

Purdue University

Purdue e-Pubs

Department of Computer Science Technical
Reports

Department of Computer Science

2007

Reliability of Krylov Subspace Methods - A Practical Perspective II

M. Naumov

M. Manguoglu

C. C. Mikkelsen

Ahmed Sameh

Purdue University - Main Campus, sameh@cs.purdue.edu

Report Number:

07-022

Naumov, M.; Manguoglu, M.; Mikkelsen, C. C.; and Sameh, Ahmed, "Reliability of Krylov Subspace Methods - A Practical Perspective II" (2007). *Department of Computer Science Technical Reports*. Paper 1686.

<https://docs.lib.purdue.edu/cstech/1686>

This document has been made available through Purdue e-Pubs, a service of the Purdue University Libraries.
Please contact epubs@purdue.edu for additional information.

**RELIABILITY OF KRYLOV SUBSPACE METHODS
A PRACTICAL PERSPECTIVE II**

**M. Naumov
M. Manguoglu
C.C. Mikkelsen
A. Sameh**

**Department of Computer Science
Purdue University
West Lafayette, IN 47907**

**CSD TR #07-022
October 2007**

Reliability of Krylov Subspace Methods

A Practical Perspective II

M. Naumov, M. Manguoglu, C. C. Mikkelsen
and A. Sameh *

October 7, 2007

Abstract

In this report authors continued to study the reliability of Krylov Subspace iterative methods. We analyzed the performance of Conjugate Gradient (CG) [1], Generalized Minimum Residual (GMRES) [2] and Stabilized Bi-Conjugate Gradient (BiCGStab) [3] algorithms, with an Improved Incomplete LU with fill-in tolerance (IILUT) preconditioner proposed by M. Benzi, J. C. Haws and M. Tuma in [5]. The comparisons with regular ILUT preconditioner studied in [8] are shown. The 65 non-symmetric test matrices were selected from the largest sparse matrices of the Tim Davis Matrix Collection. The Krylov Subspace methods still failed to produce a solution to the desired accuracy in fixed number of iterations in more than 66% of the cases. The reader should be advised that this report does not have theoretical bounds on convergence and is based purely on numerical experiments.

*Computer Science Department, Purdue University - West Lafayette
305 N. University Street, West Lafayette, IN, 47907-2107
(naumov@purdue.edu, mmanguog@cs.purdue.edu, cmikkels@cs.purdue.edu, and
sameh@cs.purdue.edu).

1 Introduction

Krylov Subspace iterative methods are among the most popular iterative methods used to solve large sparse linear systems. Although, these methods are known to be unreliable, most of the scientific community has embraced them. The simple experiments conducted in this study suggest that a search for robust alternatives is imperative and is especially important for nonsymmetric systems.

Letting $A \in \mathbb{R}^{n \times n}$ and $\mathbf{f} \in \mathbb{R}^n$ we are interested in solving

$$A\mathbf{x} = \mathbf{f} \quad (1)$$

using GMRES, BiCGStab and

$$A^T A\mathbf{x} = A^T \mathbf{f} \quad (2)$$

using CG (CGNR). In our experiment we let $\mathbf{f}^T = [1 \dots 1]$. The software implementation of Krylov Subspace Methods from Harwell Subroutine Library (HSL) and ILUT factorization from SPARSEKIT have been used in our experiments. The matrix preprocessing to reduce the chance of zero pivot and fill-in in ILUT was done as suggested in [5]. First, using a reorder and scale routine (RS) to maximize the product of diagonal entries and scale the matrix, so that these diagonal entries have absolute value of one. Second, applying the reverse CutHill-McKee reordering (RCM) [6, 7] to make the matrix more narrow banded (see Tab. 1,2). The square nonsymmetric matrices in this study were obtained from Tim Davis Matrix Collection and are listed below (see Tab. 3,4).

Table 1: Subroutines

Algorithm	Subroutine	Stopping Criteria		
		# it.	rel. residual	time (matrix #, time allowed)
CG	MI21 (HSL)	250	$\leq 10^{-4}$	1-23 $\leq$ 2h, 24-56 $\leq$ 3h, 57-65 $\leq$ 5h
GMRES	MI24 (HSL)	250	$\leq 10^{-4}$	1-23 $\leq$ 2h, 24-56 $\leq$ 3h, 57-65 $\leq$ 5h
BiCGStab	MI26 (HSL)	250	$\leq 10^{-4}$	1-23 $\leq$ 2h, 24-56 $\leq$ 3h, 57-65 $\leq$ 5h

