Linear scaling computation of the Fock matrix[1]

Matt Challacombe[2] and Eric Schwegler
Department of Chemistry and the Minnesota Supercomputer Institute,
1200 Washington Ave. South
Minneapolis, Minnesota 55415
(December 26, 1996)

Computation of the Fock matrix is currently the limiting factor in the application of Hartree-Fock and hybrid Hartree-Fock/density functional theories to larger systems. Computation of the Fock matrix is dominated by calculation of the Coulomb and exchange matrices. With conventional Gaussian-based methods, computation of the Fock matrix typically scales as $\sim N^{2.7}$, where N is the number of basis functions.

A hierarchical multipole method is developed for fast computation of the Coulomb matrix. This method, together with a recently described approach to computing the Hartree-Fock exchange matrix of insulators (J. Chem. Phys. **105**, 2726), leads to a linear scaling algorithm for calculation of the Fock matrix.

Linear scaling computation the Fock matrix is demonstrated for a sequence of water clusters at the restricted Hartree-Fock/3-21G level of theory, and corresponding accuracies in converged total energies are shown to be comparable with those obtained from standard quantum chemistry programs. Restricted Hartree-Fock/3-21G calculations on several proteins of current interest are documented, including endothelin, charybdotoxin, and the tetramerization monomer of P53. The P53 calculation, involving 698 atoms and 3836 basis functions, may be the largest Hartree-Fock calculation to date. The electrostatic potentials of charybdotoxin and the tetramerization monomer of P53 are visualized and the results are related to molecular function.

I. INTRODUCTION

The Hartree-Fock approximation continues to be a central model in quantum chemistry. In molecular orbital theory, the Hartree-Fock approximation provides a reference for methods that include the effects of dynamical electron correlation [1]. Within the density functional theory, Hartree-Fock exchange is a necessary ingredient in the most accurate functionals [2–7]. Unmodified, Hartree-Fock theory models the essential physics in systems where the effects of dynamical correlation are small, and generally yields qualitatively correct results.

The rapid growth of fast technologies is opening the door to a variety of *ab initio* electronic structure methods with dramatically extended aplicabilities. Already, the fast Fourier transform has had a tremendous impact on the plane-wave methods commonly used in solid-state local density functional theory. More recently, fast methods based on grids [8–10] and wavelets [11,12] have been introduced which are appropriate for electronic structure calculations of both finite and periodic systems. With the growing interest in fast bases, it should be pointed out that fast Gaussian methods also exist [13–18], are continuing to advance (here for example), and that Gaussian basis functions feature a number of properties lacking in many of the fast bases. In particular, to the list of Gaussian attributes recently compiled by Wilson [19] we add that (a) Gaussian basis functions are smooth and (b) Gaussian basis functions can efficiently represent both core and valence regions of the electron density. Feature (a) is important for the accurate calculation of non-Hamiltonian properties using variationally determined wavefunctions or electron densities [20]. Property (b) is vital for avoiding pseudopotentials, which do not admit the calculation of the Fermi contact terms required for the *ab initio* determination of NMR and ESR spectroscopic observables.

In quantum chemistry, the basis functions of choice are the Cartesian Gaussian-type functions (CGTFs),

$$\phi_a(\mathbf{r}) = (x-A_x)^{l_a}(y-A_y)^{m_a}(z-A_z)^{n_a}\, e^{-\zeta_a(\mathbf{r}-\mathbf{A})^2}, \tag{1}$$

which facilitate the rapid and analytic evaluation of the two-electron repulsion integrals (ERIs) [21–32]

$$(\phi_a\phi_b|\phi_c\phi_d) = \int d\mathbf{r} \int d\mathbf{r}'\; \phi_a(\mathbf{r})\,\phi_b(\mathbf{r})\,|\mathbf{r}-\mathbf{r}'|^{-1}\,\phi_c(\mathbf{r}')\,\phi_d(\mathbf{r}')\ . \tag{2}$$

Solution of the Hartree-Fock (HF) equations by the molecular orbital method [33–35], or the density matrix approach [36–42], requires computation of the Fock matrix

$$\mathbf{F} = \mathbf{h} + \mathbf{J} - \tfrac{1}{2}\mathbf{K} \tag{3}$$

where

$$h_{ab} = (\phi_a|h|\phi_b)\,, \tag{4}$$

$$J_{ab} = \sum_{cd} D_{cd}\,(\phi_a\phi_b|\phi_c\phi_d)\,, \tag{5}$$

[1] To appear in the Journal of Chemical Physics
[2] mf10111@msi.umn.edu

and

$$K_{ab} = \sum_{cd} D_{cd} \left(\phi_a \phi_c | \phi_b \phi_d\right) \quad (6)$$

are elements of the core-Hamiltonian, Coulomb, and exchange matrices for a closed-shell system. In Eqs. (5) and (6), $\mathbf{D}$ is the density matrix

$$D_{cd} = 2 \sum_{\mu} C_{d\mu} C^*_{c\mu} \quad (7)$$

obtained by the contraction of molecular orbital coefficients $C_{d\mu}$ [33–35], or found directly [36–42].

Computation of the Fock matrix formally requires CPU time proportional to N^4, where N is the number of basis functions. This N^4 dependence is associated with the evaluation and processing of the four-index ERIs. However, the product of two Gaussians is

$$e^{-\zeta_a(\mathbf{r}-\mathbf{A})^2} e^{-\zeta_b(\mathbf{r}-\mathbf{B})^2} = e^{-\xi(\mathbf{A}-\mathbf{B})^2} e^{-\zeta_p(\mathbf{r}-\mathbf{P})^2} , \quad (8)$$

where $\xi_p = \zeta_a \zeta_b / \zeta_p$, $\mathbf{P} = (\zeta_a \mathbf{A} + \zeta_b \mathbf{B})/\zeta_p$, and $\zeta_p = \zeta_a + \zeta_b$. As the system size increases, the radial overlap $e^{-\xi_p(\mathbf{A}-\mathbf{B})^2}$ falls as a Gaussian with separation of the two centers A and B, and the number of distributions $\rho_{ab}(\mathbf{r}) = \phi_a(\mathbf{r})\phi_b(\mathbf{r})$ significant to the density, $\rho(\mathbf{r}) = \sum_{ab} D_{ab}\rho_{ab}(\mathbf{r})$, approaches N from above as N becomes large. Likewise, the number of significant ERIs larger than some small threshold approaches N^2 [43,44].

With $\mathcal{O}(N)$ significant distributions $\rho_{ab}(\mathbf{r})$, the Coulomb matrix may in principle be computed in $\mathcal{O}(N)$ or $\mathcal{O}(N \log N)$ CPU time with hierarchical multipole methods related to the N-body methods of classical particle simulation [45–60]. Previously, we described a quantum chemical tree code [16,17] for the fast computation of $\mathbf{J}$, but found that it was only able to achieve a scaling of $N^{1.6}$ for α-helices and $N^{1.9}$ for water clusters.

In this paper, an entirely new development of the quantum chemical tree code (QCTC) is presented. Together with a recently described method for computing $\mathbf{K}$ in $\mathcal{O}(N)$ CPU time [61] (valid only for insulating systems), this development enables linear scaling computation of the Fock matrix, and a dramatically extended applicability of Hartree-Fock theory. This paper is organized as follows: In Section II, a new direct (N^2) method for computing $\mathbf{J}$ is described. This direct method forms a foundation for the fast evaluation of $\mathbf{J}$. In Section III, QCTC is introduced and its new formulation is developed. Then, in Section IV A, linear scaling computation of the Fock matrix is demonstrated on a sequence of water clusters at the RHF 3-21G level of theory, and in Section IV B RHF/3-21G calculations on several proteins of current interest are documented, including endothelin, charybdotoxin, and the tetramerization monomer of P53. In Section V A, the performance of these methods is discussed. Then, in Section V B the electrostatic potentials of charybdotoxin and the tetramerization monomer of P53 are visualized and the results related to molecular function. Finally, in Section VI, concluding remarks are made.

