A new method for diagonalising large matrices

1985 J. Phys. A: Math. Gen. 18 1343

(http://iopscience.iop.org/0305-4470/18/9/018)

J. Phys. A: Math. Gen. **18** (1985) 1343-1359. Printed in Great Britain

A new method for diagonalising large matrices

D M Wood and Alex Zunger

Solar Energy Research Institute, Golden, Colorado 80401, USA

Received 7 March 1984, in final form 14 November 1984

Abstract. The structure and implementation of a new general iterative method for diagonalising large matrices (the 'residual minimisation/direct inversion in the iterative subspace' method of Bendt and Zunger) are described and contrasted with other more commonly used iterative techniques. The method requires the direct diagonalisation of only a small submatrix, does not require the storage of the large matrix and provides eigensolutions to within a prescribed precision in a rapidly convergent iterative procedure. Numerical results for two rather different matrices (a real 50×50 non-diagonally dominant matrix and a complex Hermitian 181×181 matrix corresponding to the pseudopotential band structure of a semiconductor in a plane wave basis set) are used to compare the new method with the competing methods. The new method converges quickly and should be the most efficient for very large matrices in terms both of computation time and central storage requirements; it is quite insensitive to the properties of the matrices used. This technique makes possible efficient solution of a variety of quantum mechanical matrix problems where large basis set expansions are required.

1. Introduction

The matrix formulation of quantum mechanics permits essentially any quantum mechanical problem to be reduced to the diagonalisation of matrices. In modern electronic structure calculations, for instance, the problem is of the so-called 'general Hermitian eigenproblem' form, i.e.

$$\boldsymbol{H}|a\rangle = \lambda \boldsymbol{S}|a\rangle \tag{1.1}$$

where $\boldsymbol{H}$ is the $N \times N$ (Hermitian) Hamiltonian matrix, $\boldsymbol{S}$ is the $N \times N$ overlap matrix (reflecting possible non-orthogonality of the basis set, i.e. $S_{ij} = \langle \phi_i | \phi_j \rangle$ where $\{|\phi_i\rangle\}$ is the set of basis functions) and $|a\rangle$ is one of the eigenvectors sought, with eigenvalue λ. For matrices of reasonable size ($N \leqslant 500$) this eigenproblem is usually solved by the Choleski-Householder (CH) procedure; the Choleski decomposition of the overlap matrix $\boldsymbol{S}$ (Wilkinson 1965) is used to reduce the problem to the form

$$\boldsymbol{H}'|a'\rangle = \lambda |a'\rangle \tag{1.2}$$

and the Householder method (Wilkinson 1965) is then used to find eigenvectors and eigenvalues of this problem. Finally, the eigenvectors are back-transformed using the inverse Choleski decomposition to find those of equation (1.1).

These are, however, a number of practical difficulties with this approach. First, the CH process needs random access to all of $\boldsymbol{H}$ and $\boldsymbol{S}$ (and often they must be stored in central memory as well). Second, the computation time for CH diagonalisation of an $N \times N$ matrix scales as N^3, with a coefficient roughly independent of the number

0305-4470/85/091343+17$02.25

of eigenvalues sought: even if only the lowest few eigensolutions are needed, the computational effort involved is close to that required for *all* eigensolutions.

There are many problems of physical interest for which such difficulties with the CH method can become acute. Many quantum mechanical Hamiltonian problems require only the lowest few eigensolutions to very large matrices whose elements $\langle\phi_i|\boldsymbol{H}|\phi_j\rangle$ are easy to calculate, but which possess no simple systematics (e.g. they are not diagonally dominant or sparse). Such is the case, for example, in ground-state electronic structure calculations for molecules, solids and surfaces, where our relative ignorance of what constitutes a physically motivated basis set $\{|\phi_i\rangle\}$ often leads to the need for very large bases of simple functions. This is also the case when reasonable basis functions for each particle are known, but when the number of such particles (or quasiparticles) is large, e.g. in spin Hamiltonian problems. Interestingly, many well known approximations in electronic structure theory have, in fact, evolved as a result of the difficulty with direct matrix diagonalisation by the CH method. For instance, an expansion in terms of a simple plane wave basis set constitutes a reasonable and practical description of systems in a weak periodic potential (i.e. only a few Fourier components of the potential are important). The need to limit the number of plane wave basis functions in order that the CH diagonalisation method be practical has spurred a variety of physical approximations which make the atomic (pseudo)potential artificially weak (e.g. Cohen 1970), but the physical validity of these remains uncertain (see the discussion in Zunger (1979) and Zunger and Cohen (1978, 1979)).

Similarly, the need to limit the dimension of the Hamiltonian matrix in molecular electronic structure studies has often led to laborious procedures for constructing a small set of physically efficient but unintuitive basis orbitals (e.g. Huzinaga 1963, Bishop 1967, ten Hoor 1980). As a final example, in many empirical tight binding methods for studying molecules and polymers, e.g. the Hückel method and its variants, the interaction between distant particles is set to zero so that the resulting Hamiltonian matrix will possess certain special properties, permitting matrix folding and partial block diagonalisation (e.g. Harrison 1980). Clearly, many of these approximations would become unnecessary if the difficulties with the CH diagonalisation method could be circumvented.

In response to these serious drawbacks in the usual Choleski-Householder method there have developed in the past three decades a number of *iterative* methods†, all of which have the twin virtues that (*a*) they require the matrices $\boldsymbol{H}$ and $\boldsymbol{S}$ at worst only one row at a time and (*b*) the time per iteration scales only as N^2. The tradeoff, however, is that one must iterate the procedure several (typically 2-20) times until a convergence criterion is met. It should be emphasised that these iterative techniques are intended to be useful when the computational bottleneck is the actual *diagonalisation* of the large matrices and not the *generation* of the Hamiltonian and overlap matrix elements themselves.

In § 2 of this paper we briefly review the philosophy of several well known iterative procedures for diagonalising large matrices and establish the notation to be used in § 3, where a new iterative method (Bendt and Zunger 1982b) is described in detail and contrasted with its predecessors. In § 4 two numerical examples are given and discussed, and in § 5 some suggestions for implementation of the new method are presented; there the true benefits of the new method, and when it is likely to be the method of preference, are discussed. An appendix contains the program structure.

† An excellent overview of most current iterative (and other) large matrix methods appears in NRCC (1978). A good brief review is Nesbet (1981).

2. Iterative methods for large matrix diagonalisation

2.1. Background

One is generally confronted with the problem of finding the lowest n eigenvalues and eigenvectors of the $N \times N$ general Hermitian eigenproblem (1.1). Typically, for ground-state properties of a physical system, the number of eigenvalues sought, n, is a small or very small fraction of N. In this section we discusss the features common to most iterative approaches to this problem. In what follows it will be assumed that one has identified a finite dimension N such that the matrices $\boldsymbol{H}$ and $\boldsymbol{S}$ describe adequately the physics of the problem at hand. One can then use the methods of finite-dimensional Hilbert spaces to proceed.

The first point of similarity among the different iterative methods for diagonalisation is that most define or use three distinct sets of auxiliary N-dimensional vectors. These are (i) the *basis set* $\{|\phi_i\rangle, i = 1, \ldots, N\}$, a set of functions of position used to compute the numerical elements of $\boldsymbol{H}$ and $\boldsymbol{S}$; (ii) the *complete set*, $\{|x_i\rangle, i = 1, \ldots, N\}$, a set of vectors which must span the *entire* N-dimensional Hilbert space and in terms of which *any* vector may be expressed; (iii) the *expansion set*: most iterative methods share the assumption that all eigenvectors of interest may be expanded in a *small* set of N-dimensional vectors $\{|b_j\rangle, j = 1, \ldots, N_b\}$ (the expansion set); here $N_b \ll N$. The expansion set need contain only N_b elements precisely because it need span *only* that part of the N-dimensional Hilbert space corresponding to the lowest n (or smaller) eigenvectors of $\boldsymbol{H}$. Depending on the iterative method discussed, N_b may be smaller than or greater than n.

