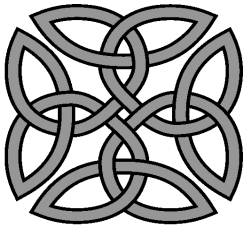

School of Mathematics

Interior Point Methods: Inexact Newton Directions, Conjugate Gradient & Preconditioners

Jacek Gondzio
Email: J.Gondzio@ed.ac.uk
URL: http://www.maths.ed.ac.uk/~gondzio

Three Lectures

1. **IPMs: Motivation and Linear Algebra**

2. **Inexact Newton Directions, Conjugate Gradient and Preconditioners**

3. **Example Applications: Design of Efficient Preconditioners**

Lecture 1

IPMs: Motivation & Linear Algebra

Outline

- **IPM Motivation**
  - Lagrangian duality, logarithmic barrier
  - first order optimality conditions, Newton method
- **Polynomial Complexity of IPM**
  - primal-dual framework
  - following the central path
  - key ideas in the proof of polynomial complexity
- **Linear Algebra in IPM**
  - from LP, via QP to NLP
  - definite, indefinite and quasidefinite systems
  - Cholesky factorization
  - exploiting sparsity

Primal-Dual Pair of Linear Programs

Primal

$$\begin{aligned} \min \quad & c^T x \\ \text{s.t.} \quad & Ax = b, \\ & x \geq 0; \end{aligned}$$

Dual

$$\begin{aligned} \max \quad & b^T y \\ \text{s.t.} \quad & A^T y + s = c, \\ & s \geq 0. \end{aligned}$$

Lagrangian

$$L(x, y) = c^T x - y^T (Ax - b) - s^T x.$$

Optimality Conditions

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe &= 0, \quad (\text{ i.e., } x_j \cdot s_j = 0 \quad \forall j), \\ (x, s) &\geq 0, \end{aligned}$$

$X = diag\{x_1, \cdots, x_n\}$, $S = diag\{s_1, \cdots, s_n\}$, $e = (1, \cdots, 1) \in \mathcal{R}^n$.

Logarithmic barrier

$$-\ln x_j$$

"replaces" the inequality

$$x_j \geq 0 \ .$$

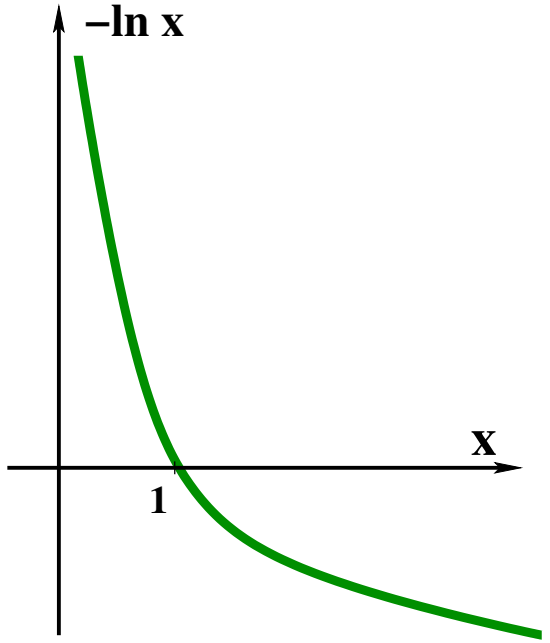

Observe that

$$\min \ \mathrm{e}^{-\sum_{j=1}^{n} \ln x_j} \quad \Longleftrightarrow \quad \max \prod_{j=1}^{n} x_j$$

The minimization of $-\sum_{j=1}^{n} \ln x_j$ is equivalent to the maximization of the product of distances from all hyperplanes defining the positive orthant: it prevents all x_j from approaching zero.

Logarithmic barrier

Replace the **primal** LP

$$\begin{aligned} \min \quad & c^T x \\ \text{s.t.} \quad & Ax = b, \\ & x \geq 0, \end{aligned}$$

with the **primal barrier program**

$$\begin{aligned} \min \quad & c^T x - \mu \sum_{j=1}^{n} \ln x_j \\ \text{s.t.} \quad & Ax = b. \end{aligned}$$

Lagrangian: $$L(x, y, \mu) = c^T x - y^T (Ax - b) - \mu \sum_{j=1}^{n} \ln x_j.$$

Conditions for a stationary point of the Lagrangian

$$\begin{aligned} \nabla_x L(x,y,\mu) &= c - A^T y - \mu X^{-1} e = 0 \\ \nabla_y L(x,y,\mu) &= \qquad\qquad Ax - b = 0, \end{aligned}$$

where $X^{-1} = diag\{x_1^{-1}, x_2^{-1}, \cdots, x_n^{-1}\}$.

Let us denote

$$s = \mu X^{-1} e, \quad \text{i.e.} \quad XSe = \mu e.$$

The **First Order Optimality Conditions** are:

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe &= \mu e, \\ (x,s) &> 0. \end{aligned}$$

The pronunciation of Greek letter μ [mi]

Robert De Niro, Taxi Driver (1976)

Central Trajectory

The first order optimality conditions for the barrier problem

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe &= \mu e, \\ (x, s) &\geq 0 \end{aligned}$$

approximate the first order optimality conditions for the LP

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe &= 0, \\ (x, s) &\geq 0 \end{aligned}$$

more and more closely as μ goes to zero.

Central Trajectory

Parameter μ controls the distance to optimality.

$$c^T x - b^T y = c^T x - x^T A^T y = x^T (c - A^T y) = x^T s = n\mu.$$

Analytic centre (μ-centre): a (unique) point

$$(x(\mu), y(\mu), s(\mu)), \quad x(\mu) > 0, \ s(\mu) > 0$$

that satisfies FOC.

The path

$$\{(x(\mu), y(\mu), s(\mu)) : \mu > 0\}$$

is called the **primal-dual central trajectory**.

Newton Method

is used to find a stationary point of the barrier problem.

Recall how to use Newton Method to find a root of a nonlinear equation

$$f(x) = 0.$$

A tangent line

$$z - f(x^k) = \nabla f(x^k) \cdot (x - x^k)$$

is a local approximation of the graph of the function $f(x)$.
Substituting $z = 0$ defines a new point

$$x^{k+1} = x^k - (\nabla f(x^k))^{-1} f(x^k).$$

Newton Method

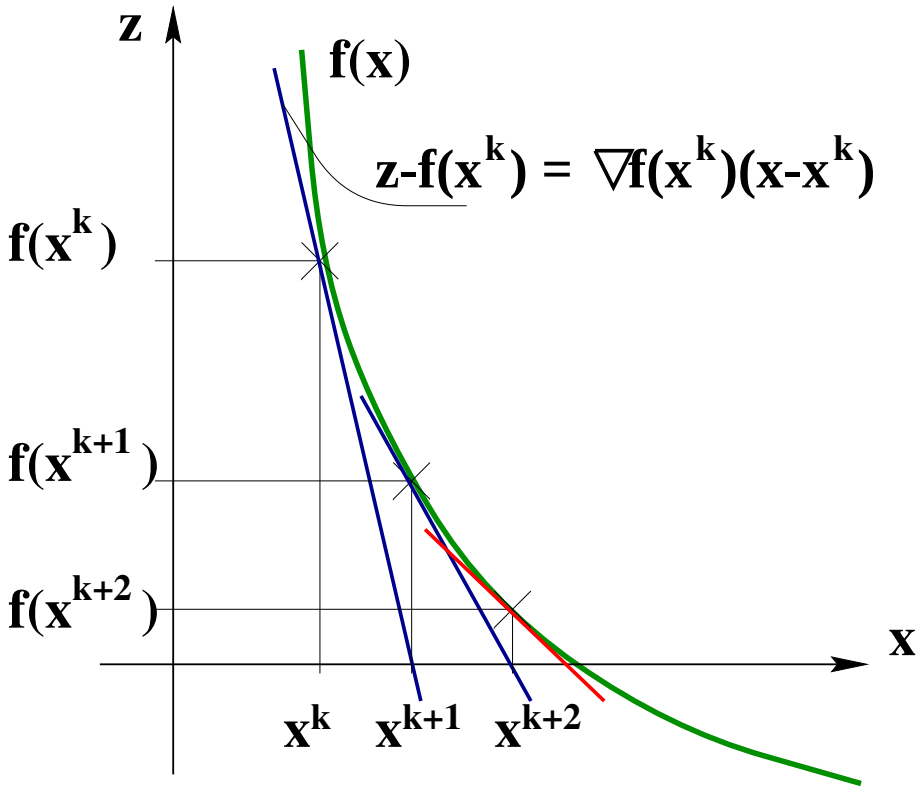

$$x^{k+1} = x^k - (\nabla f(x^k))^{-1} f(x^k).$$

Apply Newton Method to the FOC

The first order optimality conditions for the barrier problem form a large system of nonlinear equations

$$f(x, y, s) = 0,$$

where $f : \mathcal{R}^{2n+m} \mapsto \mathcal{R}^{2n+m}$ is a mapping defined as follows:

$$f(x, y, s) = \begin{bmatrix} Ax - b \\ A^T y + s - c \\ XSe - \mu e \end{bmatrix}.$$

Actually, the first two terms of it are **linear**; only the last one, corresponding to the complementarity condition, is **nonlinear**.

Newton Method (cont'd)

Note that

$$\nabla f(x,y,s) = \begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S & 0 & X \end{bmatrix}.$$

Thus, for a given point (x,y,s) we find the Newton direction $(\Delta x, \Delta y, \Delta s)$ by solving the system of linear equations:

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S & 0 & X \end{bmatrix} \cdot \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} b - Ax \\ c - A^T y - s \\ \mu e - XSe \end{bmatrix}.$$

Interior-Point Framework

The **logarithmic barrier**

$$-\ln x_j$$

"replaces" the inequality

$$x_j \geq 0.$$

We derive the **first order optimality conditions** for the primal barrier problem:

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe &= \mu e, \end{aligned}$$

and apply **Newton method** to solve this system of (nonlinear) equations.

Actually, we fix the barrier parameter μ and make only **one** (damped) Newton step towards the solution of FOC. We do not solve the current FOC exactly. Instead, we immediately reduce the barrier parameter μ (to ensure progress towards optimality) and repeat the process.

Interior Point Algorithm

Initialize

$k = 0$ $\qquad (x^0, y^0, s^0) \in \mathcal{F}^0$

$\mu_0 = \frac{1}{n} \cdot (x^0)^T s^0$ $\qquad \alpha_0 = 0.9995$

Repeat until optimality

$k = k + 1$

$\mu_k = \sigma \mu_{k-1}$, where $\sigma \in (0, 1)$

$\Delta = (\Delta x, \Delta y, \Delta s) =$ Newton direction towards μ-centre

Ratio test:

$$\alpha_P := \max \{\alpha > 0 : \ x + \alpha \Delta x \geq 0\},$$
$$\alpha_D := \max \{\alpha > 0 : \ s + \alpha \Delta s \geq 0\}.$$

Make step:

$$x^{k+1} = x^k + \alpha_0 \alpha_P \Delta x,$$
$$y^{k+1} = y^k + \alpha_0 \alpha_D \Delta y,$$
$$s^{k+1} = s^k + \alpha_0 \alpha_D \Delta s.$$

Central Path Neighbourhood

Assume a primal-dual strictly feasible solution $(x, y, s) \in \mathcal{F}^0$ lying in a neighbourhood of the central path is given; namely (x, y, s) satisfies:

$$\begin{aligned} Ax &= b, \\ A^T y + s &= c, \\ XSe &\approx \mu e. \end{aligned}$$

We define a θ**-neighbourhood** of the central path $N_2(\theta)$, a set of primal-dual strictly feasible solutions $(x, y, s) \in \mathcal{F}^0$ that satisfy:

$$\|XSe - \mu e\| \leq \theta\mu,$$

where $\theta \in (0, 1)$ and the barrier μ satisfies:

$$x^T s = n\mu.$$

Hence $N_2(\theta) = \{(x, y, s) \in \mathcal{F}^0 \mid \|XSe - \mu e\| \leq \theta\mu\}$.

Central Path Neighbourhood

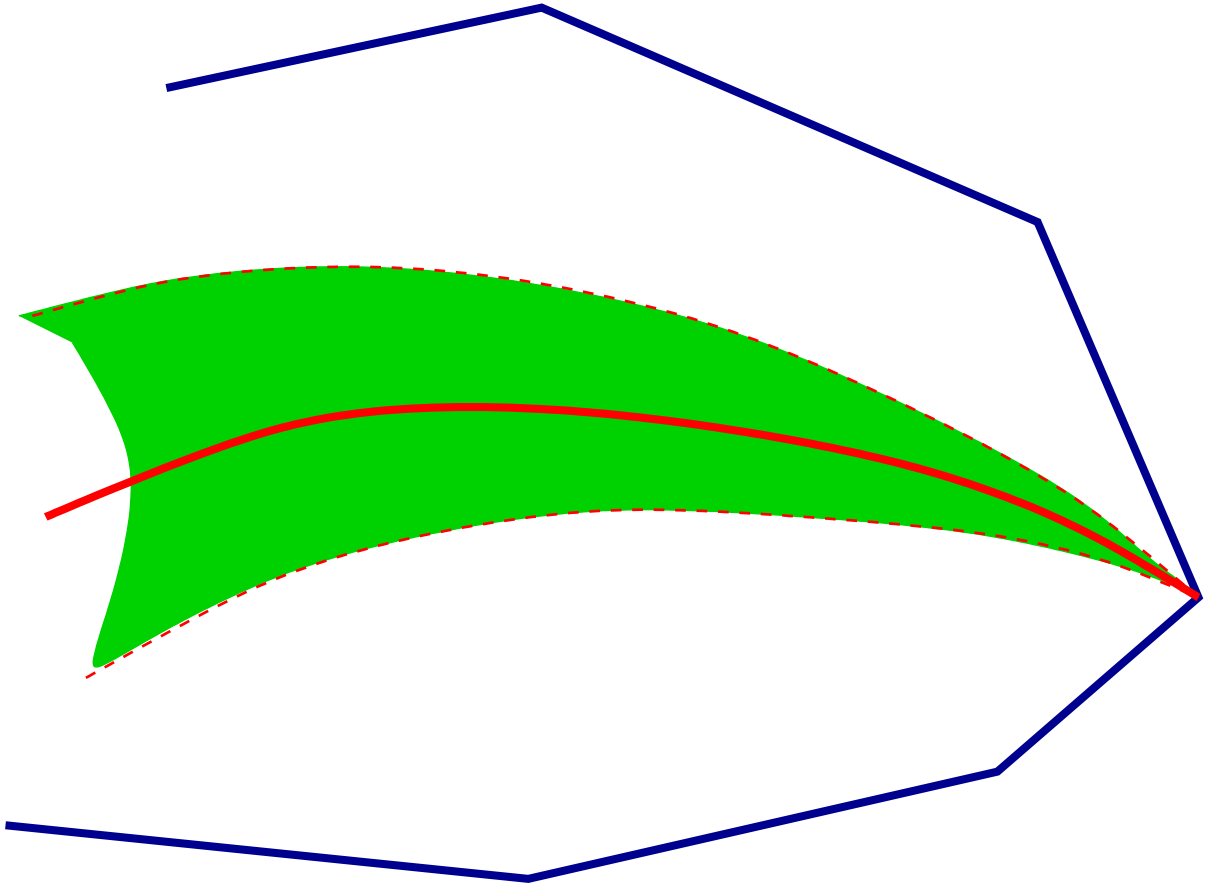

$N_2(\theta)$ neighbourhood of the central path

Progress towards optimality

Assume a primal-dual strictly feasible solution $(x, y, s) \in N_2(\theta)$ for some $\theta \in (0, 1)$ is given.

