

Linear-scaling formation of Kohn-Sham Hamiltonian: Application to the calculation of excitation energies and polarizabilities of large molecular systems

Mark A. Watson

*Department of Chemistry, University of Oslo, Box 1033 Blindern, N-0315 Oslo, Norway
and Department of Chemistry, University of Cambridge, Cambridge CB2 1EW, United Kingdom*

Paweł Satek

*Department of Chemistry, University of Oslo, Box 1033 Blindern, N-0315 Oslo, Norway
and Laboratory of Theoretical Chemistry, The Royal Institute of Technology, Roslagstullsbacken 15,
Stockholm SE-10691, Sweden*

Peter Macak

*Department of Chemistry, University of Oslo, Box 1033 Blindern, N-0315 Oslo, Norway
and The AlbaNova University Center, Institute of Physics, Stockholm S-10691, Sweden*

Trygve Helgaker^{a)}

Department of Chemistry, University of Oslo, Box 1033 Blindern, N-0315 Oslo, Norway

(Received 13 February 2004; accepted 19 May 2004)

We present calculations of excitation energies and polarizabilities in large molecular systems at the local-density and generalized-gradient approximation levels of density-functional theory (DFT). Our results are obtained using a linear-scaling DFT implementation in the program system DALTON for the formation of the Kohn-Sham Hamiltonian. For the Coulomb contribution, we introduce a modification of the fast multipole method to calculations over Gaussian charge distributions. It affords a simpler implementation than the original continuous fast multipole method by partitioning the electrostatic Coulomb interactions into “classical” and “nonclassical” terms which are explicitly evaluated by linear-scaling multipole techniques and a modified two-electron integral code, respectively. As an illustration of the code, we have studied the singlet and triplet excitation energies as well as the static and dynamic polarizabilities of polyethylenes, polyenes, polyynes, and graphite sheets with an emphasis on the trends observed with system size. © 2004 American Institute of Physics. [DOI: 10.1063/1.1771639]

I. INTRODUCTION

Density-functional theory (DFT) has emerged over the last 15 years to become the most widely used method of *ab initio* computational chemists for the calculation of energetics, geometries, and molecular properties. This is largely due to its inherently low computational cost. Combined with the surprisingly high accuracy afforded by state-of-the-art generalized-gradient approximation (GGA) functionals^{1–4} and hybrid functionals,^{5–7} DFT has enjoyed a far-reaching impact in the community. Moreover, DFT is a natural candidate for the application of linear-scaling techniques⁸ and great success has been achieved in this respect. Linear-scaling GGA calculations on systems with thousands of atoms and tens of thousands of basis functions are now possible with only modest computer resources.⁹ One might argue that the intense efforts seen today to address the algorithmic complexity of the “standard models” have been strongly encouraged by this development. Advances toward $O(N)$ implementations of wave function based methods such as Møller-Plesset second-order perturbation theory (MP2) (Refs. 10–12) and the coupled-cluster approximation have

also been demonstrated,^{13,14} indeed, they remain an active area of research. Such progress is essential if these methods are to be more widely applied. Already, it is clear that DFT has an important role to play in the study of biomolecules and this trend can only continue.

In this paper, we present a linear-scaling code based on the program system DALTON (Ref. 15) for the formation of the Kohn-Sham Hamiltonian in the generalized-gradient approximation. The new routines are also used in the construction of the electronic Hessian for the calculation of molecular properties using linear response theory. As a result, the traditional bottlenecks for the treatment of second-order static and dynamic properties in large molecular systems have been removed. We demonstrate the capabilities of our code with the examination of excitation energies and polarizabilities in long-chain hydrocarbons and graphite sheets.

In a traditional DFT implementation, there are three main bottlenecks to obtaining the unperturbed density of large molecular systems: the electronic Coulomb problem, the exchange-correlation quadrature, and the diagonalization (or otherwise) of the Kohn-Sham matrix. First, with the use of integral-prescreening techniques,^{16–18} the evaluation of the two-electron Coulomb integrals has an asymptotic $O(N^2)$ scaling, where N is the size of the one-particle basis;

^{a)}Electronic mail: trygve.helgaker@kjemi.uio.no

in most applications, this is the dominant step. Next, for the exchange-correlation contribution, traditional numerical quadrature requires an amount of work that scales formally as $O(N^3)$ since contributions from all basis-function products are evaluated at each grid point. Finally, diagonalization of the effective Hamiltonian is an $O(N^3)$ process. However, the prefactor associated with this cost can be made sufficiently small with modern library routines to not cause a problem except in the largest systems. Indeed, in this work, diagonalization has not inhibited us unduly but we refer to earlier work, from this group and others, exploring alternative approaches.^{19–24}

There is now a considerable literature on methods to tackle the Coulomb problem, beginning with the first developments in the classical physics community addressing the fast summation of the electrostatic energy of large collections of point charges.^{25–33} More recently, generalizations to treat the electronic quantum Coulomb problem abound.^{34–38} These methods have been dominated by hierarchical schemes, which reduce the asymptotic $O(N^2)$ scaling to the $O(N \ln N)$ and $O(N)$ regime (for homogeneous charge densities³⁹). In this work, we have used a generalization of Greengard and Rokhlin's fast multipole method (FMM) for point-charge systems.²⁸ Our approach is similar to the continuous fast multipole method (CFMM) of White *et al.*³⁴ in that we use an extent definition to identify two-electron integrals over Gaussians that may be accurately approximated by multipole expansions. Our scheme, which we term the branch-free multipole method (BFMM), is simpler, however, than the CFMM; in particular, no bookkeeping is required to avoid interacting charge distributions with significant overlap via multipole expansions. To achieve this, we have formulated explicit expressions for both the "classical" and "nonclassical" components of the electrostatic interaction, where the classical component is simply the multipole expansion energy, while the nonclassical part may be regarded as a correction term. In this way, we are able to obtain the total Coulomb energy with $O(N)$ work through two simple phases: a single FMM-like pass (treating all charge distributions equally) followed by a nonclassical correction phase requiring a minor modification of a standard two-electron integral code.

There now exist a number of successful implementations demonstrating near linear scaling of the numerical DFT integration.^{40–44} Standard techniques for exploiting the fast-decaying nature of the basis functions are insufficient to achieve the linear regime and it is necessary to avoid the $O(N^2)$ testing of significant grid-point/basis-function pairs.⁸ In this work, we have employed "batches" of grid points as introduced by Stratmann and co-workers.⁴⁰ This not only reduces the scaling, it also allows an efficient use of computer memory that avoids calling data with wildly different addresses; this is crucial if the prefactor associated with the cost is to be favorable.

The treatment of molecular properties, however, has received far less attention in the literature with respect to issues of computational complexity. In our implementation, the solution of the linear equations to obtain the first-order response is carried out in a direct iterative manner, as described

in Ref. 45. The time-consuming steps are the construction of the Kohn-Sham matrix (with a modified density) and the direct linear transformation of trial vectors with the DFT part of the electronic Hessian. Solving the linear response problem in $O(N)$ operations is nontrivial⁴⁶ and we do not attempt to achieve this here. However, by removing the traditional bottlenecks associated with the construction of the Kohn-Sham matrix and the two-electron contributions to the electronic Hessian, we have vastly increased the scope of possible study. As a demonstration of this improvement, we examine the singlet and triplet excitation energies, and static and dynamic polarizabilities of polyethylenes, polyenes, polyynes, and graphite sheets containing many hundreds of atoms, with an emphasis on the trends observed with system size.

Our work follows many other DFT studies of excitation energies and polarizabilities, including efficient implementations for the treatment of large molecules⁴⁷ within the framework of time-dependent density-functional theory.^{48,49} We also refer to the recent studies by Snijders and co-workers^{50–52} and by Cai, Sendt, and Reimers,⁵³ who have examined the performance of DFT as applied to the calculation of polarizabilities and excitation energies in extended π systems, identifying a failure of existing functionals to predict correctly qualitatively realistic polarizabilities and excitation energies of these systems. However, we do not intend our efforts to be a comprehensive study of these problems; rather, our emphasis is toward an illustration of the capabilities of our code with a view to future, more detailed, calculations.

The remainder of this paper is divided into four sections. First, in Sec. II, we discuss the theory underlying our linear-scaling code for the formation of the Kohn-Sham Hamiltonian. Next, after a short presentation of performance benchmark calculations in Sec. III, we consider in Sec. IV calculations of excitation energies and polarizabilities in linear polyethylene, polyene, polyene chains, and in graphite sheets. Finally, in Sec. V, we give some concluding remarks.

II. THEORY: LINEAR SCALING TECHNIQUES

We now present in detail and in a self-contained manner the necessary theory for our linear-scaling density-functional implementation, referring to the work of others as appropriate. We begin by partitioning the exact Coulomb interaction energy into classical and nonclassical components for which we derive explicit expressions. Having identified the classical contribution as the bipolar multipole expansion energy, we discuss our approach to computing this quantity with $O(N)$ work using fast multipole techniques. In contrast, the nonclassical term has an exponential decay with distance that automatically yields the linear regime for sufficiently large systems, as we also discuss. Finally, in Sec. II E, we highlight aspects of our implementation for a linear-scaling numerical DFT integration.

A. Partitioning of the Coulomb interaction energy

Consider two spherical Gaussian charge distributions of exponents p and q , centred at \mathbf{P} and \mathbf{Q} ,

$$U_{pq} = \int \int \frac{\exp(-pr_{1p}^2)\exp(-qr_{2Q}^2)}{r_{12}} d\mathbf{r}_1 d\mathbf{r}_2$$

$$= \frac{2\pi^{5/2}}{pq\sqrt{p+q}} F_0\left(\frac{pq}{p+q} R_{PQ}^2\right), \quad (1)$$

where the Boys function $F_n(x)$ with $n \geq 0$ is a special case of the Kummer confluent hypergeometric function $M(a, b, x)$,

$$F_n(x) = (2n+1)^{-1} M\left(n + \frac{1}{2}, n + \frac{3}{2}, -x\right)$$

$$= \int_0^1 \exp(-xt^2) t^{2n} dt. \quad (2)$$

Integrals involving nonspherical charge distributions are related to the integrals, Eq. (1), by differentiation with respect to the centres \mathbf{P} and \mathbf{Q} . Noting the recurrences

$$F'(x) = -F_{n+1}(x), \quad (3)$$

we find that all integrals over Gaussian orbitals can be expressed as linear combinations of Boys functions $F_n(x)$ of orders $n \geq 0$.

To decompose the integrals, Eq. (1), into classical and nonclassical contributions, we note that the interaction becomes purely classical in the limit of large exponents or large separations R_{PQ} . In this limit, x is large in Eq. (2) and we may extend the upper integration limit to infinity. We therefore decompose the Boys function in the manner

$$F_n(x) = F_n^{\text{cls}}(x) + F_n^{\text{non}}(x), \quad (4)$$

where the classical contribution is given by

$$F_n^{\text{cls}}(x) = \int_0^\infty \exp(-xt^2) t^{2n} dt = \frac{(2n-1)!!}{2^{n+1}} \sqrt{\frac{\pi}{x^{2n+1}}} > 0, \quad (5)$$

and the nonclassical correction term by

$$F_n^{\text{non}}(x) = - \int_1^\infty \exp(-xt^2) t^{2n} dt = -\frac{1}{2} E_{1/2-n}(x) < 0, \quad (6)$$

where $E_n(x)$ is the exponential integral function

$$E_n(x) = \int_1^\infty \exp(-xt) t^{-n} dt = x^{-1} \exp(-x) [1 + O(x^{-1})]. \quad (7)$$

For large separations R_{PQ} , the negative nonclassical contribution Eq. (6) decays exponentially in R_{PQ}^2 , and the Boys function becomes dominated by the positive classical contribution, Eq. (5), which is proportional to R_{PQ}^{-2n-1} . Both $F_n^{\text{cls}}(x)$ and $F_n^{\text{non}}(x)$ satisfy the same recurrence relation as the full Boys function, Eq. (3), and may therefore be calculated using a standard two-electron integral program with very minor modifications.

We now identify the expansion of a general electrostatic interaction in the derivatives $F_n^{\text{cls}}(x)$ with the classical bipolar multipole expansion—see also Eq. (60). For spherical distributions, this may be seen explicitly by substituting Eqs. (5) and Eq. (6) into Eq. (1), to obtain

$$U_{pq} = U_{pq}^{\text{cls}} + U_{pq}^{\text{non}}, \quad (8)$$

where

$$U_{pq}^{\text{cls}} = \left(\frac{\pi^2}{pq}\right)^{3/2} \frac{1}{R_{PQ}}, \quad (9)$$

$$U_{pq}^{\text{non}} = -\frac{\pi^{5/2}}{pq\sqrt{p+q}} E_{1/2}\left(\frac{pq}{p+q} R_{PQ}^2\right). \quad (10)$$

We immediately recognize the first term as the multipole expansion energy to zero order, while the second term can be regarded as the nonclassical correction accounting for charge overlap effects.