We will not dwell on the properties of the elements of $\boldsymbol{H}$ and $\boldsymbol{S}$, but will assume that they are readily available for use by a computer program. Each iterative method is characterised by its choice of the complete set $\{|x_i\rangle\}$ and the expansion set $\{|b_i\rangle\}$; these will be discussed as they arise below.

A second feature shared by the iterative methods is that, in order to provide reasonable input guesses to the iteration problem, one often partitions $\boldsymbol{H}$ into a 'small' part $\boldsymbol{H}_0$ (i.e. the unperturbed or 'zero-order' problem), of 'small' dimensions $N_0 \times N_0$, and the rest. (With the exception of the Löwdin perturbation approach (Löwdin 1951), none of the methods to be described below depends for its success (or lack thereof) on the extent to which the full $N \times N$ matrices can be *accurately* partitioned into 'big' and 'small' parts.) One directly diagonalises the $N_0 \times N_0$ problem (say, by the CH method because N_0 is 'small') and uses its output as input to the iterative procedure, which then refines eigenvectors and eigenvalues to a desired accuracy†. Since the $N_0 \times N_0$ problem is by assumption small, in what follows we will assume that we have the general hierarchy

$$n \leqslant N_0 \ll N. \tag{2.1}$$

Generally (as discussed in § 5) one should choose N_0 so that the *level structure* (i.e. the ordering and degeneracy of eigenvalues) of the lowest n states of the full $N \times N$ Hamiltonian $\mathbf{H}$ is preserved, so as to avoid possible omission of eigenvalues and the

† In electronic structure studies, for example, N_0 may be the size of the 'minimal basis set', determined by the product of the number of atoms of a given type and the minimum number of basis states needed to crudely represent the wavefunctions for that atom, summed over the number of different atomic types in a unit cell. For a spin degeneracy of two, the number of occupied states (those, n in number, whose eigenvectors and eigenvalues are sought) may be roughly $N_0/2$.

necessity for eigenvalue sorting†. The first step is to use the Choleski-Householder method to find the lowest n eigensolutions of $\boldsymbol{H}_0$. Next, one augments the N_0-dimensional eigenvectors of $\boldsymbol{H}_0$ (which we will denote $\{|a_j^0\rangle, j = 1, \ldots, N_0\}$) with $(N - N_0)$ zeros to make them, together with the set of eigenvalues of $\boldsymbol{H}_0$ (denoted $\{\lambda_j^0\}$), suitable input guesses for whatever iterative method is being used‡.

Most of the efficient iterative methods are of the so-called 'basis expansion' variety (see, for example, Davidson 1978), based on the expectation that the approximate eigenvector will converge to the exact one rapidly as one adds to the expansion set $\{|b_j\rangle\}$ one or more new N-dimensional vectors per iteration. In each iteration one finds the eigenvalues and eigenvectors of the 'expansion set-projected' matrix problem

$$\Xi|c\rangle = \varepsilon\Omega|c\rangle \tag{2.2}$$

where

$$\Xi_{ij} = \langle b_i|\boldsymbol{H}|b_j\rangle \qquad \Omega_{ij} = \langle b_i|\boldsymbol{S}|b_j\rangle. \tag{2.3}$$

Then ε_k is the current approximation to the kth eigenvalue of $\boldsymbol{H}$ and

$$|a_k\rangle = \sum_i \langle b_i|c_k\rangle|b_i\rangle \tag{2.4}$$

is the current approximation to the kth eigenvector of $\boldsymbol{H}$, where $|c_k\rangle$ is the kth eigenvector of equation (2.2).

Finally, we note that most of the methods described below are *sequential*, i.e. they are designed to find first one eigenvector (to within a specified tolerance) and its eigenvalue, then the next, etc. Several, however, are 'block' methods, i.e. they generate approximations to several eigenvalues and eigenvectors at a time.

2.2. The Newton-Nesbet step

Obviously all of the iterative methods have in common the need to generate a new vector (either for addition to the expansion set or directly as a new approximation to an eigenvector) from a current one. Hence we next discuss the so-called Newton step. A quantity of central importance is the 'residual vector'

$$|R(|A^{ap}\rangle, E^{ap})\rangle = (\boldsymbol{H} - E^{ap}\mathbf{S})|A^{ap}\rangle \tag{2.5}$$

the quantity $R = (\langle R|R\rangle/\langle A^{ap}|\boldsymbol{S}|A^{ap}\rangle)^{1/2}$ is a widely used measure of the extent to which an approximate eigenvector $|A^{ap}\rangle$ and eigenvalue E^{ap} fail to be exact. In equation (2.5) E^{ap} is directly computable from $|A^{ap}\rangle$ via the Rayleigh quotient:

$$E^{ap} = \langle A^{ap}|\boldsymbol{H}|A^{ap}\rangle/\langle A^{ap}|\boldsymbol{S}|A^{ap}\rangle. \tag{2.6}$$

The strategy adopted by most iterative methods is to determine a vector increment $|\delta A\rangle$ which, when added to the current approximate eigenvector $|A^{ap}\rangle$, yields a residual vector as small as possible (ideally, zero). We thus require that

$$|R(|A^{ap} + \delta A\rangle, E^{ap})\rangle = |R(|A^{ap}\rangle, E^{ap})\rangle + (\boldsymbol{H} - E^{ap}\boldsymbol{S})|\delta A\rangle = 0 \tag{2.7}$$

† Since $n \ll N$ and one often knows from group theory what the degeneracies should be, this is rarely a problem.

‡ Alternatively, one may use as the input guess a vector believed to be 'close' to the final eigenvector sought.

where we have used the linearity of the residual operator. Unfortunately, the formal solution

$$|\delta A\rangle = -(\boldsymbol{H} - E^{\mathrm{ap}}\boldsymbol{S})^{-1}|R(|A^{\mathrm{ap}}\rangle, E^{\mathrm{ap}})\rangle \tag{2.8}$$

is no easier to solve than the original eigenproblem because of the need for matrix inversion. One may develop a simple basis for iteration by expanding $|\delta A\rangle$ in the complete set $\{|x_j\rangle\}$ and taking the inner product with one, say $|x_i\rangle$; the result

$$0 = \langle x_i | R\rangle + \sum_j \langle x_i|\boldsymbol{H} - E^{\mathrm{ap}}\boldsymbol{S}|x_j\rangle\langle x_j | \delta A\rangle \tag{2.9}$$

however, remains equivalent to equation (2.8). To overcome this difficulty, all popular iterative methods make the 'diagonal approximation' in keeping only the term $j = i$ in equation (2.9). (For later use we note that this would be *exact* if the $\{|x_i\rangle\}$ were the true eigenvectors of $\boldsymbol{H}$; hence, the closer they are to true eigenvectors, the better the diagonal approximation will be.) The result of the diagonal approximation becomes

$$|\delta A\rangle = -\sum_i{}' \frac{\langle x_i | R\rangle |x_i\rangle}{\langle x_i|\boldsymbol{H} - E^{\mathrm{ap}}\boldsymbol{S}|x_i\rangle} \tag{2.10}$$

which we shall term the 'Newton step' because of its resemblance to the usual Newton-Raphson iterative procedure for finding the zero of a function. The similarity of this expression to low-order perturbation theory follows from its Taylor series origins†. The prime in (2.10) means 'omit any i such that the magnitude of the denominator is less than a cutoff δ'; one may choose δ to be typical of the numerical difference between eigenvalues which will be regarded as physically degenerate‡. This omission is necessary because we have effectively defined a diagonal approximation to $(\boldsymbol{H} - E^{\mathrm{ap}}\boldsymbol{S})^{-1}$ in equation (2.8); if $R_k = 0$ then the component $A_k^{\mathrm{ap}} = A_k^{\mathrm{exact}}$ so δA_k should be zero.

