Linear scaling computation of the Fock matrix. III. Formation of the exchange matrix with permutational symmetry

Eric Schwegler[1,2] and Matt Challacombe
Theoretical Chemistry Division, Los Alamos National Laboratory, Los Alamos, New Mexico 87545

A direct comparison is made between two recently proposed methods for linear scaling computation of the Hartree-Fock exchange matrix in order to determine the importance of exploiting two-electron integral permutational symmetry. Calculations on three dimensional water clusters and graphitic sheets with different basis sets and levels of accuracy are presented to identify specific cases where permutational symmetry may or may not be useful. We conclude that a reduction in integrals via permutational symmetry does not necessarily translate into a reduction in computational times. For large insulating systems and weakly contracted basis sets the advantage of permutational symmetry is found to be negligible, while for non-insulating systems and highly contracted basis sets a four fold speedup is approached.

I. INTRODUCTION

Since the appearance of the first computer programs for *ab initio* electronic structure theory more than 30 years ago, considerable effort has been spent on developing new computational methods to carry out the calculations efficiently. This effort has been driven to a large extent by the desire to investigate the properties of extended molecular systems. Recently, significant progress in this regard has come from the so called linear scaling methods. A variety of new linear scaling methods now exist for computation of the Fock matrix [1–10], as well as optimization of the density matrix [11–15] that is encountered in Hartree-Fock (HF) [16] and density functional [17] theories.

As with the introduction of any radically new computational method it is often useful to reexamine the preexisting algorithms and techniques to determine if they should be retained or modified. In Paper II of this series, ONX was introduced, which is the first rigorous method for linear scaling computation of the HF exchange matrix [9]. ONX achieves linear scaling by adopting a novel loop structure that avoids potentially quadratic prescreening of small interactions, but does not employ two-electron integral permutational symmetry as is traditional [18–24]. Despite the fact that permutational symmetry can potentially save a factor of four in construction of the exchange matrix, ONX is highly competitive with standard direct SCF methods, and has been used in the largest HF calculations reported to date [9,15].

However, in Ref. [10] Ochenfeld *et al.* introduced a modification of ONX (LinK) that retains full integral permutational symmetry. Unfortunately, differences between the integral evaluation routines prevented a direct comparison between ONX and LinK. Nevertheless, Ochenfeld *et al.* claim an approximately three-fold speedup of LinK over ONX for large water clusters.

In order to actually measure the importance of integral permutational symmetry in linear scaling methods for construction of the HF exchange matrix, we have incorporated the methods discussed in Ref. [10] into ONX, which we refer to here as symmeterized ONX (SONX). By basing both ONX and SONX on identical integral evaluation routines, an analysis of the reduction in computed integrals as well as the corresponding reduction in CPU times is possible.

This paper is organized as follows. In the next section the permutational symmetry of electron repulsion integrals over Cartesian Gaussian basis functions is discussed and several issues that may affect the observed computational savings in large scale calculations are mentioned. In Section III, our implementation of integral permutational symmetry in SONX is described. Then in Sections IV A and IV B, a direct comparison between ONX and SONX is made and differences are given in terms of both the reduction in the number of evaluated integrals as well as in the reduction of CPU time.

II. PERMUTATIONAL SYMMETRY

It is common in modern electronic structure theory to use basis functions formed by a contraction of primitive functions,

$$\phi_a(\mathbf{r}) = \sum_i^{K_a} C_{ai}\varphi_{ai}(\mathbf{r}), \tag{1}$$

where C_{ai} are contraction coefficients, φ_{ai} are primitive basis functions, and K_a is the contraction length. For molecular computations, the primitive functions of choice are the Cartesian Gaussian type functions (CGTFs),

---

[1]University of Minnesota Supercomputer Institute, 1200 Washington Avenue South, Minneapolis, MN 55415
[2]Current address: Lawrence Livermore National Laboratory, Livermore, CA 94551
Cite this preprint as: LA-UR-99-577

$$\varphi_{ai}(\mathbf{r}) = (x-A_x)^l (y-A_y)^m (z-A_z)^n \times \exp\left[-\zeta_{ai}(\mathbf{r}-\mathbf{A})^2\right], \tag{2}$$

first proposed by Boys [25]. Primitive CGTFs are preferred because, as stated by the Gaussian product theorem (GPT), a product of CGTFs can be expressed exactly as a one-center finite sum of CGTFs. For example, the product of two s-type ($l=m=n=0$) primitive CGTFs is

