PASSING THE ONE-BILLION LIMIT
IN FULL CONFIGURATION-INTERACTION (FCI) CALCULATIONS

Jeppe OLSEN

Theoretical Chemistry, Chemical Center, University of Lund, P.O. Box 124, S-22100 Lund, Sweden

Poul JØRGENSEN

Department of Chemistry, Aarhus University, DK-8000, Aarhus C, Denmark

and

Jack SIMONS

Chemistry Department, University of Utah, Salt Lake City, UT 84112, USA

Full configuration-interaction calculations have been carried out using more than one-billion determinants. Such large eigenvalue calculations are possible because of advances in the direct CI technology and in the iterative technique used to solve the eigenvalue equations. The CPU time per direct CI iteration varies approximately linearly with the dimension of the matrix from one million to more than one billion. One direct CI iteration is found to take about 1.2–1.4 min per million determinants on an IBM 3090/VF.

1. Introduction

In ab initio computational chemistry [1], the solution of the time-independent Schrödinger equation for electronic states of atoms and molecules is expressed in terms of a double basis-set expansion. The one-particle molecular orbitals (MOs) are expanded in terms of a basis of *atomic orbitals* (AOs) and the *n*-particle wave functions are expressed in terms of a *determinant basis* constructed as antisymmetric products of molecular orbitals.

In principle, the *exact solution* to the Schrödinger equation can be obtained if the basis of AOs is complete and if all determinants that can be constructed by occupying the MOs in all possible manners are considered. Such calculations are, of course, impossible. However, the variational principle provides a framework in which one can approach such exact solutions in a systematic manner.

In actual calculations, one first choses a set of N one-electron functions that can represent the one- and

two-particle densities in an efficient manner. The variationally optimal solution in this one-electron basis is determined by solving the *n*-particle configuration-interaction (CI) *matrix eigenvalue equation* in the space of *all* determinants that can be formed within this basis. The resultant solutions are termed full configuration-interaction (FCI) solutions within the finite AO basis. The only error in a FCI calculation originates from the use of a finite one-electron basis. The size of the FCI eigenvalue problem is approximately $(N!/(n/2)! (N-n/2)!)^2$. Although the number of determinants can be somewhat reduced using spatial symmetry, the factorial growth with the number of orbitals N and the number of electrons n prohibits the use of FCI calculations for systems with many orbitals and electrons. In most practical applications, approximations are introduced by imposing chemically or physically motivated constraints and function-space truncations on the treatment of the *n*-particle basis.

Ab initio electronic-structure calculations are

therefore seen to suffer from errors due to truncation in the one-particle space and in the n-particle space. It is important to evaluate the independent contributions made by these two sources of error because we need to know how much effort should be focused on improving the one- and n-particle function-space approximations. The purpose of FCI calculations is to remove *all* error from the treatment of the n-particle space. In this way the accuracy of various one-particle AO bases can be assessed.

Essentially all models that have been developed to date utilize the same finite AO basis expansions. They differ primarily in how they treat the n-electron determinant aspect of the expansion. For this reason, it is especially relevant to make available benchmark data which (i) utilize AO bases that are representative of the present state of the art (which certainly changes as years pass) and (ii) which eliminate approximation of the n-electron space by including *all* determinants in the calculation. Such benchmark data are obtained by carrying out FCI calculations within a given atomic basis. These data constitute essential information for the calibration of approximate computational methods.

It is prohibitive to realize such benchmark results on large numbers of species. However, it is important that such calculations be performed on as many representative atoms and molecules as possible and that the scale of these benchmark calculations be pushed to larger and larger size as the best available computational power continues to grow so these benchmarks can relate to larger classes of molecules. An order-of-magnitude advance in such benchmarks will be demonstrated in this paper.

The importance of and potential for FCI benchmarks were emphasized in 1981 by Saxe et al. [2] who solved the first one-million-dimension FCI eigenvalue equation. In 1984 Knowles and Handy [3] developed a new approach to solving the FCI problem that allowed Bauschlicher et al. [4] to carry out a new-generation sequence of benchmark calculations having matrix dimensions up to 28 million [5]. These FCI benchmark calculations resulted in a much more thorough understanding of the correlation contributions to energies and properties. Although somewhat larger atomic basis set were considered in these studies, it remained difficult to use realistic present-day bases in such FCI studies. Benchmark

calculations have also been carried out by Bauschlicher at al. [4] on molecular properties other than the energy. Because even larger basis sets are required for describing such molecular properties, the relevance of these FCI results is restricted to even smaller atoms and molecules when viewed as benchmarks for properties. Harrison and Zarrabian [6] have reported timings for a 77-million-dimension FCI calculation, although no energies, convergence information, or other data was given. Recently, Knowles and Handy [7] have reported controlled approximations to FCI calculations with 210-million determinants using an algorithm of Knowles [8] where the sparsity of the solution vector is exploited to reduce memory and desk requirements and to speed up the timing of the linear transformation. Although the calculations presented by Knowles and Handy [7] are significantly smaller than the largest presented here, future developments in FCI calculations have to take directions similar to the ones proposed by Knowles [8]. The algorithmic advances presented here may help in improving the convergence characteristics of algorithms where the sparsity of the solution vector is used explicitly. In the reported calculations of Knowles and Handy [7], energies have been converged to 10^{-4} au.

