Lagrange multipliers reduces the ill-conditioning inherent in the penalty formulation:

$$\mathcal{L}(x;\mu;\lambda) = f(x) - \sum \lambda_i c_i(x) + \frac{1}{2\mu} \sum c_i^2(x).$$

See e.g. Nocedal & Wright, Chapter 17.

Spectral Analysis ($W = \gamma I$)

- The eigenvalues: Strong connection between the eigenvalues of $\mathcal{K}(W)$ and the generalized eigenvalues of the problem $\lambda A x = BB^T x$. There exists an $n \times n$ matrix G such that $A = GG^T$ and $BB^T = G\Lambda G^T$, where Λ are the generalized eigenvalues, and

$$\begin{pmatrix} A + \gamma BB^T & B \\ B^T & 0 \end{pmatrix} = \begin{pmatrix} G & 0 \\ 0 & V \end{pmatrix} \begin{pmatrix} I + \gamma \Sigma\Sigma^T & \Sigma \\ \Sigma^T & 0 \end{pmatrix} \begin{pmatrix} G^T & 0 \\ 0 & V^T \end{pmatrix}.$$

- Clustering: The m negative eigenvalues of $\mathcal{K}(\gamma)$ tend to cluster near to $-\frac{1}{\gamma}$.

- Dependence of $\kappa_2(\mathcal{K})$ on γ:

$$\kappa(\mathcal{K}(\gamma)) \approx \kappa(\mathcal{K}(0)) \cdot \kappa\left(\operatorname{diag}\left[\begin{pmatrix} 1+\gamma\lambda_i & \sqrt{\lambda_i} \\ \sqrt{\lambda_i} & 0 \end{pmatrix}\right]\right),$$

$$\frac{\kappa(\mathcal{K}(\gamma))}{\gamma^2} \to \|B\|_2^2 \neq 0 \ \text{ as } \ \gamma \to \infty.$$

The Inverse

Proposition. Suppose that A is a general $n \times n$ matrix, B and C are full column rank $n \times m$ matrices ($m \leq n$), and W is a $m \times m$ matrix. Define

$$\mathcal{K}(W) := \begin{pmatrix} A + BWC^T & B \\ C^T & 0 \end{pmatrix}.$$

Then for any $W \neq 0$ such that $\mathcal{K}(W)$ is nonsingular, the following holds: (we denote $\mathcal{K} \equiv \mathcal{K}(0)$)

$$\mathcal{K}^{-1}(W) = \mathcal{K}^{-1} - \begin{pmatrix} 0 & 0 \\ 0 & W \end{pmatrix}.$$

Possible benefits: Tight bound for condition number without additional work.

Picking the Right $\mathbf{W}$:

$$\begin{pmatrix} A + BWB^T & B \\ B^T & 0 \end{pmatrix} \begin{pmatrix} x \\ \lambda \end{pmatrix} = \begin{pmatrix} c + BWd \\ d \end{pmatrix}$$

- Sparsity considerations (e.g. use only certain columns of B).
- Obtain positive definiteness of the (1,1) block.
- Scaling/Balancing. Examples:
  - Set $W = \gamma I$, where $\gamma = \frac{\|A\|}{\|B\|^2}$.
  - Pick W to be a 'scaling operator'.

Block Preconditioning

Challenge: How to precondition while having in mind the structure of the matrix.

Indefinite block preconditioning

Constraint Preconditioning: Keller, Gould and Wathen (2000), Golub and Wathen (1998), Luksan and Vlcek (1998), Perugia and Simoncini (2000), Rozloznik and Simoncini (2002), and others.

$$\mathcal{M} = \begin{pmatrix} G & B \\ B^T & 0 \end{pmatrix}.$$

- Eigenvalue 1 with multiplicity $2m$. Rest of eigenvalues strongly related to the generalized eigenvalue problem $Z^TAZx = \lambda Z^TGZx$.

- Some more results related to eigenvalue bounds and eigenvector distribution.