$$\begin{aligned}\varphi_{ai}(\mathbf{r})\varphi_{bj}(\mathbf{r}) &= \exp\left[-\zeta_{ai}(\mathbf{r}-\mathbf{A})^2\right] \times \exp\left[-\zeta_{bj}(\mathbf{r}-\mathbf{B})^2\right] \\ &= \exp\left[-\zeta_{ai}\zeta_{bj}/(\zeta_{ai}+\zeta_{bj})(\mathbf{A}-\mathbf{B})^2\right] \\ &\quad \times \exp\left[-\zeta_p(\mathbf{r}-\mathbf{P})^2\right], \end{aligned}\tag{3}$$

where $\zeta_p = \zeta_{ai}+\zeta_{bj}$ and $\mathbf{P} = (\zeta_{ai}\mathbf{A}+\zeta_{bj}\mathbf{B})/\zeta_p$. Because of the GPT, it follows that the four-center electron repulsion integrals (ERIs),

$$[a_i b_j | c_k d_l] = \int\int d\mathbf{r} d\mathbf{r}' \varphi^*_{ai}(\mathbf{r})\varphi_{bj}(\mathbf{r})\frac{1}{|\mathbf{r}-\mathbf{r}'|}\varphi^*_{ck}(\mathbf{r}')\varphi_{dl}(\mathbf{r}'), \tag{4}$$

can be reduced to two-center integrals, which may be readily evaluated with a variety of analytic methods [25–28].

In addition to the GPT, since CGTFs are real, there is an eight-fold permutational symmetry between the indices a_i, b_j, c_k, and d_l in Eq. (4), and the following integrals are equal [18]:

$$\begin{aligned}[a_ib_j|c_kd_l] &= [b_ja_i|c_kd_l] = [a_ib_j|d_lc_k] \\ &= [b_ja_i|d_lc_k] = [c_kd_l|a_ib_j] \\ &= [d_lc_k|a_ib_j] = [c_kd_l|b_ja_i] \\ &= [d_lc_k|b_ja_i]. \end{aligned}\tag{5}$$

Equation (5) is also true for contracted ERIs,

$$(ab|cd) = \sum_i^{K_a}\sum_j^{K_b}\sum_k^{K_c}\sum_l^{K_d} C_{ai}C_{bj}C_{ck}C_{dl}[a_ib_j|c_kd_l]. \tag{6}$$

In formation of the HF exchange matrix

$$K_{ab} = \sum_{cd} D_{cd}(ac|bd), \tag{7}$$

rather than compute all possible ERIs, it is common practice to loop over just the unique ERIs [19,29] and to use the equalities shown in Eq. (5) to form the exchange matrix as

$$\begin{aligned} K_{ac} &= K_{ac} + D_{bd}\ (ab|cd)_s, \\ K_{ad} &= K_{ad} + D_{bc}\ (ab|cd)_s, \\ K_{bc} &= K_{bc} + D_{ad}\ (ab|cd)_s, \quad \text{and} \\ K_{bd} &= K_{bd} + D_{ac}\ (ab|cd)_s. \end{aligned}\tag{8}$$

The subscript s in Eq.(8) indicates that the ERI is scaled by a factor of 1, 1/2, 1/4, or 1/8 depending on the possible coincidence of basis function indices [20]. Since the majority of the ERIs involve unique indices, i.e. the ERI is scaled by 1, the use of Eq. (8) offers the possibility of saving a factor of four in the computation of $\mathbf{K}$.

In the original formulation of ONX presented in Ref. [9], this permutational symmetry of ERI indices was foregone in favor of a novel loop structure that enables loop skip-out statements. These loop skip-out statements are crucial to avoiding quadratically scaling prescreening of small ERIs and enables linear scaling computation of $\mathbf{K}$ for insulating systems. However, the recently developed LinK algorithm [10] demonstrates that it is possible to use the loop skip-out statements of ONX while at the same time taking advantage of ERI permutational symmetry. Although the advantage of exploiting ERI permutational symmetry has been well established in more traditional electronic structure methods [18–23], the situation is unclear for large scale calculations of $\mathbf{K}$ with new linear scaling methods.