In this paper we use a newly developed formalism of Olsen et al. [9] to carry out FCI benchmark calculations on one-billion-dimensional eigenvalue problems. In so doing we make available, for a large family of atoms and small molecules, the capability of correlating more electrons and treating considerably more realistic atomic basis sets. These atomic basis sets are, in many cases, also large enough to describe molecular properties other than the energy. Our benchmark results include total energies for the Ne and Mg atom.

FCI eigenvalue equations are solved using iterative techniques. For this purpose, we have used a technique developed by Olsen [10] which is an inverse-iteration-based generalization of the Davidson algorithm [11]. The technique is stable, rapidly convergent, and has a significantly lower disk-space requirement than the conventional technique it replaces. In the linear-transformation step, which is the central component of all iterative eigenvalue routines, the Hamiltonian matrix H is multiplied by a trial vector C. In our method, this step is carried out

using the new approach of Olsen et al. [9] which exploits the non-banded sparseness of H and computes only non-zero quantities.

The one-billion-determinant calculation was carried out using $D_{\infty h}$ point-group symmetry taking advantages of the simplification which occurs for a single state with $M_l = 0$. A similar size calculation may be carried out on a system with lower point-group symmetry with the same memory requirements if the FCI space is treated as a restricted active subspace (RAS) composed of several sub-blocks [9]. Each sub-block can then be treated in a similar way as the symmetry blocks of the present implementation (see ref. [9]). The larger number of symmetry unique coefficients puts additional requirements to disk space if the sparsity of the solution vector is not used explicitly.

In section 2 we describe the algorithm we have used to find the eigensolutions of the FCI eigenvalue equation. We discuss in section 2.1 the inverse-iteration generalisation of the Davidson algorithm. In section 2.2, we discuss the linear transformation involving a multiplication of the Hamiltonian matrix on a trial vector, and in section 2.3 we discuss the problems pertinent to large-scale calculations. In section 3, we report the results of FCI calculations on the Ne and Mg atoms. Section 4 contains our concluding remarks.

2. Theory

In section 2.1 we summarize the iterative technique of Olsen [11] which we have used to find the lowest eigensolutions of the FCI eigenvalue equation

$$HC = EC. \tag{1}$$

The basic features of the linear transformation described by Olsen et al. [9]

$$\sigma = HC \tag{2}$$

are summarized in section 2.2, and in section 2.3, the special problems related to large scale implementation are considered.

2.1. The inverse-iteration generalized Davidson technique

Following Olsen [10] let us begin by assuming that $C^{(0)}$ is the current iteration's approximation to the eigenvector, and let us define its energy as

$$E^{(0)} = (C^{(0)T}HC^{(0)}). \tag{3}$$

We next divide H into a zeroth-order term $H^{(0)}$ and a correction term $H^{(1)}$

$$H = H^{(0)} + H^{(1)}, \tag{4}$$

and write the FCI eigenvalue equation as

$$(H^{(0)} + H^{(1)})(C^{(0)} + C^{(1)})$$
$$= (E^{(0)} + E^{(1)})(C^{(0)} + C^{(1)}), \tag{5}$$

where $C^{(1)}$ and $E^{(1)}$ are correction terms to $C^{(0)}$ and $E^{(0)}$, respectively.

Requiring $C^{(1)}$ to be orthogonal to $C^{(0)}$ and neglecting terms in eq. (5) which are quadratic in the correction terms, one obtains:

$$C^{(1)} = -(H^{(0)} - E^{(0)})^{-1}$$
$$\times [(H - E^{(0)})C^{(0)} - E^{(1)}C^{(0)}], \tag{6}$$

where

$$E^{(1)} = \frac{C^{(0)T}(H^{(0)} - E^{(0)})^{-1}(H - E^{(0)})C^{(0)}}{C^{(0)T}(H^{(0)} - E^{(0)})^{-1}C^{(0)}}. \tag{7}$$

The first term on the right-hand side of eq. (6) gives, after the $C^{(0)}$ component is projected out, the correction vector of the conventional Davidson algorithm [11]. In the algorithm implied by eq. (6), such explicit projection is unnecessary because $C^{(1)}$ is automatically orthogonal to $C^{(0)}$. The correction vector in eq. (6) can alternatively be derived as an approximate Newton–Raphson iteration [10].

