CCLRC
Rutherford Appleton Laboratory

A note on GMRES preconditioned by a perturbed LDL^T decomposition with static pivoting

M. Arioli, I. S. Duff, S. Gratton, and S. Pralet

Outline

- Multifrontal
- Static pivoting
- GMRES and Flexible GMRES
- Flexible GMRES: a roundoff error analysis
- GMRES right preconditioned: a roundoff error analysis
- Test problems
- Numerical experiments

Linear system

We wish to solve large sparse systems

$$\mathbf{Ax} = \mathbf{b}$$

where $\mathbf{A} \in \mathbf{R}^{\mathrm{N}\times\mathrm{N}}$ is symmetric indefinite

Linear system

A particular and important case arises in saddle-point problems where the coefficient matrix is of the form

$$\begin{bmatrix} H & A \\ A^T & 0 \end{bmatrix}$$

Since we want accurate solutions, we would prefer to use a direct method of solution and our method of choice uses a multifrontal approach.

Multifrontal method

ASSEMBLY TREE

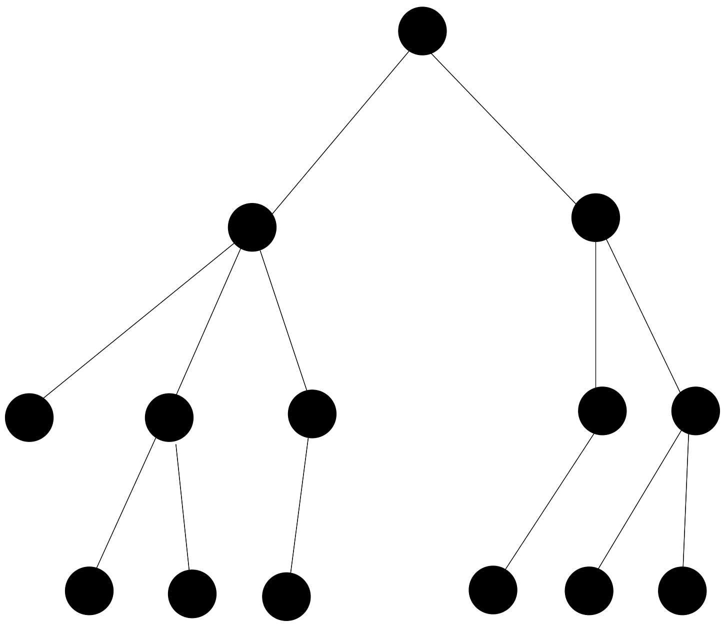

Multifrontal method

ASSEMBLY TREE

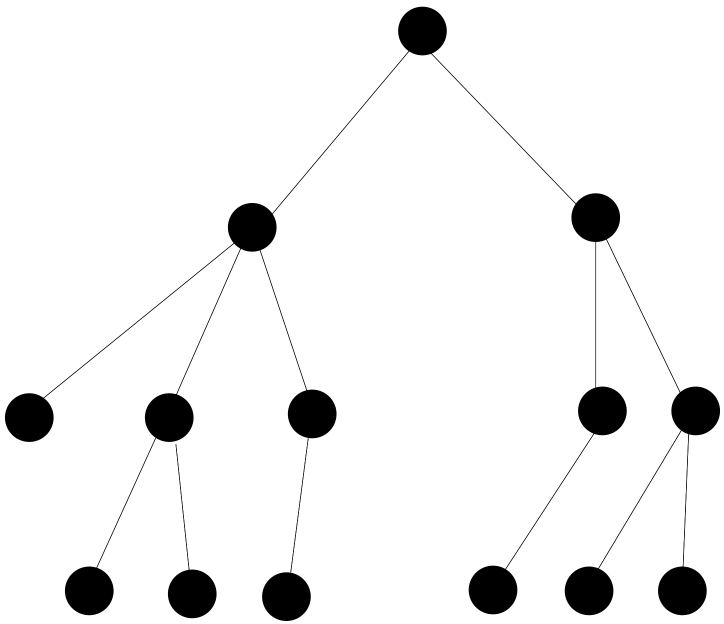

AT EACH NODE

