Linear algebra issues in Interior Point methods for bound-constrained least-squares problems

Stefania Bellavia

Dipartimento di Energetica "S. Stecco"

Università degli Studi di Firenze

Joint work with Jacek Gondzio, and Benedetta Morini

DUE GIORNI DI ALGEBRA LINEARE NUMERICA 2008,
Bologna, 6th-7th Mar 2008

Outline

1. Introduction
   - The problem
   - The Inexact Interior Point Framework
   - Focus on the Linear Algebra Phase
2. The regularized Newton-like method
3. Iterative Linear Algebra
   - The preconditioner
   - Spectral properties
   - PPCG
4. Numerical experimentation

Bound Constrained Least-Squares Problems

$$\min_{l \leq x \leq u} q(x) = \frac{1}{2}\|Ax - b\|_2^2 + \mu\|x\|^2$$

- Vectors $l \in (\mathbb{R} \cup -\infty)^n$ and $u \in (\mathbb{R} \cup \infty)^n$ are lower and upper bounds on the variables.
- $A \in \mathbb{R}^{m \times n}$, $b \in \mathbb{R}^m$, $\mu \geq 0$ are given and $m \geq n$. We expect A to be large and sparse.
- We allow the solution x^* to be degenerate:

$$x_i^* = l_i \text{ or } x_i^* = u_i, \ \nabla q_i(x^*) = 0, \quad \text{for some } i, 1 \leq i \leq n$$

.

- We limit the presentation to NNLS problems:

$$\boxed{\min_{x \geq 0} q(x) = \frac{1}{2}\|Ax - b\|_2^2}$$

- We assume A has full column rank $\Rightarrow$ there is a unique solution x^*.
- Let $g(x) = \nabla q(x) = A^T(Ax - b)$ and $D(x)$ be the diagonal matrix with entries:

$$d_i(x) = \begin{cases} x_i & \text{if} \quad g_i(x) \geq 0 \\ 1 & \text{otherwise} \end{cases}$$

- The core of our procedure is an Inexact Newton-like method applied to the First Order Optimality condition for NNLS:

$$\boxed{D(x)g(x) = 0}$$

Inexact Newton Interior Point methods for $D(x)g(x) = 0$

[Bellavia, Macconi, Morini, NLAA, 2006]

- The method uses ideas of [Heinkenschloss, Ulbrich, Ulbrich, Math. Progr., 1999]
- Let $E(x)$ be the diagonal positive semidefinite matrix with entries:

$$e_i(x) = \begin{cases} g_i(x) & \text{if} \quad 0 \le g_i(x) < x_i^2 \text{ or } g_i(x)^2 > x_i \\ 0 & \text{otherwise .} \end{cases}$$

- Let $W(x)$ and $S(x)$ be the diagonal matrices

$$W(x) = (E(x) + D(x))^{-1} \qquad S(x) = (W(x)D(x))^{\frac{1}{2}}$$

- Note that $(S(x))^2_{i,i} \in (0, 1]$ and $(W(x)E(x))_{i,i} \in [0, 1)$.

k-th iteration

- Solve the s.p.d. system:

$$Z_k \tilde{p}_k = -S_k g_k + r_k, \quad \|r_k\| \leq \eta_k \|W_k D_k g_k\|$$

where $\eta_k \in [0,1)$ and $Z_k \equiv Z(x_k)$ is given by:

$$Z_k = S_k^T (A^T A + D_k^{-1} E_k) S_k = S_k^T A^T A S_k + W_k E_k$$

- Form the step $p_k = S_k \tilde{p}_k$
- Project it onto an interior of the positive orthant:

$$\hat{p}_k = \max\{\sigma, 1 - \|P(x_k + p_k) - x_k\|\} (P(x_k + p_k) - x_k),$$

where $\sigma \in (0,1)$ is close to one.

- Globalization Phase Set:

$$x_{k+1} = x_k + (1-t)\hat{p}_k + tp_k^C \quad t \in [0,1)$$

  - where p_k^C is a constrained Cauchy step.
  - t is chosen to guarantee a sufficient decrease of the objective function $q(x)$.
- Strictly positive iterates
- Eventually $t = 0$ is taken $\Rightarrow$ up to quadratic convergence can be obtained without assuming strict complementarity at x^*.