To understand the need for improving the convergence properties of the Davidson algorithm, consider the limit in which the matrix $H^{(0)}$ approaches the full H matrix. The first term on the right side of eq. (6) then gives $C^{(1)} = -C^{(0)}$; that is, the correction vector is directed along the previous iteration's $C^{(0)}$ vector. This clearly limits this algorithm's ability to introduce independent character into the trial vector. In contrast, in this same limit, eq. (6) yields

$$C^{(1)} = -C^{(0)} + \frac{(H^{(0)}-E^{(0)})^{-1}C^{(0)}}{C^{(0)\mathrm{T}}(H^{(0)}-E^{(0)})^{-1}C^{(0)}} . \tag{8}$$

Because this $C^{(1)}$ is rigorously orthogonal to $C^{(0)}$, it retains the potential to introduce new character into the trial vector in each iteration. Eq. (6) gives the step of an inverse-iteration process in which the $C^{(0)}$ component is projected out [12]. Because the step is automatically orthogonal to $C^{(0)}$, improved convergence is obtained with the preconditioning algorithm given in eq. (6). A more detailed analysis of the convergence characteristics of the algorithm will be reported in ref. [10]. In practice, we use a preconditioning in terms of a matrix $H^{(0)}$ consisting of a block of the Hamiltonian matrix formed within a determinant space defined by the p lowest diagonal elements of H. Outside this block, we use the diagonal elements of H to define $H^{(0)}$. The dimension p is set equal to 400 in the applications presented in this paper.

In the conventional algorithm, the sequence of vectors $\{v_k\}$ generated as correction $C^{(1)}$ vectors are used, together with $C^{(0)}$, to define a "reduced-space" and the projection of H onto this space is computed to find an optimal solution vector. To implement such an approach in large-matrix applications such as those considered here places stringent demands on disk storage because several long vectors (i.e. the $\{v_k\}$) need to be stored. In the algorithm used here, the above reduced-space is not used. Instead, each time a correction vector $C^{(1)}$ is formed, it is added to $C^{(0)}$ and the result is used as the next $C^{(0)}$ (after renormalization). This process has disk requirements that are substantially less than the conventional algorithm. The algorithm is converged when the norm of the residual is smaller than a prescribed threshold; the accuracy of the total energy is then proportional to the square of this residual.

2.2. Treatment of the linear transformations

The effect of the Hamiltonian operator acting on a determinant. In this section we follow closely the derivation of Olsen et al. [9] and refer to that reference for further details.

Each of the determinants $\{I\}$ used to define the H_{IJ} elements of H consists of a product of a so-called alpha-string and a beta-string [13]

$$|\alpha(I_\alpha)\beta(I_\beta)\rangle = \alpha(I_\alpha)\beta(I_\beta)|\mathrm{vac}\rangle . \tag{9}$$

An alpha-string $\alpha(I_\alpha)$ (beta-string $\beta(I_\beta)$), is an ordered product of n_α, (n_β), operators that create occupancy of molecular spin-orbitals with alpha (beta) spin [11]. The addressing scheme for the strings is described in detail in ref. [9].

Within the string basis, a trial vector C can be written as

$$|0\rangle = \sum_{I_\alpha I_\beta} C(I_\alpha, I_\beta)|\alpha(I_\alpha)\beta(I_\beta)\rangle . \tag{10}$$

Note that the trial vector C here is written in matrix form. The result of acting with the non-relativistic Hamiltonian H on $|0\rangle$ is represented in the string basis by the set of numbers $\sigma(I_\alpha, I_\beta)$:

$$\sigma(I_\alpha, I_\beta) = \sum_{J_\alpha J_\beta} \langle \beta(J_\beta)\alpha(J_\alpha)|H|\alpha(I_\alpha)\beta(I_\beta)\rangle$$
$$\times C(J_\alpha, J_\beta) . \tag{11}$$

