2007

Reliability of Krylov Subspace Methods - A Practical Perspective

M. Naumov

M. Manguoglu

C. C. Mikkelsen

A. Arsenieva

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

Report Number:
07-021

RELIABILITY OF KRYLOV SUBSPACE METHODS
A PRACTICAL PERSPECTIVE

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

Department of Computer Science
Purdue University
West Lafayette, IN 47907

CSD TR #07-022
July 2007
Revised October 2007

Reliability of Krylov Subspace Methods
A Practical Perspective

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

July 8, 2007

Abstract

In this report the reliability of Krylov Subspace iterative methods has been studied. We analyzed the performance of Conjugate Gradient (CG) [1], Generalized Minimum Residual (GMRES) [2] and Stabilized Bi-Conjugate Gradient (BiCGStab) [3] algorithms, with No, Incomplete LU (ILU) [4] and Incomplete LU with fill-in tolerance (ILUT) preconditioning. The 65 nonsymmetric test matrices were selected from the largest sparse matrices of the Tim Davis Matrix Collection. The Krylov Subspace methods failed to produce a solution to the desired accuracy in fixed number of iterations in more than 77% 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, arsenieva@insightbb.com 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

$$A\mathbf{x} = \mathbf{f} \tag{1}$$

using GMRES, BiCGStab and

$$A^T A\mathbf{x} = A^T\mathbf{f} \tag{2}$$

using CG (CGNR). In our experiment we let $\mathbf{f}^T = [1 \ldots 1]$. The software implementation of Krylov Subspace Methods from Harwell Subroutine Library (HSL) and ILUT factorization from SPARSEKIT have been used in our experiments (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≤2h, 24-56≤3h, 57-65≤5h
GMRES	MI24 (HSL)	250	$\leq 10^{-4}$	1-23≤2h, 24-56≤3h, 57-65≤5h
BiCGStab	MI26 (HSL)	250	$\leq 10^{-4}$	1-23≤2h, 24-56≤3h, 57-65≤5h

Table 2: Preconditioning

Algorithm	Subroutine	Dropping Tolerance	Max. Storage
ILU	MA48 (HSL)	elements $\leq 10^{-3}\|\|A\|\|_\infty$	10(#nonzeros)
ILUT	ilut (SPARSEKIT)	elements $\leq 10^{-3}\|\|A\|\|_\infty$ fill-in ≤ 0.1(#nonzeros)	10(#nonzeros)

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 - ILU(T) failure due to lack of storage, time or singularity of the resulting preconditioner. Second, for every successful case we show plots of the history of residual. Finally, we comment on three typical residual histories of failures and draw the conclusions.

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.

2 Numerical Experiments

In this section we present the results of our experiments and show plots of the history of residual for everyone of the successful cases and describe three typical residual histories when the algorithms failed.

Table 5: Numerical Experiments

Matrix	CGNR			GMRES			BiCGStab		
	No	ILU	ILUT	No	ILU	ILUT	No	ILU	ILUT
1	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
2	F_1	F_3	F_1	F_1	F_3	F_1	F_2	F_3	F_1
3	F_1	F_3	F_3	F_1	F_3	F_3	F_2	F_3	F_3
4	F_1	F_3	F_3	F_1	F_3	F_3	F_2	F_3	F_3
5	F_1	F_3	F_3	F_1	F_3	F_3	F_2	F_3	F_3
6	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
7	F_1	F_3	F_1	F_1	F_3	72	F_1	F_3	49
8	F_1	F_3	F_1	189	F_3	46	144	F_3	28
9	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
10	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
11	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
12	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
13	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
14	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
15	F_1	F_3	18	226	F_3	7	F_2	F_3	4
16	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
17	F_1	F_3	F_1	F_1	F_3	F_1	F_1	F_3	F_1
18	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
19	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
20	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
21	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
22	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
23	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
24	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
25	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
26	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
27	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
28	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
29	103	100	4	18	3	3	10	3	2
30	F_1	F_3	F_1	F_1	F_3	F_1	F_1	F_3	F_1
31	F_1	F_3	F_1	F_1	F_3	195	F_1	F_3	F_1
32	F_1	F_3	F_1	F_1	F_3	58	F_1	F_3	36
33	F_1	F_3	F_1	F_1	F_3	60	F_1	F_3	63
34	F_1	F_3	F_1	F_1	F_3	28	F_1	F_3	16

Table 6: Numerical Experiments

Matrix	CGNR			GMRES			BiCGStab		
	No	ILU	ILUT	No	ILU	ILUT	No	ILU	ILUT
35	F_1	F_3	F_1	F_1	F_3	54	F_1	F_3	33
36	F_1	F_3	F_3	F_1	F_3	F_3	F_2	F_3	F_3
37	F_1^*	F_3	F_1	F_1	F_3	F_1	F_1	F_3	F_1
38	12	4	3	9	3	3	F_2	2	2
39	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
40	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
41	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
42	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
43	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
44	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
45	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
46	F_1	F_3	F_1	F_1	F_3	249	F_1	F_3	F_1
47	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
48	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
49	115	19	6	85	2	4	F_2	2	2
50	F_1	127	F_1	68	4	101	44	5	F_1
51	F_1	F_3	27	F_1	F_3	F_1	F_1	F_3	4
52	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
53	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
54	F_1	3	F_1	26	F_3	20	21	F_3	13
55	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
56	20	6	4	10	3	3	6	2	2
57	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
58	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
59	F_1	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
60	F_1^*	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
61	F_1^*	F_3	F_3	F_1	F_3	F_3	F_1	F_3	F_3
62	F_1	F_3	F_3	F_3	F_3	F_3	F_1	F_3	F_3
63	9	F_3	3	F_3	F_3	F_3	F_2	F_3	2
64	F_1	F_3	F_3	F_3	F_3	F_3	F_1	F_3	F_3
65	F_3	F_3	F_3	F_3	F_3	F_3	F_3	F_3	F_3

The failures indicated by $*$ are special in the sense that while the residual for (2) is small, the residual for (1) is large. The success of krylov subspace methods is summarized in Tab. 7.

Table 7: Number of Successful Runs

	CGNR	GMRES	BiCGStab
No	5	8	5
ILU	6	5	5
ILUT	7	15	14

For the matrices 7, 8, 15, 29, 31-35, 38, 46, 49, 50, 51, 54, 56 and 63 we show the history of residual below.

1. Matrix 7

Figure 1: epb3, GMRES (ILUT)

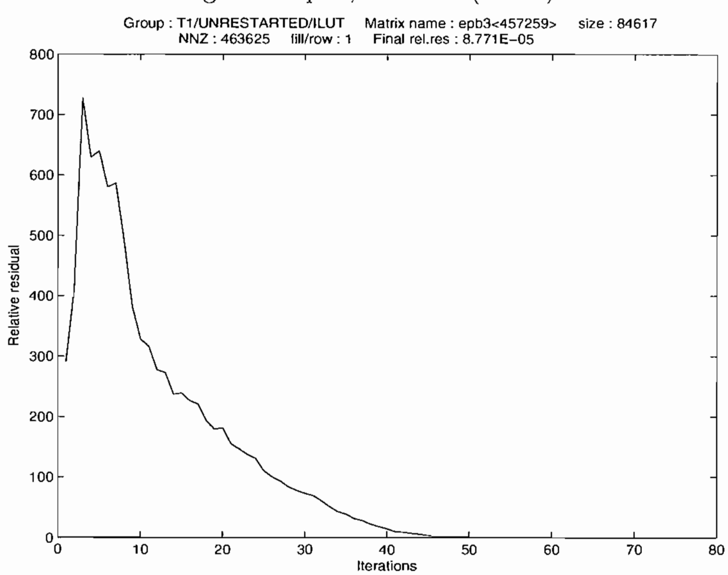

Figure 2: epb3, BiCGStab (ILUT)

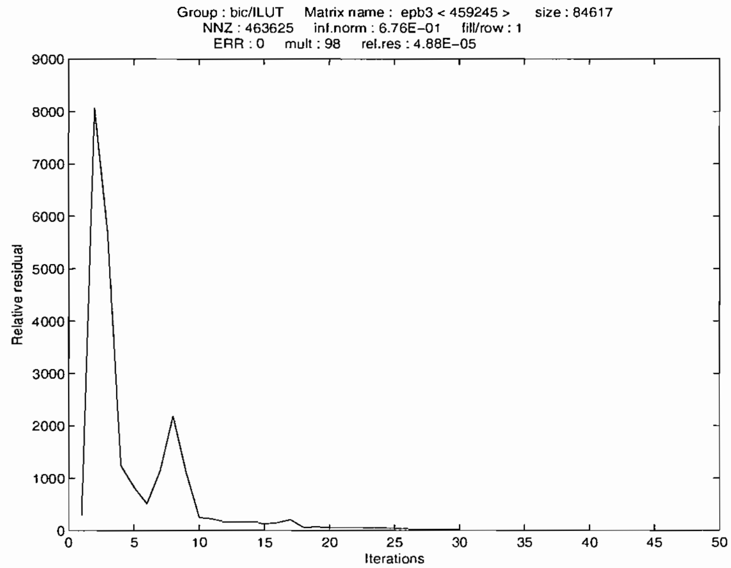

2. Matrix 8

Figure 3: poisson3Db, GMRES (No)

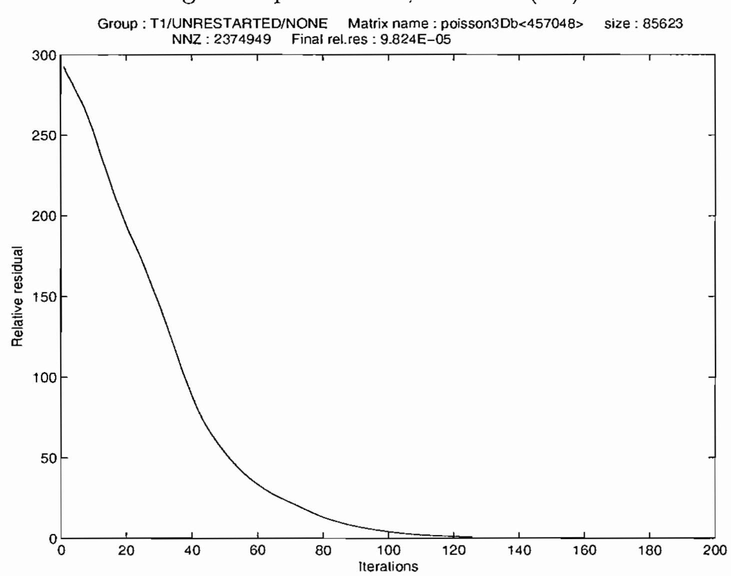

Figure 4: poisson3Db, GMRES (ILUT)

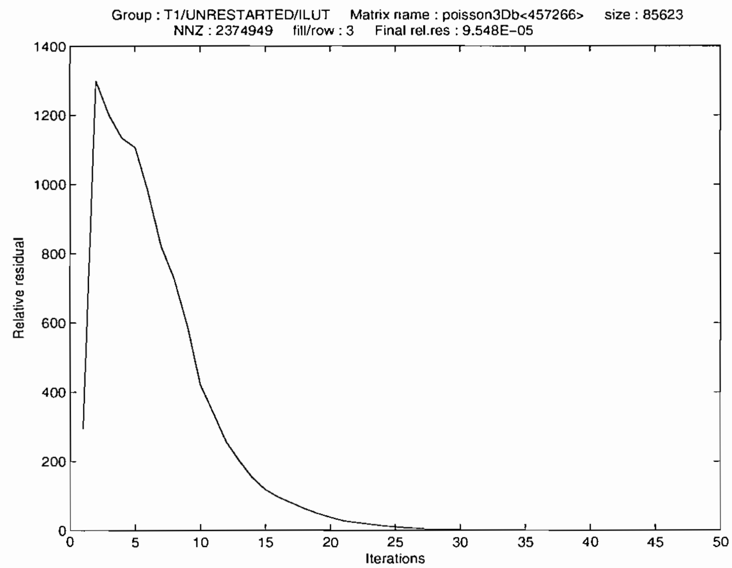

Figure 5: poisson3Db, BiCGStab (No)

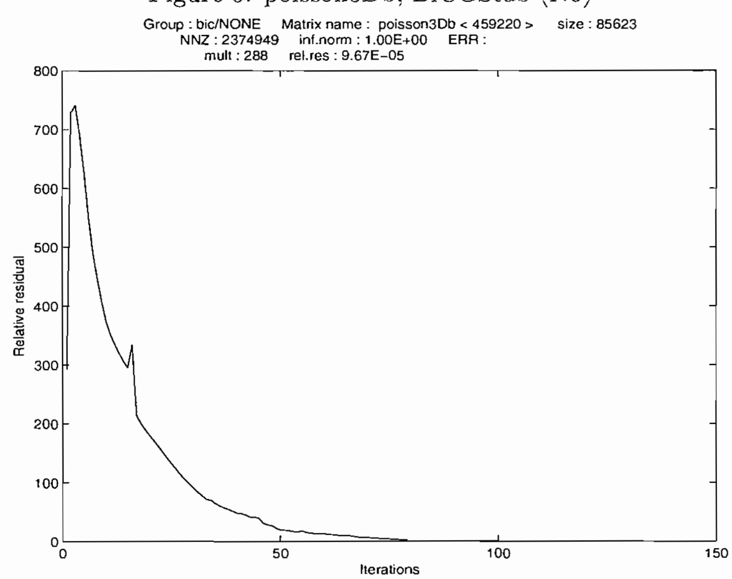

Figure 6: poisson3Db, BiCGStab (ILUT)

