Linear scaling computation of the Fock matrix. IV. Multipole accelerated formation of the exchange matrix

Eric Schwegler, and Matt Challacombe

Citation: The Journal of Chemical Physics **111**, 6223 (1999); doi: 10.1063/1.479926
View online: https://doi.org/10.1063/1.479926
View Table of Contents: http://aip.scitation.org/toc/jcp/111/14
Published by the American Institute of Physics

Linear scaling computation of the Fock matrix. IV. Multipole accelerated formation of the exchange matrix

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

(Received 9 February 1999; accepted 21 June 1999)

A new method for the multipole evaluation of contracted Cartesian Gaussian-based electron repulsion integrals is described, and implemented in linear scaling methods for computation of the Hartree–Fock exchange matrix. The new method, which relies on a nonempirical multipole acceptability criterion [J. Chem. Phys. **109**, 8764 (1998)], renders the work associated with integral evaluation independent of the basis set contraction length. Benchmark calculations on a series of three-dimensional water molecule clusters and graphitic sheets with highly contracted basis sets indicate that the new method is up to 4.6 times faster than a well optimized direct integral evaluation routine. For calculations involving lower levels of contraction a factor of 2 speedup is typically observed. Importantly, the method achieves these large gains in computational efficiency while maintaining numerical equivalence with standard direct self consistent field theory. © *1999 American Institute of Physics.* [S0021-9606(99)30435-9]

I. INTRODUCTION

A number of recent advances in computational methods has greatly extended the range of systems that can be investigated with *ab initio* electronic structure theory.[1–6] By removing several limiting bottlenecks, electronic structure calculations of unprecedented size are now possible on common workstation class computers.[5] Many of these new methods depend on the efficient use of the multipole approximation to calculate interactions. For example, hierarchical expansion methods such as treecodes[5] or fast multipole methods[3,4] use the multipole approximation to achieve a reduction in the computational complexity of Coulomb matrix formation found in Hartree–Fock (HF) and density functional theories. In addition, new approaches have been described for the rapid evaluation of individual electron repulsion integrals (ERIs) with the multipole approximation.[7–10]

Significant computational savings can be achieved in multipole based methods because expensive direct integral evaluations are replaced with less expensive multipole interactions. Since the multipole expansions are truncated at a finite order, the ability to accurately determine acceptability of a given expansion plays an important role in both the computational efficiency and numerical accuracy of the method. However, the development of error estimates needed to effectively do this have, until recently, received little attention.[7]

In the following, we describe the use of the multipole approximation in conjunction with a multipole acceptability criterion (MAC)[7] and a penetration acceptability criterion (PAC) to effectively accelerate the computation of contracted ERIs. The choice of MAC and PAC is made with a particular emphasis on balancing computational simplicity with the required numerical accuracy. Implementations in linear scaling methods for computing the HF exchange matrix[6] demonstrate the computational efficiency of this approach. By eliminating much of the work associated with basis set contraction, this new method is highly competitive for computation of the HF exchange matrix for both large systems and highly contracted basis sets.

The paper is organized as follows: In the following section the multipole limit of ERIs are reviewed. An approach for translating primitive distributions to common centers to avoid explicit ERI contraction is discussed in Sec. II A, and efficient application of MAC and PAC are described in Secs. II B and II C. An implementation of these ideas in the linear scaling methods ONX (order *N* exchange) and SONX (symmetric order *N* exchange) are presented in Sec. III. Then in Secs. III A and III B, results for a series of calculations on three-dimensional clusters of water molecules and graphitic sheets are given for a variety of basis sets. Finally, in Sec. IV we present our conclusions.

II. MULTIPOLE EXPANSION OF ERIS

In molecular electronic structure theory the basis functions of choice are the Cartesian Gaussian-type functions (CGTFs)[11]

$$\varphi_{ai}(\mathbf{r}) = (r_x - A_x)^{l_a}(r_y - A_y)^{m_a}(r_z - A_z)^{n_a} \times \exp[-\zeta_{ai}(\mathbf{r}-\mathbf{A})^2], \tag{1}$$

which are typically contracted

$$\phi_a(\mathbf{r}) = \sum_i^{K_a} C_{ai}\varphi_{ai}(\mathbf{r}). \tag{2}$$

Use of contracted basis functions leads to expensive fourfold loops

[a]University of Minnesota Supercomputer Institute, 1200 Washington Avenue South, Minneapolis, MN 55415. Current address: Lawrence Livermore National Laboratory, Livermore, CA 94551.

$$(ab|cd)=\sum_{i}^{K_a}\sum_{j}^{K_b}\sum_{k}^{K_c}\sum_{l}^{K_d} C_{ai}C_{bj}C_{ck}C_{dl}[a_ib_j|c_kd_l], \tag{3}$$

over primitive electron repulsion integrals (ERIs)

$$[a_ib_j|c_kd_l] = \int\int d\mathbf{r}d\mathbf{r}'\,\varphi_{ai}(\mathbf{r})\varphi_{bj}(\mathbf{r})|\mathbf{r}-\mathbf{r}'|^{-1}\varphi_{ck}(\mathbf{r}')\varphi_{dl}(\mathbf{r}'). \tag{4}$$

Over the past 20 years, many different methods have been developed to efficiently compute contracted ERIs.[12–18] For example, by employing specialized recursion relations some fraction of the ERI evaluation can be moved outside of the contraction summation.[16,18] In addition, it is possible to reduce the average contraction length ($\tilde{K}$) of basis function products or distributions

$$\phi_a(\mathbf{r})\phi_b(\mathbf{r})=\sum_{i}^{K_a}\sum_{j}^{K_b} C_{ai}C_{bj}\varphi_{ai}(\mathbf{r})\varphi_{bj}(\mathbf{r}) \tag{5}$$

by prescreening with the criterion[16]

$$\left|C_{ai}C_{bj}\int d\mathbf{r}\varphi_{ai}(\mathbf{r})\varphi_{bj}(\mathbf{r})\right| \leqslant \texttt{DistNeglect}, \tag{6}$$

or by modeling basis function products with expansions that involve a smaller number of primitive functions.[18,19] These methods, though, still result in a significant amount of computational work that grows as $\mathcal{O}(\tilde{K}^4)$.