Clearly we now have the framework for iteration: we replace $|A^{\mathrm{ap}}\rangle$ by $|A^{\mathrm{new}}\rangle = |A^{\mathrm{ap}}\rangle + |\delta A\rangle$, and E^{ap} by $E^{\mathrm{new}} = \langle A^{\mathrm{new}}|\boldsymbol{H}|A^{\mathrm{new}}\rangle / \langle A^{\mathrm{new}}|\boldsymbol{S}|A^{\mathrm{new}}\rangle$, and iterate. In practice, however, simple iteration of this process is not guaranteed to converge for an arbitrary Hamiltonian $\boldsymbol{H}$. In our description of the new RMM-DIIS method in § 3 we will show how the sequence of the $\{|\delta A^{(m)}\rangle\}$ generated by iterations, labelled m, using a careful choice of complete set and a powerful convergence acceleration step, can be used in an efficient technique for diagonalising large matrices.

With the notation established above we are now in a position to review several well known iterative methods commonly used in physics.

2.3. *Simple iteration*

Choosing

$$\{|x_i\rangle\} = \{|e_i\rangle\} \tag{2.11}$$

where $|e_i\rangle$ is a column vector of zeros with a one in the ith position, and using it in the Newton step (equation (2.10) above) yields one version of the well-known Nesbet method (Cooper 1948, Nesbet 1965). This is *not* a basis expansion method, but relies

† However, in contrast to perturbation theory, in deriving equation (2.10) we did not need to require orthogonality of the complete set $\{|x_i\rangle\}$, and we arrived at a form which does not become more difficult as the iteration number (order of perturbation theory) increases. See Bendt and Zunger (1982b) for details.

‡ $\delta \sim 10$ Ryd is typical, though results should be and are extremely insensitive to this choice.

instead on repetition of the Newton step. Variants of this procedure and a block form are described in detail by Raffenetti (1978, 1979). Because the $\{|x_i\rangle\}$ are not close to eigenvectors of $\boldsymbol{H}$, however, the diagonal approximation (equation (2.10)) is expected to be poor; even a convergence accelerator may not overcome the limitations of an inadequate complete set.

2.4. *The Lanczos method*

Probably the oldest and best known basis-expansion iterative method is due to Lanczos (Lanczos 1950, Parlett 1980, Scott 1981). It may be described succinctly by the identification (for the case $\boldsymbol{S} = \boldsymbol{I}$)

$$\{|b_i\rangle\}[p = 0/1/2/\ldots] = [|a_j^0\rangle/\boldsymbol{H}|a_j^0\rangle/\boldsymbol{H}^2|a_j^0\rangle/\ldots] \tag{2.12}$$

where the vectors between slashes are those added to the expansion set in the corresponding iteration number p (see Davidson (1975) for an identification of this basis set): in the zeroth iteration the expansion set consists of the input guess for the eigenvector, in the first iteration it contains two vectors $|a_j^0\rangle$ and $\boldsymbol{H}|a_j^0\rangle$, etc†. $\boldsymbol{H}$ is then diagonalised with the space of expansion vectors; see equations (2.2)-(2.5). This method typically works well for the largest eigenvalues, since repeated multiplication by $\boldsymbol{H}$ followed by normalisation converges to the eigenvector corresponding to the highest eigenvalue.

2.5. *Block Davidson*

We describe a method originally proposed by Davidson (1975)‡, as modified into block form by Liu (1978). It is described by the sets

$$\{|x_i\rangle\} = \{|e_i\rangle\} \tag{2.13}$$

$$\{|b_i\rangle\}[p = 0/1/\ldots] = \left[\{|a_j^0\rangle, j = 1, \ldots, N_0\} \Big/ \left\{|b_{N_0+k}\rangle \equiv |\delta A_k\rangle - \sum_{j=1}^{N_0+k-1} \langle b_j|\delta A_k\rangle |b_j\rangle\right\}_{k=1,\ldots,m} \Big/ \ldots\right]. \tag{2.14}$$

In words, in each iteration m new orthogonal vectors are added to the expansion set: in the first iteration each new vector consists of the Newton correction $|\delta A_k\rangle$ corresponding to the kth eigenvector of the $N_0 \times N_0$ problem (augmented with zeros to make it N-dimensional), Gram-Schmidt orthogonalised to its predecessors already added to the expansion set. (In practice, if the resulting vector has too small a norm ($< \varepsilon$), it is discarded.) This approach has the virtue that the input of a single row of $\boldsymbol{H}$ can be used to refine m eigenvalues and eigenvectors per iteration. Since there are m more expansion vectors at the end of an iteration, the rate of convergence of the iteration

† To be precise, if all the $\{|b_i\rangle\}$ are normalised (but not orthogonal) $|b_1\rangle \equiv |a_j^0\rangle$, $|b_2\rangle \equiv |T_1\rangle$, and for $k \geq 3$, $|b_k\rangle \equiv |T_k\rangle - |b_{k-2}\rangle$ where, for all k, $|T_k\rangle \equiv [\boldsymbol{H} - \langle b_{k-1}|\boldsymbol{H}|b_{k-1}\rangle \boldsymbol{I}]|b_{k-1}\rangle$. In the absence of round-off error Ξ (for $\Omega = \boldsymbol{I}$) would be *tridiagonal*. In practice, one usually Gram-Schmidt orthogonalises $|b_k\rangle$ to all earlier members of $\{|b_i\rangle\}$ to prevent error accumulation.

This scheme is an alternate formulation of the *unsymmetric* Lanczos method; the so-called 'Paige' variant (Parlett 1980) of the *symmetric* Lanczos method, which is applicable to Hermitian matrices, would probably be more efficient.

‡ See also Nesbet (1981) and Butscher and Kammer (1976) for a root-homing version.

sequence may be improved considerably. (The usual Davidson procedure (Davidson 1975, Butscher and Kammer 1976) corresponds to the choice $m = 1$.)

2.6. *The Löwdin perturbation method*

This technique (Löwdin 1951, Brust 1964, 1968) bears no close relationship to any of the approaches described above or below, but is widely used to 'fold in' the effects of a large $N \times N$ Hamiltonian into a much smaller $N_0 \times N_0$ effective Hamiltonian. For the matrix eigenproblem $(\boldsymbol{H} - \lambda I)|x\rangle = 0$, the $N \times N$ Hamiltonian $\boldsymbol{H}$ is partitioned into four blocks, where $\boldsymbol{H}_0$ is $N_0 \times N_0$ as shown schematically:

$$\left[\begin{array}{c|c} \boldsymbol{H}_0 & \boldsymbol{B} \\ \hline \boldsymbol{B}^+ & \boldsymbol{C} \end{array}\right] \tag{2.15}$$

so that, provided the number of eigenvalues sought, n, satisfies $n \leqslant N_0 \ll N$, the effective eigenproblem may be reduced to diagonalising the $N_0 \times N_0$ matrix

$$\boldsymbol{H}_0^{\text{eff}} = \boldsymbol{H}_0 - \boldsymbol{B}(\boldsymbol{C} - \lambda I)^{-1}\boldsymbol{B}^+. \tag{2.16}$$

Obviously the effective matrix elements now depend on the actual eigenvalues of the full Hamiltonian $\boldsymbol{H}$; the iterative solution of the new problem (2.16) naturally generates the so-called Brillouin–Wigner perturbation series (e.g. Ziman 1969). In practice (Brust 1964, 1968), to find eigenvalues of $\boldsymbol{H}$ near λ_j^0 (an eigenvalue of $\boldsymbol{H}_0$) one replaces λ above by λ_j^0, or even by a free-electron approximation to it. As typically implemented, however, the Löwdin scheme keeps only the lowest two orders in perturbation theory. As such, the technique may diverge from a correct eigenvalue or converge to an incorrect one, as will be illustrated below.