The Linear Algebra Phase: normal equations

- The system

$$Z_k \tilde{p}_k = -S_k g_k$$

represents the normal equations for the least-squares problem

$$\min_{\tilde{p} \in \mathbb{R}^n} \| B_\delta \tilde{p} + h \|$$

with

$$B_\delta = \begin{pmatrix} AS_k \\ W_k^{\frac{1}{2}} E_k^{\frac{1}{2}} \end{pmatrix}, \qquad h = \begin{pmatrix} Ax_k - b \\ 0 \end{pmatrix}.$$

The Linear Algebra Phase: augmented system

The step $\tilde{p}_k$ can be obtained solving:

$$\underbrace{\begin{pmatrix} I & AS_k \\ S_k A^T & -W_k E_k \end{pmatrix}}_{\mathcal{H}_\delta} \begin{pmatrix} \tilde{q}_k \\ \tilde{p}_k \end{pmatrix} = \begin{pmatrix} -(Ax_k - b) \\ 0 \end{pmatrix}$$

Note that $W_k E_k$ is positive semidefinite and

$$\boxed{v^T W_k E_k v \geq \delta v^T v, \quad \forall v \in \mathbb{R}^n,}$$

where $1 > \delta = \min_i (w_k e_k)_i$.

Conditioning issues

Let $0 < \sigma_1 \leq \sigma_2 \ldots \leq \sigma_n$, be the singular values of AS_k
Assume $\sigma_1 << 1$.

- If $\delta = 0$ then

$$\kappa_2(\mathcal{H}_0) \leq \frac{1+\sigma_n}{\sigma_1^2} \qquad \kappa_2(B_0) \leq \frac{1+\sigma_n}{\sigma_1},$$

  i.e. $\kappa_2(\mathcal{H}_0)$ may be much greater than $\kappa_2(B_0)$.

- If $\delta > 0$ (regularized system), then

$$\kappa_2(\mathcal{H}_\delta) \leq \frac{1+\sigma_n}{\delta} \qquad \kappa_2(B_\delta) \leq \frac{1+\sigma_n}{\sqrt{\delta}},$$

  i.e. If $\delta > \sigma_1$: $\kappa_2(\mathcal{H}_\delta)$ $(\kappa_2(B_\delta))$ may be considerably smaller than $\kappa_2(\mathcal{H}_0)$ $(\kappa_2(B_0))$

The Regularized I.P. Newton-like method

- If σ_1 is not small, the regularization does not deteriorate $\kappa_2(\mathcal{H}_\delta)$ with respect to $\kappa_2(\mathcal{H}_0)$.

$$\Downarrow$$

- Clear benefit from regularization (see also [Saunders, BIT, 1995][Silvester and Wathen, SINUM, 1994])
- Modification of the Affine Scaling I.P. method:

$$\tilde{Z}_k \tilde{p}_k = -S_k g_k + r_k$$

where

$$\begin{aligned} \tilde{Z}_k &= S_k^T (A^T A + D_k^{-1} E_k + \Delta_k) S_k \\ &= \underbrace{S_k^T A^T A S_k + W_k E_k}_{Z_k} + \Delta_k S_k^2 \end{aligned}$$

and Δ_k is diagonal with entries in $[0, 1)$.

- Fast convergence of the method is preserved (in presence of degeneracy, too)
- The globalization strategy of [BMM] can be applied with slight modifications.
- The least square problem and the augmented system are regularized:

$$B_\delta = \begin{pmatrix} AS_k \\ C_k^{\frac{1}{2}} \end{pmatrix}$$

$$\mathcal{H}_\delta = \begin{pmatrix} I & AS_k \\ S_k A^T & -C_k \end{pmatrix}$$

where

$$\boxed{C_k = W_k E_k + \Delta_k S_k^2}$$

Features of the method

- Let $\tau \in (0,1)$ be a small positive threshold and

$$\begin{aligned} \mathcal{I}_k &= \{ i \in \{1,2,\dots,n\},\ s.t.\ (s_k^2)_i \geq 1-\tau \}, \\ \mathcal{A}_k &= \{1,2,\dots,n\}/\mathcal{I}_k, \quad n_1 = card(\mathcal{I}_k), \end{aligned}$$

  then $S_k = diag((S_k)_{\mathcal{I}}, (S_k)_{\mathcal{A}})$

- Note that $S_k^2 + W_k E_k = I$. When x_k converges to x^*,

$$\begin{array}{ll} \lim_{k\to\infty}(S_k)_{\mathcal{I}} = I, & \lim_{k\to\infty}(S_k)_{\mathcal{A}} = 0. \\ \lim_{k\to\infty}(W_k E_k)_{\mathcal{I}} = 0, & \lim_{k\to\infty}(W_k E_k)_{\mathcal{A}} = I. \end{array}$$

