VirginiaTech
Invent the Future

Preconditioners and Solvers for Sequences of Linear Systems

Eric de Sturler
Department of Mathematics, Virginia Tech
sturler@vt.edu
http://www.math.vt.edu/people/sturler

Efficient Solution of Large Systems of Nonlinear PDEs in Science
ENS Lyon, Lyon, France, October 7 – 9, 2013

Collaborators and Applications

- Kapil Ahuja (IIT Indore)
- Shun Wang (Credit Suisse), Mike Parks, Chris Siefert, Greg Mackey (Sandia)
- Misha Kilmer, Tufts University
- Glaucio Paulino, Bob Dodds, Philippe Geubelle, UIUC
- Hector Klie, University of Texas at Austin
- David Ceperley, Jeongnim Kim, UIUC, Bryan Clark, Princeton
- Emilio Silva, Luis Motta Mello, University of San Paolo
- Jan Biermann (TU-Harburg, BMW)
- Serkan Gugercin, Chris Beattie (VT)

- Applications: crack propagation, fatigue & fracture, tomography (DOT & EIT), topology optimization, nonlinear equations, materials science and computational physics, Lattice QCD, large-scale fracture in disordered materials, model reduction, electromagnetic wave propagation, stochastic FE/uncertainty, QMC, acoustics (tire noise)
- Funded by NSF (1025327, 1217156), DOE, NASA

Overview

- Sequences of Linear Systems, Problems, Trends
- Solving Sequences of Linear Systems
- Application from Topology Optimization
- Krylov Methods Crash Course
- Recycling Krylov Subspaces
- Convergence for Recycling Approximate Invariant Subspaces
- Application from Computational Acoustics
- Application in the Quantum Monte Carlo Method – Recycling Preconditioners
- Conclusions, Extra Info, Future Work

Solving Sequences of Linear Systems

- Many applications involve a sequence or collection of systems with small or localized changes in space or structure
  - Time-dependent/time-like problems, nonlinear problems and optimization, adaptive discretizations and representations
  - Systems depend (nonlinearly) on multiple parameters
  - Inverse problems and parameter estimation, Monte Carlo and Markov Chain Monte Carlo methods, design, model reduction
  - Uncertainty quantification, reliability (with design)
- Application requires solution of hundreds to thousands of large, sparse, linear systems
- Discretization and coupling of ever more complex problems leads to very large linear systems

Sequences of Linear Systems

- Important trends:
  - Simulations increasingly part of larger analysis, including design, uncertainty/reliability, inverse problems
  - Simulations often involve parameters/parameter space
  - Simulations involve wide ranges of scales and multi-physics, leading to problems of intractable size. Need to drastically reduce the effective number of unknowns: model reduction, parameterizing problems, adaptive meshing
  - We are moving from generic models with idealized properties to realistic models individualized by parameterization (with uncertainty) – such models need to calibrated and then simulated
  - Simulation also used to find parameters that cannot be measured directly
  - New architectures for HPC require new algorithms, but significant support for solving many related problems

Sequences of Problems

- These trends have wider implications for simulation, optimization, design, and so on
- Consider sequences of nonlinear problems, optimizations, inverse problems, …
  - Directly propagate solution spaces or spaces containing/near solution (continuation)
  - Use of model reduction in inverse problems, optimization, UQ
  - Parameterizing problems and how to solve such parameterized problems efficiently

Solving Sequences of Linear Systems

- Even though we can solve linear systems efficiently, we need to reduce the amount of work across multiple systems
- Fast solution by exploiting the slowly changing nature of problem or special structural changes
- Obvious: use previous solution as starting guess for subsequent problem (or space of solutions - Fischer)
- Recycle (adapt & reuse) search spaces from previous problems to improve rate of convergence
- Recycle preconditioners ($Ax = b \rightarrow PAx = Pb$)
- Try to compute spaces of solutions for large number of problems directly
- Recycle other expensive components for linear (or nonlinear) solver – Powell&Gordon recycle components of AMG (preconditioner)

Applications and Examples

- Crack propagation
- Materials Science
- Topology optimization (optimal design of structures)
- Tomography
  - DOT
  - EIT (topology optimization)
- Uncertainty quantification
- Model reduction
- (Quasi-)Newton methods – Nonlinear Systems/Optimization
- Statistical mechanics
- Lattice QCD
- Quantum Monte Carlo
- Computational mechanics, Acoustics
- Nonlinear time-dependent PDEs, …