3. RMS–DIIS method

The 'residual minimisation/direct inversion in the iterative subspace' (RMM-DIIS or simply DIIS) method due to Bendt and Zunger (1982a, b) will be discussed in detail below; it can be described by the choices, for iteration number p and eigenvalue j,

$$\{|x_i\rangle\} = \{|a_j^0\rangle, j = 1, \ldots, N_0\} + \{|e_j\rangle, j = N_0 + 1, \ldots, N\} \tag{3.1}$$

and

$$\{|b_i\rangle\}[p = 0/1/2/\ldots] = [|a_j^0\rangle/|\delta A_j^{(1)}\rangle/|\delta A_j^{(2)}\rangle/\ldots]. \tag{3.2}$$

The first distinguishing feature of the DIIS method is its choice of complete set. As discussed after equation (2.9) above the usual diagonal approximation in the Newton step (equation (2.10)) becomes exact if the complete set selected consists of the true eigenvectors of $\boldsymbol{H}$. Thus we expect that a complete set consisting of approximate eigenvectors of $\boldsymbol{H}$ (i.e. the set $\{|a_j^0\rangle\}$) of eigenvectors of $\boldsymbol{H}_0$ augmented with zeros to make them N-dimensional) plus a set of unit vectors (to make the complete set contain N linearly independent vectors) will be a considerably better choice than the usual set $\{|e_j\rangle, j = 1, \ldots, N\}$. Moreover, the eigenvectors of $\boldsymbol{H}_0$ are immediately available since $\boldsymbol{H}_0$ was diagonalised to provide starting guesses for the eigenvalues and eigenvectors to be refined by the iterative process. For our choice (3.1), then

$$\langle e_k|\,\delta A\rangle = -\sum_{i=1}^{N_0}{}' \frac{\langle a_i^0|\,R(E_{\text{old}})\rangle\langle e_k|\,a_i^0\rangle}{(\lambda_i^0 - E_{\text{old}})\langle a_i^0|\boldsymbol{S}|a_i^0\rangle} - \sum_{i=N_0+1}^{N}{}' \frac{\langle e_i|\,R\rangle\delta_{ik}}{(H_{ii} - E_{\text{old}}S_{ii})}. \tag{3.3}$$

The expansion set $\{|b_i\rangle\}$ requires further explanation. Like the other basis expansion methods, DIIS uses a Newton step, equation (3.3), to generate a new vector $|\delta A\rangle$ which is then added to $\{|b_i\rangle\}$. The elements of this set are thus the $|\delta A\rangle$ generated in each of the preceding iterations, so that DIIS clearly incorporates information from the entire iteration history for the given eigenvector being refined. Since the vectors $\{|\delta A^{(i)}\rangle\}$ are generated by the iteration process, they are said to span the 'iterative subspace'. The great power of the DIIS method lies, however, not in the choices for $\{|b_i\rangle\}$ and $\{|x_i\rangle\}$ but in the fact that, to produce the new approximate eigenvector of $\boldsymbol{H}$ one asks: 'what linear combination of the expansion vectors *minimises* the *residual* of the resulting vector $|A^{\text{new}}_{\text{DIIS}}\rangle$?' (hence the RMM of residual minimisation).

In what follows the steps involved will be described sequentially. It will be assumed that upon entering the mth iteration one has available $E_{\text{old}} = E^{(m-1)}$ and $|A_{\text{old}}\rangle = |A^{(m-1)}\rangle$, the latest approximations to the eigenvalue and its eigenvector. At this point the expansion set consists of $\{|\delta A^{(0)}\rangle, |\delta A^{(1)}\rangle, \ldots, |\delta A^{(m-1)}\rangle\}$, where we have defined $|\delta A^{(0)}\rangle = |a^0_j\rangle$, the input guess for the eigenvector. The procedure is as follows.

(i) Generate $|\delta A^{(m)}\rangle$ via the Newton step equation (3.3) and add it to the set $\{|b_i\rangle\}$, which now contains $(m+1)$ elements. (One generally forms and stores the set of vectors $\{\boldsymbol{H}|\delta A^{(m)}\rangle\}$ and $\{\boldsymbol{S}|\delta A^{(m)}\rangle\}$ as the $\{|\delta A^{(m)}\rangle\}$ are generated, for reasons which will become obvious in step (ii)).

(ii) Perform the DIIS step: one writes

$$|A^{\text{new}}_{\text{DIIS}}\rangle = \sum_{j=0}^{m} \alpha_j |\delta A^{(j)}\| \tag{3.4}$$

and requires that $\rho^2 = \|R(|A^{\text{new}}_{\text{DIIS}}\rangle, E_{\text{old}})\rangle\|^2$ be a minimum, i.e. that

$$\frac{\delta\rho^2}{\delta\alpha^*_k} = \frac{\delta}{\delta\alpha^*_k} \frac{\sum^m_{r,s=0} \alpha^*_r \alpha_s \langle \delta A^{(r)}(\boldsymbol{H} - E_{\text{old}}\boldsymbol{S})|(\boldsymbol{H} - E_{\text{old}}\boldsymbol{S})\delta A^{(s)}\rangle}{\sum^m_{r,s=0} \alpha^*_r \alpha_s \langle \delta A^{(r)}|\boldsymbol{S}|\delta A^{(s)}\rangle} = 0 \tag{3.5}$$

for $k = 0, \ldots, m$. Letting $|\alpha\rangle$ denote the $(m+1)$-dimensional vector whose components are the α_k, this problem is equivalent to finding the eigenvector of lowest eigenvalue of the generalised Hermitian eigenproblem:

$$\boldsymbol{P}|\alpha\rangle = \rho^2 \boldsymbol{Q}|\alpha\rangle \tag{3.6}$$

where

$$P_{rs} = \langle (\boldsymbol{H} - E_{\text{old}}\boldsymbol{S})\delta A^{(r)}|(\boldsymbol{H} - E_{\text{old}}\boldsymbol{S})\delta A^{(s)}\rangle \tag{3.7}$$

and

$$Q_{rs} = \langle \delta A^{(r)}|\boldsymbol{S}|\delta A^{(s)}\rangle. \tag{3.8}$$

Since $\boldsymbol{P}$ and $\boldsymbol{Q}$ are matrices of size only $(m+1)\times(m+1)$ and the number of iterations required is small, this diagonalisation may be performed by the Choleski–Householder method in negligible time. (This step has been termed by Pulay, who introduced this step to accelerate the convergence of self-consistent solutions of the Schrödinger equation (Pulay 1980), 'direct inversion in the iterative subspace', hence RMM-DIIS.)

(iii) Substituting the coefficients $\{\alpha_k\}$ minimising the residual into (3.4), one then calculates

$$E^{\text{new}}_{\text{DIIS}} = \langle A^{\text{new}}_{\text{DIIS}}|\boldsymbol{H}|A^{\text{new}}_{\text{DIIS}}\rangle / \langle A^{\text{new}}_{\text{DIIS}}|\boldsymbol{S}|A^{\text{new}}_{\text{DIIS}}\rangle \tag{3.9}$$

and

$$|R^{\text{new}}\rangle = (\boldsymbol{H} - E^{\text{new}}_{\text{DIIS}}\boldsymbol{S})|A^{\text{new}}_{\text{DIIS}}\rangle / \langle A^{\text{new}}_{\text{DIIS}}|\boldsymbol{S}|A^{\text{new}}_{\text{DIIS}}\rangle \tag{3.10}$$

using the stored values of the $\{\boldsymbol{H}|\delta A^{(k)}\rangle\}$. Note that the new residual $R^{\text{new}} = (\langle R_{\text{new}}|R_{\text{new}}\rangle)^{1/2}$ is usually much less than $(\rho^2)^{1/2}$ in equation (3.6) because the energy (3.9) has been updated.