- $\mathcal{I}_k$ is expected to eventually settle down (inactive components and possibly degenerate components)

Solving the augmented system

- The following partition on the augmented system is induced:

$$\begin{pmatrix} I & A_{\mathcal{I}}(S_k)_{\mathcal{I}} & A_{\mathcal{A}}(S_k)_{\mathcal{A}} \\ (S_k)_{\mathcal{I}}A_{\mathcal{I}}^T & -(C_k)_{\mathcal{I}} & 0 \\ (S_k)_{\mathcal{A}}A_{\mathcal{A}}^T & 0 & -(C_k)_{\mathcal{A}} \end{pmatrix} \begin{pmatrix} \tilde{q}_k \\ (\tilde{p}_k)_{\mathcal{I}} \\ (\tilde{p}_k)_{\mathcal{A}} \end{pmatrix} = \begin{pmatrix} -(Ax_k - b) \\ 0 \\ 0 \end{pmatrix}$$

- Eliminating $(\tilde{p}_k)_{\mathcal{A}}$ we get

$$\underbrace{\begin{pmatrix} I + Q_k & A_{\mathcal{I}}(S_k)_{\mathcal{I}} \\ (S_k)_{\mathcal{I}}A_{\mathcal{I}}^T & -(C_k)_{\mathcal{I}} \end{pmatrix}}_{\mathcal{H}_k} \begin{pmatrix} \tilde{q}_k \\ (\tilde{p}_k)_{\mathcal{I}} \end{pmatrix} = \begin{pmatrix} -(Ax_k - b) \\ 0 \end{pmatrix}$$

- $\mathcal{H}_k \in \mathbb{R}^{(m+n_1)\times(m+n_1)}$

The Preconditioner

- Note that

$$\mathcal{H}_k = \underbrace{\begin{pmatrix} I & A_{\mathcal{I}}(S_k)_{\mathcal{I}} \\ (S_k)_{\mathcal{I}} A_{\mathcal{I}}^T & -(\Delta_k S_k^2)_{\mathcal{I}} \end{pmatrix}}_{\mathcal{P}_k} + \begin{pmatrix} Q_k & 0 \\ 0 & -(W_k E_k)_{\mathcal{I}} \end{pmatrix}$$

  where $Q_k = A_{\mathcal{A}}(S_k C_k^{-1} S_k)_{\mathcal{A}} A_{\mathcal{A}}^T$

- When x_k converges to x^*, $(S_k)_{\mathcal{A}} \rightarrow 0$, $(C_k)_{\mathcal{A}} \rightarrow I$, then

$$\lim_{k\rightarrow\infty} (Q_k) = 0, \qquad \lim_{k\rightarrow\infty} (W_k E_k)_{\mathcal{I}} = 0.$$

Factorization of the Preconditioner

$$\mathcal{P}_k = \begin{pmatrix} I & A_{\mathcal{I}}(S_k)_{\mathcal{I}} \\ (S_k)_{\mathcal{I}} A_{\mathcal{I}}^T & -(\Delta_k S_k^2)_{\mathcal{I}} \end{pmatrix}$$

can be factorized as

$$\mathcal{P}_k = \begin{pmatrix} I & 0 \\ 0 & (S_k)_{\mathcal{I}} \end{pmatrix} \underbrace{\begin{pmatrix} I & A_{\mathcal{I}} \\ A_{\mathcal{I}}^T & -(\Delta_k)_{\mathcal{I}} \end{pmatrix}}_{\Pi_k} \begin{pmatrix} I & 0 \\ 0 & (S_k)_{\mathcal{I}} \end{pmatrix}$$

- If $\mathcal{I}_k$ and Δ_k remain unchanged for a few iterations, the factorization of matrix Π_k does not have to be updated.
- $\mathcal{I}_k$ is expected to eventually settle down.