Example: Topology Optimization

Optimize material distribution, ρ, in design domain

Minimize compliance $u^T K(\rho) u$, where $K(\rho) u = f$

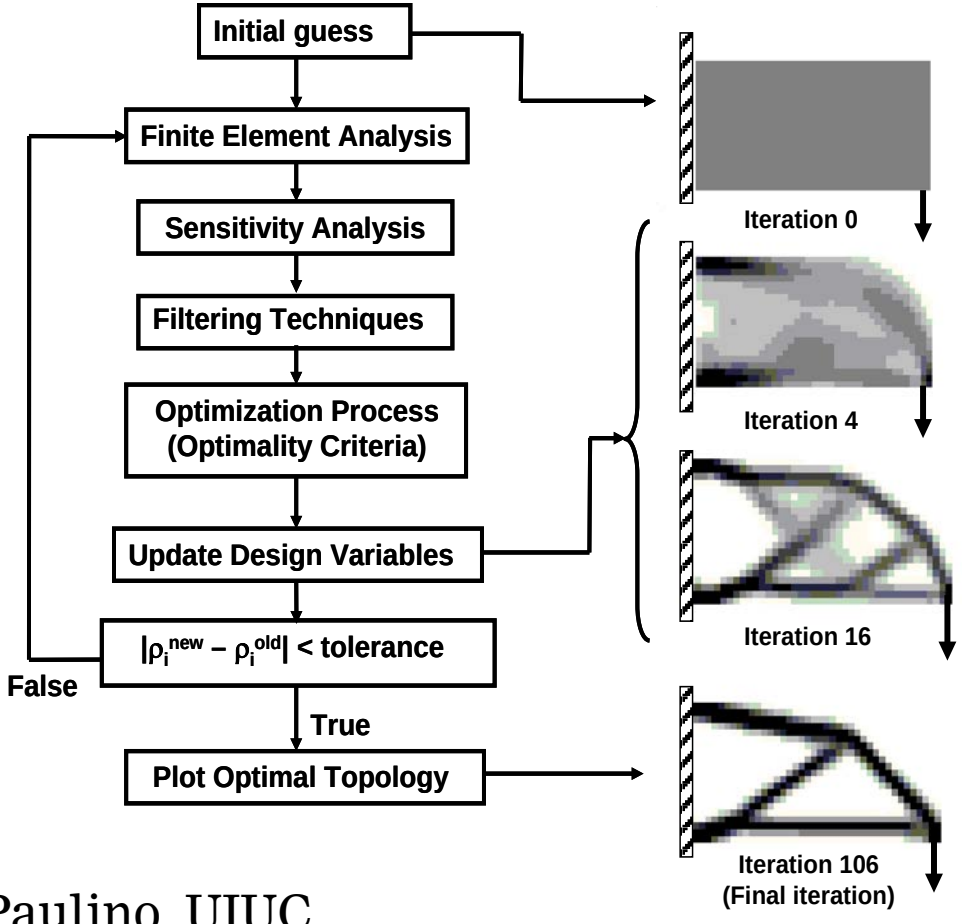

with Glaucio Paulino, UIUC

Topology Optimization - Example

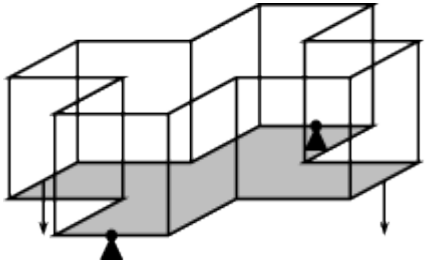

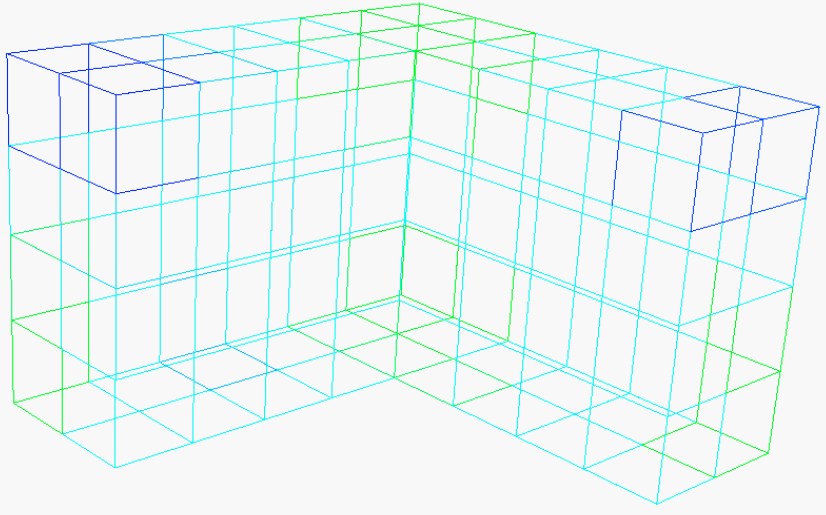

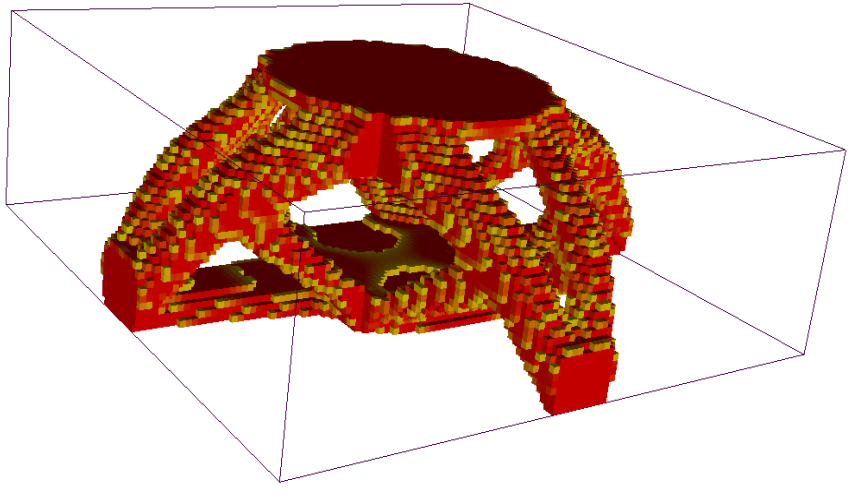

with Glaucio Paulino, UIUC

Krylov Methods Crash Course

Solve $Ax = b$, given initial solution x_0, residual $r_0 = b - Ax_0$

Solve for error, $Ae_0 = r_0$, by finding update from search space

Generate a search space: $K_m\left(A, r_0\right) = \text{span}\left\{r_0, Ar_0, A^2 r_0, \ldots, A^{m-1} r_0\right\}$

Find update $z_m \in K_m\left(A, r_0\right)$ by minimizing

- error in suitable norm (typically for special matrices only)
- residual in suitable norm

Implemented through orthogonal projection (in suitable inner product) – can be expensive.

Alternatively, give up on minimization and compute a cheap projection. Results in some of the most popular methods, but may have some robustness problems.

Krylov Methods Crash Course

Consider $Ax = b$ (or preconditioned system $PAx = Pb$)
Given x_0 and $r_0 = b - Ax_0$, compute optimal update z_m from

$$K^m(A, r_0) = \text{span}\{r_0, Ar_0, \ldots, A^{m-1}r_0\}:$$

$$\min_{z \in K^m(A,r_0)} \|b - A(x_0 + z)\|_2 \quad \Leftrightarrow \quad \min_{z \in K^m(A,r_0)} \|r_0 - Az\|_2$$

Let $K_m = \left[r_0 \; Ar_0 \; A^2 r_0 \cdots A^{m-1} r_0 \right]$, then $z = K_m \zeta$,

and we must solve the following least squares problem

$$AK_m\zeta \approx r_0 \quad \Leftrightarrow \quad \left[Ar_0 \; A^2 r_0 \cdots A^m r_0 \right] \zeta \approx r_0$$

GMRES – Saad and Schulz '86, GCR – Eisenstat, Elman, and Schulz '83

$$K_m\zeta = \zeta_0 r_0 + \zeta_1 Ar_0 + \cdots + \zeta_{m-1} A^{m-1} r_0 = p_{m-1}(A) r_0 \qquad p_{m-1} \text{ arbitrary}$$

$$r_m = r_0 - Ap_{m-1}(A) r_0 = \left(I - Ap_{m-1}(A) \right) r_0 = q_m(A) r_0 \qquad q_m(0) = 1$$

Minimum Residual Solutions: GMRES

Solve $Ax = b$: Choose x_0; set $r_0 = b - Ax_0$; $v_1 = r_0 / \|r_0\|_2$, $k = 0$.

while $\| r_k \|_2 \geq \varepsilon$ do

$k = k + 1$

$\tilde{v}_{k+1} = A v_k$;

for $j = 1 \ldots k$,

$h_{j,k} = v_j^* \tilde{v}_{k+1}$; $\tilde{v}_{k+1} = \tilde{v}_{k+1} - h_{j,k} v_j$;

end

$h_{k+1,k} = \|\tilde{v}_{k+1}\|_2$; $v_{k+1} = \tilde{v}_{k+1} / h_{k+1,k}$;

Solve LS $\min_\zeta \left\| \eta_1 \|r_0\|_2 - \underline{H}_k \zeta \right\|_2 \left(= \|r_k\|_2 \right)$ by construction

(in practice we update the residual of LS each step)

end

$x_k = x_0 + V_k \zeta$;

$r_k = r_0 - V_{k+1} \underline{H}_k \zeta = V_{k+1} \left(\eta_1 \|r_0\| - \underline{H}_k \zeta \right)$ or simply $r_k = b - Ax_k$

Why solution from projection on small space accurate?

Krylov space is a *space of polynomials in a matrix times a vector*.

Krylov space inherits the approximation properties of polynomials on the real line or in the complex plane.

Let A be diagonalizable, $A = V\Lambda V^{-1}$ (simplify explanation)
Then $A^2 = V\Lambda V^{-1}V\Lambda V^{-1} = V\Lambda^2 V^{-1}$ and generally $A^k = V\Lambda^k V^{-1}$.
So, the polynomial $p_m(t) = \alpha_0 + \alpha_1 t + \cdots + \alpha_m t^m$ applied to A gives

$$p_m(A) = V\left(\alpha_0 I + \alpha_1 \Lambda + \alpha_2 \Lambda^2 + \cdots + \alpha_m \Lambda^m\right)V^{-1} \quad \text{and hence}$$
$$p_m(A) = V p_m(\Lambda) V^{-1} = V\,\mathrm{diag}\left(p_m(\lambda_1),\ldots,p_m(\lambda_n)\right)V^{-1}$$

The polynomial is applied to the eigenvalues individually.

Approximate solutions to linear systems, eigenvalue problems, and more general problems using polynomial approximation can be analyzed/understood this way.

Approximation by Matrix Polynomials

Let $A = V \Lambda V^{-1}$, let $\Lambda(A) \subset \Omega \subset \mathbb{C}$.

If $p_{m-1}(t) \approx 1/t$ for all $t \in \Omega$, then $p_{m-1}(A) \approx A^{-1}$

Let $r_0 = V\rho$. Then $p_{m-1}(A) r_0 = \sum_i v_i p_{m-1}(\lambda_i) \rho_i \approx \sum_i v_i \frac{\rho_i}{\lambda_i} = A^{-1} r_0$

Also $r_m = q_m(A) r_0 = \left(I - A p_{m-1}(A)\right) r_0 = \sum_i v_i q_m(\lambda_i) \rho_i \approx 0$

We want $q_m(\lambda_i)$ small for all λ_i. If we can construct such polynomials for modest m, we have an efficient linear solver.

This is possible if the region Ω is nice – small region away from origin: clustered eigenvalues

If this is not the case, we improve by *preconditioning*: $PAx = Pb$ s.t. PA has clustered eigenvalues and product with P is cheap.

Convergence Bounds

Relate convergence to polynomials. Residual (error) at step m

$r_m = r_0 - Az_m = p_m(A)r_0$ and $e_m = p_m(A)e_0$

CG: $\|e_m\| / \|e_0\| \le \min_{p(0)=1} \max_i |p(\lambda_i)| \le C_m(q_1(\lambda)) / C_m(q_1(0))$

MINRES: $\|r_m\| / \|r_0\| \le \min_{p_m(0)=1} \max_i |p_m(\lambda_i)| \le C_{m/2}(q_2(\lambda)) / C_{m/2}(q_2(0))$

(as CG for HPD matrices)

GMRES: $\|r_m\| / \|r_0\| \le \kappa(V) \min_{p_m(0)=1} \max_{\lambda \in \Lambda(A)} |p_m(\lambda)| \le \kappa(V) C_m(\lambda) / C_m(\gamma)$

$\kappa(V)$ small: convergence determined by minimal polynomial

GMRES (FOV): $\|r_m\| \le 2\|r_0\| \min_{\substack{p \in \Pi_m \\ p(0)=1}} \max_{\gamma \in W(A)} |p(\gamma)|$

What to Recycle?

- Krylov methods build search space; get solution by projection
- Building search space often dominates cost
- Initial convergence often poor, reasonable size search space needed, then superlinear convergence
- Get fast convergence rate and good initial guess immediately by recycling selected search spaces from previous systems
- How to select the right subspace to recycle?
  - Invariant subspace - GMRESDR (Morgan'04, '95), with recycling: GCRODR (Parks&dSJM'06), RMINRES (Wang&dSP'07), RBiCG (2010)
  - Preconditioning based on invariant subspaces (mostly for fixed matrix): Vuik, Nabben, Saad, Erhel, Burrage, Loghin, Rey, Risler
  - Canonical angles between successive spaces - GCROT (dS'99), with recycling (Parks&dSJM'06)
  - Subspace from previous solutions - initial guess (Fischer'96), for recycling/combined with invariant subspace (Kilmer&dS'06)
- Structural update matrix in QMC – preconditioners Ahuja, dS et al '10

How to Recycle?