One requirement that needs to be met for ERI permutational symmetry to be useful is that the magnitude of each contribution an ERI makes to $\mathbf{K}$ should decay at a similar rate, which depends on the structure of $\mathbf{D}$ as seen in Eq. (8). Specifically, this is due to the different ways small contributions to $\mathbf{K}$ are avoided in methods that do not use permutational symmetry as compared to those that do use symmetry. For example, in ONX small contributions to $\mathbf{K}$ are avoided with the criterion

$$|D_{cd}||(ac|ac)|^{1/2}|(bd|bd)|^{1/2} \leq \texttt{TwoENeglect}, \tag{9}$$

where the Schwartz inequality has been used [24]. However, in methods that exploit permutational symmetry all four of the contributions in Eq. (8) are considered together, and the criterion

$$\mathrm{Max}\left[|D_{bd}|, |D_{bc}|, |D_{ad}|, |D_{ac}|\right] \times |(ab|ab)|^{1/2}|(cd|cd)|^{1/2} \leq \texttt{TwoENeglect} \tag{10}$$

must be satisfied instead.

In Ref. [10], it is argued that in order for a given ERI, $(ab|cd)$, to be considered significant it is necessary that ϕ_a and ϕ_b (and likewise ϕ_c and ϕ_d) have significant overlap. This implies that the basis function pair ϕ_a and ϕ_b are spatially close. Since the ERI is contracted with D_{bd}, D_{bc}, D_{ad}, and D_{ac} in Eq. (8), this also implies that the index pairs of these four density matrix elements share similar (but not necessarily small) spatial separations. A central premise of Ref. [10] is that since the elements of $\mathbf{D}$ involve similar separations, it is not unreasonable to assume that they also have similar magnitudes. However, this assumption is contradicted by Figures 1 and 2 of Ref. [9], which demonstrate that elements of $\mathbf{D}$ sharing identical spatial separations can vary in magnitude by up to five orders.

```
do a = 1,N
   N_a^dis = 0
   do b = 1,N
      if (overlap of φ_a with φ_b > TwoENeglect) then
         N_a^dis = N_a^dis + 1
         add φ_b to distribution list {(ab|}_a
      endif
   enddo b
   sort distribution list {(ab|}_a by decreasing overlap
enddo a
```

FIG. 1. Ordering the list of significant distributions.

```
do a = 1,N
   N_a^den = 0
   do c = 1,N
      if( |D_ac| * |(aa|aa)|^(1/2) * |(cc|cc)|^(1/2)
      > TwoENeglect) then
         N_a^den = N_a^den + 1
         add φ_c to density list {D_ac}_a
      endif
   enddo c
   sort density list {D_ac}_a by
   decreasing |D_ac| * |(cc|cc)|^(1/2)
enddo a
```

FIG. 2. Ordering the list of significant density matrix elements.

Although the evaluation of ERIs dominates the formation of $\mathbf{K}$, there are additional steps that have a significant computational expense. For instance, once an ERI has been calculated it is necessary to contract it with the appropriate elements of $\mathbf{D}$ and accumulate the contributions to form $\mathbf{K}$. In some cases these steps, referred to here as integral digestion, can represent up to 20 percent of the total time needed to form $\mathbf{K}$ [30]. This is important because if the observed permutational symmetry factor is less than four, then in methods that use permutational symmetry there will be extra integral digestions that do not occur in methods like ONX. For example, if a given calculation reveals that using permutational symmetry only reduces the number of calculated ERIs by a half rather than the full factor of four, then there will also be a two-fold *increase* in the number of integral digestions. We also note that when sparse matrix storage routines are used to represent $\mathbf{D}$ and $\mathbf{K}$, which is a necessity for large-scale calculations, the relative cost of extra integral digestions increases because accessing elements of a sparse matrix is less efficient than for a dense matrix.

III. METHODS

In order to investigate the importance of exploiting ERI permutational symmetry in formation of the HF exchange matrix, we have incorporated ERI permutational symmetry as described in Ref. [10] into ONX, which we will refer to here as symmetrized ONX (SONX) to indicate that we have used the same ERI routines as in ONX. The integral evaluation routines in both ONX and SONX are based on the HGP method [27], and have been significantly improved over our previous implementations [9,31]. In the following, we briefly outline the main components of the method as it is implemented in SONX.