Owing to the exponential decay of the exponential integral function, Eq. (7), for sufficiently large systems, the number of significant nonclassical contributions will grow just linearly with the system size. It is therefore trivial to implement an efficient linear-scaling code for the computation of U_{pq}^{non} by prescreening away negligible nonclassical contributions. This may be achieved in a number of ways. For spherical Gaussians, it is possible to establish a simple expression relating the relative sizes of U_{pq}^{cls} and U_{pq}^{non} . We note that several expressions have been suggested,^{34,35,38,54,55} but we employ the original definition of White *et al.*³⁴ where the extent r_p is given as

$$r_p = \frac{1}{\sqrt{p}} \operatorname{erfc}^{-1}(10^{-k}). \quad (11)$$

Here, p is the reduced exponent of the overlap distribution, k is an arbitrary constant, and $\operatorname{erfc}^{-1}(x)$ is the inverse complementary error function. For spherical Gaussians separated by more than the sum of their extents, it follows that

$$\left| \frac{U_{pq}^{\text{non}}}{U_{pq}^{\text{cls}}} \right| \leq 10^{-k}. \quad (12)$$

For a given choice of k , it is then possible to identify interactions that have a significant nonclassical contribution. Such interactions may be treated in total by conventional integration, or by a combination of multipole techniques and nonclassical correction, according to Eqs. (5) and (6).

However, in the limit of small exponents p and q and small separations R_{PQ} , the classical and nonclassical contributions to the integrals (and to the Boys function) tend to plus and minus infinity, respectively. In these cases, the decomposition into classical and nonclassical contributions becomes numerically unstable and should be avoided. This is an issue particular to our multipole implementation, as discussed in Sec. IID 2.

In contrast to the nonclassical term, the number of long-range classical interactions U_{pq}^{cls} grows quadratically with the system size. To achieve linear scaling, therefore, this contribution must be evaluated using a fast summation method afforded by the use of multipole expansion techniques, as we discuss in Sec. IID.

B. Evaluation of the Coulomb contribution to the Kohn-Sham matrix

In this section, we consider the evaluation of the Coulomb matrix elements by the McMurchie-Davidson integra-

tion scheme⁵⁶ in a manner that avoids the explicit construction of the electron-repulsion integrals through a presummation of the density. Moreover, we decompose our expressions into classical and nonclassical parts such that the expensive long-range contributions may be rapidly evaluated by multipole techniques. We note that this work has significant overlap with previous publications concerning the efficient evaluation of J -matrix elements and refer to Refs. 57–60 as well as our presentation in Ref. 61.

Consider a basis of real-valued solid-harmonic Gaussians,

$$G_{l\mu}(\mathbf{r}_A) = S_{l\mu}(\mathbf{r}_A) \exp(-ar_A^2). \quad (13)$$

Such functions may be expanded in linear combinations of Cartesian Gaussians,

$$G_{l\mu}(\mathbf{r}_A) = \sum_{ijk} C_{l\mu}^{ijk} G_{ijk}(\mathbf{r}_A), \quad (14)$$

$$G_{ijk}(\mathbf{r}_A) = x_A^i y_A^j z_A^k \exp(-ar_A^2).$$

Following McMurchie and Davidson,⁵⁶ a product of two such Gaussians with exponents a and b and centred at \mathbf{A} and \mathbf{B} ,

$$\Omega_{ab}(\mathbf{r}) = G_{ikm}(\mathbf{r}_A) G_{jln}(\mathbf{r}_B), \quad (15)$$

may be expressed as a linear combination of Gaussians,

$$\Omega_{ab}(\mathbf{r}) = \sum_{i=0}^{i+j} \sum_{u=0}^{k+l} \sum_{v=0}^{m+n} E_t^{ij} E_u^{kl} E_v^{mn} \frac{\partial^{t+u+v} \exp(-pr_P^2)}{\partial P_x^t \partial P_y^u \partial P_z^v}$$

$$= \sum_{tuv} E_{tuv}^{ab} \frac{\partial^{t+u+v} \exp(-pr_P^2)}{\partial P_x^t \partial P_y^u \partial P_z^v}, \quad (16)$$

where we have introduced

$$\mathbf{P} = \frac{a}{p} \mathbf{A} + \frac{b}{p} \mathbf{B}, \quad p = a + b. \quad (17)$$

The expansion coefficients E_t^{ij} are traditionally determined by a set of three-term recurrence relations presented by McMurchie and Davidson.⁵⁶ However, we here use the two-term relations⁶¹

$$E_0^{00} = \exp\left(-\frac{ab}{a+b} X_{AB}^2\right), \quad (18)$$

$$E_0^{i+1,j} = X_{PA} E_0^{ij} + E_1^{ij}, \quad (19)$$

$$E_0^{i,j+1} = X_{PB} E_0^{ij} + E_1^{ij}, \quad (20)$$

$$E_t^{ij} = \frac{1}{2pt} (iE_{t-1}^{i-1,j} + jE_{t-1}^{i,j-1}), \quad t > 0, \quad (21)$$

which are simpler and more efficient than the corresponding McMurchie-Davidson recurrence relations. Similar relations are valid for the y and z components of the expansion coefficients.

In a manner similar to Eq. (16), we arrive at the result for the expansion of the one-electron density as a linear combination of differentiated Gaussians,

$$\rho(\mathbf{r}) = \sum_{cd} D_{cd} \Omega_{cd}(\mathbf{r})$$

$$= \sum_Q \sum_{\tau\nu\phi} (-1)^{\tau+\nu+\phi} F_{\tau\nu\phi}^Q \frac{\partial^{\tau+\nu+\phi} \exp(-qr_Q^2)}{\partial Q_x^\tau \partial Q_y^\nu \partial Q_z^\phi}, \quad (22)$$

where we have introduced

$$\mathbf{Q} = \frac{c\mathbf{C} + d\mathbf{D}}{q}, \quad q = c + d, \quad (23)$$

$$F_{\tau\nu\phi}^Q = (-1)^{\tau+\nu+\phi} \sum_{cd \in Q} D_{cd} E_{\tau\nu\phi}^{cd}, \quad (24)$$

and Q is the subset of overlap distributions with exponent q centred at \mathbf{Q} . The Coulomb contribution to the Kohn-Sham matrix may then be calculated as

$$J_{ab} = \frac{2\pi^{5/2}}{pq\sqrt{p+q}} \sum_{tuv} E_{tuv}^{ab} \sum_Q \sum_{\tau\nu\phi} F_{\tau\nu\phi}^Q$$

$$\times \frac{\partial^{t+\tau+u+\nu+v+\phi}}{\partial P_x^{t+\tau} \partial P_y^{u+\nu} \partial P_z^{v+\phi}} F_0\left(\frac{pq}{p+q} R_{PQ}^2\right). \quad (25)$$

Decomposing this interaction into classical and nonclassical contributions,

$$J_{ab} = J_{ab}^{\text{cls}} + J_{ab}^{\text{non}}, \quad (26)$$

we obtain

$$J_{ab}^{\text{cls}} = \left(\frac{\pi^2}{pq}\right)^{3/2} \sum_{tuv} E_{tuv}^{ab} \sum_Q \sum_{\tau\nu\phi} F_{\tau\nu\phi}^Q$$

$$\times \frac{\partial^{t+\tau+u+\nu+v+\phi}}{\partial P_x^{t+\tau} \partial P_y^{u+\nu} \partial P_z^{v+\phi}} \frac{1}{R_{PQ}}, \quad (27)$$

$$J_{ab}^{\text{non}} = -\frac{\pi^{5/2}}{pq\sqrt{p+q}} \sum_{tuv} E_{tuv}^{ab} \sum_Q \sum_{\tau\nu\phi} F_{\tau\nu\phi}^Q$$

$$\times \frac{\partial^{t+\tau+u+\nu+v+\phi}}{\partial P_x^{t+\tau} \partial P_y^{u+\nu} \partial P_z^{v+\phi}} E_{1/2}\left(\frac{pq}{p+q} R_{PQ}^2\right), \quad (28)$$

where $E_n(x)$ is given by Eq. (7). As discussed in the preceding section, the nonclassical contribution J_{ab}^{non} may be evaluated from Eq. (28) with $O(N)$ work by prescreening away negligible terms. In the following sections, we discuss how the long-range classical contribution J_{ab}^{cls} may also be computed in linear time.

For an efficient evaluation of the Coulomb contribution, it is necessary to carry out a careful screening of negligible contributions to the Kohn-Sham matrix. We base our screening on the Schwarz inequality, which, when applied to the Hermitian integrals

$$R_{tuv}(\alpha, \mathbf{R}_{PQ}) = \frac{\partial^{t+u+v} F_0(\alpha R_{PQ}^2)}{\partial P_x^t \partial P_y^u \partial P_z^v}, \quad (29)$$

rather than to the full Gaussian integrals as in Ref. 62, provides the following upper bound:

$$\left| R_{t+\tau, u+v, \phi} \left(\frac{pq}{p+q}, \mathbf{R}_{PQ} \right) \right| \leq \sqrt{\left| R_{2t, 2u, 2v} \left(\frac{p}{2}, \mathbf{0} \right) \right|} \sqrt{\left| R_{2\tau, 2\nu, 2\phi} \left(\frac{q}{2}, \mathbf{0} \right) \right|}. \quad (30)$$

The one-center Hermite integrals in this expression are easily calculated from the formula

$$R_{2t, 2u, 2v}(\alpha, \mathbf{0}) = \frac{(2t-1)!!(2u-1)!!(2v-1)!!(-\alpha)^{t+u+v}}{2t+2u+2v+1}. \quad (31)$$

When evaluating the Coulomb expression according to Eq. (25), this upper bound is used to identify negligible contributions. For the size correction Eq. (28), estimates may be based on the properties $0 < E_n(x) < x^{-1} \exp(-x)$ of the exponential integral function in Eq. (7).

C. Efficient multipole expansion techniques

In this section, we give detailed equations for the multipole-expansion techniques as used in our implementation. These formulas are a development of the work of White³⁰ and Pérez-Jordá.⁶³ However, we have emphasized some of the symmetries and scaling relations we found useful for an efficient implementation which we did not find in previous work and have extended the real formulation of the solid spherical harmonics to include the translation of the multipole moments. We refer also to our work in Ref. 61 for further discussion.

1. Regular and irregular solid harmonics

In Racah's normalization, the complex solid harmonics are given by

$$C_{lm}(\mathbf{r}) = \sqrt{\frac{4\pi}{2l+1}} r^l Y_{lm}(\theta, \varphi), \quad (32)$$

where the $Y_{lm}(\theta, \varphi)$ are the standard spherical harmonic functions. The solid harmonics obey the relations

$$C_{lm}(a\mathbf{r}) = a^l C_{lm}(\mathbf{r}), \quad C_{lm}^*(\mathbf{r}) = (-1)^m C_{l, -m}(\mathbf{r}), \quad (33)$$

where a is some real (positive or negative) scalar variable.

We then introduce the scaled complex regular and irregular solid harmonics as

$$R_{lm}(\mathbf{r}) = \frac{1}{\sqrt{(l-m)!(l+m)!}} C_{lm}(\mathbf{r}), \quad (34)$$

$$I_{lm}(\mathbf{r}) = \sqrt{(l-m)!(l+m)!} r^{-2l-1} C_{lm}(\mathbf{r}). \quad (35)$$

These may be further decomposed into real (cosine) and imaginary (sine) parts,

$$R_{lm}(\mathbf{r}) = R_{lm}^c(\mathbf{r}) + iR_{lm}^s(\mathbf{r}), \quad (36)$$

$$I_{lm}(\mathbf{r}) = I_{lm}^c(\mathbf{r}) + iI_{lm}^s(\mathbf{r}), \quad (37)$$

which satisfy the following relations:

$$R_{l, -m}^c(\mathbf{r}) = (-1)^m R_{lm}^c(\mathbf{r}), \quad R_{l, -m}^s(\mathbf{r}) = -(-1)^m R_{lm}^s(\mathbf{r}), \quad (38)$$

$$I_{l, -m}^c(\mathbf{r}) = (-1)^m I_{lm}^c(\mathbf{r}), \quad I_{l, -m}^s(\mathbf{r}) = -(-1)^m I_{lm}^s(\mathbf{r}), \quad (39)$$

$$R_{l, m}^{c/s}(a\mathbf{r}) = a^l R_{lm}^{c/s}(\mathbf{r}), \quad I_{l, m}^{c/s}(a\mathbf{r}) = a^{-l-1} I_{lm}^{c/s}(\mathbf{r}). \quad (40)$$