Solve $Ax = b$ using recycled subspace/matrix $\tilde{U}$ (for new A):

Compute $A\tilde{U} = \tilde{C}$, $CR = \tilde{C}$ (QR), $U = \tilde{U}R^{-1}$ (implicit)
Now $AU = C$ and $C^*C = I$

Set $r_0 = \left(I - CC^*\right)b$, $x_0 = UC^*b$, and $v_1 = r_0 \,/\, \|r_0\|$
Augmented Arnoldi: $AV_m = CC^*AV_m + V_{m+1}\underline{H}_m = CB + V_{m+1}\underline{H}_m$

Minimize $\left\|b - A\left(x_0 + Uz + V_m y\right)\right\| = \left\|r_0 - Cz - CBy - V_{m+1}\underline{H}_m y\right\| =$

$$\left\|V_{m+1}\left(e_1\|r_0\| - \underline{H}_m y\right) - C\left(z + By\right)\right\| \qquad \text{(optimal)}$$

Solve $\underline{H}_m y \approx e_1\|r_0\|$ and set $z = -By$
$x_m = x_0 + Uz + V_m y$ and $r_m = V_{m+1}\left(e_1\|r_0\| - \underline{H}_m y\right)$ (GCRO, dS'95)

Recycling Krylov Subspaces Using GCRO

By construction $\left(I - CC^*\right)A = \left(I - CC^*\right)A\left(I - CC^*\right)$ over

$$\mathrm{R}\left(C\right)^{\perp} \supseteq K_*\left(\left(I - CC^*\right)A, v_1\right)$$

So, $\left(I - CC^*\right)A$ is self-adjoint if A self-adjoint and

we can use MINRES, CG, etc (Kilmer&dS'06,Wang et al'07)

Analyze the convergence of GMRES, MINRES, CG,

by considering $\left(I - CC^*\right)A\left(I - CC^*\right)$ i.s.o. $\left(I - CC^*\right)A$

Convergence bounds from eigenvalues for MINRES or CG
Meinardus'63, Kaniel&Daniel'67,books Saad, Greenbaum, vdVorst
Convergence bounds from field of values for GMRES
Eiermann'93, Greenbaum'97, Embree'99

Recycling MINRES

- More complicated than for GCRODR and RGCROT (for non-Hermitian matrices)
- For full recurrence methods we restart anyway (for a single system) with new recycle space to minimize effect of restarting on rate of convergence
- For MINRES (CG) there is no restart and we do not replace the recycle space for a single system (non-optimal)
  - We do not want to keep all the Lanczos vectors
  - Method is cheaper so overhead more detrimental

- Updating the recycle space
  - Periodically update the recycle space (for next system) with new Lanczos vectors (which can be discarded afterwards)
  - Must be done cheaply on the side

Recycling MINRES

Compute the Harmonic Ritz Vectors

- A harmonic Ritz pair $(\theta, \boldsymbol{w})$ of $\boldsymbol{K}$ with respect to range$(\boldsymbol{W})$ satisfies

$$\boldsymbol{K}\boldsymbol{w} - \theta\boldsymbol{w} \perp \text{range}(\boldsymbol{K}\boldsymbol{W})$$

  where $\boldsymbol{w} \in \text{range}(\boldsymbol{W})$, or $\boldsymbol{w} = \boldsymbol{W}\boldsymbol{p}$

- After the j-th cycle, define

$$\boldsymbol{W}_j = [\boldsymbol{U}_{j-1}\ \boldsymbol{V}_j], \qquad \widetilde{\boldsymbol{W}}_j = [\boldsymbol{C}\ \boldsymbol{C}_{j-1}\ \underline{\boldsymbol{V}}_j], \qquad \underline{\widetilde{\boldsymbol{H}}}_j = \begin{bmatrix} \boldsymbol{0} & \boldsymbol{B}_j \\ \boldsymbol{I} & \boldsymbol{0} \\ \boldsymbol{0} & \underline{\boldsymbol{T}}_j \end{bmatrix}$$

  where $\boldsymbol{B}_j = \boldsymbol{C}^T\boldsymbol{K}\boldsymbol{V}_j$. Then we have

$$\boldsymbol{K}\boldsymbol{W}_j = \widetilde{\boldsymbol{W}}_j\underline{\widetilde{\boldsymbol{H}}}_j$$

- Compute harmonic Ritz vectors by solving generalized eigenvalue problem:

$$\underline{\widetilde{\boldsymbol{H}}}_j^T\widetilde{\boldsymbol{W}}_j^T\widetilde{\boldsymbol{W}}_j\underline{\widetilde{\boldsymbol{H}}}_j\boldsymbol{p} = \theta\underline{\widetilde{\boldsymbol{H}}}_j^T\widetilde{\boldsymbol{W}}_j^T\boldsymbol{W}_j\boldsymbol{p}$$

Recycling MINRES

Straightforward setup of eigenvalue problem expensive
All but one block can be written as recursion
In newer version all C-blocks rewritten in terms of 'starting' C and recursion over Lanczos vectors

Solve the Generalized Eigenvalue Problem

$$\widetilde{\underline{H}}_j^T \widetilde{W}_j^T \widetilde{W}_j \widetilde{\underline{H}}_j p = \theta \widetilde{\underline{H}}_j^T \widetilde{W}_j^T W_j p$$

$$\widetilde{W}_j^T \widetilde{W}_j = \begin{bmatrix} I & C^T C_{j-1} & 0 \\ C_{j-1}^T C & I & C_{j-1}^T \overline{V}_j \\ 0 & \overline{V}_j^T C_{j-1} & I \end{bmatrix}$$

$$\widetilde{W}_j^T W_j = \begin{bmatrix} C^T U_{j-1} & 0 \\ C_{j-1}^T U_{j-1} & C_{j-1}^T V_j \\ \overline{V}_j^T U_{j-1} & \underline{I} \end{bmatrix}$$

Topology Optimization Convergence Results

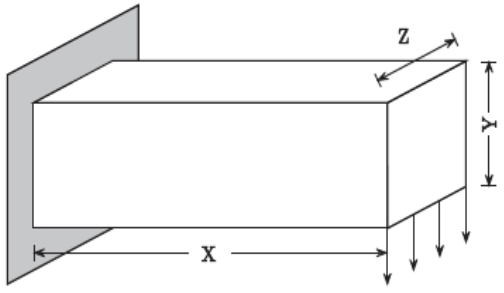

Results: 84x28x14 mesh (107K dofs) and 180x60x30 (1M dofs) (on PC)

Currently more complicated models up to 2M dof (on PC)