Interior point algorithm tries to move from this point to another one that also belongs to the θ-neighbourhood of the central path but corresponds to a smaller μ. The required reduction of μ is small:

$$\mu^{k+1} = \sigma \mu^k, \qquad \text{where} \qquad \sigma = 1 - \beta/\sqrt{n},$$

for some $\beta \in (0, 1)$.

This is a **short-step** method:
It makes short steps to optimality.

Progress towards optimality

Given a new μ-centre, interior point algorithm computes Newton direction:

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S & 0 & X \end{bmatrix} \cdot \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \sigma\mu e - XSe \end{bmatrix},$$

and makes step in this direction.

Magic numbers:

$$\theta = 0.1 \quad \text{and} \quad \beta = 0.1.$$

θ controls the proximity to the central path;
β controls the progress to optimality.

How to prove the $\mathcal{O}(\sqrt{n})$ complexity result

One has to prove the following:

- full step in Newton direction is feasible;
- the new iterate
  $(x^{k+1}, y^{k+1}, s^{k+1}) = (x^k, y^k, s^k) + (\Delta x^k, \Delta y^k, \Delta s^k)$
  belongs to the θ-neighbourhood of the new μ-centre
  (with $\mu^{k+1} = \sigma\mu^k$);
- duality gap is reduced $1 - \beta/\sqrt{n}$ times.

$\mathcal{O}(\sqrt{n})$ complexity result

Note that since at one iteration duality gap is reduced $1-\beta/\sqrt{n}$ times, after $\sqrt{n}$ iterations the reduction achieves:

$$(1-\beta/\sqrt{n})^{\sqrt{n}} \approx e^{-\beta}.$$

After $C \cdot \sqrt{n}$ iterations, the reduction is $e^{-C\beta}$. For sufficiently large constant C the reduction can thus be arbitrarily large (i.e. the duality gap can become arbitrarily small).

Hence this algorithm has complexity $\mathcal{O}(\sqrt{n})$.

This should be understood as follows:

> *"after the number of iterations proportional to $\sqrt{n}$ the algorithm solves the problem".*

Reading about IPMs

S. Wright
Primal-Dual Interior-Point Methods, SIAM Philadelphia, 1997.

Gondzio
Interior point methods 25 years later,
European J. of Operational Research 218 (2012) 587–601.
`http://www.maths.ed.ac.uk/~gondzio/reports/ipmXXV.html`

Gondzio and Grothey
Direct solution of linear systems of size 10^9 arising in optimization with interior point methods, in: *Parallel Processing and Applied Mathematics PPAM 2005*, R. Wyrzykowski, J. Dongarra, N. Meyer and J. Wasniewski (eds.), *Lecture Notes in Computer Science*, 3911, Springer-Verlag, Berlin, 2006, pp 513–525.

OOPS: Object-Oriented Parallel Solver

`http://www.maths.ed.ac.uk/~gondzio/parallel/solver.html`

Linear Algebra in IPM (Newton Directions)

LP
$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} \xi_p \\ \xi_d \\ \xi_\mu \end{bmatrix},$$

QP
$$\begin{bmatrix} A & 0 & 0 \\ -Q & A^T & I \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} \xi_p \\ \xi_d \\ \xi_\mu \end{bmatrix},$$

NLP
$$\begin{bmatrix} A(x) & 0 & I \\ -Q(x,y) & A(x)^T & 0 \\ 0 & Z & Y \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta z \end{bmatrix} = \begin{bmatrix} \xi_p \\ \xi_d \\ \xi_\mu \end{bmatrix}.$$

Take QP case for example:

Newton direction

$$\begin{bmatrix} A & 0 & 0 \\ -Q & A^T & I \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} \xi_p \\ \xi_d \\ \xi_\mu \end{bmatrix},$$

where

$$\begin{aligned} \xi_p &= b - Ax, \\ \xi_d &= c - A^T y - s + Qx, \\ \xi_\mu &= \mu e - XSe. \end{aligned}$$

Eliminate Δs to get the **Augmented System**

$$\begin{bmatrix} -Q - \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} \xi_d - X^{-1}\xi_\mu \\ \xi_p \end{bmatrix}.$$

LP case corresponds to $Q = 0$.

Augmented System vs Normal Equations

Augmented system in **QP**

$$\begin{bmatrix} -Q - \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} \xi_d - X^{-1}\xi_\mu \\ \xi_p \end{bmatrix}.$$

Eliminate Δx from the first equation and get normal equations

$$(A(Q + \Theta^{-1})^{-1}A^T)\Delta y \;=\; g.$$

One can use normal equations in LP, but not in QP. Normal equations in QP may become almost completely dense even for sparse matrices A and Q. Thus, in QP, usually the indefinite augmented system form is used.

KKT systems in IPMs for LP, QP and NLP

LP
$$\begin{bmatrix} \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

QP
$$\begin{bmatrix} Q + \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

NLP
$$\begin{bmatrix} Q(x,y) + \Theta_P^{-1} & A(x)^T \\ A(x) & -\Theta_D \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

Matrices Θ, Θ_P, Θ_D are very ill-conditioned.

Cholesky factorization

Compute a decomposition

$$LDL^T = A\Theta A^T.$$

where:
L is a lower triangular matrix; and
D is a diagonal matrix.

Cholesky factorization is simply the **Gaussian Elimination** process that exploits two properties of the matrix:

- symmetry;
- positive definiteness.

Definite & Indefinite Systems

IPMs:

For indefinite *augmented system*

$$\begin{bmatrix} -\Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} r \\ h \end{bmatrix}.$$

one needs to use some **special tricks**.

For positive definite *normal equations*

$$(A\Theta A^T)\Delta y \;=\; g.$$

one can compute the **Cholesky factorization**.

Major Cholesky

Andre-Louis Cholesky (1875-1918)
Major of French Army,
descendant from the Cholewski family of Polish imigrants.

Read: **M. A. Saunders**, Major Cholesky would feel proud,
ORSA Journal on Computing, vol 6 (1994) No 1, pp 23–27.

Symmetric Factorization

Two step solution method:

- factorization to LDL^T form,
- backsolve to compute direction Δy.

A symmetric nonsingular matrix H is **factorizable** if there exists a diagonal matrix D and unit lower triangular matrix L such that $H = LDL^T$.

A symmetric matrix H is **strongly factorizable** if for any permutation matrix P a factorization $PHP^T = LDL^T$ exists.

The general symmetric indefinite matrix **is not factorizable**.

Factoring Indefinite Matrix

Two options are possible:

1. Replace diagonal matrix D with a block-diagonal one and allow 2×2 (indefinite) pivots

$$\begin{bmatrix} 0 & a \\ a & 0 \end{bmatrix} \quad \text{and} \quad \begin{bmatrix} 0 & a \\ a & d \end{bmatrix}.$$

Hence obtain a decomp. $H = LDL^T$ with **block-diagonal** D.

2. Regularize indefinite matrix to produce a **quasidefinite** matrix

$$K = \begin{bmatrix} -E & A^T \\ A & F \end{bmatrix},$$

where
$E \in \mathcal{R}^{n \times n}$ is positive definite,
$F \in \mathcal{R}^{m \times m}$ is positive definite, and
$A \in \mathcal{R}^{m \times n}$ has full row rank.

Quasidefinite (QDF) Matrices

Symmetric matrix is called **quasidefinite** if

$$K = \begin{bmatrix} -E & A^T \\ A & F \end{bmatrix},$$

where $E \in \mathcal{R}^{n \times n}$ and
$F \in \mathcal{R}^{m \times m}$ are positive definite, and
$A \in \mathcal{R}^{m \times n}$ has full row rank.

Vanderbei (SIOPT, 1995, pp. 100–113) proved that QDF matrices are **strongly factorizable**. For any quasidefinite matrix there exists a **Cholesky-like** factorization

$$K = LDL^T,$$

where
D is **diagonal** but **not positive definite**:
n negative pivots; and m positive pivots.

From Indefinite to Quasidefinite

Indefinite matrix

$$H = \begin{bmatrix} -Q - \Theta^{-1} & A^T \\ A & 0 \end{bmatrix}.$$

in IPMs can be converted to a quasidefinite one.
Regularize indefinite matrix to produce a **quasi-definite** matrix.
Use **dynamic regularization**

$$\bar{H} = \begin{bmatrix} -Q - \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} + \begin{bmatrix} -R_p & 0 \\ 0 & R_d \end{bmatrix},$$

where $R_p \in \mathcal{R}^{n \times n}$ and $R_d \in \mathcal{R}^{m \times m}$ are the *primal* and *dual* regularizations. For any quasidefinite matrix there exists a Cholesky-like factorization

$$\bar{H} = LDL^T,$$

where D is **diagonal** but **not positive definite**:
n negative pivots and m positive pivots.

Large Problems are Sparse

Suppose a large LP is solved: $m, n \sim 10^6$ or larger.
Can all variables be linked at the same time?
No, usually only a subset of them is linked.

There are usually only *several* nonzeros per row in an LP.
Large problems are almost always **sparse**.
Exploiting sparsity in computations leads to huge savings.
Exploiting sparsity means mainly avoiding doing useless computations: the computations for which the result is known, as for example multiplications with zero.

General Sparse Systems

Single step in Gaussian Elimination

$$A = \begin{bmatrix} \mathbf{p} & v^T \\ u & A_1 \end{bmatrix}$$

produces the following Schur complement

$$A_1 - p^{-1} u v^T.$$

Markowitz Pivot Choice

Let r_i and $c_i, i = 1, 2, ..., n$ be numbers of nonzero entries in row and column i, respectively. The elimination of the pivot a_{ij} needs

$$f_{ij} = (r_i - 1)(c_j - 1)$$

flops to be made. This step creates at most f_{ij} new nonzero entries in the Schur complement.

General Sparse Systems

The effect of pivot elimination on the sparsity pattern

	1	2	3	4	5	6	7	8
1	**p**			x	x		x	x
2		x				x	x	x
3	x	x	x		x			
4			x	x		x		
5	x				x	x		
6				x			x	
7		x			x	x		
8	x				x		x	x

pivot : **p**
nonzero : x

	1	2	3	4	5	6	7	8
1	**p**			**x**	**x**		**x**	**x**
2		x				x	x	x
3	**x**	x	x	**f**	**f**		**f**	**f**
4			x	x		x		
5	**x**			**f**	**f**	x	**f**	**f**
6				x			x	
7		x			x	x		
8	**x**			**f**	**f**		**f**	**f**

pivot : **p**
nonzero : x
fill − in : **f**

Markowitz Pivot Choice: Example

Markowitz: Choose the pivot with $\min_{i,j} f_{ij}$.

	1	2	3	4	5	6	7	8
1	x			x	x		x	x
2		x				x	x	x
3	x	x	x		x			
4			x	x		x		
5	x				x	x		
6				**p**			x	
7		x			x	x		
8	x				x		x	x

	1	2	3	4	5	6	7	8
1	x			**x**	x		**f**	x
2		x				x	x	x
3	x	x	x		x			
4			x	**x**		x	**f**	
5	x				x	x		
6				**p**			**x**	
7		x			x	x		
8	x				x		x	x

Minimum Degree Ordering (MDO)

In symmetric positive definite case:
pivots are chosen from the diagonal and $r_i = c_i$
hence choose the pivot with $\min_i \ r_i$

Minimum degree ordering:
choose an element with the minimum number of nonzeros in a row,

that is, choose a node with the minimum number of neighbours
(a node with the *minimum degree*)
in a graph related to sparsity pattern of the matrix.

Minimum Degree Ordering (MDO)

Sparse Matrix **Pivot h_{11}** **Pivot h_{22}**

$$H = \begin{bmatrix} x & & x & x & x & \\ & x & & & x & \\ x & & x & & & x \\ x & & & x & & x \\ x & x & & & x & \\ & & x & x & & x \end{bmatrix} \quad \begin{bmatrix} \mathbf{p} & & \mathbf{x} & \mathbf{x} & \mathbf{x} & \\ & x & & & x & \\ \mathbf{x} & & x & \mathbf{f} & \mathbf{f} & x \\ \mathbf{x} & & \mathbf{f} & x & \mathbf{f} & x \\ \mathbf{x} & x & \mathbf{f} & \mathbf{f} & x & \\ & & x & x & & x \end{bmatrix} \quad \begin{bmatrix} x & & x & x & x & \\ & \mathbf{p} & & & \mathbf{x} & \\ x & & x & & & x \\ x & & & x & & x \\ x & \mathbf{x} & & & x & \\ & & x & x & & x \end{bmatrix}$$

Minimum degree ordering:
choose a diagonal element corresponding to a row with the *minimum* number of nonzeros.
Permute rows and columns of H accordingly.

