The First Annual Large Dense Linear System Survey

Alan Edelman
Department of Mathematics
University of California
Berkeley, California 94720

Abstract

In the March 24, 1991 issue of NA Digest, I submitted a questionnaire asking who was solving large dense linear systems of equations. Based on the responses, nearly all large dense linear systems today arise from either the benchmarking of supercomputers or applications involving the influence of a two dimensional boundary on three dimensional space. Not surprisingly, the area of computational aerodynamics or aero-electromechanics represents an important commercial application requiring the solution of such systems. The largest unstructured matrix that has been factored using Gaussian Elimination was a complex matrix of size 55,296. The largest dense matrix solved on a Sun using an iterative method was a real matrix of size 20,000. It is unclear at this time whether dense methods are truly needed at all for huge matrices. It is intended to survey users every year with the hope of including more applications as I am made aware of them.

1 Introduction

The idea to poll solvers of large dense linear system began when I delivered lectures on random matrices, and was asked without fail whether anyone in practice was solving large dense linear systems. Though large sparse systems solving is far more common, there are a number of applications that do indeed require the solution of large dense systems. This is an informal report of the answers that I received.

After a preliminary version of this survey over a year ago, it became clear that a survey of this style can serve a number of valuable purposes:

1. to inform software library developers what kinds of software might be needed
2. similarly, to inform researchers in supercomputing and numerical analysis what kinds of hardware and algorithms are needed
3. to communicate ideas among researchers who have little else in common other than that they are solving large dense linear systems

4. technology transfer to industry
5. more personally, to satisfy my own curiosity (and those of my colleagues, students, and perhaps even of history.)

Every basic numerical methods course discusses the solution of linear equations, but there is little motivation based on "cutting-edge" applications.

The key pieces of information that I asked for appear below. (The actual request as it appeared in the NA Digest is given in Appendix A.)

1. Size of largest matrix solved
2. Solution method: (LU or iterative?)
3. Time for solving
4. Machine used
5. Source of matrix
6. Did the solver come from a package?
7. How accurate was the solution?

The quality of a survey of this type most obviously depends on 1) the quality of the responses received and 2) how well the questionnaire gets distributed. Regarding the latter point, it is impossible for me to know the fraction of the community of large dense equation solvers I have reached, but I did receive over 50 pieces of electronic (and ordinary) mail including 15 of a very detailed nature. Of these, the breakup was

1. 6 universities
2. 4 aircraft industry
3. 2 research labs
4. 3 software/hardware industry

It is hoped with time, however, as more and more people receive the NA digest, or become aware of this activity, this survey could only improve. It is already clear that I must be more precise when using words such as "large" (maybe should mean greater than 10,000), "dense" (should I exclude (block) Toeplitz matrices and least squares problems?), and "time" (should be broken up into setup time and solve time).

One final question that is difficult to formulate would examine the trickle down effect of computing power. In the immediate future access to the most computing power will be in the hands of relatively few. Thus when I ask about the largest matrices that are being solved, I am by necessity restricting my audience. One can legitimately wonder whether many more people would desire to solve systems of equations of size 50,000 using dense methods if the solution could take minutes on a machine that for economic and logistical reasons was readily accessible to many users. Presumably this will be the situation some years down the road.

2 Highlights

All truly dense linear systems arising from scientific applications come from the solutions of **boundary integral equations.** Boundary integral equations are integral equations defined on the boundary of a region of interest. All examples of practical interest compute some intermediate quantity on a two dimensional boundary, and then use this information to compute the final desired quantity in three dimensional space. The price that one pays for replacing three dimensions with two is that what started out as a sparse problem in $O(n^3)$ variables is replaced by a dense problem in $O(n^2)$.

Not all respondents specified a clear application, but the ones who did mentioned

- stealth airplane technology
- airflow past an airplane wing in commercial aircraft
- supercomputer benchmarking
- flow around ships and other off-shore constructions
- diffusion of solid bodies in a liquid (block Toeplitz)
- noise reduction (block Toeplitz)
- diffusion of light through small particles (block Toeplitz)
- tomography (sparse least squares)

In the case of stealth airplane technology, the underlying differential equation is the Helmholtz equation, and the boundary integral solution is known as the **method of moments** [7, 10]. In the case of fluid flow, the problem is often Laplace's equation or Poisson's equation, and the boundary integral solution is known as the **panel method**[8, 9] Generally these methods are called **boundary element methods**.

The panel method gets it name from the quadrilaterals that discretize and approximate a structure such as an airplane. In fact, the airplane model was the first application of these methods. Later, it was realized that these methods could more generally be used to solve differential equations, and thus the general name of boundary element methods. For the case of an airplane, it is not sufficient to merely panel the plane itself, also one must panel the wake behind the plane. So long as the speeds are not too high, these methods will model the flow of air in three dimensional space with a large degree of success. The same ideas have been attempted on automobiles with the observation that agreement between computed quantities and numerical quantities was good everywhere but the back of the car.