The loop skip-out statements of ONX can be combined with ERI permutational symmetry through ordered distribution lists in conjunction with ordered lists of density matrix elements. The formation of these ordered density lists, which were first proposed in Ref. [10], are repeated here in Figs. 1 and 2 for clarity.

As shown in Fig. 3, the construction of $\mathbf{K}$ begins with an outer loop over the sorted distribution list $\{(ab|\}_a$, where the subscript $\{...\}_a$ indicates the index that is held fixed. The density ordered list $\{D_{ac}\}_a$ is then used to select an index c, for which the ordered distribution list $\{|cd)\}_c$ is iterated over. Each distribution in the list $\{|cd)\}_c$ that satisfies the criterion

$$|D_{ac}||(ab|ab)|^{1/2}|(cd|cd)|^{1/2} \leq \texttt{TwoENeglect} \qquad (11)$$

is stored in a temporary ERI list specific to a. When Eq. (11) is no longer satisfied, the inner-most loop can be aborted safely because the distribution list has been ordered as in ONX. Due to the sorted density lists, the loop over index c may also be aborted safely when the inner loops do not alter the temporary ERI list. This process is repeated in a second set of loops except now the density ordered list corresponding to index b is used to select index d, and the resulting distributions $|cd)$ are stored in a temporary ERI list specific to b.

Once the temporary ERI lists have been formed, they are sorted and merged together so that all redundancies are removed. The combined ERI list is then iterated over and the ERIs $(ab|cd)$ are computed and contracted with elements of $\mathbf{D}$ to form $\mathbf{K}$.

In order to prevent memory requirements from increasing quadratically with system size, it is necessary to avoid the storage of insignificant elements of $\mathbf{D}$ and $\mathbf{K}$. Small elements of $\mathbf{D}$ and $\mathbf{K}$ are removed when the Frobenius norm of the corresponding atom block is less than `TrixNeglect` [15]. In ONX and SONX, an atom blocked sparse matrix representation is used, which enables the use of a highly optimized library of sparse matrix algebra routines [15]. This representation also minimizes the cost of accessing individual matrix elements because the ratio of matrix elements to atom block indices is large. Also, an efficient binary search over column indices is used, where the search is performed only over atom blocks rather than matrix elements.

```
do a = 1,N
   do b = 1,N_a^dis

      do c = 1,N_a^den
         do d = 1,N_c^dis
            if (|D_ac| * |(ab|ab)|^1/2 * |(cd|cd)|^1/2
            < TwoENeglect) then
               goto 100
            else
               add |cd) to ERI list a
            endif
         enddo d
100      continue
         if( ERI list a has not changed ) goto 200
      enddo c
200   continue

      do d = 1,N_b^den
         do c = 1,N_d^dis
            if (|D_bd| * |(ab|ab)|^1/2 * |(cd|cd)|^1/2
            < TwoENeglect) then
               goto 300
            else
               add |cd) to ERI list b
            endif
         enddo c
300      continue
         if( ERI list b has not changed ) goto 400
      enddo d
400   continue

      combine ERI list a and b into single ERI list
      sort and remove duplicates in ERI list
      do c,d = 1, Length of ERI list
         compute (ab|cd)
         contract (ab|cd) with D_bd, D_bc, D_ad, and D_ac
         add contributions to K_ac, K_ad, K_bc, and K_bd
      enddo c,d

   enddo b
enddo a
```

FIG. 3. The LinK algorithm as implemented in SONX.

TABLE I. The total number of ERIs calculated by ONX and SONX in a RHF/STO-2G calculation of a cluster of 30 water molecules for different levels of `TwoENeglect`. In each case, small elements in **D** and **K** are retained by setting `TrixtNeglect`=0.

	Number of ERIs		
`TwoENeglect`	ONX	SONX	Ratio
1.0e-5	1.7e6	1.0e6	1.7
1.0e-7	6.3e6	3.3e6	1.9
1.0e-9	1.6e7	7.5e6	2.1
1.0e-11	3.3e7	1.4e7	2.3
1.0e-13	5.7e7	2.2e7	2.6
1.0e-15	8.7e7	3.1e7	2.8

IV. RESULTS AND DISCUSSION

In the following sections we present several ONX and SONX benchmark calculations, which illustrate the speedups achieved by exploiting ERI permutational symmetry in computation of the HF exchange matrix. All calculations have been carried out within the MondoSCF [32] suite of programs on a single 332 MHz 604e PowerPC processor. Unless otherwise stated, the calculations have been performed with `TrixNeglect`$= 10^{-6}$, which results in errors in the total energy that are commensurate with those resulting from `TwoENeglect`$=10^{-7}$.