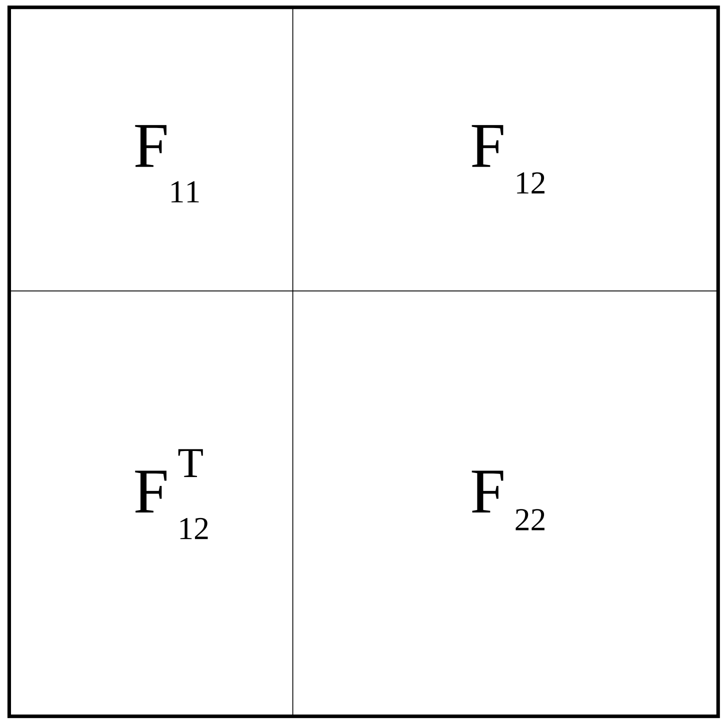

Multifrontal method

ASSEMBLY TREE

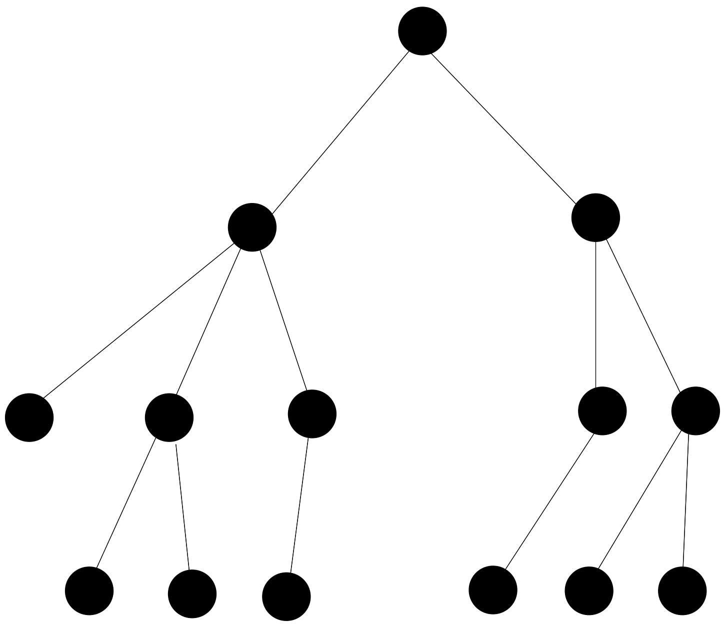

AT EACH NODE

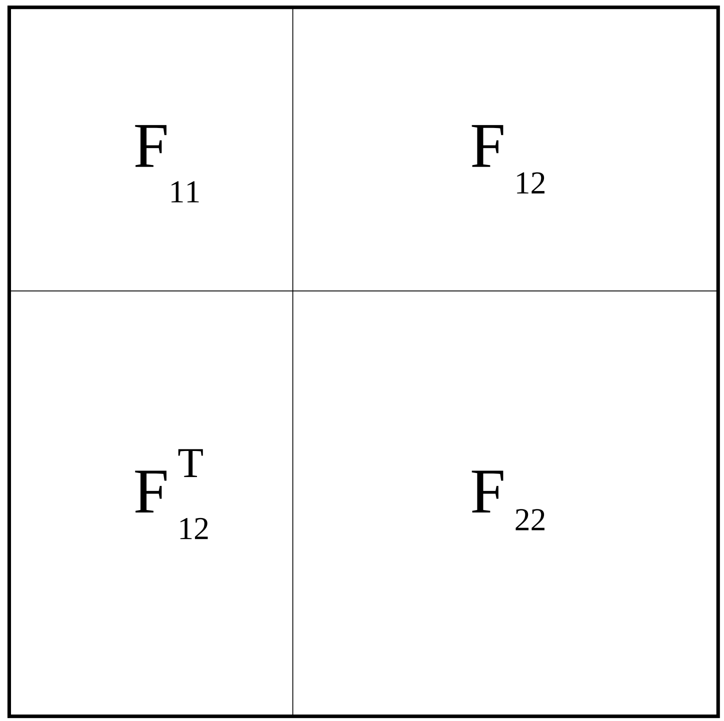

$$F_{22} \leftarrow F_{22} \; - \; F_{12}^T F_{11}^{-1} F_{12}$$

Multifrontal method

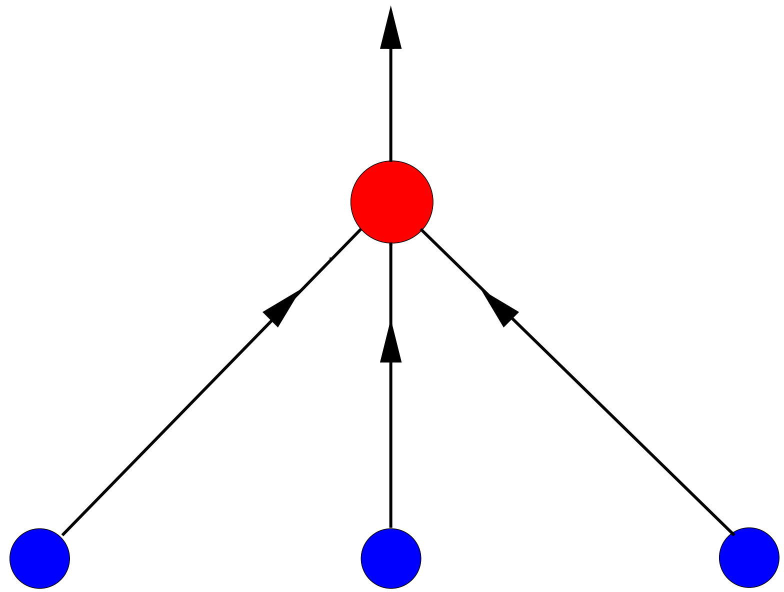

- From children to parent

Multifrontal method

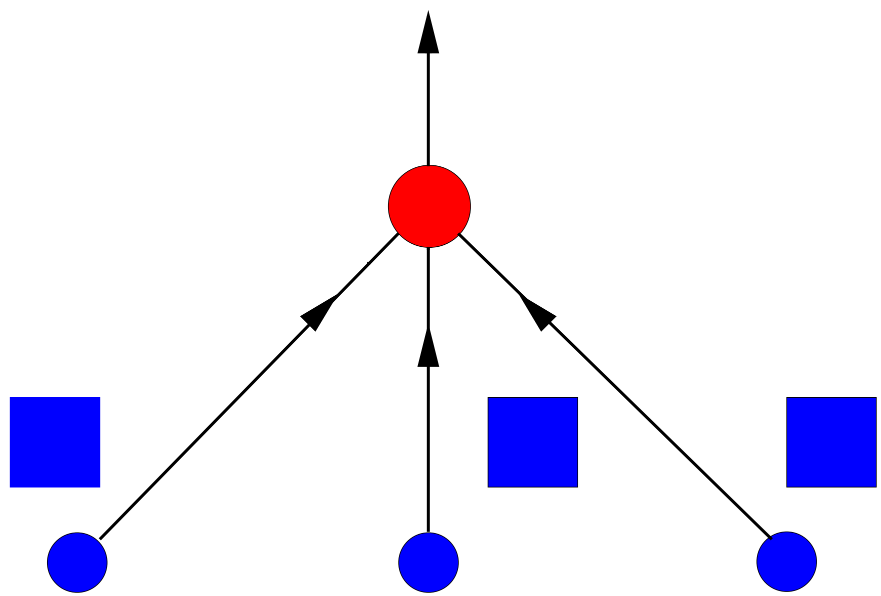

- From children to parent
- **ASSEMBLY** Gather/Scatter operations (indirect addressing)

Multifrontal method

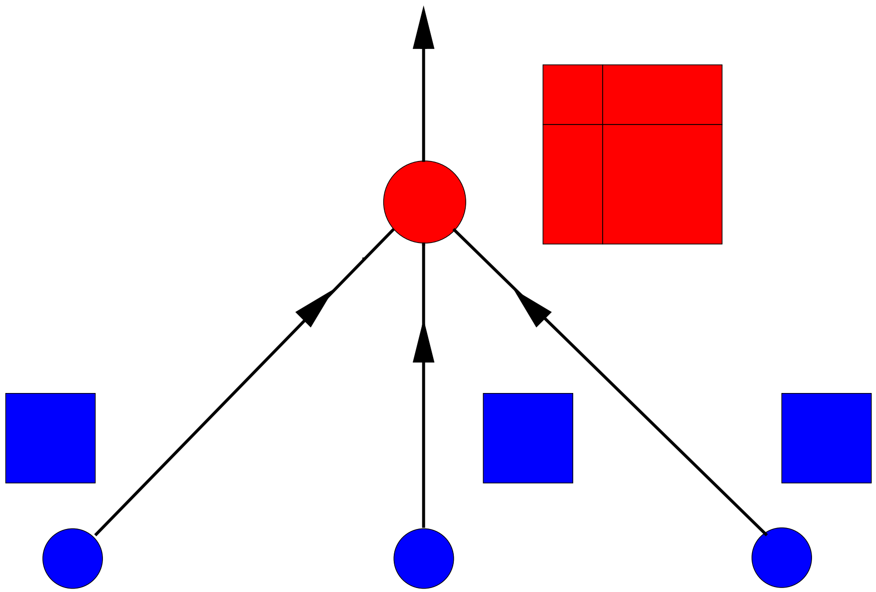

- From children to parent
- **ASSEMBLY** Gather/Scatter operations (indirect addressing)
- **ELIMINATION** Full Gaussian elimination, Level 3 BLAS (TRSM, GEMM)

Multifrontal method

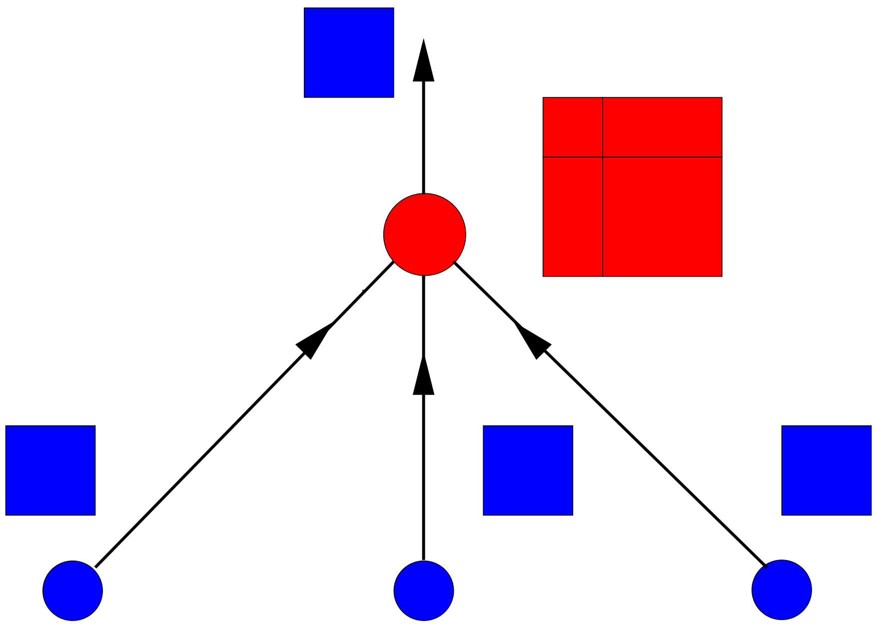

- From children to parent
- **ASSEMBLY** Gather/Scatter operations (indirect addressing)
- **ELIMINATION** Full Gaussian elimination, Level 3 BLAS (TRSM, GEMM)

Multifrontal method

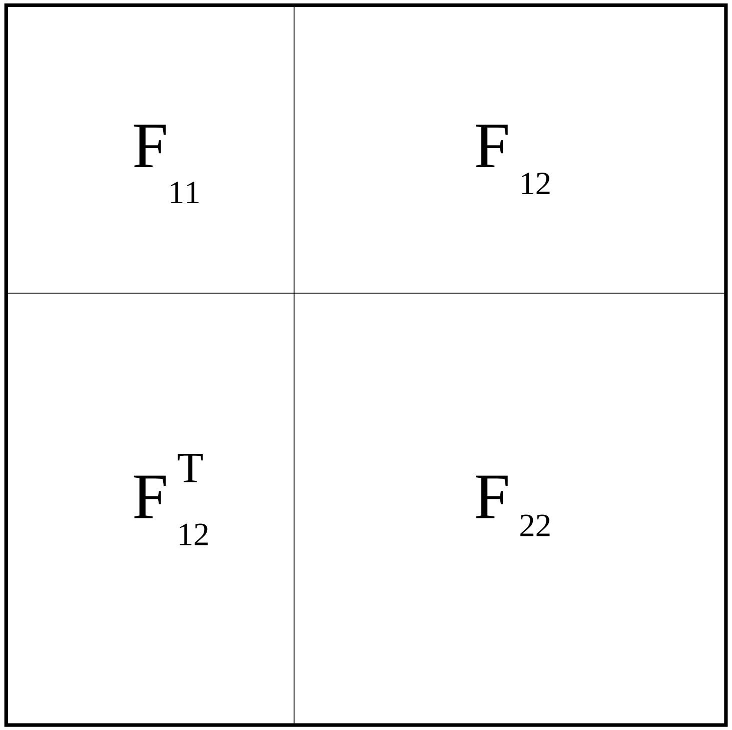

Pivot can only be chosen from F_{11} block since F_{22} is **NOT** fully summed.