The scaled harmonics may be conveniently constructed from recurrence relations. For the regular solid harmonics, we find

$$R_{00}^c = 1, \quad (41)$$

$$R_{00}^s = 0, \quad (42)$$

$$R_{l+1, l+1}^c = -\frac{xR_{ll}^c - yR_{ll}^s}{2l+2}, \quad (43)$$

$$R_{l+1, l+1}^s = -\frac{yR_{ll}^c + xR_{ll}^s}{2l+2}, \quad (44)$$

$$R_{l+1, m}^{c/s} = \frac{(2l+1)zR_{lm}^{c/s} - r^2 R_{l-1, m}^{c/s}}{(l+m+1)(l-m+1)}, \quad 0 \leq m < l, \quad (45)$$

while the corresponding relations for the irregular harmonics are given by

$$I_{00}^c = \frac{1}{r}, \quad (46)$$

$$I_{00}^s = 0, \quad (47)$$

$$I_{l+1, l+1}^c = -(2l+1) \frac{xI_{ll}^c - yI_{ll}^s}{r^2}, \quad (48)$$

$$I_{l+1, l+1}^s = -(2l+1) \frac{yI_{ll}^c + xI_{ll}^s}{r^2}, \quad (49)$$

$$I_{l+1, m}^{c/s} = \frac{(2l+1)zI_{lm}^{c/s} - (l^2 - m^2)I_{l-1, m}^{c/s}}{r^2}, \quad 0 \leq m < l. \quad (50)$$

For brevity of notation, we shall sometimes employ a common notation for the cosine and sine components by the use of greek letters, adopting the following convention to distinguish between them:

$$R_{l\mu}(\mathbf{r}) = \begin{cases} R_{l\mu}^c(\mathbf{r}), & \mu \geq 0 \\ R_{l\mu}^s(\mathbf{r}), & \mu < 0 \end{cases} \quad (51)$$

2. Multipole moments and their translations

Let us now consider some charge distribution $\rho_p(\mathbf{r})$. We may associate with this distribution the multipole moments centred on \mathbf{P}

$$q_{l\mu}^p(\mathbf{P}) = \int \rho_p(\mathbf{r}) R_{l\mu}(\mathbf{r}_P) d\mathbf{r}. \quad (52)$$

Different centres of expansion give rise to different multipole expansions. However, all such expansions are related by a linear transformation of the form

$$q_{l\mu}^p(\bar{\mathbf{P}}) = \sum_{j=0}^l \sum_{\kappa=-j}^j W_{l\mu, j\kappa}(\mathbf{R}_{\bar{P}P}) q_{j\kappa}^p(\mathbf{P}), \quad (53)$$

where we have used the notation

$$\mathbf{R}_{AB} = \mathbf{A} - \mathbf{B}. \quad (54)$$

The elements of the translation matrix $\mathbf{W}(\mathbf{r})$ are given by

$$W_{lm,jk}^{cc}(\mathbf{r}) = (-1)^{l-j} \left(\frac{1}{2}\right)^{\delta_{k0}} [R_{l-j,m-k}^c(\mathbf{r}) + (-1)^k R_{l-j,m+k}^c(\mathbf{r})], \quad (55)$$

$$W_{lm,jk}^{cs}(\mathbf{r}) = (-1)^{l-j} [-R_{l-j,m-k}^s(\mathbf{r}) + (-1)^k R_{l-j,m+k}^s(\mathbf{r})], \quad (56)$$

$$W_{lm,jk}^{sc}(\mathbf{r}) = (-1)^{l-j} \left(\frac{1}{2}\right)^{\delta_{k0}} [R_{l-j,m-k}^s(\mathbf{r}) + (-1)^k R_{l-j,m+k}^s(\mathbf{r})], \quad (57)$$

$$W_{lm,jk}^{ss}(\mathbf{r}) = (-1)^{l-j} [R_{l-j,m-k}^c(\mathbf{r}) - (-1)^k R_{l-j,m+k}^c(\mathbf{r})], \quad (58)$$

and satisfy the scaling relation

$$W_{l\mu,j\kappa}(a\mathbf{r}) = a^{l-j} W_{l\mu,j\kappa}(\mathbf{r}). \quad (59)$$

Note that $\mathbf{W}(\mathbf{r})$ constitutes a lower unit triangular matrix of infinite dimension with the inverse $\mathbf{W}^{-1}(\mathbf{r}) = \mathbf{W}(-\mathbf{r})$. Therefore, only multipole moments $q_{j\kappa}^p(\mathbf{P})$ with $j\kappa \leq l\mu$ contribute to $q_{l\mu}^p(\bar{\mathbf{P}})$ in Eq. (53). In practice, the expansion at $\bar{\mathbf{P}}$ is truncated at some sufficiently high order $l \leq L$.

3. The bipolar multipole expansion

Let us now consider two nonoverlapping charge distributions, represented by the multipole moments $q_{l\mu}^p(\mathbf{P})$ and $q_{j\kappa}^q(\mathbf{Q})$ associated with the centers \mathbf{P} and \mathbf{Q} , respectively. The energy of interaction between these charges is then given by the following bipolar multipole expansion:

$$U_{pq}^{\text{cls}} = \sum_{l\mu} \sum_{j\kappa} q_{l\mu}^p(\mathbf{P}) T_{l\mu,j\kappa}(\mathbf{R}_{QP}) q_{j\kappa}^q(\mathbf{Q}), \quad (60)$$

where the elements of the interaction matrix $\mathbf{T}(\mathbf{r})$ may be calculated from the irregular solid harmonics in the following manner:

$$T_{lm,jk}^{cc}(\mathbf{r}) = (-1)^j \left(\frac{1}{2}\right)^{\delta_{m0} + \delta_{k0} - 1} [I_{l+j,m+k}^c(\mathbf{r}) + (-1)^k I_{l+j,m-k}^c(\mathbf{r})], \quad (61)$$

$$T_{lm,jk}^{cs}(\mathbf{r}) = (-1)^j \left(\frac{1}{2}\right)^{\delta_{m0} + \delta_{k0} - 1} [I_{l+j,m+k}^s(\mathbf{r}) - (-1)^k I_{l+j,m-k}^s(\mathbf{r})], \quad (62)$$

$$T_{lm,jk}^{sc}(\mathbf{r}) = (-1)^j \left(\frac{1}{2}\right)^{\delta_{m0} + \delta_{k0} - 1} [I_{l+j,m+k}^s(\mathbf{r}) + (-1)^k I_{l+j,m-k}^s(\mathbf{r})], \quad (63)$$

$$T_{lm,jk}^{ss}(\mathbf{r}) = (-1)^j \left(\frac{1}{2}\right)^{\delta_{m0} + \delta_{k0} - 1} [-I_{l+j,m+k}^c(\mathbf{r}) + (-1)^k I_{l+j,m-k}^c(\mathbf{r})]. \quad (64)$$

We note the symmetry and scaling relations,

$$T_{j\kappa,l\mu}(\mathbf{r}) = (-1)^{l-j} T_{l\mu,j\kappa}(\mathbf{r}), \quad (65)$$

$$T_{l\mu,j\kappa}(a\mathbf{r}) = a^{-l-j-1} T_{l\mu,j\kappa}(\mathbf{r}). \quad (66)$$

Thus, although the matrix \mathbf{T} itself is not symmetric, it is related to the symmetric matrix $\tilde{\mathbf{T}}$ by the relation

$$\tilde{T}_{l\mu,j\kappa} = (-1)^{-j} T_{l\mu,j\kappa}. \quad (67)$$

Introducing the modified moments

$$\tilde{q}_{j\kappa}^q(\mathbf{Q}) = (-1)^j q_{j\kappa}^q(\mathbf{Q}), \quad (68)$$

we may write the energy of interaction as

$$U_{pq}^{\text{cls}} = \mathbf{q}^p \tilde{\mathbf{T}}(\mathbf{R}_{QP}) \tilde{\mathbf{q}}^q. \quad (69)$$

We may also calculate the energy of interaction, Eq. (60), in a two-step manner:

$$V_{l\mu}^q(\mathbf{P}) = \sum_{j=0}^l \sum_{\kappa=-j}^j T_{l\mu,j\kappa}(\mathbf{R}_{QP}) q_{j\kappa}^q(\mathbf{Q}), \quad (70)$$

$$U_{pq}^{\text{cls}} = \sum_{l\mu} q_{l\mu}^p(\mathbf{P}) V_{l\mu}^q(\mathbf{P}), \quad (71)$$

where $V_{l\mu}^q(\mathbf{P})$ is the potential at \mathbf{P} generated by the charge distribution associated with the multipole moments $q_{j\kappa}^q(\mathbf{Q})$.

In the course of the FMM evaluation, it is often necessary to generate the potentials $V_{l\mu}^q(a\mathbf{P})$ arising from a given multipole expansion $q_{j\kappa}^q(\mathbf{Q})$ in Eq. (70) for many values of a . In such cases, we may use the scaling relation, Eq. (66), to avoid recalculating the interaction matrix for each value of a since

$$\begin{aligned} V_{l\mu}^q(a\mathbf{P}) &= \sum_{j\kappa} T_{l\mu,j\kappa}(a\mathbf{R}_{QP}) q_{j\kappa}^q(\mathbf{Q}) \\ &= a^{-l-1} \sum_{j\kappa} T_{l\mu,j\kappa}(\mathbf{R}_{QP}) a^{-j} q_{j\kappa}^q(\mathbf{Q}). \end{aligned} \quad (72)$$

Thus, rather than constructing a new interaction matrix, we calculate the potential from the scaled moments $a^{-j} q_{j\kappa}^q(\mathbf{R}_Q)$. Use of this scaling relation significantly reduces the number of \mathbf{T} matrices that need to be explicitly constructed. This is an important reduction since each such construction involves $O(L^4)$ operations, whereas the scaling of the multipoles and their potentials involves just $O(L^2)$ work.

D. Multipole method for large systems

The theory for the rapid summation of the total Coulomb energy of a large collection of point charges by means of the FMM of Greengard and Rokhlin is now well established.²⁸ The generalization to continuous distributions involves some complexity since the conventional criteria for constructing a convergent multipole expansion requires that the charge distributions are nonoverlapping. In the CFMM, this generalization is achieved by defining an extent for each distribution; Gaussians separated by more than the sum of their extents are considered to interact “classically,” such that the effects of charge overlap can be safely neglected. The remaining interactions are evaluated by explicit integration. In the context of the CFMM algorithm, a coarse graining of the extent is then required when the moments are grouped into boxes. This introduces a “well-separatedness” index Ω , describing the necessary separation between boxes that ensures that only classical Coulomb interactions are evaluated. Moments are further grouped into branches corresponding to common values of Ω . A standard point-charge FMM algorithm is then

applied to all pairs of branches to yield an asymptotic $O(N)$ scaling of computational cost for large systems.

By contrast, our extension of the FMM for quantum chemical calculations avoids the book keeping complexity of the original CFMM by avoiding the introduction of branches. It is very much a pure generalization of the point-charge FMM, requiring only minor changes, and we refer to it as the BFMM. This is possible through the use of Eqs. (27) and (28) which allows us the flexibility to treat all interactions using multipole techniques, regardless of the overlap, and to correct for the nonclassical components in a separate phase; in particular, no special book keeping is required to avoid the nonclassical contributions in the multipole phase.

1. Evaluation of the far-field classical Coulomb contribution

We now describe our implementation for computing the far-field classical Coulomb contribution of the Kohn-Sham matrix by the BFMM. Our discussion is very similar to previous presentations of the original FMM since we have deliberately designed our approach to be a simple generalization of the point-charge scheme. We note that we have also implemented a tree algorithm with $O(N \ln N)$ scaling as a special case of our linear scaling code, which requires less memory, but due to the greater number of multipole contractions involved, is generally slower than the $O(N)$ method with respect to CPU time. In addition, all our algorithms also work within the conventional CFMM framework, using branches to ensure that only purely classical interactions are treated by multipole expansions. We presently discuss only our $O(N)$ BFMM algorithm, but note that our CFMM implementation uses the same engine with some additional book keeping to take care of the branchlike structures.

To describe the BFMM, it is necessary to introduce the concept of a hierarchy of boxes: a three-dimensional mesh of boxes generated recursively to create an *oct-tree* such that each parent box is split into eight children per level in the hierarchy. Termination is defined by an input parameter which specifies the size of the smallest box, rather than by an explicit occupancy criterion. We note that the choice of box size is arbitrary with respect to achieving linear scaling and will only affect the prefactor of the overall cost. Our default volume of $8a_0^3$ has proved to be near optimal for a wide range of molecular systems and basis sets, such as those used in the present work.

With respect to a given reference box, the boxes are grouped into three classes: the *nearest-neighbor* (NN) boxes; the *local far-field* (LFF) boxes, which correspond to the children of the NN boxes of the *parent* of the reference box (and are not NN boxes); and the *remote far-field* (RFF) boxes, which include all the remaining boxes. We note that a child's RFF corresponds to its parent's *total* far field (FF), where the FF space corresponds to the LFF and RFF boxes combined.

The purpose of the NN space is to keep multipole moments well separated: the half width of the space corresponds to the well-separatedness index Ω , which is a coarse graining of the extent. In conventional CFMM, this width is distinct for each branch in the CFMM tree, but in the original point-charge FMM algorithm it is a constant, usually taken to be

two boxes; that is, no charges within a box pair separated by less than two intervening boxes are to be interacted by multipole expansions.