Insertion of H into eq. (11) gives [9]

$$\sigma(I_\alpha, I_\beta) = \sigma_1(I_\alpha, I_\beta) + \sigma_2(I_\alpha, I_\beta) + \sigma_3(I_\alpha, I_\beta) , \tag{12}$$

where

$$\sigma_1(I_\alpha, I_\beta) = \sum_{J_\beta} \sum_{kl} \langle \beta(J_\beta)|E_{kl}^\beta|\beta(I_\beta)\rangle$$
$$\times \left(h_{kl} - \tfrac{1}{2}\sum_j (kj|jl)\right) C(I_\alpha, J_\beta)$$
$$+ \tfrac{1}{2}\sum_{J_\beta}\sum_{ijkl} \langle \beta(J_\beta)|E_{ij}^\beta E_{kl}^\beta|\beta(I_\beta)\rangle$$
$$\times (ij|kl)C(I_\alpha, J_\beta) , \tag{13a}$$

$$\sigma_2(I_\alpha, I_\beta) = \sum_{J_\alpha} \sum_{kl} \langle \alpha(J_\alpha)|E_{kl}^\alpha|\alpha(I_\alpha)\rangle$$
$$\times \left(h_{kl} - \tfrac{1}{2}\sum_j (kj|jl)\right) C(J_\alpha, I_\beta)$$
$$+ \tfrac{1}{2}\sum_{J_\alpha}\sum_{ijkl} \langle \alpha(J_\alpha)|E_{ij}^\alpha E_{kl}^\alpha|\alpha(I_\alpha)\rangle$$
$$\times (ij|kl)C(J_\alpha, I_\beta) , \tag{13b}$$

$$\sigma_3(I_\alpha, I_\beta) = \sum_{J_\alpha J_\beta}\sum_{ijkl} \langle \beta(J_\beta)|E_{ij}^\beta|\beta(I_\beta)\rangle$$
$$\times \langle \alpha(J_\alpha)|E_{kl}^\alpha|\alpha(I_\alpha)\rangle (ij|kl)C(J_\alpha, J_\beta) . \tag{13c}$$

and E_{kl} is a generator of the linear group which defines the string excitation operators

$$E_{kl} = E_{kl}^\alpha + E_{kl}^\beta . \qquad (14)$$

For states which have $M_s = 0$, the coefficient matrix obeys [9]

$$C(I_\alpha, I_\beta) = (-1)^S C(I_\beta, I_\alpha) , \qquad (15)$$

which implies that

$$\sigma_2(I_\alpha, I_\beta) = (-1)^S \sigma_1(I_\beta, I_\alpha) \qquad (16a)$$

$$\sigma_3(I_\alpha, I_\beta) = (-1)^S \sigma_3(I_\beta, I_\alpha) . \qquad (16b)$$

Eq. (16a) eliminates the calculation of σ_2 and eq. (16b) imposes the restriction $I_\alpha \geqslant I_\beta$ on the evaluation of σ_3. As a result, the computer operation count is reduced by a factor of two compared to the $M_s \neq 0$ case.

A vectorizable algorithm for setting up the σ_1 vector in eq. (13a) uses the fact that the Hamiltonian matrix elements are independent of I_α and is described in detail in ref. [9]. The second part of the σ vector, σ_2 of eq. (13b), can be constructed by an algorithm which is similar to the one for σ_1. The third part of the σ vector, σ_3 of eq. (13c), is less trivial to vectorize. A gather operation and a scatter operation can be introduced to avoid use of indirect addressing in the time-consuming part of the algorithm. The result is outlined below:

Loop over kl
    set up $L(I)$, $R(I)$, and $\mathrm{sgn}(I)$ defined by
    $|\alpha[L(I)]\rangle = E_{kl}^\alpha |\alpha[R(I)]\rangle \, \mathrm{sgn}(I)$
    $C'(I, J_\beta) = C[L(I), J_\beta] \, \mathrm{sgn}(I)$; vectorized
    gathering loop over I_β
    loop over excitations E_{ij}^β from $|\beta(I_\beta)\rangle$
      $|\beta(J_\beta)\rangle = \mathrm{sgn}\ (ij) E_{ij}^\beta |\beta(I_\beta)\rangle$
      $F(J_\beta) = F(J_\beta) + \mathrm{sgn}(ij) \, (ij|kl)$
    end of loop over E_{ij}^β
    $V(I) = \sum_{J_\beta} F(J_\beta) C'(I, J_\beta)$; vectorized over I
    $\sigma_3[R(I), I_\beta] = \sigma_3[R(I), I_\beta] + V(I)$; vectorized
    scattering
    end of loop over I_β
end of loop over kl.   (17)

The construction of $V(I)$ from F and C' is usually the most operation-intensive part of the construction of σ_3. A comparison with the computer opera-
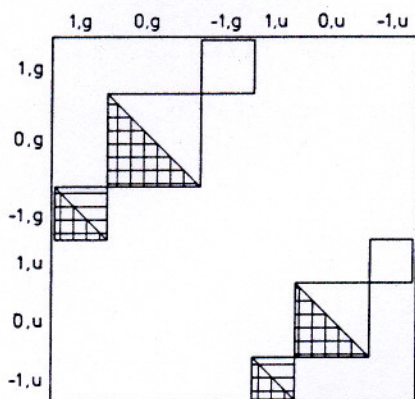

Fig. 1. The structure of the $C(I_\alpha, I_\beta)$ and $\sigma(I_\alpha, I_\beta)$ matrices for a $^1\Sigma_g$ state. The columns and rows label the symmetries of the alpha- and beta-strings. The symmetry labels are gerade (g) and ungerade (u) and M_l is restricted to 1, 0, −1. The empty blocks are determined by symmetry for a $M_s = 0$ state. If the state, in addition, has $M_l = 0$, only the squared triangles are symmetry unique.

tion count of previous FCI schemes is given in ref. [9].