II. A DIRECT METHOD FOR COMPUTING J

A new algorithm for the direct computation of $\mathbf{J}$ is introduced that is based on thresholding primitive ERIs. The method begins with a previously described decomposition of the density into families of Hermite Gaussian-type functions (HGTFs) [17,16,62]. In this representation, insignificant contributions to $\mathbf{J}$ and ρ are eliminated, yielding an $\mathcal{O}(N)$ complexity for both. Then, by structuring the density, a competitive algorithm for the thresholding and evaluation of primitive ERI contributions to $\mathbf{J}$ is obtained.

A. Hermite Gaussian-basis functions and the classical McMurchie-Davidson method for the evaluation of J

In the McMurchie-Davidson (MD) algorithm [63–65] for the evaluation of CGTF-based ERIs, products of CGTFs describing elementary charge distributions are expanded exactly in HGTFs as

$$\rho_{ab}(\mathbf{r}) \equiv \phi_a(\mathbf{r})\, \phi_b(\mathbf{r}) = \sum_{L}^{l_a+l_b} \sum_{M}^{m_a+m_b} \sum_{N}^{n_a+n_b} e^{ab}_{LMN} \Lambda^p_{LMN}(\mathbf{r}) \ , \quad (9)$$

where e^{ab}_{LMN} is an expansion coefficient obtained with the MD two-term recurrence relation and multiplied with the corresponding radial overlap. The HGTFs are

$$\Lambda^p_{LMN}(\mathbf{r}) = \frac{\partial^L}{\partial P_x^L} \frac{\partial^M}{\partial P_y^M} \frac{\partial^N}{\partial P_z^N} \exp\left[-\zeta_p (\mathbf{r}-\mathbf{P})^2\right] . \quad (10)$$

The ERIs

$$\left(\Lambda^p_{LMN} \middle| \Lambda^q_{L'M'N'}\right) = (-1)^{L'+M'+N'} u_{pq}\, R^{pq}_{L+L',M+M',N+N'} \quad (11)$$

may be obtained efficiently with the MD three-term recursion relations for computing the auxiliary functions

$$R^{pq}_{lmn0} = \frac{\partial^l}{\partial PQ_x^l} \frac{\partial^m}{\partial PQ_y^m} \frac{\partial^n}{\partial PQ_z^n} F_0(\tau_{pq}) \ . \quad (12)$$

In Eq. (12),

$$u_{pq} = \frac{2\pi^{5/2}}{\zeta_p \zeta_q \sqrt{\zeta_p + \zeta_q}} \ , \quad (13)$$

and PQ_x, PQ_y and PQ_z are the Cartesian components of $\mathbf{P}-\mathbf{Q}$. The auxiliary functions are related to the 0^{th} order reduced incomplete gamma function [66,67], which is

$$F_0(\tau_{pq}) = \sqrt{\frac{\pi}{4\tau_{pq}}} \mathrm{erf}\left(\sqrt{\tau_{pq}}\right) = \int_0^1 e^{-\tau_{pq} u^2} du \ , \quad (14)$$

with argument

$$\tau_{pq} = \omega_{pq} |\mathbf{P} - \mathbf{Q}|^2 \; , \tag{15}$$

where

$$\omega_{pq} = \frac{\zeta_p \zeta_q}{\zeta_p + \zeta_q} \; . \tag{16}$$

Differentiation of F_0 leads to higher order reduced incomplete gamma functions $F_{j+1} = -F_j'$, which we refer to hereafter simply as "gamma" functions.

The conventional MD method for evaluating the Coulomb matrix is

$$J_{ab} = \sum_{cd} D_{cd} \, u_{pq} \sum_{L}^{l_a+l_b} \sum_{M}^{m_a+m_b} \sum_{N}^{n_a+n_b} e^{ab}_{LMN} \tag{17}$$
$$\times \sum_{L'}^{l_c+l_d} \sum_{M'}^{m_c+m_d} \sum_{N'}^{n_c+n_d} (-1)^{L'+M'+N'} e^{cd}_{L'M'N'} R^{pq}_{L+L',M+M',N+N'} \; ,$$

which involves four-fold loops over basis functions a, b, c and d, inside of which an expensive HGTF to CGTF transformation over LMN and $L'M'N'$ takes place.

B. ρ in a Hermite Gaussian basis

The expansion of CGTF distributions leads to primitive families, labeled by q, of HGTFs that are all related to a "mother" (s-type) Gaussian $e^{-\zeta_q (\mathbf{r}-\mathbf{Q})^2}$ by differentiation. It is possible to obtain a compact representation of the density in terms of HGTF families by accumulation of contributions to each family q [17,16,62] via

$$E^q_{L'M'N'} = (-1)^{L'+M'+N'} \sum_{q \in cd} D_{cd} \, e^{cd}_{L'M'N'} \; . \tag{18}$$

In this way, the redundancies associated with equivalent combinations of CGTF angular symmetries may be eliminated. Also, the evaluation of ERIs in a family HGTF basis is greatly simplified as the three-term MD recursion relations generate all auxiliary functions R_{LMN} for which $L + M + N \leq \mathcal{L}$ in CPU time proportional to $\binom{\mathcal{L}+4}{4}$.

Importantly, this representation facilitates the characterization and ordering of primitive distributions on the basis of the length scales $1/\sqrt{\zeta_q}$ and radial overlaps $e^{-\xi_q (\mathbf{C}-\mathbf{D})^2}$ associated with the mother Gaussian. The advantages of this decomposition are elaborated upon in the following.

C. Thresholding contributions to ρ and $\mathbf{J}$

Our objective is to compute each element J_{ab} to within an absolute error $\sim \mathcal{O}(\texttt{thresh})$. An essential difference between our approach and conventional direct SCF methods is the use of statistical error estimates that are size dependent. Previously [16,17], we used heuristic and statistical arguments to arrive at the criteria

$$\frac{2\pi}{\zeta_p} e^{-\xi_p (\mathbf{A}-\mathbf{B})^2} \leq \frac{\texttt{thresh}}{N_{\rm el}} \; , \tag{19}$$

and

$$\zeta_q^{-5/4} e^{-\xi_q (\mathbf{C}-\mathbf{D})^2} \max_{cd \in q}(|D_{cd}|) \leq \left(\frac{\sqrt{2}\, \zeta_{\rm min}^{5/4}}{\pi^{5/2} N} \right) \texttt{thresh} \tag{20}$$

for the independent thresholding of primitive family HGTF contributions to $\mathbf{J}$ and ρ, respectively. The size dependence of these criteria is important as it prevents the build up of small but collectively significant errors that might otherwise accrue with fixed tolerances.

D. Thresholding primitive ERIs

For each distribution ρ_{cd}, there is an associated length scale $1/\sqrt{\zeta_q}$ and corresponding radial overlap $e^{-\xi_q (\mathbf{C}-\mathbf{D})^2}$. Likewise, for each distribution ρ_{ab} contributing to $\mathbf{J}$, there is an associated length scale $1/\sqrt{\zeta_p}$ and corresponding radial overlap $e^{-\xi_p (\mathbf{A}-\mathbf{B})^2}$. Clearly, when $e^{-\xi_p (\mathbf{A}-\mathbf{B})^2}$ is small, an acceptable absolute error in J_{ab} may be maintained while omitting a large but insignificant fraction of the density. To exploit this fact, we structure the density by first grouping the family HGTFs q by primitive exponent, and then sorting each group, sub-indexed by Q, on the basis of their significance. Specifically, for each primitive exponent ζ_q the elements $E^Q_{L'M'N'}$ are grouped and then sorted in decreasing order on the basis of the associated value

$$\texttt{rad}(Q) = \max_{cd \in Q}(|D_{cd}|) e^{-\xi_q (\mathbf{C}-\mathbf{D})^2} \; . \tag{21}$$

With this ordering, the HGTF-based Coulomb matrix is evaluated as

$$\mathcal{J}_{p,LMN} = \sum_q u_{pq} {\sum_Q}' \sum_{L'M'N'} E^Q_{L'M'N'} R^{pQ}_{L+L',M+M',N+N'} \; , \tag{22}$$

with abortion of the primed sum (loop) when

$$\texttt{rad}(Q) \leq \frac{\texttt{thresh}}{u_{pq} N_{\rm ket}} e^{\xi_p (\mathbf{A}-\mathbf{B})^2} \; . \tag{23}$$

CGTF-based matrix elements corresponding to the bra family functions p are then evaluated by contraction with the basis function MD expansion coefficients as

$$J_{ab} = \mathbf{e}^{ab} \cdot \boldsymbol{\mathcal{J}}_p \, , \quad J_{a'b'} = \mathbf{e}^{a'b'} \cdot \boldsymbol{\mathcal{J}}_p \, , \quad \ldots \quad (ab,\, a'b',\, \ldots \, \in p) \; . \tag{24}$$

A key feature of this approach is that the computation of effectively zero primitive ERIs is completely avoided. Thresholding in this way is efficient since only distributions that are significant at some stage are considered, and because the number of effectively zero primitive ERIs grows rapidly with system size.