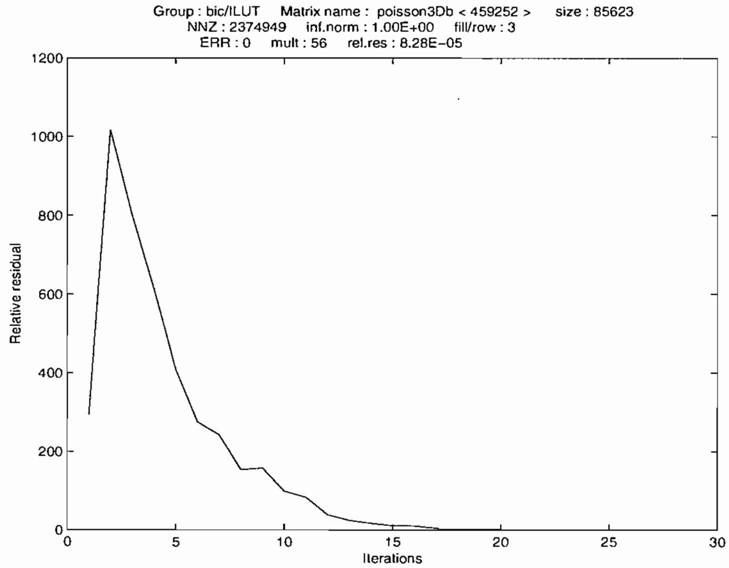

3. Matrix 15

Figure 7: ASIC_100ks, CGNR (ILUT)

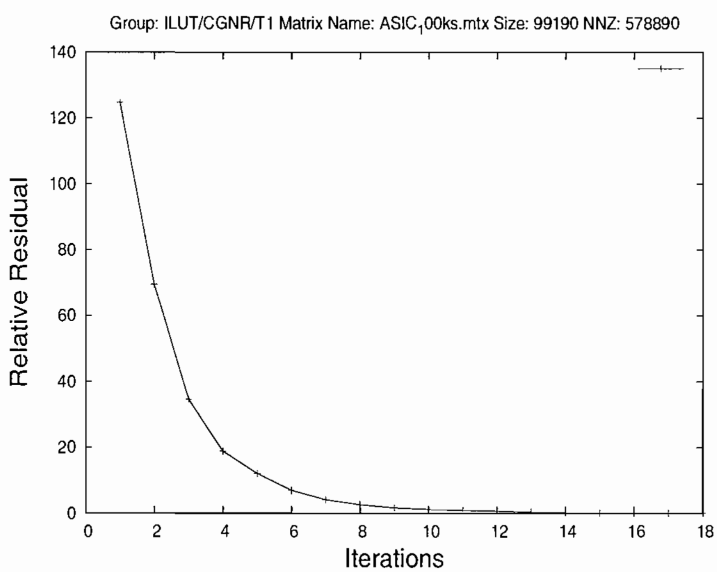

Figure 8: ASIC_100ks, GMRES (No)

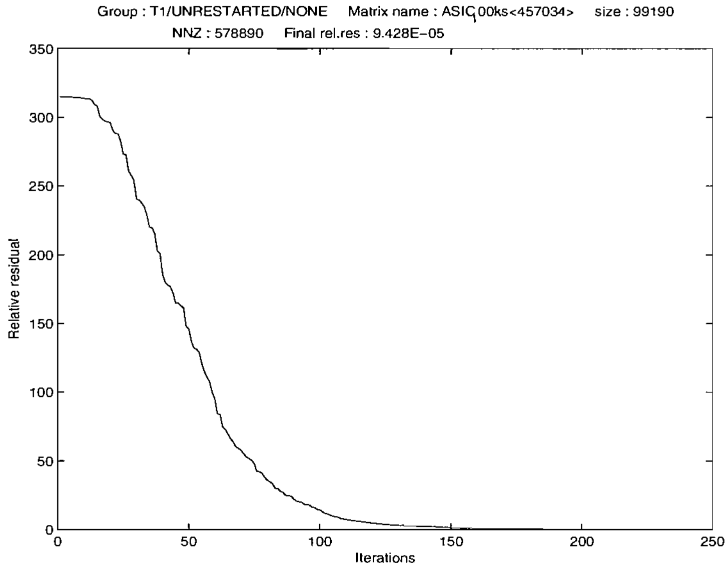

Figure 9: ASIC_100ks, GMRES (ILUT)

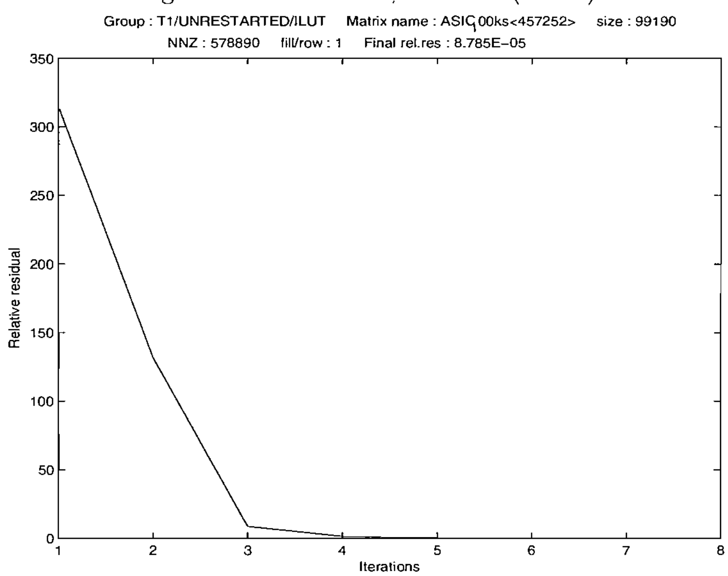

Figure 10: ASIC_100ks, BiCGStab (ILUT)

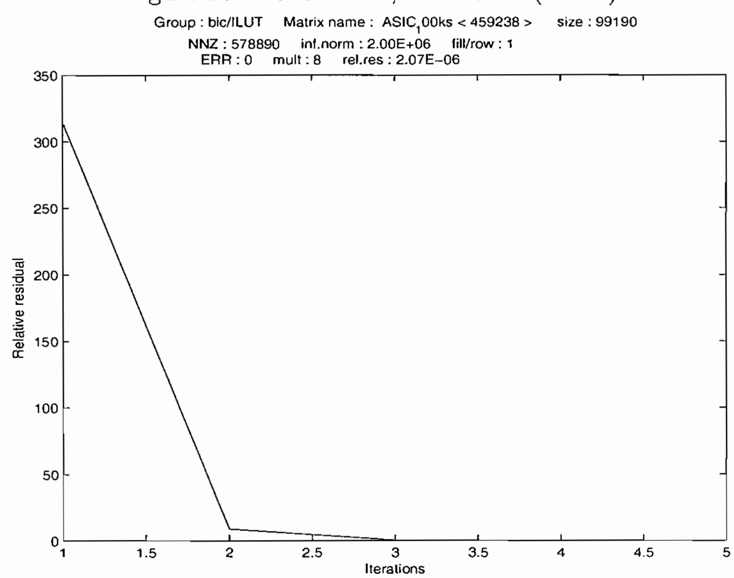

4. Matrix 29

Figure 11: torso2, CGNR (No)

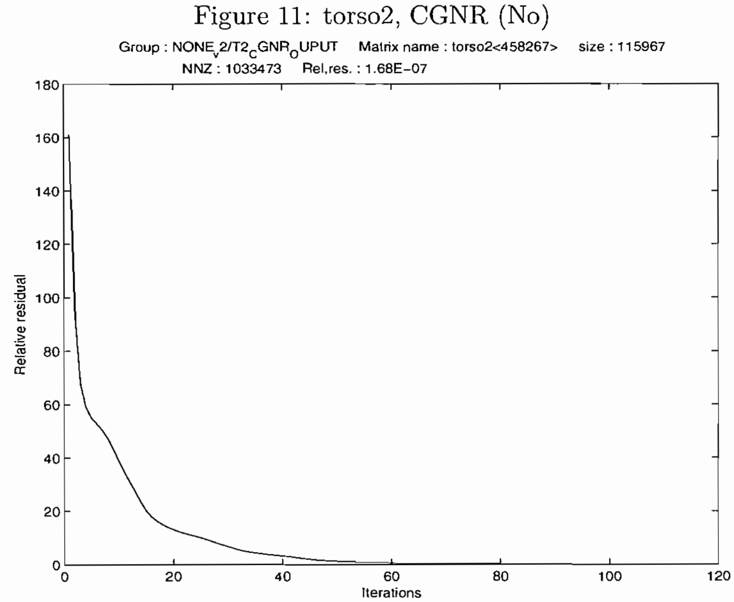

Figure 12: torso2, CGNR (ILU)

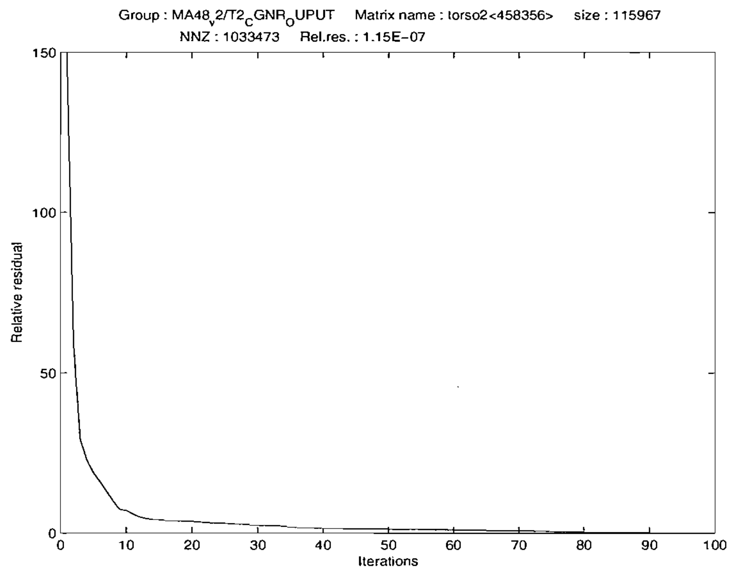

Figure 13: torso2, CGNR (ILUT)

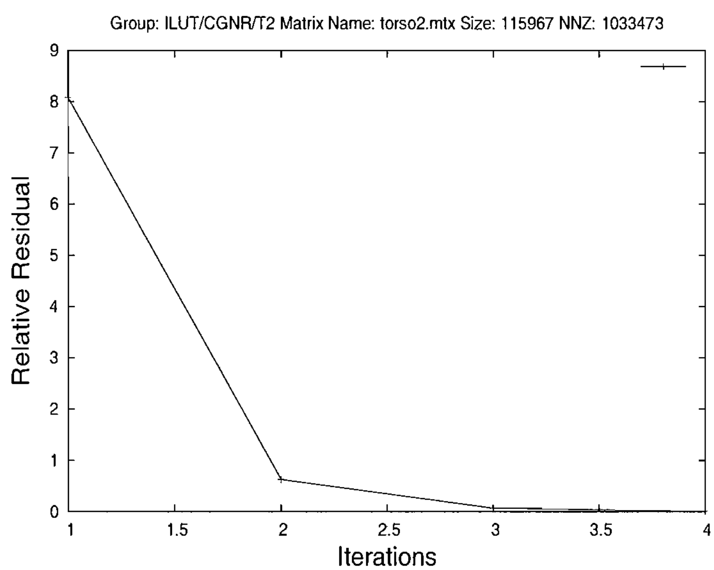

Figure 14: torso2, GMRES (No)

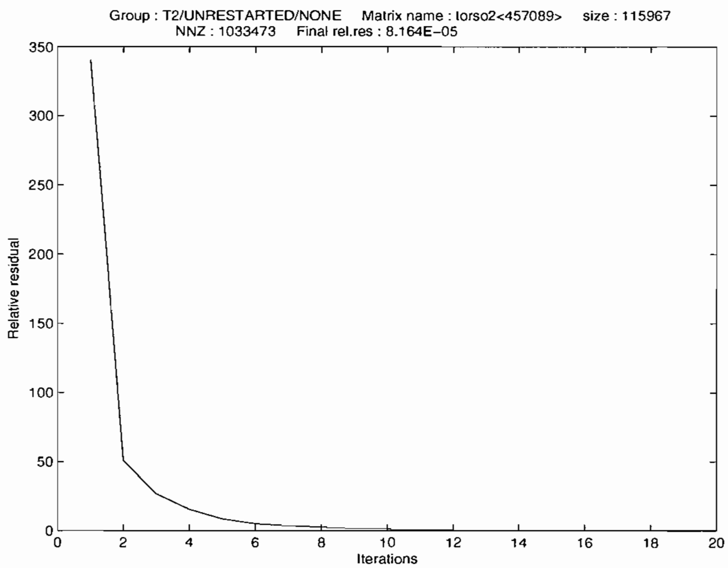

Figure 15: torso2, GMRES (ILU)