MDO is simply the symmetric version of Markowitz pivot rule.

Nested Dissection:

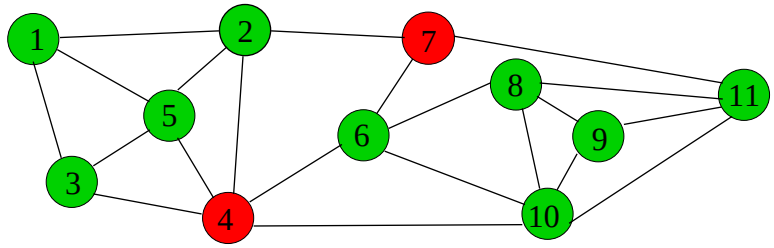

Original Matrix

	1	2	3	4	5	6	7	8	9	10	11
1	x	x	x		x						
2	x	x		x	x		x				
3	x		x	x	x						
4		x	x	x	x	x				x	
5	x	x	x	x	x						
6				x		x	x	x		x	
7		x				x	x				x
8						x		x	x	x	x
9								x	x	x	x
10				x		x		x	x	x	x
11							x	x	x	x	x

Reordered Matrix

	1	2	3	5	6	8	9	10	11	4	7
1	x	x	x	x							
2	x	x		x						**x**	**x**
3	x		x	x						**x**	
5	x	x	x	x						**x**	
6					x	x		x		**x**	**x**
8					x	x	x	x	x		
9						x	x	x	x		
10					x	x	x	x	x	**x**	
11						x	x	x	x		**x**
4		**x**	**x**	**x**	**x**			**x**		**x**	
7		**x**			**x**				**x**		**x**

Conclusions:

Interior Point Methods

→ **are well-suited to large scale optimization**

Direct Methods of Linear Algebra

→ **are able to exploit sparsity very well**

Use IPMs in your research!

Implementation of IPMs

Andersen, Gondzio, Mészáros and Xu
Implementation of IPMs for large scale LP,
in: *Interior Point Methods in Mathematical Programming*,
T. Terlaky (ed.), Kluwer Academic Publishers, 1996, pp. 189–252.

Altman and Gondzio
Regularized symmetric indefinite systems in interior point methods for linear and quadratic optimization, *Optimization Methods and Software*, 11-12 (1999), pp 275–302.

Survey on IPMs (easy reading)

Gondzio
Interior point methods 25 years later,
European J. of Operational Research 218 (2012) 587–601.

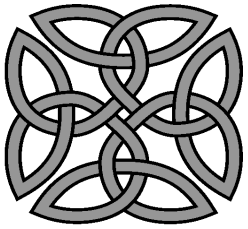

School of Mathematics

Interior Point Methods: Inexact Newton Directions, Conjugate Gradient & Preconditioners

Jacek Gondzio
Email: J.Gondzio@ed.ac.uk
URL: http://www.maths.ed.ac.uk/~gondzio

Lecture 2

Inexact Newton Directions, Conjugate Gradient and Preconditioners

Outline

- **Computational Challenges**
  - ill-conditioned linear systems
  - prohibitively expensive direct method
  - access to A as operator only
- **Inexact Newton Directions**
  - theory: how much of inexactness is allowed?
  - practice: iterative methods
- **Inexact^2 IPM**
- **Conjugate Gradient algorithm**
- **Preconditioners**
  - ideal preconditioner
  - general preconditioners in IPMs for LP/QP
  - splitting preconditioner (back to simplex?)

Complementarity $x_j \cdot s_j = 0 \quad \forall j = 1, 2, ..., n.$

Simplex Method makes a guess of optimal partition:

For *basic* variables, $s_B = 0$ and

$$(x_B)_j \cdot (s_B)_j = 0 \quad \forall j \in \mathcal{B}.$$

For *non-basic* variables, $x_N = 0$ hence

$$(x_N)_j \cdot (s_N)_j = 0 \quad \forall j \in \mathcal{N}.$$

Interior Point Method uses ε-mathematics:

Replace $x_j \cdot s_j = 0 \quad \forall j = 1, 2, ..., n$
by $x_j \cdot s_j = \boldsymbol{\mu} \quad \forall j = 1, 2, ..., n.$

Force convergence $\boldsymbol{\mu} \to 0$.

Optimality Conditions

$$\begin{aligned} Ax &= b \\ A^T y + s &= c \\ XSe &= \mu e \\ x,\, s &\geq 0. \end{aligned}$$

Newton Direction

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} \xi_p \\ \xi_d \\ \xi_\mu \end{bmatrix}.$$

Eliminate Δs to get Augmented System, then
eliminate Δx to get Normal Equations (Schur complement)

Augmented System

$$\begin{bmatrix} \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

Normal Equations

$$(A\Theta A^T)\Delta y = g$$

Ill-conditioned scaling matrix $\Theta = XS^{-1}$.

For **"basic"** variables: $\Theta_j = x_j/s_j \to \infty \quad \Theta_j^{-1} \to 0;$

For **"non-basic"** variables: $\Theta_j = x_j/s_j \to 0 \quad \Theta_j^{-1} \to \infty.$

From LP via QP to NLP

LP
$$\begin{bmatrix} \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

QP
$$\begin{bmatrix} Q + \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

NLP
$$\begin{bmatrix} Q(x,y) + \Theta_P^{-1} & A(x)^T \\ A(x) & -\Theta_D \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

Matrix Θ poses a numerical challenge!

Inefficient Direct Approach

Cholesky factors get sometimes hopelessly dense.
QAP (Quadratic Assignment Problems).

Problem	Dimensions		
	rows	columns	nonzeros
qap12	3192	8856	38304
qap15	6330	22275	94950

Problem	**Normal Equations**			**Augmented System**		
	nz(AAt)	nz(LLt)	Flops	nz(A)	nz(LLt)	Flops
qap12	74592	2135388	2.38e+9	38304	1969957	2.05e+9
qap15	186075	8191638	1.79e+10	94950	7374972	1.52e+10

Constraints accessed only as "Operators"

There exist applications where constraint matrix A is very large, possibly dense and too large to store. However, operations with A such as $y := Ax$ and $z := A^T y$ can be executed very efficiently.

In many signal or image processing applications the constraint matrix A has the form $A = RW$, where

- R is a low-rank randomised sensing matrix
- W is a basis over which the signal has a sparse representation (columns of W form this basis, for example wavelet basis)
- Operations $y := Ax$ and $z := A^T y$ might involve fast transforms, such as Radon, FFT, etc.

Factorization (direct method) is not an option.

Iterative Methods Needed

Normal Equations or **Augmented System**:

- NE is positive definite: can use conjugate gradients;
- AS is indefinite: can use BiCGSTAB, GMRES, QMR;

Augmented System

$$\begin{bmatrix} Q+\Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}$$

Normal Equations

$$(A(Q+\Theta^{-1})^{-1}A^T)\Delta y = g$$

Ill-conditioned scaling matrix $\Theta = XS^{-1}$.

For *"basic"* variables: $\Theta_j = x_j/s_j \to \infty$ $\quad \Theta_j^{-1} \to 0$;

For *"non-basic"* variables: $\Theta_j = x_j/s_j \to 0$ $\quad \Theta_j^{-1} \to \infty$.

Main Tool: Inexact Newton Method

Replace an *exact* Newton direction

$$\nabla^2 f(x)\Delta x = -\nabla f(x)$$

with an *inexact* one:

$$\nabla^2 f(x)\Delta x = -\nabla f(x) + \boldsymbol{r},$$

where the *residual* $\boldsymbol{r}$ is small: $\|\boldsymbol{r}\| \leq \boldsymbol{\eta}\|\nabla f(x)\|,\ \boldsymbol{\eta} \in (0,1)$.

The NLP community usually writes it as:

$$\|\nabla^2 f(x)\Delta x + \nabla f(x)\|_2 \leq \boldsymbol{\eta}\|\nabla f(x)\|_2, \quad \boldsymbol{\eta} \in (0,1).$$

Bellavia,
Inexact Interior Point Method, *JOTA* 96 (1998) 109–121.

Dembo, Eisenstat & Steihaug,
Inexact Newton Methods, *SIAM J. on Numerical Analysis* 19 (1982) 400–408.

Theorem: Suppose the feasible IPM for QP is used.

If the method operates in the *small* neighbourhood

$$\mathcal{N}_2(\theta) := \{(x,y,s) \in \mathcal{F}^0 : \|XSe - \mu e\|_2 \leq \theta\mu\}$$

and uses the *inexact* Newton direction with $\boldsymbol{\eta} = 0.3$, then it converges in at most

$$K = \mathcal{O}(\sqrt{\boldsymbol{n}}\, \ln(1/\epsilon)) \quad \text{iterations.}$$

If the method operates in the *symmetric* neighbourhood

$$\mathcal{N}_S(\gamma) := \{(x,y,s) \in \mathcal{F}^0 : \gamma\mu \leq x_i s_i \leq (1/\gamma)\mu\}$$

and uses the *inexact* Newton direction with $\boldsymbol{\eta} = 0.05$, then it converges in at most

$$K = \mathcal{O}(\boldsymbol{n}\, \ln(1/\epsilon)) \quad \text{iterations.}$$

Gondzio, Convergence Analysis of an Inexact Feasible IPM for Convex Quadratic Programming, *SIAM Journal on Optimization* 23 (2013) No 3, pp. 1510–1527.

Inexact^2 IPM

Standard *Inexact Newton Method*

$$\nabla^2 f(x)\Delta x = -\nabla f(x) + \boldsymbol{r},$$

where the *residual* $\boldsymbol{r}$ is small: $\|\boldsymbol{r}\| \leq \boldsymbol{\eta}\|\nabla f(x)\|$, $\boldsymbol{\eta} \in (0,1)$
is *disappointingly conservative* when applied in IPMs!

Newton direction $(\Delta x, \Delta y, \Delta s)$ comes from the system of linear equations:

$$\begin{bmatrix} A & 0 & 0 \\ 0 & A^T & I \\ S & 0 & X \end{bmatrix} \cdot \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} \xi_P \\ \xi_D \\ \xi_\mu \end{bmatrix} = \begin{bmatrix} b - Ax \\ c - A^T y - s \\ \sigma\mu e - XSe \end{bmatrix}.$$

Full step in Newton direction ($\alpha = 1$) would immediately reach primal feasibility and dual feasibility.
In practice such steps rarely happen.
Why should we waste time on computing accurate directions?

Intriguing observation

What happens at a particular IPM iteration?

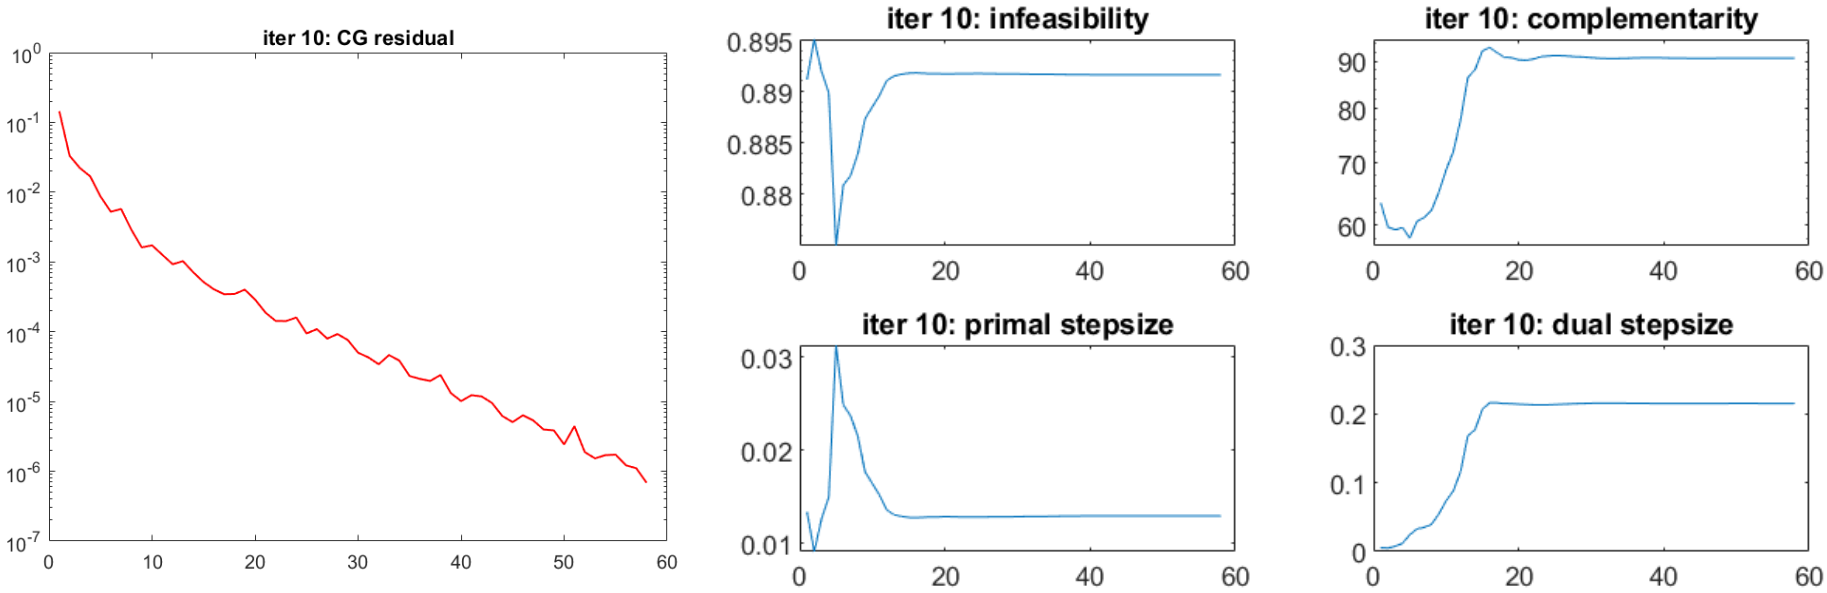

The accuracy required from the inner solver **does not change** the quality of Newton direction!