III. THE QUANTUM CHEMICAL TREE-CODE

A common feature of hierarchical multipole methods is that the system is recursively subdivided into a hierarchy of cells; cubes, boxes, polyhedra, etc. This hierarchy establishes a tree-like data structure with a complete representation of the system at each tier, or level of spatial refinement. Reductions in computational complexity are achieved by approximating particle distributions within a cell as a series (multipole) expansion that converges rapidly in the far-field. Computations involving the hierarchy of cells are effective because the series expansions become increasingly more accurate with separation; larger cells may be used as the interaction distance increases.

Broadly, hierarchical multipole methods may be classified as tree-codes [45–51] or fast multipole methods [52–58]. Tree-codes involve particle-cell interactions, while fast multipole methods are characterized by cell-cell multipole-to-local transformations of the potential. Tree-codes apply the multipole approximation more aggressively than fast multipole methods because tighter error estimates (accessibility criteria) may be obtained for the particle-cell interaction than for cell-cell interactions. On the other hand, fast multipole methods can achieve an $\mathcal{O}(N)$ complexity, while tree-codes are at best $\mathcal{O}(N \log N)$, although there is some confusion about this [68–71].

Hierarchical multipole methods may be used to compute far-field components of the Coulomb matrix [13–18]. However, unlike the classical N-body problem, quantum chemical methods involve continuous distributions. Integrals that involve non-penetrating charge distributions may be efficiently evaluated with multipole methods, while integrals with overlapping charge distributions are near-field and must be computed by direct methods. The problem is complicated by the fact that the near/far distinction depends on the extent (rate of decay) of the charge distributions and their separation, as well as the error requested and the order of multipole expansion employed.

In the quantum Coulomb problem, penetration effects and the range of distribution magnitudes (~ 10 orders) may increase the advantage of employing distribution-cell error estimates. In an initial implementation of a quantum chemical tree code (QCTC_A) [16,17], we attempted to exploit this advantage, and were able to achieve sub-N^2 scalings with relatively crude methods. In particular, QCTC_A employed low order Cartesian multipole expansions, used static penetration accessibility criteria, and did not threshold primitive integrals. Perhaps most limiting, QCTC_A accessed multipole information with tier indexed arrays rather than linked lists, restricting the maximum number of tiers to 5 due to memory constraints. Here, we report a completely new implementation of the quantum chemical tree code (QCTC) that is based on a variety of new technologies developed in the following Sections.

A. The two-center multipole expansion in real arithmetic

The real arithmetic, two-center (bipolar) multipole expansion introduced in the seminal work of Carlson and Rushbrooke [72] can be simply re-written in terms of the scaled associated Legendre polynomials. These polynomials, which were introduced by White and Head-Gordon [56], may be computed very efficiently using recursion relations [56,73,74].

Carlson and Rushbrooke's Eq. IIIa [72] can be expressed as

$$\begin{aligned}\frac{1}{|\mathbf{r}-\mathbf{r}'|} = \sum_{l_p}^{\infty}(-1)^{l_p}\sum_{l_q}^{\infty}\sum_{m_p=-l_p}^{l_p}\sum_{m_q=-l_q}^{l_q}\Big(& \\ C_{l_p}^{m_p}[\mathbf{r}-\mathbf{P}]\mathcal{C}_{l_p+l_q}^{m_p+m_q}[\mathbf{P}-\mathbf{Q}]C_{l_q}^{m_q}[\mathbf{r}'-\mathbf{Q}] & \\ + S_{l_p}^{m_p}[\mathbf{r}-\mathbf{P}]\mathcal{S}_{l_p+l_q}^{m_p+m_q}[\mathbf{P}-\mathbf{Q}]C_{l_q}^{m_q}[\mathbf{r}'-\mathbf{Q}] & \\ - S_{l_p}^{m_p}[\mathbf{r}-\mathbf{P}]\mathcal{C}_{l_p+l_q}^{m_p+m_q}[\mathbf{P}-\mathbf{Q}]S_{l_q}^{m_q}[\mathbf{r}'-\mathbf{Q}] & \\ + C_{l_p}^{m_p}[\mathbf{r}-\mathbf{P}]\mathcal{S}_{l_p+l_q}^{m_p+m_q}[\mathbf{P}-\mathbf{Q}]S_{l_q}^{m_q}[\mathbf{r}'-\mathbf{Q}]\Big) & ,\end{aligned} \tag{25}$$

where

$$C_l^m[\mathbf{r}] = |\mathbf{r}|^l \tilde{P}_l^m(\cos\theta_{\mathbf{r}})\cos m\phi_{\mathbf{r}} \tag{26}$$

and

$$S_l^m[\mathbf{r}] = |\mathbf{r}|^l \tilde{P}_l^m(\cos\theta_{\mathbf{r}})\sin m\phi_{\mathbf{r}} \tag{27}$$

are multipole tensors, and

$$\mathcal{C}_l^m[\mathbf{R}] = |\mathbf{R}|^{-(l+1)}\ \ \tilde{\tilde{P}}_l^m(\cos\theta_{\mathbf{R}})\cos m\phi_{\mathbf{R}} \tag{28}$$

and

$$\mathcal{S}_l^m[\mathbf{R}] = |\mathbf{R}|^{-(l+1)}\ \ \tilde{\tilde{P}}_l^m(\cos\theta_{\mathbf{R}})\sin m\phi_{\mathbf{R}} \tag{29}$$

are interaction tensors. The scaled associated Legendre polynomials introduced by White and Head-Gordon [56] are

$$\tilde{P}_l^m(x) = \frac{1}{(l+m)!}P_l^m(x)\ . \tag{30}$$

and

$$\tilde{\tilde{P}}_l^m(x) = (l-m)!P_l^m(x)\ . \tag{31}$$

A key property of these polynomials is that they may be computed very efficiently with recursion relations [56,73,74].

In addition to avoiding complex arithmetic, Eq. (25) can be greatly simplified with the identities

$$C_l^{-m} = (-1)^m C_l^m\ , \tag{32}$$

$$\mathcal{C}_l^{-m} = (-1)^m \mathcal{C}_l^m \ , \tag{33}$$

$$S_l^{-m} = (-1)^{(m+1)} S_l^m \ , \tag{34}$$

and

$$\mathcal{S}_l^{-m} = (-1)^{(m+1)} \mathcal{S}_l^m \ . \tag{35}$$

In the case of a one-center multipole expansion, it is straightforward to employ these identities and to extract a factor of 2 from the consolidation of redundant terms [74]. In the two-center case, a 2 may may also be factored out, but the resulting summations cannot be simply expressed without conditionals or bounds on indices. We have therefore constructed multipole contraction libraries that perform the secondary sum in Eqs. (57) and (58), as well as multipole translation libraries that carry out the transformations (41) and (42), in explicit fortran (*ie.* without control structures). These libraries were built using the symbolic manipulation features of *Mathematica* [75] and the `FortranAssign` utility in `Format.m` [76].

In addition to factoring out a 2 and symbolically eliminating elements that are zero by symmetry, the resulting code does not encumber loop overheads and fully exploits floating-point multiply-add parallelism (a factor of two on the IBM RS6000s).