In the BFMM, we employ the same conceptual hierarchy of boxes as in point-charge FMM, defining a single, constant, well-separatedness index Ω , whose value is discussed below. In this manner, it follows that the BFMM is identical to the point-charge FMM algorithm in its treatment of the FF interactions. For reference, we now summarize our procedure for obtaining the FF interaction energy (for Gaussian distributions) according to the original FMM scheme.

(1) Generate $O(N)$ multipole moments \mathbf{q}_i^p corresponding to the $O(N)$ significant Gaussian overlap distributions centred at \mathbf{P}_i .

(2) Associate each moment with a box at the deepest level in the hierarchy corresponding to the smallest boxes.

(3) Translate and sum the \mathbf{q}_i^p to their box centers to generate the first set of "boxed" moments at the deepest level.

(4) Recursively translate and sum the boxed moments from child boxes to the center of the parent boxes through the whole box hierarchy. There are $O(N)$ boxes in the hierarchy so this step is also $O(N)$, both in CPU time and in the memory required to store the moments.

(5) The far-field potential is now generated at the center of all boxes at all levels with $O(N)$ work as follows.

LOOP over all levels L in the box hierarchy, starting with the largest boxes.

LOOP over all boxes C at level L in the hierarchy.

(a) Generate the LFF potential $\mathbf{V}_C^{\text{LFF}}$ at the center of box C due to all LFF boxes associated with C .

(b) Translate the FF potential of the parent of box C , \mathbf{V}_P^{FF} , to the center of box C to obtain the RFF potential $\mathbf{V}_C^{\text{RFF}}$ as

$$\mathbf{V}_C^{\text{RFF}} = \mathbf{W}_{PC}^T \mathbf{V}_P. \quad (73)$$

(c) Generate the total FF potential at the center of box C , \mathbf{V}_C^{FF} , as

$$\mathbf{V}_C^{\text{FF}} = \mathbf{V}_C^{\text{LFF}} + \mathbf{V}_C^{\text{RFF}}. \quad (74)$$

END LOOP over boxes at level L .

END LOOP over levels in the box hierarchy.

(6) The total FF multipole contribution to the Coulomb energy may now be obtained in $O(N)$ time by contracting the FF potentials at the center of the boxes at the deepest level with the boxed moments at the deepest level.

The remaining interactions to be considered are those between charges in NN boxes at the deepest level. These local interactions may all be evaluated conventionally in linear time. However, many interactions will be purely classical; we prefer to evaluate as many classical interactions as possible using multipole techniques. In our CFMM implementation, therefore, all nearest-neighbor interactions are tested by extent and evaluated by multipole expansion only when purely classical. This testing requires a quadratic loop over all branches, each pair defining its own NN space. Finally, the remaining interactions are recovered in the non-

classical phase. In the BFMM, however, the situation is slightly different; in particular, only a single NN space needs to be considered, as we now discuss.

2. Evaluation of the local classical Coulomb contribution

To compute the total Coulomb energy within the BFMM, it is necessary to add on to the FF classical contribution the following two contributions: all nonclassical corrections neglected in the FF phase, and all total (i.e., classical and nonclassical) NN contributions. Clearly, the nonclassical part U_{pq}^{non} of every two-electron integral must be evaluated by conventional means. However, it is unacceptable to take the sum of the total U_{pq}^{non} evaluated from Eq. (28) and all NN and FF contributions evaluated by multipole expansion according to Eq. (26). For reasons of numerical stability, as discussed in Sec. II A, some of these contributions must be evaluated in total by conventional means according to Eq. (25). We outline this requirement first.

We define the *classical cutoff radius* R_{\min} as the distance between charge distribution centers within which all interactions are evaluated, in total, by conventional integration Eq. (25). For such interactions, it is expected that the classical and nonclassical contributions will each be very large, but of different sign, when accurate cancellation is not guaranteed. Some of these interactions may be dominantly classical (perhaps if R_{\min} is chosen to be very large), but in our default implementation we compute all such interactions conventionally in any case. It should be noted that we require

$$R_{\min} < \Omega L_{\text{box}}, \quad (75)$$

where Ω is the well-separated index previously discussed and L_{box} is the dimension of a single box at the deepest level. In this way, we ensure that the far field does not include any “forbidden” multipole expansion contributions. Our default value for R_{\min} is $2a_0$, although numerous tests suggest that the performance is not strongly affected by this choice, for sensible values of the smallest box size.

Since the only role of Ω in the BFMM is to ensure that FF interactions contracted by multipole expansion do not contain contributions from very closely spaced moments, it may be much smaller than in the CFMM, where many boxes are required to keep moderately diffuse Gaussians well separated. Indeed, we typically take Ω to be as small as possible while remaining consistent with Eq. (75). In this way, we may optimize the savings of grouping moments into larger boxes as efficiently as possible, without large NN and LFF spaces. In Sec. III, we briefly discuss some of the performance benefits we have observed from this improvement.

It is therefore necessary in the final pass of the multipole phase to test all pairs of multipoles in the NN space with respect to R_{\min} ; only if they lie outside the cutoff are they contracted by multipole expansion. Similarly, in the nonclassical pass, if an interaction lies inside the cutoff then it must be evaluated in total regardless of the result of the extent test, else only the nonclassical component will be evaluated, or otherwise skipped.

This completes the generation of the Coulomb energy using the linear scaling BFMM algorithm which can be seen

to involve an FMM-like pass which has no more complexity than the original point-charge FMM algorithm, followed by a nonclassical correction phase.

E. Linear-scaling evaluation of the exchange-correlation contribution

The DFT exchange-correlation evaluation consists of four main steps: the grid generation, the evaluation of the electron density, the evaluation of the Kohn-Sham matrix elements, and the evaluation of the effective density contributing to the linear response trial vectors. Each of these steps has been implemented in a linear-scaling manner, by exploiting the sparsity of the density matrix.

1. Grid generation

The molecular grid for the numerical integration is generated by a superposition of atomic grids, followed by a weight correction in the overlapping regions by means of space partitioning. If w_i^A is the weight of an atomic grid point \mathbf{r}_i associated with atom A , the corrected weight w_i is given by

$$w_i = w_i^A \frac{\prod_{B \neq A}^N f(\mu_i^{AB})}{\sum_{C=1}^N \prod_{B \neq C} f(\mu_i^{BC})},$$

$$\mu_i^{AB} = \frac{|\mathbf{r}_i - \mathbf{R}_A| - |\mathbf{r}_i - \mathbf{R}_B|}{|\mathbf{R}_A - \mathbf{R}_B|}, \quad (76)$$

where N is the number of atoms and where the partitioning function $f(\mu)$ has the properties $f(0)=0.5$, $f(-1)=0$, $f(1)=0$. Common choices of $f(\mu)$ are the Becke partitioning function⁶⁴

$$f(\mu) = \frac{1}{2}[1 - g(g(\mu))], \quad g(\mu) = \frac{1}{2}\mu(3 - \mu^2) \quad (77)$$

as well as the Stratmann-Scuseria-Frisch (SSF) partitioning function.⁴⁰

From Eq. (76), it is clear that a simple-minded implementation of the partitioning algorithm scales as $O(N^3)$, since we must perform $O(N^2)$ operations on the grid. However, the scaling may be reduced to the linear regime by observing that the corrected weights generated from points associated with atom A are essentially unaffected by excluding contributions from atoms B far from A . Therefore, we distribute the N atoms among a set of L boxes such that all atoms in one box are local in space. The corrected weights may then be obtained in linear time by evaluating Eq. (76) separately for each box. Each box $1 \leq n \leq L$ has an associated weight M_n such that $\sum_n M_n = M$, where M is the total number of (linear-scaling) grid points.

For each box, we also determine a list of the nonvanishing basis functions, thereby reducing the number of basis functions to be included in subsequent summations. The number of nonvanishing functions K_n in box n depends on the density of basis functions and their typical extent—diffuse functions contribute to many boxes, while steep functions describing core orbitals may be present in one box only. To achieve linear scaling in the following discussion, we require that the number of nonvanishing functions per box becomes saturated as the system size increases. In general,

this is achieved rapidly on the molecular scale, when the system in question is large in comparison to the extent of the basis functions used.

2. Evaluation of the electron density

At each of the M grid points \mathbf{r}_i , the electron density is evaluated as

$$\rho(\mathbf{r}) = \sum_{i=1}^M w_i \sum_{p,q=1}^K D_{pq} \phi_p(\mathbf{r}_i) \phi_q(\mathbf{r}_i), \quad (78)$$

where w_i is the quadrature weight, D_{pq} an element of the atomic-orbital density matrix, $\phi_p(\mathbf{r}_i)$ a basis function, and K the total number of basis functions in the system. Formally, the number of operations scales as $O(N^3)$ since there are K^2 pairs of functions to be evaluated at M points. However, noting that the density at \mathbf{r}_i vanishes if $\phi_p(\mathbf{r}_i)$ or $\phi_q(\mathbf{r}_i)$ vanish, we can carry out the summation in the manner

$$\rho(\mathbf{r}) = \sum_{n=1}^L \sum_{p_n, q_n=1}^{K_n} \sum_{i=1}^{M_n} w_i D_{p_n q_n} \phi_{p_n}(\mathbf{r}_i) \phi_{q_n}(\mathbf{r}_i), \quad (79)$$

where, for large systems, the summation is divided into $L \sim O(N)$ terms, each of which requires a constant amount of work, whose prefactor depends on the number of nonvanishing basis functions and grid points in each box.

Note that, in Eq. (79), we have changed the summation order relative to Eq. (78) such that the inner rather than outer loop is now over grid points. Depending on the implementation, the perhaps more natural summation order with basis functions first would lead either to short inner loops or to indirect memory accesses when addressing the density matrix D_{pq} , both of which are inefficient on modern computers. The summation order in Eq. (79) turns the innermost loop into a simple $ax+y$ operation with a stride of one, which can be unrolled by compilers and which has a better memory locality. The length of this loop is chosen to be in order of a few hundreds. The loop is then sufficiently long to take advantage of vectorizing compilers and sufficiently short to fit all data into the fast cache memory of most modern computers.

3. Evaluation of the Kohn-Sham matrix

For the evaluation of the exchange-correlation contribution to the Kohn-Sham matrix, we use the same approach as that discussed in Sec. II E 2. For a local-density approximation (LDA) functional with an exchange-correlation energy density $F(\rho)$, the Kohn-Sham contribution is given by the following expression:

$$f_{pq} = \sum_{i=1}^M w_i \left. \frac{\partial F(\rho)}{\partial \rho} \right|_{\rho=\rho(\mathbf{r}_i)} \phi_p(\mathbf{r}_i) \phi_q(\mathbf{r}_i). \quad (80)$$

Again, since many basis functions vanish at a given grid point \mathbf{r}_i , this point will contribute only to a few matrix elements f_{pq} . To exploit this sparsity, the outer loop is over boxes $1 \leq n \leq L$, each of which has M_n grid points and is associated with K_n basis functions. Inside this loop, we compute contributions to the associated subset of matrix ele-

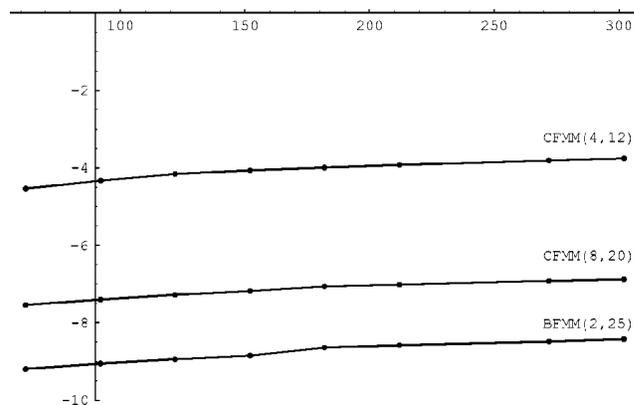


FIG. 1. Errors in the total molecular energy as computed using the CFMM and BFMM algorithms, with different orders of multipole expansion for a linear polyethylene chain at the Hartree/3-21G level. The logarithm of the absolute error in E_n is plotted against the number of atoms in the chain.

ments f_{pq} in a simple stride-one loop over the M_n grid points. The same $O(N)$ approach is used for GGA-type functionals.