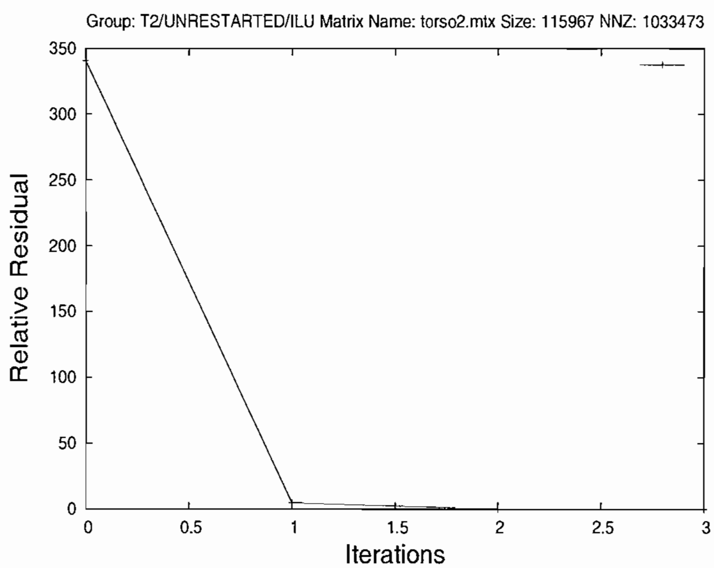

Figure 16: torso2, GMRES (ILUT)

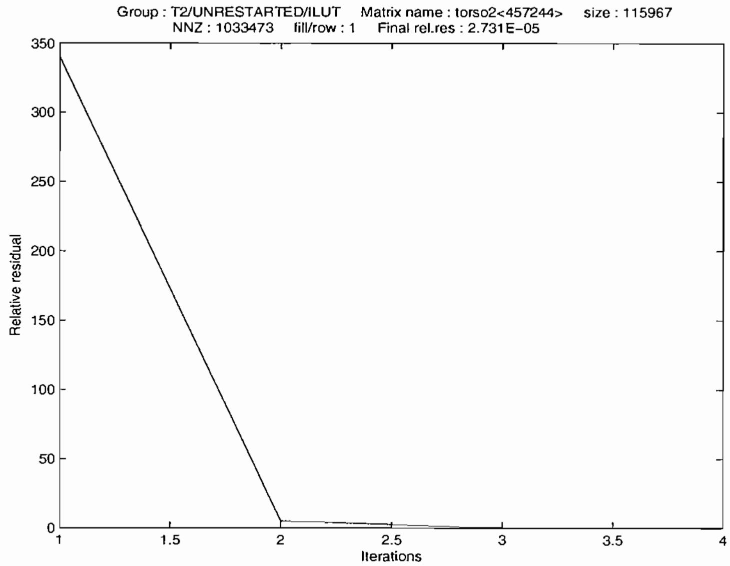

Figure 17: torso2, BiCGStab (No)

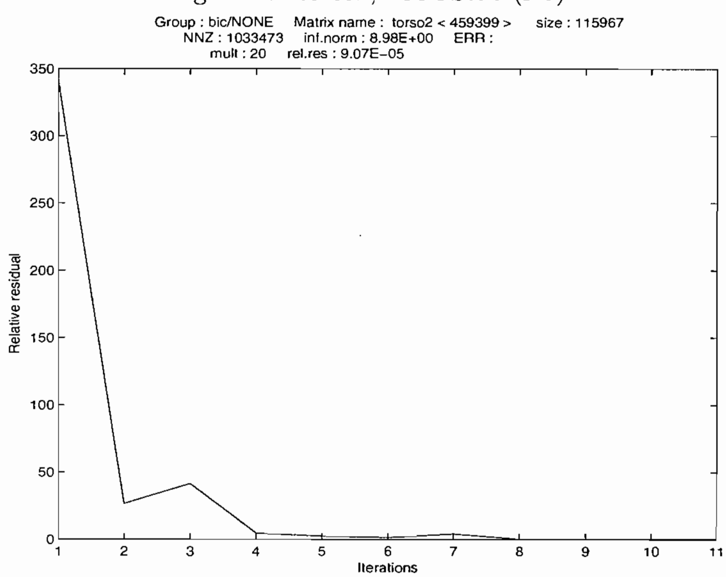

Figure 18: torso2, BiCGStab (ILU)

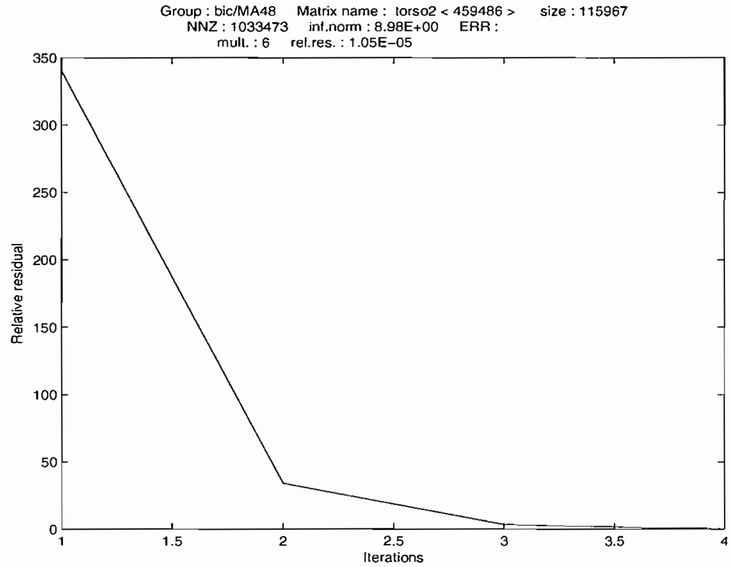

Figure 19: torso2, BiCGStab (ILUT)

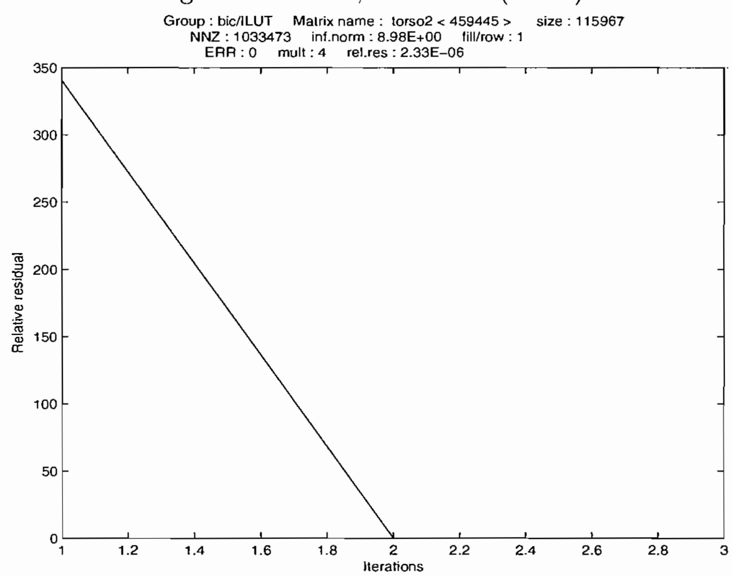

5. Matrix 31

Figure 20: dc1, GMRES (ILUT)

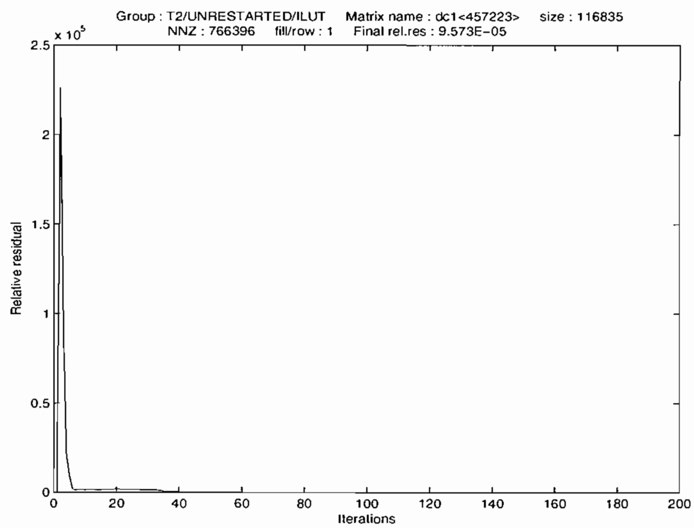

6. Matrix 32

Figure 21: dc2, GMRES (ILUT)

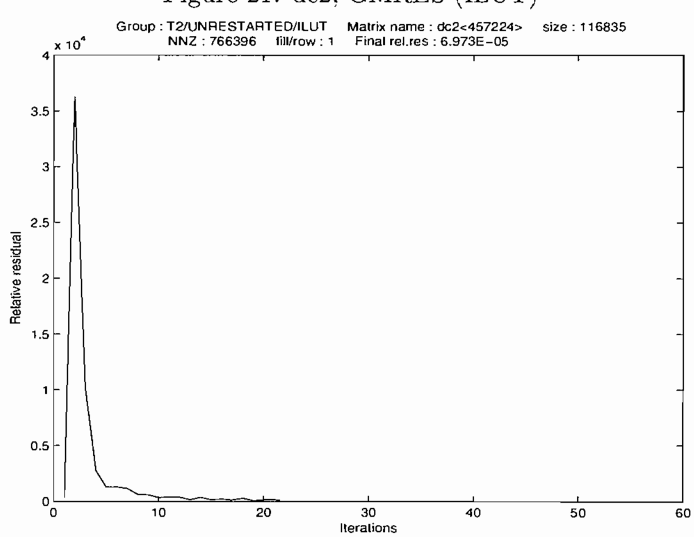

Figure 22: dc2, BiCGStab (ILUT)

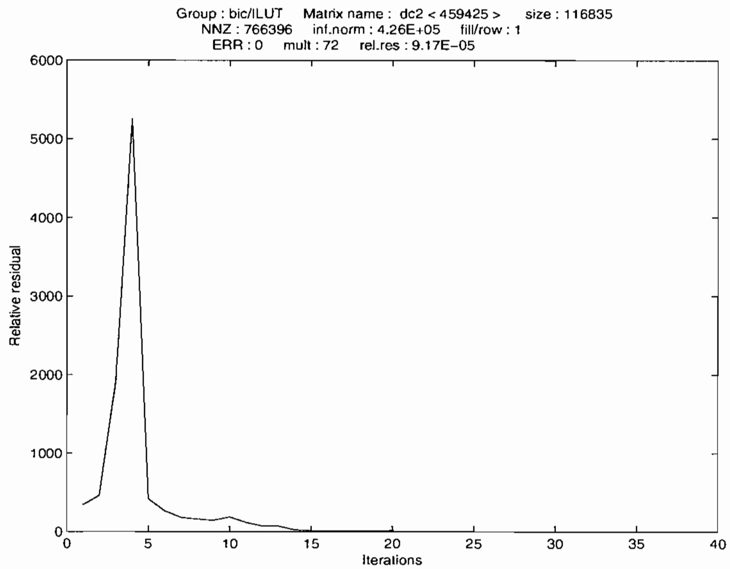

7. Matrix 33

Figure 23: dc3, GMRES (ILUT)

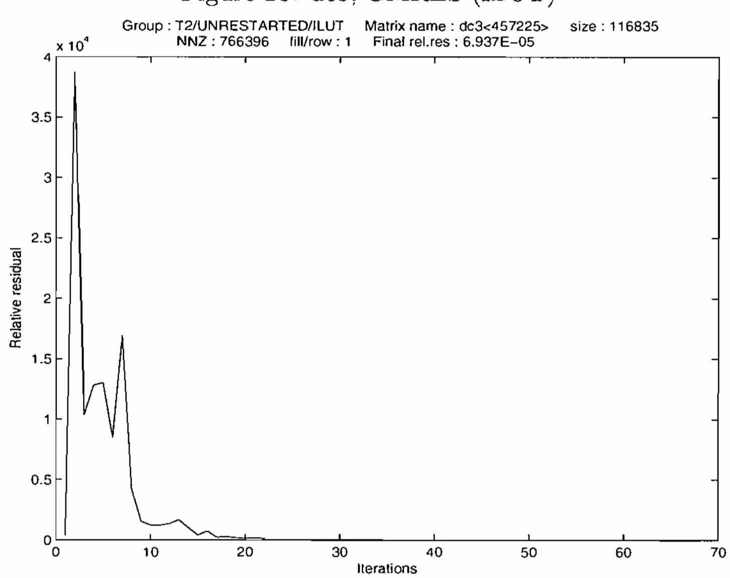

Figure 24: dc3, BiCGStab (ILUT)

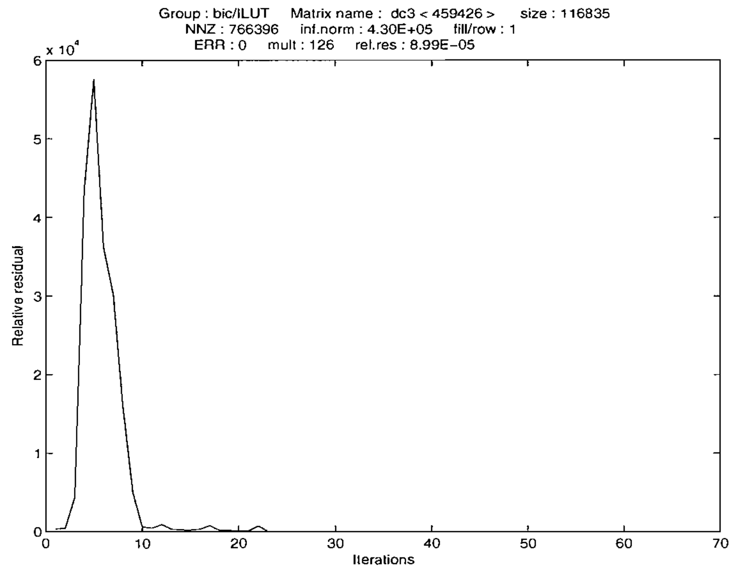

8. Matrix 34

Figure 25: trans4, GMRES (ILUT)