Multifrontal method

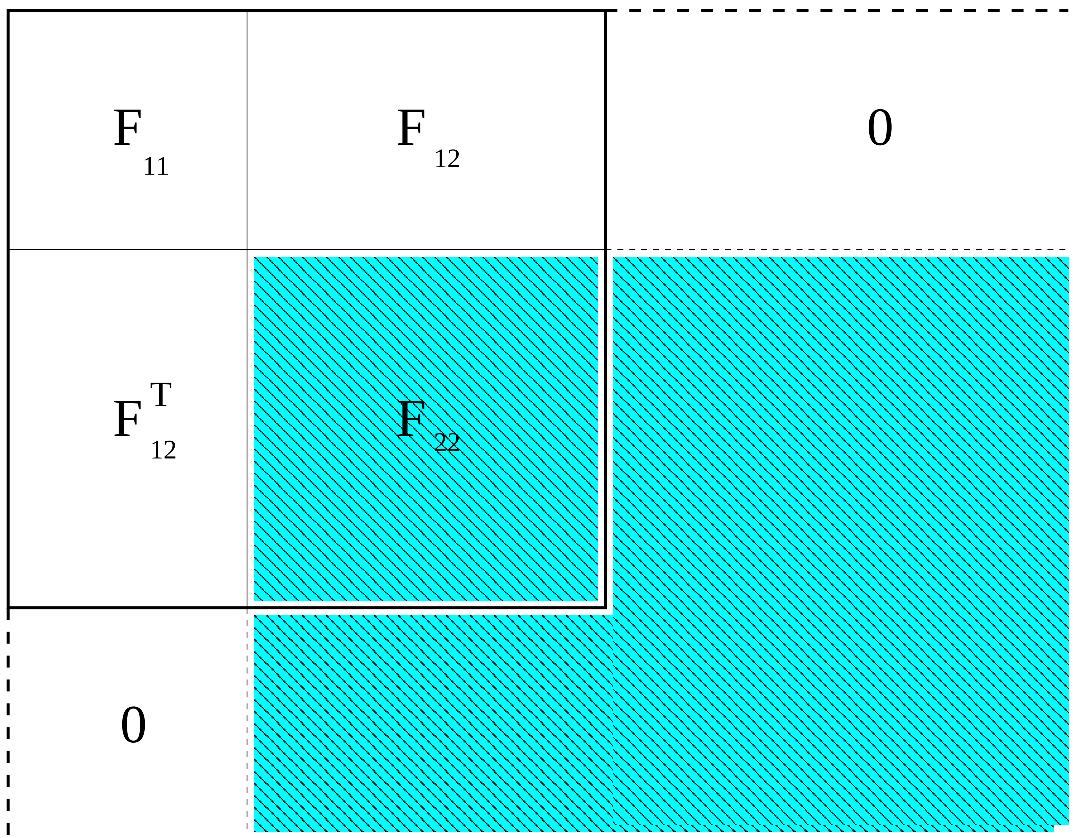

Situation wrt rest of matrix

Pivoting (1×1)

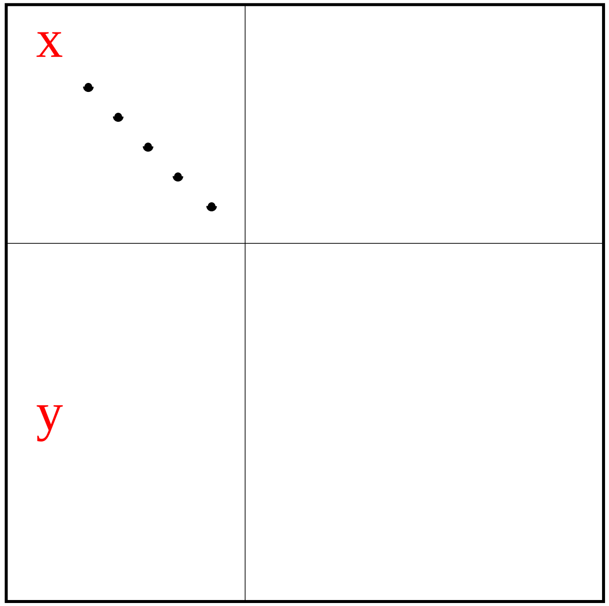

Choose x as 1×1 **pivot** if $|x| > u|y|$
where $|y|$ is the largest in column.

Pivoting (2×2)

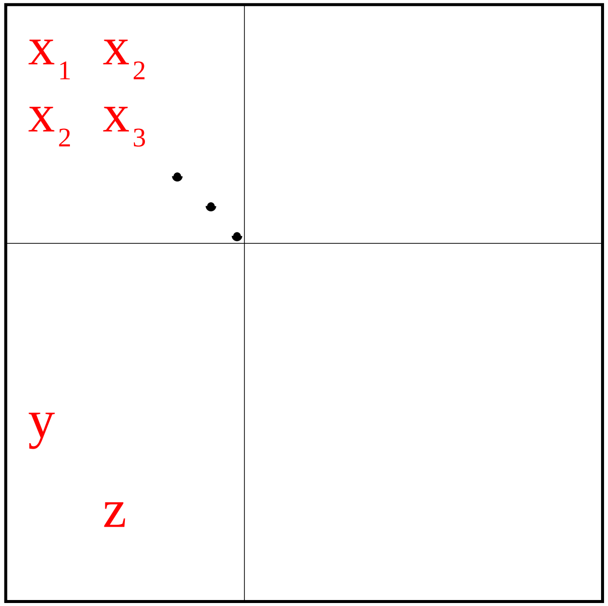

For the indefinite case, we can choose 2×2 **pivot** where we require

$$\left| \begin{bmatrix} x_1 & x_2 \\ x_2 & x_3 \end{bmatrix}^{-1} \right| \begin{bmatrix} |y| \\ |z| \end{bmatrix} \leq \begin{bmatrix} \frac{1}{u} \\ \frac{1}{u} \end{bmatrix}$$

where again $|y|$ and $|z|$ are the largest in their columns.

Pivoting

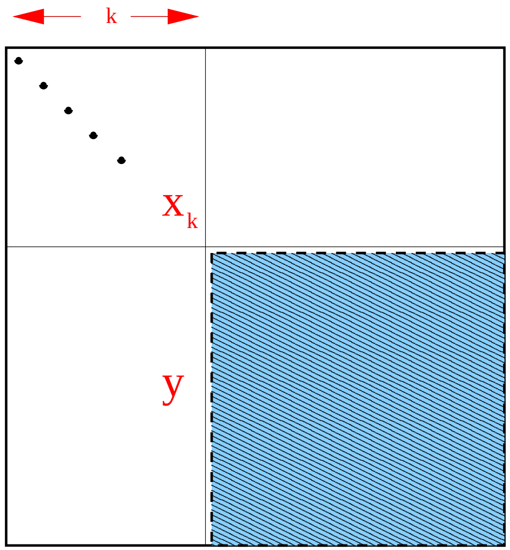

If we assume that $k-1$ pivots are chosen but $|x_k| < u|y|$:

Pivoting

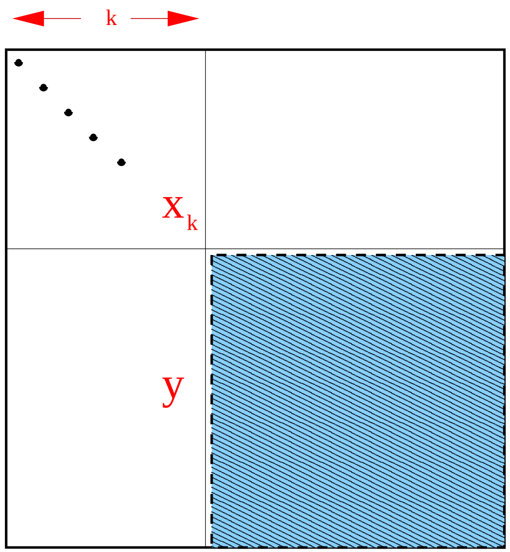

If we assume that $k-1$ pivots are chosen but $|x_k| < u|y|$:

- we can either take the **RISK** and use it or

Pivoting

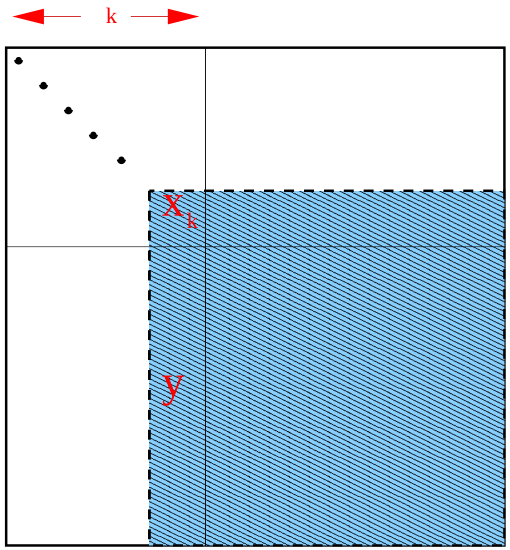

If we assume that $k-1$ pivots are chosen but $|x_k| < u|y|$:

- we can either take the **RISK** and use it or
- **DELAY** the pivot and then send to the parent a larger Schur complement.

Pivoting

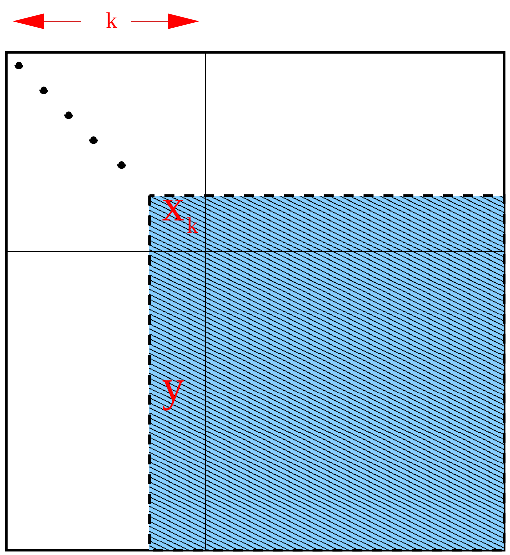

If we assume that $k-1$ pivots are chosen but $|x_k| < u|y|$:

- we can either take the **RISK** and use it or
- **DELAY** the pivot and then send to the parent a larger Schur complement.

This can cause more work and storage

Static Pivoting

An **ALTERNATIVE** is to use **Static Pivoting**, by replacing x_k by

$$x_k + \tau$$

and CONTINUE.

Static Pivoting

An **ALTERNATIVE** is to use **Static Pivoting**, by replacing x_k by

$$x_k + \tau$$

and CONTINUE.

This is even more important in the case of parallel implementation where static data structures are often preferred

Static Pivoting

Several codes use (or have an option for) this device:

- SuperLU (Demmel and Li)
- PARDISO (Gärtner and Schenk)
- MA57 (Duff and Pralet)

Static Pivoting

We thus have factorized

$$A + E = LDL^T = M$$

where $|E| \leq \tau I$

Static Pivoting

We thus have factorized

$$A + E = LDL^T = M$$

where $|E| \leq \tau I$

The three codes then have an **Iterative Refinement** option.
IR will converge if $\rho(M^{-1}E) < 1$

Static Pivoting

Choosing τ

Static Pivoting

Choosing τ

Increase $\tau \Longrightarrow$ increase stability of decomposition

Static Pivoting

Choosing τ

Increase $\tau \Longrightarrow$ increase stability of decomposition

Decrease $\tau \Longrightarrow$ better approximation of the original matrix, reduces $||E||$

Static Pivoting

Choosing τ

Increase $\tau \Longrightarrow$ increase stability of decomposition

Decrease $\tau \Longrightarrow$ better approximation of the original matrix, reduces $||E||$

Trade-off

- $\approx \varepsilon \Longrightarrow$ big growth in preconditioning matrix M
- $\approx 1 \Longrightarrow$ huge error $||E||$.

Static Pivoting

Choosing τ

Increase $\tau \Longrightarrow$ increase stability of decomposition

Decrease $\tau \Longrightarrow$ better approximation of the original matrix, reduces $||E||$

Trade-off

- $\approx \varepsilon \Longrightarrow$ big growth in preconditioning matrix M
- $\approx 1 \Longrightarrow$ huge error $||E||$.

Conventional wisdom is to choose

$$\tau = \mathcal{O}(\sqrt{\varepsilon})$$

Static Pivoting

Choosing τ

Increase $\tau \Longrightarrow$ increase stability of decomposition

Decrease $\tau \Longrightarrow$ better approximation of the original matrix, reduces $||E||$

Trade-off

- $\approx \varepsilon \Longrightarrow$ big growth in preconditioning matrix M
- $\approx 1 \Longrightarrow$ huge error $||E||$.

Conventional wisdom is to choose

$$\tau = \mathcal{O}(\sqrt{\varepsilon})$$

In real life $\rho(M^{-1}E) > 1$

Right preconditioned GMRES and Flexible GMRES

procedure [x] = right_Prec_GMRES(A,M,b)
  $x_0 = M^{-1}b, r_0 = b - Ax_0$ and $\beta = ||r_0||$
  $v_1 = r_0/\beta$; k = 0;
  while $||r_k|| > \mu(||b|| + ||A||\ ||x_k||)$
    $k = k + 1$;
    $z_k = M^{-1}v_k; w = Az_k$;
    for $i = 1, \ldots, k$ do
      $h_{i,k} = v_i^T w$;
      $w = w - h_{i,k}v_i$;
    end for;
    $h_{k+1,k} = ||w||$;
    $v_{k+1} = w/h_{k+1,k}$;
    $V_k = [v_1, \ldots, v_k]$;
    $H_k = \{h_{i,j}\}_{1 \le i \le j+1; 1 \le j \le k}$;
    $y_k = \arg\min_y ||\beta e_1 - H_k y||$;
    $x_k = x_0 + M^{-1}V_k y_k$ and $r_k = b - Ax_k$;
  end while ;
end procedure.

procedure [x] =FGMRES(A,M,b)
  $x_0 = M^{-1}b, r_0 = b - Ax_0$ and $\beta = ||r_0||$
  $v_1 = r_0/\beta$; k = 0;
  while $||r_k|| > \mu(||b|| + ||A||\ ||x_k||)$
    $k = k + 1$;
    $z_k = M^{-1}v_k; w = Az_k$;
    for $i = 1, \ldots, k$ do
      $h_{i,k} = v_i^T w$;
      $w = w - h_{i,k}v_i$;
    end for;
    $h_{k+1,k} = ||w||$;
    $v_{k+1} = w/h_{k+1,k}$;
    $Z_k = [z_1, \ldots, z_k]$; $V_k = [v_1, \ldots, v_k]$;
    $H_k = \{h_{i,j}\}_{1 \le i \le j+1; 1 \le j \le k}$;
    $y_k = \arg\min_y ||\beta e_1 - H_k y||$;
    $x_k = x_0 + Z_k y_k$ and $r_k = b - Ax_k$;
  end while ;
end procedure.

Roundoff error 1

The computed $\hat{L}$ and $\hat{D}$ in floating-point arithmetic satisfy

$$\begin{cases} A + \delta A + \tau E = M \\ ||\delta A|| \leq c(n)\varepsilon || \, |\hat{L}| \, |\hat{D}| \, |\hat{L}^T| \, || \\ ||E|| \leq 1. \end{cases}$$

The perturbation δA must have a norm smaller than τ, in order to not dominate the global error.

Roundoff error 1

The computed $\hat{L}$ and $\hat{D}$ in floating-point arithmetic satisfy

$$\begin{cases} A + \delta A + \tau E = M \\ ||\delta A|| \leq c(n)\varepsilon|| \, |\hat{L}| \, |\hat{D}| \, |\hat{L}^T| \, || \\ ||E|| \leq 1. \end{cases}$$

The perturbation δA must have a norm smaller than τ, in order to not dominate the global error.

