A Comparison of Preconditioned Krylov Subspace Methods for Nonsymmetric Linear Systems

Aditi Ghai, Cao Lu, Xiangmin Jiao

Department of Applied Mathematics & Statistics, Stony Brook University

1. Introduction: Motivation and Goal

Krylov subspace methods are widely used in solving large sparse linear systems from PDE discretizations. For symmetric systems, CG and MINRES are typically the best. For nonsymmetric systems, which often arise in practice from PDE discretizations, boundary/jump conditions, irregular meshes, no KSP method is apparently optimal. The goal of this work is to perform a systematic comparison and in turn establish some practical guidelines in choosing the best combinations of the pre- conditioned KSP solvers. We consider four KSP methods, restarted GMRES, TFQMR, BiCGSTAB, QMRCGSTAB, coupled with three preconditioners, Gauss-Seidel, incomplete LU factorization (ILU), algebraic multigrid (AMG).

2. Comparison Setup

For 3D tests, we generated three unstructured meshes, using PDE discretizations (FEM, AES-FEM, GFD) of a cube at different resolutions using TetGen, to facilitate the scalability study of the preconditioned KSP methods with respect to the number of unknowns. For the finite difference method, we consider a matrix obtained from an unequally spaced structured mesh for the Helmholtz equation with Neumann boundary conditions, so the matrix has a very large condition number.

Results were obtained using the high-performance LI-RED computing system at the Institute for Advanced Computational Science of Stony Brook University.

3. Comparison of KSP Methods

Table 1: Comparisons of KSP methods based on Krylov subspaces.

Method	Iteration	Matrix-Vector Prod.		Recurrence
		A^T	A	
GMRES	Arnoldi	0	1	k
BiCG	bi-Lanczos	1	1	3
QMR				
CGS	transpose-free	0	2	
TFQMR	bi-Lanczos 1			
BiCGSTAB	transpose-free			
QMRCGSTAB	bi-Lanczos 2			

Table 2: Comparison of operations per iteration and memory requirements of various KSP methods. n denotes the number of rows, ℓ the average number of non-zeros per row, and k the iteration count.

Method	Min.	Mat-vec Prod.	axpy	Inner Prod.	FLOPs	Stored vectors
GMRES	$\|\boldsymbol{r}_k\|$	1	$k+1$	$k+1$	$2n(\ell+2k+2)$	$k+5$
BiCGSTAB	$\|\boldsymbol{r}_k(\omega_k)\|$	2	6	4	$4n(\ell+5)$	10
TFQMR	$\|\boldsymbol{r}_k\|_{\boldsymbol{W}_{k+1}^T}$		10	4	$4n(\ell+7)$	8
QMRCGSTAB	$\|\boldsymbol{r}_k(\omega_k)\|$ & $\|\boldsymbol{r}_k\|_{\boldsymbol{W}_{k+1}^T}$		8	6		13