In Ref. 7, an attractive approach for avoiding the contraction problem was introduced that is based on the multipole approximation. By decoupling well separated distributions $[a_ib_j|$ and $|c_kd_l]$ in Eq. (4), the multipole approximation enables basis set contractions to be performed independently of ERI evaluation.

In order to develop the multipole limit of Eq. (4), it is convenient to represent primitive distributions (products of CGTFs) with exact one-center expansions in Hermite Gaussian type functions (HGTFs)

$$C_{ai}C_{bj}\varphi_{ai}(\mathbf{r})\varphi_{bj}(\mathbf{r})=\sum_{L=0}^{l_a+l_b}\sum_{M=0}^{m_a+m_b}\sum_{N=0}^{n_a+n_b} e^{ab}_{LMN}\Lambda^{r_p}_{LMN}(\mathbf{r}), \tag{7}$$

where e^{ab}_{LMN} are the expansion coefficients obtained by the McMurchie–Davidson two term recurrence relations[20] and multiplied by the appropriate radial overlap and contraction coefficients. The HGTFs centered at $\mathbf{r}_p$ are

$$\Lambda^{\mathbf{r}_p}_{LMN}(\mathbf{r})=\frac{\partial^L}{\partial r^L_{p_x}}\frac{\partial^M}{\partial r^M_{p_y}}\frac{\partial^N}{\partial r^N_{p_z}}\exp[-\zeta_p(\mathbf{r}-\mathbf{r}_p)^2], \tag{8}$$

where $\zeta_p=\zeta_{ai}+\zeta_{bj}$ and $\mathbf{r}_p=(\zeta_{ai}\mathbf{A}+\zeta_{bj}\mathbf{B})/\zeta_p$. In the multipole limit, HGTFs are related to derivatives of the Dirac delta function by

$$\delta^{\mathbf{r}_p}_{LMN}(\mathbf{r})=\frac{\partial^L}{\partial r^L_{p_x}}\frac{\partial^M}{\partial r^M_{p_y}}\frac{\partial^N}{\partial r^N_{p_z}}\delta(\mathbf{r}-\mathbf{r}_p)=\lim_{\zeta\to\infty}\left(\frac{\zeta_p}{\pi}\right)^{3/2}\Lambda^{\mathbf{r}_p}_{LMN}, \tag{9}$$

which enables ERIs over HGTFs to be evaluated with the Cartesian multipole interaction tensor[21,22] as

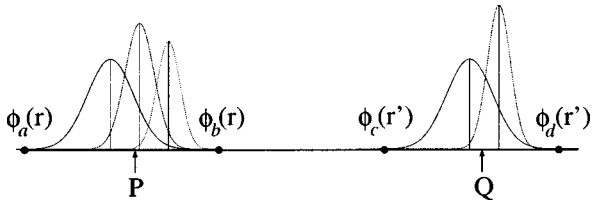

FIG. 1. The expansion centers **P** and **Q**. Primitive distributions are translated to the common centers **P** and **Q** in order to avoid contraction of primitive ERIs.

$$(\Lambda^{\mathbf{r}_p}_{LMN}|\Lambda^{\mathbf{r}_q}_{L'M'N'})=(-1)^{L+M+N}\left(\frac{\pi}{\zeta_p}\right)^{3/2}\left(\frac{\pi}{\zeta_q}\right)^{3/2} \times\frac{\partial^{L+L'}}{\partial r^{L+L'}_{pq_x}}\frac{\partial^{M+M'}}{\partial r^{M+M'}_{pq_y}}\frac{\partial^{N+N'}}{\partial r^{N+N'}_{pq_z}}\frac{1}{|\mathbf{r}_p-\mathbf{r}_q|}. \tag{10}$$

This representation, which is accurate when the effects of penetration are small due to the separation of $\mathbf{r}_p$ and $\mathbf{r}_q$, decouples the ERI allowing the multipole moments, $(\pi/\zeta_p)^{3/2}$ and $(\pi/\zeta_q)^{3/2}$, of primitive distributions located on the same center to be summed prior to ERI evaluation.

A. Translation of primitive distributions

Computing individual ERIs with a multipole representation can be advantageous for a variety of reasons. For example, spherical multipole interaction tensors can be used that only involve $\mathcal{O}(\mathcal{L}^2)$ work as compared to $\mathcal{O}(\mathcal{L}^4)$ work for a Cartesian representation, where $\mathcal{L}$ is the total angular symmetry. However, since most basis function products involve primitive distributions located on more than one center, the use of a multipole representation alone will not reduce the complexity of ERI evaluation with respect to the contraction length $\tilde{K}$. In order to avoid expensive ERI contractions, the multipole moments of primitive distributions need to be re-expressed through translation and accumulation about common centers, as illustrated by the simple one-dimensional system shown in Fig. 1. The centers **P** and **Q** are chosen so that the average translation distance within each set of primitive distributions is minimized. For those contracted distributions $(ab|$ and $|cd)$ that only involve well separated primitive distributions, the corresponding ERIs can be evaluated with a bipolar multipole expansion about **P** and **Q**. *Since multipole moments corresponding to contracted distributions can be precomputed prior to the ERI evaluation loops, the computational complexity of ERI evaluation is reduced from* $\mathcal{O}(\tilde{K}^4)$ *to* $\mathcal{O}(1)$.