Table 2: Preconditioning

Algorithm	Subroutine	Dropping Tolerance	Max. Storage
RS	MC64 (HSL)	not applicable	not applicable
RCM	genrcm (J. Burkardt)	not applicable	not applicable
ILUT	ilut (SPARSEKIT)	elements $\leq 10^{-3} \ A\ _\infty$ fill-in $\leq 0.1(\# \text{ nonzeros})$	10($\# \text{ nonzeros}$)

Table 3: Matrices

#	Matrix	Size	# Nonzeros	Type	Application
1	Raju/laminar_duct3D	67,173	3,788,857	Real	Physical Processes
2	Hamm/bcircuit	68,902	375,558	Real	Circuit Simulation
3	Mallya/lhr71	70,304	1,494,006	Real	Chemical Processes
4	Mallya/lhr71c	70,304	1,528,092	Real	Chemical Processes
5	Shyy/shyy161	76,480	329,762	Real	Comput. Fluid Dyn.
6	Bomhof/circuit_4	80,209	307,604	Real	Circuit Simulation
7	Averous/epb3	84,617	463,625	Real	Physical Processes
8	FEMLAB/poisson3Db	85,623	2,374,949	Real	Physical Processes
9	Rajat/rajat20	86,916	604,299	Real	Circuit Simulation
10	Rajat/rajat25	87,190	606,489	Real	Circuit Simulation
11	Rajat/rajat28	87,190	606,489	Real	Circuit Simulation
12	Rajat/rajat16	94,294	476,766	Real	Circuit Simulation
13	Rajat/rajat17	94,294	479,246	Real	Circuit Simulation
14	Rajat/rajat18	94,294	479,151	Real	Circuit Simulation
15	Sandia/ASIC_100ks	99,190	578,890	Real	Circuit Simulation
16	Sandia/ASIC_100k	99,340	940,621	Real	Circuit Simulation
17	Schenk_IBMSDS/matrix_9	103,430	1,205,518	Real	Semicon. Dev. Sim.
18	Hamm/hcircuit	105,676	513,072	Real	Circuit Simulation
19	Norris/lung2	109,460	492,564	Real	Biological Proces.
20	Rajat/rajat23	110,355	555,441	Real	Circuit Simulation
21	Schenk_ISEI/barrier2-1	113,076	2,129,496	Real	Semicon. Dev. Sim.
22	Schenk_ISEI/barrier2-2	113,076	2,129,496	Real	Semicon. Dev. Sim.
23	Schenk_ISEI/barrier2-3	113,076	2,129,496	Real	Semicon. Dev. Sim.
24	Schenk_ISEI/barrier2-4	113,076	2,129,496	Real	Semicon. Dev. Sim.
25	Schenk_ISEI/barrier2-9	115,625	2,158,759	Real	Semicon. Dev. Sim.
26	Schenk_ISEI/barrier2-10	115,625	2,158,759	Real	Semicon. Dev. Sim.
27	Schenk_ISEI/barrier2-11	115,625	2,158,759	Real	Semicon. Dev. Sim.
28	Schenk_ISEI/barrier2-12	115,625	2,158,759	Real	Semicon. Dev. Sim.
29	Norris/torso2	115,967	1,033,473	Real	Biological Proces.
30	Norris/torso1	116,158	8,516,500	Real	Biological Proces.
31	IBM_EDA/dc1	116,835	766,396	Real	Circuit Simulation
32	IBM_EDA/dc2	116,835	766,396	Real	Circuit Simulation
33	IBM_EDA/dc3	116,835	766,396	Real	Circuit Simulation
34	IBM_EDA/trans4	116,835	749,800	Real	Circuit Simulation
35	IBM_EDA/trans5	116,835	749,800	Real	Circuit Simulation
36	ATandT/twotone	120,750	1,206,265	Real	Circuit Simulation
37	Schenk_IBMSDS/matrix-new_3	125,329	893,984	Real	Semicon. Dev. Sim.
38	vanHeukelum/cage12	130,228	2,032,536	Real	Biological Proces.