4. Convergence Results

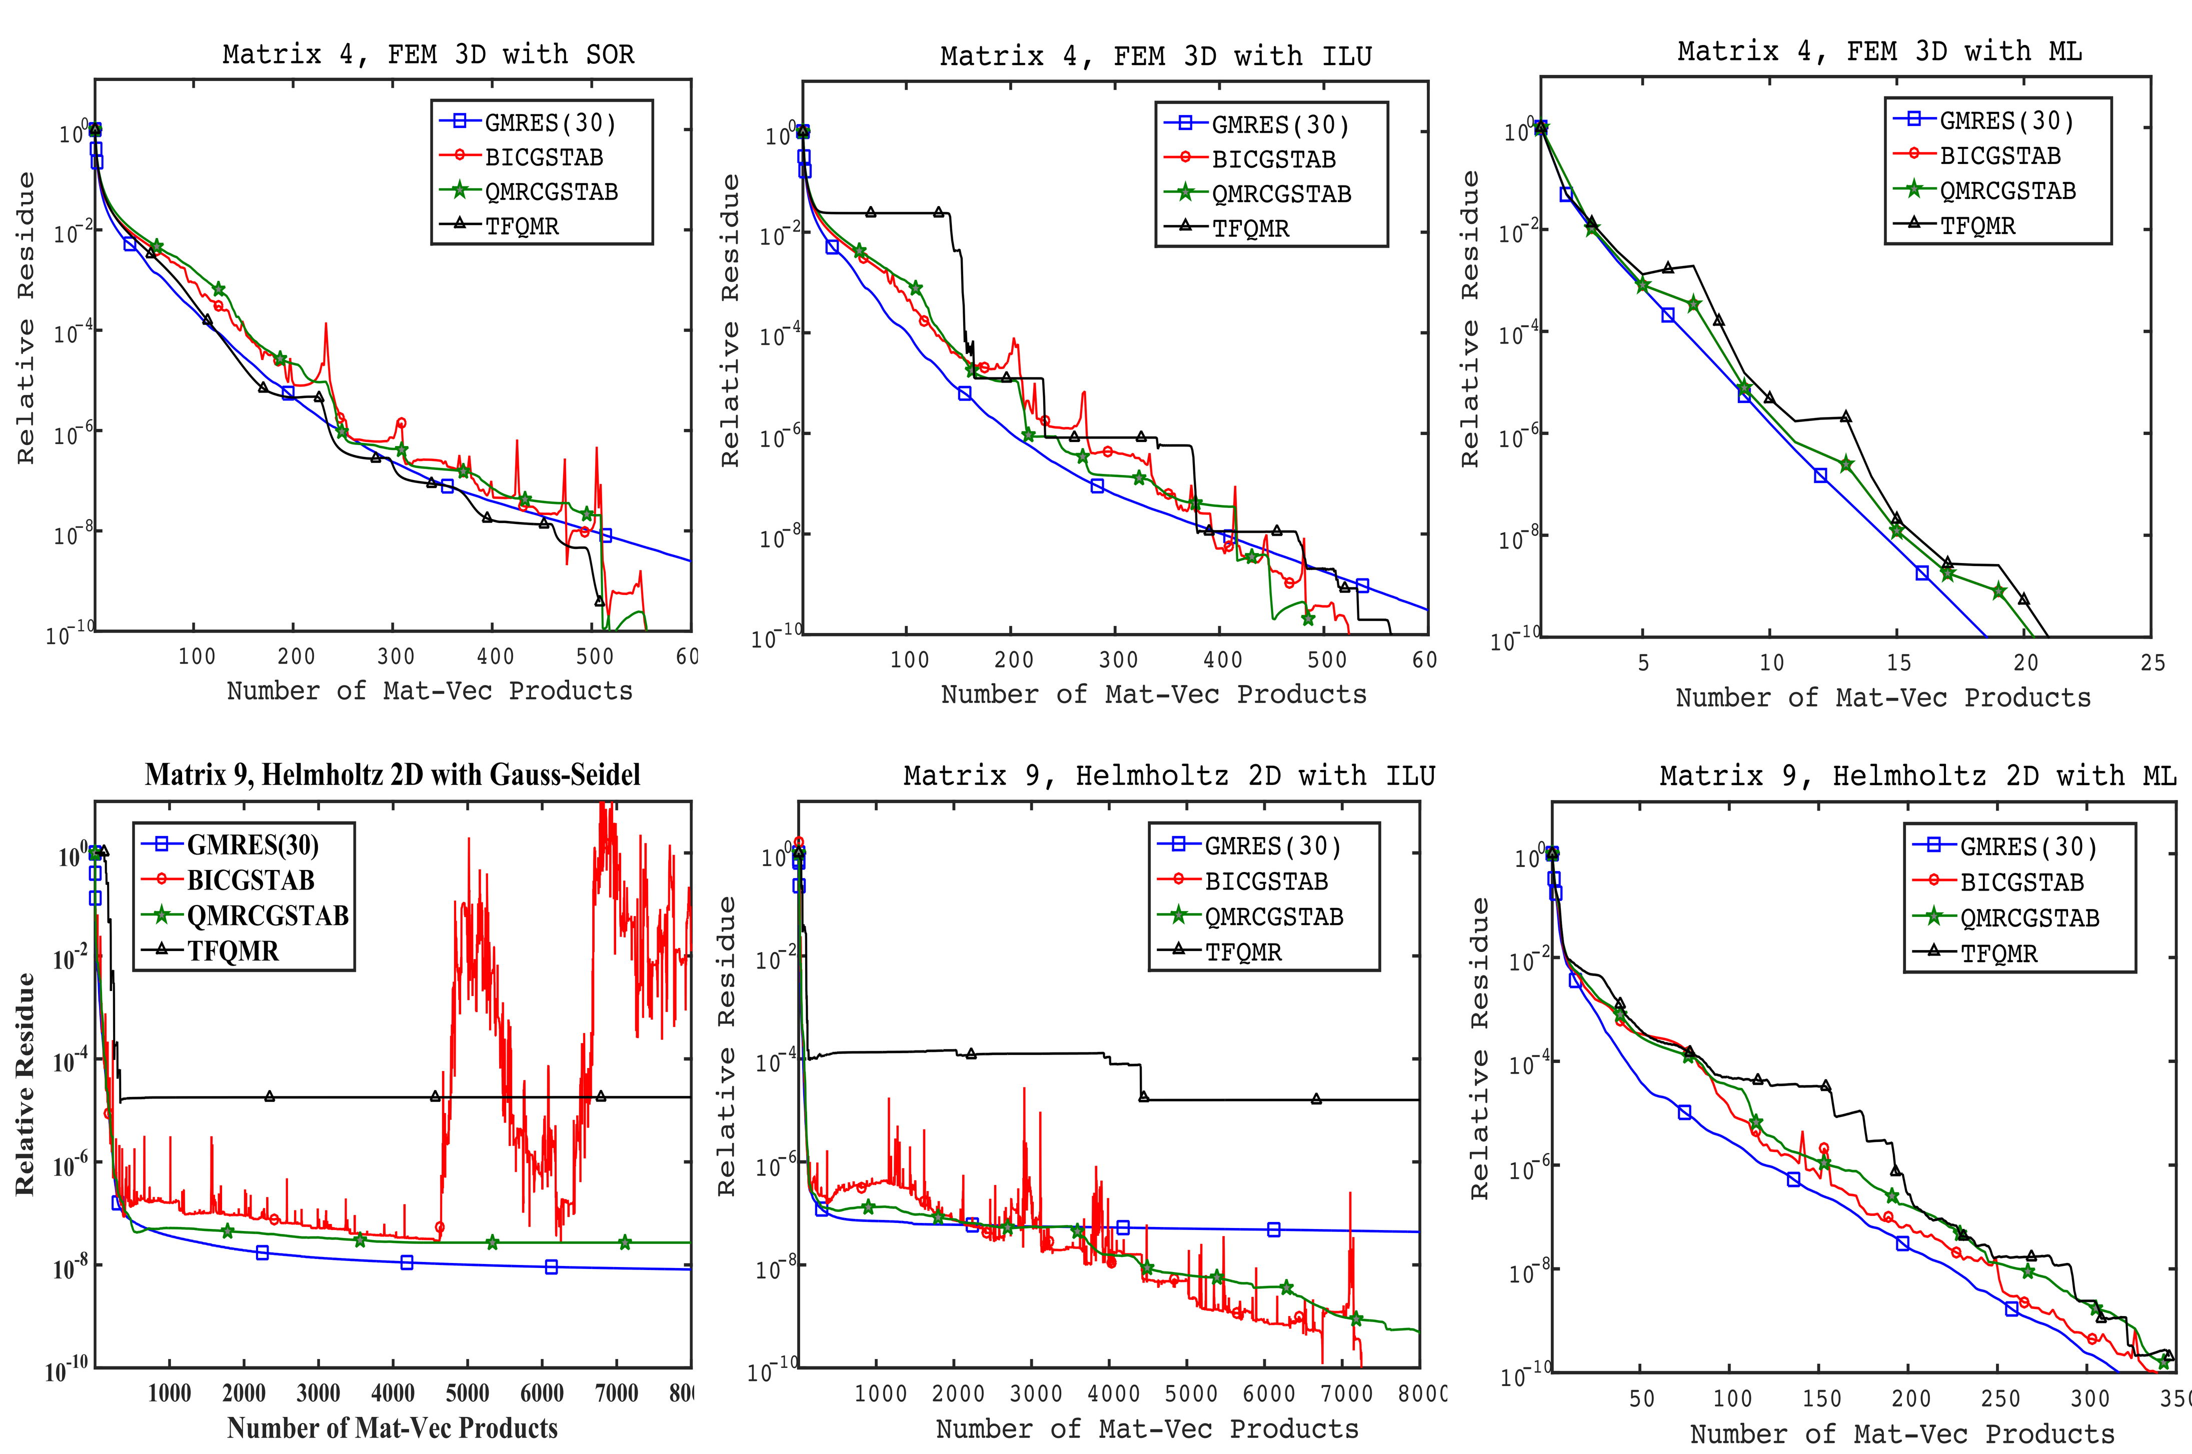

Figure 1: Relative residual versus iteration count for Gauss-Seidel, ILU and ML preconditioners for FEM 3D (197,881,373 unknowns) and Helmholtz Equation 2D (6,694,058 unknowns).