4. Evaluation of the DFT linear-response contribution

DFT linear-response functions are computed in a manner analogous to the Hartree-Fock case. In particular, the linear-response eigenvalue solver uses a direct method, avoiding the explicit construction of the electronic Hessian matrix $\mathbf{E}^{[2]}$. The eigenvectors and eigenvalues are determined iteratively, carrying out a series of multiplications of the Hessian with trial vectors: $\tilde{\mathbf{X}} = \mathbf{E}^{[2]} \mathbf{X}$. In LDA theory, the exchange-correlation contribution to the transformed trial vector is given by the expression

$$\tilde{X}_{pq} = \int \left. \frac{\partial^2 F(\rho)}{\partial \rho^2} \right|_{\rho=\rho(\mathbf{r})} \times \left[\sum_{rs} X_{rs} \phi_r(\mathbf{r}) \phi_s(\mathbf{r}) \right] \phi_p(\mathbf{r}) \phi_q(\mathbf{r}) d\mathbf{r}. \quad (81)$$

In addition to the expectation value $\sum_{rs} X_{rs} \phi_r(\mathbf{r}) \phi_s(\mathbf{r})$, this expression involves the same elements as the exchange-correlation contribution to the Kohn-Sham matrix, Eq. (80), and may be evaluated in the same manner.

III. PERFORMANCE BENCHMARKS

In Figs. 1 and 2, we present illustrative benchmarks of our new code with respect to accuracy and scalability of the computational cost as applied to the set of polyethylene chains used in Sec. IV.

Figure 1 addresses the accuracy of both our CFMM and BFMM implementations with respect to the order of the multipole expansions. We study linear polyethylene chains containing up to 302 atoms using a 3-21G split-valence basis. We have calculated the total molecular energy in the Hartree approximation—that is, we have performed a single DFT iteration, but with the two-electron term containing no exchange or correlation contributions. We compare the total energy as computed using the CFMM or BFMM with the

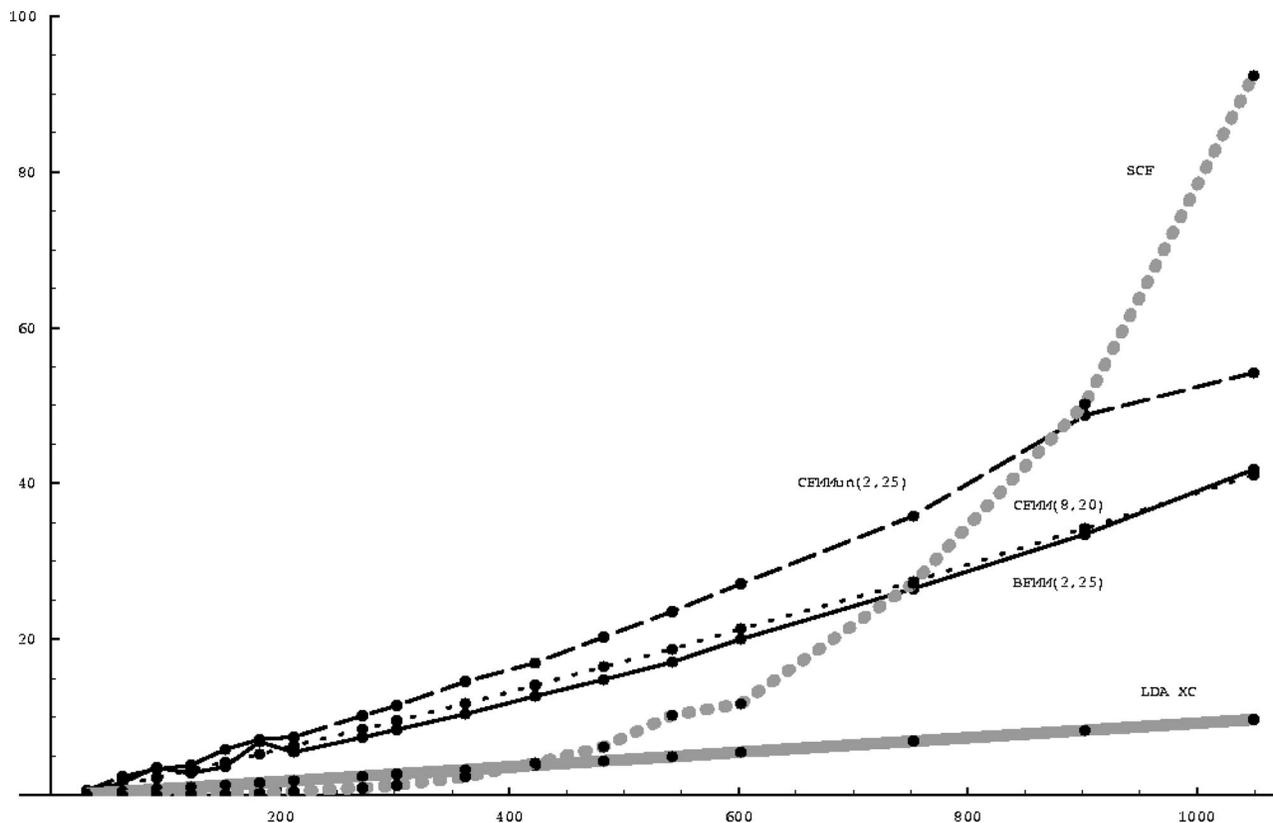


FIG. 2. Timing breakdown for a single Kohn-Sham iteration for a linear polyethylene chain at the LDA/3-21G level including the evaluation of the (classical+nonclassical) Coulomb contributions via the BFMM(2,25) (full black line), CFMM(8,20) (dotted black line), and CFMMun(2,25) (dashed black line) algorithms, the exchange-correlation contributions (full gray line), and the SCF optimization (dotted gray line); times in minutes are plotted against the number of atoms in the chain. In a given calculation, only one FMM algorithm is used.

total energy as computed using conventional integration where we may assume an accuracy that is at least $10^{-10}E_h$. Our choice of the extent definition also ensures an accuracy of at least $10^{-10}E_h$ in the multipole phase and we do not explore the variation of this parameter here.

We include three curves. The first two, denoted CFMM(L_1, L_2), relate to the CFMM algorithm to different orders of multipole expansion. Here L_1 is the order to which moments are generated from the overlap distributions and to which they are contracted in the nearest-neighbor phase, whereas L_2 is the order to which the moments are translated and subsequently contracted in the recursive far-field phase. These numbers are chosen such that an increase in L_2 does not improve the accuracy and L_1 is limiting. In this case, we observe an accuracy of approximately $10^{-4}E_h$ for the CFMM(4,12) evaluation and of approximately $10^{-7}E_h$ for the CFMM(8,20) evaluation, which we typically choose as the default setting in our DFT calculations.

In the third curve, we employ the same notation, but for the BFMM. In this case, we internally decontract the basis for the purpose of evaluating the multipole contribution only. (We do this to ensure numerical stability with small choices of P_{\min}). The moments are then generated about the charge center of overlap distributions of *primitive* Gaussian functions, such that the moments truncate exactly by symmetry. We therefore expect to achieve higher accuracy using this approach; indeed, without translations, we should recover the exact result to order 2. We therefore present data for $L_1=2$

and $L_2=25$. In this case, L_2 is limiting; we choose order 25 to achieve a better accuracy than we obtained above. Then, we observe an accuracy of better than $10^{-8}E_h$ for all sizes of molecule tested.

In Fig. 2, we give a summary of the component costs in a single DFT iteration at the LDA/3-21G level for the linear polyethylenes containing up to 1049 atoms. There are five curves: the three curves labeled CFMM(8,20), CFMMun(2,25), and BFMM(2,25) denote the evaluation of the Coulomb contributions (including the total classical and nonclassical evaluations) using different FMM algorithms; the curve labeled “LDA XC” denotes the evaluation of the exchange-correlation Kohn-Sham matrix elements, and the curve labeled “SCF” denotes the total time for the SCF iteration with the Coulomb and DFT integrations subtracted; the latter is dominated by the cost of diagonalization, performed using the standard LAPACK routine DSPEV(). For the DFT integration, we chose a radial integration accuracy of $10^{-9}E_h$ as the upper limit for the integration error for an atom. All calculations required less than 2 GB of memory.

Concerning the CFMM(8,20), CFMMun(2,25), and BFMM(2,25) curves, we note that, in the second and third curves, we have internally decontracted the basis for purposes of generating the multipole expansions. For a given box size, it is clear that this procedure will generate more moments to be contracted in the nearest-neighbor phase. Although this may be ameliorated in the far-field phase by grouping moments (and thus effectively “recontracting” the

basis), for given order, we observe a significant increase in the total work required. However, for a given accuracy, we would expect to need lower orders of multipole expansion, as observed in Fig. 1.

It is encouraging that our BFMM implementation in Fig. 2 performs better than our CFMM implementation—that is, for the same CPU time, we observe higher accuracy with the BFMM than when running the conventional CFMM scheme in a contracted basis. Moreover, the BFMM is faster than the CFMM when running with internal decontraction. We believe that this is an inherent improvement of the simpler BFMM algorithm, since our CFMM and BFMM codes share all of the time-critical routines involving contraction and translation of the moments.

From Fig. 2, it is clear that the Coulomb evaluation dominates the calculations for short chains. For chains with 800 atoms, however, the Coulomb bottleneck has been overcome as the $O(N^3)$ diagonalization cost starts to dominate the calculation, and near linear scaling is clearly observed in the DFT and Coulomb integrations. The DFT exchange-correlation integration is comparatively inexpensive, costing less than the diagonalization for systems containing more than 400 atoms.

IV. APPLICATIONS: EXCITATION ENERGIES AND POLARIZABILITIES OF LARGE HYDROCARBONS

As a simple illustration of the capabilities of our program, we have carried out a DFT study of the excitation energies and polarizabilities of the following extended hydrocarbon systems: polyethylenes C_nH_{2n+2} , linear polyenes C_nH_{n+2} , polyynes C_nH_2 , and D_{6h} graphite sheets $C_{6b}H_{6b}$, where b is the number of radial carbon atoms. The calculations were carried out using the LDA and Becke-Lee-Yang-Parr (BLYP)^{1,2} exchange-correlation functionals, in the STO-3G, 4-31G, and 6-31G basis sets. The 20 lowest singlet excitation energies and the 20 lowest triplet excitation energies were calculated using linear response theory, for a series of systems with $2 \leq n \leq 225$ (677 atoms) for polyethylenes, $2 \leq n \leq 300$ (602 atoms) for polyenes, $2 \leq n \leq 600$ (602 atoms) for polyynes, and $1 \leq b \leq 9$ (540 atoms) for graphites—see Table I for details about system size for the different basis sets. The static dipole polarizability was calculated for polyethylenes with up to 1049 atoms and for polyenes and polyynes with up to 1002 atoms—see Table II. For selected systems, we also calculated the frequency-dependent polarizability.

The geometries were generated in the following manner. The polyethylenes were generated using the default geometry of the MACROMODEL (version 5.5) builder, with no optimization. For the polyenes, we carried out DFT Becke-3-parameter-Lee-Yang-Parr (B3LYP)⁵ geometry optimizations using the correlation-consistent valence-double-zeta (cc-pVDZ) basis set⁶⁵ for C_2H_4 , C_4H_6 , and for C_6H_8 . Longer polyenes were then generated from C_6H_8 by duplicating the middle section. The polyynes were generated similarly. Thus, all polyynes with ten or less carbon atoms were optimized at the B3LYP/cc-pVDZ level; the longer polyynes were generated from $C_{10}H_2$ by duplicating the middle single-triple-single bond section. Finally, the D_{6h} graphite sheets were

generated with a fixed CC bond length of 142 pm, with no optimization. Hydrogens were added to the edge carbon atoms, using a bond length of 112 pm.

A. Excitation energies

In Fig. 3, we have plotted the lowest singlet (dotted lines) and triplet (full lines) LDA excitation energies of the polyethylenes, polyenes, polyynes, and graphite sheets against n^{-1} , where n is the number of carbon atoms. Considering the chains first, we note that their excitation energies are approximately linear in n^{-1} (in particular, the lowest triplet excitation energy) but converge to a finite value for large n . There is a striking difference, however, between the behavior of the saturated chains on the one hand and that of the unsaturated chains on the other hand. Thus, at the LDA/4-31G level of theory, the polyethylene singlet excitation energy decreases from $369 mE_h$ for $n=2$ to $255 mE_h$ for $n=\infty$, while the first triplet excitation energy decreases from 352 to $255 mE_h$ —that is, the polyethylene excitation energies are reduced by about 30% going from ethane to an infinite chain. By contrast, in the polyenes and polyynes, the lowest excitation energies are reduced by as much as 80% or 90%, depending on method and basis set. The LDA/4-31G polyene triplet excitation energy, for example, is reduced from 178 to $30 mE_h$, while the corresponding polyene triplet excitation energy is reduced from 208 to $37 mE_h$.

In Table I, we have listed the lowest singlet and triplet excitation energies of the linear hydrocarbons, for different basis sets and exchange-correlation functionals. Although the general behavior of the excitation energies is much the same at all levels of theory, we note that the excitation energies are much greater in the small STO-3G basis set than in the two larger basis sets. In general, the excitation energies are slightly reduced as we go from 4-31G to 6-31G, although we note that the polyene and polyene triplet excitation energies are either unaffected or very slightly increased by this change.

