Naval Research Laboratory

Washington, DC 20375-5320

NRL/FR/7130--04-10,072

An Improved Method for Solving Systems of Linear Equations in Frequency Response Problems

MARTIN H. MARCUS

Physical Acoustics Branch
Acoustics Division

April 15, 2004

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*
15-04-2004

2. REPORT TYPE
Interim

3. DATES COVERED *(From - To)*

4. TITLE AND SUBTITLE

An Improved Method for Solving Systems of Linear Equations in Frequency Response Problems

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER
61153N

6. AUTHOR(S)

Martin H. Marcus

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Naval Research Laboratory
4555 Overlook Avenue, SW
Washington, DC 20375-5320

8. PERFORMING ORGANIZATION REPORT NUMBER

NRL/FR/7130--04-10,072

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Office of Naval Research
800 N. Quincy Street
Arlington, VA 22217-5660

10. SPONSOR / MONITOR'S ACRONYM(S)

11. SPONSOR / MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution is unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

Vibration and acoustics problems that require solving matrix equations for many frequencies have solutions that vary very little from one frequency to the next. This suggests that an iterative method using the solution at one frequency could be used to obtain the next frequency solution. In this report, the matrix inverse at the first frequency serves as the preconditioner for the preconditioned conjugate gradient method to obtain subsequent frequencies. For the problem examined in this report, in which the matrix comes from the finite-element analysis of a cube, the iterative method is more than an order of magnitude faster than Gaussian elimination.

15. SUBJECT TERMS

Numerical Methods

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			Martin H. Marcus
Unclassified	Unclassified	Unclassified	SAR	10	19b. TELEPHONE NUMBER *(include area code)* (202) 767-7217

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

CONTENTS

EXECUTIVE SUMMARY

Vibration and acoustics problems that require solving matrix equations for many frequencies have solutions that vary very little from one frequency to the next. This suggests that an iterative method using the solution at one frequency could be used to obtain the next frequency solution. In this report, the matrix inverse at the first frequency serves as the preconditioner for the preconditioned conjugate gradient method to obtain subsequent frequencies. For the problem examined in this report, in which the matrix comes from the finite-element analysis of a cube, the iterative method is more than an order of magnitude faster than Gaussian elimination.

AN IMPROVED METHOD FOR SOLVING SYSTEMS OF LINEAR EQUATIONS IN FREQUENCY RESPONSE PROBLEMS

INTRODUCTION

When solving vibration and acoustics problems, investigators often perform calculations in the frequency domain. Since Gaussian elimination is typically used [1], as many matrix equations must be solved as the number of frequencies chosen. It becomes very important to use the fastest numerical method possible. One choice is to compute the eigenvalues and eigenvectors of the system and obtain the coefficients of the eigenvectors for each frequency. This is done in NISA [2] and other commercial finite-element codes. This method has recently been enhanced with a new technique called Automated Multi-Level Substructuring [3-7], in which the modes of individual substructures are obtained and only some of these modes are used for the global solution. This method is very promising. One drawback in its versatility, however, is that it is currently applicable to in-vacuo structures and to structures in a light acoustic fluid [4]. This means that it can be used for air acoustics, but not for underwater acoustics. Also, Bathe [1] on page 336 notes that modal methods are not efficient for point loads or shock loading because this type of load requires many modes.

Another promising numerical method is FETI (finite-element tearing and interconnect) [8]. The problem is broken into many substructures that are solved with Gaussian elimination. The solutions are then combined into one large sparse matrix that is solved with an iterative technique. The Gaussian elimination can be solved in parallel, with one processor per substructure, but the parallel optimization of the sparse matrix cannot be as opportunistic. Nevertheless, FETI has demonstrated good efficiency in a parallel environment. FETI's efficiency is important to note because the method discussed in this report can speed up the Gaussian elimination part of FETI.