The derivation can be modified also to incorporate spatial symmetries. In our program, we have implemented $D_{\infty h}$ point-group symmetry. The strings with a given M_l value are stored together. The strings with $M_l' \neq 0$ are ordered and scaled so that σ_v reflection of string I with $M_l = M_l'$ gives string I with $M_l = -M_l'$. This will be used below to identify the symmetry-unique coefficients for a $M_l = 0$ state. The CI coefficient matrix $C(I_\alpha, I_\beta)$, and the sigma vector $\sigma(I_\alpha, I_\beta)$ have non-vanishing elements for blocks where the sum of the M_l values for the alpha and beta strings are zero. In fig. 1 we have shown, for the simple case of a gerade state with M_l equal to zero and where the M_l values of the strings are restricted to +1, 0 and −1, the blocks that are different from zero. This block structure will be explored below. More details about the implementation of $D_{\infty h}$ symmetry can be found in ref. [14].

2.3. Large-scale calculations

Memory and disk requirements. In large-scale implementations of our algorithm, it is not possible to have all non-vanishing blocks of the coefficient matrix C and the sigma vector σ in main memory.

Without degredation in performance, the calculation can be organized so that only one symmetry block of the sigma vector is needed at any time. The contributions to a sigma vector block from the various blocks of the C matrix are completed, and written to disk before the next sigma block is processed.

The algorithm can be further modified to require also only one symmetry block of the coefficient matrix C in memory at any time. This requires that the coefficient matrix is processed block-by-block for each sigma block (so save input/output for as many sigma blocks as can be in main memory). The net result is that less main memory is required but at the expense of increased input/output (I/O). The maximum amount of necessary main memory is given in terms of the maximum dimension of one symmetry block of the coefficient matrix and the space for the corresponding block of the sigma vector, plus the memory required to store the gathered-coefficient matrix C'. In our calculations the C' matrix is never larger than half the size of a symmetry block of the coefficient matrix.

The storage requirements can be diminished even further for gerade states with $M_l=0$ and/or $M_s=0$ by keeping only the symmetry-unique coefficients on disk. For $M_s=0$ states, the fact that the coefficient matrix satisfies the symmetry relation in eq. (15) reduces the computer operation count by a factor of two as described in eq. (16). The symmetry relation may also be used to store the symmetry blocks of the C and σ arrays that are at or below the diagonal (see fig. 1). This reduces the requirement for input/output by nearly a factor of two.

For $M_l=0$ states, each block of the coefficient matrix has the additional symmetry

$$C(I_\alpha, I_\beta) = \Lambda C(I_\beta, I_\alpha), \tag{18}$$

where Λ is one for a Σ^+ state and minus one for a Σ^- state. This symmetry relation may also be used to reduce the operation count by two. For a gerade state with $M_s=0$ and $M_l=0$, all symmetry blocks of the C and σ arrays can be determined from the lower triangles of the symmetry blocks below the diagonal. For example, for the simple case considered in fig. 1, $M_s=0$ symmetry eliminates, for the gerade–gerade block, the need to store the $(1, g; -1, g)$ symmetry block and the upper half of the $(0, g; 0, g)$ symmetry block. The $M_l=0$ symmetry further elim-

inates the need for storing the upper half of the $(-1, g; 1, g)$ block.

The main memory requirement for a gerade $M_s=0$, $M_l=0$ state is therefore a matrix of half the maximum dimension of one symmetry block for the coefficient matrix and a similar matrix for the σ matrix. As the coefficient matrix is read in from disk, it is expanded to form a full square matrix to improve vectorization and to reduce paging. In addition, memory is required for the coefficient matrix C' (see eq. (17)). In fig. 2 we have sketched how the central part (σ_3) of the linear transformation is implemented in our computer program. The total requirement of main memory is $\frac{3}{2}$ matrices of the maximum dimension of one symmetry block of the coefficient matrix, plus storage for the C' matrix plus a minor amount originating from storing string information, etc. To illustrate the space requirements, we report the salient features of our one-billion-determinant calculation described later. The largest symmetry block contains 100280196 elements and the largest C' matrix contains 19647468 elements. Together with string information, etc., this adds up to a total memory of 197679895 double-precision words. Since the arrays, in general, are referenced sequentially, a large part of this memory only needs to be page-addressable extended storage. All of the calculations reported here were done with 64 Mbyte of directly addressable central storage. The symmetry-unique parts of the coefficient and sigma matrices are stored
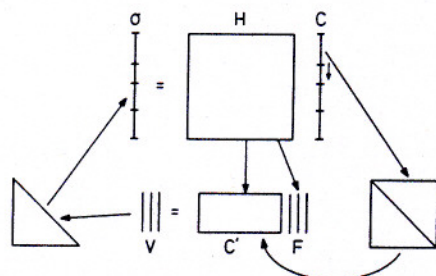