Table 4: Matrices

#	Matrix	Size	# Nonzeros	Type	Application
39	Schenk.ISEI/para-4	153,226	2,930,882	Real	Semicon. Dev. Sim.
40	Schenk.ISEI/para-5	155,924	2,094,873	Real	Semicon. Dev. Sim.
41	Schenk.ISEI/para-6	155,924	2,094,873	Real	Semicon. Dev. Sim.
42	Schenk.ISEI/para-7	155,924	2,094,873	Real	Semicon. Dev. Sim.
43	Schenk.ISEI/para-8	155,924	2,094,873	Real	Semicon. Dev. Sim.
44	Schenk.ISEI/para-9	155,924	2,094,873	Real	Semicon. Dev. Sim.
45	Schenk.ISEI/para-10	155,924	2,094,873	Real	Semicon. Dev. Sim.
46	Ronis/xenon2	157,464	3,866,688	Real	Physical Processes
47	Hamm/scircuit	170,998	958,936	Real	Circuit Simulation
48	Schenk.ISEI/ohne2	181,343	6,869,939	Real	Semicon. Dev. Sim.
49	Norris/stomach	213,360	3,021,648	Real	Biological Proces.
50	Norris/torso3	259,156	4,429,042	Real	Biological Proces.
51	Sandia/ASIC_320ks	321,671	1,316,085	Real	Circuit Simulation
52	Sandia/ASIC_320k	321,821	1,931,828	Real	Circuit Simulation
53	Rajat/rajat24	358,172	1,946,979	Real	Circuit Simulation
54	Tromble/language	399,130	1,216,334	Real	Natural Lang. Proc.
55	Rajat/rajat21	411,676	1,876,011	Real	Circuit Simulation
56	vanHeukelum/cage13	445,315	7,479,343	Real	Biological Proces.
57	Rajat/rajat29	643,994	3,760,246	Real	Circuit Simulation
58	Rajat/rajat30	643,994	6,175,244	Real	Circuit Simulation
59	ATandT/pre2	659,033	5,834,044	Real	Circuit Simulation
60	Sandia/ASIC_680ks	682,712	1,693,767	Real	Circuit Simulation
61	Sandia/ASIC_680k	682,862	2,638,997	Real	Circuit Simulation
62	Hamrle/Hamrle3	1,447,360	5,514,242	Real	Circuit Simulation
63	vanHeukelum/cage14	1,505,785	27,130,349	Real	Biological Proces.
64	Rajat/rajat31	4,690,002	20,316,253	Real	Circuit Simulation
65	vanHeukelum/cage15	5,154,859	99,199,551	Real	Biological Proces.

This report is organized as follows. First, we present the numerical experiments distinguishing between success, stating the number of iterations required for convergence, and three types of failure: F_1 - maximum number of iterations reached, F_2 - algorithm broke down and F_3 - IILUT failure due to lack of storage, time or singularity of the resulting preconditioner. Second, for every successful case we show logscale plots of the history of relative residual for the IILUT preconditioner (plots for ILUT preconditioner can be found in [8]). Finally, we comment on three typical residual histories of failures and draw the conclusions.

2 Numerical Experiments

Table 5: Numerical Experiments

Matrix	CGNR		GMRES		BiCGStab	
	ILUT	ILUT	ILUT	ILUT	ILUT	ILUT
1	F_3	F_3	F_3	F_3	F_3	F_3
2	F_1	F_1	F_1	F_1	F_1	F_1
3	F_3	F_3	F_3	F_3	F_3	F_3
4	F_3	F_3	F_3	F_3	F_3	F_3
5	F_1	F_3	F_1	F_3	F_1	F_3
6	F_1	F_3	F_1	F_3	F_1	F_3
7	F_1	F_1	103	72	75	49
8	F_1	F_1	31	46	23	28
9	F_3	F_3	F_3	F_3	F_3	F_3
10	F_3	F_3	F_3	F_3	F_3	F_3
11	F_1^*	F_3	31	F_3	18	F_3
12	F_1	F_3	F_1	F_3	F_1	F_3
13	F_1	F_3	F_1	F_3	F_1	F_3
14	F_1	F_3	F_1	F_3	F_1	F_3
15	F_1	18	7	7	4	4
16	F_1	F_3	7	F_3	4	F_3
17	F_1	F_1	F_1	F_1	F_1	F_1
18	F_1	F_3	47	F_3	35	F_3
19	8	F_3	2	F_3	1	F_3
20	F_1	F_3	F_1	F_3	F_1	F_3
21	F_1	F_3	F_1	F_3	F_1	F_3
22	F_1	F_3	F_1	F_3	F_1	F_3
23	F_1	F_3	F_1	F_3	F_1	F_3
24	F_1	F_3	F_1	F_3	F_1	F_3
25	F_1	F_3	F_1	F_3	F_1	F_3
26	F_1	F_3	F_1	F_3	F_1	F_3
27	F_1	F_3	F_1	F_3	F_1	F_3
28	F_1	F_3	F_1	F_3	F_1	F_3
29	4	4	3	3	2	2
30	F_1	F_1	F_1	F_1	F_1	F_1
31	F_1	F_1	182	195	F_1	F_1
32	F_1	F_1	61	58	40	36
33	F_1	F_1	69	60	112	63
34	F_1	F_1	32	28	20	16
35	F_1	F_1	58	54	35	33
36	F_3	F_3	F_3	F_3	F_3	F_3
37	F_1	F_1	F_1	F_1	F_1	F_1