One conclusion that could be made from this brief look at other matrix methods is that if the analyst has just one processor, a method faster than Gaussian elimination would be useful. If one has many processors, a method to enhance the Gaussian elimination part of FETI would also be useful. It is now appropriate to discuss what enhancements can be made.

In a typical frequency response analysis, the solutions from one frequency to the next look quite similar to each other. It would be opportunistic to take advantage of this fact by using the solution of the matrix at one frequency to speed up the computation of the solution at the next frequency. Matrix iterative methods [1, 9-17] can take advantage of a good initial guess. Of the matrix iterative methods, the preconditioned conjugate gradient method [9] appears to be the most efficient. It allows for an initial guess of the solution, as well as a guess of the matrix inverse. This report shows that using the solution and the matrix inverse from the previous frequency as the guesses performs much better than does Gaussian elimination for some problems with large bandwidth.

REVIEW OF MATRIX ITERATIVE METHODS

In Jacobi's method [16-17], a solution is assumed and each equation is solved individually to provide improved values. These improved values are not used until every equation in the matrix is solved, in other

Manuscript approved July 8, 2003.

words, when one iteration is completed. In contrast, the Gauss Seidel method uses these values as soon as they are available. This allows for faster convergence. Yet another improvement can be made by using the successive overrelaxation method. Here, the error in the current solution of each equation is multiplied by the relaxation parameter. This is added to the current solution to make the new solution. If the relaxation parameter is one, then this method is identical to the Gauss Seidel method. Typically, the parameter is greater than one, which is why the term overrelaxation is used. By computing the relaxation parameter from an estimate of the spectral radius for this method, convergence is greatly superior to that of the Gauss Seidel method.

In the method of steepest descent [9], each iteration seeks to minimize a function of the residual (error). Unfortunately, the solution converges very slowly, because each iteration tends to resemble previous iterations. In contrast, with the conjugate gradient method, each residual is chosen to be orthogonal to the previous ones. By so doing, the method is guaranteed to converge in as many steps as the size of matrix. Whether or not it takes this long depends on the condition number of the matrix. An ill-conditioned matrix leads to very slow convergence. The way around this problem is to precondition the matrix (preconditioned conjugate gradient). Since the best conditioned matrix is the identity matrix, the best preconditioner is the matrix inverse. Such a choice leads to convergence in just one iteration. Obviously, we typically do not know the inverse of the matrix that we are inverting, so the inverse of the matrix diagonal is often selected as the preconditioner. It is also practical to invert the three middle diagonals of the matrix.

It is important to contrast the speed of iterative methods vs Gaussian elimination. Since one iteration requires that a matrix be multiplied by a vector, if the matrix is sparse, the iteration can be performed quite quickly. In contrast, Gaussian elimination causes many of the zeros in a sparse matrix to become nonzero, which makes extra work, so iterative methods are preferred for sparse matrices. For banded matrices, which are typical of the finite-element method, Gaussian elimination is usually superior [1]. If one defines the matrix size as n and the half bandwidth as b, then the number of multiplies for Gaussian elimination is

$$g = nb^2. \tag{1}$$

For many iterative methods, including preconditioned conjugate gradient, the number of multiplies is

$$p = 2bnm, \tag{2}$$

where m is the number of iterations for convergence. Combining Eqs. (1) and (2), we get the number of iterations where the two methods are equal as

$$m = b/2. \tag{3}$$

For a tightly banded matrix, such as from a beam in bending in a finite-element problem, the number of iterations would have to be extremely low for the iterative method to be competitive. In contrast, for a shell with complicated internal structure or for a solid, the bandwidth would be much higher and an iterative method could be allowed to have more iterations before being inefficient. Even under these circumstances, the above iterative methods converge too slowly. This necessitates the improvement discussed in the next section.

IMPROVEMENT TO PRECONDITIONED CONJUGATE GRADIENT METHOD FOR FREQUENCY RESPONSE PROBLEMS

For frequency response problems, the investigator typically selects a sufficiently refined frequency spacing such that the results vary smoothly from frequency to frequency. The matrix inverse and solution for one frequency look very much like the inverse and solution for the next frequency. Consequently, the solution for one frequency makes a very good initial guess in an iterative method seeking the solution for the next frequency. Similarly, the matrix inverse for one frequency makes a very good preconditioner for the conjugate gradient method applied to the next frequency.