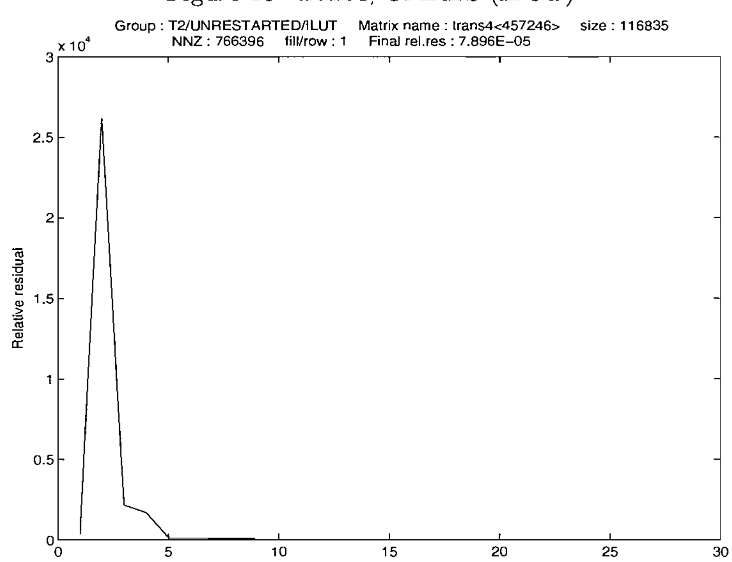

Figure 26: trans4, BiCGStab (ILUT)

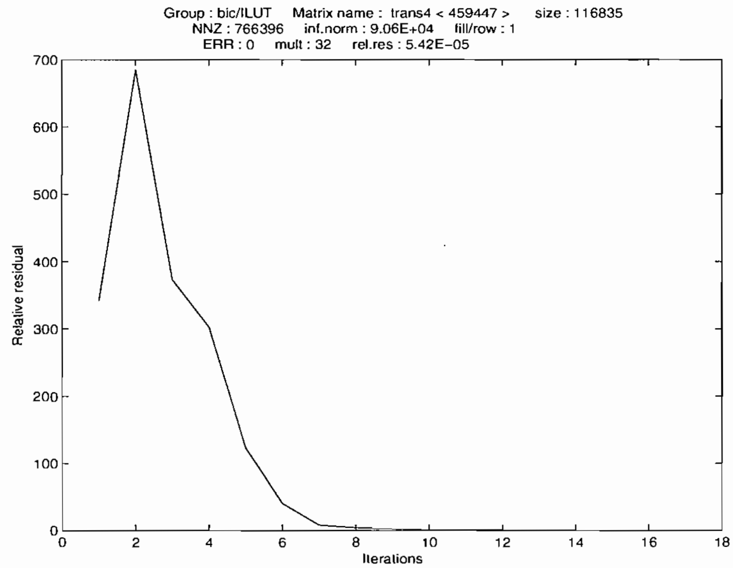

9. Matrix 35

Figure 27: trans5, GMRES (ILUT)

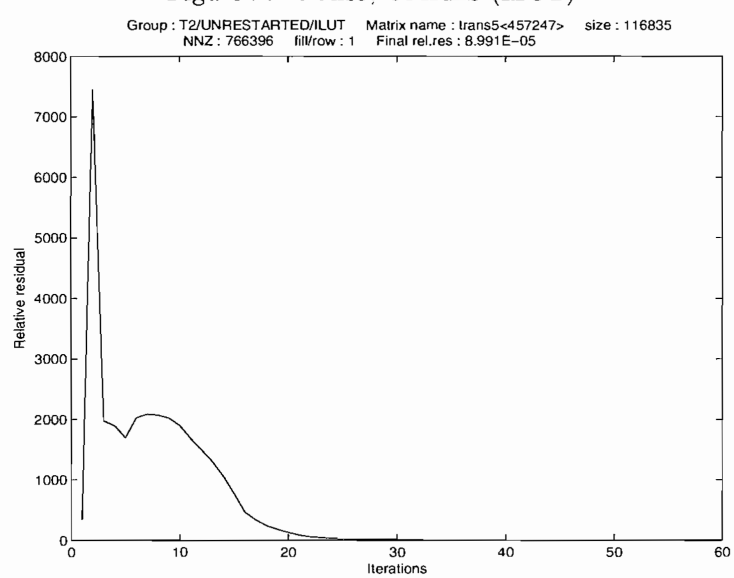

Figure 28: trans5, BiCGStab (ILUT)

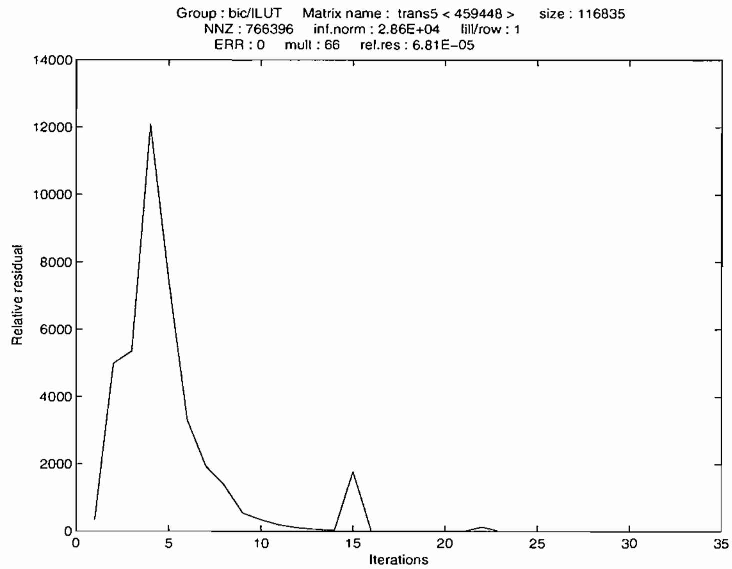

10. Matrix 38

Figure 29: cage12, CGNR (No)

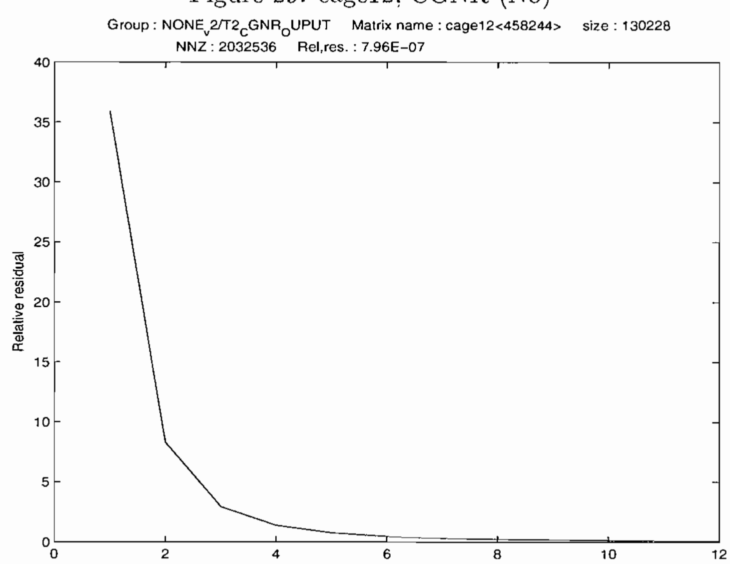

Figure 30: cage12, CGNR (ILU)

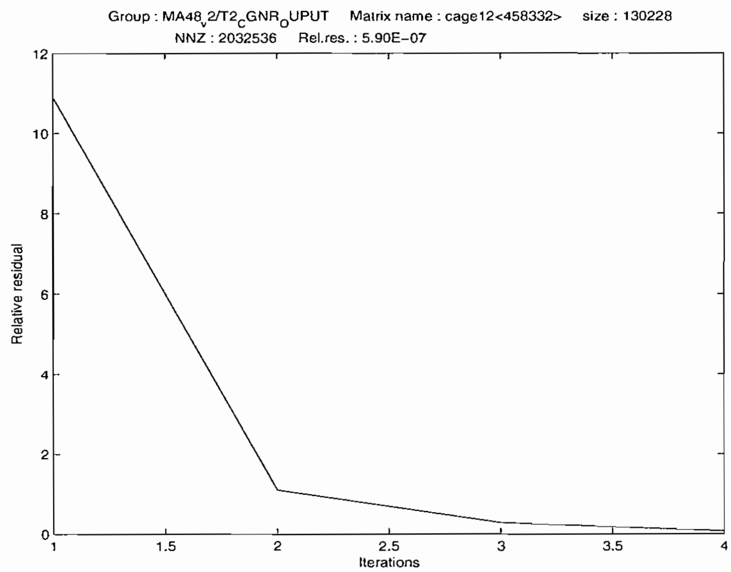

Figure 31: cage12, CGNR (ILUT)

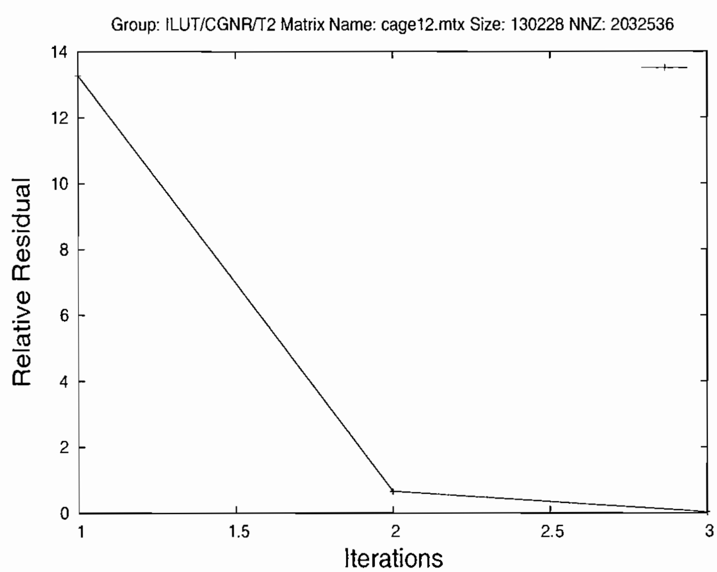

Figure 32: cage12, GMRES (No)

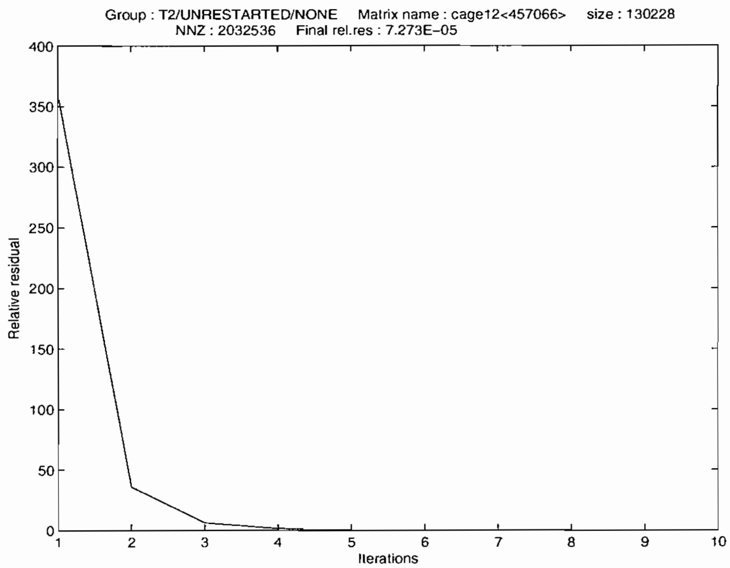

Figure 33: cage12, GMRES (ILU)

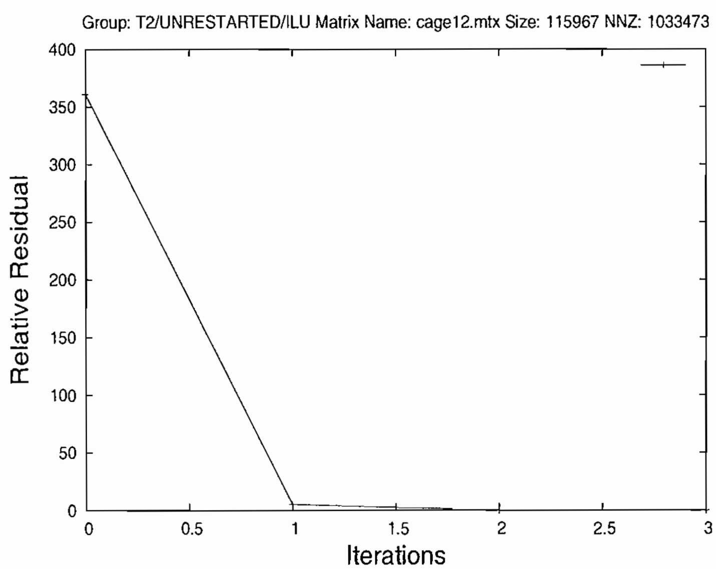

Figure 34: cage12, GMRES (ILUT)