A truncated bipolar expansion of the term $|\mathbf{r}_p-\mathbf{r}_q|^{-1}$ in Eq. (10) about the centers **P** and **Q** can be carried out in many different representations.[22–26] We use a real arithmetic implementation involving solid spherical harmonics, which has been simplified by removing redundancies with the symbolic manipulation features of *Mathematica*[27] along with the FORTRANASSIGN utility in FORMAT.M.[28] The resulting expansions are expressed as explicit FORTRAN77 statements, which operate at peak efficiencies on modern computational platforms.[5]

B. Multipole acceptability criterion

In order to ensure that a given contracted ERI can be safely represented with the multipole approximation, it is necessary to apply a multipole acceptability criterion (MAC) to estimate the error involved in translating basis function products with truncated multipole expansions.[5,7,21] In Ref. 7, the first numerically tight MAC was developed that is directly applicable to moments of angular symmetry beyond spherical. For ERIs that involve CGTFs of angular symmetry l_{ab} and l_{cd}, where $l_{ab}=l_a+m_a+n_a+l_b+m_b+n_b$, the MAC is computed as

$$\mathrm{MAC}[l_{ab},l_{cd}]=\sum_{i=0}^{l_{ab}}\sum_{j=0}^{l_{cd}} c_i^{ab}c_j^{cd}(|\mathcal{M}_{r_p}^{l_p}[i,j]|+|\mathcal{M}_{r_q}^{l_q}[j,i]|), \tag{11}$$

with

$$c_i^{ab}=\max_{L+M+N=i}[|e_{LMN}^{ab}|]\left(\frac{\pi}{\zeta_p}\right)^{3/2}, \tag{12}$$

and where l_p and l_q are truncation orders of the bipolar multipole expansion. The terms $\mathcal{M}_{r_p}^{l_p}[i,j]$ in Eq. (11) can be evaluated efficiently by recurrence with

$$\begin{aligned}
\mathcal{M}_{r_p}^{l_p}[0,0]&=\frac{|\mathbf{r}_p-\mathbf{P}|^{l_p+1}}{|\mathbf{r}_q-\mathbf{P}|^{l_p+2}},\\
\mathcal{M}_{r_p}^{l_p}[i,0]&=\mathcal{M}_{r_p}^{l_p}[i-1,0]\frac{(l_p-i+2)}{|\mathbf{r}_p-\mathbf{P}|},\quad i>0\\
\mathcal{M}_{r_p}^{l_p}[i,j]&=\mathcal{M}_{r_p}^{l_p}[i,j-1]\frac{(l_p+j+1)}{|\mathbf{r}_q-\mathbf{P}|},\quad j>0.
\end{aligned} \tag{13}$$

Although numerically tight, Eq. (11) is too expensive to apply to every ERI because it requires computation of the distance between interacting basis function products. We have adopted a scheme to reduce the cost of applying MAC by clustering ERIs into groups. The clustering method is described below in terms of the linear scaling method ONX for computation of the Hartree–Fock exchange matrix.[6]

In ONX, elements of the exchange matrix

$$K_{ab}=\sum_{cd} D_{cd}(ac|bd), \tag{14}$$

are formed through ordered lists of distributions $\{(ac|\}_c$ and $\{|bd)\}_d$ where the indices c and d are fixed and a and b run over all possible distributions of a particular angular symmetry and contraction length. In the calculation of large molecular systems, when the centers of the basis functions ϕ_c and ϕ_d are far apart, the sets $\{(ac|\}_c$ and $\{|bd)\}_d$ form two spatially distinct groups. Instead of applying a different MAC to each pair $(ac|$ and $|bd)$, a single MAC can be computed between each $(ac|$ and the entire set $\{|bd)\}_d$. This is accomplished by determining the smallest sphere that encompasses all of the centers in $\{|bd)\}_d$. In the evaluation of Eq. (11), the following substitutions are then made:

$$|\mathbf{r}_q-\mathbf{P}|,\quad |\mathbf{r}_p-\mathbf{Q}|\Rightarrow|\mathbf{P}-\mathbf{S}|-r_s, \tag{15}$$

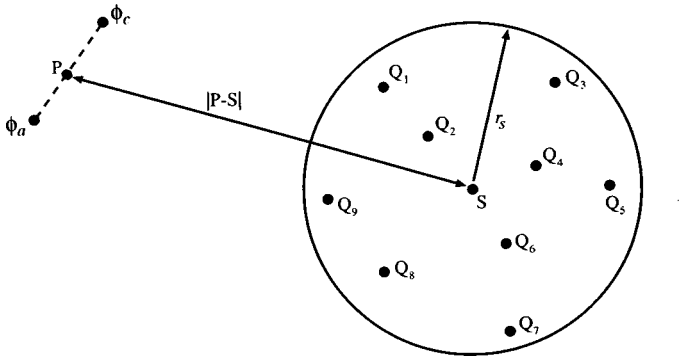

FIG. 2. The clustering scheme used for application of the MAC. **P** and $\mathbf{Q}_n$ are expansion centers corresponding to distributions $(ac|$ and $\{|bd)\}_d$, respectively. The sphere with center **S** and radius r_s is used in the MAC to compute the worst case error estimate for the entire set $\{|bd)\}_d$.

where **S** is the center of the sphere and r_s is its radius. In addition, the translation distance $|\mathbf{r}_q-\mathbf{Q}|$ and the moment c_i^{bd} are replaced by their maximum values over $\{|bd)\}_d$. This clustering, which is illustrated in Fig. 2, enables calculation of a single MAC for each $(ac|$, which represents the worst-case error in the interactions with $\{|bd)\}_d$.

