

A Fast Robust Sparse Mixed Precision Solver

Jonathan Hogg J.D.Hogg@rl.ac.uk
Jennifer Scott J.A.Scott@rl.ac.uk

Rutherford Appleton Laboratory

CERFACS Sparse Days 2008



Science & Technology Facilities Council
Rutherford Appleton Laboratory

Introduction

Aim to solve

$$A\mathbf{x} = \mathbf{b}$$

where A is

- Square
- Symmetric
- Sparse
- Not Small

Using a direct method



Introduction

Aim to solve

$$A\mathbf{x} = \mathbf{b}$$

where A is

- Square
- Symmetric
- Sparse
- Not Small

Using a direct method

- Achieve double precision accuracy as fast as possible.



Introduction

Aim to solve

$$A\mathbf{x} = \mathbf{b}$$

where A is

- Square
- Symmetric
- Sparse
- Not Small

Using a direct method

- Achieve double precision accuracy as fast as possible.
- Achieve double precision accuracy under tight memory constraints.




Test Set

Set of 330 problems drawn from the UFL Test Collection

- Symmetric
- $n > 1000$
- $nnz(A) \leq 10^7$
- Integer or real valued




Single Precision is faster than double...



Goto BLAS on a single core of an Intel E5420 2.50GHz core



Exploit with modern direct solvers...



MA57 factorize on a single core.



But...

We want double precision accuracy.

Aim for

$$\beta = \frac{\|\mathbf{b} - A\mathbf{x}\|_{\infty}}{\|A\|_{\infty}\|\mathbf{x}\|_{\infty} + \|\mathbf{b}\|_{\infty}} \leq 5 \times 10^{-15}$$

Typically a single precision factor-solve $O(10^{-8})$



A solution...

Attempt 1

Iterative Refinement.

Iterative Refinement

Input: ϵ , maxitr , δ

Solve $A\mathbf{x}_1 = \mathbf{b}$ and set $\mathbf{r}_1 = \mathbf{b} - A\mathbf{x}_1$

Set $k = 1$ and compute scaled residual β_1 .

while $\beta_k > \text{accuracy}$ **and** $k \leq \text{maxitr}$ **do**

 Solve $A\mathbf{y}_k = \mathbf{r}_k$

 Set $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k + \mathbf{y}_k$

 Compute $\mathbf{r}_{k+1} = \mathbf{b} - A\mathbf{x}_{k+1}$ and β_{k+1}

if $\beta_{k+1} > \delta\beta_k$ **then exit** {Insufficient Reduction}

$k = k + 1$

end while



Iterative Refinement

Input: ϵ , maxitr , δ

Solve $A\mathbf{x}_1 = \mathbf{b}$ and set $\mathbf{r}_1 = \mathbf{b} - A\mathbf{x}_1$

Set $k = 1$ and compute scaled residual β_1 .

while $\beta_k > \text{accuracy}$ **and** $k \leq \text{maxitr}$ **do**

 Solve $A\mathbf{y}_k = \mathbf{r}_k$

 Set $\mathbf{x}_{k+1} \leftarrow \mathbf{x}_k + \mathbf{y}_k$

 Compute $\mathbf{r}_{k+1} = \mathbf{b} - A\mathbf{x}_{k+1}$ and β_{k+1}

if $\beta_{k+1} > \delta\beta_k$ **then exit** {Insufficient Reduction}

$k = k + 1$

end while



Iterative Refinement Termination

Choice of δ

Big $\delta \Rightarrow$ Good convergence

Small $\delta \Rightarrow$ Quick termination on failure

δ	Converged		δ	Converged
0.001	194		0.3	261
0.01	223		0.4	263
0.05	250		0.5	263
0.1	255		∞	264
0.2	258		Failed	63



A solution...

Attempt 1

Iterative Refinement.

A solution...

Attempt 1

Iterative Refinement.

Attempt 2

Preconditioned FGMRES.



Preconditioned FGMRES

Theoretical Result from Arioli and Duff

Using FGMRES to obtain backward stability in mixed precision

FGMRES can recover precision when iterative refinement fails



Preconditioned FGMRES

Theoretical Result from Arioli and Duff
Using FGMRES to obtain backward stability in mixed precision

FGMRES can recover precision when iterative refinement fails

Assumption

Condition number $< 1/\epsilon$.