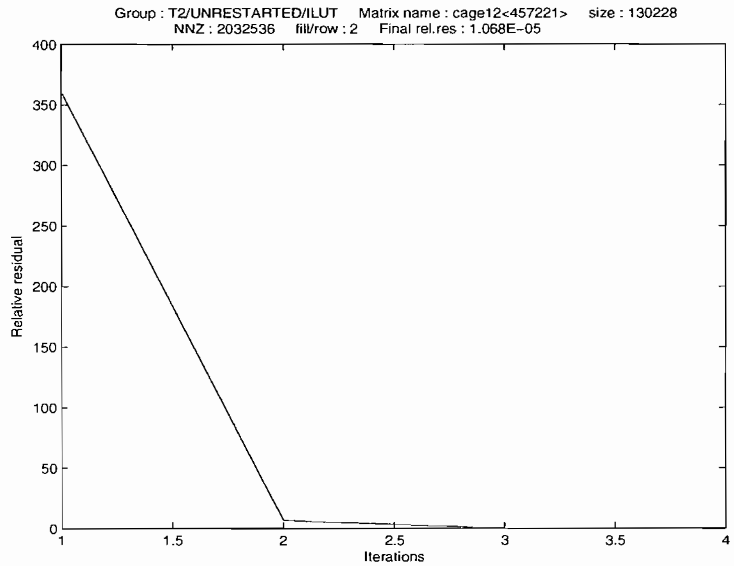

Figure 35: cage12, BiCGStab (ILU)

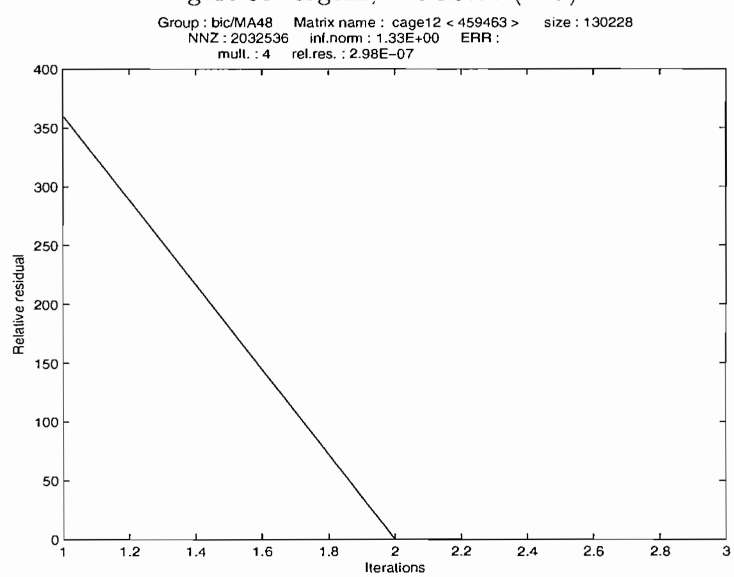

Figure 36: cage12, BiCGStab (ILUT)

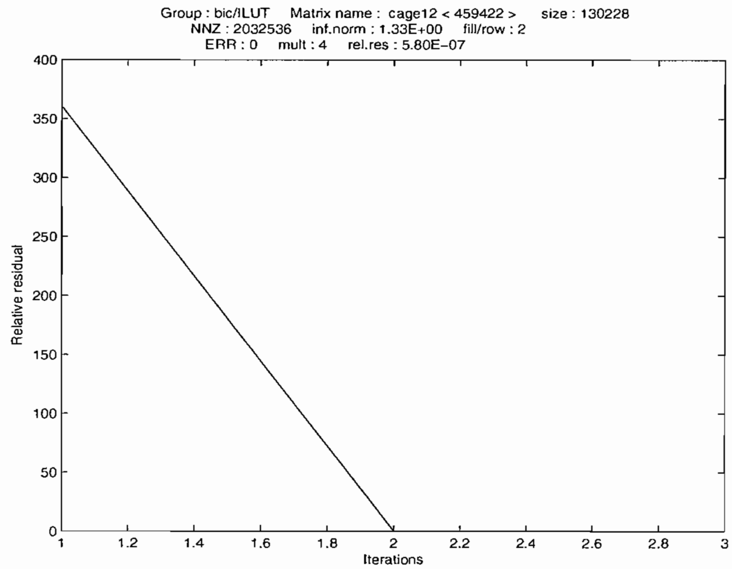

11. Matrix 46

Figure 37: xenon2, GMRES (ILUT)

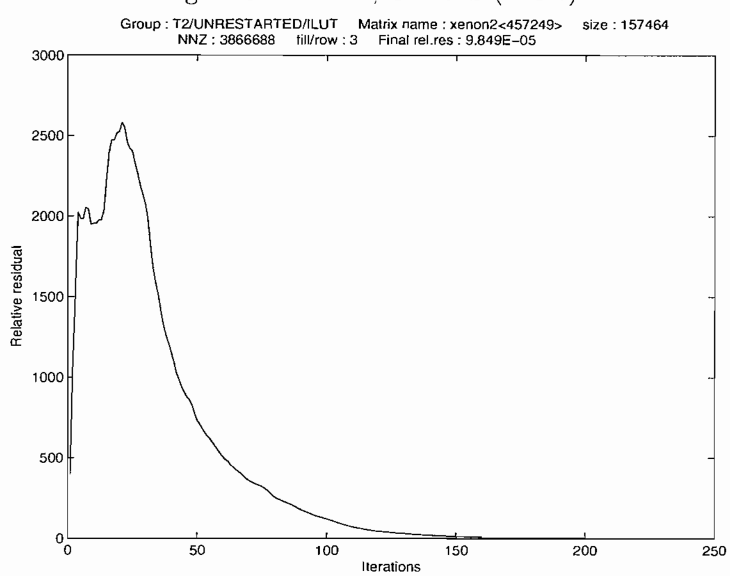

12. Matrix 49

Figure 38: stomach, CGNR (No)

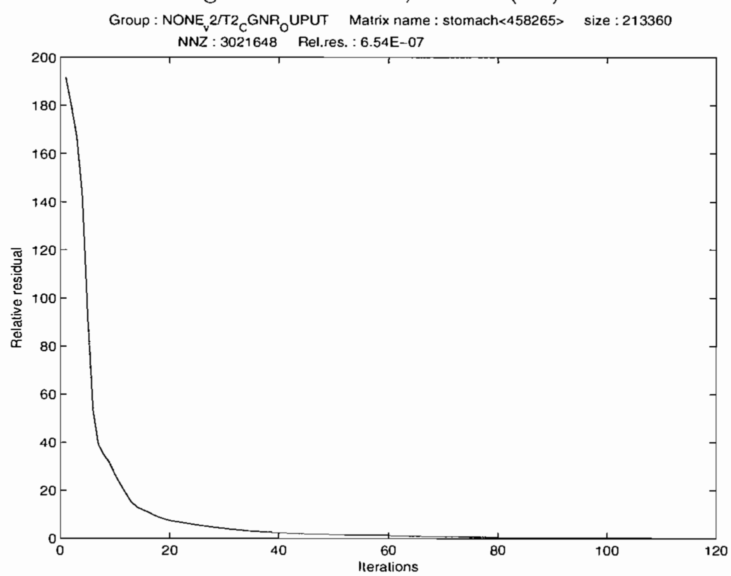

Figure 39: stomach, CGNR (ILU)

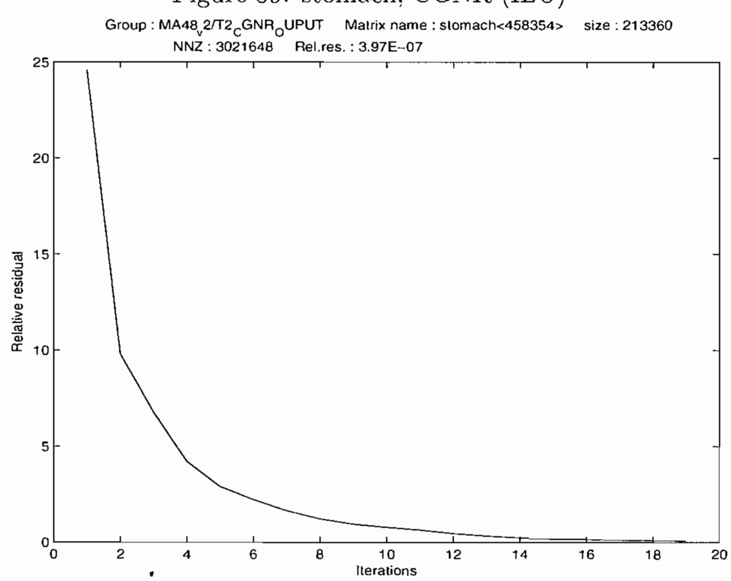

Figure 40: stomach, CGNR (ILUT)

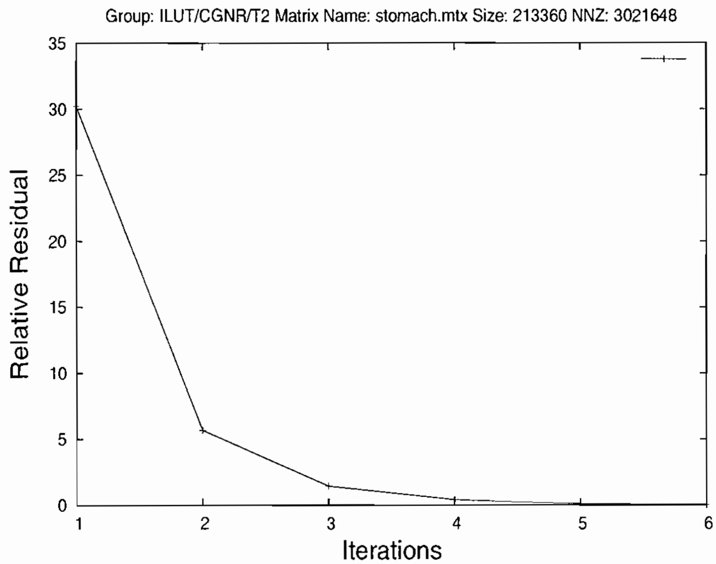

Figure 41: stomach, GMRES (No)

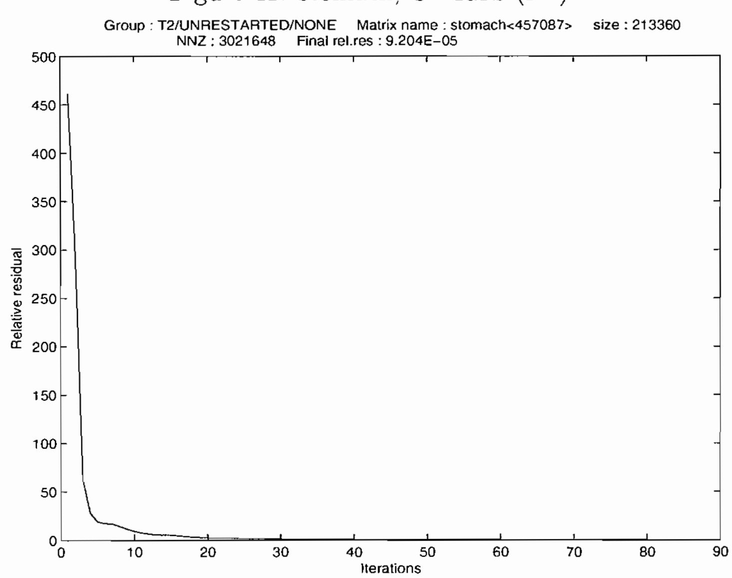

Figure 42: stomach, GMRES (ILU)

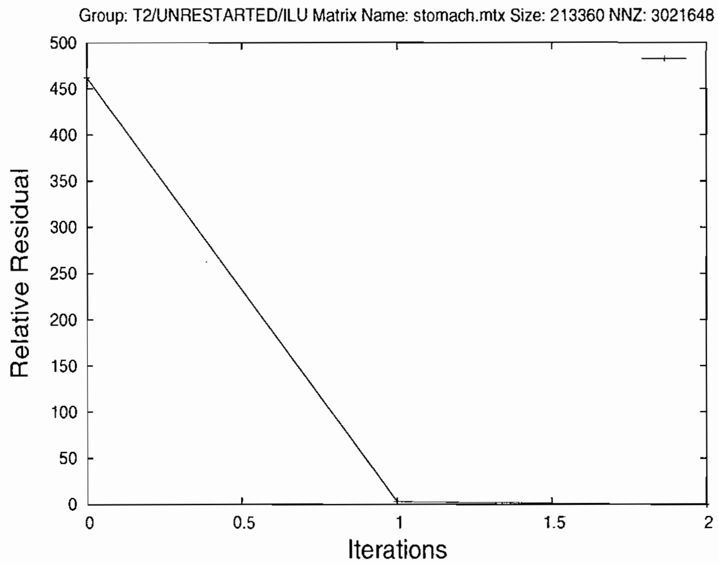

Figure 43: stomach, GMRES (ILUT)