Fig. 2. A diagrammatic illustration of the construction of the σ_3 part of the linear transformation in eq. (12) for a gerade $M_S=0$ and $M_l=0$ state. The coefficients C are read in by symmetry blocks and expanded to square form. They are then gathered together with part of the elements of the Hamiltonian matrix to form the matrix C' which then are multiplied on the sparse vector F. The resulting vector V elements are scattered into one symmetry block of the σ vector.

blockwise on disk. The total length of a symmetry matrix for the one-billion-dimensional calculation was 303845939.

In the implementation of the inverse-iteration generalized Davidson technique, the C matrix, the σ matrix, and the block-diagonal Hamiltonian matrix $H^{(0)}$ were read from disk, in symmetry blocks requiring at no stage more than $\frac{3}{2}$ symmetry blocks in memory. The construction of the correction vector $C^{(1)}$ required, in total, 10 reads or writes of the symmetry-unique part of one of these vectors. In total, we thus have four vectors on disk of which one is a scratch file. One file can be eliminated by recalculating $H^{(0)}$ in each iteration.

3. Results

In table 1 we report pertinent data for FCI calculations on Ne and Mg. These data include the number of active orbitals and electrons and the number of determinants in each calculation. The data also include the iteration times for one direct CI and the memory requirements of each calculation. More details are given below for the individual calculations.

The new generation of benchmark calculations of Bauschlicher et al. [4] was initiated with FCI calculations of the ground-state total energy of the Ne atom [15]. The 1s orbital was frozen to a canonical Hartree–Fock orbital in these calculations and the largest calculation used a contracted [5s3p2d] basis with Dunning's [5s3p] contraction [16] of the Huzinaga (9s5p) primitive basis set [17] and with exponents of 4.5 and 1.3 for the two primitive d func-

tions. The largest calculation [15] was carried out in D_{2h} point-group symmetry and gave a total of 9805897 determinants. It took 22.8 min per direct CI iteration on the NAS CRAY2 and gave a total energy −128.768877 au. We carried out the same calculation in $D_{\infty h}$ spatial symmetry giving 5502533 determinants. Our calculation took 6.0 min per iteration on the IBM 3090/VF at the Utah Supercomputer Institute; thus, our calculation runs more than three times faster than the one reported by Bauschlicher et al., and less than 16 Mbyte of central storage was necessary.

As an example of the factorial growth in the number of determinants as the basis is extended, Bauschlicher et al. reported [15] that the addition of a single f-type basis function to their basis gives approximately 20-million configuration state functions in D_{2h} symmetry. We carried out this particular calculation using an f function with exponent 2.5 in $D_{\infty h}$ symmetry (see table 1). This produced 41201651 determinants and took 46.5 min per direct CI iteration to yield a total energy of −128.794567 au.

To obtain a more accurate total energy, it was reasonable to extend the basis of Bauschlicher et al. with an f function. However to describe excited states and polarizabilities, it is more important to add further diffuseness to the basis than to add the f function. To pursue this end, we used Dunning's [4s2p] contraction [16] of the Huzinaga (9s5p) primitive set [17] and added two diffuse s functions with exponents of 0.11 and 0.03, two diffuse p functions with exponents of 0.11 and 0.03, and three d functions with exponents of 2.15, 0.7, and 0.2. This basis is large

Table 1

Pertinent data for the reported FCI calculations

Atom	Ne	Ne	Mg
active electrons	8	8	10
active orbitals	4s3p2d1f	5s4p3d	4s3p2d1f
determinants in $D_{\infty h}$	41201651	86776694	1016018176
combinations in $D_{\infty h}$	12533335	27359950	303845939
determinants in D_{2h}	93896448	161650624	2538603520
iteration time (min)	47	107	1450
storage (Mbyte)	64	64	64
ext. storage (Mbyte)	0	100	1600

Table 2
Distribution of CI coefficients in the 86776694-determinant calculation on Ne