FGMRES with adaptive restarting

Input: ϵ , maxitr, restart, max_restart, δ

Solve $\mathbf{Ax} = \mathbf{b}$. Set $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$. Calculate scaled residual β_1 .

while $\beta > \epsilon$ **and** $k < \text{maxitr}$ **do**

 Perform FGMRES(restart)

 Set $\mathbf{x} = \mathbf{x} + \mathbf{Z}_k \mathbf{y}_k$ and $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$

 Recalculate β_{k+1}

if $\beta_{k+1} \geq \delta \beta_k$ **then**

 restart = $2 \times$ restart

if restart > max_restart **then** fail.

end if

end while

FGMRES(restart) exits after restart iterations or when the estimated 2-norm of the absolute error is at machine precision.



FGMRES with adaptive restarting

Input: ϵ , maxitr, restart, max_restart, δ

Solve $\mathbf{Ax} = \mathbf{b}$. Set $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$. Calculate scaled residual β_1 .

while $\beta > \epsilon$ **and** $k < \text{maxitr}$ **do**

 Perform FGMRES(restart)

 Set $\mathbf{x} = \mathbf{x} + Z_k \mathbf{y}_k$ and $\mathbf{r} = \mathbf{b} - \mathbf{Ax}$

 Recalculate β_{k+1}

if $\beta_{k+1} \geq \delta \beta_k$ **then**

 restart = $2 \times$ restart

if restart > max_restart **then** fail.

end if

end while

FGMRES(restart) exits after restart iterations or when the estimated 2-norm of the absolute error is at machine precision.



Solve in high precision

Practical experience

- Accumulate solution in solve phase in double precision.
- Reaches desired accuracy for an extra 20-30% of problems on which iterative refinement fails
- Decreases iteration count for FGMRES by factor of two



Solve in high precision

Practical experience

- Accumulate solution in solve phase in double precision.
 - Reaches desired accuracy for an extra 20-30% of problems on which iterative refinement fails
 - Decreases iteration count for FGMRES by factor of two
-
- Offset by longer solve times for double than single
 - Good reason to have mixed precision BLAS



FGMRES Results

- Success/Failure to reach desired accuracy not affected by choice of initial restart
- **BUT** iteration count varies significantly
- No clear winner, `restart= 4` or `8` is good
- Only works for 60% of cases where Iterative Refinement Fails



A solution...

Attempt 1

Iterative Refinement.

Attempt 2

Preconditioned FGMRES.



A solution...

Attempt 1

Iterative Refinement.

Attempt 2

Preconditioned FGMRES.

Fail

Give up and use double.




Putting it all together

Solved with Iterative Refinement	263	80%
Solved with FGMRES	299	91%
Solved with double (with IR and FGMRES)	328	>99%
Failed in double	2	< 1%



Speedup



Small and Medium sized problems using MA57

Out-of-core

HSL_MA77


- A new **out-of-core** solver for HSL2007.
- Sparse Symmetric Solver (Unsymmetric HSL_MA78)
- Addresses fronts with 64-bit integers
- Experimental OpenMP variant under development

Recently solved a very large mine planning problem

- 3.6×10^6 square matrix
- 1.5×10^{10} entries in L
- Maximum front of order 60121
- Disk Space for factors about 70GB
- 170 minutes using OpenMP on 8 core machine



Speedup 2



Medium and Large problems using HSL_MA77 out-of-core



Science & Technology Facilities Council
 Rutherford Appleton Laboratory

Big Problems

Very large mine planning

- Maximum front of order 60121
- ⇒ 1.8×10^9 entries to address in front
- 6.8 GB for Single Precision
 - 13.5 GB for Double Precision

Our machine only has 8GB of RAM!



Big Problems

Very large mine planning

- Maximum front of order 60121
- ⇒ 1.8×10^9 entries to address in front
- 6.8 GB for Single Precision
- 13.5 GB for Double Precision

Our machine only has 8GB of RAM!

We successfully recovered double precision accuracy through iterative refinement and FGMRES.



Code availability

Our mixed precision code, HSL_MA79, will be added to HSL2007 when it is ready, offering an interface to either MA57 or HSL_MA77.

HSL2007 for Researchers

HSL recently became available to academics free of charge [previously free only to UK academics].

- Google for “HSL 2007”
- <http://hsl.rl.ac.uk/hsl2007/hsl20074researchers.html>