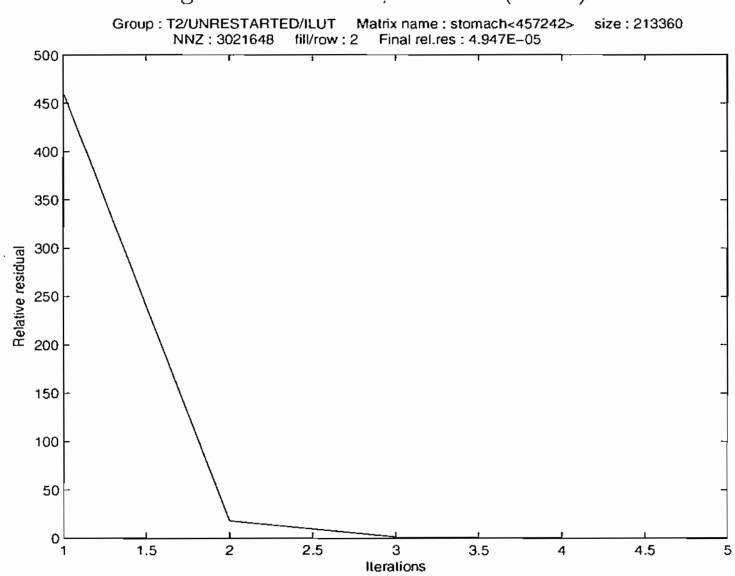

Figure 44: stomach, BiCGStab (ILU)

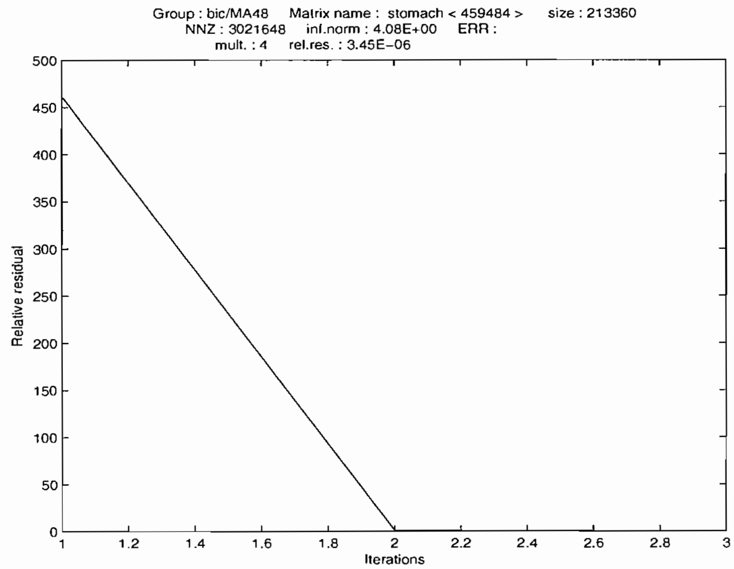

Figure 45: stomach, BiCGStab (ILUT)

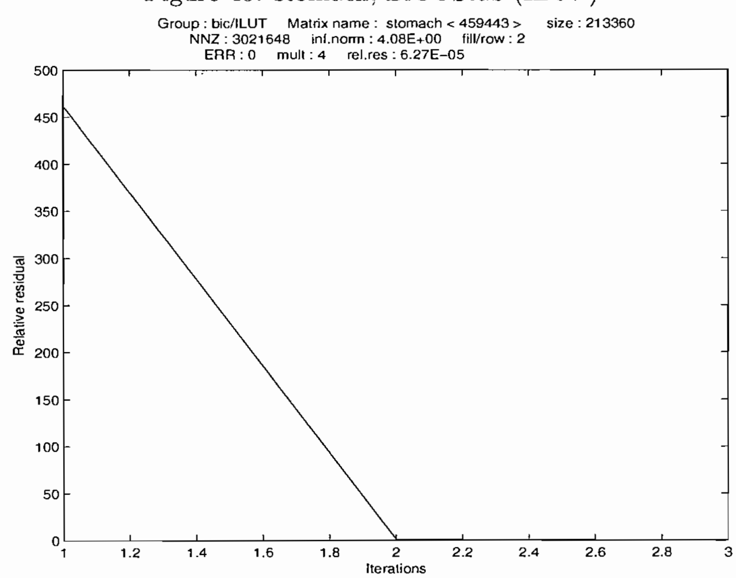

13. Matrix 50

Figure 46: torso3, CGNR (ILU)

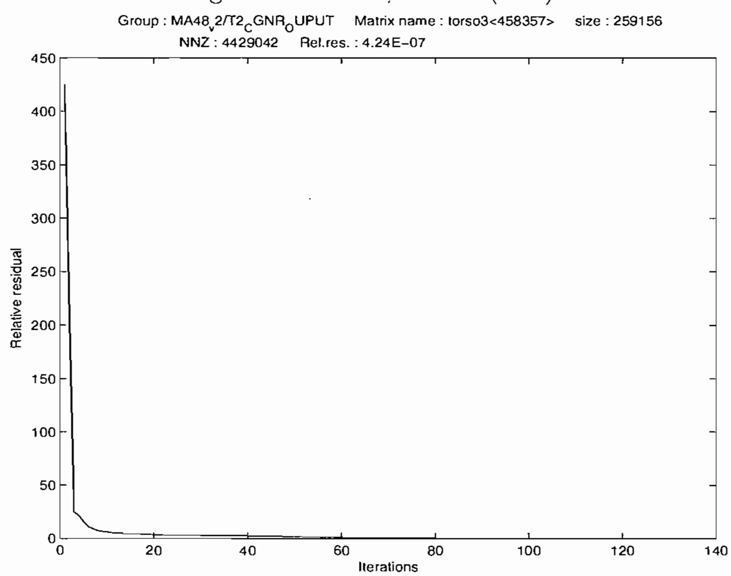

Figure 47: torso3, GMRES (No)

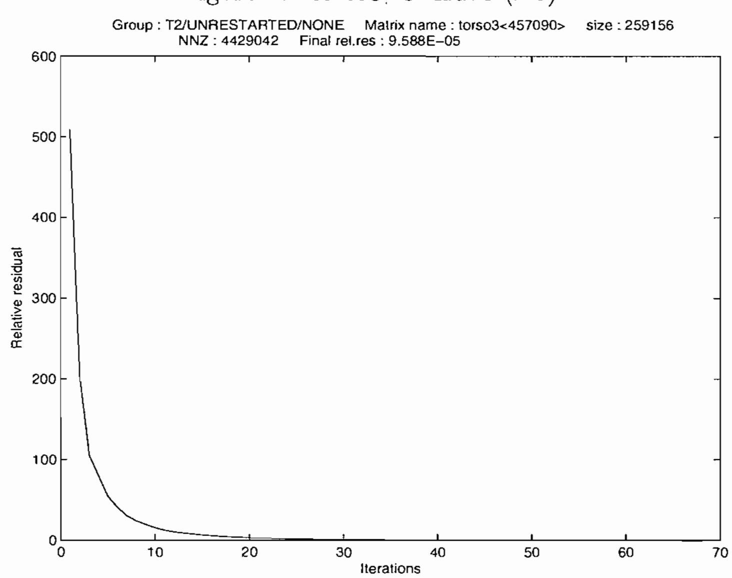

Figure 48: torso3, GMRES (ILU)

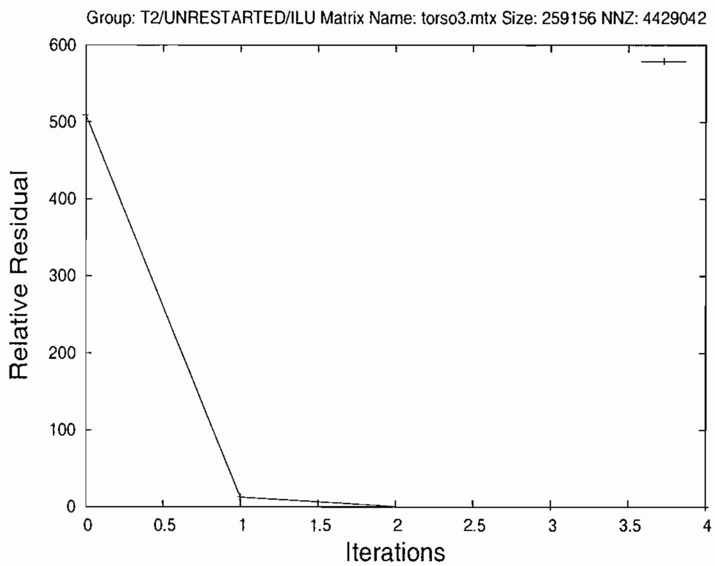

Figure 49: torso3, GMRES (ILUT)

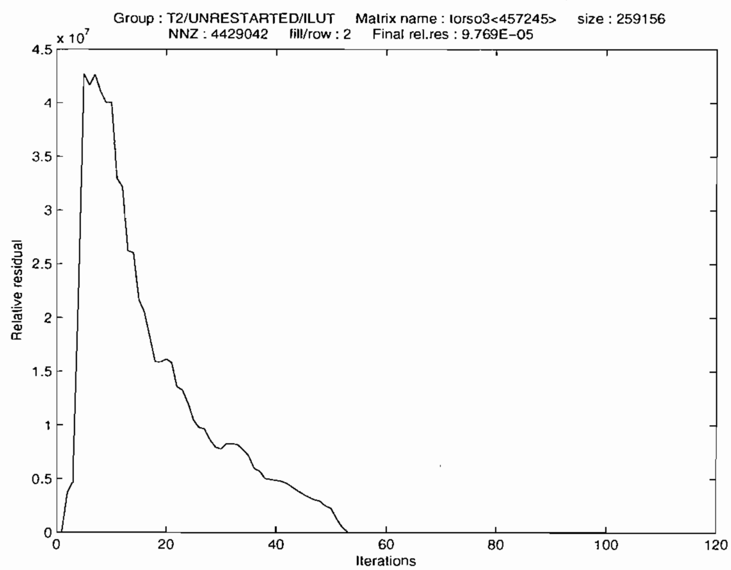

Figure 50: torso3, BiCGStab (No)

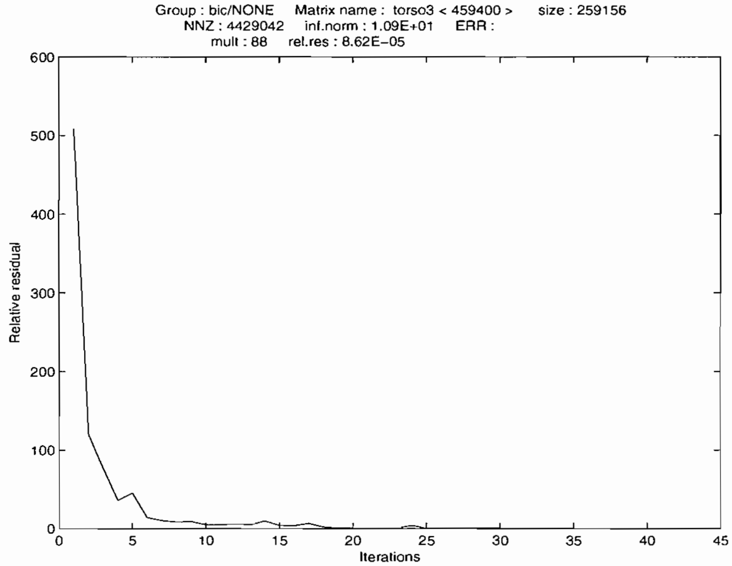

Figure 51: torso3, BiCGStab (ILU)

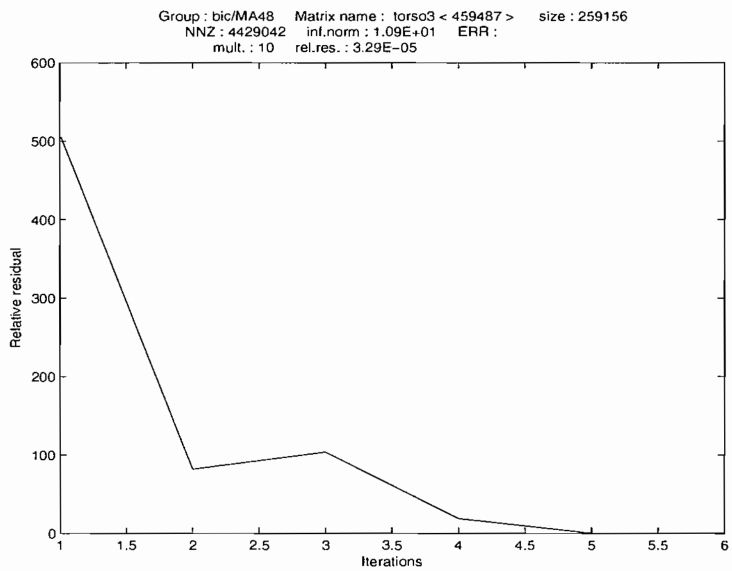

14. Matrix 51

Figure 52: ASIC_320ks, CGNR (ILUT)

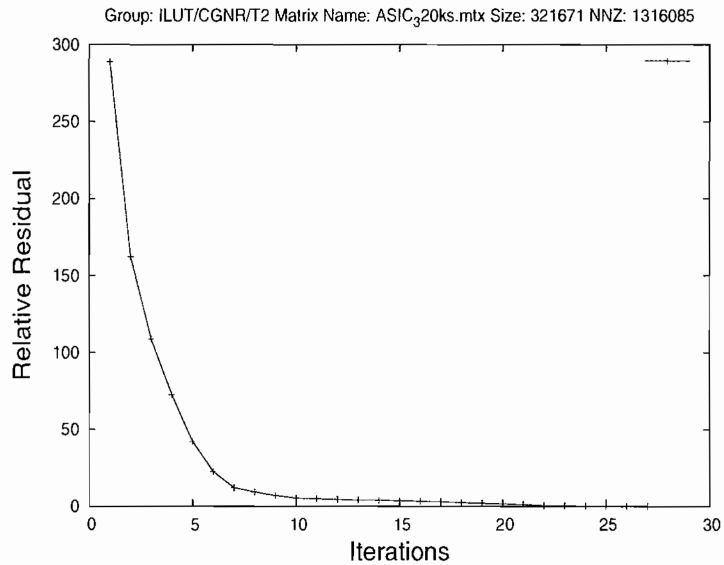

Figure 53: ASIC_320ks, BiCGStab (ILUT)