1. Multipole moments

The two-center expansion given by Eq. (25) is valid for calculation of the matrix elements $(\Lambda^p_{LMN}|\Lambda^q_{L'M'N'})$ in the (multipole) limit of large exponents, and may be truncated at some finite l_p and l_q with increasing accuracy as $|\mathbf{P}-\mathbf{Q}|$ becomes large. In the limit of large exponents, HGTFs are related to derivatives of the Dirac delta function by

$$\delta^p_{LMN}(\mathbf{r}) = \lim_{\zeta_p \to \infty} \left(\frac{\zeta_p}{\pi}\right)^{3/2} \Lambda^p_{LMN}(\mathbf{r}) \ , \tag{36}$$

which obey the identity [77]

$$\int d\mathbf{r}\, f(\mathbf{r})\ \delta^p_{LMN}(\mathbf{r}) = (-1)^{L+M+N} \frac{\partial^L}{\partial P_x^L} \frac{\partial^M}{\partial P_y^M} \frac{\partial^N}{\partial P_z^N} f(\mathbf{P}) \ . \tag{37}$$

This identity leads naturally to multipole moments of basis functions,

$$c^{ab}_{m,l}[\mathbf{P}] = \lim_{\zeta_p \to \infty} \sum_{LMN} e^{ab}_{LMN} \int d\mathbf{r}\ C^m_l[\mathbf{r}-\mathbf{P}]\ \Lambda^p_{LMN}(\mathbf{r}) \tag{38}$$
$$= \left(\frac{\pi}{\zeta_p}\right)^{3/2} \sum_{\substack{LMN \\ L+M+N \le l}} e^{ab}_{LMN} \left.\frac{\partial^{L+M+N} C^l_m[\mathbf{X}]}{\partial X_x^L \partial X_y^M \partial X_z^N}\right|_{\mathbf{X}\to\mathbf{0}} ,$$

and to multipole moments of the density,

$$C^m_l[\mathbf{Q}] = \sum_{Q\in q} \left(\frac{\pi}{\zeta_q}\right)^{3/2} \times \sum_{\substack{LMN \\ L+M+N \le l}} E^q_{LMN} \left.\frac{\partial^{L+M+N} C^l_m[\mathbf{X}]}{\partial X_x^L \partial X_y^M \partial X_z^N}\right|_{\mathbf{X}\to\mathbf{0}} . \tag{39}$$

The multipole moments $s^{ab}_{m,l}[\mathbf{P}]$ and $S^m_l[\mathbf{Q}]$ may be obtained in an analogous way.

2. Trees and translations

The use of truncated multipole expansions alone will not reduce the computational complexity of Coulomb matrix evaluation to below N^2, as the number of multipoles, $C^m_l[\mathbf{Q}]$ and $S^m_l[\mathbf{Q}]$, is proportional to the number of distributions that constitute the density. Our quantum chemical tree-code (QCTC) achieves a reduced computational complexity with a hierarchical multipole decomposition of the density. This algorithm begins by dividing a box containing all distributions (the root) into 8 equal boxes or cells (children). Recursive sub-division of this type leads to an oct-tree data-structure. Each tier of the tree represents the entire density, but involves cells of increasing spatial resolution formed by sub-division of their parents. Sub-division is stopped when a direct cell containing less than a certain number of distributions is reached. Cells that do not contain any distributions are not added to the tree.

Starting at the finest level of spatial discretization, the distribution moments contained in each cell O are translated to a common "center of mass",

$$\mathbf{O} = \frac{1}{N_O} \sum_{Q\in O} |C^0_0[\mathbf{Q}]|\ \mathbf{Q} \ , \tag{40}$$

which minimizes the associated errors. In Eq. (40), N_O is the number of distributions in cell O, and the sum is taken over these. Each translation is carried out exactly with the identities

$$C^m_l[\mathbf{Q}] = \sum_{k=0}^{l} \sum_{\substack{-k\ge n\ge k \\ |m-n|\le l-k}} \Big(C^n_k[\mathbf{Q}] C^{m-n}_{l-k}[\mathbf{O}-\mathbf{Q}] - S^n_k[\mathbf{Q}] S^{m-n}_{l-k}[\mathbf{O}-\mathbf{Q}] \Big) \tag{41}$$

and

$$S^m_l[\mathbf{Q}] = \sum_{k=0}^{l} \sum_{\substack{-k\ge n\ge k \\ |m-n|\le l-k}} \Big(C^n_k[\mathbf{Q}] S^{m-n}_{l-k}[\mathbf{O}-\mathbf{Q}] + S^n_k[\mathbf{Q}] C^{m-n}_{l-k}[\mathbf{O}-\mathbf{Q}] \Big) \ , \tag{42}$$

and each translated moment is summed to obtain the total cell moments. Analogous translations are performed recursively up the tree (to the root), in which parent multipole moments are formed from their children, establishing a hierarchical multipole representation of the density.

B. Admissibility criteria

The efficiency of hierarchical multipole methods is highly dependent on the aggressive partition of near- and far-field interactions. Coulomb matrix elements involving distributions for which τ_{pq} is sufficiently large avoid errors due to penetration, and are said to be in the far-field regime. When penetration effects are small, Coulomb interactions between a distribution p and a multipole moment translated from $\mathbf{Q}$ to $\mathbf{O}$ incur negligible errors (and are likewise said to be far-field) when for a given expansion length ℓ

$$\left(\frac{|\mathbf{Q}-\mathbf{O}|}{|\mathbf{P}-\mathbf{O}|}\right)^{\ell} \tag{43}$$

is sufficiently small.

Obviously, the near-field/far-field distinction depends on both the absolute error requested and the strength of a specific interaction (ℓ dependent in the case of multipoles). This observation is significant as distributions, both bra and ket, involve radial overlaps that span many ($\sim$ 10) orders of magnitude. QCTC exploits this fact by partitioning near- and far-field interactions on the basis of multipole and penetration admissibility criteria that depend dynamically on the magnitude of each interaction.

1. Multipole admissibility criteria

The truncation error for a family of matrix elements p interacting with the moments of cell O is

$$\epsilon_p(\ell) \approx c^p \sum_{l=\ell+1}^{\infty} \sum_{m=-l}^{l} \Big(\mathcal{C}_l^m[\mathbf{P}-\mathbf{O}] C_l^m[\mathbf{O}] + \mathcal{S}_l^m[\mathbf{P}-\mathbf{O}] S_l^m[\mathbf{O}] \Big) , \tag{44}$$

where ℓ is the expansion order, and

$$c^p = \max_{ab\in p} c_{00}^{ab}[\mathbf{P}] \, . \tag{45}$$

If the family p does not include s-type functions, the max is still taken with the assumption of spherical symmetry, but includes the appropriate contraction coefficients.

With Unsöld's theorem [72,78],

$$\sum_{m=-l}^{l} \frac{(l-m)!}{(l+m)!} \left[P_l^m(x)\right]^2 = 1, \tag{46}$$

and the Schwartz inequality [79], it is straightforward to show

$$\sum_{m=-l}^{l} \left[|\mathbf{R}|^{-(l+1)} \;\; \tilde{\tilde{P}}_l^m(\cos\theta_{\mathbf{R}})\right] \left[|\mathbf{r}|^l \tilde{P}_l^m(\cos\theta_{\mathbf{r}})\right] \tag{47}$$
$$\leq \frac{1}{|\mathbf{R}|^{l+1}} \sqrt{\sum_{m=-l}^{l} (l-m)!(l+m)! \; |\mathbf{r}|^{2l} \left[\tilde{P}_l^m(\cos\theta_{\mathbf{r}})\right]^2} \, .$$

Based on this inequality, we introduce the multipole strength

$$\bigcirc_\ell[\mathbf{O}] = \sum_{l=\ell}^{\ell+1} \sqrt{\sum_{m=-l}^{l} (l-m)!(l+m)! \left[(C_m^l[\mathbf{O}])^2 + (S_m^l[\mathbf{O}])^2\right]} \tag{48}$$

with which the truncation error is safely estimated as

$$\epsilon_p(\ell) \approx \mathtt{ff}\frac{c^p \bigcirc_\ell[\mathbf{O}]}{|\mathbf{P}-\mathbf{O}|^{\ell+1}} \, . \tag{49}$$

Equating Eq. (49) with `thresh`, the multipole admissibility criterion (MAC)

$$|\mathbf{P}-\mathbf{O}|^2 > [c^p]^{2/(\ell+1)} \left(\frac{\mathtt{ff}\bigcirc_\ell[\mathbf{O}]}{\mathtt{thresh}}\right)^{2/(\ell+1)} , \tag{50}$$

obtains after minor rearrangement. This MAC is computationally efficient, incorporates the magnitude of specific interactions, and limits the absolute error in matrix elements of $\mathbf{J}$ rather than the relative error in the electrostatic potential. Because Eq. (50) is approximate, and because errors will accumulate due to the repeated use of the multipole approximation ($\sim \log N$ times), the fudge factor `ff` has been introduced to allow fine tuning.