(iv) Check for convergence; if the residual R^{new} is less than a prescribed tolerance, stop the iterations and store the new eigenvector and eigenvalue. If the desired residual has not been achieved, replace E^{old} by $E^{\text{new}}_{\text{DIIS}}$ and $|A_{\text{old}}\rangle$ by $|A^{\text{new}}_{\text{DIIS}}\rangle$ and return to step (i).

It should be clear from the discussion above that the DIIS method in no way relies on any special features of the matrices treated, e.g. they need not be sparse, or diagonally dominant, etc. The new RMM-DIIS method clearly builds on a number of existing ideas, and it may be though of schematically as

{DIIS} = {better complete set} + {Newton step} + {convergence accelerator}

the last of which is the real key to its success. As such, it seems clear that there may be a whole class of new iterative methods based on such 'DIIS' steps, differing in details from the method described above. Some suggestions as to implementation are described in § 5 and the appendix.

4. Numerical examples

Several questions arise immediately in the context of a new matrix diagonalisation iterative method. Among these are: (i) how does its convergence *rate* compare with existing methods?; (ii) how does the computation *time* spent to obtain a set of eigenvalues and eigenvectors to desired residual accuracy compare with competitors?; (iii) to what features of the matrices to be diagonalised is the method sensitive? In what follows we will address these questions by means of specific examples. Although this method can be easily applied to huge matrices (see, for example, Bendt and Zunger 1982c, Jaffe and Zunger 1983), for illustrative purposes we consider below relatively simple matrices of moderate size.

It cannot be overemphasised that the convergence rate (i.e. the decrease in residual per iteration) *alone* is an inadequate index of the efficiency of an iterative method; if the time taken per iteration is excessive, a given iterative method may not be practical. Nonetheless, the convergence rate provides a useful demarkation between several classes of iterative methods.

4.1. Real matrix convergence rates

In figure 1 we show the iteration histories for the lowest eigenvector of the 50×50 so-called 'modified Nesbet matrix' (Raffenetti 1978, 1979) defined by

$$H_{ij} = H_{ji} = 1 \qquad i \neq j \qquad i,j = 1,2,\ldots,50$$
$$H_{ii} = \begin{cases} 1+0.1(i-1) & 1 \leqslant i \leqslant 5 \\ 2i-1 & 6 \leqslant i \leqslant 50. \end{cases} \tag{4.1}$$

Here we have used as input guesses the eigenvectors and eigenvalues of the leading 5×5 sub-block of $\boldsymbol{H}$, i.e. $N_0 = 5$; this is not so for curves, D, E and F. This matrix has been used recently by Raffenetti (1979) for tests of other iterative methods. While not large, it formally presents a fairly difficult test because the leading 5×5 sub-block is clearly not diagonally dominated; we may, in fact, regard the close spacings of the

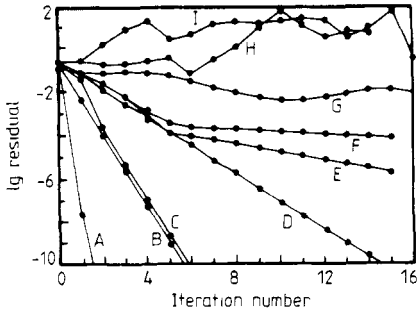

Figure 1. Iteration histories for the modified Nesbet matrix: A, block Davidson ($m = n = 4$), $\varepsilon = 10^{-7}$; B, Davidson; C, DIIS; D, E, F, simultaneous coordination relaxation for $m = 5$, 3, 1, respectively (see Raffenetti 1979); G, our version of unsymmetric Lanczos (equation (2.12)); H, simple iteration using the DIIS complete set (equation (3.1)); I, simple iteration using the Nesbet complete set (see equation (2.13)).

first five diagonal matrix elements as an imitation of the effects of five-fold degeneracy. Inspection of figure 1 shows the following.

(i) The three most rapidly convergent methods are the block Davidson method (with $m = n = 4$, curve A), the regular Davidson method ($m = 1$, curve B) and DIIS, curve C.

(ii) The 'simultaneous coordinate relaxation' method (Raffenetti 1978) (not discussed here), applied here in a block form, converges at a rate which is fairly sensitive to the number of eigenvectors being refined simultaneously ($m = 1, 3, 5$; curves F, E, and D, data of Raffenetti (1978)).

(iii) Our variant (equation (2.12)) of the unsymmetric Lanczos procedure (curve G) appears to converge, but slowly and non-monotonically.

(v) Simple iteration of the Newton step (curve H), even using the improved DIIS complete set (equation (3.1)), converges extremely slowly, if at all. Simple iteration of the Newton step using the diagonal approximation and a unit vector complete set (the 'Nesbet' approximation, curve I) also fails to converge.

Since the first three methods have the best convergence rates, we will confine discussions below to these methods. The $m = 1$ (regular Davidson) and DIIS results converge at essentially the same rate, though there is an offset because of different handling of the first iteration. The $m = n = 4$ block Davidson method, however, converges considerably faster (about $3\frac{1}{2}$ times) than the other two†.

4.2. Comparison of block Davidson and DIIS approaches

The reason for the difference in convergence rates between the block Davidson (BD) and DIIS methods is implicit in their different structures. The BD method chooses to add many (in the usual BD method, $m = n$) new vectors to the expansion set in one iteration. The benefits this provides are (i) faster or even much faster convergence than the DIIS or unblocked ($m = 1$) Davidson methods and (ii) decreased input/output requirements and increased convenience because many (n) eigenvectors and eigenvalues are being refined simultaneously in the same loop. The drawback of the BD method is most pronounced, however, when the 'small' matrix dimension N_0 is relatively

† It would have converged four times faster (because $m = 4$) except that one of the added vectors had a post Gram-Schmidt norm ε of $< 10^{-7}$ (our threshold for acceptance) and so was discarded.

large. It is that the size of the expansion set Hamiltonian Ξ increases by m with each iteration—(N_0+m) square after the first iteration, etc. If $m \sim n$ and $n \sim N_0$ (cases of typical physical interest) the time spent in the Choleski-Householder operations grows rapidly. If the number of new vectors added to the expansion set per iteration is reduced below n, however, the convergence *rate* is generally reduced accordingly.

The great virtue of the DIIS procedure is that, after the zeroth iteration diagonalisation of the small $\boldsymbol{H}_0$ problem, it need never diagonalise a matrix larger than the maximum number of iterations, typically $\leqslant 10$. Its drawback is that it is in essence a sequential process: for each eigenvector to be refined, the full Hamiltonian (and overlap matrix, if present) must be read in one row at a time†. The discussion of computation time scaling, in which the true benefits of the DIIS method are demonstrated, is postponed to § 5.

4.3. Convergence rates for a complex Hermitian matrix

It may be objected that the modified Nesbet matrix of figure 1 is rather artificial; we consider next the numerical results for a 181×181 complex Hermitian Hamiltonian matrix describing slightly expanded ZnSe, at the Γ point of the Brillouin zone, within a plane wave basis ($\boldsymbol{S} = \boldsymbol{I}$). The matrix elements are computed using the pseudopotential parametrisation‡ of Louie *et al* (1977) and were written into external files from which they were read, as usual, one row at a time. This Hamiltonian is fairly typical in both form and numerical content of those encountered in band structure calculations. It presents the added difficulty, typical of solid-state problems, that many of the eigenvalues are degenerate; the exact level sequence is shown on the right-hand side of figure 3, to be discussed later.