A sufficient condition for this is $\boxed{n \, \varepsilon|| \, |\hat{L}| \, |\hat{D}| \, |\hat{L}^T| \, || \leq \tau}$

Roundoff error 1

The computed $\hat{L}$ and $\hat{D}$ in floating-point arithmetic satisfy

$$\begin{cases} A + \delta A + \tau E = M \\ ||\delta A|| \leq c(n)\varepsilon|| \, |\hat{L}| \, |\hat{D}| \, |\hat{L}^T| \, || \\ ||E|| \leq 1. \end{cases}$$

The perturbation δA must have a norm smaller than τ, in order to not dominate the global error.

A sufficient condition for this is $\boxed{n \, \varepsilon|| \, |\hat{L}| \, |\hat{D}| \, |\hat{L}^T| \, || \leq \tau}$

$|| \, |\hat{L}| \, |\hat{D}| \, |\hat{L}^T| \, || \approx \frac{n}{\tau} \Longrightarrow \varepsilon \leq \frac{\tau^2}{n^2}$

CCLRC Rutherford Appleton Laboratory

Roundoff error 1

The computed $\hat{L}$ and $\hat{D}$ in floating-point arithmetic satisfy

$$\begin{cases} A + \delta A + \tau E = M \\ \|\delta A\| \leq c(n)\varepsilon \|\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,\| \\ \|E\| \leq 1. \end{cases}$$

The perturbation δA must have a norm smaller than τ, in order to not dominate the global error.

A sufficient condition for this is $\boxed{n\,\varepsilon \|\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,\| \leq \tau}$

$\|\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,\| \approx \frac{n}{\tau} \Longrightarrow \varepsilon \leq \frac{\tau^2}{n^2}$

Moreover, we assume that $\boxed{\max\{\|M^{-1}\|, \|\bar{Z}_k\|\} \leq \frac{\tilde{c}}{\tau}}$.

Roundoff error 2

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

Roundoff error 2

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

1. Error analysis of the Arnoldi-Krylov process (Giraud and Langou, Björck and Paige, and Paige, Rozložník, and Strakoš).

Roundoff error 2

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

1. Error analysis of the Arnoldi-Krylov process (Giraud and Langou, Björck and Paige, and Paige, Rozložník, and Strakoš).
2. Error analysis of the Givens process used on the upper Hessenberg matrix H_k in order to reduce it to upper triangular form.

CCLRC Rutherford Appleton Laboratory

Roundoff error 2

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

1. Error analysis of the Arnoldi-Krylov process (Giraud and Langou, Björck and Paige, and Paige, Rozložník, and Strakoš).
2. Error analysis of the Givens process used on the upper Hessenberg matrix H_k in order to reduce it to upper triangular form.
3. Error analysis of the computation of x_k in FGMRES and GMRES.

Roundoff error 2

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

1. Error analysis of the Arnoldi-Krylov process (Giraud and Langou, Björck and Paige, and Paige, Rozložník, and Strakoš).
2. Error analysis of the Givens process used on the upper Hessenberg matrix H_k in order to reduce it to upper triangular form.
3. Error analysis of the computation of x_k in FGMRES and GMRES.
4. Use of the static pivoting properties and of $A + E = LDL^T$ in order to have the final expressions.

Roundoff error 2

The roundoff error analysis of both FGMRES and GMRES can be made in four stages:

1. Error analysis of the Arnoldi-Krylov process (Giraud and Langou, Björck and Paige, and Paige, Rozložník, and Strakoš).
2. Error analysis of the Givens process used on the upper Hessenberg matrix H_k in order to reduce it to upper triangular form.
3. Error analysis of the computation of x_k in FGMRES and GMRES.
4. Use of the static pivoting properties and of $A + E = LDL^T$ in order to have the final expressions.

The first two stages of the roundoff error analysis are the same for both FGMRES and GMRES. the last two stages are specific to each one of the two algorithms.

CCLRC Rutherford Appleton Laboratory

Roundoff error FGMRES

We can prove that Flexible GMRES computes a $\bar{x}_k$ s.t.

$$||b - A\bar{x}_k|| \leq c(n,k)\varepsilon(||b|| + ||A||\,||\bar{x}_0|| + ||A||\,||\bar{Z}_k||\,||M(\bar{x}_k - \bar{x}_0)||) + \mathcal{O}(\varepsilon^2)$$

Roundoff error FGMRES

We can prove that Flexible GMRES computes a $\bar{x}_k$ s.t.

$$||b - A\bar{x}_k|| \leq c(n,k)\varepsilon(||b|| + ||A||\,||\bar{x}_0|| + ||A||\,||\bar{Z}_k||\,||M(\bar{x}_k - \bar{x}_0)||) + \mathcal{O}(\varepsilon^2)$$

If $\boxed{c(n,k)\varepsilon||A||\,||\bar{Z}_k|| < 1 \quad \forall k}$

Roundoff error FGMRES

We can prove that Flexible GMRES computes a $\bar{x}_k$ s.t.

$$||b - A\bar{x}_k|| \leq c(n,k)\varepsilon(||b|| + ||A||\,||\bar{x}_0|| + ||A||\,||\bar{Z}_k||\,||M(\bar{x}_k - \bar{x}_0)||) + \mathcal{O}(\varepsilon^2)$$

If $\boxed{c(n,k)\varepsilon||A||\,||\bar{Z}_k|| < 1 \quad \forall k}$

$$||b - A\bar{x}_k|| \leq 2\mu\varepsilon(||b|| + ||A||\,(||\bar{x}_0|| + ||\bar{x}_k||)) + \mathcal{O}(\varepsilon^2).$$

$$\mu = \frac{c(n,k)}{1 - c(n,k)\varepsilon||A||\,||\bar{Z}_k||}$$

Roundoff error right preconditioned GMRES

As we did for FGMRES, we can prove that right preconditioned GMRES computes a $\bar{x}_k$ s.t.

$$\begin{aligned} ||b - A\bar{x}_k|| \leq \quad & c_1(n,k)\chi\varepsilon \Big\{ ||b|| + ||A||\,||\overline{\mathbf{x}}_0|| + ||AM^{-1}||\,||M||\,||\bar{x}_k - \overline{\mathbf{x}}_0|| + \\ & \Big[||A||\,||\bar{Z}_k|| + ||AM^{-1}||\,||M^{-1}||\,||\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,|| \Big] \times \\ & \Big[||M(\bar{x}_k - \overline{\mathbf{x}}_0)|| + n\varepsilon\,||M||\,(||\bar{x}_k - \overline{\mathbf{x}}_0|| + ||\overline{\mathbf{x}}_0||) \Big] \Big\} + \mathcal{O}(\varepsilon^2). \end{aligned}$$

Roundoff error right preconditioned GMRES

As we did for FGMRES, we can prove that right preconditioned GMRES computes a $\bar{x}_k$ s.t.

$$\begin{aligned} ||b - A\bar{x}_k|| \leq \; & c_1(n,k)\chi\varepsilon \Big\{ ||b|| + ||A||\,||\mathbf{\overline{x}}_0|| + ||AM^{-1}||\,||M||\,||\bar{x}_k - \mathbf{\overline{x}}_0|| + \\ & \Big[\, ||A||\,||\bar{Z}_k|| + ||AM^{-1}||\,||M^{-1}||\,||\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,||\,\Big] \times \\ & \Big[\, ||M(\bar{x}_k - \mathbf{\overline{x}}_0)|| + n\varepsilon\,||M||\,(||\bar{x}_k - \mathbf{\overline{x}}_0|| + ||\mathbf{\overline{x}}_0||)\,\Big] \Big\} + \mathcal{O}(\varepsilon^2). \end{aligned}$$

If $\boxed{\zeta = c_2(n,k)\,\varepsilon \Big[\, ||A||\,||\bar{Z}_k|| + ||M^{-1}||\,||\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,||\,\Big] < 1 \quad \forall k}$

Roundoff error right preconditioned GMRES

As we did for FGMRES, we can prove that right preconditioned GMRES computes a $\bar{x}_k$ s.t.

$$\begin{aligned} ||b - A\bar{x}_k|| \leq \; & c_1(n,k)\chi\varepsilon \Big\{ ||b|| + ||A||\,||\overline{\mathbf{x}}_0|| + ||AM^{-1}||\,||M||\,||\bar{x}_k - \overline{\mathbf{x}}_0|| + \\ & \Big[\, ||A||\,||\bar{Z}_k|| + ||AM^{-1}||\,||M^{-1}||\,||\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,||\,\Big] \times \\ & \Big[\, ||M(\bar{x}_k - \overline{\mathbf{x}}_0)|| + n\varepsilon\,||M||\,(||\bar{x}_k - \overline{\mathbf{x}}_0|| + ||\overline{\mathbf{x}}_0||)\Big] \Big\} + \mathcal{O}(\varepsilon^2). \end{aligned}$$