Stop inner solver as soon as the **stagnation** occurs.

Inexact^2 IPM

Accept the direction produced by the inner solver as soon as

$$\max_j \left| \frac{\Delta x_j^k}{x_j^k} \right| \leq M, \qquad \max_j \left| \frac{\Delta s_j^k}{s_j^k} \right| \leq M$$

and

$$\|\xi_P^{k+1}\| \leq \eta_k \|\xi_P^k\|, \quad \|\xi_D^{k+1}\| \leq \eta_k \|\xi_D^k\|,$$

where $\eta_k \geq 1 - \alpha_k$.

- Implemented with CG and MINRES;
- Prevents IPM from "over-solving" of the linear systems
  $\rightarrow$ 70%-90% reduction of the number of Krylov iterations;
- Worst-case complexity drops from $\mathcal{O}(n)$ to $\mathcal{O}(n^2)$.

F. Zanetti and J. Gondzio,
A new stopping criterion for Krylov solvers applied in interior point methods,
SIAM Journal on Scientific Computing 45 (2023) No 2, pp A703–A728.

Inexact^2 IPM

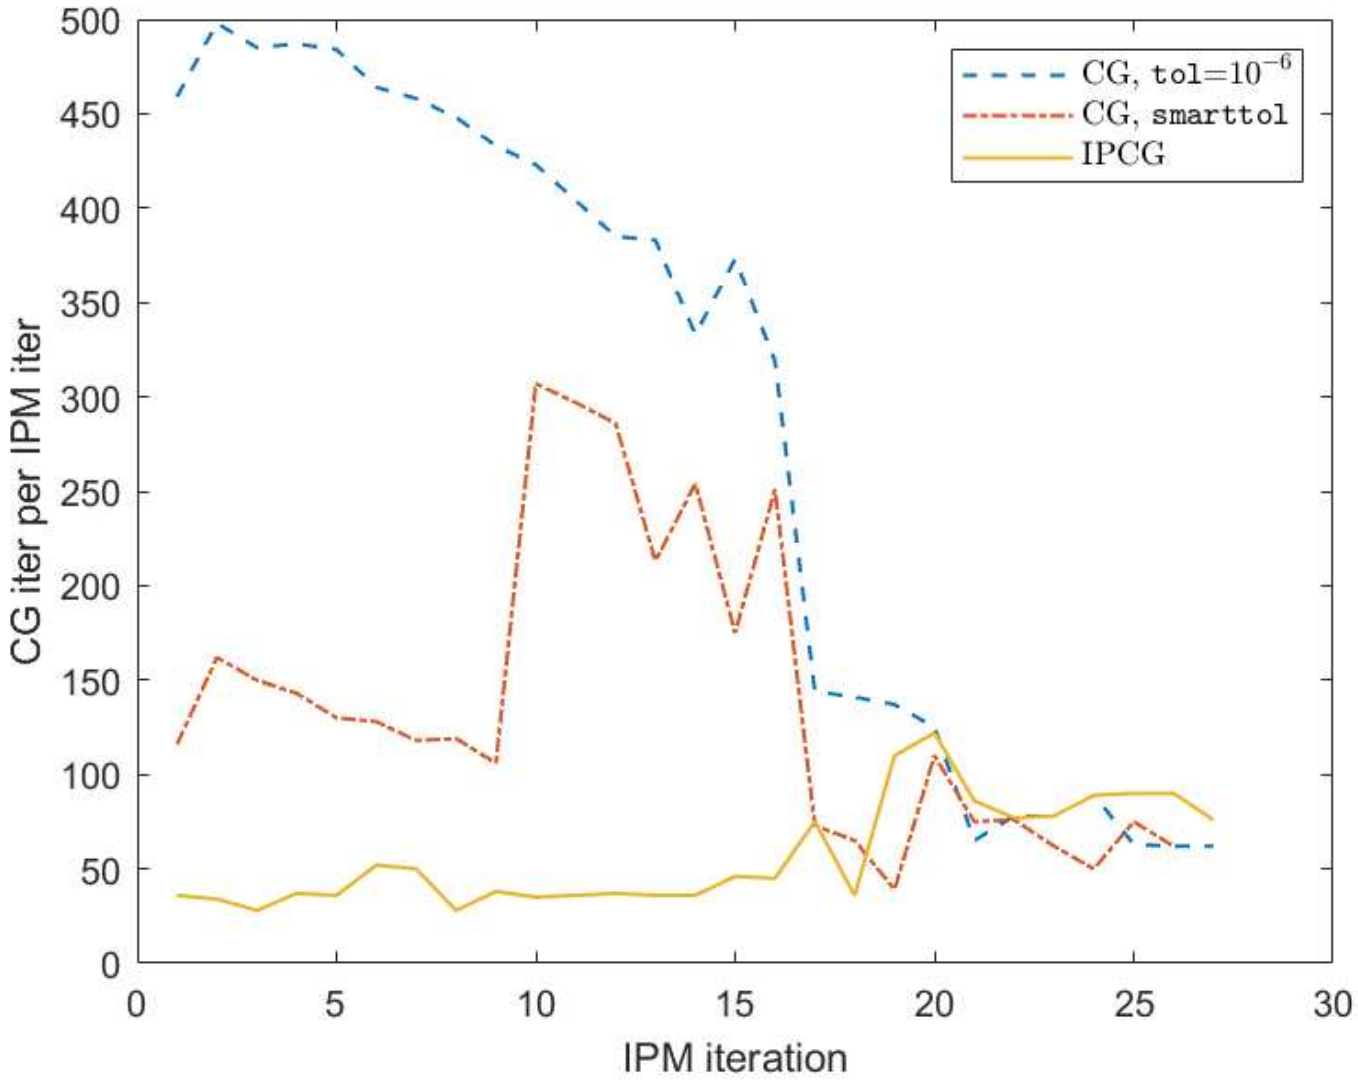

Conjugate Gradient Algorithm

Conjugate Gradient Method

Let $H \in \mathcal{R}^{n \times n}$ be a *symmetric positive definite matrix*.
Consider a linear equation

$$Hx = b.$$

Consider an unconstrained quadratic minimization problem

$$\min_x \quad f(x) := \frac{1}{2} x^T H x - b^T x$$

and observe that

$$\nabla f(x) = Hx - b = 0$$

is its necessary and sufficient optimality condition.

Conjugate Directions

Let $H \in \mathcal{R}^{n \times n}$ be a symmetric positive definite matrix.
Use it to define the scalar product

$$\langle u, v \rangle_H = u^T H v.$$

Scalar product induces orthogonality (conjugacy):
$u \perp_H v$ iff $\langle u, v \rangle_H = 0$ that is $u^T H v = 0$.
We say that u and v are H-orthogonal.

Scalar product induces the norm:

$$\|u\|_H^2 = \langle u, u \rangle_H.$$

Conjugate gradients use H-orthogonality to explore the space.
Two vectors d_i and d_j are H-orthogonal, or *conjugate*, if

$$d_i^T H d_j = 0.$$

Conjugate Directions

Solving equation with a symmetric positive definite matrix

$$Hx = b.$$

Define the error at iteration i: $e_i = x_i - \hat{x}$,
and the residual at iteration i: $r_i = b - Hx_i = -He_i$.

CG Algorithm:

$$\begin{aligned}
d_0 &= r_0 = b - Hx_0 \\
\alpha_i &= \frac{r_i^T r_i}{d_i^T H d_i} \\
x_{i+1} &= x_i + \alpha_i d_i \\
r_{i+1} &= r_i - \alpha_i H d_i \\
\beta_i &= \frac{r_{i+1}^T r_{i+1}}{r_i^T r_i} \\
d_{i+1} &= r_{i+1} + \beta_{i+1} d_i
\end{aligned}$$

Properties

Residuals are orthogonal to each other:

$$r_i^T r_j = 0 \qquad \forall\ 0 \le i < j.$$

Directions are H-orthogonal (*conjugate*) to each other:

$$d_i^T H d_j = 0 \qquad \forall\ 0 \le i < j.$$

CG makes step from x_i to $x_{i+1} = x_i + \alpha d_i$;
it chooses α so as to *minimize f along direction d_i*.
The new error e_{i+1} is H-orthogonal to d_i. Indeed:

$$\begin{aligned} \frac{d}{d\alpha} f(x_{i+1}) &= 0 \\ \nabla f(x_{i+1})^T \frac{d}{d\alpha} x_{i+1} &= 0 \\ -r_{i+1}^T d_i &= 0 \\ e_{i+1}^T H d_i &= 0. \end{aligned}$$

Krylov subspace

Define the *Krylov* subspace $\mathcal{K}_k$:

$$\mathcal{K}_k := span(r_0, Hr_0, H^2 r_0, \ldots, H^{k-1} r_0)$$

for $k \geq 1$, where r_0 is the initial residual: $r_0 = b - Hx_0$.
The k-th iterate x_k of CG minimizes

$$f(x) := \frac{1}{2} x^T H x - b^T x$$

over $x_0 + \mathcal{K}_k$.
Observe that if $f(\tilde{x})$ is the minimal value (in $\mathcal{R}^n$) then

$$\nabla f(\tilde{x}) = H\tilde{x} - b = 0$$

and hence $\tilde{x} = x^*$.

Minimization Property

Lemma: Let $\mathcal{S} \subset \mathcal{R}^n$. If x_k minimizes f over $\mathcal{S}$ then x_k also minimizes $\|x - x^*\|_H = \|r\|_{H^{-1}}$ over $\mathcal{S}$.

Proof: Observe that, since H is symmetric and $Hx^* = b$, we have

$$\|x - x^*\|_H^2 = (x - x^*)^T H(x - x^*) = 2f(x) + (x^*)^T Hx^*.$$

Since $(x^*)^T Hx^*$ is independent of x, minimizing f is equivalent to minimizing $\|x - x^*\|_H^2$ and hence to minimizing $\|x - x^*\|_H$.

Since $e = x - x^*$ we also have

$$\|e\|_H^2 = e^T He = (H(x - x^*))^T H^{-1}(H(x - x^*)) = \|b - Hx\|_{H^{-1}}^2$$

and hence the H-norm of the error is also the H^{-1}-norm of the residual.

Convergence of CG Algorithm

Let κ be the condition number of H

$$\kappa = \frac{\lambda_{max}}{\lambda_{min}},$$

where λ_{max} and λ_{min} are the largest and the smallest eigenvalues of H, respectively. Recall that $\lambda_{max} \geq \lambda_{min} > 0$.
After k iterations of CG, the error satisfies

$$\|e_k\| \leq 2 \left(\frac{\sqrt{\kappa} - 1}{\sqrt{\kappa} + 1} \right)^k \|e_0\|.$$

Usual behaviour of CG algorithm:

Fast convergence if κ is reasonably small, say, $\kappa = 10^2$ or $\kappa = 10^4$.
Slow convergence when κ is large, say, $\kappa = 10^8$ or $\kappa = 10^{12}$.

Preconditioning

Pre- and post-multiply H (to preserve its symmetry) with a matrix E^{-1} such that the condition number of the **preconditioned** matrix

$$\tilde{H} = E^{-1} H E^{-T},$$

is small.

Instead of solving equation

$$Hx = b,$$

solve the preconditioned equation

$$(E^{-1} H E^{-T})(E^T x) = E^{-1} b.$$

If $\kappa(\tilde{H})$ is small, the CG method applied to preconditioned equation will converge fast.

Preconditioner $P = EE^T$. The matrices

$$P^{-1}H = E^{-T}E^{-1}H \quad \text{and} \quad E^{-1}HE^{-T}.$$

are similar.

Ideal Preconditioner

The Preconditioner $P = EE^T$ should:

- be easy to compute
  (significantly less expensive than Cholesky factor of H)
- be easy to invert
- produce good spectral properties of $E^{-1}HE^{-T}$ (that is $P^{-1}H$):
  either have few distinct eigenvalues;
  or have all eigenvalues in a small cluster: $\lambda_{min} \leq \lambda \leq \lambda_{max}$.

Preconditioned CG Algorithm

Apply **CG Algorithm** to the **preconditioned** system:

$$(E^{-1}HE^{-T})(E^Tx) = E^{-1}b.$$

PCG Algorithm:

$$\begin{aligned}
r_0 &= b - Hx_0 \\
d_0 &= P^{-1}r_0 \\
\alpha_i &= \frac{r_i^T P^{-1} r_i}{d_i^T H d_i} \\
x_{i+1} &= x_i + \alpha_i d_i \\
r_{i+1} &= r_i - \alpha_i H d_i \\
\beta_i &= \frac{r_{i+1}^T P^{-1} r_{i+1}}{r_i^T P^{-1} r_i} \\
d_{i+1} &= P^{-1}r_{i+1} + \beta_{i+1} d_i
\end{aligned}$$

where $P = EE^T$.

Preconditioner

Observations:

The matrix H is used only in a matrix-vector multiplication:

$$y = Hd.$$

The preconditioner $P = EE^T$ is used only to compute

$$z = P^{-1}r,$$

that is to solve equation

$$Pz = r.$$

Preconditioner should be:

- easy to compute;
- easy to invert;
- a "good" approximation of H.

Preconditioning

What makes a good preconditioner?

We would like

$$E^{-1}HE^{-T} \approx I,$$

or, more generally,

$$\kappa(E^{-1}HE^{-T}) \approx 1.$$

Straightforward mathematical interpretation:
keep

$$\|E^{-1}HE^{-T} - I\|$$

small (or keep

$$\|H - EE^T\|$$

small).

Different properties of the preconditioned matrix may be desireable.

Good Preconditioner

Minimize the number of distinct eigenvalues

$$E^{-1}HE^{-T} = \begin{bmatrix} \lambda_1 & & & & & \\ & \lambda_1 & & & & \\ & & \lambda_2 & & & \\ & & & \lambda_2 & & \\ & & & & \lambda_2 & \\ & & & & & \lambda_2 \end{bmatrix}.$$

CG will converge in 2 iterations!