A. Reduction in ERIs

It is first of interest to determine to what extent utilizing permutational symmetry actually reduces the number of ERIs that need to be calculated. This can be determined by taking the ratio of the number of ERIs that are calculated by ONX versus the number calculated by SONX for identical systems and levels of accuracy.

To examine this ratio in detail we have performed a series of RHF/STO-2G calculations on a three dimensional cluster of 30 water molecules and on circumcoronene ($C_{25}H_{12}$) at different levels of integral prescreening, denoted by `TwoENeglect` [33]. As can be seen from Table I, the actual reduction in the number of ERIs for the water cluster with reasonable values of `TwoENeglect` (10^{-5} to 10^{-10}) is only one half of the full permutational factor. For the circumcoronene calculation, shown in Table II, the reduction in ERIs is increased and approaches the full permutational factor of four with similar values of `TwoENeglect`.

It appears that the reductions in ERIs for the cluster of 30 water molecules and circumcoronene differ substantially because the long-range behavior of **D** is controlled by the system's HOMO-LUMO gap [34–37]. We note that the choice of basis set can also have an effect on the HOMO-LUMO gap, which in turn alters the long-

TABLE II. The total number of ERIs calculated by ONX and SONX in a RHF/STO-2G calculation of circumcoronene, $C_{24}H_{12}$, for different levels of `TwoENeglect`. In each case, small elements in **D** and **K** are retained by setting `TrixtNeglect`=0.

	Number of ERIs		
`TwoENeglect`	ONX	SONX	Ratio
1.0e-5	1.6e7	5.8e6	2.8
1.0e-7	3.3e7	1.1e7	3.0
1.0e-9	5.0e7	1.5e7	3.3
1.0e-11	6.7e7	1.9e7	3.5
1.0e-13	8.3e7	2.3e7	3.6
1.0e-15	9.6e7	2.6e7	3.7

TABLE III. The total CPU time required to form one RHF exchange matrix for a cluster of 50 water molecules by ONX and SONX with different basis sets and with `TwoENeglect` $= 10^{-7}$. Speedup is the ratio of the ONX and SONX CPU times, and $\tilde{K}$ is the average basis set contraction length of ERIs.

	CPU Time (sec)			
Basis Set	ONX	SONX	Speedup	$\tilde{K}$
3-21G	91	86	1.1	1.4
6-31G**	979	854	1.1	1.4
6-31G*	599	491	1.2	1.5
Dunning DZ	621	474	1.3	1.8
6-31G	290	209	1.4	1.8
Dunning-Hay SV	697	462	1.5	1.9
STO-2G	20	13	1.5	1.9
STO-3G	90	50	1.8	2.6
STO-6G	1163	613	1.9	4.8

range exchange interactions [15]. However, similar calculations show that the observed reduction in ERIs is less dependent on the choice of basis sets than on the physical properties of the system.

B. Reduction in CPU time

Although ONX and SONX involve somewhat different approaches in terms of how the ERIs are looped over, both are based on the same implementation of the HGP method for computing ERIs [27]. As a result, a meaningful comparison can be made between the CPU times that each method requires.

In Table III, the total CPU time incurred by ONX and SONX to compute the RHF exchange matrix of a cluster of 50 water molecules is shown for the same series of basis sets as in Table I. The corresponding set of calculations for the circumcoronene system is given in Table IV. A comparison of the observed speedups obtained for these systems reveals that the reduction in CPU time is fairly dependent on the choice of basis sets. Given that the ratio of ERIs does not change significantly between different basis sets, this observation is somewhat surprising. However, it may be explained by noting that when contracted basis functions are used the computation of ERIs involves a four-fold summation over $\tilde{K}^4$ primitive ERIs, where $\tilde{K}$ is the average basis set contraction length. $\tilde{K}$ is less than K in Eq. (6) because products of primitive basis functions are discarded based on the criterion [27]

$$\left| C_{ai} C_{bj} \int d\mathbf{r}\, \varphi_{ai}(\mathbf{r}) \varphi_{bj}(\mathbf{r}) \right| \leq \texttt{DistNeglect}, \qquad (12)$$