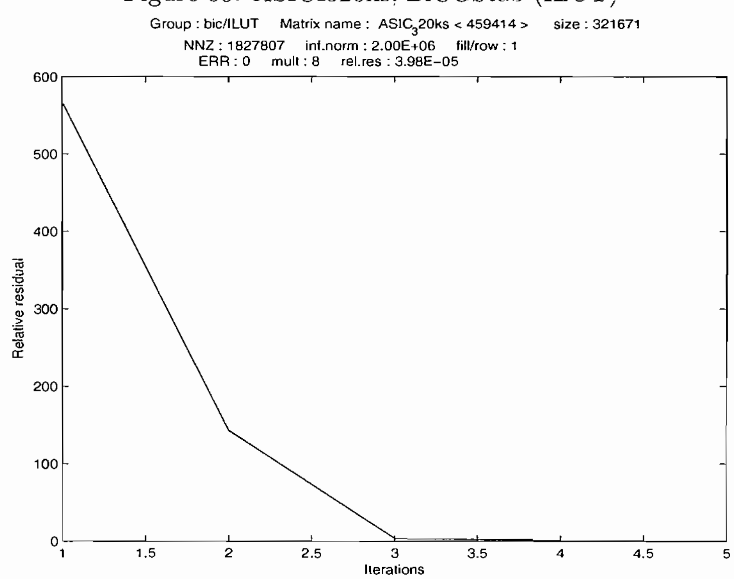

15. Matrix 54

Figure 54: language, CGNR (ILU)

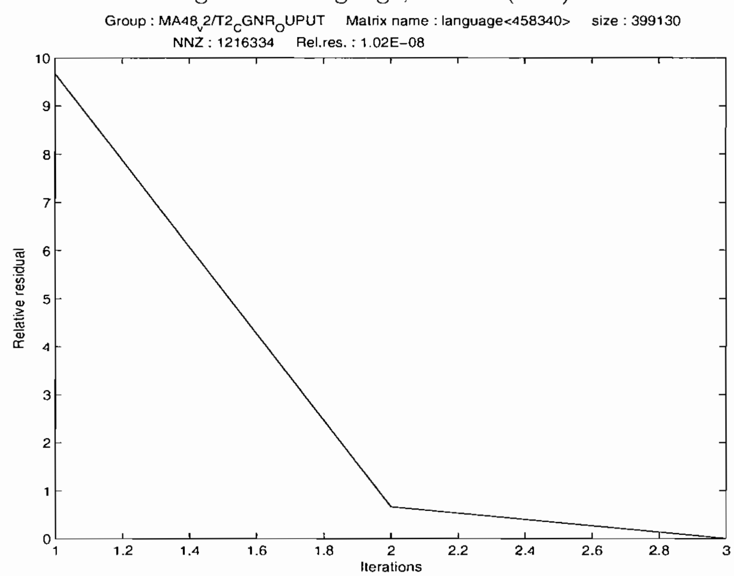

Figure 55: language, GMRES (No)

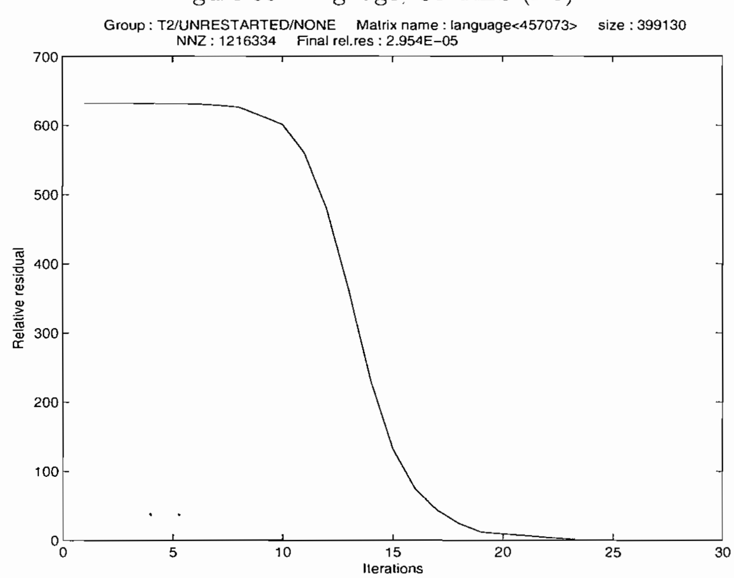

Figure 56: language, GMRES (ILUT)

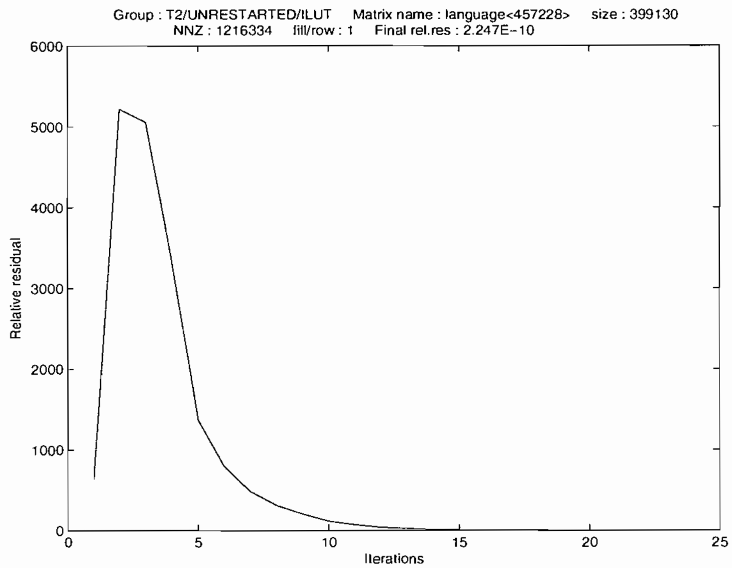

Figure 57: language, BiCGStab (No)

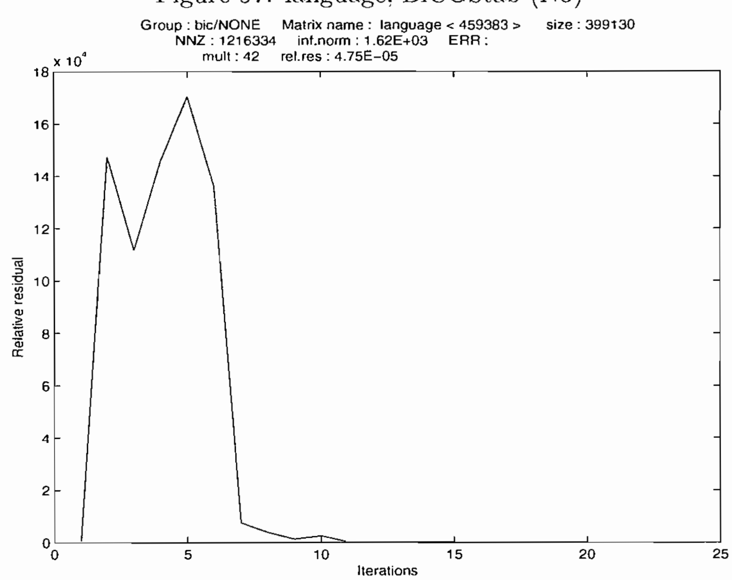

Figure 58: language, BiCGStab (ILUT)

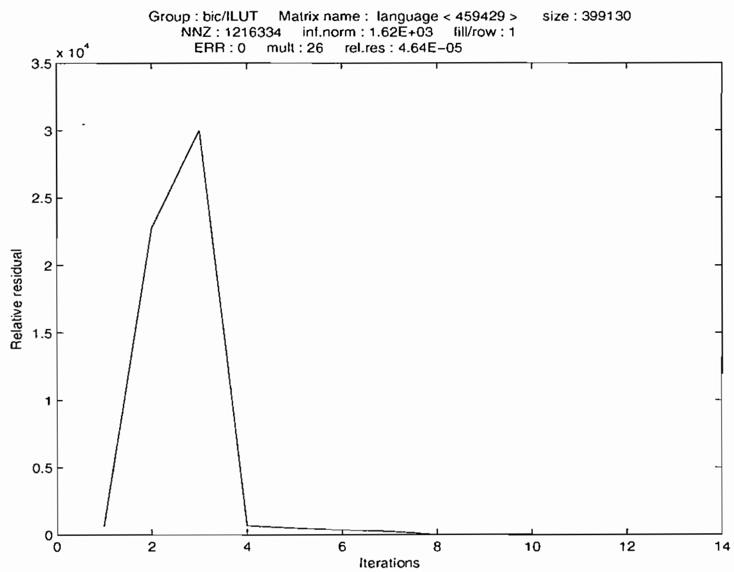

16. Matrix 56

Figure 59: cage13, CGNR (No)

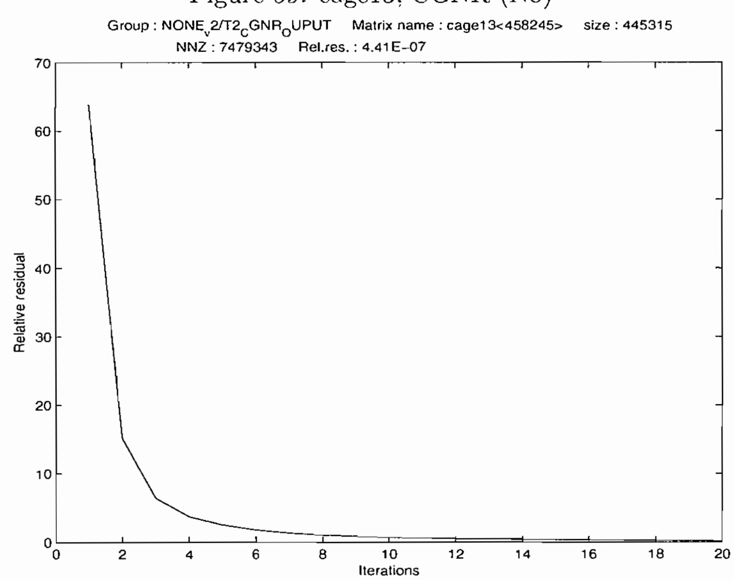

Figure 60: cage13. CGNR (ILU)

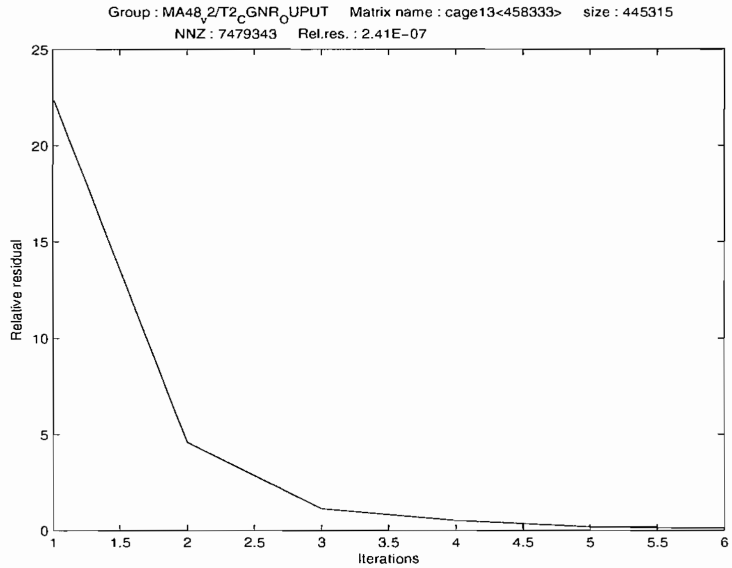

Figure 61: cage13, CGNR (ILUT)

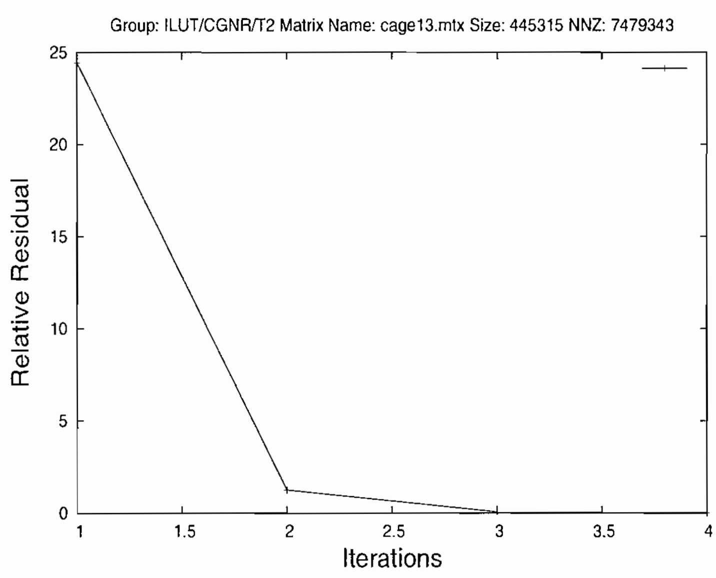

Figure 62: cage13. GMRES (No)