The way this would work in practice is that the first frequency would be solved via Gaussian elimination. The solution and the LU decomposition of the matrix are saved. The second frequency is solved using the solution to the first frequency as the initial guess in the conjugate gradient method, while the LU decomposition serves as the preconditioner. The solution to the second frequency is the initial guess for the third frequency, but because the conjugate gradient method does not provide an inverse, the LU decomposition from the first frequency must be used again. After several frequencies, the LU decomposition becomes a poor preconditioner and the computer time for the iterative method takes longer than it would for Gaussian elimination. At this point, Gaussian elimination is performed for the next frequency and the new LU decomposition can serve as the preconditioner. In practice, a new LU decomposition is needed only if the mesh is too coarse for the frequency range being considered.

This method should work for any matrix inversion problem in which many runs must be made with one parameter that changes very gradually. An example might be the static deformation of a structure whose elastic modulus varies in increments of 0.01 from 1 to 2. Since frequency response problems are much more popular, the next section demonstrates this method on the frequency response of a shear load on a cube.

DEMONSTRATION PROBLEM

In this section, a uniform shear load is applied at one face of a cube that has all of its degrees of freedom constrained at the opposite face. It is solved via the finite-element method with linear hexahedron elements. The cube measures 8 m per side and has a density of 8 kg/m^3. The Poisson's ratio is 0.3 and the elastic modulus is 10,000 Pa. It is solved for frequencies of 0.1 through 9.2 Hz in increments of 0.1 Hz.

The mesh is found to converge with 12 elements per side (1728 elements altogether). Its highest eigenvalue in the frequency range of interest is within 0.72% of the highest eigenvalue for a mesh with 16 elements per side.

The problem is solved by using Gaussian elimination and the preconditioned conjugate gradient in which the initial guess comes from the previous frequency and the preconditioner comes from the last time Gaussian elimination was used. The latter will henceforth be called improved preconditioned conjugate gradient (ipcg). Its convergence criterion is when the rms average of the difference between the solution vector and the solution from the previous iteration is less than 10^{-6} of the maximum value in the previous solution. It is separately checked to ensure that it converged with the Gaussian elimination solution.

The ipcg algorithm computes the 92 solutions at 13.3 times the speed of Gaussian elimination. To examine the performance more closely, Fig.1 shows the breakdown by frequency. Figure 1(a) displays the number of multiplies vs frequency. Using Gaussian elimination (the dashed curve), about 1.9 billion multiplies are needed for each frequency. With ipcg, the first frequency is solved with Gaussian elimination, then the matrix inverse is used on the next 91 frequencies. The last of these frequencies requires the most multiplies, but far fewer than Gaussian elimination would take. Solutions to more frequencies could be computed until ipcg becomes less efficient than Gaussian elimination. Other computations, however, reveal that before one gets to this frequency, the mesh is too coarse. One need only ensure that the mesh is sufficiently refined at the highest frequency to know that only one Gaussian elimination will be needed.

Figure 1(b) shows the distribution of eigenvalues in the frequency range of interest. In general, as the frequency increases, the eigenvalue density goes up. From Fig.1(a), we observe that the computations per