where `DistNeglect`=`TwoENeglect`$\times 10^{-2}$. For basis sets that involve large values of $\tilde{K}$, the cost of evaluating individual ERIs rapidly increases as $\mathcal{O}(\tilde{K}^4)$. This is in contrast to other steps in formation of the exchange matrix, such as integral digestion, that do not depend on $\tilde{K}$ since they operate solely on contracted quantities. This is relevant to the comparison of ONX to SONX because when $\tilde{K}$ is large, any reduction in the number of ERIs via permutational symmetry will have a much larger effect on the total CPU time than when $\tilde{K}$ is small. In Tables III and IV, $\tilde{K}$ for each basis set is shown. Inspection of Tables III and IV reveals that the speedups obtained from using ERI permutational symmetry correlates well with $\tilde{K}$.

TABLE IV. The total CPU time required to form one RHF exchange matrix for circumcoronene by ONX and SONX with different basis sets and with `TwoENeglect` $= 10^{-7}$. Speedup is the ratio of the ONX and SONX CPU times, and $\tilde{K}$ is the average basis set contraction length of ERIs.

	CPU Time (sec)			
Basis Set	ONX	SONX	Speedup	$\tilde{K}$
3-21G	144	72	2.0	1.4
6-31G**	1971	783	2.5	1.4
6-31G*	1838	706	2.6	1.4
Dunning DZ	1316	588	2.2	1.9
6-31G	350	148	2.4	1.9
Dunning-Hay SV	1415	550	2.6	2.0
STO-2G	24	9	2.7	1.9
STO-3G	94	30	3.1	2.6
STO-6G	1076	315	3.4	4.7

V. CONCLUSIONS

In this paper, the effect of ERI permutational symmetry on linear scaling methods for construction of the HF exchange matrix has been examined. By incorporating the use of permutational symmetry in SONX, this effect has been measured in terms of the number of computed ERIs as well as the total CPU time required. An unexpected finding was that, depending on the basis set contraction length, a reduction in ERIs does not necessarily translate into a corresponding reduction in CPU time. A similar conclusion was not reached in Ref. [10] because their comparison was based solely on the number of computed ERIs rather than on the relative CPU times. Such a comparison can not take into account other computational steps such as integral digestion, which can represent a significant expense.

In systems that involve long range exchange interactions, such as graphitic sheets, the reduction in ERIs readily approaches a full factor of four. Although the observed CPU times remain basis set dependent, the use of permutational symmetry is advantageous in the case of quadratic scaling due to a vanishing HOMO-LUMO gap.

In systems such as three-dimensional clusters of water molecules, permutational symmetry is shown to reduce the number of ERIs computed by a factor of two. For basis sets that have long contraction lengths, a factor of two reduction in the CPU time can also be approached. However, when basis sets involving low levels of contraction are used any reduction in the number of computed ERIs is masked by the increased overhead needed to incorporate the permutational symmetry. These results point out the importance of comparing observed CPU times. In particular, for identical water clusters computed with the 3-21G basis set, the factor of three difference claimed in Ref. [10] is not observed, and in fact the two methods perform equally well.

Admittedly, use of a sparse matrix representation, not employed in Ref. [10], may contribute somewhat to the different conclusions reached here. Nevertheless, a sparse matrix representation is necessary for true linear scaling of the SCF method, and overhead of the atom blocked sparse matrix representation [15] has been minimized as much as possible through optimization of the atom block lookup. Thus, the conclusions reached here have been obtained with the same integral code, and within the context of an optimized and fully linear scaling SCF program [32].

Finally, we conclude by noting that forgoing the use of permutational symmetry may be preferred for a number of additional reasons not discussed here. For example, even with highly contracted basis sets, use of the multipole approximation to avoid contraction can dramatically alter the balance between the cost of ERI evaluation and digestion [38]. Moreover, parallel implementations of ONX are expected to preserve data locality of **D** and **K** to a much higher extent than in SONX.

VI. ACKNOWLEDGMENTS

This work has been supported by the ASCI program at Los Alamos National Laboratory, and by the University of Minnesota Supercomputer Institute.