The iteration histories for the lowest eigenvalue at Γ for six different methods are shown in figure 2. Here $N_0 = 15$ was selected because this was the smallest set of the plane waves for the ZnSe structure for which the eigenvalue structure of the lowest eight levels of H was preserved. The structure of the Hamiltonian here is sufficiently well behaved that all methods used converged, albeit some better than others.

Once again there is a clear division between the more sophisticated basis expansion techniques: block Davidson (curve A), Davidson (curve B) and DIIS (curve C) and the simple iterative methods (curves D, E and F). This time, however, the presence of degeneracy enormously reduces the factor by which the block Davidson (here, $m = n = 8$) converges faster than the DIIS and regular Davidson methods. As before, DIIS and regular Davidson methods converge at about the same rate, as the parallelism of curves B and C indicate.

4.4. Complex Hermitian matrix: Löwdin method

Before turning to the computation time discussion, we briefly remark on the inadequacy

† To reduce input/output operations, one could easily 'block' the DIIS step, so that m eigenvalues and eigenvectors were refined simultaneously at the expense of increased central memory storage. The residual for each eigenvalue, however, would converge as if it were being refined independently of the other $(m-1)$.

One may also exploit the Hermiticity of $\boldsymbol{H}$ and $\boldsymbol{S}$ by reading only up to and including the diagonal of $\boldsymbol{H}$ and $\boldsymbol{S}$: see Shavitt (1970). We thank Dr Sverre Froyen for pointing out how the original DIIS algorithm could be modified to take advantage of this simplification.

‡ The form is $V_{\text{emp}}(q) = b_1(q^2 - b_2)\{\exp[b_3(q^2 - b_4)] + 1\}^{-1}$ where the coefficients for Zn(Se) are $b_1 = 6.7008(0.2334)$, $b_2 = 1.4983(3.3858)$, $b_3 = 0.6696(0.7266)$ and $b_4 = -4.7128(+2.2012)$. We scaled the reciprocal lattice vectors by a factor of about 1.06 (i.e. with a lattice constant of 6.002 Å) to simulate dilation of the lattice.

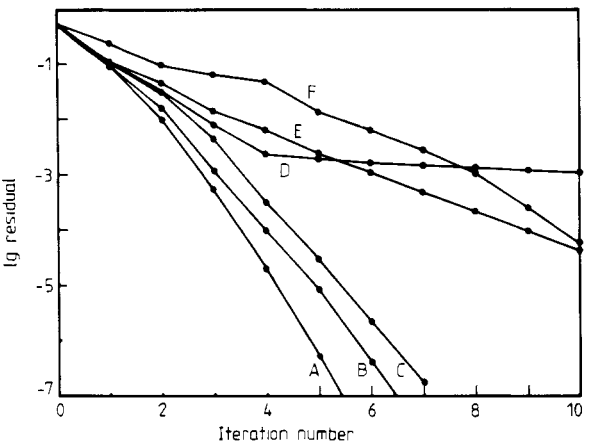

Figure 2. Iteration histories for the complex Hermitian ZnSe matrix: A, block Davidson ($m = n = 8$); B, Davidson; C, DIIS; D, simple iteration using the DIIS complete set equation (3.1); E, our version of the unsymmetric Lanczos method (equation (2.12)); F, simple iteration using the Nesbet complete set equation (2.13).

of the well known Löwdin perturbation scheme (see § 2). In figure 3 we show, for fixed 'small" matrix dimension $N_0 = 9$, the eigenvalues which result from the DIIS method (or, in fact, any method which gives correct eigenvalues) and the Löwdin perturbation calculation (using the eigenvalues of $\boldsymbol{H}_0$ in equation (2.17)) as one increases the size N of the large matrix from 9 to 181 (the full reference Hamiltonian described above). Curves belonging to the same multiplet (indicated on the right-hand margin of the figure) are connected with braces†. It is clear that for the valence band (the lowest two multiplets, of degeneracy 1 and 3, respectively) the Löwdin perturbation

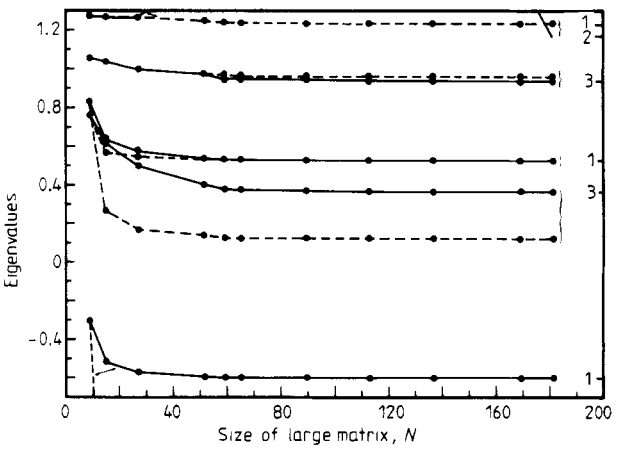

Figure 3. Comparison of Löwdin (- - -) and DIIS (——) convergence for fixed $N_0 = 9$ as a function of N.

† The pathology for DIIS and the Löwdin method for the highest eigenvalues shown (i.e. complete omission of the second level from the top) is a manifestation of the inadequacy of $N_0 = 9$ in representing the lowest six multiplets of the 181×181 problem.

method gives results which are convergent, but wrong: the valence band is *not* described by the eigenvalues of $\boldsymbol{H}_0$ used in equation (2.17). By contrast, the lowest two eigenvalues of the *conduction* band (also of degeneracy 1 and 3, respectively) found by the Löwdin method are quite close to the exact values, since free-electron conduction band levels are frequently less perturbed by the crystal potential.

4.5. Computation time considerations and constraints

As argued in § 4.2, the standard block Davidson method rapidly becomes unwieldy if (a) $N_0 \gg 1$, so that the starting matrix $\boldsymbol{H}_0$ is already large enough to make its Choleski–Householder diagonalisation time ($\sim N_0^3$) a significant constraint, and/or if (b) many new vectors are added to the expansion basis set $\{|b_i\rangle\}$ per iteration. To make more quantitative the time trade-offs in DIIS and BD, we adopt a simple scaling model. It is based on several observations: (i) both methods (indeed, all iterative methods) must calculate vectors of the form $\boldsymbol{H}|q\rangle$ for N-dimensional vectors $|q\rangle$. Such multiplications take a time proportional to N^2 (and are generally the dominant time step). There is one such computation per DIIS iteration, but m per iteration for BD; (ii) in the BD method there is an *additional* time associated with Choleski–Householder diagonalisation of the matrix Ξ in the basis expansion set, proportional to $(N_0+mI)^3$, where I is the iteration number and m is the number of new vectors added to $\{|b_i\rangle\}$ per iteration. Thus the times required to find the lowest n eigenvalues of $\boldsymbol{H}$ are, within the model,

$$T_n^{\mathrm{BD}} \simeq I_m^{\mathrm{BD}} m(aN^2) + b \sum_{I=1}^{I_m^{\mathrm{BD}}} (N_0 + mI)^3$$

$$T_n^{\mathrm{DIIS}} \simeq I_m^{\mathrm{DIIS}} n(aN^2). \tag{4.2}$$

Here a and b are numerical constants to be determined from actual computer runs, I_m^{BD} is the number of BD iterations and I_m^{DIIS} is the *average* number of iterations per eigenvector (averaged over the n eigenvalues of $\boldsymbol{H}$ found), needed to achieve some fixed residual. One can thus define a dimensionless ratio R by the equation

$$T_n^{\mathrm{BD}} / T_n^{\mathrm{DIIS}} = (I_m^{\mathrm{BD}} / I_m^{\mathrm{DIIS}}) R \tag{4.3}$$

where

$$R = \frac{(m/N_0)}{(n/N_0)} \left(1 + \frac{(b/a)}{(m/N_0)} (N_0/N)^2 G(I_m^{\mathrm{BD}}, m/N_0) \right)$$

and

$$G(I, x) = 1 + \tfrac{3}{2}(I+1)x + (I+1)(I+\tfrac{1}{2})x^2 + \tfrac{1}{4}I(I+1)^2 x^3. \tag{4.4}$$