Since the excitation energies are approximately proportional to n^{-1} , each increment in n in Table I reduces the difference between the calculated excitation energy of the finite chain and the extrapolated excitation energy of the corresponding infinite chain by about one-third. The fastest convergence is observed for the polyethylene excitations; the slowest convergence for the polyene singlet excitations. To recover 99% of the change in the lowest excitation energy from $n=2$ to $n=\infty$, we must consider polyethylenes with about 60 carbon atoms for both singlet and triplet excitations; polyenes with about 110 and 80 carbon atoms for singlets and triplets, respectively; and polyynes with about 90 and 80 carbon atoms, respectively, for singlets and triplets.

As seen from Fig. 3, the higher excitation energies are affected more strongly than the lowest excitation energy by changes in the chain length—in particular, for the unsaturated chains, where the highest calculated excitation energy is reduced by a factor of 20 as n tends to infinity. This behavior is to be expected, of course, since, for an infinite chain, the lowest excited states form a continuous band of states. To illustrate the formation of bands, we have, in Fig. 4, plotted

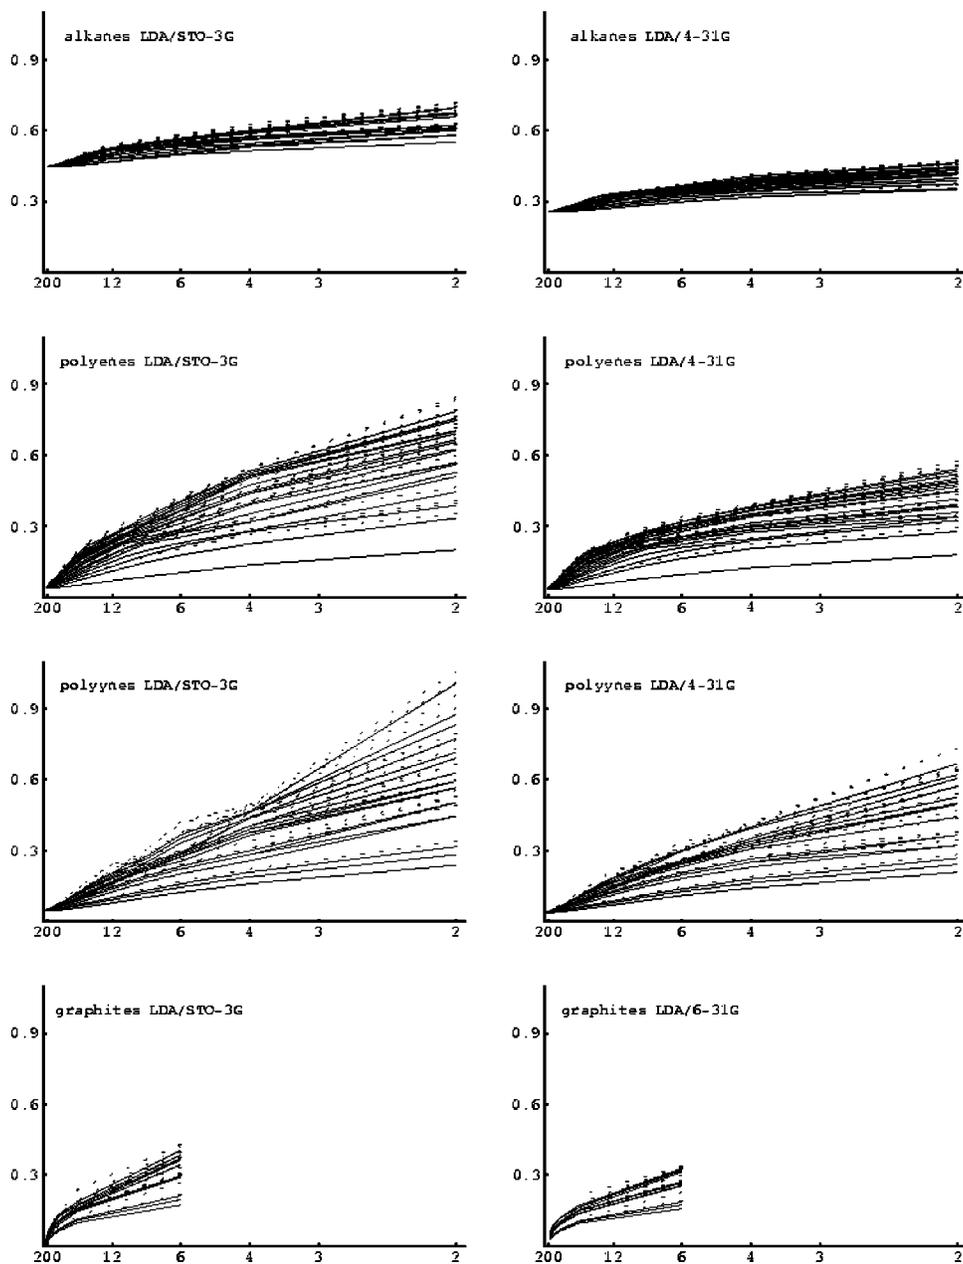


FIG. 3. The 20 lowest singlet (dotted lines) and 20 lowest triplet (full lines) excitation LDA energies (E_h) of polyethylenes, linear polyenes, polyynes, and graphite sheets plotted against the inverse number of carbon atoms.

the density of the calculated excited states as a function of the number of carbon atoms.

Whereas the density of states increases quadratically with n for the polyethylene chains, it increases less dramatically for the polyenes and polyene chains. Somewhat surprisingly, the increase in the density of states seems to flatten out for long polyynes. This behavior arises since, for large n , two distinct bands of excitation energies are formed for the polyynes at the LDA/STO-3G and LDA/4-31G levels of theory. The lowest band consists of only three or four triplet levels; the upper band contains 17 triplet levels and 20 singlet levels (some of which are degenerate). However, without carrying out more extended calculations, we cannot claim that this is the true behavior of polyynes rather than an artifact of our calculations, although we note that a similar gap is being formed also at the BLYP level of theory. By contrast, for the polyethylenes and polyenes, there is no clear gap separating the lowest triplet states from the higher ex-

cited states; rather, the calculations indicate the formation of a single triplet band that overlaps with the (slightly higher) singlet band. Without entering into further speculations, we note that we cannot guarantee that, in our calculations, we have recovered all excited states within a given energy range. Since our calculations were carried out without the explicit use of point-group symmetry, this is especially true for the polyynes, which exhibit a higher symmetry than do the polyenes and, in particular, the polyethylenes.

In general, the BLYP functional performs in much the same manner as the LDA functional. From Table I, we see that the lowest BLYP/6-31G singlet excitation energies are nearly all equal to or slightly higher (by 1 or 2 mE_h) than the corresponding LDA/6-31G singlet excitation energies, the only exception occurring for C_6H_8 . The same is true for the lowest polyethylene triplet excitation energies. By contrast, for the triplet states of the unsaturated chains, the BLYP/6-31G excitation energies are slightly below the corresponding

TABLE I. The lowest singlet and triplet excitation energies of polyethylenes C_nH_{2n+2} , polyenes C_nH_{n+2} , and polyynes C_nH_2 (mE_h). The excitation energies have been obtained by interpolation for $n \leq n_{\max}$ and by extrapolation for $n > n_{\max}$, where n_{\max} is the largest number of carbon atoms in the DFT calculations.

		n_{\max}	n (singlet)							n (triplet)						
			2	6	20	60	200	600	∞	2	6	20	60	200	600	∞
Polyethylenes	LDA STO-3G	200	608	507	459	447	446	446	446	550	497	454	446	445	445	445
	4-31G	200	369	302	264	256	255	255	255	352	297	262	256	255	255	255
	6-31G	225	367	297	258	250	249	249	249	349	292	256	250	249	249	249
	BLYP 6-31G	225	369	298	259	252	251	251	251	348	292	257	251	251	250	250
Polyenes	LDA STO-3G	300	352	205	91	50	41	40	40	199	102	53	40	37	37	37
	4-31G	300	295	175	81	43	33	32	32	178	94	47	34	31	30	30
	6-31G	200	292	174	80	43	33	32	32	177	94	48	34	31	30	30
	BLYP 6-31G	250	301	173	81	43	34	33	32	162	85	43	32	29	29	29
Polyynes	LDA STO-3G	600	317	160	76	54	49	49	48	239	123	62	47	45	45	45
	4-31G	400	267	138	65	45	40	40	40	208	109	54	40	38	37	37
	6-31G	300	264	137	65	45	41	40	40	207	109	54	40	38	38	38
	BLYP 6-31G	300	264	138	66	46	41	41	41	194	102	51	38	36	36	36

LDA/6-31G excitation energies. As a result, the BLYP functional increases the triplet-singlet separation in the extended systems—see Fig. 5, where we compare the LDA and BLYP singlet and triplet excitation energies in polyynes as functions of n .

Turning our attention to the graphite excitation energies in Fig. 3, we note that these behave very differently from those of the chains in that they do not seem to converge to a positive value for large n ; the triplet excitations may even converge to a negative value, suggesting the existence of instabilities for large n . This observation is in agreement with the fact that, for an infinite graphite sheet, there is no energy gap between the occupied and virtual orbitals. We also note that, unlike for the linear chains, the excitation energies no longer appear to be linear in n^{-1} . Indeed, a closer examination of the excitation energies in graphite suggests that these may be proportional to $n^{-1/2}$ (i.e., to the inverse radius of the D_{6h} graphite systems). We should keep in mind, however, that the true behavior of the excitation energies in systems

such as those studied here probably cannot be accurately described by simple expressions such as n^{-1} and $n^{-1/2}$.

B. Polarizabilities

In Fig. 6, we have plotted the static polarizability α of the hydrocarbons as a function of the number of carbon atoms n , as well as the polarizability per carbon atom α/n (the group polarizability). Again, there is a clear difference between the behavior of the saturated chains on the one hand and that of the unsaturated ones on the other hand. Whereas the polyethylene polarizability depends linearly on the chain length even for very small systems, the polyene and polyyne polarizability initially depends nonlinearly (quadratically) on the chain length, before the onset of linearity with a constant group polarizability for systems containing several hundred atoms. Moreover, the polarizabilities of the polyenes and polyynes are about an order of magnitude larger than those of the polyethylenes. For the graphite systems, no linearity is

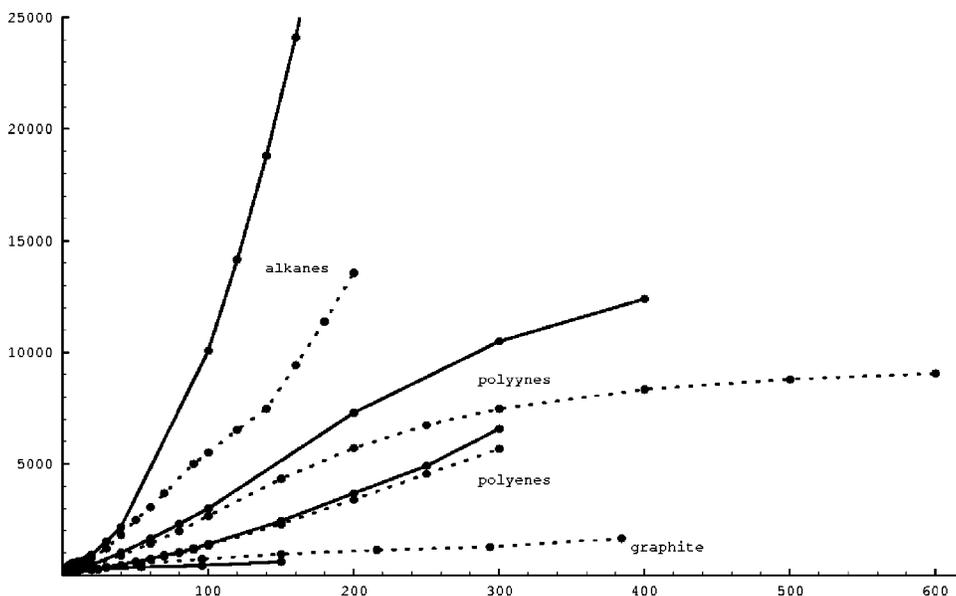


FIG. 4. The density of the 40 lowest excited states (20 singlet and 20 triplet) at the LDA/STO-3G (dotted lines) and LDA/4-31G (full lines) levels of theory of polyethylenes, linear polyenes, polyynes, and graphite sheets as functions of the number of carbon atoms.