[1] C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, Chem. Phys. Letters **230**, 8 (1994).
[2] M. C. Strain, G. E. Scuseria, and M. J. Frisch, Science **271**, 51 (1996).
[3] M. Challacombe, E. Schwegler, and J. Almlöf, in *Computational Chemistry: Review of Current Trends*, edited by J. Leczszynski (World Scientific, Singapore, 1996), pp. 53–107.
[4] M. Challacombe, E. Schwegler, and J. Almlöf, J. Chem. Phys. **104**, 4685 (1996).
[5] E. Schwegler and M. Challacombe, J. Chem. Phys. **105**, 2726 (1996).
[6] J. C. Burant, G. E. Scuseria, and M. J. Frisch, J. Chem. Phys. **105**, 8969 (1996).
[7] J. P. Dombroski, S. W. Taylor, and P. M. W. Gill, J. Phys. Chem. **100**, 6272 (1996).
[8] M. Challacombe and E. Schwegler, J. Chem. Phys. **106**, 5526 (1997).
[9] E. Schwegler, M. Challacombe, and M. Head-Gordon, J. Chem. Phys. **106**, 9703 (1997).
[10] C. Ochsenfeld, C. A. White, and M. Head-Gordon, J. Chem. Phys. **109**, 1663 (1998).
[11] X. P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10891 (1993).
[12] W. Yang and T. S. Lee, J. Chem. Phys. **103**, 5674 (1995).
[13] E. Hernández, M. J. Gillan, and C. M. Goringe, Phys. Rev. B **53**, 7147 (1996).
[14] G. Galli, Cur. Op. Sol. State Mat. Sci. **1**, 864 (1996).
[15] M. Challacombe, J. Chem. Phys. **110**, 2332 (1999).
[16] A. Szabo and N. S. Ostlund, *Modern Quantum Chemistry* (Mc Graw–Hill Inc., New York, 1989).
[17] R. G. Parr and W. Yang, *Density-functional theory of atoms and molecules* (Oxford University Press, Oxford, 1989).
[18] H. Preuss and G. Diercksen, Int. J. Quant. Chem. **1**, 605 (1967).
[19] R. Raffenetti, Chem. Phys. Letters **20**, 335 (1973).
[20] P. Redington, Comp. and Chem. **8**, 43 (1984).
[21] A. Duke, Chem. Phys. Letters **13**, 76 (1972).
[22] F. Billingsley, Int. J. Quant. Chem. **6**, 617 (1972).
[23] J. Almlöf, K. Faegri, and K. Korsell, J. Comp. Chem. **3**, 385 (1982).
[24] M. Häser and R. Ahlrichs, J. Comp. Chem. **10**, 104 (1989).
[25] S. F. Boys, Proc. Roy. Soc. **A200**, 542 (1950).
[26] L. McMurchie and E. R. Davidson, J. Comp. Phys. **26**, 218 (1978).
[27] M. Head-Gordon and J. A. Pople, J. Chem. Phys. **89**, 5777 (1988).
[28] P. M. W. Gill, Adv. Quantum Chem. **25**, 141 (1994).
[29] For example, the indicies can be restricted to: $a \geq b$, $c \geq d$, and $a(a-1)/2 + b \geq c(c-1)/2 + d$.
[30] J. Almlöf and P. Taylor, in *Advanced Theories and Computational Approaches to the Electronic Structure of Molecules*, edited by C. Dykstra (D. Reidel Publishing Company, Boston, 1984), pp. 107–125.
[31] B.G. Johnson, P.M.W. Gill, M. Head-Gordon, C.A. White, J. Baker, D.R. Maurice, M. Challacombe, E. Schwegler, T.R. Adams, J. Kong, M. Oumi, C. Ochsenfeld, N. Ishikawa, R.D. Adamson, J.P. Dombroski, R.L.

Graham, and E.D. Fleischmann, Q-CHEM, Version 1.0.2, *Q*-Chem Inc., Pittsburgh, PA (1997).
[32] M. Challacombe and E. Schwegler, MondoSCF, a suite of programs for linear scaling SCF theory, unpublished.
[33] Coordinate files available upon request.
[34] The difference between the eigenvalues of the highest occupied and the lowest unoccupied molecular orbitals.
[35] W. Kohn, Int. J. Quantum Chem. **56**, 229 (1995).
[36] P. Maslen *et al.*, J. Phys. Chem. A **102**, 2215 (1998).
[37] S. Goedecker, Rev. Mod. Phys., submitted.
[38] E. Schwegler and M. Challacombe, 1999, submitted.