Cluster the eigenvalues

$$\lambda_{min} \leq \lambda_i \leq \lambda_{max}, \quad \forall i,$$

and keep $\lambda_{max}/\lambda_{min}$ small; or
make sure eigenvalues stay in a few tight clusters:

$$\forall i \; \exists k \;\; \text{such that} \;\; \lambda^k_{min} \leq \lambda_i \leq \lambda^k_{max},$$

and keep all $\lambda^k_{max}/\lambda^k_{min}$ small.

Incomplete Cholesky Preconditioner

Incomplete Cholesky: $H \approx LL^T$
"drop small entries on the fly"

- dynamic dropping excludes the use of static data structures
  $\longrightarrow$ significant increase of the factorization cost

Structural dropping

- accept nonzeros only at positions of nonzeros in H;
- accept level-one fill-in,
- accept level-k fill-in,
- accept up to k nonzero fill-in per column, etc.

Kershaw, *Journal of Computational Physics* 26(1978) 43-65.

Example Preconditioners in IPMs

LP & QP Problems

$$\begin{aligned} \min \quad & c^T x + \tfrac{1}{2}\, x^T Q\, x \\ \text{s.t.} \quad & Ax = b, \\ & x \geq 0, \end{aligned}$$

where $A \in \mathcal{R}^{m \times n}$ has full row rank

and $Q \in \mathcal{R}^{n \times n}$ is symmetric positive semidefinite.

m **and** n **may be large.**

Assumption: A and Q are *"operators"* $A \cdot u$, $A^T \cdot v$, $Q \cdot u$

Expectation: Low complexity of these operations

Augmented System or Normal Equations

$$\begin{bmatrix} Q+\Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix} \qquad (A(Q+\Theta^{-1})^{-1}A^T)\Delta y = g$$

Advantages of Augmented System

Oliveira, *PhD Thesis*, Rice University, 1997

Oliveira & Sorensen, A New Class of Preconditioners for Large-Scale Linear Systems from Interior Point Methods for Linear Programming, *Linear Algebra and its Applications* 394 (2005) 1-24.

$\rightarrow$ It is better to precondition AS.

O, OS show that **all preconditioners for the NE have an equivalent for the AS while the opposite is not true**.

After all, NE is equivalent to a restricted order of pivoting in AS.

Augmented System

$$\begin{bmatrix} Q & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} f \\ d \end{bmatrix}.$$

- Optimization: *KKT System*
- PDE: *Saddle Point Problem*

Benzi, Golub & Liesen,
"Numerical Solution of Saddle Point Problems",
Acta Numerica 14 (2005) 1-137.

Major difference in tackling the problem:

- Optimization: *Structure of Q and A is not known*
- PDE: *Expolit features of Q and A to design preconditioners*

Indefinite Matrix H
$\rightarrow$ Indefinite Preconditioner P

Rozložník & Simoncini, *SIMAX* 24 (2002) 368-391.

RS consider the preconditioner P which guarantees that all eigenvalues of the preconditioned matrix $P^{-1}H$ are positive and bounded away from zero.

Although $P^{-1}H$ is indefinite

- the CG can be applied to this problem,
- the asymptotic rate of convergence of CG is approximately the same as that obtained for a positive definite matrix with the same eigenvalues as the original system.

Indefinite Block Preconditioner

KKT matrix and its preconditioner

$$H = \begin{bmatrix} Q & A^T \\ A & 0 \end{bmatrix} \qquad \text{and} \qquad P = \begin{bmatrix} G & A^T \\ A & 0 \end{bmatrix},$$

$Q \in \mathcal{R}^{n \times n}$ is positive definite, and $A \in \mathcal{R}^{m \times n}$ has full row rank.

$G \in \mathcal{R}^{n \times n}$ is a positive definite approximation of Q.

Keller, Gould & Wathen, *SIMAX* 21 (2000) 1300-1317.

Theorem. *Assume that A has rank m $(m < n)$.*
Then, $P^{-1}H$ has at least $2m$ unit eigenvalues, and the other eigenvalues are positive and satisfy

$$\lambda_{min}(G^{-1}Q) \;\leq\; \lambda \;\leq\; \lambda_{max}(G^{-1}Q).$$

How to choose G?

Bergamaschi, G. & Zilli, *COAP* 28 (2004) 149-171.
Augmented system in **QP, NLP**

$$H = \begin{bmatrix} \mathbf{Q} + \Theta^{-1} & A^T \\ A & 0 \end{bmatrix}.$$

Drop off-diagonal elements from Q:

Replace $\mathbf{Q} + \Theta^{-1}$ by $D = diag(\mathbf{Q}) + \Theta^{-1}$.

- With diagonal matrix D we have a choice between $\begin{bmatrix} D & A^T \\ A & 0 \end{bmatrix}$ and $AD^{-1}A^T$.
- It is important to keep Θ^{-1} in the preconditioner.
  Recall that Θ is ill-conditioned:
  "basic" j: $\Theta_j = x_j/s_j \to \infty$; *"non-basic"* j: $\Theta_j = x_j/s_j \to 0$.

Motivation: Sparsity issues: irreducible blocks in QP.

Consider the matrices

$$Q = \begin{bmatrix} \mathbf{x} & \mathbf{x} & & & \\ \mathbf{x} & \mathbf{x} & & & \\ & & x & & \\ & & & x & \\ & & & & x \end{bmatrix} \quad \text{and} \quad A = \begin{bmatrix} x & & & x \\ x & & x & \\ & x & x & x \\ & x & & x \end{bmatrix}.$$

$$H = \left[\begin{array}{cccc|cccc} \mathbf{x} & \mathbf{x} & & & x & x & & \\ \mathbf{x} & \mathbf{x} & & & & & x & x \\ & & x & & & x & x & \\ & & & x & & & & x \\ & & & & x & & x & \\ \hline x & & & x & & & & \\ x & & x & & & & & \\ & x & x & & x & & & \\ & x & & x & & & & \end{array}\right] \rightarrow H^{(2)} = \left[\begin{array}{cccc|cccc} \mathbf{x} & \mathbf{x} & & & x & x & \mathbf{f} & \mathbf{f} \\ \mathbf{x} & \mathbf{x} & & & \mathbf{f} & \mathbf{f} & x & x \\ & & x & & & x & x & \\ & & & x & & & & x \\ & & & & x & & x & \\ \hline x & \mathbf{f} & & x & \mathbf{f} & \mathbf{f} & \mathbf{f} & \mathbf{f} \\ x & \mathbf{f} & x & & \mathbf{f} & \mathbf{f} & \mathbf{f} & \mathbf{f} \\ \mathbf{f} & x & x & x & \mathbf{f} & \mathbf{f} & \mathbf{f} & \mathbf{f} \\ \mathbf{f} & x & & x & \mathbf{f} & \mathbf{f} & \mathbf{f} & \mathbf{f} \end{array}\right].$$

Spectral Analysis:

Eigenvalues of $P^{-1}H$ satisfy:

$$\begin{array}{lll} Qx & +A^T y = \lambda Dx & +\lambda A^T y \\ Ax & = \lambda Ax. & \end{array}$$

If $\lambda = 1$, we are done. If $\lambda \neq 1$ the second equation yields $Ax = 0$. After multiplying the first equation with x^T, we get:

$$x^T Qx = \lambda x^T Dx \quad \Rightarrow \quad \lambda = \frac{x^T Qx}{x^T Dx} = q(D^{-1}Q).$$

The Rayleigh quotient of the generalized eigenproblem: $Dv = \mu Qv$. Since both D and Q are positive definite we have

$$0 < \lambda_{\min}(D^{-1}Q) \leq \lambda \leq \lambda_{\max}(D^{-1}Q).$$

Conclusion: The preconditioner satisfies the requirements of **Rozloznik & Simoncini**.

Splitting Preconditioner (only for LPs)

- Consider the Augmented System $\begin{bmatrix} \Theta^{-1} & A^T \\ A & 0 \end{bmatrix}$.
- Keep Θ^{-1} in the preconditioner.
  Recall that Θ is ill-conditioned:
  "basic" j: $\Theta_j = x_j/s_j \to \infty$;
  "non-basic" j: $\Theta_j = x_j/s_j \to 0$.
- Based on the magnitude of Θ^{-1},
  guess "basic/non-basic" partition:
  $A = [B \; N]$.
- Use the fact that $\Theta_B \gg \Theta_N$ (and $\Theta_B^{-1} \ll \Theta_N^{-1}$).
- The inverse of the preconditioner only needs B^{-1}.

LP Case

For *"basic"* variables, $x_B \to \hat{x}_B > 0$ and $s_B \to \hat{s}_B = 0$ hence

$$\Theta_j = x_j/s_j \approx (x_j^2)/(x_j s_j) = \mathcal{O}(\mu^{-1}) \quad \forall j \in \mathcal{B}.$$

For *"non-basic"* variables, $x_N \to \hat{x}_N = 0$ and $s_N \to \hat{s}_N > 0$ hence

$$\Theta_j = x_j/s_j \approx (x_j s_j)/(s_j^2) = \mathcal{O}(\mu) \quad \forall j \in \mathcal{N}.$$

Convert a difficulty into an advantage

$\to$ Exploit the property:

As $\mu \to 0$ then:

$$\Theta_B \to \infty, \qquad \Theta_B^{-1} \to 0;$$
$$\Theta_N \to 0, \qquad \Theta_N^{-1} \to \infty.$$

Oliveira & Sorensen, *LAA* 394 (2005) 1-24.

Guess basic/nonbasic partition $A = [B|N]$ with **nonsingular** B.

$$
\begin{aligned}
& E^{-1} H E^{-T} \\
&= \left[\begin{array}{cc|c} \Theta_B^{1/2} & & \Theta_B^{-1/2} B^{-1} \\ & \Theta_N^{1/2} & \\ \hline \Theta_B^{1/2} & & \end{array}\right] \left[\begin{array}{cc|c} -\Theta_B^{-1} & & B^T \\ & -\Theta_N^{-1} & N^T \\ \hline B & N & 0 \end{array}\right] \left[\begin{array}{cc|c} \Theta_B^{1/2} & & \Theta_B^{1/2} \\ & \Theta_N^{1/2} & \\ \hline B^{-T}\Theta_B^{-1/2} & & \end{array}\right] \\
&= \left[\begin{array}{cc|c} I_m & W^T & \\ W & -I_{n-m} & \\ \hline & & -I_m \end{array}\right], \quad \text{where} \quad W = \Theta_N^{1/2} N^T B^{-T} \Theta_B^{-1/2}.
\end{aligned}
$$

With $\mu \to 0$ we have $\Theta_B^{-1} \to 0,\ \Theta_N \to 0$ hence

$$W = \Theta_N^{1/2} N^T B^{-T} \Theta_B^{-1/2} \approx 0.$$

Observation:

- The inverse of the preconditioner needs B^{-1}.
- Guessing a stable "basic/non-basic" partition $A = [B\ N]$ is a challenge!
- Use **maximum volume** concept to identify stable B.
  Def: $Vol(B) = |det(B)|$.
  Maximum volume of B implies that entries of $B^{-1}N$ cannot exceed 1, hence $W^T = \Theta_B^{-1/2} B^{-1} N \Theta_N^{1/2}$ is small.

Finding the maximum volume basis in NP-hard.
But finding a (relaxed) ρ-maximum volume basis can be done very efficiently. Its implementation is a part of an open source IPM code in HiGHS: `https://github.com/ERGO-Code/HiGHS`

L. Schork and J. Gondzio, Implementation of an interior point method with basis preconditioning, *Mathematical Programming Computation* 12 (2020) pp. 603–635.

Conclusions:

Direct Methods $\rightarrow$ sometimes suffer from prohibitive fill-in.

Iterative Methods $\rightarrow$ strongly depend upon preconditioners.

Challenge:

A development of new preconditioners for IPMs.

School of Mathematics

Interior Point Methods: Inexact Newton Directions, Conjugate Gradient & Preconditioners

Jacek Gondzio
Email: J.Gondzio@ed.ac.uk
URL: http://www.maths.ed.ac.uk/~gondzio

Lecture 3

Example Applications:

Design of Efficient Preconditioners

Outline

- Homotopy:

  "replace a difficult problem with a sequence of (easy to solve) auxilliary problems"

  - Interior Point Methods
  - Primal-Dual Newton Conjugate Gradient
- Example Applications and Suitable Preconditioners
  - Compressed Sensing
  - Compressed Sensing (Coherent and Redundant Dict.)
  - X-ray Tomography Material Separation
- Conclusions

Introduction

New applications:

- signal/image processing
  (compression, reconstruction, deblurring)
- modern statistics
  (inverse problems, classification, machine learning)

Overarching mathematical problem:

- dimension reduction

Optimization faces new challenges:

huge scale of data available
needs new processing algorithms

Regularized Optimization Problems

Consider a simple (unconstrained) optimization problem

$$\min_x \quad \tau\psi(x) + \phi(x)$$

where τ is a parameter,
ϕ is a convex function to be minimized, and
ψ is a (convex) regularization function.