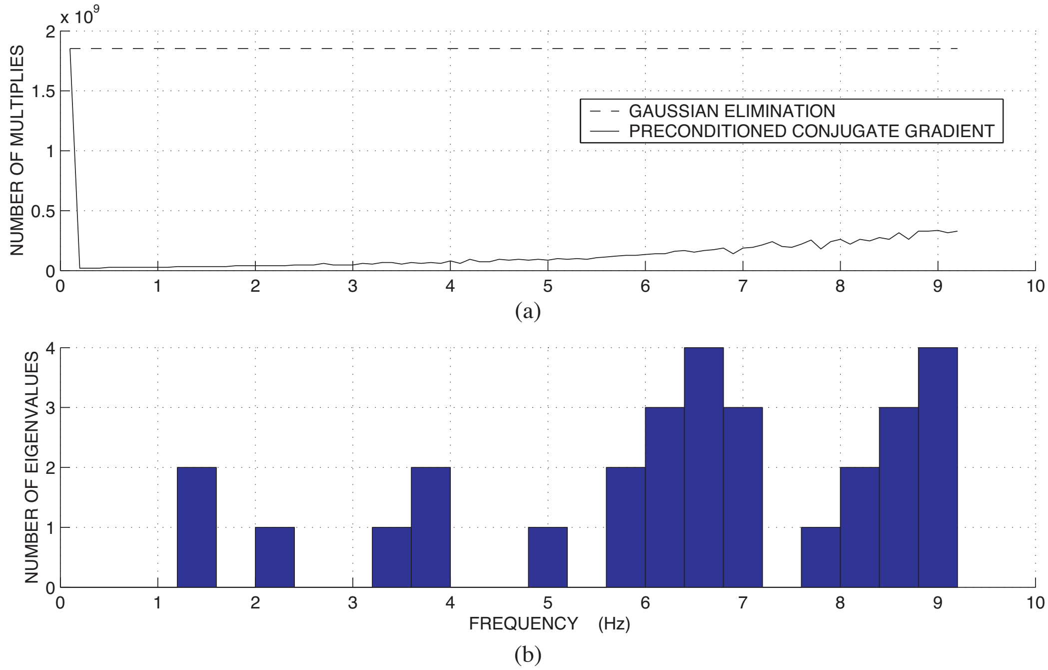

Fig. 1 — Performance of the matrix inversion algorithms vs frequency for a shear load on a 1728-element finite-element model of a cube. (a) Number of multiplies at each frequency for Gaussian elimination (dashed) and improved preconditioned conjugate gradient (solid). (b) Histogram of the matrix eigenvalues.

frequency also go up. This suggests that the iterative method performs more poorly when there are many eigenvalues, but for every frequency and for several other ipcg runs made by the author, the method converged to the exact solution. The eigenvalues indicate that the matrix is not positive definite, but the conjugate gradient method is guaranteed to converge only when the matrix is positive definite [9]. This method in this situation, however, is known to fail very rarely [18]. Adaptations to the conjugate gradient method can be implemented [19-20] so that it will always converge for indefinite matrices.

Although the demonstration problem does not represent an exhaustive proof that ipcg outperforms Gaussian elimination, it shows more than an order of magnitude improvement in speed for a problem with 6084 degrees of freedom and a frequency range containing many resonances.

SUMMARY

Vibration and acoustics problems that require solving matrix equations for many frequencies have solutions that vary very little from one frequency to the next. This suggests that an iterative method using the solution at one frequency as the initial guess for the next frequency may show promise. For problems in which Gaussian elimination outperforms existing iterative methods, a well-chosen initial guess is not sufficient for an iterative method to be faster. However, the matrix inverse at a nearby frequency serves as an excellent preconditioner for the preconditioned conjugate gradient method. Specifically, one can perform Gaussian elimination for the first frequency and obtain a matrix inverse (as an LU decomposition) and solution. For the next frequency, the solution just obtained is the initial guess in the preconditioned conjugate gradient method. The LU decomposition for the first frequency serves as the preconditioner. For subsequent frequencies, the solution from the frequency just solved is the initial guess, but the LU decomposition from the first frequency is used.

For matrices that are banded but not sparse, an iterative method is faster than Gaussian elimination only if the number of iterations is less than half of the half bandwidth. In general, this means that problems worth

considering for the above iterative method must have a large bandwidth (for example, the finite-element analysis of a shell with complicated internal structure or the analysis of a solid). For problems with a sufficiently high bandwidth, the iterative method outperforms Gaussian elimination. In this report, a finite-element analysis of a cube is performed using Gaussian elimination and the above iterative method. For this problem, the iterative method is shown to be 13 times faster than Gaussian elimination.