For purposes of numerical comparison we have taken $I_m^{\mathrm{BD}} = 3$ (fairly typical for the matrices we have examined) to achieve a residual of $\leqslant 10^{-4}$ (corresponding, in our experience, to ~eight decimal places in the eigenvalue and four in the eigenvector components); our numerical results indicate that $b/a \sim 1$. Contours of constant R are displayed as a function of N_0/N and n/N_0 ($= m/N_0$ for the usual block Davidson) in figure 4. If, for example, it were known that $I_m^{\mathrm{BD}} / I_m^{\mathrm{DIIS}} = \frac{1}{2}$ (e.g. for the band structure problem of figure 2) then, within the model, the regime of parameters above the $R = 2$ contour would favour use of the DIIS method. It should be noted that the situation of typical physical interest is that one knows the total number of eigenvalues sought

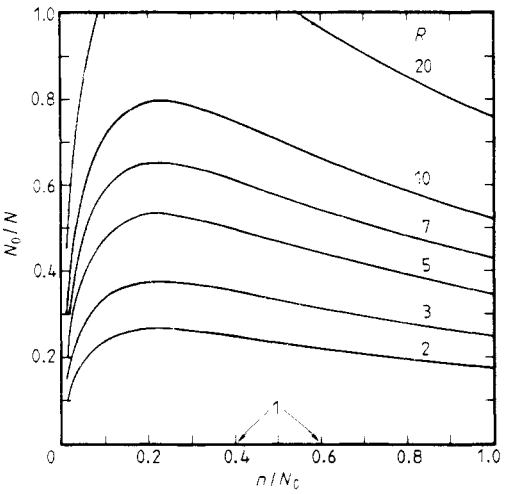

Figure 4. Block Davidson to DIIS time ratio for model computation times (equation (4.2)).

n, and the size of the large Hamiltonian N, so that the intersection of the hyperbola $n/N = \text{constant}$ with the contours R would determine the choice of method. The model results are very likely an oversimplification, but should be useful when $N_0 \gg 1$ and $N \gg 1$.

5. Implementation of the DIIS method

We discuss below some important considerations in the practical use of the RMM-DIIS method. As has been noted above, central to the implementation of most of the iterative methods is the identification of an appropriate 'small' matrix $\boldsymbol{H}_0$, subject to the constraint that its eigenvalues preserve the level structure of the lowest n eigenvalues of $\boldsymbol{H}$. It might be argued, however, that since a better initial guess for an eigenvector of $\boldsymbol{H}$ (from the $\boldsymbol{H}_0$ problem) will result in fewer iterations to achieve a given residual, one should pick N_0 as large as possible. In fact, the trade-offs in starting residual R_0 against the required number of iterations are shown for the lowest eigenvalue of the 181×181 complex Hermitian band structure matrix in figure 5. It should be noticed that the reduction in the starting residual R_0 is less than a factor of 4.5 as one increases N_0 from 1 (corresponding physically to a *free-electron* input eigenvalue) to 113†. Moreover, to achieve $R \leqslant 10^{-4}$ requires three iterations for $N_0 = 113$, but only six for $N_0 = 1$. However, the time needed to diagonalise H_0 scales as N_0^3, so that very rapidly (using arguments similar to those leading to equation (4.2)), it becomes more efficient simply to iterate more rather than use a larger N_0; the convergence *rate* $\Delta(\log R)$/iteration is essentially independent of N_0. (In fact, $\mathrm{d}(\log R_0)/\mathrm{d}N_0$ is not always monotonic.)

Thus our rules of thumb for the choice of the small matrix dimension N_0 are: (*a*) pick N_0 large enough to reproduce the level ordering and degeneracies of the lowest

† Even if we approximate the bottom of the valence band by a 1×1 matrix, corresponding to a homogeneous electron gas, we find after five iterations the correct *strongly inhomogeneous* state to eight decimals for the energy and four decimals for the wavefunction components.

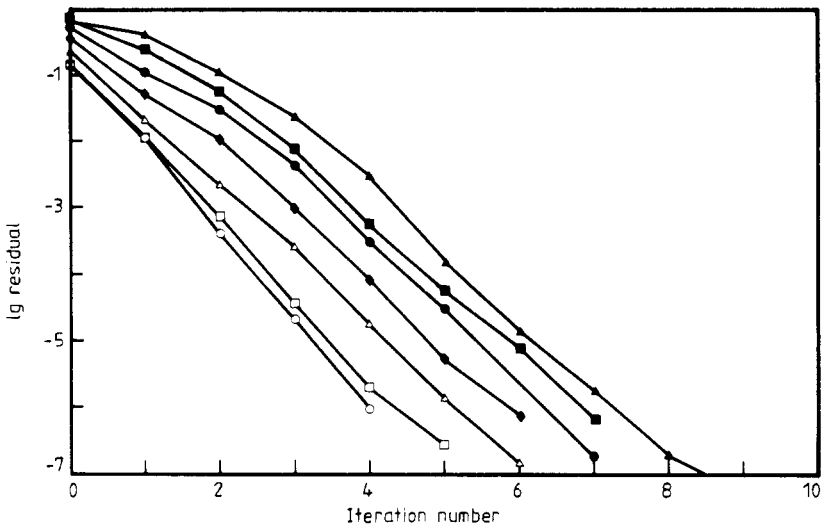

Figure 5. Dependence of iteration hisyory on small matrix dimension N_0 for 181×181 complex Hermitian matrix. Values of N_0: ▲, 1; ■, 9; ●, 15; ◆, 27; △, 51; □, 65; ○, 113.

n eigenvalues of $\boldsymbol{H}$, in order not to waste time on level crossings and eigenvalue sorts; (b) pick the smallest N_0 consistent with (a) to minimise time spent in Choleski-Householder diagonalisation of $\boldsymbol{H}_0$.

The DIIS procedure would seem to be most valuable under the following circumstances: (i) the size of the 'small' matrix, N_0, is itself large enough so that the time spent per iteration for the other efficient methods (e.g. the Davidson methods) becomes unwieldy, and/or (ii) there is significant degeneracy in the level structure of the eigenvalues sought (which often confuses or slows down competing methods).

The explicit sequence of steps in the DIIS procedure is given in the appendix.

6. Summary and conclusions

We have presented above a description of a new iterative method for diagonalising very large matrices. The structure and philosophy of the method were compared with other currently used methods; we have displayed numerical comparisons for two rather different test matrices which illustrate the strengths and weaknesses of the new and older methods, both in terms of convergence rates and computation times (within a simple model). A brief description of its implementation was given with suggestions for the choice of convergence and other parameters.

The new RMM-DIIS method has already been used successfully for a large number of electronic structure problems (e.g. Bendt and Zunger 1982c, Jaffe and Zunger 1983). It is efficient both in terms of computation time and central memory storage requirements, and holds the promise of pushing back a number of obstacles in the path of the calculation of, for example, the electronic structure of complex crystalline and amorphous materials. Subject to the constraint that one needs to provide reasonable input guesses for eigenvalues and eigenvectors, it may help remove the great premium which has customarily been placed on the choice of efficient basis functions.

Acknowledgments

In preparing this paper, we have benefitted from notes by Paul Bendt and from Bendt and Zunger (1982b). We gratefully acknowledge R Nesbet for pointing out NRCC (1978) and Nesbet (1965), and Ahri Brandt for useful discussions about the DIIS method. One of us (AZ) is indebted to E S Robinson for drawing his attention to the work of Pulay (1980), which has made the current method possible.