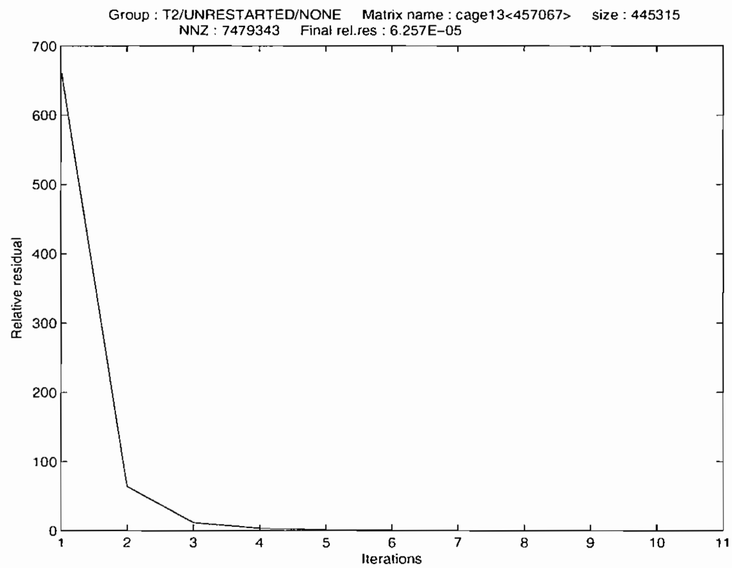

Figure 63: cage13. GMRES (ILU)

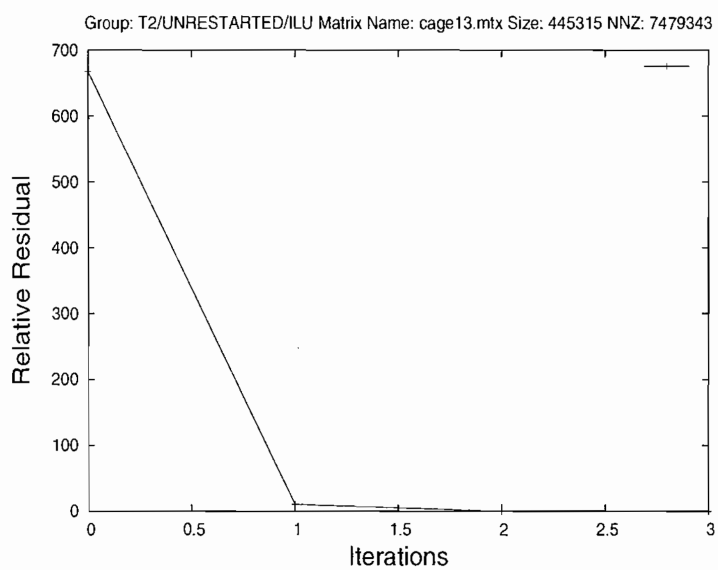

Figure 64: cage13, GMRES (ILUT)

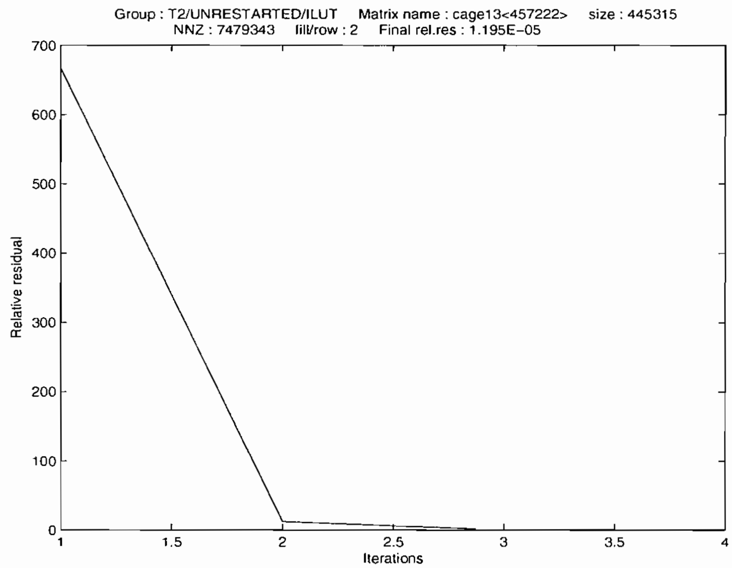

Figure 65: cage13, BiCGStab (No)

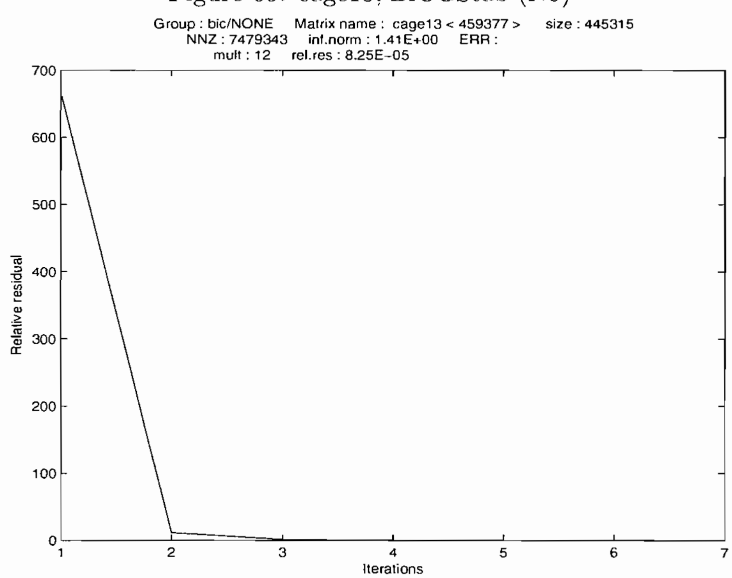

Figure 66: cage13, BiCGStab (ILU)

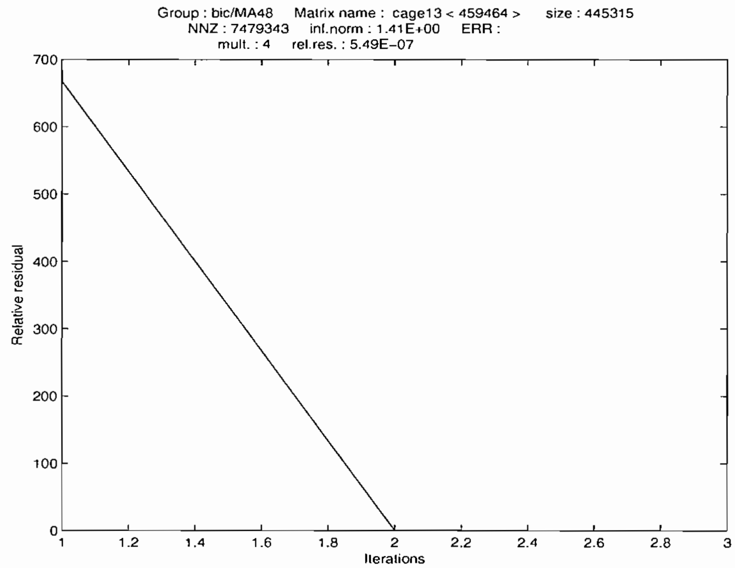

Figure 67: cage13, BiCGStab (ILUT)

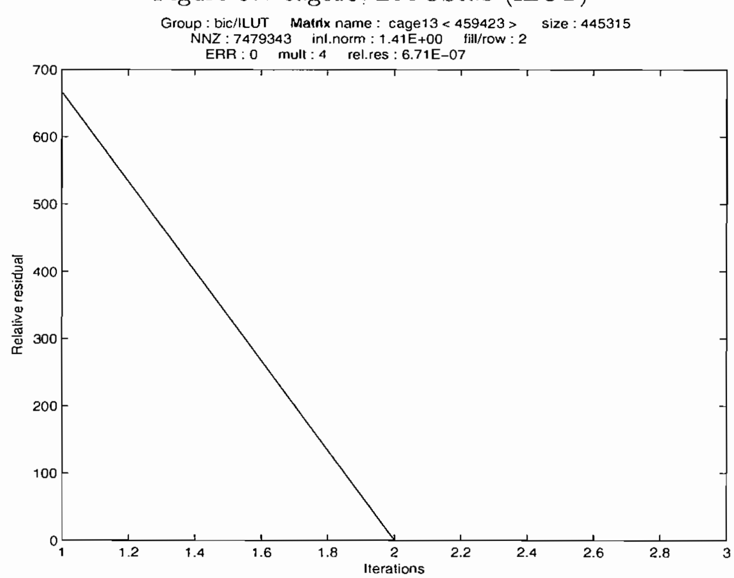

17. Matrix 63

Figure 68: cage14. CGNR (No)

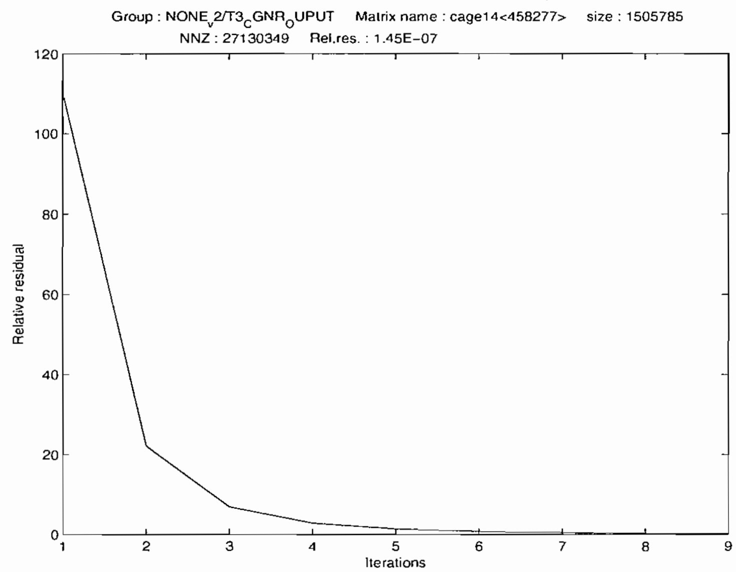

Figure 69: cage14. CGNR (ILUT)

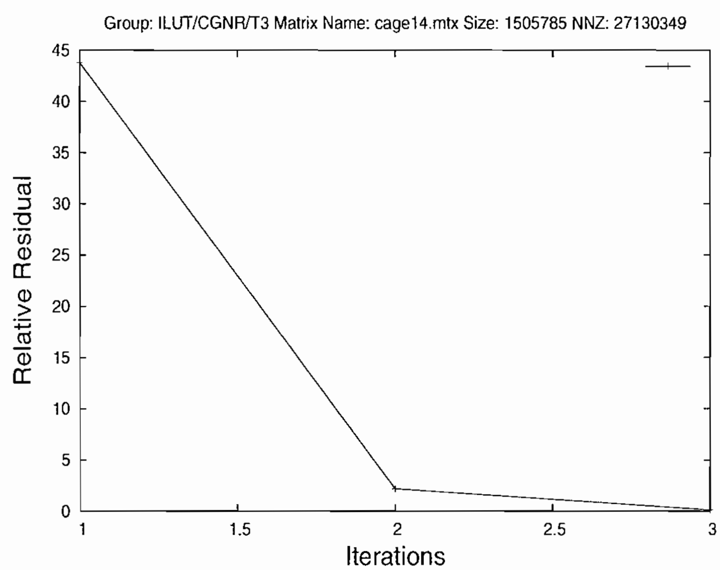

Figure 70: cage14, BiCGStab (ILUT)

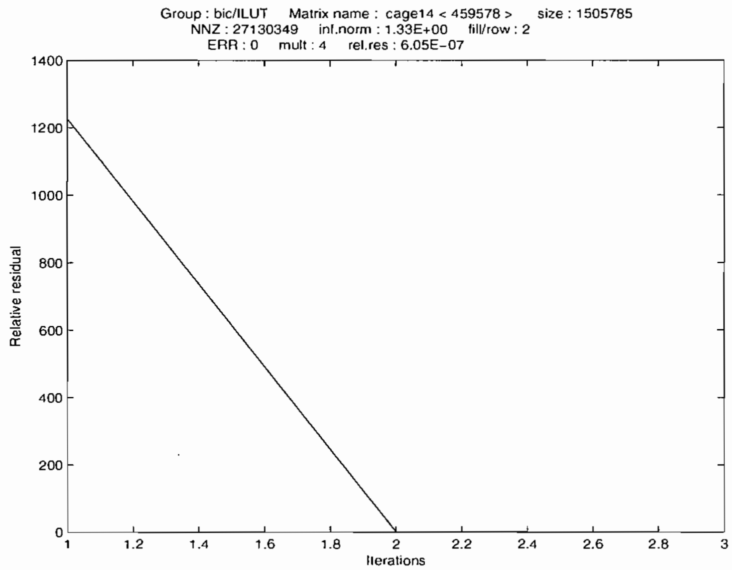

The three typical failures are the following.

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 77% of the cases. Many of them are also not amenable to parallelism, relying only on the parallelization of matrix vector multiplication. It is very important for these reasons to search for more robust (parallel) 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] G. H. Golub and C. F. Van Loan, "Matrix Computations", Third Edition, The John Hopkins University Press, 1996.