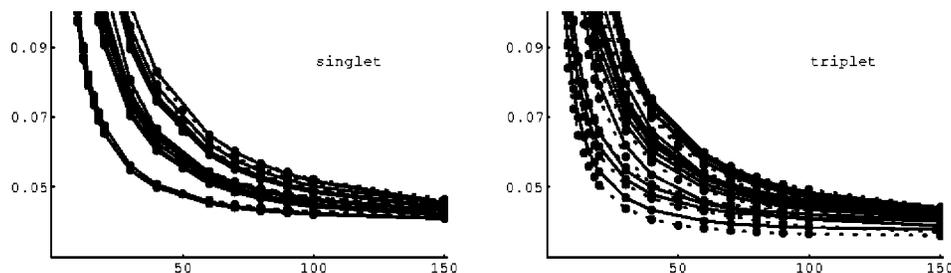


FIG. 5. The calculated LDA/6-31G (full lines) and BLYP/6-31G (dotted lines) excitation energies of linear polyynes as a function of the number of carbon atoms (E_h).

observed, even for the largest systems. Instead, the polarizability is proportional to $n^{3/2}$, as seen from the very constant value of $\alpha/n^{3/2}$ plotted in Fig. 6. It is also noteworthy that, in the graphite systems, the polarizability is about an order of magnitude smaller than in polyynes with the same number of carbon atoms. As expected, the STO-3G basis is inadequate, yielding polarizabilities that are much too small.

In Table II, we have listed the polarizability per carbon atom for different basis sets and functionals. Again, we note the contrast between the very nearly constant group polarizability of the polyethylenes and the highly variable group polarizability of the unsaturated systems, which increases by more than an order of magnitude as we extend the chains. Indeed, for the unsaturated chains, convergence has not yet

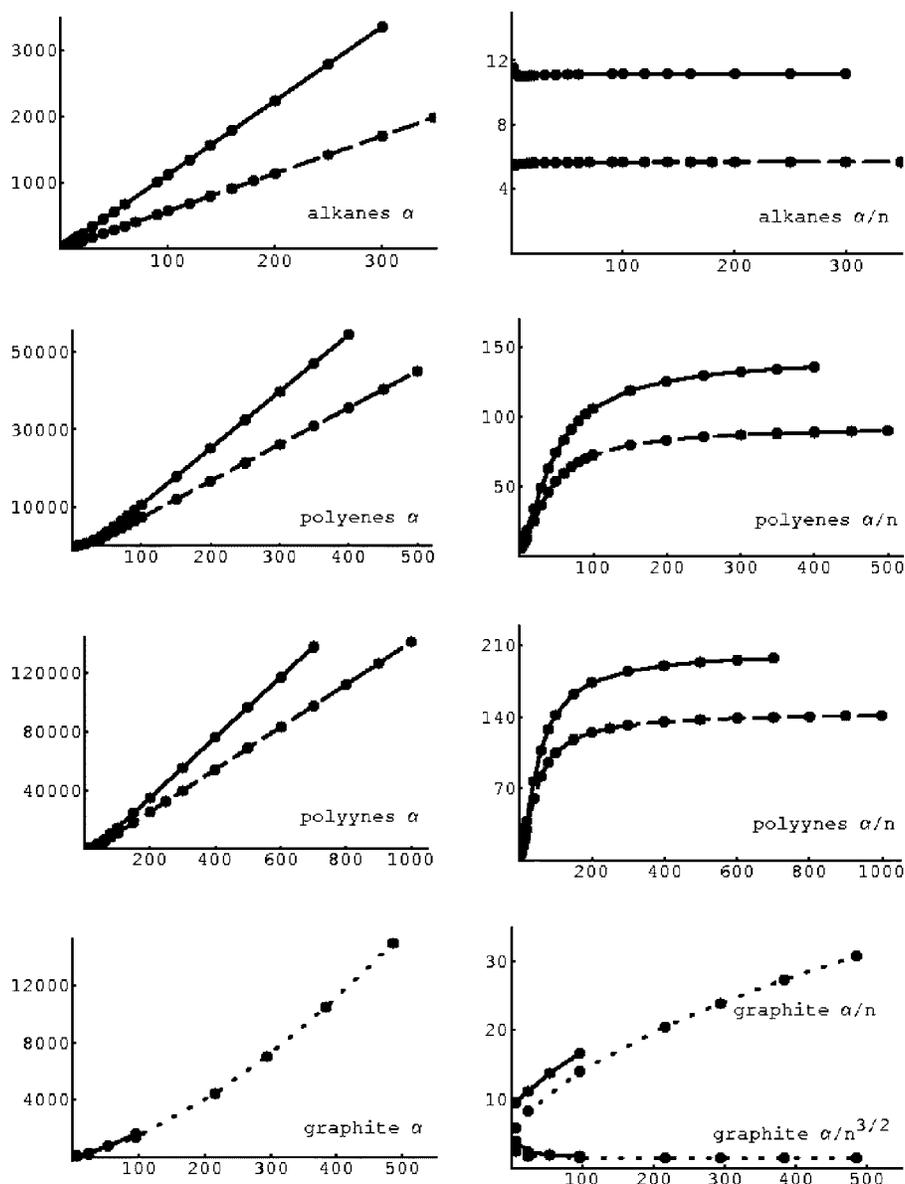


FIG. 6. The LDA static polarizabilities α in the STO-3G basis (dotted line) and the 4-31G basis (6-31G for graphite) (full line), as functions of the number of carbon atoms n .

TABLE II. The static group polarizability (polarizability per carbon atom) of polyethylenes C_nH_{2n+2} , polyenes C_nH_{n+2} , and polyynes C_nH_2 (atomic units). The group polarizabilities have been obtained by interpolation for $n \leq n_{\max}$ and by extrapolation for $n > n_{\max}$, where n_{\max} is the largest number of carbon atoms in the DFT calculations.

			n_{\max}	n						
				2	6	20	60	200	600	∞
Polyethylenes	LDA	STO-3G	349	5.5	5.5	5.6	5.7	5.7	5.7	5.7
		4-31G	300	11.6	11.0	11.1	11.1	11.2	11.2	11.2
	BLYP	6-31G	200	11.9	11.3	11.3	11.4	11.4	11.4	11.4
		6-31G	200	11.8	11.1	11.1	11.2	11.2	11.2	11.2
Polyenes	LDA	STO-3G	500	5.1	8.2	24.9	59.3	83.2	90.6	94.1
		4-31G	400	9.6	12.9	33.5	83.3	125.4	139.1	146.0
	BLYP	6-31G	250	9.9	13.1	33.9	84.0	126.4	140.1	147.0
		6-31G	300	9.8	13.1	33.8	81.3	122.1	135.0	141.5
Polyynes	LDA	STO-3G	1000	3.6	7.5	29.2	81.7	124.8	138.6	145.6
		4-31G	700	6.7	11.1	37.6	107.2	173.5	195.7	207.0
	BLYP	6-31G	350	7.0	11.4	38.0	107.5	173.4	195.5	206.8
		6-31G	350	7.0	11.4	37.7	105.0	167.3	188.1	198.8

been achieved even for systems containing as many as 600 carbon atoms. As expected, there is hardly any difference between the static polarizabilities calculated at the LDA and GGA levels of theory, the BLYP functional producing slightly smaller group polarizabilities than does the LDA functional. For graphite systems, the LDA/STO-3G value of $\alpha/n^{3/2}$ is 1.67 a.u. for C_6H_6 , 1.39 a.u. for $C_{24}H_{12}$, and 1.40 a.u. for all larger systems; in the 6-31G basis, the values are 2.27 a.u. for C_6H_6 , 1.56 a.u. for $C_{24}H_{12}$, 1.46 a.u. for $C_{54}H_{18}$, and 1.45 a.u. for all larger systems.

It is possible to rationalize the dependence of the polarizability on the number of carbon atoms in a qualitative manner, with reference to the excitation energies. Thus, from general considerations, we would expect the polarizability in large systems to depend inversely on the excitation energies but linearly on the number of atoms. To illustrate, we have, to the right in Fig. 7, plotted the inverse excitation energies as functions of the number of carbon atoms n ; to the left in the same figure, we have made similar plots of the inverse excitation energies multiplied by n . Individually, each plot in Fig. 7 bears a striking resemblance to the corresponding plot in Fig. 6, demonstrating the validity of this simple model.

We may now understand the linearity of the polarizability in the extended linear systems as a reflection of the converged, constant excitation energies; by the same token, the nonlinearity of the polarizability for small chains arises from the large reduction in the excitation energies with each added carbon unit for such systems. The polarizability of the polyethylene chains is much smaller than that of the unsaturated chains because of the much larger excitation energies in the polyethylenes. Since the polyethylene excitation energies are only reduced by 30% as we extend the chain, the polyethylene group polarizability is independent of n . By contrast, the excitation energies of the unsaturated hydrocarbons change by almost 90% as we extend the chains, resulting in a much more pronounced nonlinear behavior of these systems.

Since the inverse graphite excitation energies in Fig. 7 do not converge to a finite value, the graphite polarizability in Fig. 6 never becomes proportional to the size of the sys-

tem. Instead, the $n^{3/2}$ dependence of the polarizability in graphite may be seen as an indication that its inverse excitation energy is proportional to $n^{1/2}$. Although this explanation fits very well with the $n^{-1/2}$ dependence of the excitation energies in graphite as proposed in Sec. IV A, it should be treated with some scepticism since, as seen from Fig. 7, the calculated inverse excitation energies in graphite do not follow a $n^{1/2}$ curve accurately.

In Fig. 8, we have plotted the polarizabilities of the polyethylenes, polyenes, and polyynes at two applied frequencies, superimposed on the corresponding static polarizabilities. At the chosen frequencies of $23.88 mE_h$ (5241 cm^{-1}) and $42.82 mE_h$ (9398 cm^{-1}), the dispersion of the polyethylenes is hardly noticeable on the scale of the plot, whereas that of the unsaturated hydrocarbons is large because of their smaller excitation energies. As expected, the calculated dispersion increases with increasing n (i.e., with decreasing excitation energy), in much the same manner as it increases when we increase the applied frequency for a single molecule. The dispersion is largest for the polyenes, which have the smallest excitation energies. Indeed, at the highest applied frequency of about $43 mE_h$, the dispersion goes through a maximum at $n=40$ and becomes negative at $n=50$, indicating that the applied frequency exceeds the lowest singlet excitation energy of about $40 mE_h$ for these systems.

V. CONCLUSION

We have presented a new linear-scaling implementation in the program system DALTON for the formation of the Kohn-Sham Hamiltonian at the GGA level of density-functional theory. Our approach introduces a generalization of the FMM to the treatment of Gaussian charge distributions for the evaluation of the troublesome Coulomb matrix elements; it affords a simpler algorithm than the original CFMM. In particular, we have established explicit expressions for the decomposition of the Coulomb energy into a "classical" point-charge multipole expansion contribution

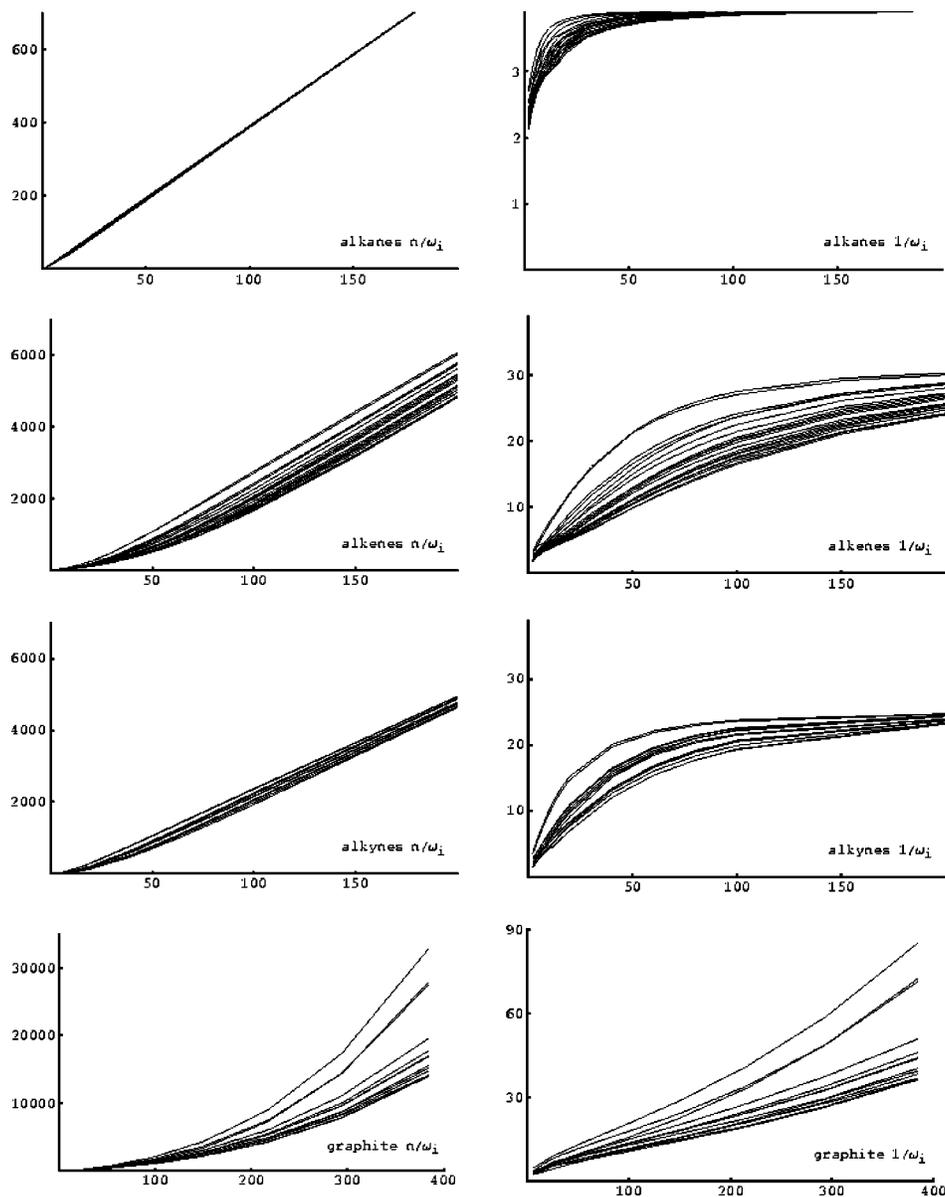


FIG. 7. The inverse excitation energies ω_i^{-1} plotted against the number of carbon atoms n (on the right) and $n\omega_i^{-1}$ plotted against n (on the left) at the LDA/STO-3G level of theory.

and a “nonclassical” size-correction term which may be easily evaluated using a modified two-electron integral package. We have presented benchmark calculations that illustrate for polyethylene chains that our algorithm is at least as efficient as the CFMM for high-accuracy work.