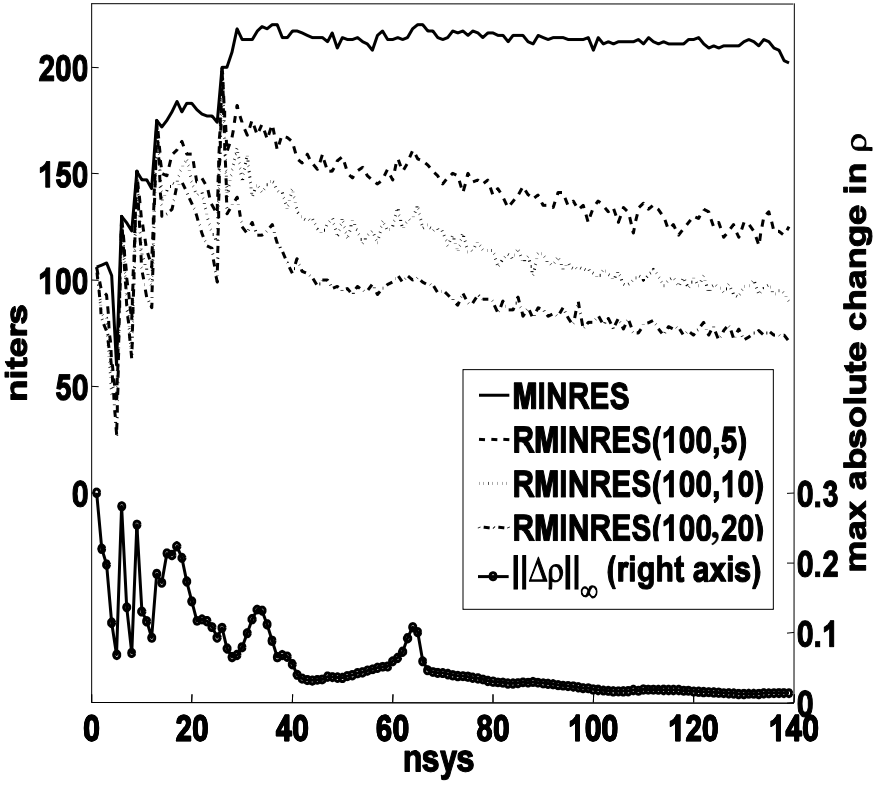

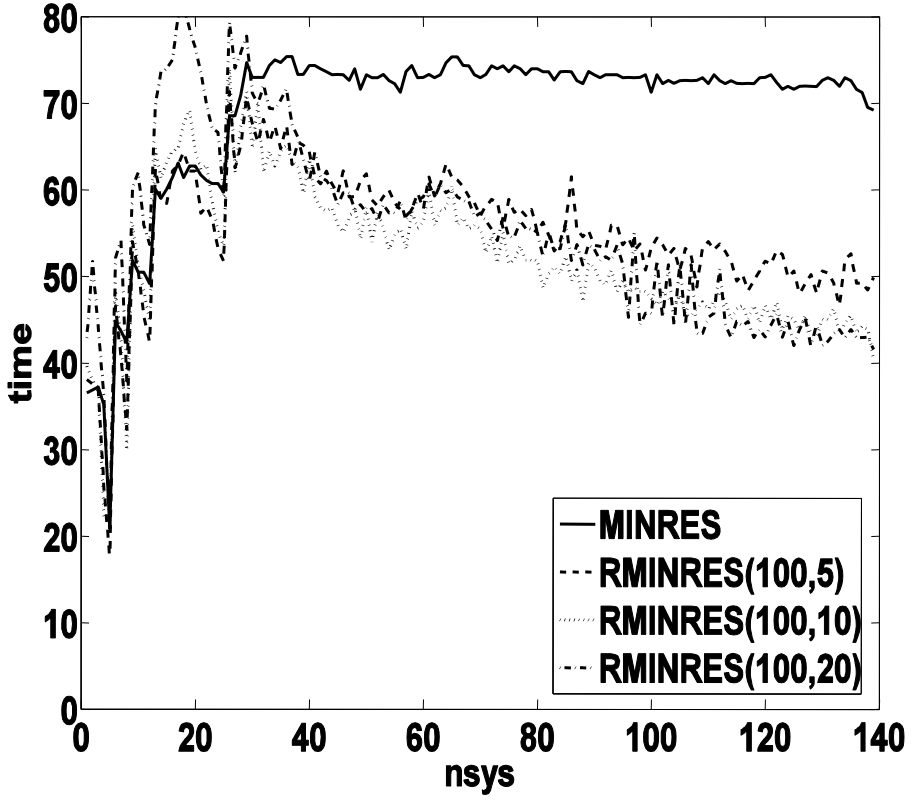

Convergence for Crack Propagation Problem

From Philippe Geubelle and Spandan Maiti UIUC - AE

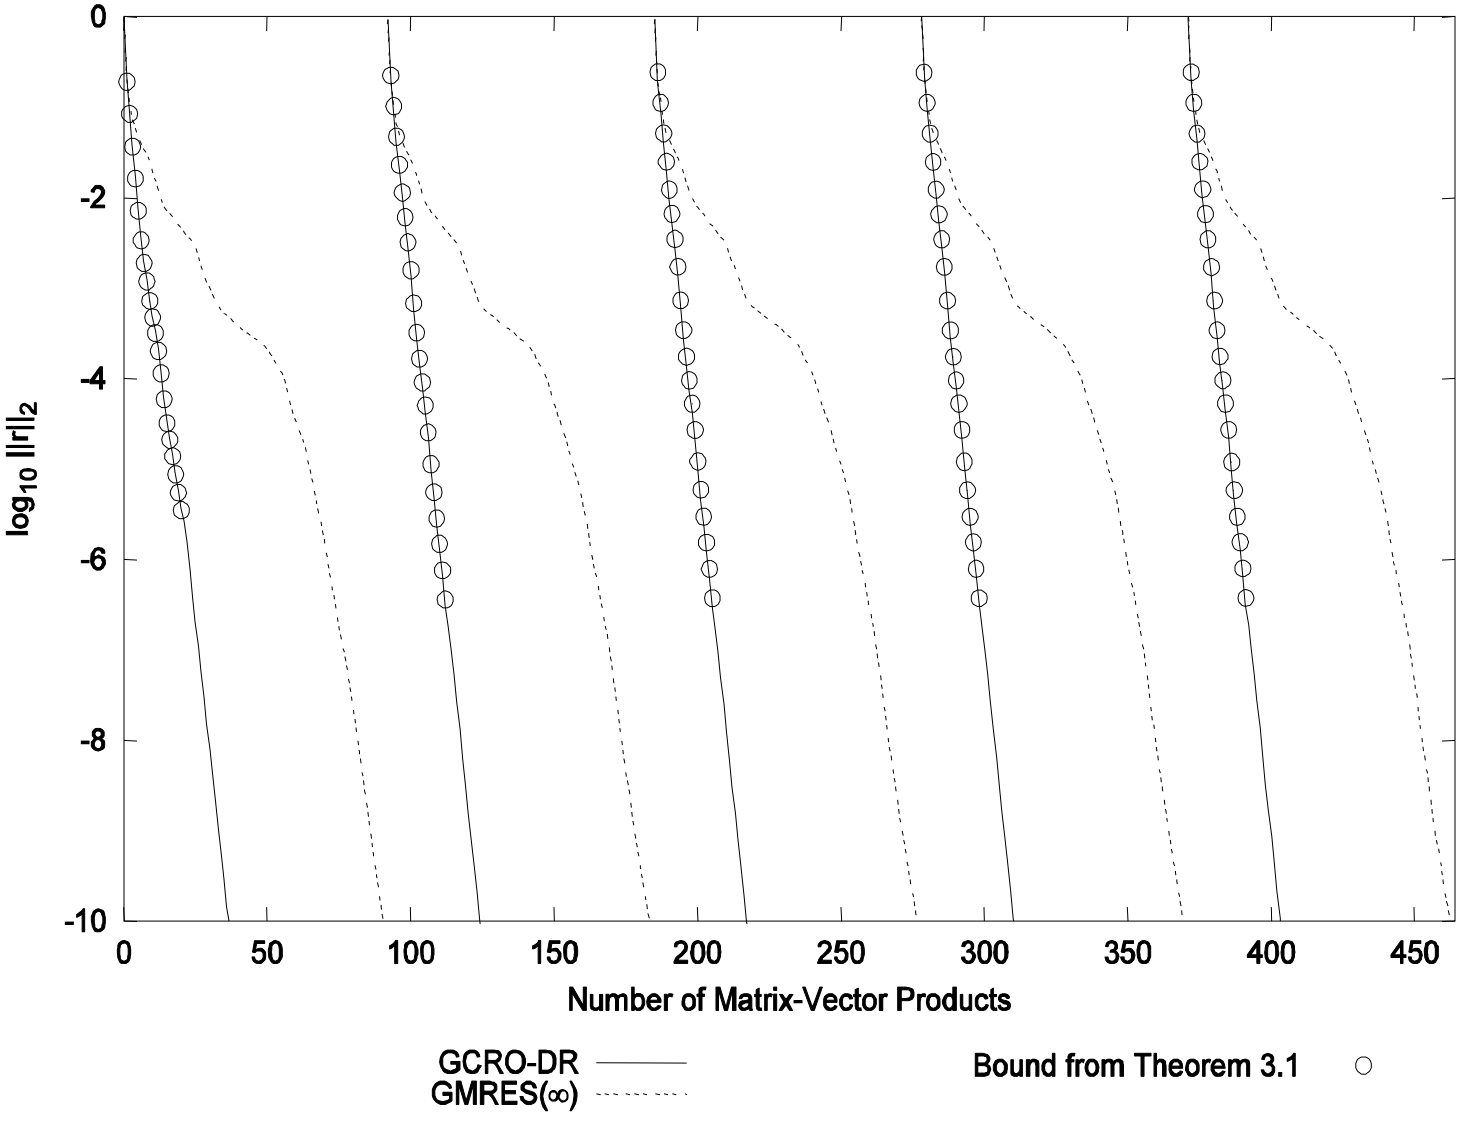

Convergence Recycling Approximate Invariant Subspace

Interested in convergence of $\left(I - CC^*\right) A \left(I - CC^*\right)$

CG for HPD: effective condition number (Erhel'00) and (Saad et al'00), but no bound for approx. invariant subspace

Saad'97: $\left\|r_m\right\| \leq \left\|q_m\left(A\right) r_d\right\| + \varepsilon \left\|q_m\left(A\right) P_U r_0\right\|$, with $q_m\left(A\right)$ optimal for r_d and second term small due to ε

Simoncini&Szyld'05 similar approach for convergence GMRES
Can adapt analysis for starting with approx. invar. subspace (Parks'05 and Parks et al'06)

For bound to prove fast convergence to very small residual, ε must be very small – approximate invariant subspace very accurate.
For recycling this would be annoying – we don't get very small ε.
We do get fast convergence …

Convergence for Hermitian Case

Solve $Ax = b$, A Hermitian: $A = \begin{bmatrix} Q & Y \end{bmatrix} \operatorname{diag}\left(\Lambda_Q, \Lambda_Y\right) \begin{bmatrix} Q & Y \end{bmatrix}^*$

Q_l is invariant subspace of dim l. C_k 'approximates' subspace Q_l where $k \geq l$, and $\delta = \left\| \left(I - \Pi_C\right) Q \right\|_2 < 1$ small. $\begin{bmatrix} C_k & W_{n-k} \end{bmatrix}$ unitary

$$m_Q = \min\left\{\Lambda_Q\right\},\ M_Q = \max\left\{\Lambda_Q\right\},\ m_Y = \min\left\{\Lambda_Y\right\},\ M_Y = \max\left\{\Lambda_Y\right\}$$

Theorem Eigenvalue Bounds:

$$m_Y + \delta^2 \min\left(0, m_Q - m_Y\right) \leq z^* W^* A\, W z \leq M_Y + \delta^2 \max\left(0, M_Q - M_Y\right)$$