2. Penetration admissibility criteria

The neglect of penetration in the interaction between HGTF family members p and q leads to errors in matrix elements J_p that are dominated by

$$\epsilon_p(\tau_{pq}) \approx e_{000}^p E_{000}^q \left(u_{pq} F_0(\tau_{pq}) - \frac{\pi^3}{(\zeta_p\zeta_q)^{3/2}} \frac{1}{|\mathbf{P}-\mathbf{Q}|}\right) . \tag{51}$$

Equating this error with `thresh`, introducing the tuning parameter `ff`, and making the assumption $|\mathbf{P}-\mathbf{O}| \sim |\mathbf{P}-\mathbf{Q}|$, leads to the penetration admissibility criterion (PAC)

$$|\mathbf{P}-\mathbf{O}|^2 > (c^p)^2 \left(\frac{C_0^0[\mathbf{O}]\,\mathtt{ff}}{\mathtt{thresh}}\right)^2 \left[\mathrm{erfc}\left(\sqrt{\tau_{po}}\right)\right]^2 . \tag{52}$$

In Eq. (52),

$$\tau_{po} = \frac{\zeta_p \zeta_o}{\zeta_p + \zeta_o} |\mathbf{P} - \mathbf{O}|^2 \tag{53}$$

and

$$\zeta_o = \min_{q \in O} \zeta_q \tag{54}$$

is the minimum exponent in cell O. In practice, the approximation [80]

$$[\mathrm{erfc}(\sqrt{x}\,)]^2 = \frac{1}{2x} e^{-2x} \quad (x \to \infty) \tag{55}$$

is employed.

C. Matrix elements

Elements of the Coulomb matrix are evaluated as

$$J_{ab} = \mathbf{e}^{ab} \cdot \boldsymbol{\mathcal{J}}^{\mathrm{near}}_p + \mathbf{s}^{ab} \cdot \boldsymbol{\mathcal{S}}^{\mathrm{far}}[\mathbf{P}] + \mathbf{c}^{ab} \cdot \boldsymbol{\mathcal{C}}^{\mathrm{far}}[\mathbf{P}] \;, \tag{56}$$

where

$$\mathcal{C}^{\mathrm{far}}_{l,m}[\mathbf{P}] = \sum_O \sum_{l_o m_o}{}' \Big(\mathcal{C}^{m+m_o}_{l+l_o}[\mathbf{P}-\mathbf{O}] C^{m_o}_{l_o}[\mathbf{O}] + \mathcal{S}^{m+m_o}_{l+l_o}[\mathbf{P}-\mathbf{O}] S^{m_o}_{l_o}[\mathbf{O}] \Big) \tag{57}$$

and

$$\mathcal{S}^{\mathrm{far}}_{l,m}[\mathbf{P}] = \sum_O \sum_{l_o m_o}{}' \Big(\mathcal{S}^{m+m_o}_{l+l_o}[\mathbf{P}-\mathbf{O}] C^{m_o}_{l_o}[\mathbf{O}] - \mathcal{C}^{m+m_o}_{l+l_o}[\mathbf{P}-\mathbf{O}] S^{m_o}_{l_o}[\mathbf{O}] \Big) \tag{58}$$

are components of the far-field potential. In Eqs. (57) and (58), the primes signify that these summations include only acceptable far-field contributions, as determined by MAC and PAC, and that the largest cells and the most truncated expansions (smallest l_o) are employed.

When a direct cell is encountered while walking the tree (as explained below in Section III C 2), its distributions are added to a list of near-field integrals that must be computed. As was explained in Section II D, many of these integrals are insignificant when the bra moment is small. To take advantage of this in the evaluation of near-field integrals, the distributions in each direct cell are ordered on the basis of

$$\mathtt{rad}(q) = \max_{cd \in q}(|D_{cd}|) \frac{e^{-\xi_q (\mathbf{C}-\mathbf{D})^2}}{\zeta_q^{5/4}} \;. \tag{59}$$

Looping over the family distributions q, building the list of near-field integrals is stopped when

$$\mathtt{rad}(q) \le \frac{\pi^{5/2} e^{-\xi_p (\mathbf{A}-\mathbf{B})}}{\zeta_p^{5/4}\; N_{\mathrm{ket}}} \;, \tag{60}$$

and the next cell in the list is considered.

1. Multiresolution

When a HGTF family p is far from cell O, use of the multipole approximation is controlled by the MAC. When p and O are close, both the MAC and PAC may control admissibility. As we previously argued [16,17], an aggressive application of the multipole approximation may be achieved by partitioning distributions both spatially and by extent. This multiresolution of the density is achieved simply by treating the distribution exponents ζ_q as a fourth dimension.

Of course, it does not make sense to always partition the density in this way, as it leads to a larger tree and since the admissibility of large cells is controlled by the MAC. Therefore, multiresolution is begun after a certain tier has been reached. If the root is tier 0, we have found that excellent results are obtained if multiresolution begins at tier 8. Typically, with multiresolution beginning at tier 8 and direct cells containing less than 50 distributions, QCTC will employ trees with 11-13 tiers.

Note that one-center distributions (those centered on an atom) cannot be further decomposed by spatial subdivision. In addition to overcoming penetration effects, multiresolution allows the decomposition of atom-centered distributions, which may be profitable with extended basis sets.

2. Walking the tree

The hierarchy of multipole moments representing the density form an oct-tree, with some hexadeca branches (due to multiresolution). This tree is efficiently traversed with two lists that (1) point to the next cell that must be checked if either MAC or PAC are not met, and (2) the next cell that must be visited if MAC and PAC are met. With these linked lists, it is a simple matter to traverse the tree, dynamically partitioning near-field and far-field interactions and computing matrix elements.

Controlling the largest proportion of interactions, the MAC is checked first with the maximum expansion length ℓ. If the MAC is not met, that cells children are visited. If the MAC is met, then the simple PAC $\tau_{pq} > 21.0$ is checked, which guarantees penetration errors that are negligible. If this preliminary PAC is not met, the more expensive PAC Eq. (52) is applied unless the cell is direct.

When both MAC and PAC are met, the multipole expansion is truncated at some $l_o \le \ell$ using Eq. (50), and the tensor contractions Eqs. (57) and (58) are carried out with highly optimized routines as described in Section III A. The current version of QCTC employs moments up to $\ell = 15$.

When direct cells are encountered, pointers are added to the near-field list until insignificant distributions are encountered. After each tree-traversal, near-field integrals are evaluated by methods analogous to those described in Section II, except that the distributions are no longer ordered.

IV. RESULTS

Subroutines ONX [61] and QCTC (Section III) have been interfaced with a modified version of the *ab initio* quantum chemistry code HONDO 95.3 [81] (hereafter referred to as MONDO). MONDO achieves absolute accuracies of $\sim 10^{-5}$ au in elements of $\mathbf{F}$ with the appropriate ONX(-5) thresholding parameters [61] and the QCTC thresholding parameters $\mathtt{thresh} = 10^{-6}$ and $\mathtt{ff} = 500$. Cells are considered direct when they contain less than 50 distributions. A multipole expansion length of $\ell = 15$ is employed by QCTC. Elements of the core-Hamiltonian are computed by HONDO, and have a negligible impact on the CPU time required to build $\mathbf{F}$. In all the calculations reported in the following, total energies were converged to within 10^{-5} hartree.

A. Scaling

MONDO SCF calculations were performed at the RHF/3-21G level of theory on a series of three-dimensional water clusters taken from an equilibrated molecular dynamics simulation with a density corresponding to standard temperature and pressure. A breakdown of SCF CPU times as a function of cluster size is shown in Fig. 1. The diagonalization times correspond to subroutine GIVEIS, which is an EISPACK based modification of the QCPE routine GIVENS [82]. All timings were obtained on a single "thin" node of an IBM SP2 (66 MHz CPU, 256 MB RAM).

To examine the adherence of our methods to a linear regime, lines were fit to timings of the $\mathbf{K}$ and $\mathbf{J}$ builds, with resulting correlation coefficients (r) equal to 0.999, and 0.998 respectively. These fits are shown in Fig. 1.