Range	Number of symmetry-unique coefficients	Weights	Accumulated weights
10^{-1} to 1	1	0.96449451	0.96449451
10^{-2} to 10^{-1}	88	$0.26294204 \times 10^{-1}$	0.99078871
10^{-3} to 10^{-2}	535	$0.85476007 \times 10^{-2}$	0.99933631
10^{-4} to 10^{-3}	8397	$0.57235998 \times 10^{-3}$	0.99990867
10^{-5} to 10^{-4}	65340	$0.87070165 \times 10^{-4}$	0.99999574
10^{-6} to 10^{-5}	344550	$0.40254007 \times 10^{-5}$	0.99999977
10^{-7} to 10^{-6}	1561430	$0.22545632 \times 10^{-6}$	0.99999999
10^{-8} to 10^{-7}	3991606	$0.67115194 \times 10^{-8}$	1.00000000

enough to give a reasonable description of the ground state, the lowest excited states of ^1S, ^3S, ^1P and ^3P symmetry, and of the polarizability. The number of determinants in $D_{\infty h}$ symmetry for this basis is 86776694 (see table 1). Each direct CI iteration took 106.5 min. and the total energy was -128.725161 au. These calculations thus demonstrate the feasibility of carrying out benchmark studies using basis sets that provide realistic descriptions of a multitude of physical properties; the results of our property calculations will be reported elsewhere [18].

In tables 2 and 3 some characteristics are given for the 86776694 determinant wave function. The wave function was converged to a residual norm 10^{-4}. The distribution of the symmetry-unique coefficients with norms in intervals from $1-10^{-1}$ to $10^{-7}-10^{-8}$ are given in table 2. It is seen that the 9021 symmetry-unique coefficients that are larger than 10^{-4} span 99.990807% of the wave function while the 74361 symmetry-unique coefficients larger than 10^{-5} span 99.999977% of the wave function. The number of symmetry-unique coefficients with values above a given threshold thus increases slowly when the magnitude of the threshold becomes smaller. The distribution of the coefficients in table 2 strongly suggests that future developments in FCI calculations should use the sparsity of the solution vector as suggested by Knowles and Handy [7,8]. In table 3 we report the weights for the various excitation levels from the Hartree–Fock determinant. As expected the doubly excitations give the major contributions. It is worth noticing that the contributions from the singles, triples and quadruples all are of the same magnitude.

To extend significantly the state of the art in FCI

Table 3
The weight of the excitation levels from the Hartree–Fock state in the 86776694-determinant calculation on Ne

Excitation level	Weights
0	0.9644945073
1	0.0009804929
2	0.0336865893
3	0.0003662339
4	0.0004517826
5	0.0000185090
6	0.0000017447
7	0.0000001393
8	0.0000000011

benchmark calculations, we performed a FCI study on the X ^1S ground state of the Mg atom. We used a contracted [5s3p2d1f] atomic orbital basis. This basis consists of McLean and Chandler's [19] contracted [53] sp basis to which we added two uncontracted d functions with exponents of 3.5 and 0.16 and an uncontracted f function with an exponent of 4.5. The exponents of the tighter d function and the f function were chosen so these orbitals would have radial extents close to those of the 2s and 2p SCF orbitals of Mg. The more diffuse d function's exponent was chosen to this orbital would have a radial extent close to that of the SCF 3s orbital. The 1s orbital was frozen with the expansion coefficients [1]. The frozen 1s orbital differs slightly from a canon-

---

[1] Ordering the s orbitals with decreasing maximal exponent, the 1s orbital has the coefficients (0.593091, 0.545060, −0.193317, 0.017461, 0.027162).

Table 4
Convergence characteristics of the 1016018176-determinant calculation for the X ^{1}S state of Mg (the SCF energy is -199.585214 au)

Iteration number	Total energy (au)	Residual norm
1	-199.612811	1.5419
2	-199.724369	0.2278
3	-199.726121	0.0681

ical Hartree–Fock 1s orbital.

The number of determinants in $D_{\infty h}$ symmetry is 1016018176, and each direct CI iteration required 1450 min of CPU time on the Utah IBM 3090 (see table 1). We have carried out three such direct CI iterations, and we report in table 4 the convergence characteristics of these iterations. The accuracy of the total energy is proportional to the square of the residual norm multiplied with an error constant that we, for CI calculations on states of this type, always have found to be smaller than 0.1. The error constants which are obtained in the corresponding single–double calculation (total energy -199.721401 au) indicate an accuracy of 0.0002 au in the total energy of iteration 3. Two or three additional iterations are needed to converge the energy to 10^{-6} au. Although each CI iteration needed 1450 min of CPU time, approximately one week of "wall time" elapsed per iteration because of paging among the page-addressable extended storage and the central storage. If we had had access to more than 64 Mbyte of central storage, we would have been able to exploit more efficiently the CPU; there simply was not enough central storage to allow us to do so when the blocks of the C matrix that needed to be handled contained 100-million double-precision words. In the smaller 86.8-million-dimension FCI calculation, where we also had 64 Mbyte of directly addressable main memory available, the corresponding block of the C matrix was of dimension 11.4 million, and hence the "wall" and CPU times were nearly identical. Nevertheless, even within the 64 Mbyte allocated to our calculation, we have been able to demonstrate that eigenvalue problems involving large randomly sparse matrices of dimension *one billion* are feasible using our algorithms.

