JOURNAL OF COMPUTATIONAL PHYSICS 89, 241–245 (1990)

Note

Davidson's Method and Preconditioning for Generalized Eigenvalue Problems

I. Introduction

The generalized eigenvalue problem

$$Az = \lambda Sz,$$

where A is symmetric and S is symmetric positive definite, arises in a number of different applications. In computational quantum chemistry, generalized eigenvalue problems occur when nonorthogonal basis functions are used [1]. We will look at methods for solution of large generalized eigenvalue problems and will give an approach that can yield rapid convergence.

The standard eigenvalue problem is $Az = \lambda z$. When A is a large symmetric matrix, the Lanczos algorithm [2] is one of the most popular ways of computing a few eigenvalues. Lanczos is essentially the Rayleigh–Ritz procedure used to extract approximate solutions from a Krylov subspace [2]. However, many quantum chemists have found that Davidson's method [3] gives much faster convergence for their problems.

Davidson's method also uses the Rayleigh–Ritz procedure, but the new trial vectors for the subspace are $(D-\theta I)^{-1}(A-\theta I)y$, where D is the diagonal of A, and where θ and y are the most recent approximations to the desired eigenvalue λ and eigenvector z. $D-\theta I$ can be viewed as a preconditioner [4, 5] to $A-\theta I$. The subspace is not Krylov, but it asymptotically resembles a Krylov subspace generated by $(D-\lambda I)^{-1}(A-\lambda I)$ [6]. Generally $(D-\lambda I)^{-1}(A-\lambda I)$ has a more favorable distribution of eigenvalues than does A. If the diagonal of A is large and also the differences in diagonal elements are fairly large relative to the size of the off-diagonal elements, then $D-\lambda I$ is a good approximation to $A-\lambda I$. In that situation, Davidson's method is usually very effective. In other problems a diagonal approximation to A is unsatisfactory, and it may be helpful to use a better approximation in place of D. The preconditioning is thus improved. This is called the GD (generalized Davidson's) method [6].

Davidson's method can also be applied to the generalized eigenvalue problem. The new trial vectors are $(D_1-\theta D_2)^{-1}(A-\theta S)y$, where D_1 and D_2 are the diagonals of A and S, respectively. Approximations are extracted from the subspace with the Rayleigh–Ritz procedure for the generalized eigenvalue problem [2].

However, Gallup [7] and others have reported poor results when using David-

0021-9991/90 $3.00

son's method on generalized eigenvalue problems. Gallop claims that Davidson's method is not appropriate for generalized eigenvalue problems, because $(S^{-1}A)^n y$ converges to the eigenvector of largest magnitude instead of $A^n y$. But this fact is not relevant. In fact simply $(A-\theta S)y$ can be used to generate the subspace [8]. Asymptotically the subspace will resemble the Krylov subspace generated by $A-\lambda S$. Since $A-\lambda S$ has z as an eigenvector, the method will be effective if the corresponding eigenvalue 0.0 is well separated from the rest of the spectrum. But there is no reason to expect a more favorable distribution of eigenvalues than that of the original generalized eigenvalue problem. Davidson's method can be effective for generalized problems because $(D_1-\lambda D_2)^{-1}(A-\lambda S)$ also has z as an eigenvector. But it is necessary that $D_1-\lambda D_2$ be a good approximation to $A-\lambda S$ so that the spectrum is improved by the preconditioning. The reason that Davidson's method was found ineffective is that generalized eigenvalue problems in quantum physics usually have strongly nondiagonal matrices [1].

Gallup [7] suggests factoring S and using a Krylov subspace generated by the operator $S^{-1}A$. But there is no preconditioning, so convergence will be slow if the eigenvalue distribution of the generalized problem is unfavorable. $(A-\sigma S)^{-1}S$, for σ near the desired eigenvalues, has an inverted spectrum and will give much faster convergence for some problems [9]. However factoring $A-\sigma S$ can be very expensive for large matrices.

We will give a method that avoids factoring either $A-\sigma S$ or S and yet can give rapid convergence. This method is a generalization of the GD method.

II. The GD Method for Generalized Eigenvalue Problems

The GD method [6] can be extended to generalized eigenvalue problems. Let M_1 be an approximation to A and M_2 be an approximation to S. Then choose the new trial vector to be

$$(M_1-\sigma M_2)^{-1}(A-\theta S)y.$$

$M_1-\sigma M_2$ is a preconditioner for $A-\theta S$. Here σ can be chosen to be the same as θ, except that it should not be changed at every step if factorization of $M_1-\sigma M_2$ is expensive. If it is desirable to factor only once, σ can be fixed to some value near the desired eigenvalues.

Suppose the pencil is of dimension n and the trial space is currently of dimension j. Then the reduced eigenvalue problem is

$$Q^{\mathrm{T}}AQg=\theta Q^{\mathrm{T}}SQg,$$

where Q is n by j with columns spanning the trial space. We denote the reduced problem by $H_1 g=\theta H_2 g$, where H_1 and H_2 are j by j matrices. The trial vectors can be orthogonalized with respect to the S inner product so that $H_2=I$. If instead Q is orthonormal in the regular sense, then a small generalized eigenvalue problem

must be solved. The appropriate θ must be chosen. The new approximate eigenvalue and eigenvector are θ and $y = Qg$.

There are many possible choices for the preconditioner. M_1 and M_2 can be block diagonal portions of A and S. Band portions can also be used. For sparse matrices, incomplete factorization [5] of $M_1 - \sigma M_2$ may be best. Care must be taken because incomplete factorization was developed for positive definite matrices. Generally there is a trade-off between the effectiveness of the preconditioning and the cost of factoring $M_1 - \sigma M_2$.