If $\boxed{\zeta = c_2(n,k)\,\varepsilon\Big[\, ||A||\,||\bar{Z}_k|| + ||M^{-1}||\,||\,|\hat{L}|\,|\hat{D}|\,|\hat{L}^T|\,||\,\Big] < 1 \quad \forall k}$

$$||b - A\bar{x}_k|| \leq \; \xi\,\varepsilon\Big[\, ||b|| + \big(||A|| + ||M^{-1}||\big)\big(\,||\bar{x}_k|| + ||\overline{\mathbf{x}}_0||\,\big)\Big] + \mathcal{O}(\varepsilon^2).$$

$$\xi = \frac{c_2(n,k)}{1-\zeta}$$

CCLRC Rutherford Appleton Laboratory

Test Problems

	n	nnz	Description
CONT_201	80595	239596	KKT matrix Convex QP (M2)
CONT_300	180895	562496	KKT matrix Convex QP (M2)
TUMA_1	22967	76199	Mixed-Hybrid finite-element

Test problems

Test Problems: TUMA 1

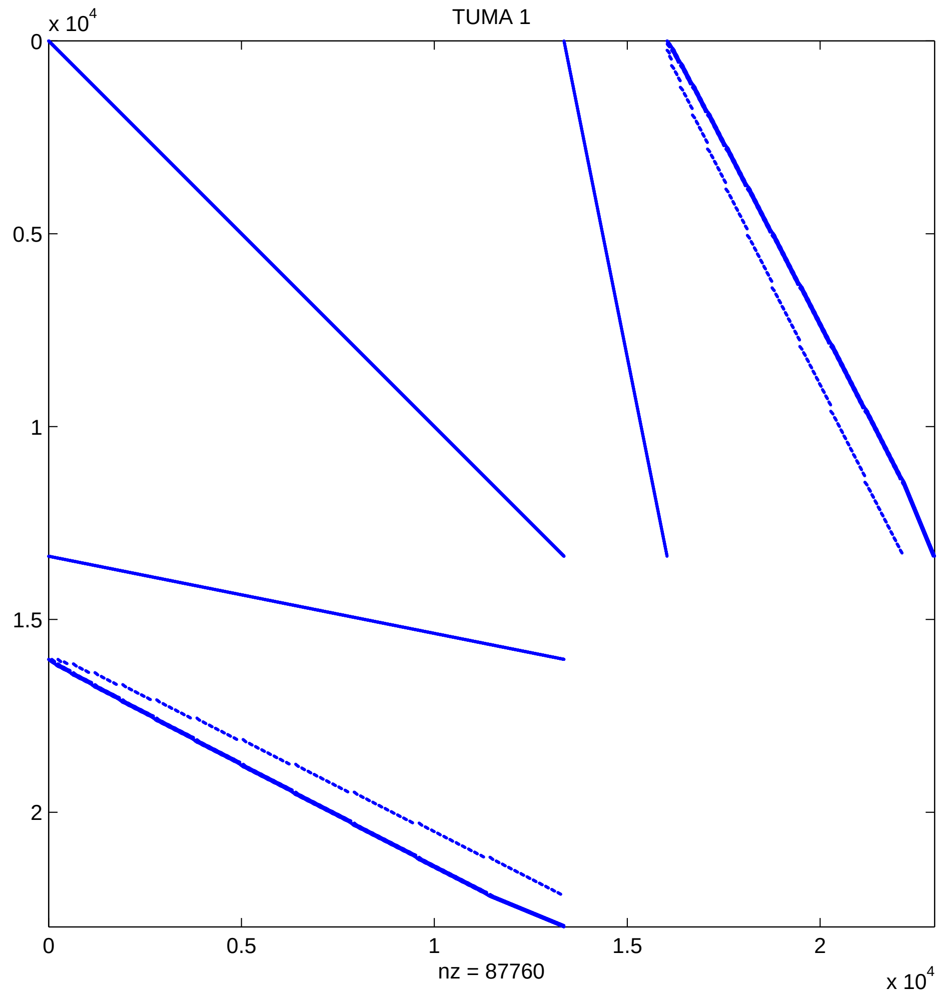

Test Problems: CONT-201

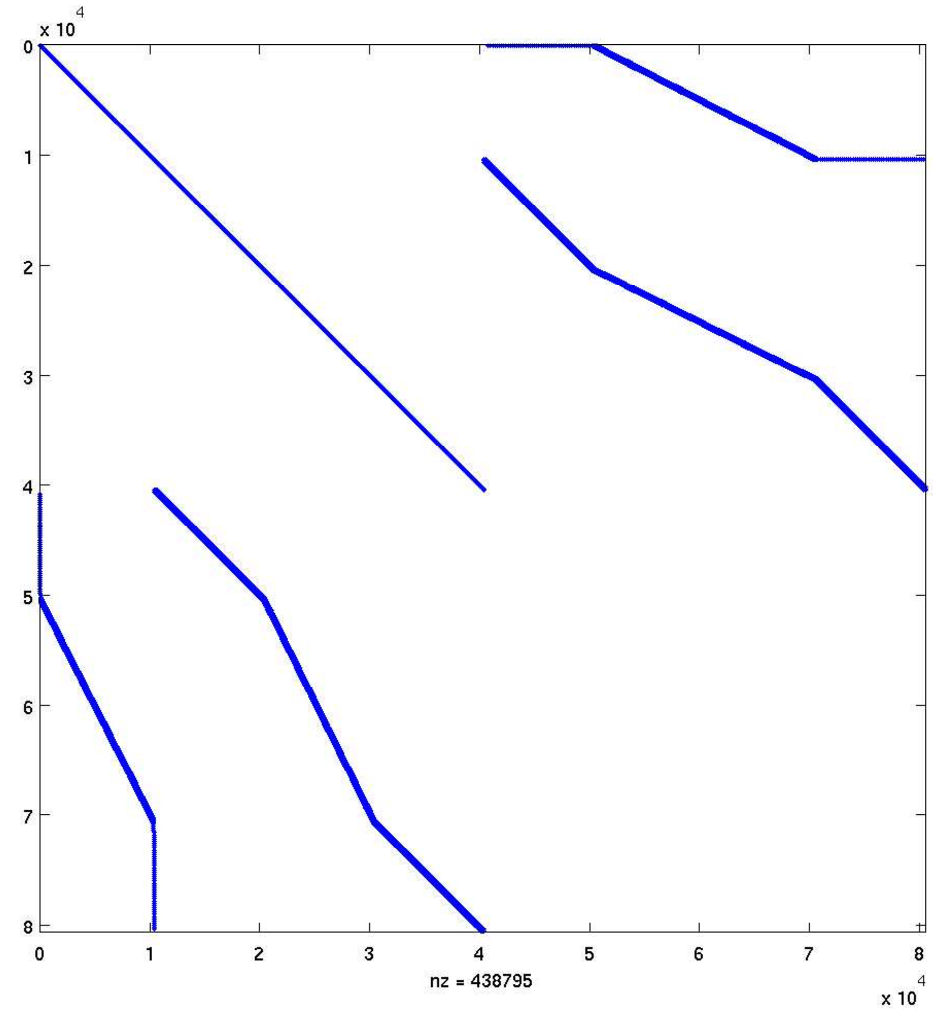

Numerical experiments: TUMA 1

	$\frac{\|\|b - A\bar{x}_k\|\|}{\|\|b\|\| + \|\|A\|\|\|\|\bar{x}_k\|\|}$							
τ	FGMRES	GMRES	$\|\|M\|\|$	$\|\|M^{-1}\|\|$	$\|\|x_0\|\|$	$\|\| \ \|\hat{L}\|\ \|\hat{D}\|\ \|\hat{L}^T\| \ \|\|$	$\|\|\bar{Z}_k\|\|$	$\|\|\overline{\mathbf{x}}_k - \overline{\mathbf{x}}_0\|\|_2$
1e-02	9.3e-17	5.7e-15	4.89	623.57	336.4	3.9e+04	186.14	315.6
1e-04	2.8e-17	9.4e-17	4.89	625.34	89.01	1.5e+05	53.69	2.8
1e-06	2.8e-17	2.0e-16	4.89	625.36	87.70	1.5e+07	42.97	2.6e-02
1e-08	2.8e-17	7.7e-17	4.89	625.76	87.69	1.5e+09	42.95	2.6e-04
1e-10	2.8e-17	3.7e-14	4.89	625.03	87.69	1.5e+11	4.44	1.6e-04
1e-12	2.8e-17	9.0e-10	4.89	626.55	87.69	1.5e+13	26.54	1.4e-02
1e-14	5.7e-17	4.7e-06	4.89	621.60	87.68	1.5e+15	232.13	1.3
1e-15	4.3e-12	9.0e-04	22.98	931.62	91.99	2.2e+17	443.34	19.8