Example regularizations:

- $\psi(x) = \|x\|_2$ (or $\|x\|_2^2$)
  promotes "equal spread of mass in vector x"
- $\psi(x) = \|x\|_1$ (it is a good proxy for $\|x\|_0$)
  promotes "few nonzero entries in vector x"

Binary Classification

$$\min \tau\|x\|_1 + \sum_{i=1}^{m} \log(1+e^{-b_i x^T a_i}) \qquad \min \tau\|x\|_2^2 + \sum_{i=1}^{m} \log(1+e^{-b_i x^T a_i})$$

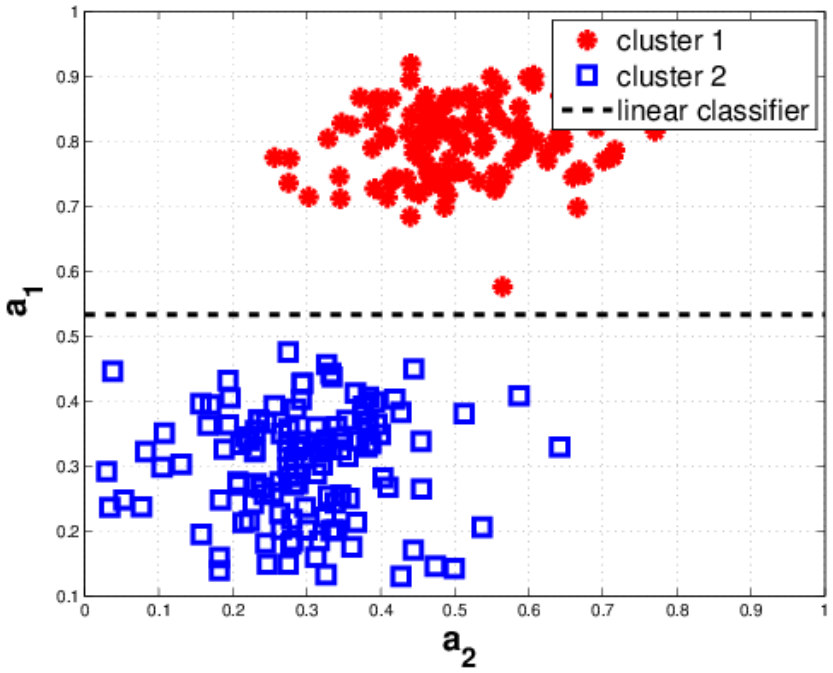

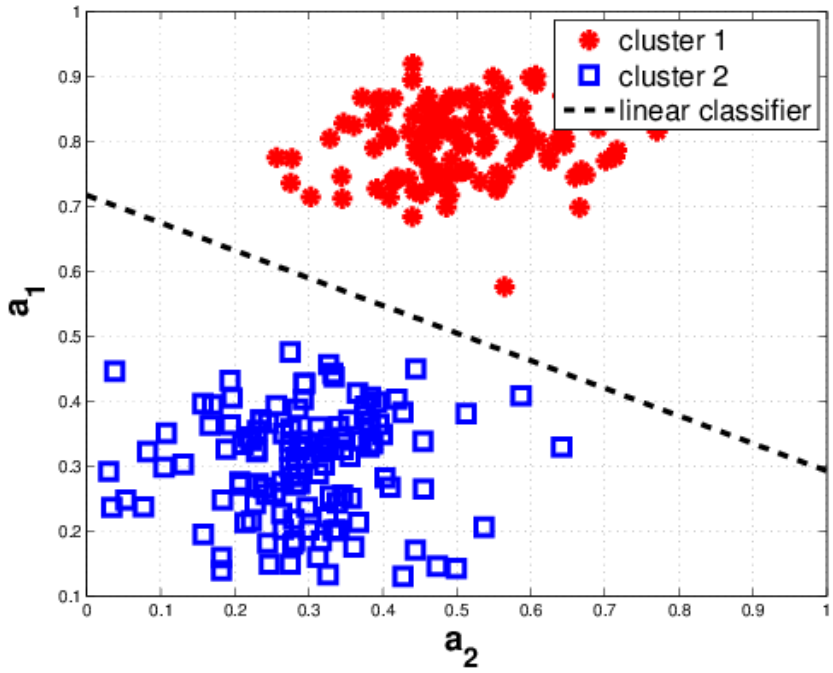

Bayesian Statistics Viewpoint

Estimate x from observations

$$y = Ax + e,$$

where y are observations and e is the Gaussian noise.

$\rightarrow \min_x \|y - Ax\|_2^2$

If the prior on x is Laplacian ($\log p(x) = -\lambda\|x\|_1 + K$) then

$$\min_x \tau\|x\|_1 + \|Ax - b\|_2^2$$

Tibshirani,
J. of Royal Statistical Soc B 58 (1996) 267-288.

ℓ_1-regularization

$$\min_x \quad \tau\|x\|_1 + \phi(x).$$

Unconstrained optimization $\Rightarrow$ easy
Serious Issue: nondifferentiability of $\|.\|_1$

Two possible tricks:

- Splitting $x = u - v$ with $u, v \geq 0$
  $\longrightarrow$ use **IPM**

  replace $z \geq 0$ with $-\mu \, log z$ and drive μ to zero.

- Smoothing with pseudo-Huber approximation
  $\longrightarrow$ use **continuation (Newton CG)**

  replace $|x_i|$ with $\mu(\sqrt{1+\frac{x_i^2}{\mu^2}} - 1)$ and drive μ to zero.

Inexact Primal-Dual IPM:

G., Convergence Analysis of an Inexact Feasible IPM for Convex QP, *SIAM J Opt*, 23 (2013) No 3, 1510–1527.

G., Matrix-Free Interior Point Method, *Computational Opt and Appls*, 51 (2012) 457–480.

Primal-Dual Newton Conjugate Gradient Method:

Fountoulakis and G., A Second-order Method for Strongly Convex ℓ_1-regularization Problems, *Mathematical Programming*, 156 (2016) 189–219.

Dassios, Fountoulakis and G., A Preconditioner for a Primal-Dual Newton Conjugate Gradient Method for Compressed Sensing Problems, *SIAM J on Sci Comput*, 37 (2015) A2783–A2812.

Continuation & Preconditioners

(three examples)

Three examples

- Compressed Sensing
  with **K. Fountoulakis and P. Zhlobich**

$$\min_x \ \tau\|x\|_1 + \frac{1}{2}\|Ax-b\|_2^2, \quad A \in \mathcal{R}^{m\times n}$$

- Compressed Sensing (Coherent and Redundant Dict.)
  with **I. Dassios and K. Fountoulakis**

$$\min_x \ \tau\|W^*x\|_1 + \frac{1}{2}\|Ax-b\|_2^2, \quad W \in \mathcal{C}^{n\times l},\ A \in \mathcal{R}^{m\times n}$$

  think of Total Variation

- X-ray Tomography Material Separation
  with **S. Latva-Äijö, S. Siltanen, M. Lassas, F. Zanetti**

$$\min_{x\geq 0} \ \|h-\mathcal{A}x\|_2^2 + \alpha\|x\|_2^2 + \beta x^T S x,$$

Example 1: Compressed Sensing

with **K. Fountoulakis and P. Zhlobich**

Large dense quadratic optimization problem:

$$\min_x \ \tau\|x\|_1 + \frac{1}{2}\|Ax - b\|_2^2,$$

where $A \in \mathcal{R}^{m\times n}$ is a **very special matrix**.

Fountoulakis, G., Zhlobich
Matrix-free IPM for Compressed Sensing Problems,
Mathematical Programming Computation 6 (2014), pp. 1–31.

Software available at `http://www.maths.ed.ac.uk/ERGO/`

Restricted Isometry Property (RIP)

- *rows* of A are orthogonal to each other (A is built of a subset of rows of an othonormal matrix $U \in \mathcal{R}^{n \times n}$)

$$AA^T = I_m.$$

- small subsets of *columns* of A are nearly-orthogonal to each other: *Restricted Isometry Property (RIP)*

$$\|\bar{A}^T \bar{A} - \frac{m}{n} I_k\| \leq \delta_k \in (0, 1).$$

Candès, Romberg & Tao,
Comm on Pure and Appl Maths 59 (2005) 1207-1233.

Restricted Isometry Property

Matrix $\bar{A} \in \mathcal{R}^{m \times k}$ $(k \ll n)$ is built of a subset of columns of $A \in \mathcal{R}^{m \times n}$.

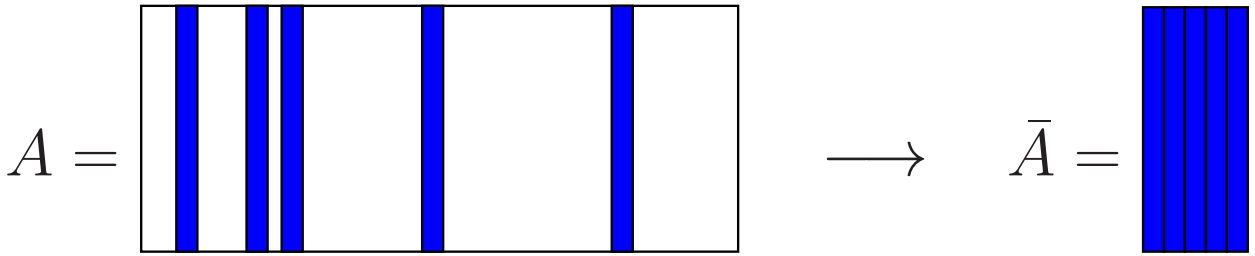

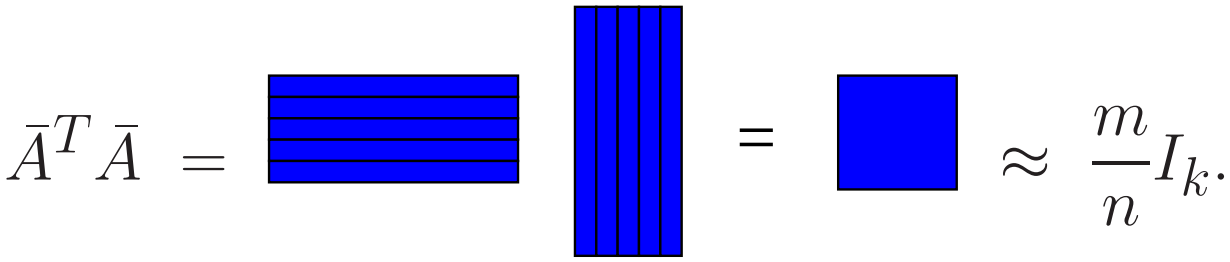

This yields a very well conditioned optimization problem.

Problem Reformulation

$$\min_x \ \tau\|x\|_1 + \frac{1}{2}\|Ax - b\|_2^2$$

Replace $x = x^+ - x^-$ to be able to use $|x| = x^+ + x^-$.
Use $|x_i| = z_i + z_{i+n}$ to replace $\|x\|_1$ with $\|x\|_1 = 1_{2n}^T z$.
(Increases problem dimension from n to $2n$.)

$$\min_{z \geq 0} \ c^T z + \frac{1}{2} z^T Q z,$$

where

$$Q = \begin{bmatrix} A^T \\ -A^T \end{bmatrix} [\, A \ -A \,] = \begin{bmatrix} A^T A & -A^T A \\ -A^T A & A^T A \end{bmatrix} \in \mathcal{R}^{2n \times 2n}$$

Preconditioner

Approximate

$$\mathcal{M} = \begin{bmatrix} A^TA & -A^TA \\ -A^TA & A^TA \end{bmatrix} + \begin{bmatrix} \Theta_1^{-1} & \\ & \Theta_2^{-1} \end{bmatrix}$$

with

$$\mathcal{P} = \frac{m}{n} \begin{bmatrix} I_n & -I_n \\ -I_n & I_n \end{bmatrix} + \begin{bmatrix} \Theta_1^{-1} & \\ & \Theta_2^{-1} \end{bmatrix}.$$

We expect (*optimal partition*):

- k entries of $\Theta^{-1} \to 0, \quad k \ll 2n,$
- $2n - k$ entries of $\Theta^{-1} \to \infty.$

Spectral Properties of $\mathcal{P}^{-1}\mathcal{M}$

Theorem

- Exactly n eigenvalues of $\mathcal{P}^{-1}\mathcal{M}$ are 1.
- The remaining n eigenvalues satisfy

$$|\lambda(\mathcal{P}^{-1}\mathcal{M}) - 1| \leq \delta_k + \frac{n}{m\delta_k L},$$

where δ_k is the RIP-constant, and
L is a threshold of "large" $(\Theta_1 + \Theta_2)^{-1}$.

Fountoulakis, G., Zhlobich
Matrix-free IPM for Compressed Sensing Problems,
Mathematical Programming Computation 6 (2014), pp. 1–31.