- Detailed convergence analysis for Krylov solvers.

Positive Definite Block Preconditioning

Motivation: (Murphy, Golub & Wathen [1998])

$$\mathcal{M} = \begin{pmatrix} A & 0 \\ 0 & B^T A^{-1} B \end{pmatrix}.$$

$\mathcal{M}^{-1}\mathcal{K}$ has at most four nonzero distinct eigenvalues: $0, 1, \frac{1}{2} \pm \frac{\sqrt{5}}{2}$. Hence a minimum residual Krylov solver will terminate within four iterations.

However:

- Computing the Schur complement may be very expensive.
- The Schur complement may be 'inappropriate' in terms of the corresponding differential operator.
- The Schur complement may not exist altogether.

A natural way to generalize the above while avoiding the difficulties

We consider

$$\mathcal{M} = \begin{pmatrix} M & 0 \\ 0 & B^T N^{-1} B \end{pmatrix} ,$$

where M is much easier to invert.

See also De Sturler and Liesen, 2003.

A natural choice here: $M = N = A + BWB^T$.

Spectrum analysis

Let λ be an eigenvalue of the preconditioned matrix $\mathcal{M}^{-1}\mathcal{K}$, whose associated eigenvector is $\begin{pmatrix} u \\ v \end{pmatrix}$. Then setting $w = M^{\frac{1}{2}}u$ we have:

$$\left[\lambda^2 I - \lambda M^{-\frac{1}{2}} A M^{-\frac{1}{2}} - M^{-\frac{1}{2}} B (B^T M^{-1} B)^{-1} B^T M^{-\frac{1}{2}}\right] w = 0.$$

The matrix $P = P^2 = M^{-\frac{1}{2}} B (B^T M^{-1} B)^{-1} B^T M^{-\frac{1}{2}}$ is an orthogonal projector onto range$(M^{-1/2}B)$.

Denoting $C = M^{-1/2}B$, $K = M^{-1/2}AM^{-1/2}$, we have a quadratic eigenvalue problem of the form

$$(\lambda^2 I - \lambda K - P)z = 0.$$

The choice $M = A + \gamma B B^T$

$$\mathcal{M} = \begin{pmatrix} A + \gamma B B^T & 0 \\ 0 & B^T (A + \gamma B B^T)^{-1} B \end{pmatrix} .$$

Theorem. (With Varah.) The matrix P is a polynomial of degree m in K, and the explicit mapping is given by

$$I - P = f(K),$$

where $\{\lambda_i\}$ are the p nonzero (finite) eigenvalues of the generalized eigenvalue problem $Ax = \lambda BB^T x$, and f is a Lagrange interpolant of degree p:

$$f(t) = \frac{\prod_i \left(t - \frac{\lambda_i}{\lambda_i + \gamma}\right)}{\prod_i \left(1 - \frac{\lambda_i}{\lambda_i + \gamma}\right)}.$$

Corrolary. $n - m$ eigenvalues of the preconditioned matrix $\mathcal{M}^{-1}\mathcal{K}$ are equal to 1. (Therefore, fast convergence of minimum residual Krylov subspace solvers.)

Corollary. If A is positive semidefinite, the eigenvalues of the preconditioned matrix are bounded within the two intervals:

$$[-1, \frac{1}{2} - \frac{\sqrt{5}}{2}] \bigcup [1, \frac{1}{2} + \frac{\sqrt{5}}{2}].$$

Example: A geophysical inverse problem. (Provided by Eldad Haber.)

minimize

$$\phi(u,m) = \frac{1}{2}\|Qu - b\|^2 + \frac{\beta}{2}\|W(m - m_0)\|^2$$

$$\text{subject to} \;\; A(m)u = f.$$

The Lagrangian:

$$\mathcal{L}(u,m,\lambda) = \frac{1}{2}\|Qu - b\|^2 + \frac{\beta}{2}\|W(m - m_0)\|^2 + \lambda^T[A(m)u - f].$$

where λ is a vector of Lagrange multipliers.