TUMA 1 results

Numerical experiments: CONT_201

	$\frac{\|\|b - A\bar{x}_k\|\|}{\|\|b\|\| + \|\|A\|\|\|\|\bar{x}_k\|\|}$							
τ	FGMRES	GMRES	$\|\|M\|\|$	$\|\|M^{-1}\|\|$	$\|\|x_0\|\|$	$\|\| \|\hat{L}\| \|\hat{D}\| \|\hat{L}^T\| \|\|$	$\|\|\bar{Z}_k\|\|$	$\|\|\overline{\mathbf{x}}_k - \overline{\mathbf{x}}_0\|\|_2$
1e-02	1.1e-03	1.1e-03	8.11	1.9e+04	296.5	5.4e+06	1.7e+04	267.6
1e-04	5.7e-07	5.7e-07	8.08	5.4e+05	394.9	1.9e+08	5.3e+05	372.7
1e-06	7.0e-17	4.6e-13	8.08	4.4e+06	159.2	1.4e+10	4.9e+06	137.0
1e-08	4.5e-17	1.1e-11	8.08	3.6e+07	241.8	2.5e+12	3.5e+07	113.4
1e-10	3.5e-17	7.5e-10	8.08	1.5e+07	163.0	8.7e+14	1.0e+06	1.4
1e-12	3.7e-17	1.8e-08	8.08	1.5e+07	199.6	1.7e+16	5.4e+02	58.5
1e-14	2.6e-15	5.0e-04	197.6	1.2e+07	16504	3.4e+18	1.7e+03	16585
1e-15	1.4e-12	6.8e-02	421.1	4.9e+06	36941	2.5e+19	5.1e+02	37022

CONT_201 results

Numerical experiments: CONT_300

	$\frac{\|\|b - A\bar{x}_k\|\|}{\|\|b\|\| + \|\|A\|\|\|\|\bar{x}_k\|\|}$							
τ	FGMRES	GMRES	$\|\|M\|\|$	$\|\|M^{-1}\|\|$	$\|\|x_0\|\|$	$\|\| \|\hat{L}\| \|\hat{D}\| \|\hat{L}^T\| \|\|$	$\|\|\bar{Z}_k\|\|$	$\|\|\overline{\mathbf{x}}_k - \overline{\mathbf{x}}_0\|\|_2$
1e-02	1.2e-03	1.2e-03	8.1	1.7e+04	345.9	4.2e+07	1.2e+04	286.4
1e-04	1.2e-06	1.2e-06	8.1	3.1e+05	357.0	2.3e+09	2.3e+05	300.2
1e-06	7.9e-15	1.1e-12	8.1	2.5e+06	309.0	2.3e+11	2.1e+06	252.2
1e-08	4.4e-17	4.4e-12	8.1	9.1e+07	266.8	4.0e+13	7.1e+07	284.0
1e-10	3.5e-17	3.4e-08	8.0	7.7e+07	246.0	2.5e+15	6.1e+05	1.8
1e-12	4.1e-17	1.4e-06	16.3	7.6e+07	372.9	3.9e+17	7.3e+02	181.6
1e-14	2.0e-15	9.2e-04	3031.6	7.4e+07	20176	3.5e+19	5.1e+02	20159
1e-15	7.9e-15	2.0e-02	25391	6.1e+07	70577	2.6e+20	5.5e+02	70670