Eigenvalues

- $\mathcal{P}_k^{-1}\mathcal{H}_k$ has
  - at least $m - n + n_1$ eigenvalues at 1
  - the other eigenvalues are positive and of the form

$$\lambda = 1 + \mu, \quad \mu = \frac{u^T Q_k u + v^T (W_k E_k)_{\mathcal{I}}\, v}{u^T u + v^T (\Delta_k S_k^2)_{\mathcal{I}}\, v},$$

  where $(u^T, v^T)^T$ is an eigenvector associated to λ.
- if μ is small: the eigenvalues of $\mathcal{P}_k^{-1}\mathcal{H}_k$ are clustered around one. This is the case when x_k is close to the solution.

Eigenvalues (x_k far away from x^*)

- The eigenvalues of $\mathcal{P}_k^{-1}\mathcal{H}_k$ have the form $\lambda = 1 + \mu$ and
  - If $(\Delta_k)_{i,i} = \delta > 0$ for $i \in \mathcal{I}_k$,

$$\mu \leq \frac{\|A_{\mathcal{A}}(S_k)_{\mathcal{A}}\|^2}{\tau} + \frac{\tau}{\delta(1-\tau)}$$

  - If $(\Delta_k)_{i,i} = \begin{cases} (w_k)_i(e_k)_i & \text{for } i \in \mathcal{I}_k \text{ and } (w_k)_i(e_k)_i \neq 0 \\ \delta > 0 & \text{for } i \in \mathcal{I}_k \text{ and } (w_k)_i(e_k)_i = 0 \end{cases}$

$$\mu \leq \frac{\|A_{\mathcal{A}}(S_k)_{\mathcal{A}}\|^2}{\tau} + \frac{1}{1-\tau}$$

  $\Rightarrow$ Better distribution of the eigenvalues.

- A scaling of A at the beginning of the process is advisable.

Solving the augmented system by PPCG

- We can adopt the Projected Preconditioned Conjugate-Gradient (PPCG) [Gould, 1999], [Dollar, Gould, Schilders, Wathen, SIMAX, 2006]
- It is a CG procedure for solving indefinite systems:

$$\begin{pmatrix} H & A \\ A^T & -C \end{pmatrix} \begin{pmatrix} p \\ q \end{pmatrix} = \begin{pmatrix} -g \\ 0 \end{pmatrix}$$

with $H \in \mathbb{R}^{m \times m}$ symmetric, $C \in \mathbb{R}^{n \times n}$ $(n \leq m)$ symmetric, $A \in \mathbb{R}^{m \times n}$ full rank, using preconditioners of the form:

$$\begin{pmatrix} G & A \\ A^T & -T \end{pmatrix}$$

with G symmetric, T nonsingular.

- When C is nonsingular, PPCG is equivalent to applying PCG to the system

$$(H + AC^{-1}A^T)p = g$$

with preconditioner:

$$G + AT^{-1}A^T$$

- In our case, it is equivalent to applying PCG to the system:

$$\underbrace{(I + Q_k + A_{\mathcal{I}}(S_k C_k^{-1} S_k)_{\mathcal{I}} A_{\mathcal{I}}^T)}_{\mathcal{F}_k} \tilde{q}_k = -(Ax_k - b),$$

using ther preconditioner

$$\mathcal{G}_k = I + A_{\mathcal{I}}(\Delta_k)_{\mathcal{I}}^{-1} A_{\mathcal{I}}^T,$$

Eigenvalues of $\mathcal{G}_k^{-1}\mathcal{F}_k$

.

- If $(\Delta_k)_{i,i} = (w_k)_i(e_k)_i$ for $i \in \mathcal{I}_k$, then the eigenvalues of $\mathcal{G}_k^{-1}\mathcal{F}_k$ satisfy:

$$1 - \frac{1}{2-\tau} \leq \lambda \leq 1 + \frac{\|A_{\mathcal{A}}(S_k)_{\mathcal{A}}\|^2}{\tau}.$$

- Drawback: Differently from the previous results, no cluster of eigenvalues at 1 is guaranteed
- Advantage: PPCG is characterized by a minimization property and requires a fixed amount of work per iteration

Implementation issues

- Dynamic regularization:

$$(\Delta_k)_{i,i} = \begin{cases} 0, & \text{if} \quad i \notin \mathcal{I}_k \quad (i.e. (w_k)_i (e_k)_i > \tau) \\ \min\{\max\{10^{-3}, (w_k)_i (e_k)_i\}, 10^{-2}\}, & \text{otherwise.} \end{cases}$$