5. Timing and Scalability Results

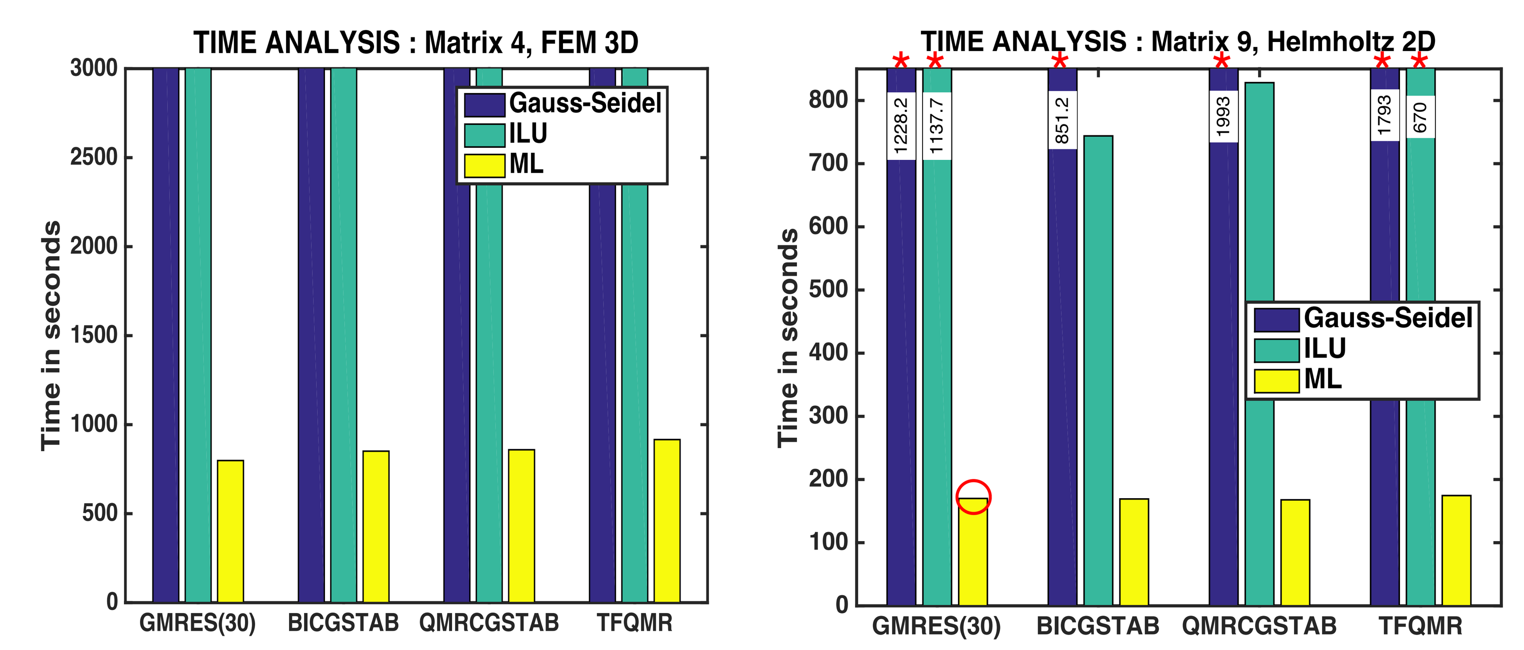

Figure 2: Timing results for FEM3D on the left and Helmholtz 2D on the right, encircled bars indicate the fastest solver-preconditioner combination in timing results; star (*) indicates stagnation.