CONT_300 results

Numerical experiments

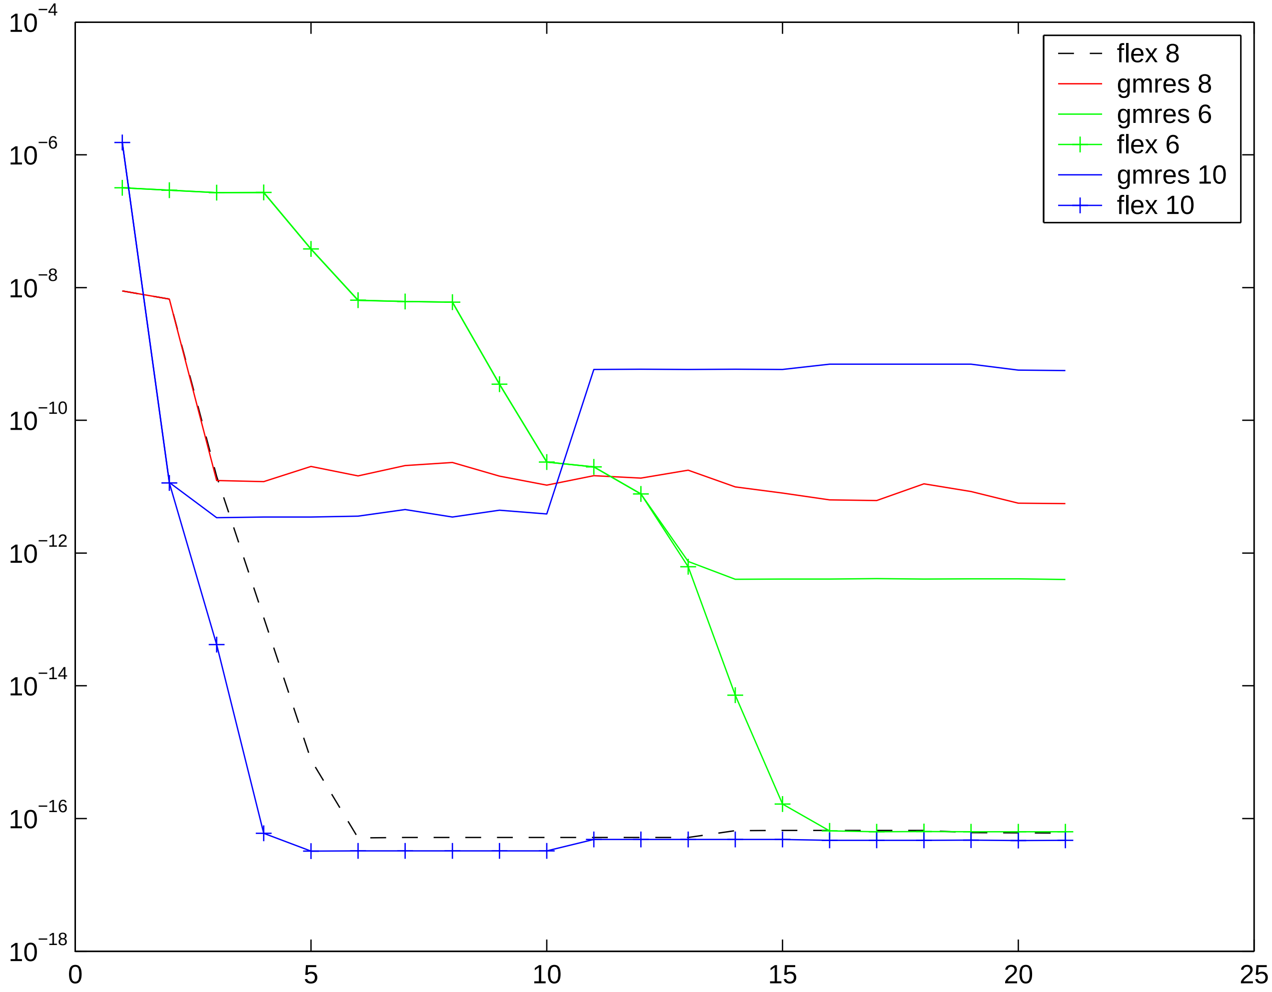

GMRES vs. FGMRES on CONT-201 test example:
$\tau = 10^{-6}, 10^{-8}, 10^{-10}$

Numerical experiments

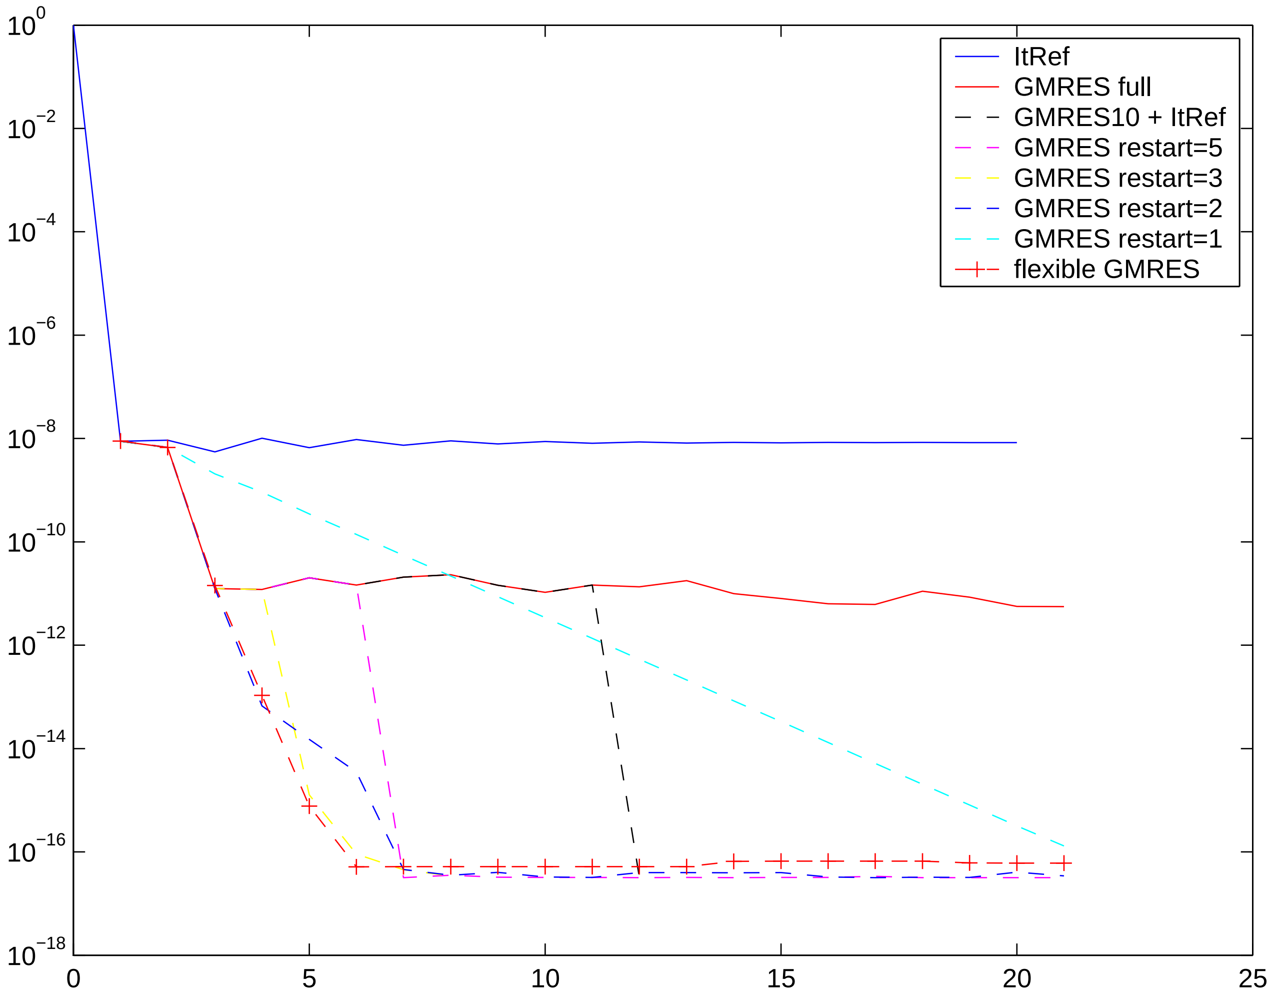

Restarted GMRES vs. FGMRES on CONT-201 test example: $\tau = 10^{-8}$

Numerical experiments

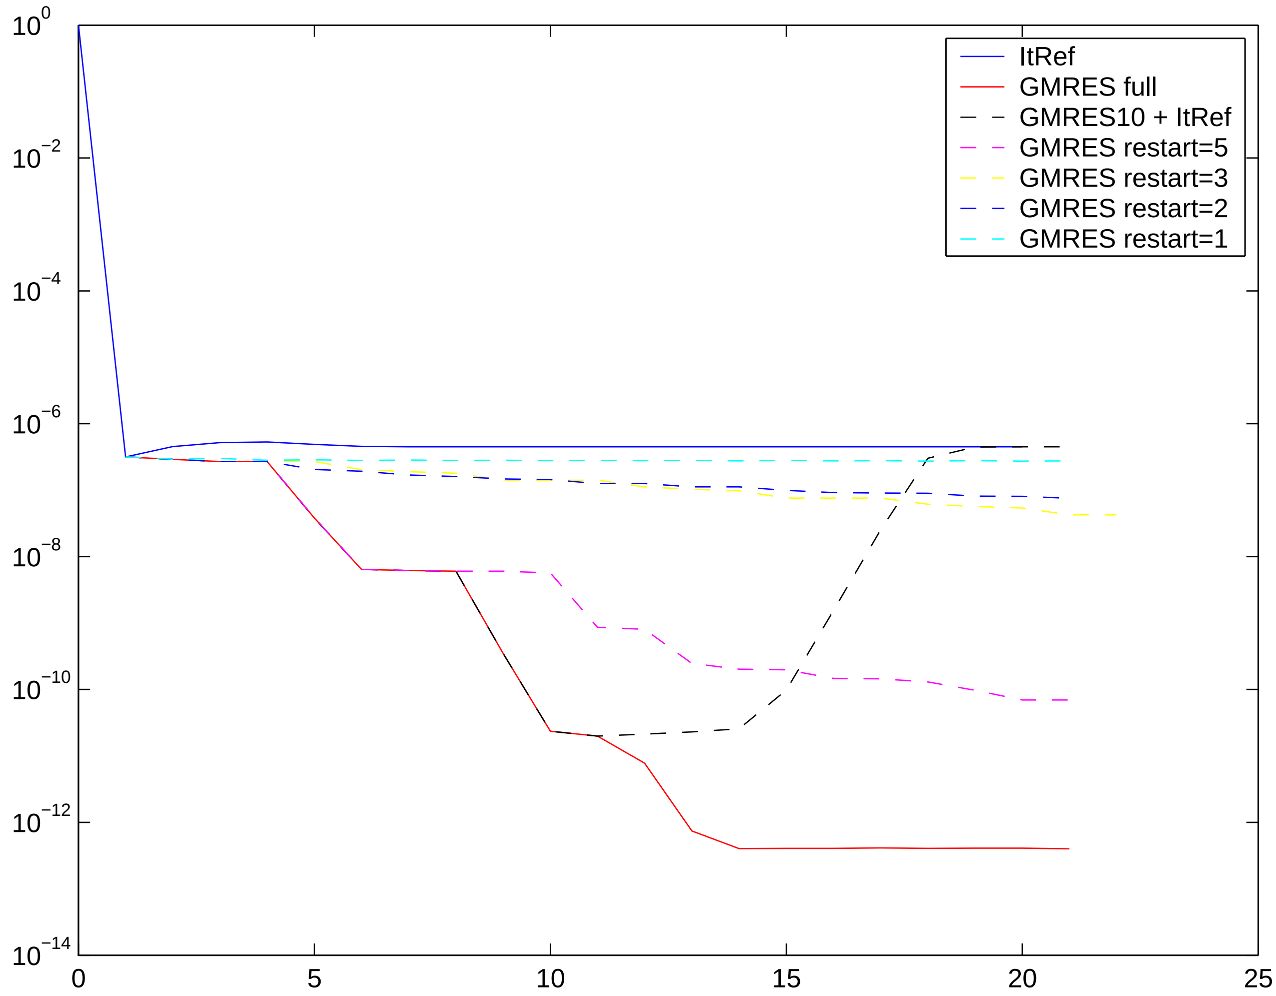

Restarted GMRES on CONT-201 test example: $\tau = 10^{-6}$

CCLRC Rutherford Appleton Laboratory

Summary

- IR with static pivoting is very sensitive to τ and not robust

Summary

- IR with static pivoting is very sensitive to τ and not robust
- GMRES is also sensitive and not robust

Summary

- IR with static pivoting is very sensitive to τ and not robust
- GMRES is also sensitive and not robust
- FGMRES is robust and less sensitive (see roundoff analysis)

CCLRC Rutherford Appleton Laboratory

Summary

- IR with static pivoting is very sensitive to τ and not robust
- GMRES is also sensitive and not robust
- FGMRES is robust and less sensitive (see roundoff analysis)
- Gains from restarting. Makes GMRES more robust, saves storage in FGMRES (but not really needed)

Summary

- IR with static pivoting is very sensitive to τ and not robust
- GMRES is also sensitive and not robust
- FGMRES is robust and less sensitive (see roundoff analysis)
- Gains from restarting. Makes GMRES more robust, saves storage in FGMRES (but not really needed)
- Understanding of why $\tau \approx \sqrt{\varepsilon}$ is best.