This method always produces a dense linear system of size $O(N)$ by $O(N)$ where N is the number of boundary points or panels that are being used. It is not unusual to see size $3N$ by $3N$, because of three physical quantities of interest at every boundary element. Each entry of the matrix is computed as an interaction of two boundary elements, often requiring the computations of many integrals. In many instances, the time required to compute the matrix is considerably larger than the time for solution.

Only the builders of stealth technology who are interested in radar cross-sections are considering using direct Gaussian elimination methods for solving their system. Their systems are always symmetric and complex, but not Hermitian. The two large solutions of which I am aware is a system that was 55,296 by 55,296 which required 4.4 days to solve in 64 bit precision on a Connection Machine CM2, and a system of around size 40,000 by 40,000 where the time was not specified. In the first case, the solution was known to be good, but in the second case the author writes

> By and large, we don't know how good our answers are. They seem to be good enough for what we're doing, and certainly better than the traditional methods of antenna engineering.

However, such computing power is by no means widely available. Consider the contrasting statement published in 1991 [10] p.16 that

> ...such a large matrix ...is often too large for ordinary mainframe computers, whose central memory cannot handle a matrix of more than, say, 240×241 elements.

In any event, it seems likely that the world has not yet seen the solution of a system of equations of size 100,000 by 100,000 by direct methods, and perhaps more interestingly, it is not clear anyone will ever need to solve such a system using direct methods. However, computing power is currently such that if it were deemed important enough, it could certainly be done.

3 Some Details

The questionnaire as it appeared in the March 24, 1991 issue of NA Digest appears in Appendix 1. I received fifteen replies that contained a great wealth of information as well as a larger number of smaller comments. Some authors included more details than others. None of the authors requested anonymity.

John Richardson at Thinking Machines Corporation in conjunction with Lockheed Corporation solved the double precision complex system of size 55,296 in 4.4 days on a Connection Machine CM-2 with 1024 Weitek (floating point) processors. The matrices arise from moment methods and the solution method was an out of core LU factorization. In evaluation tests using matrices with known solutions, the accuracy was good.

Francis Canning of Rockwell solved a complex symmetric (non-Hermitian) matrix on a Cray 2 in order to solve the Helmholtz equation in two dimensions generated from Moment Methods. The program required over 7 hours to generate the matrix and preconditioner, while only 100 seconds were required to solve each right hand side. Canning [3, 4, 5] found a transformation of his matrix which could reasonably be approximated by a sparse matrix of the same size. Then incomplete LU preconditioning was used along with orthomin to actually solve the problem. The accuracy was reported to be clearly good since it agrees with the high frequency asymptotic results.

Yoshiki Seo at NEC in Kanagawa, Japan has LU factored a system of size 32,768 in under 11 hours on an experimental HPP-LHS Supercomputer developed under Japan's MITI project. His purpose was evaluating a 4 processor computer. He comments that the accuracy seemed okay.

Woody Lichtenstein of Thinking Machines Corporation benchmarked a double precision system of size 28,672 in core in under half an hour on a CM-200. His matrix was uniformly distributed in $(-2, 2)$ with 100 added to a random permutation of the diagonal. He took the right hand side to be the sum of the columns and found results that were close to all ones.

Anne Greenbaum at the Courant Institute of NYU solved an integral equation for Laplace's equation on a complicated 2D multiply-connected domain. Her system was of size 20,000 [6]. The inovation exhibited here was the use of a fast multipole method to compute the matrix-vector products that were needed for her GMRES solver. The computation required 54 minutes on a Sun Sparcstation. She set up a problem with known solution for a continuous case and obtained a small residual for the discrete problem with good agreement to the continuous case.

Brian Whitney of FPS Computing in Beaverton, Oregon benchmarked systems of size 20,000 on an FPS M64/145 with 15 MAX boards. The matrices were random with a normal distribution, and the solution required more than a day. He chose a solution, then multiplied by the matrix, and then solved. He reported good accuracy.

A. Yeremin at the USSR Academy of Sciences solved a system of size 12,088 arising from panel methods for calculating transonic lifting potential flow. He used GMRES with preconditioner. He reported good accuracy after ten iterations.

Ken Atkinson at the University of Iowa has been developing methods for the solutions of boundary integral equations, though the matrices he has solved were not huge (under 10,000). [1] B. Alpert at Lawrence Berkeley Laboratories, Per Lotstdedt at the Aircraft Division of SAAB-Scania in Sweden, and Hans Munthe-Kaas who was working for the Norwegian Hydrodynamic laboratories in Trondheim (now at the University of Bergen, Norway) were all working on panel methods (or something similar) and they also solved systems of size smaller than 10,000. Also Richard Lehoucq at IMSL in Houston reported results on benchmarking matrices of size 1000.