Since the basis function products within the sphere shown in Fig. 2 are restricted to have the same index d, the MAC can be significantly improved by weighting it with the magnitude of D_{cd}, similar to density weighted ERI estimates found in direct self consistent field (SCF) methods.[29] By including the magnitude of D_{cd}, the MAC is adaptive and becomes more permissive when the size of the contribution to K_{ab} decreases.

C. Penetration acceptability criteria

In addition to MAC, a penetration acceptability criterion (PAC) must be satisfied in order for an ERI to be accurately calculated in a multipole representation.[21] Unlike point charges common in classical physics, CGTF charge distributions have extents, which can lead to contributions that the multipole expansion cannot reproduce. These errors are significant when the tails of interacting distributions overlap.

As first suggested in Ref. 7, PAC that account for contributions from higher order penetration terms have recently been developed by Kudin and Scuseria[30] in the context of the fast multipole method, where high order multipole expansions lead to the dominance of penetration errors. In the present context of individual ERI evaluation using low order expansions (equal to the original ERI angular symmetry), we have found that multipole errors dominate penetration errors, and that a simple s-s type PAC, outlined in the following, provides sufficient error control. These assertions are born out in Sec. III, where error control in converged total energies is demonstrated in several numerical experiments.

Using the same notation as in Ref. 7, the s-s type penetration error is

$$\epsilon=e_{000}^{ac}e_{000}^{bd}\left[\frac{2\pi^{5/2}}{\zeta_p\zeta_q\sqrt{\zeta_p+\zeta_q}}F_0(\tau)-\frac{\pi^3}{(\zeta_p\zeta_q)^{3/2}}\frac{1}{|\mathbf{r}_p-\mathbf{r}_q|}\right], \tag{16}$$

where

```
do a = a_f, a_l
   P = expansion center of (ac|
   S = center of sphere around {|bd)}_d
   r_s = radius of sphere
   if( |P - S| - r_s < 0 ) goto 352
   compute MAC
   if(MAC >TwoENeglect) goto 352
   τ = ζ_p ζ_q/(ζ_p + ζ_q)(|P - S| - r_s)^2
   if(τ >20.0) goto 351
   compute PAC
   if(PAC >TwoENeglect) goto 352
351 continue
   mark (ac| for multipole ERI evaluation
   goto 353
352 continue
   mark (ac| for direct ERI evaluation
353 continue
enddo a
```

FIG. 3. Applying the MAC and PAC in MAONX. The above code is evaluated immediately following the loop skip-out statement in ONX, the values of **P**, **S**, r_s, and the majority of the MAC can be precomputed at the level of distributions.

$$\tau = \frac{\zeta_p \zeta_q}{\zeta_p + \zeta_q} |\mathbf{r}_p - \mathbf{r}_q|^2. \tag{17}$$

In Eq. (16)

$$F_0(\tau) = \frac{1}{2}\sqrt{\frac{\pi}{\tau}}\mathrm{erf}(\sqrt{\tau}) = \int_0^1 e^{-\tau u^2} du, \tag{18}$$

is the zeroth order reduced incomplete gamma function[31,32] and e^{ac}_{000} is the same as in Eq. (7). By introducing the complementary error function, the s-s PAC is

$$\mathrm{PAC} = c_0^{ac} c_0^{bd} \left(\frac{1}{|\mathbf{r}_p - \mathbf{r}_q|} \right) \mathrm{erfc}(\sqrt{\tau}), \tag{19}$$

where c_0^{ac} are the distribution moments defined in Eq. (12). It should be reiterated that while this simple PAC has been found to provide sufficient error control over a range of systems and basis sets, the methods outlined here may be readily enhanced with PACs that account for angular symmetry.[30]

For the same reasons mentioned in Sec. II B for the MAC, the PAC is too expensive to apply to every ERI. We again use the clustering scheme described above by making the substitution

$$|\mathbf{r}_p - \mathbf{r}_q| \Rightarrow |\mathbf{r}_p - \mathbf{S}| - r_s, \tag{20}$$

in Eqs. (17) and (19). Figure 3 illustrates how the MAC and PAC are incorporated into ONX. The pseudo-code shown in Fig. 3 is evaluated immediately after the line "`301 continue`" in Fig. 4 of Ref. 6.

To further reduce the computational cost, simple estimates are placed before the actual MAC and PAC that screen out ERIs, which are obviously either in the multipole or direct regime. For example, in Fig. 3, the conditional

$$\texttt{if}(|\mathbf{P}-\mathbf{S}| - r_s < 0)\,\texttt{goto}\,352, \tag{21}$$

is used to skip the MAC and PAC when $(ac|$ is inside the $\{|bd)\}_d$ sphere and

$$\texttt{if}(\tau > 20.0)\,\texttt{goto}\,351, \tag{22}$$

is used to skip the PAC when separations are too large for penetration to be an issue.