The Hessian is given by

$$H(m,u,\lambda) = \begin{pmatrix} Q^TQ & K(m,\lambda)^T & A(m)^T \\ K(m,\lambda) & \beta W^TW + R(m,u,\lambda) & G(m,u)^T \\ A(m) & G(m,u) & 0 \end{pmatrix}.$$

Sparsity pattern of matrix

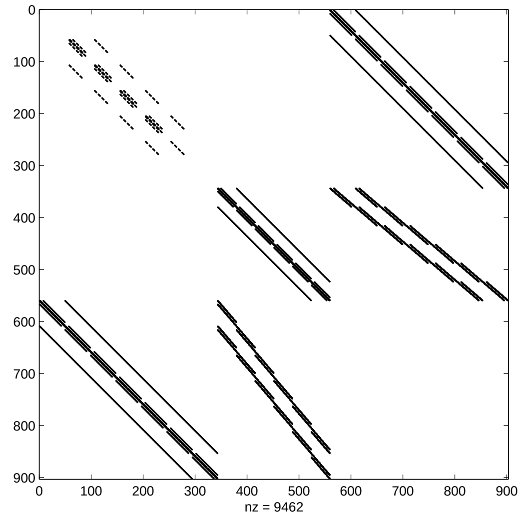

Matrix, again

$$H(m,u,\lambda) = \begin{pmatrix} Q^TQ & K(m,\lambda)^T & A(m)^T \\ K(m,\lambda) & \beta W^T W + R(m,u,\lambda) & G(m,u)^T \\ A(m) & G(m,u) & 0 \end{pmatrix}.$$

$A(m)$ - a discretization of a PDE, typically 2nd order.
f - source.
b - measured data.
R - Regularization operator.
m - Distributed parameter (conductivity, seismic velocity, porocity, etc.)
Q - Singular diagonal matrix, perhaps.
G - first order differential operator.

Gauss-Newton plus null-space: $(J^T J + \beta L)p = -g$.
L is a differential operator and J is an integral operator.

Set R and K to zero.

Illustration of Clustering Effect

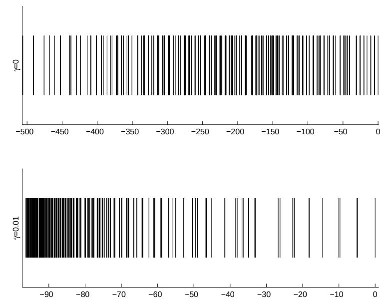

Figure 1: The 729 negative eigenvalues of the original matrix, i.e. for $\gamma = 0$, vs. those of a modified matrix with $\gamma = 0.01$. The eigenvalues closest to zero are approximately -0.0019 in both cases.