This leads to well-known bound for improved condition number if deflated problem is definite (otherwise need some additional work)

Change in Bounds on Spectrum

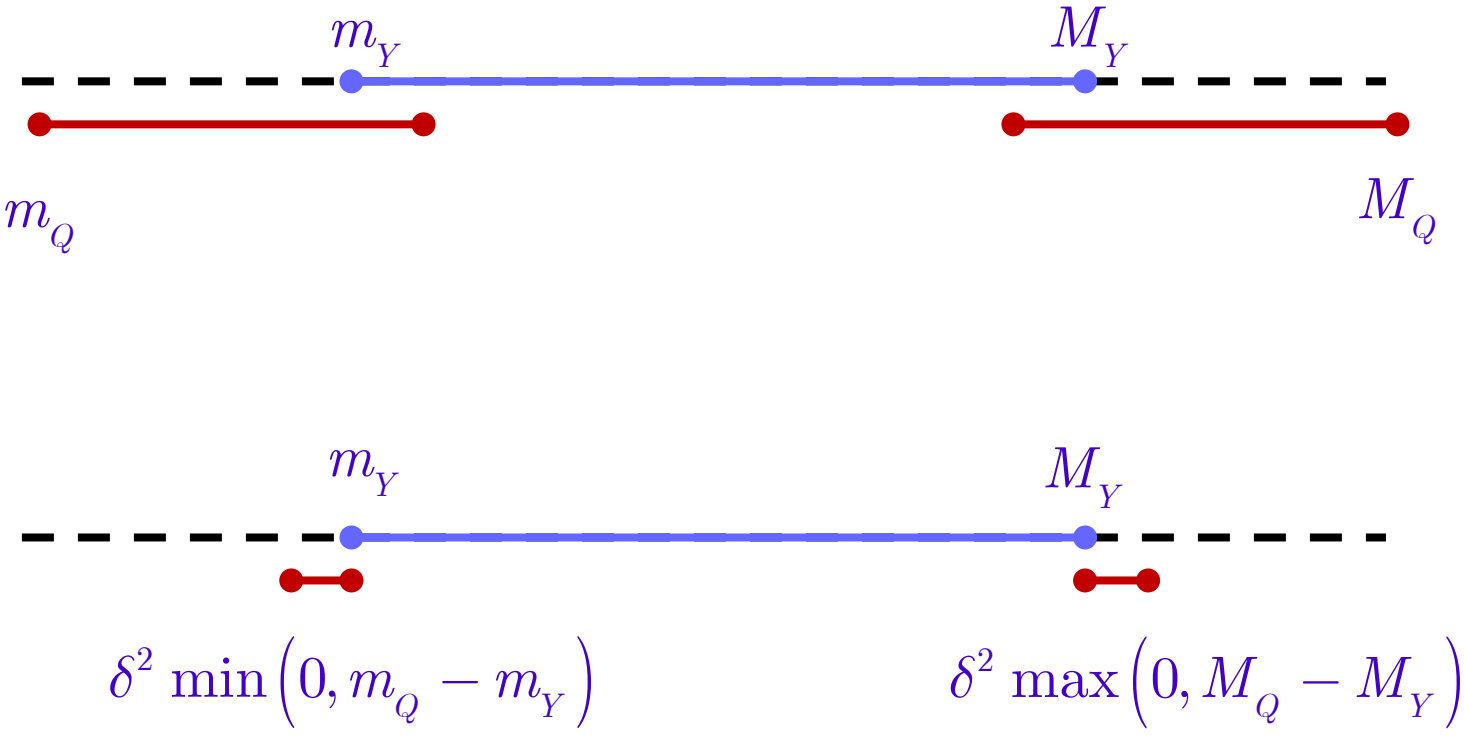

RMINRES for Preconditioned Laplace Equation

$\Delta u = 0$ - Dirichlet boundary conditions

202 x 202 regular grid, IC(0) preconditioner

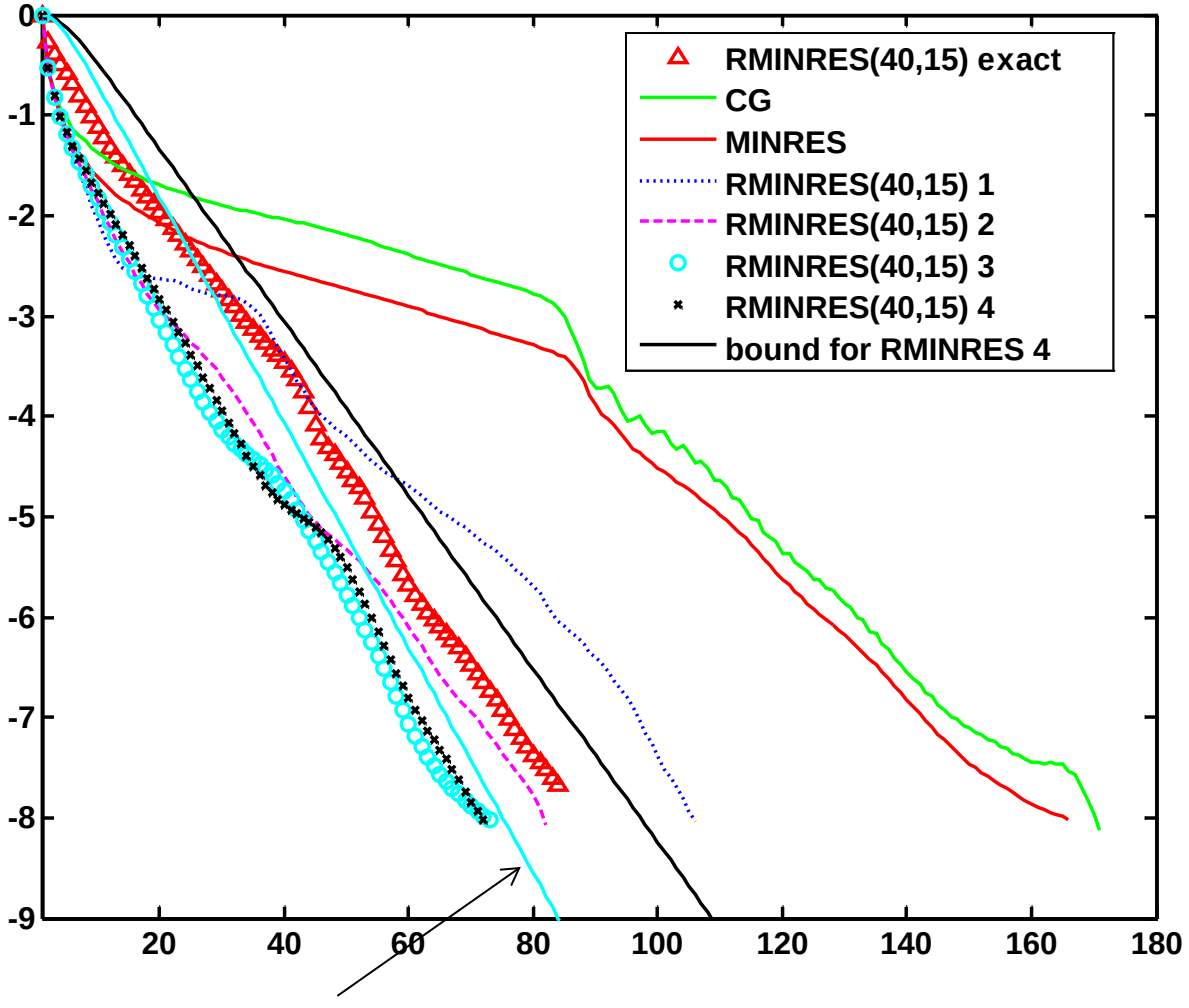

Bound from first 25 eigenvalues

RMINRES for Indefinite Problem

$A = \mathrm{diag}(-9, -7, -5, \ldots, 187, 189)$

After first solve, first 4 eigenvectors accurate, 5th eigenvector poor:

$1 - \delta^2 \left(m_Q - m_Y \right)$ positive but small

After second solve: $\delta \approx 0.08 \rightarrow \lambda_{\min} \geq 0.94$

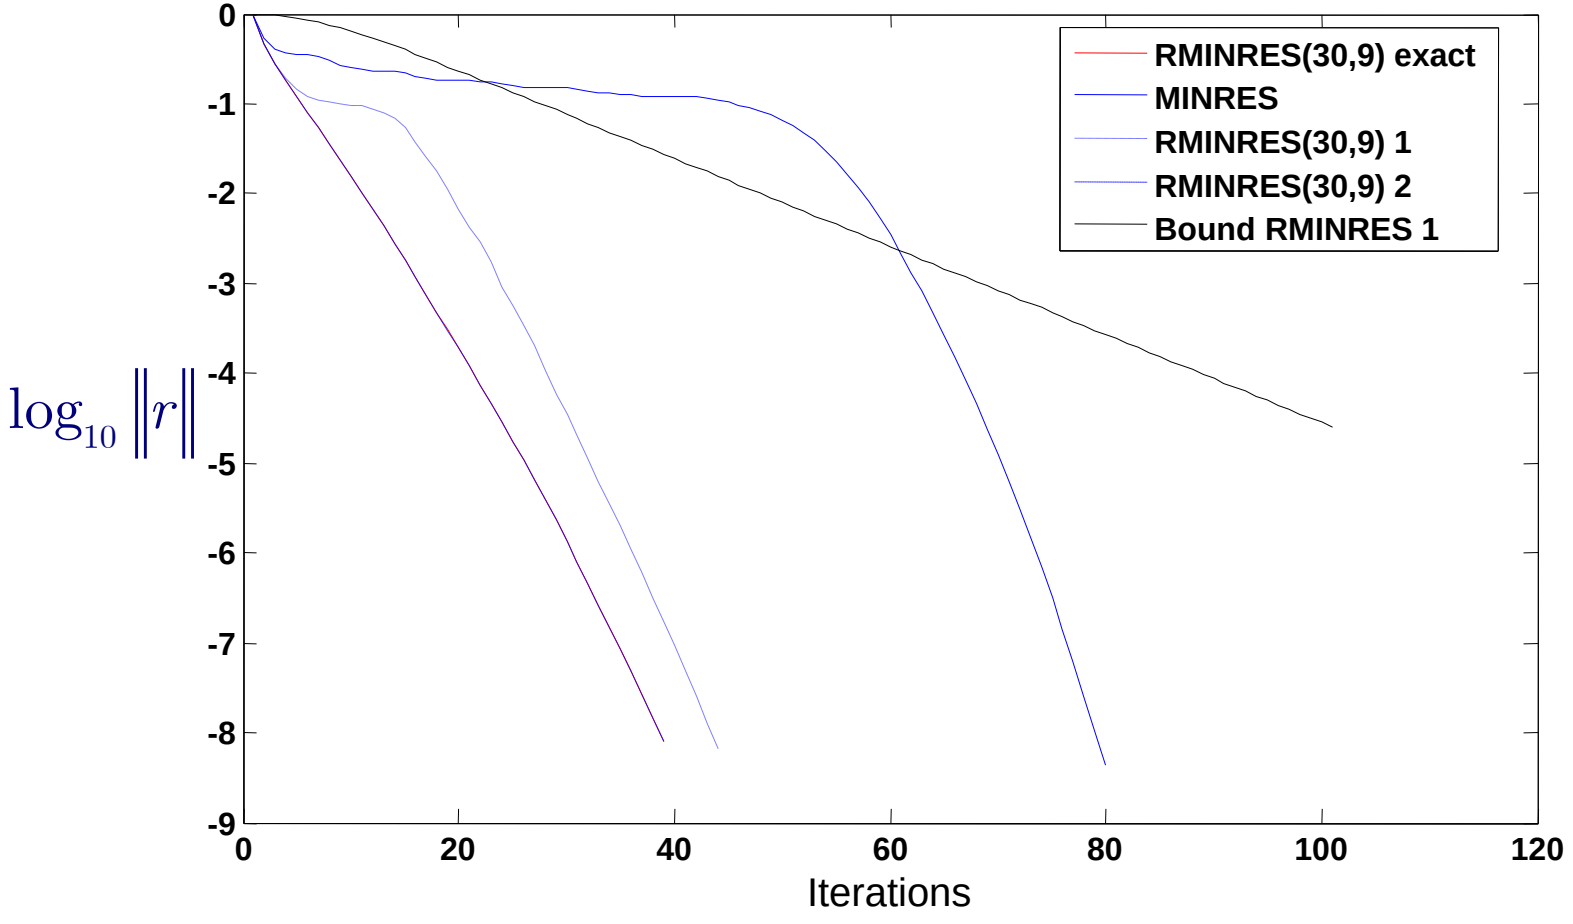

Acoustics Problems

with Jan Bierman (BMW)

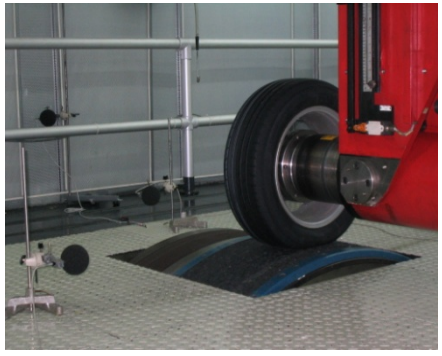

Real experiment

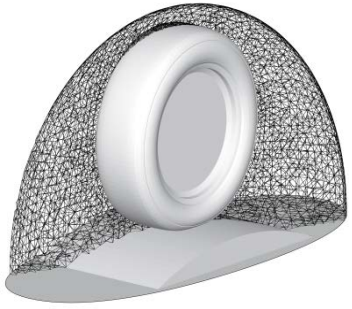

Part of the acoustic FE mesh

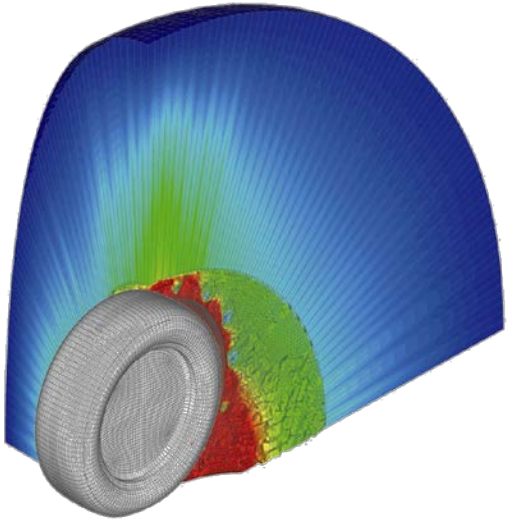

Acoustic FE/IFE mesh with solution

Details:

2nd order acoustic Finite Elements

6th order acoustic Infinite Elements

~10,000 degrees of freedom

about 150 frequencies to be evaluated

Discretization

Variational form and resulting matrix components:

$$\int_\Omega (\nabla p \cdot \nabla \bar{q} - k^2 p \bar{q})\,\mathrm{d}V - \int_{S_N} \mathrm{i}\rho\omega v_n^o p\bar{q}\,\mathrm{d}S - \int_{S_R} \mathrm{i}k\alpha p\bar{q}\,\mathrm{d}S = 0$$

$$K_{ij} = \int_{\Omega^e} \nabla N_i \nabla N_j \,\mathrm{d}V,$$

$$K_{ij} = \int_{\Omega^e} (\nabla D\Phi_i + \nabla\Phi_i D)\nabla\Phi_j \,\mathrm{d}V,$$

$$M_{ij} = 1/c^2 \int_{\Omega^e} N_i N_j \,\mathrm{d}V,$$

$$M_{ij} = 1/c^2 \int_{\Omega^e} (1 - (\nabla\mu\nabla\mu))\Phi_i\Phi_j D \,\mathrm{d}V,$$

$$C_{ij} = \rho \int_{S_R^e} \alpha N_i N_j \,\mathrm{d}S.$$

$$C_{ij} = 1/c \int_{\Omega^e} (\nabla\mu\nabla\Phi_j)D\Phi_i - (\nabla D\nabla\mu)\Phi_i\Phi_j - (\nabla\Phi_i\nabla\mu)D\Phi_j \,\mathrm{d}V, + 1/c \int_{S_R^e} \alpha\Phi_i\Phi_j D \,\mathrm{d}S.$$

$$f_i = -\mathrm{i}\omega \int_{S_N^e} \rho v_n^o N_i \,\mathrm{d}S.$$

Tire Rolling Noise Modeling

Equations for the interior and exterior acoustics simulation

$$A(\omega)\,p = \left(K + i\omega C - \omega^2 M\right) p = f(\omega)$$

Right hand side depends on excitation frequency (from road texture)

Problem needs to be solved for $\omega = 100, \ldots, 1500$ and $\Delta\omega = 10$
So we need to solve 140 large linear systems (for small model problem)