- Iterative solver: PPCG with adaptive choice of the tolerance in the stopping criterion.
  - Linear systems are solved with accuracy that increases as the solution is approached.
  - PPCG is stopped when the preconditioned residual drops below

$$tol = \max(10^{-7}, \frac{\eta_k \|W_k D_k g_k\|}{\|A^T S_k\|_1}).$$

- To avoid preconditioner factorizations: at iteration $k+1$ freeze the set $\mathcal{I}_k$ and the matrix Δ_k if

$$\#(IT\_PPCG)_k \leq 30 \quad \& \quad |card(\mathcal{I}_{k+1}) - card(\mathcal{I}_k)| \leq 10.$$

- If $\mathcal{I}_k$ is empty (i.e. $\|S_k\| \leq 1-\tau$):
  - we apply PCG to the normal system

$$(S_k^T A^T A S_k + C_k)\tilde{p}_k = -S_k A^T (Ax_k - b).$$

- `Matlab` code, $\epsilon_m = 2.\ 10^{-16}$.
- The threshold τ is set to 0.1
- Initial guess $x_0 = (1, \dots, 1)^T$.
- Succesfull termination:

$$\begin{cases} q_{k-1} - q_k < \epsilon\ (1 + q_{k-1}), \\ \|x_k - x_{k-1}\|_2 \leq \sqrt{\epsilon}\ (\ 1 + \|x_k\|_2\ ) \\ \|P(x_k - g_k) - x_k\|_2 < \epsilon^{\frac{1}{3}}\ (\ 1 + \|g_k\|_2\ ) \end{cases}$$

or

$$\|P(x_k - g_k) - x_k\|_2 \leq \epsilon$$

with $\epsilon = 10^{-9}$.

- A failure is declared after 100 iterations.

Test Problems

- The matrix A is the transpose of the matrices in the `LPnetlib` subset of The University of Florida Sparse Matrix Collection. We discarded the matrices with $m < 1000$ and the matrices that are not full rank.
- A total of 56 matrices.
- Dimensions ranges up to 10^5
- The vector b is set equal to $b = -A(1, 1, \ldots, 1)^T$
- When $\|A\|_1 > 10^3$, we scaled the matrix using a simple row and column scaling scheme.

Numerical Results

- On a total of 56 test problems we succesfully solve 51 tests:
- 41 test problems are solved with less than 20 nonlinear iterations.
- In 40 tests the average number of PPCG iterations does not exceed 40.
- In 8 tests the solution is the null vector. At each iteration $\mathcal{I}_k = \emptyset$, $S_k^T A^T A S_k + C_k \simeq I$ and the convergence of the linear solver is very fast.

Savings in the number of preconditioner factorizations

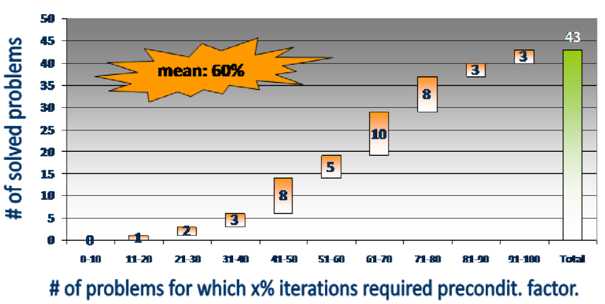

Percent Reduction in the dimension n

- We solve augmented system of reduced dimension $m + n_1$

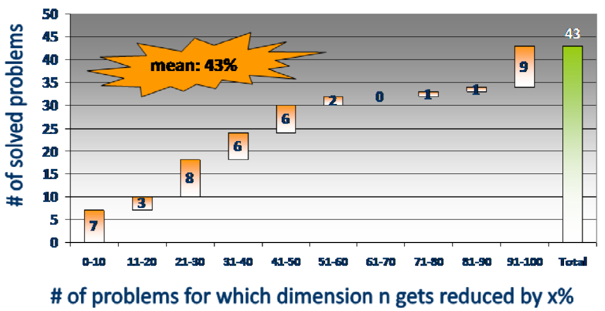

Future work

- More experimentation, using also QMR and GMRES
- Develop a code for the more general problem:

$$\min_{l \le x \le u} q(x) = \frac{1}{2}\|Ax - b\|_2^2 + \mu\|x\|^2$$

If $\mu > 0$:

  - A may also be rank deficient
  - the augmented systems are regularized "naturally"

- Comparison with existing codes (e.g. BCLS (Fiedlander), PDCO (Saunders))