Table 6: Numerical Experiments

Matrix	CGNR		GMRES		BiCGStab	
	IILUT	ILUT	IILUT	ILUT	IILUT	ILUT
38	4	3	3	3	2	2
39	F_1	F_3	F_1	F_3	F_1	F_3
40	F_1	F_3	F_1	F_3	F_1	F_3
41	F_1	F_3	F_1	F_3	F_1	F_3
42	F_1	F_3	F_1	F_3	F_1	F_3
43	F_1	F_3	F_1	F_3	F_1	F_3
44	F_1	F_3	F_1	F_3	F_1	F_3
45	F_1	F_3	F_1	F_3	F_1	F_3
46	F_2	F_1	F_1	249	F_1	F_1
47	F_1	F_3	F_1	F_3	F_1	F_3
48	F_1	F_3	116	F_3	82	F_3
49	16	6	5	4	3	2
50	107	F_1	16	101	13	F_1
51	F_1	27	F_1	F_1	7	4
52	F_1	F_3	F_1	F_3	6	F_3
53	F_1	F_3	F_1	F_3	F_1	F_3
54	F_1	F_1	11	20	9	13
55	F_1	F_3	F_1	F_3	F_1	F_3
56	4	4	3	3	2	2
57	F_3	F_3	F_3	F_3	F_3	F_3
58	F_3	F_3	F_3	F_3	F_3	F_3
59	F_3	F_3	F_3	F_3	F_3	F_3
60	F_2	F_3	16	F_3	13	F_3
61	F_2	F_3	12	F_3	13	F_3
62	F_3	F_3	F_3	F_3	F_3	F_3
63	F_3	3	F_3	F_3	F_3	2
64	F_3	F_3	F_3	F_3	F_3	F_3
65	F_3	F_3	F_3	F_3	F_3	F_3

The convergence entry indicated by * is special in the sense that while the residual for (2) is small, the residual for (1) is relatively large. It is worth noticing that using IILUT preconditioning significantly improves the success of Krylov Subspace methods (see Tab. 7), however it is not enough for us to be able to call them robust.

Table 7: Number of Successful Runs

	CGNR	GMRES	BiCGStab
ILUT	7	15	13
ILUT	6	21	22

For the matrices 7, 8, 11, 15, 16, 18, 19, 29, 31-35, 38, 48-52, 54, 56, 60 and 61 we show the logscale history of relative residual below.

1. Matrix 7

Figure 1: epb3, GMRES (ILUT)





Figure 2: epb3, BiCGStab (ILUT)



2. Matrix 8

Figure 3: poisson3Db, GMRES (ILUT)





Figure 4: poisson3Db, BiCGStab (ILUT)



3. Matrix 11

Figure 5: rajat28, GMRES (ILUT)





Figure 6: rajat28, BICGSTAB (ILUT)



4. Matrix 15

Figure 7: ASIC_100ks, GMRES (ILUT)





Figure 8: ASIC_100ks, BiCGStab (ILUT)



5. Matrix 16

Figure 9: ASIC_100k, GMRES (ILUT)





Figure 10: ASIC_100k, BiCGStab (IILUT)



6. Matrix 18

Figure 11: hcircuit, GMRES (IILUT)





Figure 12: hcircuit, BICGSTAB (ILUT)



7. Matrix 19

Figure 13: lung2, CGNR (ILUT)




Figure 14: lung2, GMRES (ILUT)





Figure 15: lung2, BICGSTAB (ILUT)



8. Matrix 29

Figure 16: torso2, CGNR (ILUT)




Figure 17: torso2, GMRES (ILUT)






Figure 18: torso2, BiCGStab (ILUT)



9. Matrix 31

Figure 19: dc1, GMRES (ILUT)



10. Matrix 32

Figure 20: dc2, GMRES (ILUT)





Figure 21: dc2, BiCGStab (ILUT)



11. Matrix 33

Figure 22: dc3, GMRES (IILUT)





Figure 23: dc3, BiCGStab (IILUT)



12. Matrix 34

Figure 24: trans4, GMRES (ILUT)




Figure 25: trans4, BiCGStab (ILUT)




Figure 26: trans5, GMRES (IILUT)