Errors in converged total energies of selected clusters are listed in Table I, and correspond to an error of 0.4 μhartree per atom. For reference, GAUSSIAN 94 [83] was found to achieve a precision of 0.6 μhartree per atom with default parameters. These values were determined by comparison with GAUSSIAN 94 calculations employing the tight option.

The major CPU components of the fast $\mathbf{J}$ build are shown in Fig. 2, along with times for the direct method described in Section II. It is interesting to note that the QCTC CPU time for the $(H_2O)_{150}$ cluster on a single IBM SP2 node corresponds almost exactly to the timing obtained with QCTC_A for the same cluster on the Cray C-90 [17]. QCTC_A is a fully vectorized code achieving performances greater than 370 MFlops on the C-90.

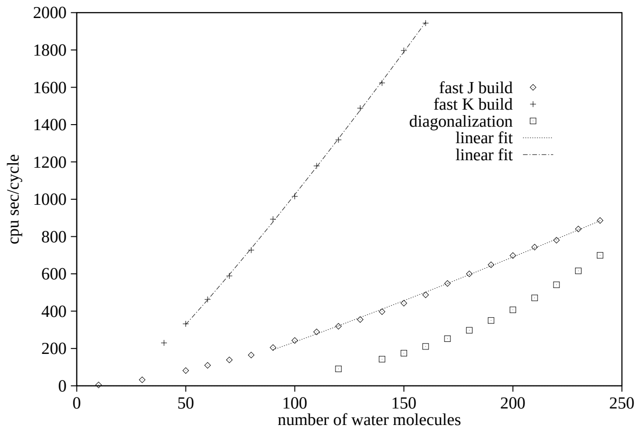

FIG. 1. Timings for construction of the RHF/3-21G Fock matrix two-electron components corresponding to a sequence of water clusters at STP.

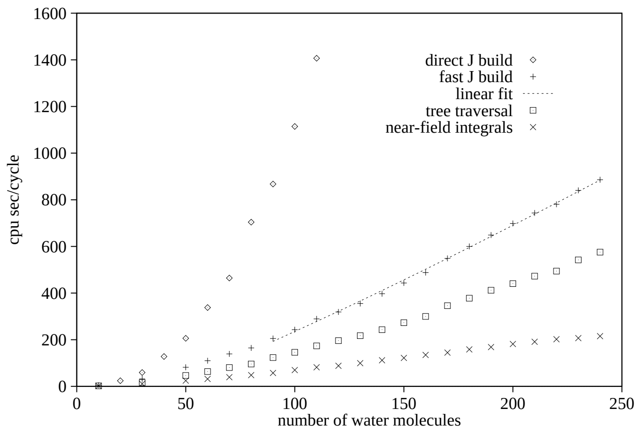

FIG. 2. Timings for construction of the RHF/3-21G Coulomb matrix by direct and fast methods for the sequence of water clusters.

B. Proteins

MONDO SCF calculations at the RHF/3-21G level of theory were performed on several proteins of current interest, including endothelin (EDP) [84], charybdotoxin (CRD) [85], and the tetramerization monomer of P53 [86]. System sizes, total energies, and cpu timings are listed in Table II. These MONDO SCF calculations were carried out on a single "wide" node of an IBM SP2 (77 MHz, 512 MB RAM) and correspond to ONX(-5) and $\mathtt{thresh} = 10^{-6}$.

TABLE I. Errors in converged RHF/3-21G total energies of selected water clusters computed with MONDO.

Water Molecules	Error (mhartree)
50	0.053
70	0.078
90	0.090
110	0.081
130	0.10
150	0.12

TABLE II. Results of single point MONDO RHF/3-21G calculations on selected proteins. Structures were obtained from the protein data bank and capped where appropriate to obtain a closed shell system.

Protein	Energy (hartree)	Atoms	Basis Functions	Time/Cycle (hours) **K**	**J**
EDP	-8154.72370	255	1461	0.85	0.11
	-8154.72323[a]				
CRD	-16737.90643	572	3237	2.70	0.38
P53	-17115.30474	698	3836	2.10	0.38

[a]Energy obtained with GAUSSIAN 94 default settings

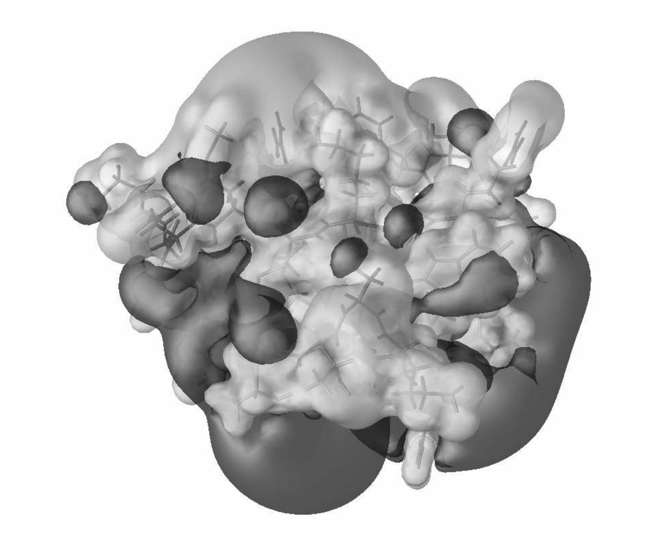

FIG. 3. Iso-surfaces of the CRD electrostatic potential at the RHF/3-21G level of theory. The light surface corresponds to a value of 0.035 au, and the darker one to -0.035 au.

V. DISCUSSION

A. Performance

Although much is often made of a fast algorithm's complexity, *e.g.* $\mathcal{O}(N)$ vs $\mathcal{O}(N \log N)$, it is often possible to observe a strict adherence to these behaviors only with ideal (very large, classical, homogeneous, flat or square) systems. Nevertheless, evidence ($r = 0.998$) for linear scaling of the quantum chemical tree-code has been obtained for the 3-21G basis set and water clusters corresponding to an STP density. A claim for an $\mathcal{O}(N)$ dependence of QCTC is *not* being made, although this point has been argued for classical tree-codes [70,71]. Rather, over the range of sizes considered (limited by memory constraints), and with the noise in the data, differences between N or $N \log N$ behavior may not be manifest, and from a practical point of view, may be irrelevant.

Interestingly, QCTC achieves a time dependence dominated by tree-traversal, rather than the evaluation of two-electron integrals. However, this partitioning is somewhat artificial as timings for the tree-traversal include thresholding at the level of primitives and building the list of near-field integrals. A breakdown of tree-traversal timings reveals that slightly more time (~ 10 %) is spent walking the tree (checking MAC and PAC, building near-field integral lists, and finding minimum multipole expansions) than in actually evaluating multipoles. This approximate balance can be shifted by employing larger multipole expansions. Significant additional improvements may then be possible through higher multipole expansions in combination with Chebyshev economiztion [87] and coordinate rotations [88]. It may also be possible to improve the efficiency of tree-traversal through advanced partitioning strategies [89] which aim to recursively subdivide workload, and/or by reusing information related to common or closely spaced bra centers.

For the series of water clusters, the use of multiresolution improves timings by approximately 5%. Without a dynamic PAC, this number increases to about 10-15 % (of much less favorable timings). For large basis set calculations, especially those supplemented with diffuse functions, the use of multiresolution is expected to play a more important role. In this context, the fast Gauss transform [16,17], may also be useful.

Although the crossover of diagonalization with fast **J** builds appears imminent from Fig. 1, the default Givens algorithm may be non-optimal. A factor of two improvement may certainly be obtained with alternative diagonalization schemes [90]. Significantly more limiting than the CPU time associated with diagonalization or computation of the exchange matrix is the $\mathcal{O}(N^2)$ memory required for scratch, MO, and density matrices, as well as the CPU time associated with unoptimized matrix multiplications.

In Table II, the CPU time required to compute the exchange matrix for P53 is less than that for charybdotoxin, which is smaller. Also, the QCTC timings for these two systems are the same. This is because charybdotoxin is more compact than P53, and the performance of ONX and QCTC depends on both the number and the density of distributions.