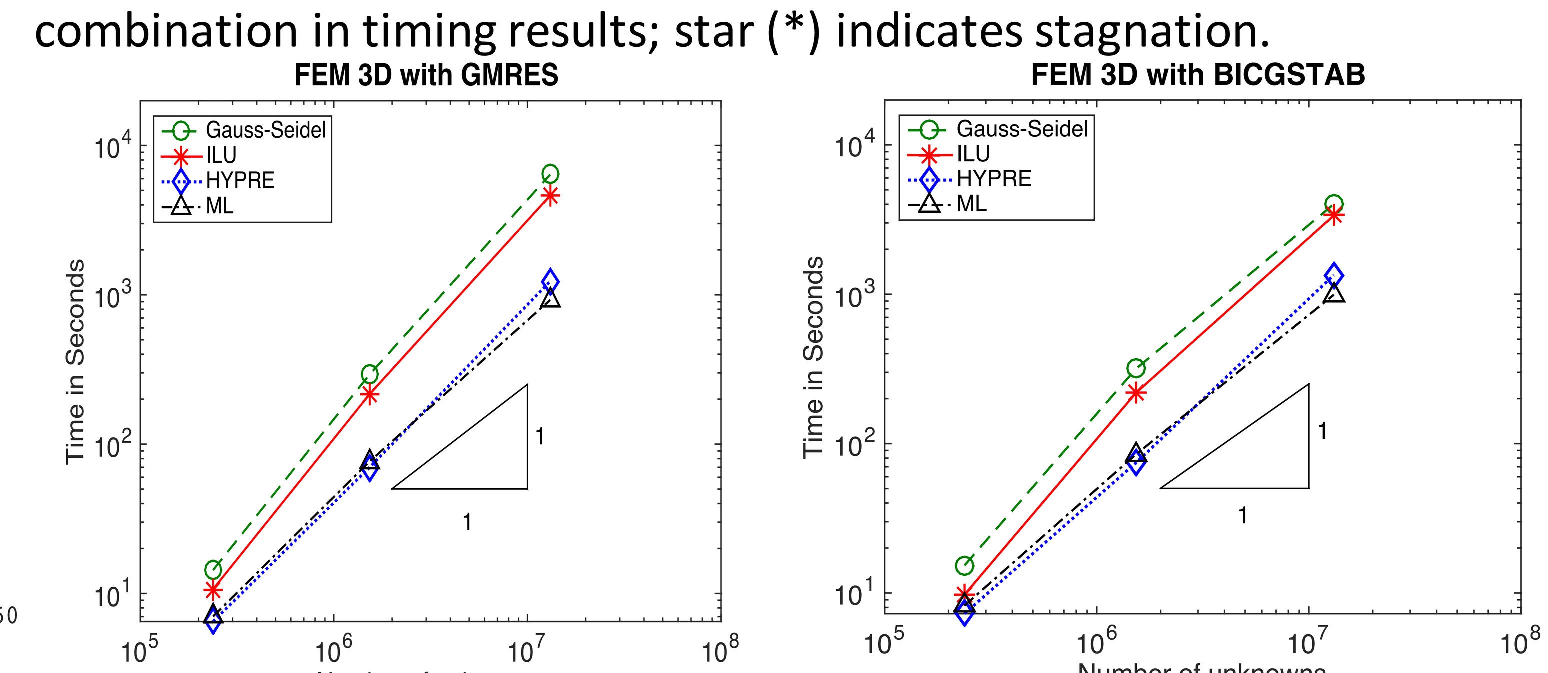

Figure 3: Scalability result of the preconditioned solvers in terms of number of unknowns.