III. RESULTS AND DISCUSSION

The multipole based ERI methods described above have been implemented in both ONX[6] and SONX,[33] which we will refer to as multipole accelerated ONX (MAONX) and multipole accelerated SONX (MASONX). SONX[33] is a linear scaling method for computing the HF exchange matrix similar to ONX, but includes ERI permutational symmetry as outlined in Ref. 34. The direct ERI evaluation routines used in ONX and SONX, which are based on the HGP (Head-Gordon–Pople) method,[16] have been significantly improved over our previous implementations.[6,35] In particular, the evaluation of ERIs involving d-type basis functions has been highly optimized.

MAONX and MASONX calculations within MONDOSCF[36] were performed with a variety of basis sets on a series of three-dimensional water clusters and graphitic sheets. In each case, the multipole expansions have been truncated at the same order as the total angular symmetry of the distributions involved (l_{ac} and l_{bd}), and the thresholding values used in the MAC and PAC were set equal to `TwoENeglect`, which is the parameter also used in thresholding ERIs by the Schwartz inequality.[37] For primitive distribution prescreening, `DistNeglect=TwoENeglect` $\times 10^{-2}$ was used in Eq. (6). In each case, the contracted distribution moments have been precomputed and are included in the reported CPU timings. All CPU timings were obtained on a single 332 MHz 604e PowerPC processor.

There are two key factors that determine to what extent formation of the exchange matrix can be accelerated by use of the multipole approximation. These factors are the percentage of all interactions that can be accurately computed in a multipole representation, and the relative speed of calculating an interaction with multipole versus direct ERI evaluation.

A. Percentage of multipole based ERIs

The percentage of ERIs computed with the multipole approximation in MAONX is shown in Fig. 4 for a series of three-dimensional water clusters. *Both the MAC and the PAC used in Fig. 4 do not include the magnitude of* D_{cd}. The relative ordering of the percentage of multipole based ERIs with different basis sets is a result of the differences in contraction lengths. As the basis set contraction length increases, the multipole representation involves the translation of a

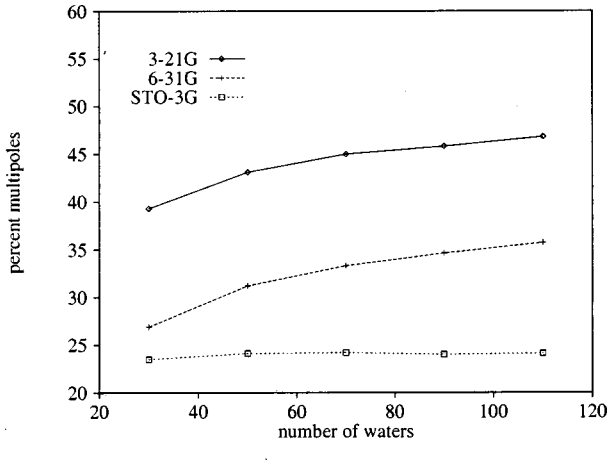

FIG. 4. The percentage of ERIs that can be accurately computed with the multipole approximation in a RHF calculation of the exchange matrix for a series of three-dimensional water clusters. The MAC and PAC *do not* include the magnitude of the density matrix, and the multipole truncation error corresponds to `TwoENeglect` $=10^{-7}$.

larger number of primitive distributions to common centers, which results in longer average translation distances and fewer ERIs that satisfy the MAC.

As mentioned in Sec. II B, the MAC and PAC can be improved by weighting each estimate with the magnitude of D_{cd} that the ERI is contracted with. The percentage of multipole based ERIs that are allowed by a density weighted MAC and PAC for the same series of three-dimensional water clusters is shown in Fig. 5. This percentage reaches 75 to 80 percent of the ERIs as system sizes are increased. Similar results are obtained with MASONX. Based on these observed trends, if the time needed to compute a multipole interaction can be made small relative to direct ERI evaluation then, at best, a four to five fold decrease in the total CPU time can be expected.

The percentages shown in Fig. 5 were obtained with `TwoENeglect` $=10^{-7}$. When smaller values of `TwoENeglect` are used the percentages only slightly decrease. For example, when `TwoENeglect` $=10^{-9}$, the percentage of multipole based ERIs decreases to 65 percent in the cluster of 110 water molecules with the 6-31G basis.

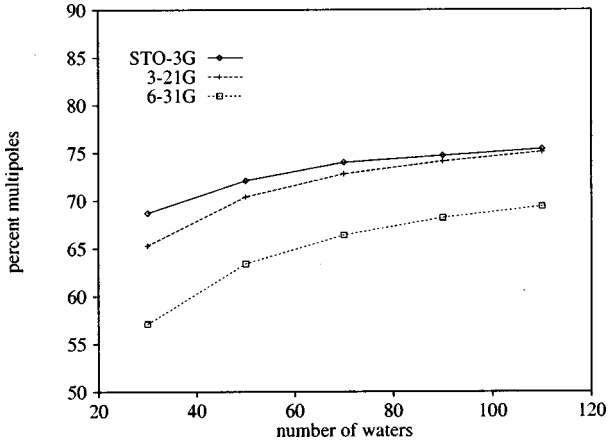

FIG. 5. The percentage of ERIs that can be accurately calculated with the multipole approximation in a RHF calculation of the exchange matrix for a series of three-dimensional water clusters. The MAC and PAC are weighted by the magnitude of the corresponding density matrix element, and the multipole truncation error corresponds to `TwoENeglect` $=10^{-7}$.

TABLE I. The observed speedup of multipole ERI evaluation with ERI permutational symmetry (MASONX) and without permutational symmetry (MAONX) in the calculation of the RHF (restricted Hartree–Fock) exchange matrix of a three-dimensional cluster of 50 water molecules. The average basis set contraction length is $\tilde{K}$.