Preconditioning

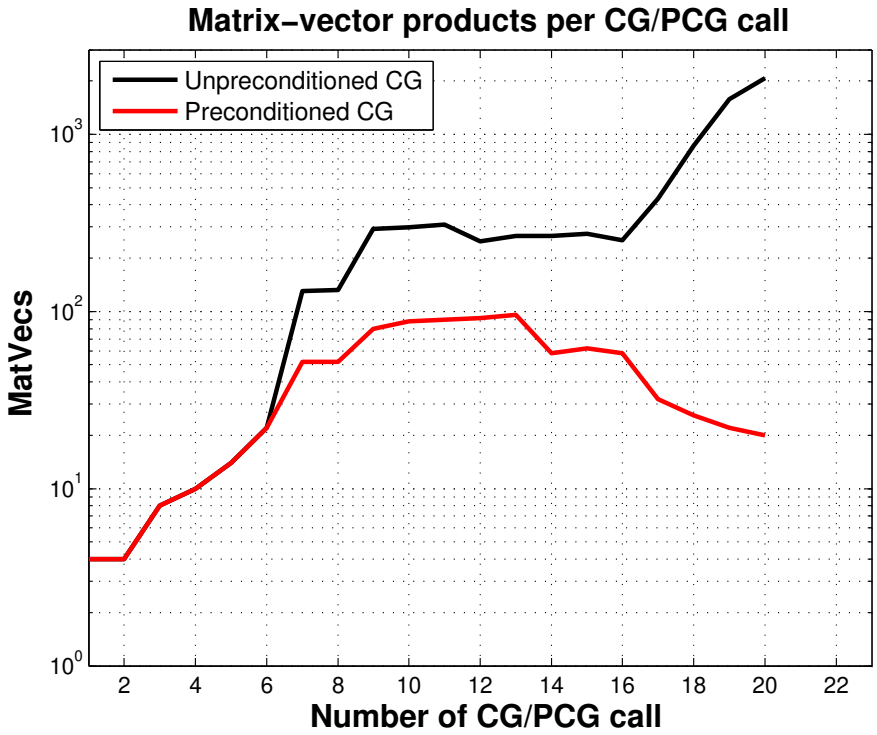

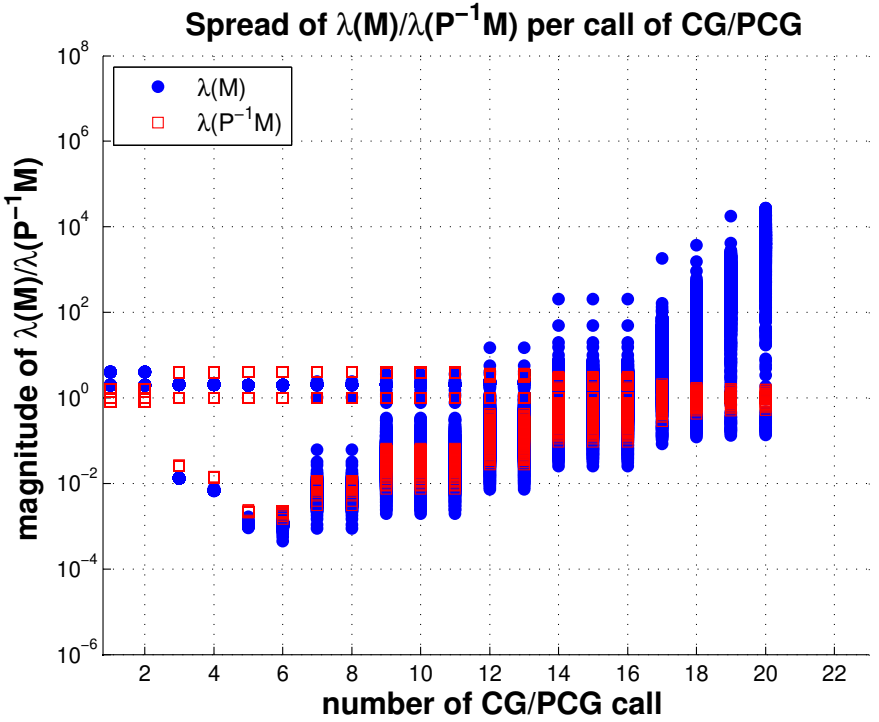

$\longrightarrow$ good clustering of eigenvalues

`mf-IPM` compares favourably with `NestA` on easy probs
(`NestA`: Becker, Bobin and Candés).

SPARCO problems

Comparison on 18 out of 26 classes of problems
(all but 6 complex and 2 installation-dependent ones).

Solvers compared:

PDCO, *Saunders and Kim*, Stanford,
$\ell_1 - \ell_s$, *Kim, Koh, Lustig, Boyd, Gorinevsky*, Stanford,
FPC-AS-CG, *Wen, Yin, Goldfarb, Zhang*, Rice,
SPGL1, *Van Den Berg, Friedlander*, Vancouver, and
mf-IPM, *Fountoulakis, G., Zhlobich*, Edinburgh.

On 36 runs (noisy and noiseless problems), **mf-IPM**:

- is the fastest on 11,
- is the second best on 14, and
- overall is very robust.

Linear Algebra Perspective

Convert a difficulty into an advantage

Interestingly the same trick works:
in IPMs and in Newton Conjugate Gradient!

Linear Algebra of IPMs for LP/QP

Newton direction

$$\begin{bmatrix} A & 0 & 0 \\ -Q & A^T & I \\ S & 0 & X \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \\ \Delta s \end{bmatrix} = \begin{bmatrix} \xi_p \\ \xi_d \\ \xi_\mu \end{bmatrix}.$$

Eliminate Δs from the second equation and get

$$\begin{bmatrix} -Q - \Theta^{-1} & A^T \\ A & 0 \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta y \end{bmatrix} = \begin{bmatrix} \xi_d - X^{-1}\xi_\mu \\ \xi_p \end{bmatrix},$$

where $\boldsymbol{\Theta = XS^{-1}}$ is a diagonal scaling matrix.

Eliminate Δx from the first equation and get

$$(A(Q + \Theta^{-1})^{-1}A^T)\Delta y \;=\; g.$$

IPM Linear Algebra: Splitting Preconditioner

For *"basic"* variables, $x_B \to \hat{x}_B > 0$ and $s_B \to \hat{s}_B = 0$ hence

$$\Theta_j = x_j/s_j \approx (x_j^2)/(x_j s_j) = \mathcal{O}(\mu^{-1}) \quad \forall j \in \mathcal{B}.$$

For *"non-basic"* variables, $x_N \to \hat{x}_N = 0$ and $s_N \to \hat{s}_N > 0$ hence

$$\Theta_j = x_j/s_j \approx (x_j s_j)/(s_j^2) = \mathcal{O}(\mu) \quad \forall j \in \mathcal{N}.$$

Convert a difficulty into an advantage

$\to$ Exploit the property:

As $\mu \to 0$ then:

$$\Theta_B \to \infty, \quad \Theta_B^{-1} \to 0;$$
$$\Theta_N \to 0, \quad \Theta_N^{-1} \to \infty.$$

Oliveira & Sorensen, *LAA* 394 (2005) 1-24.

Guess basic/nonbasic partition $A = [B|N]$, *invertible* B.

$$E^{-1}HE^{-T}$$

$$= \left[\begin{array}{cc|c} \Theta_B^{1/2} & & \Theta_B^{-1/2}B^{-1} \\ & \Theta_N^{1/2} & \\ \hline & & \end{array}\right] \left[\begin{array}{cc|c} -\Theta_B^{-1} & & B^T \\ & -\Theta_N^{-1} & N^T \\ \hline B & N & 0 \end{array}\right] \left[\begin{array}{cc|c} & \Theta_N^{1/2} & \Theta_B^{1/2} \\ & & \\ \hline B^{-T}\Theta_B^{-1/2} & & \end{array}\right]$$

$$= \left[\begin{array}{cc|c} I_m & W^T & \\ W & -I_{n-m} & \\ \hline & & -I_m \end{array}\right], \quad \text{where} \quad W = \Theta_N^{1/2} N^T B^{-T} \Theta_B^{-1/2}.$$

With $\mu \to 0$ we have $\Theta_B^{-1} \to 0$, $\Theta_N \to 0$ hence

$$W = \Theta_N^{1/2} N^T B^{-T} \Theta_B^{-1/2} \approx 0.$$

Property of Sparse Approximations

$$\min_x \; f(x) = \tau \|x\|_1 + \|Ax - b\|_2^2$$

Assume a sparse solution exists $\hat{x} = [\hat{x}_B \,|\, \hat{x}_Z]$ with $\hat{x}_Z = 0$. Partition matrix $A = [A_B \,|\, A_Z]$ accordingly.

Then only a small subset of the Hessian $A^T A$ is "relevant"

$$A^T A = \begin{bmatrix} A_B^T A_B & A_B^T A_Z \\ A_Z^T A_B & A_Z^T A_Z \end{bmatrix}$$

Splitting $x = u - v$ and IPMs

There is a need to solve equations with $A^T A + \Theta^{-1}$ and

- $\Theta_j^{-1} \to 0$, for j in the sparse part B ($\hat{x}_j > 0$),
- $\Theta_j^{-1} \to \infty$, for j in the zero part Z ($\hat{x}_j \approx 0$).

Then

$$A^T A + \Theta^{-1} = \begin{bmatrix} A_B^T A_B + \Theta_B^{-1} & A_B^T A_Z \\ A_Z^T A_B & A_Z^T A_Z + \Theta_Z^{-1} \end{bmatrix}$$

$$\approx \begin{bmatrix} A_B^T A_B & \\ & \Theta_Z^{-1} \end{bmatrix}$$

Smoothing with pseudo-Huber approximation

There is a need to solve equations with $A^TA + \nabla^2\psi_\mu(x)$ and

- $\psi_\mu(x_j) \gg 0$ and $\nabla^2\psi_\mu(x_j) \to 0$, for j in part B,
- $\psi_\mu(x_j) \approx 0$ and $\nabla^2\psi_\mu(x_j) \to \frac{1}{\mu}$, for j in part Z.

Then

$$A^TA + \nabla^2\psi_\mu(x) = \begin{bmatrix} A_B^TA_B + \nabla^2\psi_\mu(x_B) & A_B^TA_Z \\ A_Z^TA_B & A_Z^TA_Z + \nabla^2\psi_\mu(x_Z) \end{bmatrix}$$

$$\approx \begin{bmatrix} A_B^TA_B & \\ & \frac{1}{\mu}I \end{bmatrix}$$

Example 2: CS, Coherent & Redundant Dict.

with **I. Dassios and K. Fountoulakis**.

Large dense quadratic optimization problem:

$$\min_x \ \tau\|W^*x\|_1 + \frac{1}{2}\|Ax-b\|_2^2,$$

where $A \in \mathcal{R}^{m\times n}$ and $W \in \mathcal{C}^{n\times l}$ is a *dictionary*.

Dassios, Fountoulakis and G.
A Preconditioner for a Primal-Dual Newton Conjugate Gradient Method for Compressed Sensing Problems,
SIAM J on Sci. Comput. 37 (2015) A2783–A2812.

Software available at `http://www.maths.ed.ac.uk/ERGO/`

Compressed Sensing and Continuation

Replace

$$\min_x \quad f(x) = \tau \|W^* x\|_1 + \frac{1}{2}\|Ax - b\|_2^2, \qquad \longrightarrow \boldsymbol{x_\tau}$$

with

$$\min_x \quad f_\mu(x) = \tau \psi_\mu(W^* x) + \frac{1}{2}\|Ax - b\|_2^2, \qquad \longrightarrow \boldsymbol{x_{\tau,\mu}}$$

Solve approximately a family of problems for a (short) decreasing sequence of μ's: $\boldsymbol{\mu_0 > \mu_1 > \mu_2 \cdots}$

Theorem (Brief description)

There exists a $\tilde{\mu}$ such that $\forall \mu \leq \tilde{\mu}$ the difference of the two solutions satisfies

$$\|x_{\tau,\mu} - x_\tau\|_2 = \mathcal{O}(\mu^{1/2}) \quad \forall\, \tau, \mu.$$

A better linearization

$$\tau \underbrace{Dx}_{\nabla\psi_\mu(x)} + A^T(Ax-b) = 0,$$

where $D{:=}diag(D_1,...,D_n)$ with $D_i{:=}(\mu^2{+}x_i^2)^{-\frac{1}{2}} \quad \forall i{=}1,...,n$

Set $g = Dx$. Use the easier form of the equations.

Difficult:

$$\tau g + A^T(Ax-b) = 0,$$
$$g = Dx.$$

Easy:

$$\tau g + A^T(Ax-b) = 0,$$
$$D^{-1}g = x.$$

Chan, Golub, Mulet,
SIAM J. on Sci. Comput. 20 (1999) 1964–1977.

A better linearization

Example: $g_i = 0.99$

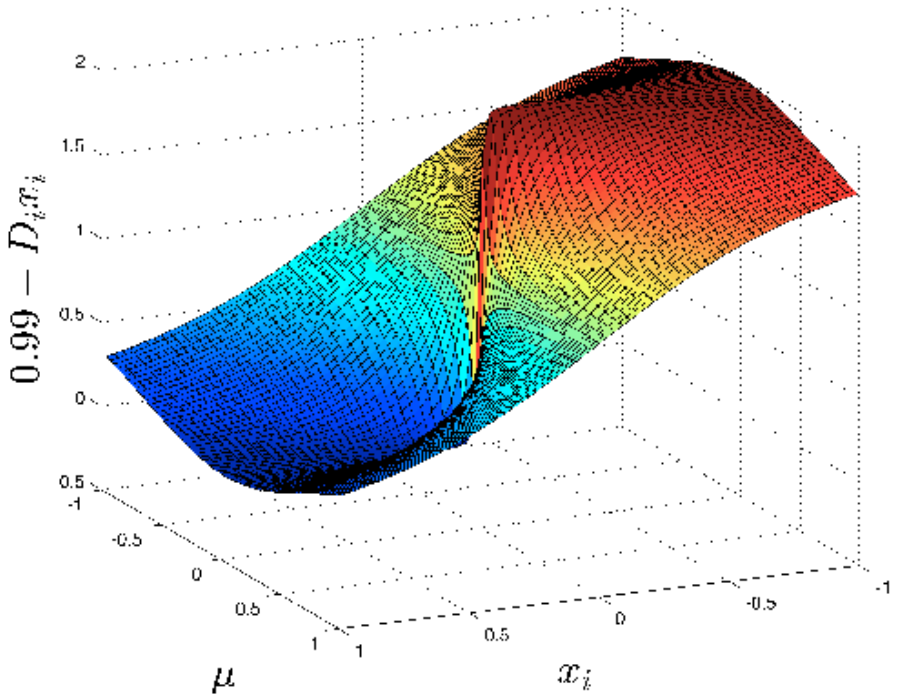

bad: $g_i = D_i x_i$

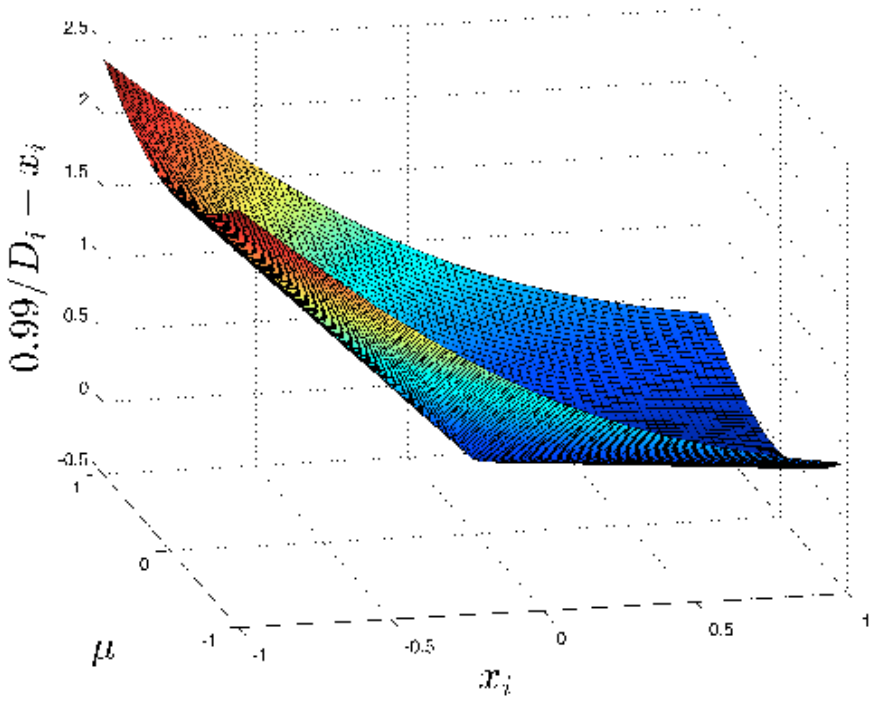

good: $D_i^{-1} g_i = x_i$

W-Restricted Isometry Property (W-RIP)

- *rows* of A are nearly-orthogonal to each other, i.e., there exists a small constant δ such that

$$\|AA^T - I_m\| \leq \delta.$$

- *W-Restricted Isometry Property (W-RIP)*:
there exists a constant δ_q such that

$$(1-\delta_q)\|Wz\|_2^2 \leq \|AWz\|_2^2 \leq (1+\delta_q)\|Wz\|_2^2$$

for all at most q-sparse $z \in \mathcal{C}^n$.

