



Experience of Linear Solvers in a Nonlinear Interior Point Method

Jonathan Hogg

(joint work with Jennifer Scott)

STFC Rutherford Appleton Laboratory

Software

IPOPT[1]

Interior point optimization code for nonlinear problems

Filter line-search approach

Written by Wächter and Biegler

Part of COIN-OR project

Open Source Eclipse Licence

Probably most widely used open-source IPM solver

Interfaces to a number of linear solvers

[1] A. Wächter and L. T. Biegler, *On the Implementation of a Primal-Dual Interior Point Filter Line Search Algorithm for Large-Scale Nonlinear Programming*, *Mathematical Programming* 106(1), pp. 25-57, 2006

What's happening

Very simplistically...

while not converged **do**

Find a descent direction.

Conduct a line search for next trial point.

(Repeat with second-order corrections).

Take the step; update parameters.

end while

Find a descent direction

Solve

$$Ax = b$$

where

$$A = \begin{pmatrix} W + \Sigma_k & J \\ J^T & -\delta_c I \end{pmatrix}$$

By sparse direct method i.e. factorize with pivoting

$$A = LDL^T$$

Find a descent direction

Solve

$$Ax = b$$

where

$$A = \begin{pmatrix} W + \Sigma_k + \delta_w I & J \\ J^T & -\delta_c I \end{pmatrix}$$

By sparse direct method i.e. factorize with pivoting

$$A = LDL^T$$

Requirements and options

- ▶ Get the “right” answer. Inaccuracy \Rightarrow more IPM iterations.
- ▶ Report correct inertia — required for filter line search to work

Two main options:

1. Static pivoting — if a pivot is too small, add something to it.
Faster, less accurate
2. Threshold pivoting — if a pivot is small, delay until later.
Slower, more accurate

HSL_MA86

Recently developed a new multicore code (OpenMP)
Originally aimed at large finite element and general indefinite problems ($> 10^5$ variables)

Problem	n	MA57	HSL_MA86		
		1	1	8	Speedup
Boeing/bcsstk38	8032	0.152	0.087	0.144	0.60
Oberwolfach/t3dh	79171	20.2	12.1	2.17	5.58
GHS_psdef/s3dkq4m2	90449	2.94	1.94	0.440	4.40
ND/nd12k	36000	164	88.5	15.2	5.83
Oberwolfach/bone010	986703	877	590	88.3	6.68

(MA57 is popular choice for IPOPT; it is also used in MATLAB)

Does HSL_MA86 work well if dropped into IPOPT?

Comparison details

Use 913 solvable problems from CUTEr test set.

529 Trivial $< 0.1s$ to optimum (with MA57)

353 Easy $< 0.5s$ per iteration (with MA57), but not trivial

31 Difficult $> 0.5s$ per iteration (with MA57)

Pitfalls

- ▶ 100% accurate inertia not possible
- ▶ Different paths to optimum
- ▶ Different optima
- ▶ Different matrices

Does scaling help?

Good scalings speed up factorizations of poorly scaled matrices.

MC64 Find weighted maximum matching.

Good; can be slow.

MC77 Several matrix-vector multiplies.

Fast; can be insufficient.

None Free?

Fastest; can cause many delayed pivots.

Best approach varies by problem.

Does scaling help?

Good scalings speed up factorizations of poorly scaled matrices.

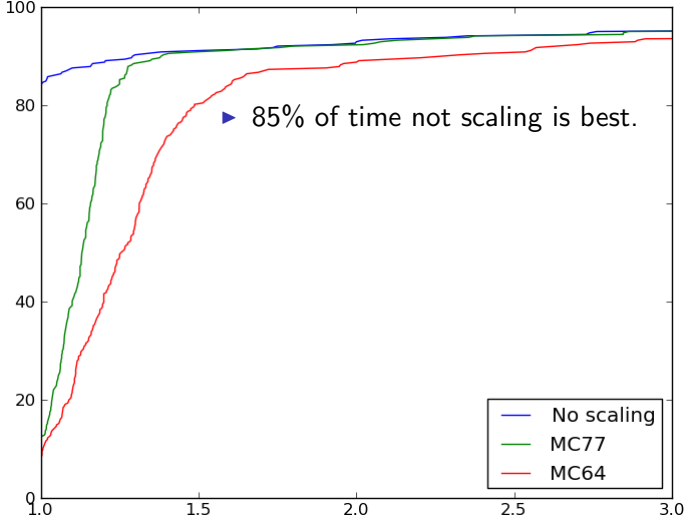
MC64 Find weighted maximum matching. **Not parallel**
Good; can be slow.

MC77 Several matrix-vector multiplies. **Parallelisable**
Fast; can be insufficient.


None Free? **Trivial**
Fastest; can cause many delayed pivots.

Best approach varies by problem.