	Speedup		
Basis set	MAONX	MASONX	$\tilde{K}$
3-21G	1.3	1.0	1.4
6-31G**	1.5	1.1	1.4
6-31G*	1.6	1.2	1.4
STO-2G	1.5	1.2	1.9
6-31G	1.9	1.3	1.9
Dunning DZ	2.1	1.4	1.9
Dunning-Hay SV	2.3	1.6	2.0
STO-3G	2.6	2.1	2.6
STO-6G	4.6	4.1	4.7

It is interesting to note that the basis set ordering in Fig. 5 is different than in Fig. 4. In particular, the percentage of multipole based ERIs is greatly increased in the STO-3G (Slater-type orbitals) calculations when density weighted MAC and PAC are used. This can be explained by observing that the HOMO–LUMO (highest occupied molecular-orbital–lowest unoccupied molecular orbitals) gap[38] is overestimated when a minimal basis, such as STO-3G, is employed.[39] Because a larger HOMO–LUMO gap corresponds to a faster decay in the density matrix,[2,40,41] which is included in the density weighted MAC and PAC, the percentage of ERIs that fall in the multipole regime is dramatically increased.

B. Speed of multipole evaluation

The observed speed of multipole versus direct ERI evaluation depends on a number of issues, including the various implementational details of each method. However, the most significant factor is the number of primitive ERIs that are avoided in a multipole representation. This, of course, is related to the average basis set contraction length $\tilde{K}$. Since the computational cost of direct ERI evaluation grows as $\mathcal{O}(\tilde{K}^4)$, the benefits of multipole evaluation, which is $\mathcal{O}(1)$ with respect to $\tilde{K}$, can be quite large.

MAONX and MASONX results for the cluster of 50 water molecules calculated with different basis sets are summarized in Table I. Speedup refers to the ratio of ONX CPU time to MAONX CPU time and SONX CPU time to MASONX CPU time. For the highest contracted basis set used, STO-6G, both MAONX and MASONX approach the expected four to fivefold speedup. In Figs. 6 and 7, accelerations commensurate with those in Table I are observed for the series of water clusters for the STO-3G and 6-31G* basis sets.

Systems with small HOMO–LUMO gaps, such as graphitic sheets, are particularly challenging because interactions persist for much longer distances than in water

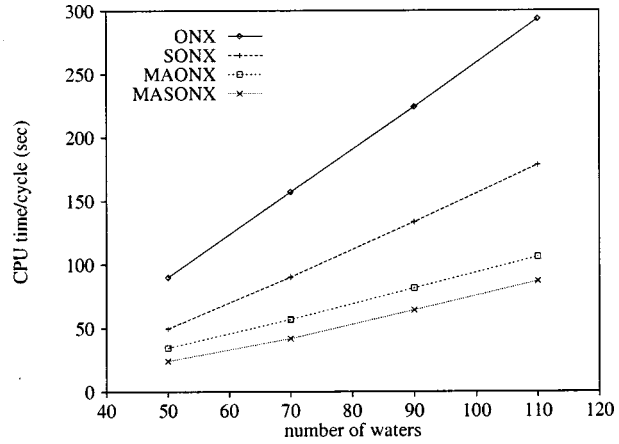

FIG. 6. RHF/STO-3G water cluster timings with `TwoENeglect` $=10^{-7}$.

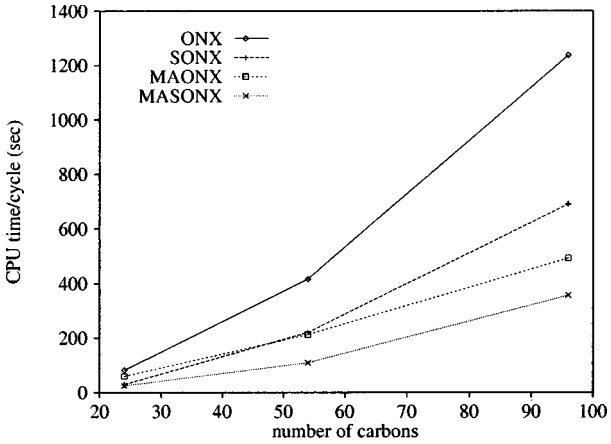

FIG. 8. RHF/STO-3G graphitic sheet timings with `TwoENeglect` $=10^{-7}$. Note that linear scaling is not observed for systems such as graphitic sheets, because their small HOMO–LUMO gaps lead to long-range exchange interactions.

clusters.[6] For these types of systems, ONX and SONX no longer exhibit linear scaling, and the computational expense increases rapidly with system size. However, multipole acceleration is still applicable, and as shown in Figs. 8 and 9, can greatly increase the efficiency of exchange matrix formation.

Inspection of Table I reveals that the observed MAONX speedup is always larger than in MASONX, and for basis sets with small contraction lengths MAONX is often the fastest method available. This is caused by various overheads and extra computations that are unavoidable when ERI permutational symmetry is employed.[33]

In Figs. 10 and 11, the absolute error in the total energy is shown for the series of three-dimensional water clusters.[42] It is important to note that the errors incurred by MAONX are indistinguishable from the integral neglect errors associated with the Schwartz inequality used by ONX. These results demonstrate the power of a numerically tight MAC, the sufficiency of using a *s*-*s* type PAC in this context, and establishes the equivalence of MAONX and MASONX with standard methods of direct SCF.[29,37]

IV. CONCLUSIONS

Multipole accelerated (MA) versions of the linear scaling exchange methods ONX and SONX that employ a nonempirical multipole acceptability criterion (MAC) have been developed. MAONX and MASONX effectively combine multipole and direct integral evaluation while maintaining numerical equivalence with standard direct SCF.