A drawback of the GD method is that one must have access to portions of the large matrices A and S when forming the preconditioner. The difficulty in this depends on how the matrices are generated and stored.

The method is effective because as θ converges to an eigenvalue λ with associated eigenvector z, the operator generating the trial space approaches $(M_1 - \sigma M_2)^{-1}(A - \lambda S)$. This matrix has z as an eigenvector with eigenvalue 0.0. Due to the preconditioning, there is a tendency for this eigenvalue to stand out in the spectrum. Thus the trial space rapidly develops an approximation to z. The Rayleigh–Ritz procedure is applied to the pencil (A, S) and extracts this eigenvector.

It is also possible to apply this method to nonsymmetric generalized eigenvalue problems. The trial vectors should probably be orthogonalized in the regular sense, because S will not define an inner product. Of course the reduced eigenvalue problem will be nonsymmetric and complex numbers must be dealt with [10].

III. Example

Gallup [7] gives a test example that is easy to generate. The model has N^2 atoms in a square lattice. The interaction between atoms decreases exponentially as the atoms get farther apart. With $N = 20$, this gives a pencil of dimension 400. The off-diagonal elements of A of largest magnitude are -0.49 and the diagonal elements range between -7.0 and 12.0. S has largest off-diagonal element of 0.35, and all of the diagonal elements are 1.0.

The three smallest eigenvalues, -15.9370, -12.9093, and -12.2770, and their associated eigenvectors are computed. A block method [11] is used with three new trial vectors added at each iteration. Three initial trial vectors are found by solving the leading 25 by 25 subpencil. The trial space is restarted with three approximate eigenvectors when the space reaches dimension 45. The restarting reduces orthogonalization and other Rayleigh–Ritz costs. The convergence test used is the norm of the residual vector going below 10^{-8}. This gives eigenvalues accurate to about 16 digits. The number of steps given is the total number of matrix-vector products with A required for convergence (i.e., the total number of trial vectors used). The number of iterations is one-third of that amount.

Generating the subspace with $(A - \theta S)y$ takes 276 steps. Davidson's method converges in 204 steps. For GD, two different preconditioners are used. First M_1

has its leading 100 by 100 submatrix the same as A's, and its main diagonal is the same as A's, and the rest are all zeroes. M_2 has the same nonzero portions from S. The value of σ is fixed at -14.0, so that only one factorization is needed. This version of GD requires 120 steps. Next M_1 and M_2 are block diagonal with four 100 by 100 blocks from A and S, respectively. Convergence is in 45 steps. This is a big improvement over the diagonal preconditioning in Davidson's method. It should also be noted that the cost of factoring this particular $M_1 - \sigma M_2$ is only one-sixteenth of the cost of factoring either S or $A - \sigma S$.

We now reorder the atoms in the lattice and show that the preconditioning is favorably affected. The square lattice is divided into four square quarters. To generate the matrices, the atoms in one quarter are numbered first (by rows), then one by one each of the other quarters are numbered. Davidson's method again reaches convergence in 204 steps. GD with M_1 and M_2 each having one 100 by 100 block requires 66 steps. This is a considerable improvement compared to 120 before reordering. With four 100 by 100 blocks in the preconditioner, convergence comes in 36 steps. So the trial space is of dimension 36 and only 12 iterations of this block method are needed.

Of course, there are many other possible ways to order the atoms. And there are also many choices for preconditioner.

IV. Conclusion

The GD method can be much better than standard Davidson's method for generalized eigenvalue problems. If elements of the matrices are accessible and if there are nondiagonal approximations to the matrices that are better than diagonal, then GD is worth trying. In quantum chemistry, GD seems to have more potential for generalized eigenvalue problems than for standard problems. This is because of the non-diagonally dominant form of generalized eigenvalue problems. The construction and ordering of the matrices can make a difference in the effectiveness of the preconditioning.

References

1. E. R. Davidson, "Super-matrix Methods," Workshop on Practical Iterative Methods for Large Scale Computations, Minnesota Supercomputer Institute, Minneapolis, 1988; *Comput. Phys. Commun.* **53**, 49 (1989).
2. B. N. Parlett, *The Symmetric Eigenvalue Problem* (Prentice-Hall, Englewood Cliffs, NJ, 1980).
3. E. R. Davidson, *J. Comput. Phys.* **17**, 87 (1975).
4. P. Concus, G. H. Golub, and G. Meurant, *SIAM J. Sci. Stat. Comput.* **6**, 220 (1985).
5. J. A. Meijerink and H. A. van der Vorst, *Math. Comput.* **31**, 148 (1977).
6. R. B. Morgan and D. S. Scott, *SIAM J. Sci. Stat. Comput.* **7**, 817 (1986).
7. G. Gallup, *Comput. Chem.* **3**, 127 (1982).

8. T. Z. KALAMBOUKIS, *J. Phys. A* **13**, 57 (1980).
9. T. ERICSSON AND A. RUHE, *Math. Comput.* **35**, 1251 (1980).
10. R. B. MORGAN, "Generalizations of Davidson's method for Computing Eigenvalues of Large Nonsymmetric Matrices," submitted (unpublished).
11. B. LIU, in *Numerical Algorithms in Chemistry: Algebraic Methods*, edited by C. Moler and I. Shavitt (LBL-8158, Lawrence Berkeley Lab., 1978), p. 49 (unpublished).

RECEIVED: February 14, 1989; REVISED: June 12, 1989

RONALD B. MORGAN

Department of Mathematics
University of Missouri
Columbia, Missouri 65211