For full problem up to 500 frequencies

Matrix components from interior domain are symmetric

Matrix components from exterior domain are nonsymmetric

In general, the exterior domain component is not a low rank update

Acoustics – GCROT vs BiCGStab in # Matvecs

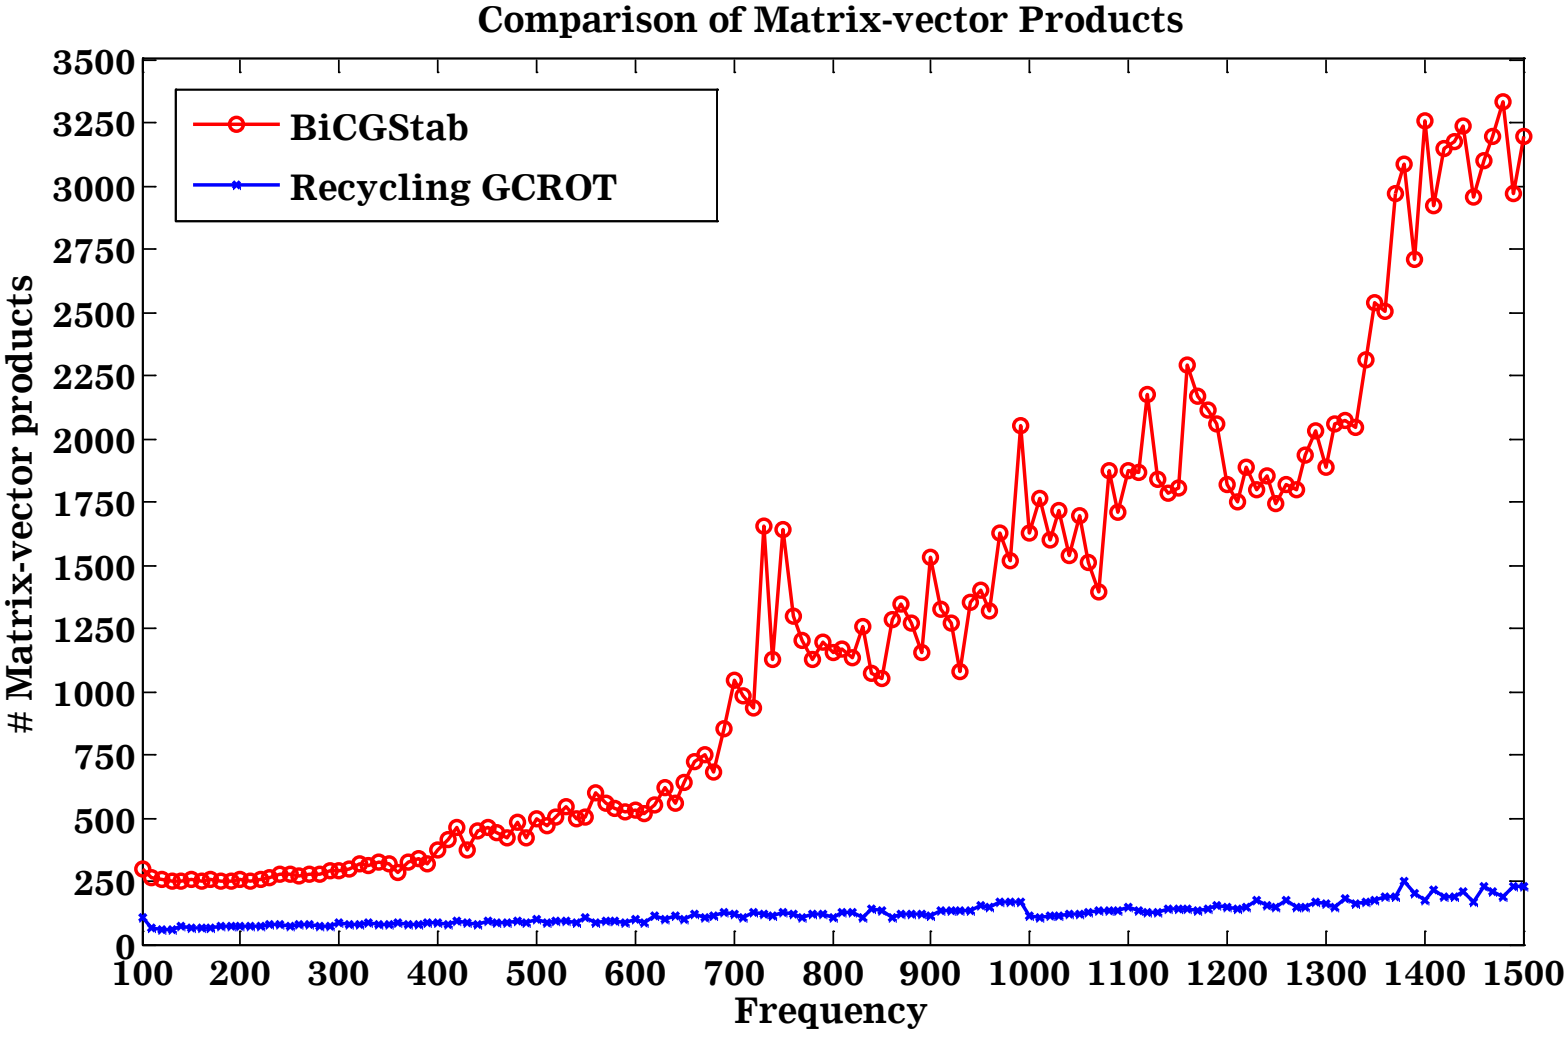

Real World Acoustics Problems

ACCELERATION OF ACOUSTICAL COMPUTATIONS BY

V SUBSE

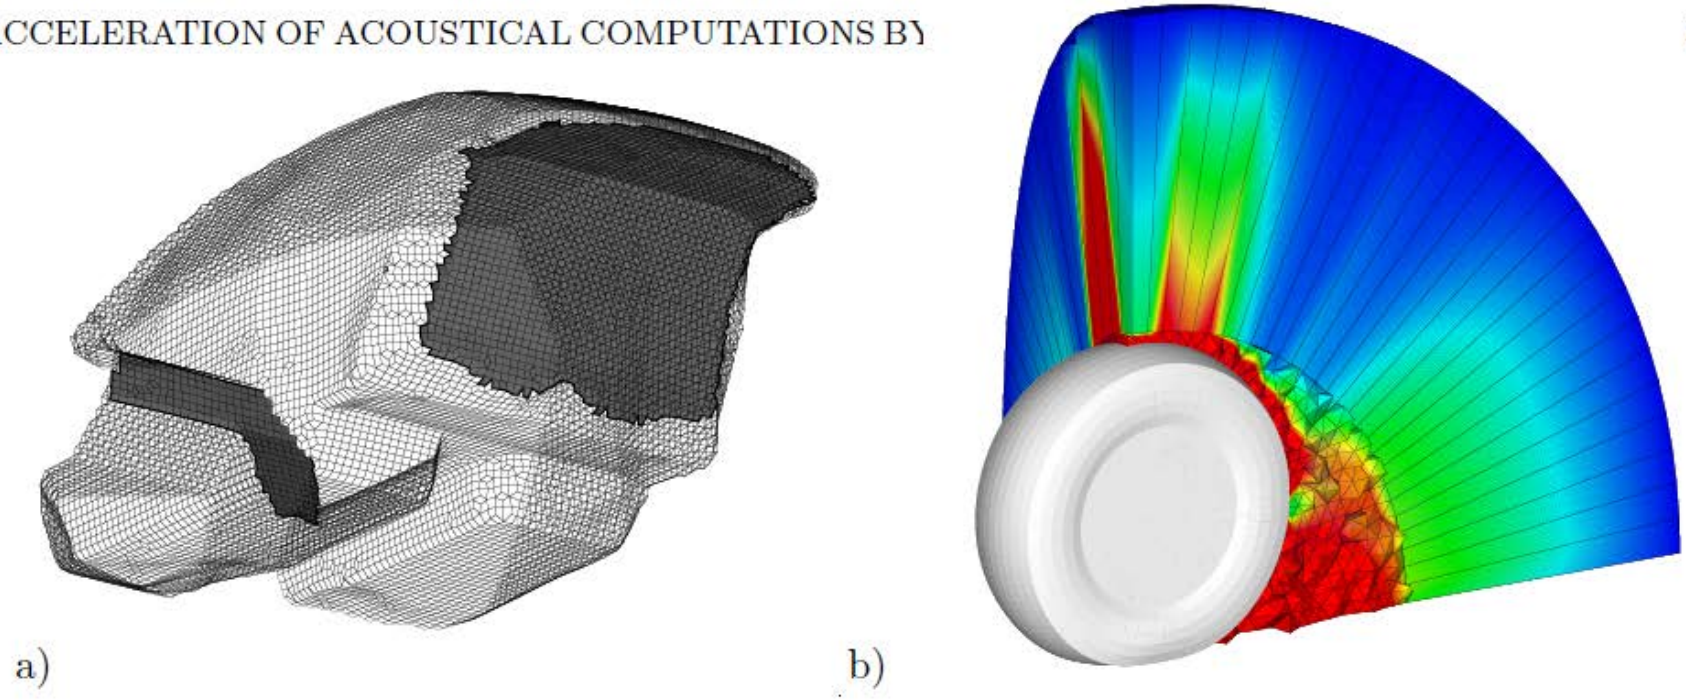

a) b)

Figure 30. FE model of a car interior acoustic cavity (a) and a quarter of a FE/IFE model for computing the sound radiation of tires (b).