An important and novel finding of this work is that adaptive, density weighted error estimates must be used with the MAC and PAC to obtain significant acceleration in linear scaling evaluation of exchange ERIs.

As expected, numerical experiments with MAONX and MASONX result in large accelerations when highly contracted basis sets are used. For example, MAONX is 4.6 times faster than ONX for computing the exchange matrix of water molecule clusters with highly contracted basis sets. For typical basis sets, a factor of two speedup is observed. While MAONX and MASONX will scale quadratically for systems with a vanishing gap like the graphitic sheets, it is expected that for sufficiently large systems an asymptotic regime will be entered in which the number of ERIs computed with the multipole approximation approaches 100%. In this regime an acceleration proportional to $\tilde{K}^4$ can be expected.

While the nonempirical MAC has demonstrated its use here for the evaluation of exchange ERIs, the methods outlined are general and may be readily extended to the compu-

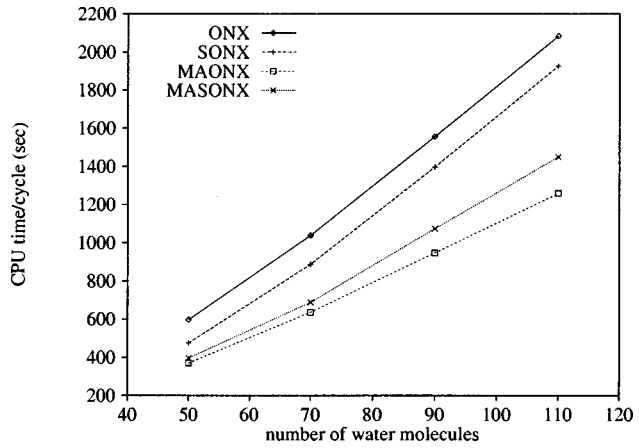

FIG. 7. RHF/6-31G* water cluster timings with `TwoENeglect` $=10^{-7}$.

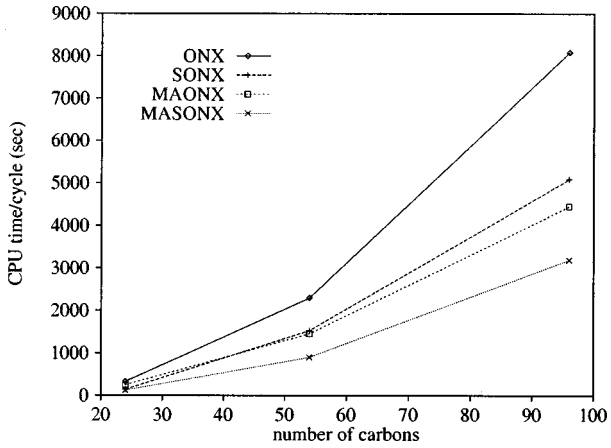

FIG. 9. RHF/6-31G graphitic sheet timings with `TwoENeglect` $=10^{-7}$.

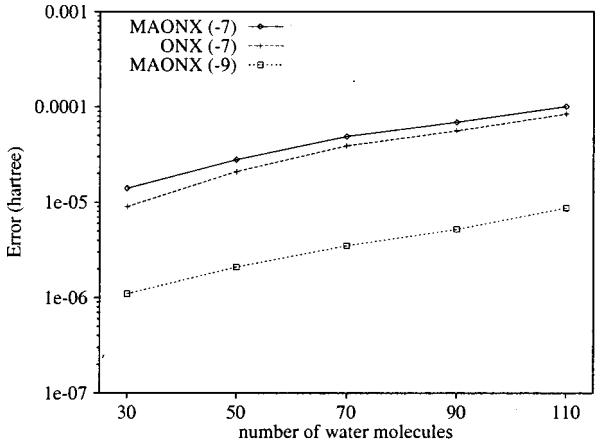

FIG. 10. Absolute errors in the converged RHF total energies for a series of water clusters with the 3-21G basis set. The value within the parenthesis, n, corresponds to the level of $\texttt{TwoENeglect} = 10^{-n}$ used.

tation of ERIs in other contexts, such as MP2 (second-order Moller–Plesset) theory,[43] or the computation of Coulomb sums with the quantum chemical tree code (QCTC).[5] The only caveat is that a simple s-s type PAC may or may not be adequate in these contexts, depending certainly on the order of multipole expansion used. Nevertheless, the necessary machinery for constructing more sophisticated PACs is established,[30] and we see no impediment for an increasing prevalence of tightly bounded multipole algorithms in quantum chemistry.

Finally, we conclude by noting that although the multipole approximation can be used to greatly accelerate the computation of contracted integrals, it does not reduce the overhead needed to incorporate permutational symmetry. As a result, multipole acceleration has a relatively larger effect in methods that do not use permutational symmetry than in those that do. In some cases, such as the calculation of water clusters with the 3-21G or 6-31G* basis, this effect results in MAONX speeds that are faster than MASONX. Similar but less pronounced effects were observed in a comparison of ONX and SONX in Paper III of this series.[33] It will be interesting to further explore the differences between MAONX and MASONX in parallel implementations where the penalty for nonlocality of density and exchange matrices should grow.