6. HYPRE Versus ML Preconditioner

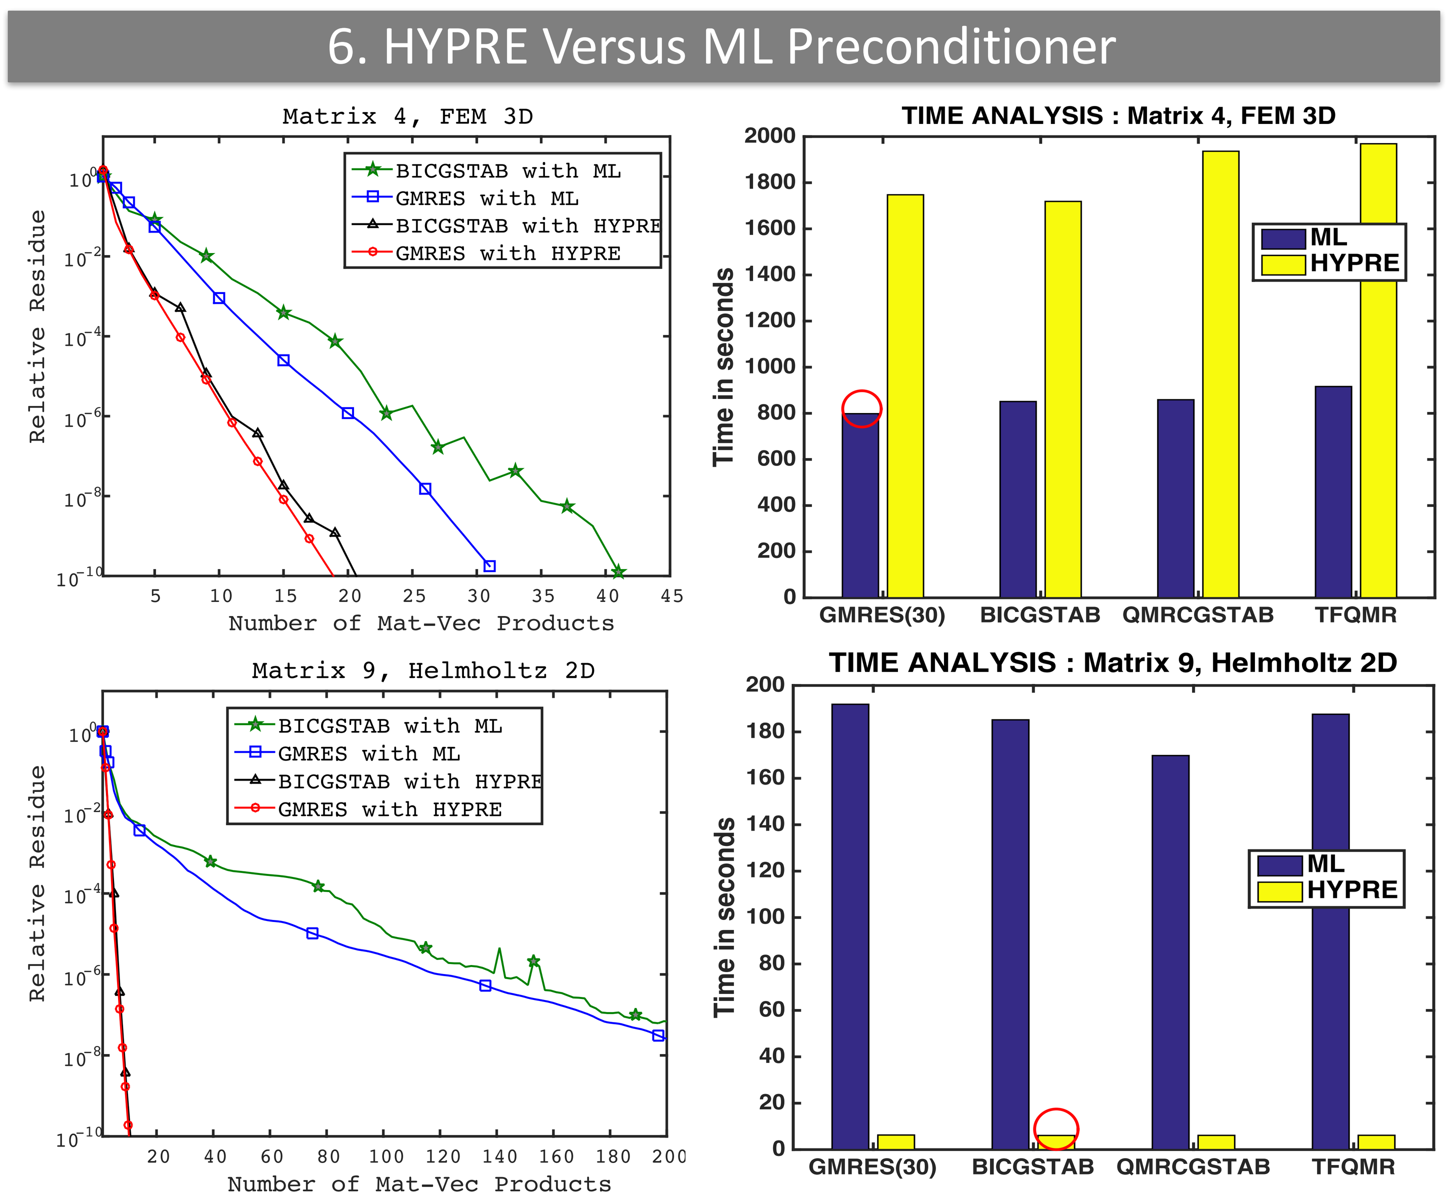

7. Conclusions and Future Work

- For a very large, reasonably well-conditioned linear system, use GMRES with smoothed-aggregation AMG as right preconditioner. If AMG is unavailable and the problem size is moderate, BiCGSTAB with ILU as right preconditioner is a reasonable choice.
- For ML or Hypre, the scalability for the four KSP methods is nearly linear, whereas Gauss-Seidel and ILU are less scalable. Therefore, the performance advantage of multigrid preconditioners would become even larger as the problem size increases.
- Hypre performs better than ML for ill-conditioned systems, indicating no clear winner. These results also indicate that more research into multigrid preconditioners are needed.

8. References

Ghai A, Lu C, Jiao X. Comparison of Preconditioned Krylov Subspace Methods for Nonsymmetric Linear Systems, 2016. Preprint: http://arxiv.org/abs/1607.00351

9. Acknowledgements

The second and third authors were partially supported by DoD-ARO under contract #W911NF0910306 and also in part by a subcontract to Stony Brook University from Argonne National Laboratory under Contract DE-AC02-06CH11357 for the SciDAC program of the U.S. Department of Energy. The second author was also supported by NSF under grant #ags1418309.