Appendix. Structure of DIIS calculation

It is assumed that the Hamiltonian $\boldsymbol{H}$ (and overlap $\boldsymbol{S}$, if present) are available for reading from external devices.

(1) Define $\boldsymbol{H}_0$ and use Choleski-Householder (e.g., in the EISPACK library) to find its lowest n eigenvalues $\{\lambda_j^0, j = 1, \ldots, N_0\}$ and eigenvectors $\{|a_j^0\rangle, j = 1, \ldots, N_0\}$; the $\{|a_j^0\rangle\}$ are augmented with zeroes to make them suitable N-dimensional input guesses.

(2) Pick an input eigenvalue $E_{\text{old}} = \lambda_j^0$ and an input eigenvector $|A_{\text{old}}\rangle = |a_j^0\rangle$ to refine:

(i) Initialise: store $|\delta A^{(0)}\rangle$; calculate and store $\boldsymbol{H}|\delta A^{(0)}\rangle$ and the input residual vector.

(ii) Begin iteration m.

(a) Use the Newton step to generate $|\delta A^{(m)}\rangle$; store it.

(b) Read in a row of $\boldsymbol{H}$ (and $\boldsymbol{S}$ if present) and compute $\boldsymbol{H}|\delta A^{(m)}\rangle$ (and $\boldsymbol{S}|\delta A^{(m)}\rangle$) and store them.

DIIS step:

(c) Calculate the new elements of $\boldsymbol{P}$ and $\boldsymbol{Q}$ and find the lowest eigenvector using CH.

(d) Form the DIIS new best eigenvector $|A_{\text{DIIS}}^{\text{new}}\rangle$: using the $\{\boldsymbol{H}|\delta A^{(i)}\rangle\}$ (and $\{\boldsymbol{S}|\delta A^{(i)}\rangle\}$) find the new DIIS expectation energy $E_{\text{DIIS}}^{\text{new}}$ and the corresponding new DIIS residual vector and its norm R. (This R is generally less than the $(\rho^2)^{1/2}$ found as the lowest eigenvalue of the DIIS eigenproblem since the energy $E_{\text{DIIS}}^{\text{new}}$ is better than E_{old}.)

(e) If $R <$ (stopping criterion), stop; if not, let $E_{\text{old}} = E_{\text{DIIS}}^{\text{new}}$ and $|A_{\text{old}}\rangle = |A_{\text{DIIS}}^{\text{new}}\rangle$ and return to (a).

It should be apparent that the principal central core memory storage requirements for use of the DIIS method are (i) $\boldsymbol{H}_0$ (and $\boldsymbol{S}_0$), for use by CH; (ii) $\{|a_j^0\rangle\}$, the N_0 (N_0-dimensional) eigenvectors of $\boldsymbol{H}_0$; (iii) various N-dimensional vectors: $|A_{\text{new}}\rangle$, $|R\rangle$, $|h\rangle$ (one row of $\boldsymbol{H}$), $|s\rangle$ (one row of $\boldsymbol{S}$, if present), and several other subsidiary vectors: $\boldsymbol{P}$ and $\boldsymbol{Q}$ require negligible storage, being $\leqslant 15 \times 15$.

As mentioned above, for use on a CDC 7600 (single precision $\sim$15 significant figures) we have selected δ (the Newton step 'skipping parameter', equation (2.10)) to be 10^{-5} and the output of the $\boldsymbol{H}_0$ Choleski-Householder step is assumed to give eigenvalues good to 10^{-16}. It should be noted that round-off error does eventually destroy positive-definiteness of the overlap matrix $\boldsymbol{Q}$ in equation (3.8), but this typically occurs for $\geqslant 10$ iterations, for which the residual is typically $\leqslant 10^{-6}$ (10-12) significant figures in the eigenvalue).

References

Bendt P and Zunger A 1982a *Bull. Am. Phys. Soc.* **27** 248

—— 1982b *Solar Energy Research Institute Tech. Rep.* TP-212-1698

—— 1982c *Phys. Rev.* B **26** 3114

Bishop D M 1967 *Advances in Quantum Chemistry* vol 3, ed P Löwdin (New York: Academic) p 25
Brust D 1964 *Phys. Rev.* A **134** 1337
—— 1968 in *Energy Bands in Solids* ed B Alder, S Fernbach and M Rotenberg (New York: Academic) § III, p 40
Butscher W and Kammer W E 1976 *J. Comput. Phys.* **20** 313
Cohen M L 1970 *Solid State Physics* vol 24, ed H Ehrenreich, F Seitz and D Turnbull (New York: Academic)
Cooper J L B 1948 *Q. Appl. Math.* **6** 179
Davidson E R 1975 *J. Comput. Phys.* **17** 87
—— 1978 *Report on the Workshop 'Numerical Algorithms in Chemistry: Algebraic Methods' of the National Resource for Computation in Chemistry, Lawrence Berkeley Laboratory, 9-11 August* LBL-8158, UC-32, CONF-780878, p 15
Harrison W A 1980 *Electronic Structre and the Properties of Solids* (San Francisco: Freeman) chap 3.C, p 71
Huzinaga S 1963 *J. Chem. Phys.* **42** 1293
Jaffe J and Zunger A 1983 *Phys. Rev.* B **27** 1882
Lanczos C 1950 *J. Res. NBS* **45** 255
Liu B 1978 *Report on the Workshop 'Numerical Algorithms in Chemistry: Algebraic Methods' of the National Resource for Computation in Chemistry, Lawrence Berkeley Laboratory, 9-11 August* LBL-8158, UC-32, CONF-780878, p 49
Louie S G, Chelikowsky J R and Cohen M L 1977 *Phys. Rev.* B **15** 2154
Löwdin P 1951 *J. Chem. Phys.* **19** 1396
Nesbet R K 1965 *J. Chem. Phys.* **43** 311
—— 1981 in *Sparse Matrices and their Uses* ed I S Duff (New York: Academic) p 161
NRCC 1978 *Report on the Workshop 'Numerical Algorithms in Chemistry: Algebraic Methods', of the National Resource of Computation in Chemistry, Lawrence Berkeley Laboratory 9-11 August*, LBL-8158, UC-32, CONF-780878. (Available from the National Technical Service, Department of Commerce, 5285 Port Royal Road, Springfield, VA 22161; price code A08.)
Parlett B N 1980 *The Symmetric Eigenvalue Problem* (Englewood Cliffs, NJ: Prentice Hall)
Pulay P 1980 *Chem. Phys. Lett.* **73** 393
Raffenetti R C 1978 *Report on the Workshop 'Numerical Algorithms in Chemistry: Algebraic Methods', of the National Resource for Computation in Chemistry, Lawrence Berkeley Laboratory, 9-11 August* LBL-8158, UC-32, CONF-780878, p 29
—— 1979 *J. Comput. Phys.* **32** 403
Scott S 1981 in *Sparse Matrices and their Uses* ed I S Duff (New York: Academic) p 139
Shavitt I 1970 *J. Comput. Phys.* **6** 124
ten Hoor M J 1980 *Int. J. Quantum Chem.* **17** 1099
Wilkinson J H 1965 *The Algebraic Eigenvalue Problem* (Oxford: Clarendon) chap 4, § 42, p 229, chap 5, § 29, p 290
Ziman J M 1969 *Elements of Advanced Quantum Theory* (Cambridge: CUP) § 3.1, p 53
Zunger A 1979 *J. Vac. Sci. Technol.* **16** 1337
Zunger A and Cohen M L 1978 *Phys. Rev.* B **18** 5449
—— 1979 *Phys. Rev.* B **20** 4082