2nd order acoust. finite elements
~350,000 degrees of freedom
20 frequencies

2nd order finite element
8th order infinite elements
~650,000 degrees of freedom
335 frequencies

Results Car Interior Acoustic Cavity

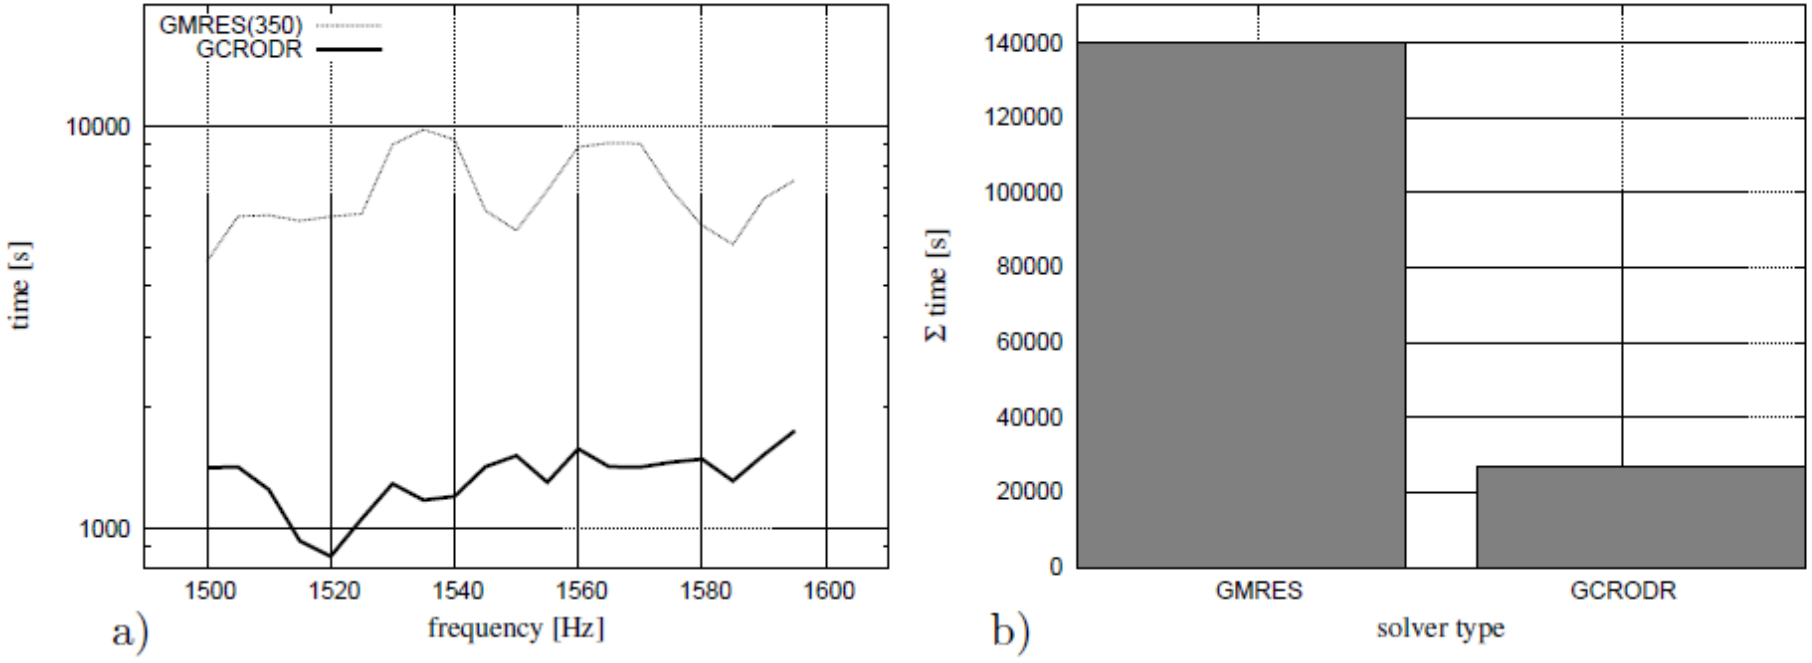

Figure 31. Comparison of the solution times of the standard solvers and using GCRODR for the car cavity example, a) frequency dependent solution times, and b) accumulated solution times.

BiCGStab did not converge for this problem

Results Sound Radiation of Tires

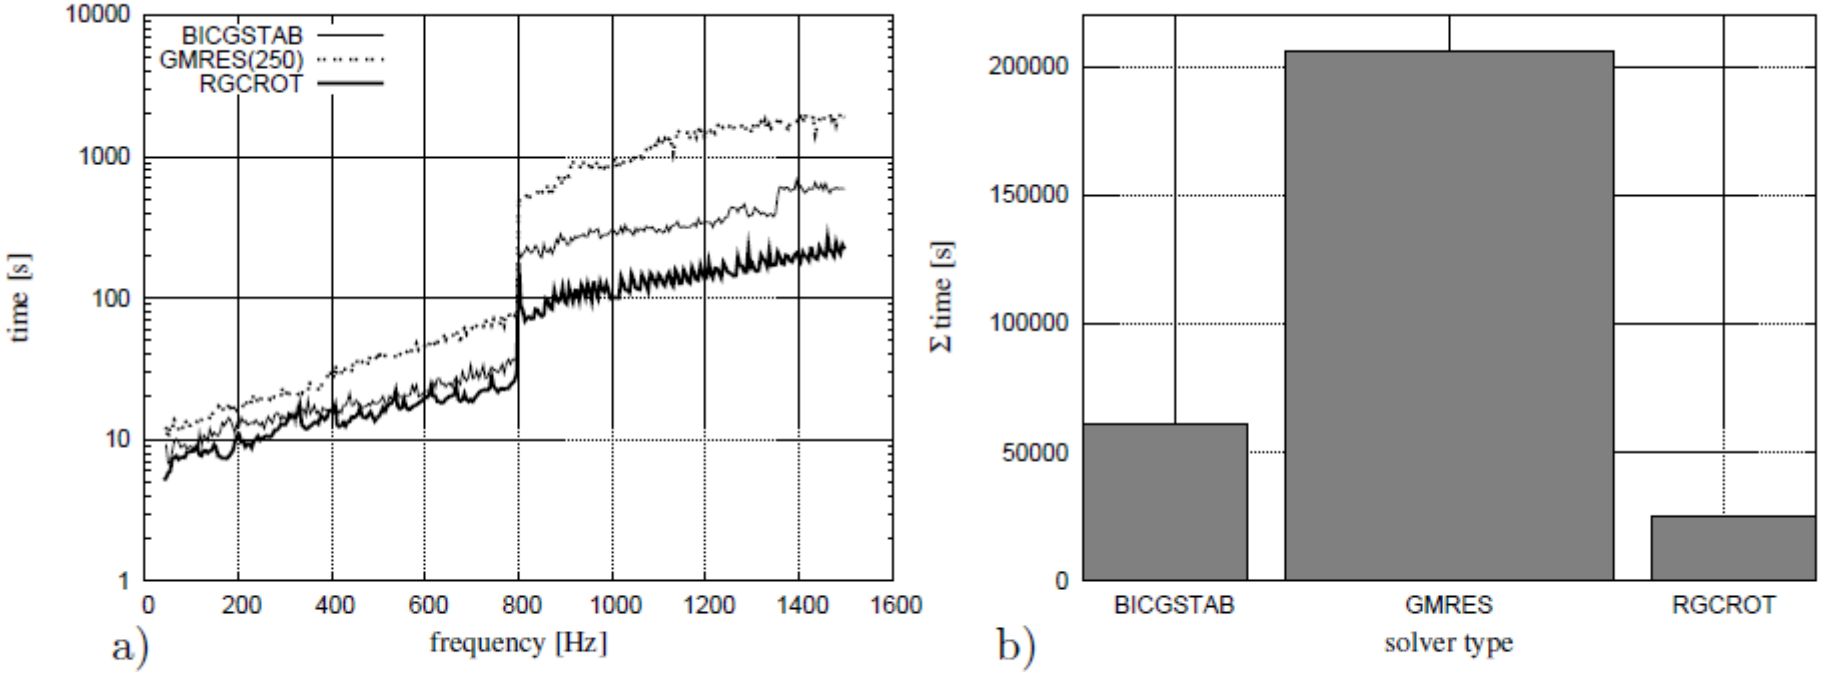

Figure 32. Comparison of the solution times of the standard solvers and using RGCROT for the tire radiation example, a) frequency dependent solution times, and b) accumulated solution times.

RGCROT vs BiCGStab

~ factor 10 cheaper in matrix-vector products

~ factor 2.5 faster in time (factor 10 better than GMRES(250))

Conclusions, Extra Info, and Future Work

- Recycling search spaces and preconditioners is very effective for range of applications
- Techniques for recycling subspaces are fairly cheap
- Effective convergence theory for recycling/projecting out approximate invariant subspace
  - Modest accuracy needed for fast convergence
- Best recycle space is open question (even Herm. case)
- Need regular updating to track changes in problem
- Recycling based directly on FOV or pseudospectra for highly nonnormal problems
- Recycling preconditioners can take many forms
- Software available (or soon): Trilinos at Sandia, or ask me

Good reading

- Recycling Krylov Subspaces for Sequences of Linear Systems, Parks, de Sturler, Mackey, Johnson, Maiti, SIAM SISC 28(5), 2006
- Recycling Subspace Information for Diffuse Optical Tomography, Kilmer, de Sturler, SIAM SISC 27(6), 2006
- Large-scale topology optimization using preconditioned Krylov subspace methods with recycling, Wang, de Sturler, Paulino, IJNME 69, 2007
- Multilevel Sparse Approximate Inverse Preconditioners for Adaptive Mesh Refinement, Wang, de Sturler, Lin. Alg. Appl. vol. 431, 2009
- Recycling Krylov Subspaces and Preconditioners, PhD Thesis, Ahuja
- Improved Scaling for Quantum Monte Carlo on Insulators, Ahuja, Clark, de Sturler, Ceperley, and Kim, SIAM SISC 33(4), 2011
- Krylov subspace methods for topology optimization and adaptive meshes, PhD Thesis, Shun Wang, UIUC/CS 2007
- Topology Optimization with Adaptive Mesh Refinement, de Sturler, Paulino, Wang, Proc. 6th Intl Conf on Comput. of Shell and Spatial Structures IASS-IACM 2008
- www.math.vt.edu/people/sturler/index.html