ACKNOWLEDGMENTS

The author gratefully acknowledges the Office of Naval Research for support of this work.

REFERENCES

1. K. Bathe and E. Wilson, *Numerical Methods in Finite Element Analysis* (Prentice Hall, Englewood Cliffs, NJ, 1976) p. 291.

2. K.S. Kothawala, *Users Manual for NISA II, Version 7.0* (Engineering Mechanics Research Corporation, Troy, MI, 1997).

3. J.K. Bennighof and M.F. Kaplan, "Frequency Sweep Analysis Using Multilevel Substructuring, Global Modes and Iteration," *Proceedings of 39th AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 1998 (http://www.ae.utexas.edu/research/amls/sweep.ps).

4. J.K. Bennighof, "Vibroacoustic Frequency Sweep Analysis Using Automated Multi-level Substructuring," *Proceedings of the AIAA Modeling and Simulation Technologies Conference and Exhibit*, 1999 (http://www.ae.utexas.edu/research/amls/acous.ps).

5. J.K. Bennighof, M.F. Kaplan, and M.B. Muller, "Extending the Frequency Response Capabilities of Automated Multi-level Substructuring," *Proceedings of 41st AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference*, 2000 (http://ae.utexas.edu/research/amls/extend.ps).

6. J.K. Bennighof, M.F. Kaplan, M.B. Muller, and M. Kim, "Meeting the nvh Computational Challenge: Automated Multi-level Substructuring," *Proceedings of International Modal Analysis Conference*, 2000 (http://www.ae.utexas.edu/research/amls/imac.ps).

7. J.K. Bennighof and M.F. Kaplan, "Frequency Window Implementation of Adaptive Multi-level Substructuring," *J. Vib. Acoust.* **120**, 409-418 (1998).

8. R. Tezaur, A. Macedo, and C. Farhat, "Iterative Solution of Large-scale Acoustic Scattering Problems with Multiple Right-hand Sides by a Domain Decomposition Method with Lagrange Multipliers," *Int. J. Num. Meth. Eng.* **5**, 1175-1193 (2001).

9. G.H. Golub and C.F. Van Loan, *Matrix Computations* (Johns Hopkins University Press, Baltimore, MD, 1983) pp. 352-379.

10. W.H. Press, *Numerical Recipes* (Cambridge University Press, NY, 1986) pp. 301-307.

11. A. Bjorck, *Large Scale Matrix Problems* (North Holland, NY, 1981) pp. 85-194.

12. R.S. Varga, *Matrix Iterative Analysis* (Springer, NY, 2000).

13. C.W. Ueberhuber, *Numerical Computation 2* (Springer, NY, 1997).

14. L.A. Hageman and D.M. Young, *Applied Iterative Methods* (Academic Press, NY, 1981).

15. D.M. Young, *Iterative Solution of Large Linear Systems* (Academic Press, NY, 1971).

16. G. Dahlquist and A. Bjorck, *Numerical Methods* (Prentice Hall, Englewood Cliffs, NJ, 1974), pp. 188-196.

17. K.E. Atkinson, *An Introduction to Numerical Methods* (Wiley, NY, 1978) pp. 471-495.

18. E. Chow and Y. Saad, "Experimental Study of ILU Preconditioners for Indefinite Matrices," *J. Comput. Appl. Math.* **86**(2), 387-414 (1997).

19. Y. Lai, W. Lin, and D. Pierce, "Conjugate Gradient and Minimal Residual Methods for Solving Symmetric Indefinite Systems," *J. Comput. Appl. Math.* **84**(2), 243-256 (1997).

20. P. Kettil, T. Ekevid, and N.E. Wiberg, "Towards Fully Mesh Adaptive fe Simulations in 3d Using Multigrid Solver," *Comput. Struct.* **81**(8-11), 735-746 (2003).