B. Molecular electrostatic potentials

An important molecular property readily obtained from Hartree-Fock theory is the electrostatic potential,

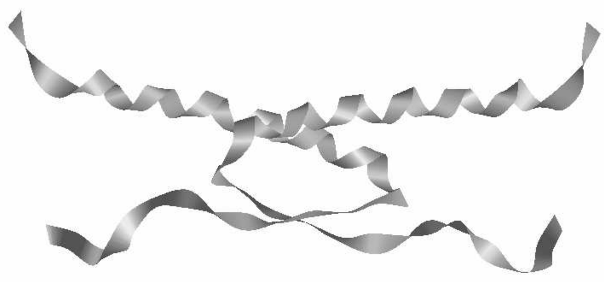

FIG. 4. Ribbons representation of the dimer formed by complexation of two P53 tetramerization monomers.

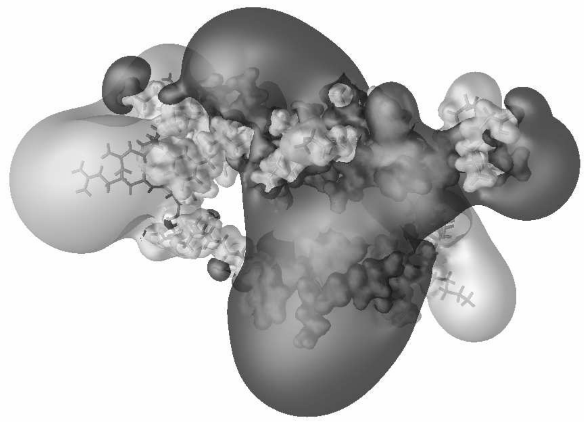

FIG. 5. Iso-surfaces of a P53 tetramerization monomer electrostatic potential at the RHF/3-21G level of theory. The light surface corresponds to a value of 0.06 au, and the darker one to -0.06 au. The orientation of this monomer is approximately that of the rightmost ribbon in Fig. 4.

which is widely implicated in molecular recognition, binding, and the enhanced diffusion of charged substrates [91]. Iso-surfaces of the electrostatic potential of charybdotoxin are shown in Fig. 3, and those of P53 in Fig. 5. A modified QCTC was used to compute the electrostatic potential, and IBM Visualization Data Explorer [92] was used to compute and render the surfaces.

Charybdotoxin belongs to a family of scorpion toxins that inhibit voltage–gated and Ca^{++} activated K^+ channels. The large, hemi-spherical positive surface (light grey) at the top of the protein is associated with the amino acid residues known to mediate the toxin–receptor interaction through mutational studies [93]; from upper left to right they are M29, R34, and K27. The association of these amino acids with this large region of positive potential corroborates a proposed mechanism of K^+ channel inhibition in which the channel mouth is competitively blocked by positively charged moieties [93,94].

Mutations in the P53 tumor suppressor are the most frequently observed genetic alterations in human cancer. The P53 protein forms a dimer of dimers (a tetramer) [86], mediated by the tetramerization domain. A ribbons diagram of the tetramerization domain dimer is shown in Fig. 4. The coordinates used in the P53 calculation were obtained from those of the tetramer, which were taken from the protein data bank. The electrostatic potential of the P53 tetramerization monomer is shown in Fig. 5, and corresponds to the rightmost ribon in Fig. 4. Starting with the monomer shown in Fig. 5, the electrostatic potential of the leftmost monomer in Fig. 4 can be visualized with a $\approx$ 160 degree rotation about the short axis, followed by a $\approx$ 90 degree rotation about the long axis, and translation to the right. This transformation roughly superimposes the two large peanut shaped lobes of opposite potential seen in Fig. 5, at the left and middle. This observation is in keeping with the hypothesis that molecular recognition and binding often involve complimentary regions of opposite potential [91].

VI. CONCLUSION

In this paper, new methods for computing the Coulomb matrix have been developed and linear scaling computation of the Hartree-Fock Fock matrix has been demonstrated for the first time. Importantly, the accuracies achieved with these methods are comparable to those obtained with standard electronic structure programs. An initial application of these methods has been made to several proteins in which the P53 calculation, involving 698 atoms and 3836 basis functions with the 3-21G basis set, may be the largest Hartree-Fock calculations to date. To our knowledge, the next largest is the RHF/4-21G (3-3-21G for S) calculation on crambin [95], involving 642 atoms and 3597 basis functions, achieved with the multiplicative integral approximation [96] implemented in the program BRABO [97].

While most of this paper has focused on the development of linear scaling methods for computation of the Coulomb matrix, it is hoped that this demonstrated success, and those of others [14,18], will help to focus attention on what are now the larger problems in SCF theory. In particular, improving the rigor and efficiency of $\mathbf{K}$ builds is an important aim of our current research. In addition, with the current hardware we have employed, memory and CPU efficient methods for solving the SCF equations are required for serial calculations significantly larger than 4000 basis functions.

ACKNOWLEDGMENTS

The advice and encouragement of Dr.s Martin Head-Gordon, Vince Ortiz, Julia Rice, Jim Chelikowsky, Bill Gleason and Bob Hill is greatly appreciated and acknowledged. Special thanks go to Dr. Shuxia Zhang for suggesting the use of real arithmetic in the bipolar expansion and to Dr. Barry Schaudt for graphical support. MC recognizes a CISE postdoctoral fellowship from the NSF with matching support from the Minnesota Supercomputer Institute. Computational resources were provided by the University of Minnesota-IBM Shared Research Project.