Figure 27: trans5, BiCGStab (IILUT)



14. Matrix 38

Figure 28: cage12, CGNR (IILUT)




Figure 29: cage12, GMRES (IILUT)





Figure 30: cage12, BiCGStab (ILUT)



15. Matrix 48

Figure 31: ohne2, GMRES (ILUT)





Figure 32: ohne2, BICGSTAB (ILUT)



16. Matrix 49

Figure 33: stomach, CGNR (ILUT)




Figure 34: stomach, GMRES (IILUT)





Figure 35: stomach, BiCGStab (IILUT)



17. Matrix 50

Figure 36: torso3, CGNR (ILUT)




Figure 37: torso3, GMRES (ILUT)






Figure 38: torso3, BiCGStab (ILUT)




18. Matrix 51

Figure 39: ASIC_320ks, BiCGStab (ILUT)



19. Matrix 52

Figure 40: ASIC_320k, BiCGStab (ILUT)



20. Matrix 54

Figure 41: language, GMRES (ILUT)





Figure 42: language, BiCGStab (IILUT)



21. Matrix 56

Figure 43: cage13, CGNR (IILUT)




Figure 44: cage13, GMRES (ILUT)





Figure 45: cage13, BiCGStab (ILUT)



22. Matrix 60

Figure 46: ASIC_680ks, GMRES (ILUT)




Figure 47: ASIC_680ks, BICGSTAB (ILUT)




Figure 48: ASIC_680k, GMRES (ILUT)





Figure 49: ASIC_680k, BICGSTAB (ILUT)



The three typical failures are the following.

1. **Very Slow Convergence or Stagnation:** The residual is decreasing very slowly and does not achieve convergence in the fixed number of iterations. In our experience this would usually indicate that the algorithm will stagnate in the future, because Krylov Subspace methods have a tendency to either converge fast or do not converge at all.
2. **Blow up of the Residual:** The residual increases as the iterations progress, which usually indicates the near singularity of the preconditioner.
3. **Large Oscillations of the Residual:** The residual oscillates significantly between large and small values as iterations progress. This can be caused by the singularity of the preconditioner as well.

3 Conclusions

Finally, as can be seen from the experiments, Krylov Subspace methods are not reliable and have failed in more than 66% of the cases. This is true even with the state of the art black-box IILUT preconditioner. Most of the time methods direct failure, due to for example singularity of the preconditioner F_3 , was avoided. However, the same methods still failed to converge, due to for example stagnation or blow up of the residual F_1 , still indicating problems in the preconditioning. The authors again restate that it is very important for these reasons to search for more robust iterative linear system solvers.

4 Acknowledgments

This work was partially supported by NSF grant NSF-CCF-0635169, Department of Air Force - Air Force Research Laboratory grant FA8750-06-1-0233 and Intel gift.

References

- [1] G. H. Golub and D. P. O’Leary, "Some history of the Conjugate Gradient and Lanczos algorithms: 1948-1976", SIAM Review , Vol. 31, pp. 50-102, 1989.
- [2] Y. Saad and M. H. Schultz, "GMRES: a generalized minimal residual algorithm for solving non symmetric linear systems", SIAM Journal on Scientific and Statistical Computing, Vol. 7, pp. 856-869, 1986.
- [3] H. A. van der Vorst, "Bi-CGSTAB: A fast and smoothly converging variant of Bi-CG for the solution of nonsymmetric linear systems", SIAM Journal on Scientific and Statistical Computing, Vol. 13(2),pp. 631-644, 1992.
- [4] T. C. Chan and H. A. van der Vorst, "Approximate and incomplete factorizations, Parallel Numerical Algorithms", Vol. 4, pp. 167-202, 1997.
- [5] M. Benzi, J. C. Haws and M. Tuma, "Preconditioning highly indefinite and nonsymmetric matrices", SIAM Journal on Scientific and Statistical Computing, Vol. 22(4), pp. 1333-1353, 2000.
- [6] E. Cuthill and J. McKee, "Reducing the bandwidth of sparse symmetric matrices", In Proceedings of the 1969 24th National Conference, pp. 157-172, 1969.
- [7] J. Burkardt, Fortran 90 RCM implementation, January, 2007.
- [8] M. Naumov, M. Manguoglu, C. C. Mikkelsen, A. Arsenieva and A. Sameh, "Reliability of Krylov subspace methods a practical perspective", Technical Report, CSD TR # 07-021, July, 2007.
- [9] G. H. Golub and C. F. Van Loan, "Matrix Computations", Third Edition, The John Hopkins University Press, 1996.