Perhaps a bit off my intended subject, but no less interesting, Linda Kaufman at Bell Labs in Murray Hill, New Jersey was solving a least squares problem arising from tomography of size 12,000 by 16,0000 using conjugate gradient on $A^T A$. She reported that 20 iterations were required to obtain a good picture. The solution took 10 seconds on a Cray machine and 5 minutes on a Silicon Graphics Machine. Julia Olkin solved a block Toeplitz matrix of size 28,576 using preconditioned conjugate gradient. Olkin's problem required 28 minutes on a Multiflow Computer. A larger block Toeplitz matrix, (size 375,000) was solved by Goodman, Draine, and Flatau at the Princeton Observatory on a Convex in a day. Rudnei Dias da Cunha at the Computing Lab of the University of Kent at Canterbury reported results of a parallel triangular solver on T800 Transputers that required 220 seconds for a system of size 32,768.

4 Afterthoughts

I was a bit surprised that nobody is currently using condition estimators to evaluate accuracy, despite the importance that we stress on conditioning to our students of Numerical analysis. A quick "grep" of my survey responses showed me that the word "condition" was only found inside the word "preconditioned." It seems that the accuracy of software for solving large systems of equations is always being evaluated on the basis of tests on systems with known or approximately known solutions.

Perhaps a bit less surprising, but interesting nonetheless, is that the scientific and engineering calculations are most effectively being solved via iterative methods, while supercomputer experts are devoting time to direct methods. I presume the full benefits of these efforts may not be seen until what we currently call large dense matrices will actually be small dense matrices that will arise from really large sparse matrices.

At this time, the number of people reporting system solving of equations in the 20,000–50,000 range is quite small. I am looking forward to seeing how this will change with time.

A Text of Questionnaire

THE FIRST ANNUAL LARGE DENSE LINEAR SYSTEM SURVEY

Without realizing it, about a year ago, I initiated the 0th annual large dense linear system survey here in NANET. I've had so many requests for a repeat survey that I decided to formalize the process by making it a yearly event. (My calendar file should remind me to repeat this next year.) I understand the NANET list has grown considerably since last time, so this survey should reach many more people. By default, none of the information you supply will be anonymous, however I will keep any information strictly confidential upon request. All of these questions relate to large DENSE linear systems. Feel free to interject any comments between the lines, etc. Results will be tallied into a LaTeXpaper and will be available by anonymous FTP from math.berkeley.edu.

```
Name _________________
Address _______________
        _______________
        _______________
        _______________
Type of Institution  ___ University        ___ Independent Research Lab
                     ___ Aircraft industry ___ Supercomputer Manufacturer
                     ___ Other

Largest matrix size that you solved  n=_________
Length of time   ___________ (seconds, hours, weeks, ...)
on which machine ________________

Matrix was generated from ___ Moment Methods
                          ___ Panel Methods for Lifting Potential Flow
                          ___ Panel Methods for Potential Flow
                          ___ Other

Solution method used was  ____ LU factorization
                          ____ An iterative Method (Please specify___________)
            (If an iterative method was used, did you take advantage of
             symmetry, diagonal dominance, or any property at all?)

The solver was ___ your own
               ___ from a package (Please specify____________________)

The accuracy of the solution obtained
               ____ was clearly good (Specify how you know ____________________)
               ____ seems okay, but you are not really sure
               ____ is unknown
```

Any other comments, suggested questions for next year, etc?

I'm aware that aircraft manufacturers and supercomputing companies would be most interested in these results, but might be reluctant to reveal their own secrets. I would like to urge such manufacturers to feel free to mail me anonymous responses by surface mail even without return addresses and names. I will guarantee anonymity in any case upon request. Everyone will so benefit.

I trust that the academics out there who are doing this will be more than happy to be forthcoming.

References

[1] K. Atkinson, A survey of boundary integral equation methods for the numerical solution of Laplace's equation in three dimensions, in *Numerical Solutions of Integral Equations*, ed. by M. Golberg, Plenum Press, New York, 1990.

[2] C.A. Brebbia, ed., *Topics in Boundary Element Research, Volume 1: Basic Principles and Applications*, Springer-Verlag, New York, 1984.

[3] F.X. Canning, Sparse approximation for solving integral equations with oscillatory kernels, to appear, *SIAM J. Sci. Stat. Comp.*

[4] F.X. Canning, The impedance matrix localization (IML) method for moment-method calculations, *IEEE Antennas and Propagation Magazine* (October 1990).

[5] F.X. Canning, Sparse matrix approximations to an integral equation of scattering, *Communications in applied numerical methods, 6* (1990), 543–548.

[6] A. Greenbaum, L. Greengard, and G.B. McFadden, Laplace's equation and the Dirichlet-Neumann map in multiply connected domains, *preprint*.

[7] R. Harrington, Origin and development of the method of moments for field computation, *IEEE Antennas and Propagation Magazine* (June 1990).

[8] J.L. Hess, Panel methods in computational fluid dynamics, *Annual Rev. Fluid Mechanics, 22* (1990), 255-274.

[9] J.L. Hess and M.O. Smith, Calculation of potential flows about arbitrary bodies, in *Progress in Aeronautical Sciences, Volume 8*, ed. by D. Küchemann, Pergamon Press, London, 1967.

[10] J.J.H. Wang, *Generalized Moment Methods in Electromagnetics*, John Wiley & Sons, New York, 1991.