Scaling helps... sometimes



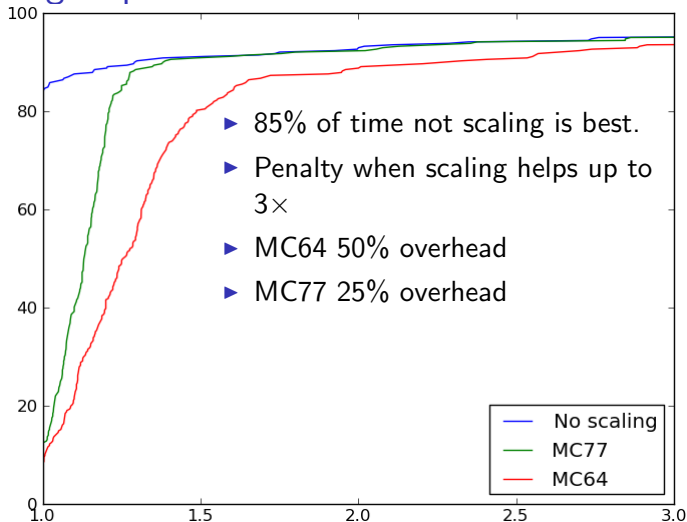
Scaling helps... sometimes




- ▶ 85% of time not scaling is best.
- ▶ Penalty when scaling helps up to 3x

— No scaling
— MC77
— MC64

Scaling helps... sometimes



Scaling helps... sometimes



Scaling examples

Times

Problem	None	MC64	MC77
ELEC	82.7	61.8	74.4
HADAMALS	16.6	24.2	24.7
HVYCRASH	25.0	11.4	42.1
LUKVLE12	6.4	5.8	4.9

Iterations

Problem	None	MC64	MC77
ELEC	324	185	262
HADAMALS	311	307	386
HVYCRASH	748	298	1427
LUKVLE12	30	30	30

Explanations

- ▶ Different paths (Inertia!)
- ▶ More work (Delayed Pivots!)

What should we do?

Based on problem A2ENSNDL

$$\begin{pmatrix} 1e-22 & & & & & & \\ & 1e-22 & & & & & \\ & & 1e-22 & & & & \\ & & & \ddots & & & \\ & 0.3 & 0.2 & 0.4 & \dots & 1e10 & \\ & 0.2 & 0.3 & 0.5 & \dots & 0.5 & 1e10 \end{pmatrix}$$

What is the inertia?

No scaling: Inertia (2,2,996), Maximum front 1000×1000

Parallel?

Modern Intel machine: 2 quad-cores — eight threads

Easy

208 slower

142 faster

3 failed

Min/max/avg speedup = 0.22 / 30.6 / 1.30

Difficult

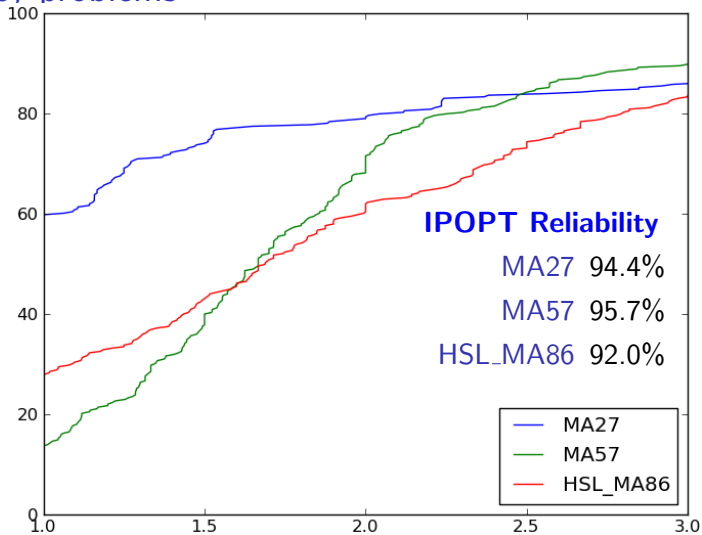
3 slower

24 faster

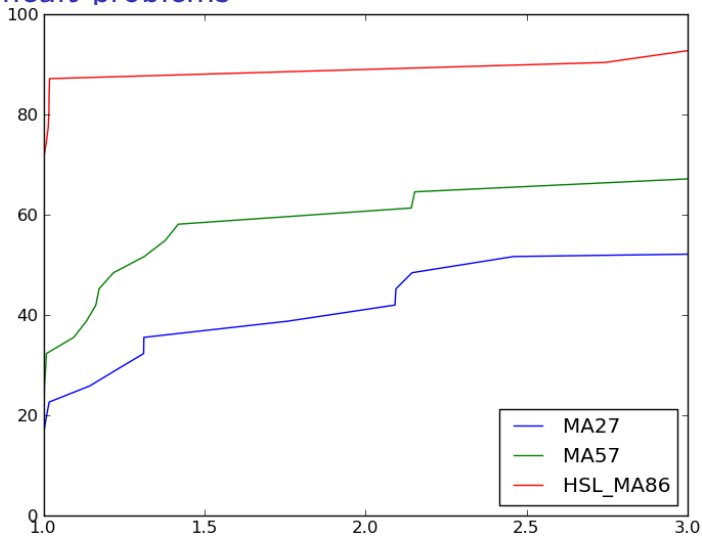
4 failed

Min/max/avg speedup = 0.27 / 5.35 / 1.76

Easy problems



Difficult problems



Open questions?

- ▶ Can a code perform well on both **small** and **large** problems?
- ▶ How should **inertia detection** be handled with respect to **scaling**? **What is zero?**
- ▶ Can we get better parallel performance by driving parallelism up into the IPM somehow?



General HSL: <http://www.hsl.rl.ac.uk>
HSL IPOPT: <http://www.hsl.rl.ac.uk/ipopt>

HSL is freely available to academics



Questions?

General HSL: <http://www.hsl.rl.ac.uk>

HSL IPOPT: <http://www.hsl.rl.ac.uk/ipopt>

HSL is freely available to academics