Candès, Eldar & Needell,
Appl and Comp Harmonic Anal 31 (2011) 59-73.

Preconditioner

Approximate

$$\mathcal{H} = \tau \nabla^2 \psi_\mu(W^* x) + A^T A$$

with

$$\mathcal{P} = \tau \nabla^2 \psi_\mu(W^* x) + \rho I_n.$$

We expect (*optimal partition*):

- k entries of $W^* x \gg 0, \quad k \ll l,$
- $l - k$ entries of $W^* x \approx 0.$

The preconditioner approximates well the 2nd derivative of the pseudo-Huber regularization.

Spectral Properties of $\mathcal{P}^{-1}\mathcal{H}$

Theorem

- The eigenvalues of $\mathcal{P}^{-1}\mathcal{H}$ satisfy

$$|\lambda(\mathcal{P}^{-1}\mathcal{H}) - 1| \leq \frac{\eta(\delta, \delta_q, \rho)}{\rho},$$

where δ_q is the W-RIP constant,
δ is another small constant, and
$\eta(\delta, \delta_q, \rho)$ is some simple function.

Dassios, Fountoulakis and G.
A Preconditioner for a Primal-Dual Newton Conjugate Gradient Method for Compressed Sensing Problems,
SIAM J on Sci. Comput. 37 (2015) A2783–A2812.

CS: Coherent and Redundant Dictionaries

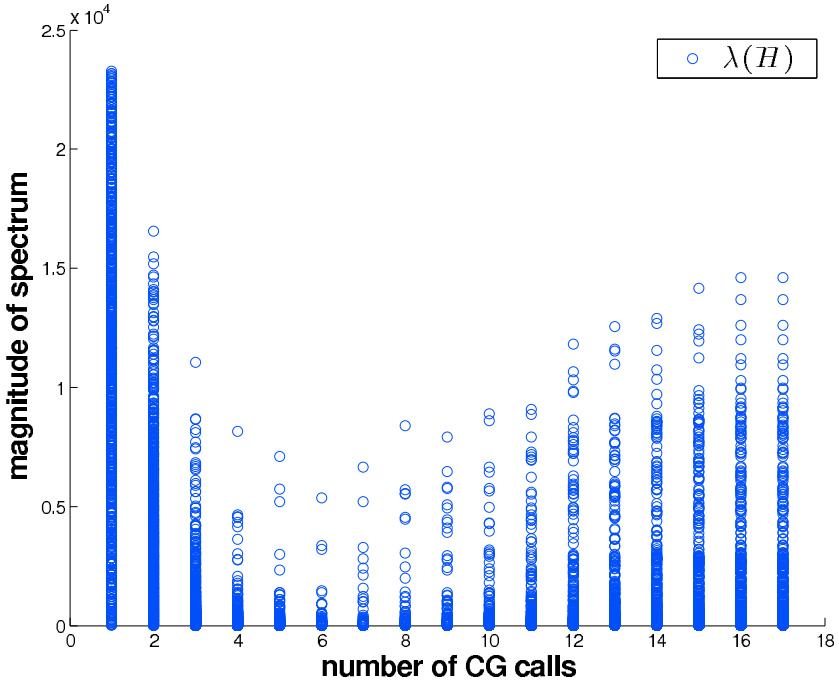

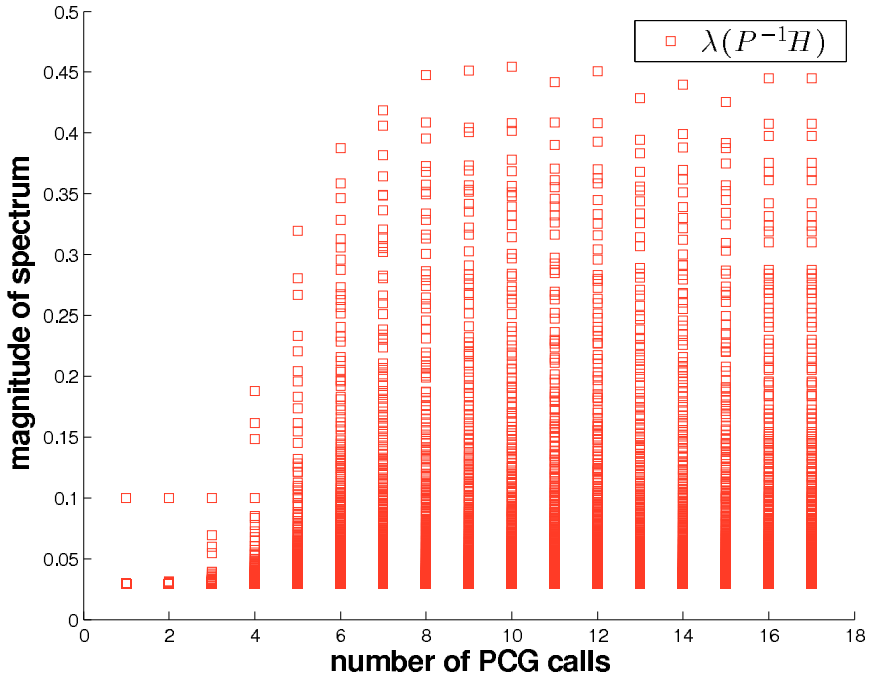

$\longrightarrow$ good clustering of eigenvalues

pdNCG outperforms TFOCS on several examples
(TFOCS: Becker, Candés and Grant).

Example 3: Multi-energy X-ray Tomography

with **S. Latva-Äijö, S. Siltanen, M. Lassas, F. Zanetti**

Inverse problem:

$$\min_{x \geq 0} \ \|h - \mathcal{A}x\|_2^2 + \alpha \|x\|_2^2 + \beta x^T S x,$$

where $S = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}$ is an *inner product regularizer*.

S promotes material separation. Indeed, minimizing

$$[x_1, x_2]^T \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix} \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} = 2 x_1^T x_2$$

with both $x_1 \geq 0$ and $x_2 \geq 0$ forces at least one of components (either $(x_1)_j$ or $(x_2)_j$) to be zero.

J. Gondzio, S.-M. Latva-Äijö, S.M Siltanen, M. Lassas, F. Zanetti,
Material-separating regularizer for multi-energy X-ray tomography, *Inverse Problems*, 38 (2022) 025013.

Multi-energy X-ray Tomography (cont'd)

This is a QP:

$$\min_{x \geq 0} \ \frac{1}{2} x^T \mathcal{Q} x + d^T x,$$

where

$$\mathcal{Q} = \begin{bmatrix} c_{11}^2 + c_{21}^2 & c_{11}c_{12} + c_{21}c_{22} \\ c_{11}c_{12} + c_{21}c_{22} & c_{12}^2 + c_{22}^2 \end{bmatrix} \otimes R^T R + \begin{bmatrix} \rho & \eta \\ \eta & \rho \end{bmatrix} \otimes I.$$

Here c_{11}, c_{12}, c_{21} and c_{22} are the attenuation constants.
R describes information about the geometry of the measurements.
It can only be accessed via *matrix-vector products* performed using the Radon transform.
Preconditioner:

$$\mathcal{P} = \begin{bmatrix} (c_{11}^2 + c_{21}^2)\nu I + \rho I & (c_{11}c_{12} + c_{21}c_{22})\nu I + \eta I \\ (c_{11}c_{12} + c_{21}c_{22})\nu I + \eta I & (c_{12}^2 + c_{22}^2)\nu I + \rho I \end{bmatrix} + X^{-1}S,$$

where ν approximates the main diagonal of the blocks in $R^T R$.

Toeplitz-like structure of $\mathcal{A}^T\mathcal{A}$

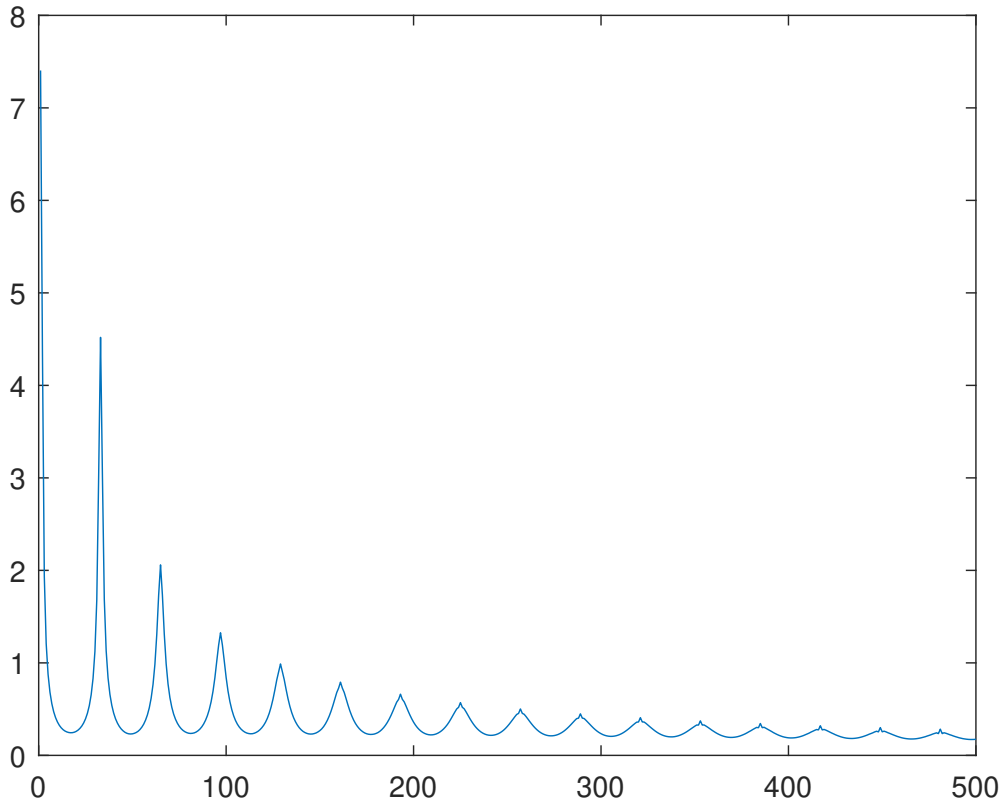

Fig: Magnitude of the mean element along a specific diagonal of $\mathcal{A}^T\mathcal{A}$ against the distance from the main diagonal.

Spectral Properties of $\mathcal{P}^{-1}\mathcal{Q}$

Theorem

Let $\mathcal{P}$ and $\mathcal{Q}$ be defined as before.
The eigenvalues of $\mathcal{P}^{-1}\mathcal{Q}$ satisfy

$$\frac{\rho-\eta}{\rho\Lambda_F+\rho+\eta} \leq \lambda \leq \frac{\sigma_{max}(A)\Lambda_F+\rho+\eta}{\rho\lambda_F+\rho-\eta},$$

where $\Lambda_F \geq \lambda_F$ are the two eigenvalues of matrix

$$F = \begin{bmatrix} c_{11}^2+c_{21}^2 & c_{11}c_{12}+c_{21}c_{22} \\ c_{11}c_{12}+c_{21}c_{22} & c_{12}^2+c_{22}^2 \end{bmatrix}.$$

J. Gondzio, S.-M. Latva-Äijö, S.M Siltanen, M. Lassas, F. Zanetti,
Material-separating regularizer for multi-energy X-ray tomography, *Inverse Problems*, 38 (2022) 025013.

Multi-Energy X-ray Tomography

$$\min_{x \geq 0} \ \|h - \mathcal{A}x\|_2^2 + \alpha \|x\|_2^2 + \beta x^T S x,$$

where $S = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}$ is an *inner product regularizer* which promotes material separation (note $x_1 \geq 0, x_2 \geq 0$, keep $x_1^T x_2$ small).

	CG, $tol = 10^{-6}$			IPCG, $\varepsilon = 10^{-2}$		
Size	IPM	PCG	Time	IPM	PCG	Time
2,048	18	3,810	7.46	19	586	1.44
8,192	20	6,301	35.04	24	1,149	6.29
32,768	23	9,249	140.91	26	1,366	23.02
131,072	26	15,115	817.45	32	1,763	106.36
524,288	29	25,112	5,174.26	49	2,639	639.92

J. Gondzio, S.-M. Latva-Äijö, S.M Siltanen, M. Lassas, F. Zanetti,
Material-separating regularizer for multi-energy X-ray tomography, *Inverse Problems*, 38 (2022) 025013.

Conclusions

2nd-order methods for optimization (including IPMs):

- employ **inexact Newton method**
- rely on **preconditioners**
- enjoy **matrix-free** implementation

Computational practice:

Such methods need:

- **few** iterations
- with $\mathcal{O}(nz(A))$ cost per iteration.

Use IPMs in your research!