---

[1] D. K. W. Mok, R. Neumann, and N. C. Handy, J. Phys. Chem. **100**, 6225 (1996).
[2] A. D. Becke, J. Chem. Phys. **98**, 5648 (1993).
[3] A. D. Becke, J. Chem. Phys. **104**, 1040 (1996).
[4] B. Johnson, P. Gill, and J. Pople, J. Chem. Phys. **98**, 5612 (1993).
[5] C. W. Bauschlicher, Chem. Phys. Lett. **246**, 40 (1995).
[6] G. S. Tschumper, J. T. Fermann, and H. F. Schaefer III, J. Chem. Phys. **104**, 3676 (1996).
[7] V. Barone, C. Adamo, and F. Mele, Chem. Phys. Letters **249**, 290 (1996).
[8] J. R. Chelikowsky, N. Troullier, and Y. Saad, Phys. Rev. Lett. **72**, 1240 (1994).
[9] J. R. Chelikowsky, N. Troullier, K. Wu, and Y. Saad, Phys. Rev. B **50**, 11355 (1994).
[10] E. L. Briggs, D. J. Sulivan, and J. Bernholc, Phys. Rev. B **52**, R5471 (1995).
[11] K. Cho, T. A. Arias, J. D. Joannopoulos, and P. K. Lam, Phys. Rev. Lett. **71**, 1808 (1993).
[12] S. Q. Wei and M. Y. Chou, Phys. Rev. Lett. **76**, 2650 (1996).
[13] C. A. White, B. Johnson, P. Gill, and M. Head-Gordon, Chem. Phys. Letters **230**, 8 (1994).
[14] C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, Chem. Phys. Lett. **253**, 268 (1996).
[15] R. Kutteh and J. B. Nicholas, Chem. Phys. Lett. **238**, 173 (1995).
[16] 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.
[17] M. Challacombe, E. Schwegler, and J. Almlöf, J. Chem. Phys. **104**, 4685 (1996).
[18] M. C. Strain, G. E. Scuseria, and M. J. Frisch, Science **271**, 51 (1996).
[19] S. Wilson, Int. J. Quant. Chem. **60**, 47 (1996).
[20] P. O. Löwdin, Ann. Rev. Phys. Chem. **11**, 107 (1960).
[21] S. Obara and A. Saika, J. Chem. Phys. **84**, 3963 (1985).
[22] M. Head-Gordon and J. A. Pople, J. Chem. Phys. **89**, 5777 (1988).
[23] P. M. W. Gill, M. Head-Gordon, and J. A. Pople, Int. J. Quant. Chem. Symp. 269 (1989).
[24] T. P. Hamilton and H. F. Schaefer, Chem. Phys. **150**, 163 (1991).
[25] T. P. Hamilton and H. F. Schaefer, Can. J. Chem. **70**, 416 (1992).
[26] P. M. W. Gill and J. A. Pople, Int. J. Quant. Chem. **40**, 753 (1991).
[27] I. Panas, Chem. Phys. Lett. **184**, 86 (1991).
[28] U. Ryu, Y. S. Lee, and R. Lindh, Chem. Phys. Lett. **185**, 562 (1991).
[29] R. Lindh, Theor. Chim. Acta **85**, 423 (1993).
[30] S. Ten-no and S. Iwata, Chem. Phys. Lett. **211**, 259 (1993).
[31] B. Johnson, P. Gill, and J. Pople, Chem. Phys. Lett. **206**, 229 (1993).
[32] P. M. W. Gill, Adv. Quantum Chem. **25**, 141 (1994).
[33] C. Roothaan, Rev. Mod. Phys. **23**, 69 (1951).
[34] G. G. Hall, Proc. Roy. Soc. **208 A**, 328 (1951).
[35] J. A. Pople and R. K. Nesbet, J. Chem. Phys. **22**, 571 (1954).
[36] P. O. Löwdin, Phys. Rev. **97**, 1490 (1955).
[37] R. McWeeny, Rev. Mod. Phys. **126**, 1028 (1962).
[38] L. Cohen and C. Frishberg, J. Chem. Phys. **65**, 4234 (1976).
[39] X. P. Li, R. W. Nunes, and D. Vanderbilt, Phys. Rev. B **47**, 10891 (1993).
[40] E. Hernández, M. J. Gillan, and C. M. Goringe, 1995, available at http://www.chem.brown.edu/chem-ph.html.
[41] E. Hernández and M. J. Gillan, Phys. Rev. B **51**, 10157 (1995).
[42] X. Chen, J. M. Langlois, and W. A. Goddard III, Phys. Rev. B **54**, 2348 (1995).
[43] V. Dyczmons, Theoret. Chim. Acta **28**, 307 (1973).
[44] R. Ahlrichs, Theoret. Chim. Acta **33**, 157 (1974).
[45] J. Barnes and P. Hut, Nature **324**, 446 (1986).
[46] L. Hernquist, Ap. J. Suppl. **64**, 715 (1987).
[47] J. E. Barnes and P. Hut, Ap. J. Suppl. **70**, 389 (1989).
[48] M. S. Warren and J. K. Salmon, in *Proceedings SUPERCOMPUTING '93* (IEEE Computer Society Press, Los Alamitos, CA, 1993), pp. 12–21.
[49] S. Pfalzner and P. Gibbon, Comp. Phys. Comm. **79**, 24 (1994).
[50] J. K. Salmon, Int. J. Super. Appl. **8**, 129 (1994).
[51] M. S. Warren and J. K. Salmon, Comp. Phys. Comm. **87**, 266 (1995).
[52] L. Greengard and V. Rokhlin, J. Comp. Phys. **73**, 325 (1987).
[53] K. E. Schmidt and M. A. Lee, J. Stat. Phys. **63**, 1223 (1991).
[54] H. Petersen, D. Soelvason, J. W. Perram, and E. R. Smith, J. Chem. Phys. **101**, 8870 (1994).
[55] H. Petersen, E. R. Smith, and D. Soelvason, Proc. Roy. Soc. A **448**, 401 (1994).
[56] C. A. White and M. Head-Gordon, J. Chem. Phys. **101**, 6593 (1994).
[57] C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, Chem. Phys. Lett. **257**, 647 (1996).
[58] C. A. White and M. Head-Gordon, J. Chem. Phys. **105**, (1996).
[59] W. D. Elliot and J. A. Board Jr., Technical Report No. 94-001, Duke University Dept. of Electrical Engineering, (unpublished), available at http://www.ee.duke.edu/Research/SciComp/papers.html.
[60] J. A. Board Jr., Z. S. Hakura, W. D. Elliot, and W. T. Rankin, Technical Report No. 94-006, Duke University Dept. of Electrical Engineering, (unpublished), available at

http://www.ee.duke.edu/Research/SciComp/papers.html.
[61] E. Schwegler and M. Challacombe, J. Chem. Phys. **105**, 2726 (1996).
[62] G. R. Ahmadi and J. Almlöf, Chem. Phys. Letters **246**, 364 (1995).
[63] L. McMurchie, Ph.D. thesis, University of Seattle, 1977.
[64] L. McMurchie and E. R. Davidson, J. Comp. Phys. **26**, 218 (1978).
[65] V. R. Saunders, in *Methods in Compuational Molecular Physics*, edited by G. H. F. Diercksen and S. Wilson (D. Reidel Publishing Company, Boston, 1983), pp. 1–36.
[66] I. Shavitt, Meth. Comp. Phys. **2**, 1 (1963).
[67] P. J. Davis, in *Handbook of Mathematical Functions*, edited by M. Abramowitz and I. A. Stegun (Dover Publications Inc., New York, 1987), pp. 260–2.
[68] S. Aluru, G. M. Prabhu, and J. Gustafson, in *Proceedings SUPERCOMPUTING '94* (IEEE Computer Society Press, Los Alamitos, CA, 1995), pp. 420–6.
[69] S. Aluru, J. Sci. Comp. **17**, 773 (1996).
[70] K. Esselink, Inf. Proc. Lett. **41**, 141 (1992).
[71] K. Esselink, Comput. Phys. Comm. **87**, 375 (1995).
[72] B. C. Carlson and G. S. Rushbrooke, Proc. Cambridge Philos. Soc. **46**, 626 (1950).
[73] H. Y. Wang and R. LeSar, J. Chem. Phys. **104**, 4173 (1996).
[74] J. M. Pérez-Jordá and W. Yang, J. Chem. Phys. **104**, 8003 (1996).
[75] S. Wolfram, *Mathematica, A System for Doing Mathematics by Computer*, 2nd ed. (Addison-Wesley, Redwood City, CA., 1991).
[76] M. Sofroniou, The Mathematica Journal **3**, 74 (1993).
[77] C. Cohen-Tannoudji, B. Diu, and F. Laloë, *Quantum Mechanics* (Wiley and Sons, New York, 1977), Vol. II, see Appendix II for a useful discussion of the Dirac delta function and its derivatives.
[78] J. Applequist, J. Phys. A **22**, 4303 (1989).
[79] I. S. Gradshteyn and I. M. Ryzhik, *Table of Integrals, Series, and Products*, 2nd ed. (Academic Press, San Diego, 1980).
[80] *Handbook of Mathematical Functions*, ninth ed., edited by M. Abramowitz and I. A. Stegun (Dover Publications Inc., New York, 1987).
[81] M. Dupuis, A. Marquez, and E. R. Davidson, *HONDO 95.3 from CHEM-Station*, IBM Corporation, Neighborhood Road, Kingston, NY. 12401, 1995.
[82] H. Michels and F. Prosser, QCPE **11**, 62 (1965).
[83] M. J. Frisch *et al.*, *Gaussian 94, Revision D.1*, Gaussian Inc., Pittsburgh PA, 1995.
[84] N. H. Andersen *et al.*, Biochemistry **31**, 1280 (1992).
[85] F. Bontems *et al.*, Biochemistry **31**, 7756 (1992).
[86] G. M. Clore *et al.*, Nat. Struct. Biol. **2**, 321 (1995).
[87] S. Lustig, S. Rastogi, and N. Wagner, J. Comp. Phys. **122**, 317 (1995).
[88] C. A. White and M. Head-Gordon, J. Chem. Phys. **104**, (1996).
[89] J. R. Pilkington and S. B. Baden, IEEE Trans. Parallel and Distributed Systems. **7**, 288 (1996).
[90] W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipies in FORTRAN* (Cambridge University Press, Port Chester, New York, 1992), see page 463.
[91] B. Honig and A. Nicholls, Science **268**, 1144 (1995).
[92] *IBM Visualization Data Explorer Version 3.1*, IBM Corporation, NY 10577, 1995, DX chemistry extensions were used and are freely available from the Cornell Theory Center DX repository.
[93] S. A. Goldstein, D. J. Pheasant, and C. Miller, Neuron **12**, 1377 (1994).
[94] C. Miller, Neuron **1**, 1003 (1988).
[95] C. Van Alsenoy, A. Peeters, and J. Martin, 1996, personal communication.
[96] C. Van Alsenoy, J. Comp. Chem. **9**, 620 (1988).
[97] C. Van Alsenoy and A. Peeters, J. Mol. Struct. (THEOCHEM) **286**, 19 (1993).