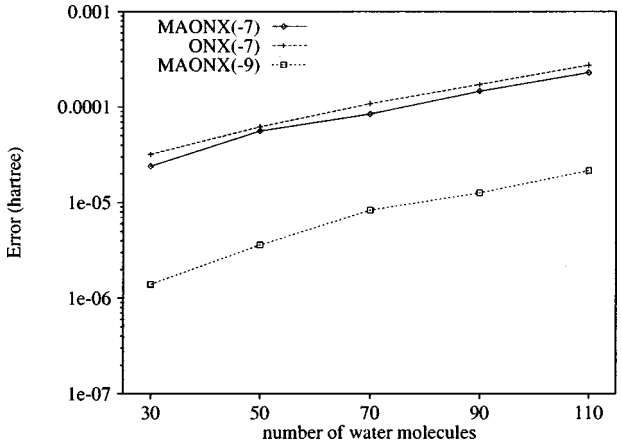

FIG. 11. Absolute errors in the converged RHF total energies for a series of water clusters with the 6-31G* basis set. The value within the parenthesis, n, corresponds to the level of $\texttt{TwoENeglect} = 10^{-n}$ used.

ACKNOWLEDGMENTS

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

[1] G. Galli, Curr. Opin. Solid State Mater. Sci. **1**, 864 (1996).
[2] S. Goedecker, Rev. Mod. Phys. (submitted).
[3] C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, Chem. Phys. Lett. **230**, 8 (1994).
[4] M. C. Strain, G. E. Scuseria, and M. J. Frisch, Science **271**, 51 (1996).
[5] M. Challacombe and E. Schwegler, J. Chem. Phys. **106**, 5526 (1997).
[6] E. Schwegler, M. Challacombe, and M. Head-Gordon, J. Chem. Phys. **106**, 9703 (1997).
[7] E. Schwegler, M. Challacombe, and M. Head-Gordon, J. Chem. Phys. **109**, 8764 (1998).
[8] J. Cioslowski, Comput. Phys. Commun. **53**, 117 (1989).
[9] I. Panas, J. Almlöf, and M. W. Feyereisen, Int. J. Quantum Chem. **40**, 797 (1991).
[10] T. R. Adams, R. D. Adamson, and P. M. W. Gill, J. Chem. Phys. **107**, 124 (1997).
[11] S. F. Boys, Proc. R. Soc. London, Ser. A **200**, 542 (1950).
[12] M. Dupuis, J. Rys, and H. F. King, J. Chem. Phys. **65**, 111 (1976).
[13] J. A. Pople and W. J. Hehre, J. Comput. Phys. **27**, 161 (1978).
[14] L. McMurchie and E. R. Davidson, J. Comput. Phys. **26**, 218 (1978).
[15] S. Obara and A. Saika, J. Chem. Phys. **84**, 3963 (1985).
[16] M. Head-Gordon and J. A. Pople, J. Chem. Phys. **89**, 5777 (1988).
[17] P. Gill, B. Johnson, and J. Pople, Int. J. Quantum Chem. **40**, 745 (1991).
[18] P. M. W. Gill, Adv. Quantum Chem. **25**, 141 (1994).
[19] P. M. W. Gill, B. G. Johnson, J. A. Pople, and S. W. Taylor, J. Chem. Phys. **96**, 7178 (1992).
[20] L. McMurchie, Ph.D. thesis, University of Seattle, 1977.
[21] M. Challacombe, E. Schwegler, and J. Almlöf, J. Chem. Phys. **104**, 4685 (1996).
[22] M. Challacombe, E. Schwegler, and J. Almlöf, Chem. Phys. Lett. **241**, 67 (1995).
[23] B. C. Carlson and G. S. Rushbrooke, Proc. Cambridge Philos. Soc. **46**, 626 (1950).
[24] J. Cipriani and B. Silvi, Mol. Phys. **45**, 259 (1982).
[25] I. Panas, Int. J. Quantum Chem. **53**, 255 (1995).
[26] W. W. Schulz, Chem. Phys. Lett. **254**, 337 (1996).
[27] S. Wolfram, *Mathematica, A System for Doing Mathematics by Computer*, 2nd ed. (Addison-Wesley, Redwood City, CA, 1991).
[28] M. Sofroniou, The Mathematica Journal **3**, 74 (1993).
[29] J. Almlöf, K. Faegri, and K. Korsell, J. Comput. Chem. **3**, 385 (1982).
[30] K. Kudin and G. E. Scuseria, J. Chem. Phys. **111**, 2351 (1999).
[31] I. Shavitt, Methods Comput. Phys. **2**, 1 (1963).
[32] P. J. Davis, in *Handbook of Mathematical Functions*, 9th ed., edited by M. Abramowitz and I. A. Stegun (Dover, New York, 1987), pp. 260–262.
[33] E. Schwegler and M. Challacombe, Theor. Chem. Accts. (submitted).
[34] C. Ochsenfeld, C. A. White, and M. Head-Gordon, J. Chem. Phys. **109**, 1663 (1998).
[35] Q-CHEM, Version 1.0.2., 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 Inc., Pittsburgh, PA (1997).
[36] MONDOSCF, M. Challacombe and E. Schwegler, a suite of programs for linear scaling SCF theory (unpublished).
[37] M. Häser and R. Ahlrichs, J. Comput. Phys. **10**, 104 (1989).
[38] The difference between the eigenvalues of the highest occupied and the lowest unoccupied molecular orbitals.
[39] M. Challacombe, J. Chem. Phys. **110**, 2332 (1999).
[40] W. Kohn, Int. J. Quantum Chem. **56**, 229 (1995).
[41] P. Maslen *et al.*, J. Phys. Chem. A **102**, 2215 (1998).
[42] All reported errors were computed by comparison with ONX where $\texttt{TwoENeglect} = 10^{-10}$.
[43] G. Hetzer, P. Pulay, and H. J. Werner, Chem. Phys. Lett. **290**, 143 (1998).