Comparison of Condition Numbers: Random matrix

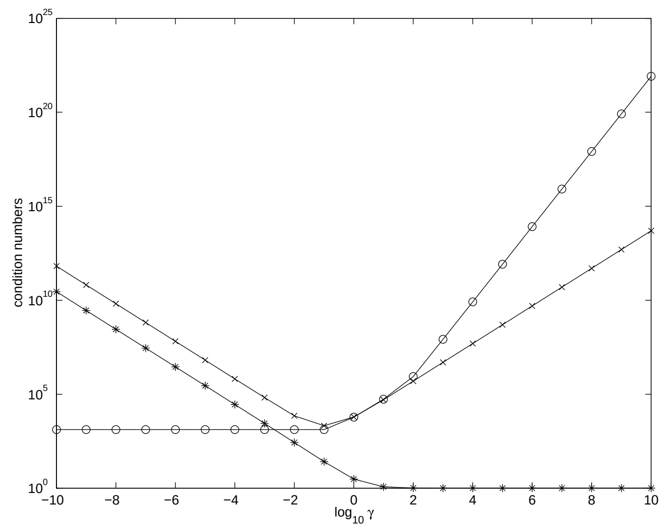

Comparison of Condition Numbers: Inverse problem

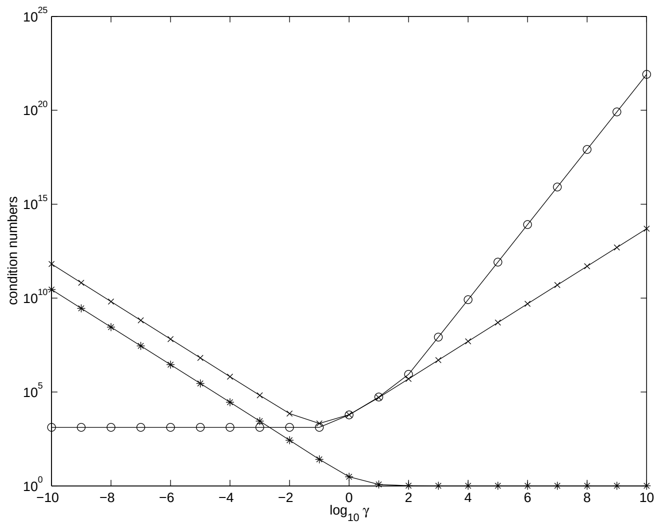

Important observation: $\gamma = \|A\|_2 / \|B\|_2^2$ very close to optimal.

Eigenvalues of preconditioned matrix

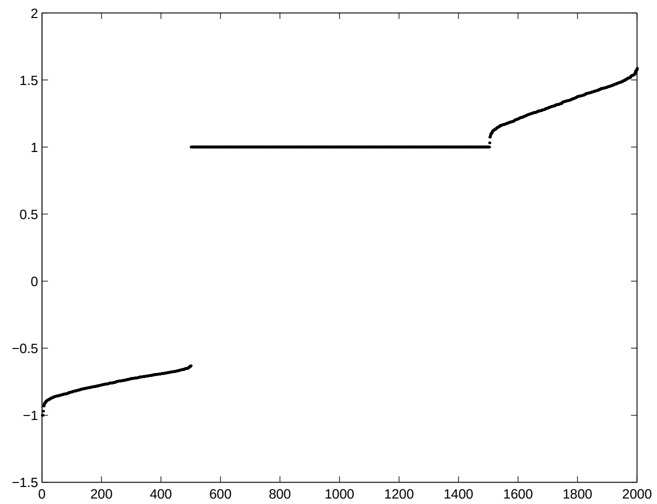

Dependence on regularization parameter

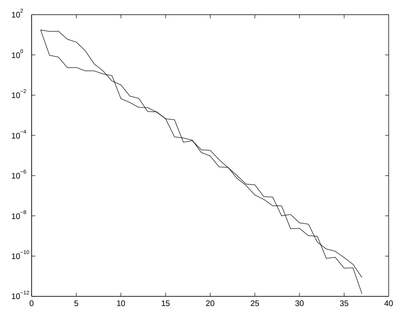

Figure 2: Convergence of preconditioned MINRES with a positive definite block preconditioner, for $\beta = 10^{-4}$ and for $\beta = 10^{-8}$.

Insensitivity to regularization parameter is looking promising. But are we just sweeping all the ill-conditioning under the rug?

Matrix, yet again

$$H = \begin{pmatrix} Q^TQ & 0 & A^T \\ 0 & \beta W^TW & G^T \\ A & G & 0 \end{pmatrix} .$$

A – Second order operator ; G – first order operator.

Two approaches:

1. Take A^{-1} as *weight matrix*.

2. Add $[A; G]$ to first block row: Obtain decoupling (block triangular) but lose symmetry.

Results (un-optimized code)

β	iter1	time1 (sec)	iter2	time2 (sec)
1	14	80.7	14	43.4
1e-2	14	78.2	14	43.8
1e-4	14	72.5	15	51.8
1e-6	15	78.1	17	58.7
1e-8	12	74.4	88	335.5

Table 1: 1970×1970 matrix, with $n = 1241$, m=729, GMRES with tolerance of 10^{-8}.

The End