It is interesting to note that our CPU times, for a few million to over one-billion determinants, are approximately proportional to the number of determinants. The number of operations increases faster than the number of determinants [9], but this increase is almost entirely offset by the increased vector lengths.

4. Discussion

We have shown that FCI eigenvalue calculations can be carried out using *more than one-billion* determinants. Such large matrix eigenvalue calculations are now possible because of the advances in the direct CI technology described here, which exploit the non-banded sparseness of the matrix, as well as improvements in the algorithm used to converge these iterative calculations. We make use of $D_{\infty h}$ symmetry while previous FCI benchmark calculations were able to use at most D_{2h} symmetry; this allows us to consider approximately half the number of determinants as in past work.

We find that the CPU time per direct CI iteration varies approximately linearly with the dimension of the matrix for dimensions from one million to more than one billion. The time required for one direct CI iteration is found to be approximately 1.2 to 1.4 min per million determinants. Previous FCI benchmark calculations [15] report needing approximately 2.3 min per million determinants on a Cray 2 computer, which has a cycle time of one third the IBM 3090s.

The memory requirement in the current implementation is a number of double-precision records about equal to one fifth of the number of determinants. The disk requirement is a number of double-precision words about equal to the number of determinants. FCI calculations with one-billion determinants can thus only be carried out on supercomputers with large memory, whereas calculations in the range of 100-million determinants can be carried out straightforwardly on most minisupercomputers.

Prior to this work, the largest FCI benchmark calculations calculating and storing the complete vectors treated 8 electrons distributed among 23 orbitals, yielding 28-million determinants. We have been able to handle 10 electrons in 30 orbitals with more than 1016-million determinants, and have thus ushered in a new era in FCI benchmarks where "exact"

energies can be obtained for more electrons and more realistic atomic-orbital bases. As a result, a more thorough and detailed understanding of the electron-correlation problem is at hand.

Acknowledgement

References

[1] K.P. Lawley, ed., Advances in chemical physics, Vols. 67, 69. Ab initio methods in quantum chemistry (Wiley, New York, 1987).

[2] P. Saxe, H.F. Schaefer III and N.C. Handy, Chem. Phys. Letters 79 (1981) 202.

[3] P.J. Knowles and N.C. Handy, Chem. Phys. Letters 111 (1984) 313;
P.E.M. Siegbahn, Chem. Phys. Letters 109 (1984) 417.

[4] C.W. Bauschlicher Jr., S.R. Langhoff and P.R. Taylor, Advances in chemical physics, in press.

[5] C.W. Bauschlicher Jr. and P.R. Taylor, J. Chem. Phys. 85 (1986) 2779.

[6] R.J. Harrison and S. Zarrabian, Chem. Phys. Letters 158 (1989) 393.

[7] P.J. Knowles and N.C. Handy, J. Chem. Phys. 91 (1989) 2396.

[8] P.J. Knowles, Chem. Phys. Letters 155 (1989) 513.

[9] J. Olsen, B.O. Roos, P. Jørgensen and H.J. Aa. Jensen, J. Chem. Phys. 89 (1988) 2185.

[10] J. Olsen, Chem. Phys. Letters, to be published.

[11] E.R. Davidson, J. Comput. Phys. 17 (1975) 87.

[12] B.N. Parlett, The symmetric eigenvalue problem (Prentice-Hall, New York, 1980).

[13] N.C. Handy, Chem. Phys. Letters 74 (1980) 280.

[14] J. Olsen and D. Sundholm, to be published.

[15] C.W. Bauschlicher Jr., S.R. Langhoff, P.R. Taylor and H. Partridge, Chem. Phys. Letters 126 (1986) 436.

[16] T.H. Dunning, J. Chem. Phys. 53 (1970) 2823.

[17] S. Huzinaga, J. Chem. Phys. 42 (1965) 1293.

[18] J. Olsen, P. Jørgensen and J. Simons, Chem. Phys. Letters, to be submitted for publication.

[19] A.D. McLean and G.S. Chandler, J. Chem. Phys. 72 (1980) 5639.