Our approach also removes the major traditional bottlenecks to the calculation of second-order response properties in large molecular systems. To demonstrate the capabilities

of our code, we have considered the singlet and triplet excitation energies as well as the static and dynamic polarizabilities of polyethylenes, polyenes, polyynes, and graphite sheets with an emphasis on the trends observed with system size. For the linear hydrocarbons, the excitation energies were found to decay approximately linearly with $1/n$, where n is the number of carbon atoms, and to converge to a finite value in the limit $n \sim \infty$. In contrast, the excitations in the

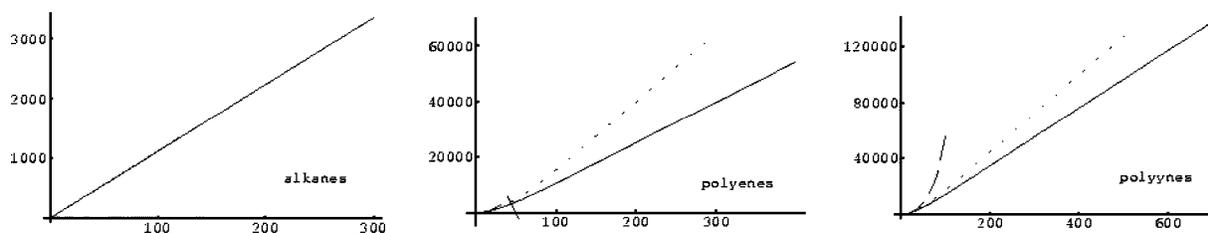


FIG. 8. The polarizabilities of polyethylenes, polyenes, and polyynes plotted against the number of carbon atoms n . The polarizabilities have been calculated at the LDA/4-31G level of theory at zero frequency (full lines), at $23.88 mE_h$ (dotted lines), and at $42.82 mE_h$ (dashed lines).

graphite sheets appeared to decay slower than $1/n$ and more like $n^{-1/2}$; they did not converge to a finite value, consistent with the metallic character of the system. For the linear hydrocarbons, the polarizability per carbon atom converged to a constant value, whereas in the graphite sheets, the polarizability increased like $n^{3/2}$. For both the excitation energies and polarizabilities, the BLYP functional performed in much the same manner as the LDA functional. Moreover, little difference was observed between the use of the 4-31G and 6-31G basis sets in these cases. Finally, in assessing the results, we should keep in mind the documented problems of LDA and GGA functionals in predicting polarizabilities and excitation energies of extended conjugated systems.^{50–53} For quantitative studies of these systems, therefore, more general (nonlocal) functionals are needed, including either a dependence on current densities or some proportion of exact exchange.

In summary, we believe we have an efficient and robust code for linear-scaling generalized-gradient approximation density-functional calculations of electromagnetic properties. In a parallel paper, we have published results using the same code where we study nuclear spin-spin coupling constants in biomolecules.⁶⁶ We are most excited by our ability to use DALTON to compute molecular properties of large systems.

ACKNOWLEDGMENTS

This work was supported by the EU training and research network MOLPROP (“Molecular Properties and Molecular Materials,” Contract No. HPRN-CT-2000-00013), by the Norwegian Research Council through a Strategic University Program in Quantum Chemistry (Grant No. 154011/420), and by the Norwegian Program for Supercomputing (nn1118k).

- ¹A. D. Becke, *Phys. Rev. A* **38**, 3098 (1988).
- ²C. Lee, W. Yang, and R. G. Parr, *Phys. Rev. B* **37**, 785 (1988).
- ³B. Johnson, P. M. W. Gill, and J. A. Pople, *J. Chem. Phys.* **98**, 5612 (1993).
- ⁴N. C. Handy and A. J. Cohen, *Mol. Phys.* **99**, 403 (2001).
- ⁵P. J. Stephens, F. Devlin, C. F. Chabalowski, and M. J. Frisch, *J. Phys. Chem.* **98**, 11623 (1994).
- ⁶A. D. Becke, *J. Phys. Chem.* **107**, 8554 (1997).
- ⁷F. A. Hamprecht, A. J. Cohen, D. J. Tozer, and N. C. Handy, *J. Chem. Phys.* **109**, 6264 (1998).
- ⁸W. Yang and J. M. Pérez-Jordá, in *Encyclopedia of Computational Chemistry*, edited by P. V. R. Schleyer, N. L. Allinger, T. Clark, J. Gasteiger, P. A. Kollman, H. F. Schaefer III, and P. R. Schreiner (Wiley, New York, 1998), p. 1496.
- ⁹G. E. Scuseria, *J. Phys. Chem. A* **103**, 4782 (1999).
- ¹⁰C. Möller and M. S. Plesset, *Phys. Rev.* **46**, 618 (1934).
- ¹¹P. Y. Ayala and G. E. Scuseria, *J. Chem. Phys.* **110**, 3660 (1999).
- ¹²M. Shütz, G. Hetzer, and H. J. Werner, *J. Chem. Phys.* **111**, 5691 (1999).
- ¹³C. Hampel and H. J. Werner, *J. Chem. Phys.* **104**, 6286 (1996).
- ¹⁴M. Shütz, *J. Chem. Phys.* **113**, 9986 (2000).
- ¹⁵T. Helgaker, H. J. A. Jensen, P. Jørgensen *et al.*, DALTON, an *ab initio* electronic structure program, Release 1.2. See <http://www.kjemi.uio.no/software/dalton/dalton.html> (2001).
- ¹⁶J. Almlöf, K. F. Faegri, Jr., and K. Korsell, *J. Comput. Chem.* **3**, 385 (1982).
- ¹⁷H. Horn, H. Weiß, M. Häser, M. Ehrig, and R. Ahlrichs, *J. Comput. Chem.* **12**, 1058 (1991).
- ¹⁸M. Häser and R. Ahlrichs, *J. Comput. Chem.* **10**, 104 (1989).
- ¹⁹A. D. Daniels, J. M. Millam, and G. E. Scuseria, *J. Chem. Phys.* **107**, 425 (1997).
- ²⁰J. M. Millam and G. E. Scuseria, *J. Chem. Phys.* **106**, 5569 (1997).
- ²¹A. D. Daniels and G. E. Scuseria, *J. Chem. Phys.* **110**, 1321 (1999).
- ²²M. Challacombe, *J. Chem. Phys.* **110**, 2332 (1999).
- ²³M. Challacombe, *Comput. Phys. Commun.* **128**, 93 (2000).
- ²⁴H. Larsen, J. Olsen, P. Jørgensen, and T. Helgaker, *J. Chem. Phys.* **115**, 9685 (2001).
- ²⁵A. W. Appel, *SIAM (Soc. Ind. Appl. Math.) J. Sci. Stat. Comput.* **6**, 85 (1985).
- ²⁶V. Rokhlin, *J. Comput. Phys.* **60**, 187 (1985).
- ²⁷J. Barnes and P. Hut, *Nature (London)* **324**, 446 (1986).
- ²⁸L. Greengard and V. Rokhlin, *J. Comput. Phys.* **73**, 325 (1987).
- ²⁹L. Greengard, *The Rapid Evaluation of Potential Fields in Particle Systems* (MIT, Cambridge, MA, 1988).
- ³⁰C. A. White and M. Head-Gordon, *J. Chem. Phys.* **101**, 6593 (1994).
- ³¹H. Q. Ding, N. Karasawa, and W. A. Goddard, *Chem. Phys. Lett.* **196**, 6 (1992).
- ³²J. M. Pérez-Jordá and W. Yang, *Chem. Phys. Lett.* **241**, 469 (1995).
- ³³L. Greengard, *Science* **265**, 909 (1994).
- ³⁴C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, *Chem. Phys. Lett.* **230**, 8 (1994).
- ³⁵C. A. White, B. G. Johnson, P. M. W. Gill, and M. Head-Gordon, *Chem. Phys. Lett.* **253**, 268 (1996).
- ³⁶M. Challacombe, E. Schwegler, and J. Almlöf, *J. Chem. Phys.* **104**, 4685 (1996).
- ³⁷M. C. Strain, G. E. Scuseria, and M. J. Frisch, *Science* **271**, 51 (1996).
- ³⁸J. M. Pérez-Jordá and W. Yang, *J. Chem. Phys.* **107**, 1218 (1997).
- ³⁹S. Aluru, *SIAM J. Sci. Comput. (USA)* **177**, 773 (1996).
- ⁴⁰R. E. Stratmann, G. E. Scuseria, and M. J. Frisch, *Chem. Phys. Lett.* **257**, 213 (1996).
- ⁴¹J. M. Pérez-Jordá and W. Yang, *Chem. Phys. Lett.* **247**, 484 (1995).
- ⁴²O. Treutler and R. Ahlrichs, *J. Chem. Phys.* **102**, 346 (1995).
- ⁴³K. Eichkorn, F. Weigend, O. Treutler, and R. Ahlrichs, *Theor. Chem. Acc.* **97**, 119 (1997).
- ⁴⁴M. Challacombe, *J. Chem. Phys.* **113**, 10037 (2000).
- ⁴⁵P. Jørgensen, H. J. A. Jensen, and J. Olsen, *J. Chem. Phys.* **89**, 3654 (1988).
- ⁴⁶C. Ochsenfeld and M. Head-Gordon, *Chem. Phys. Lett.* **270**, 399 (1997).
- ⁴⁷R. E. Stratmann, G. E. Scuseria and M. J. Frisch, *J. Chem. Phys.* **109**, 8218 (1998).
- ⁴⁸R. Bauernschmitt and R. Ahlrichs, *Chem. Phys. Lett.* **256**, 454 (1996).
- ⁴⁹M. E. Casida, K. C. Casida, and D. R. Salahub, *Int. J. Quantum Chem.* **70**, 933 (1998).
- ⁵⁰B. Champagne, E. A. Perète, S. J. A. van Gisbergen, E.-J. Baerends, J. G. Snijders, C. Soubra-Ghaoul, K. A. Robins, and B. Kirtman, *J. Chem. Phys.* **109**, 10489 (1998).
- ⁵¹M. van Faassen, P. L. de Boeij, R. van Leeuwen, J. A. Berger, and J. G. Snijders, *Phys. Rev. Lett.* **88**, 186401 (2002).
- ⁵²M. van Faassen, P. L. de Boeij, R. van Leeuwen, J. A. Berger, and J. G. Snijders, *J. Chem. Phys.* **118**, 1044 (2003).
- ⁵³Z.-L. Cai, K. Sendt, and J. R. Reimers, *J. Chem. Phys.* **117**, 5543 (2002).
- ⁵⁴M. Sierka, A. Hogekamp, and R. Ahlrichs, *J. Chem. Phys.* **118**, 9136 (2003).
- ⁵⁵K. N. Kudin and G. E. Scuseria, *J. Chem. Phys.* **111**, 2351 (1999).
- ⁵⁶L. E. McMurchie and E. R. Davidson, *J. Comput. Phys.* **26**, 218 (1978).
- ⁵⁷G. R. Ahmadi and J. Almlöf, *Chem. Phys. Lett.* **246**, 364 (1995).
- ⁵⁸C. A. White and M. Head-Gordon, *J. Chem. Phys.* **104**, 2630 (1996).
- ⁵⁹W. Klopper, *J. Mol. Struct.* **388**, 175 (1996).
- ⁶⁰Y. Shao and M. Head-Gordon, *Chem. Phys. Lett.* **323**, 425 (2000).
- ⁶¹T. Helgaker, P. Jørgensen, and J. Olsen, *Molecular Electronic-Structure Theory* (Wiley, New York, 2000).
- ⁶²M. Häser and R. Ahlrichs, *J. Comput. Chem.* **10**, 104 (1989).
- ⁶³J. M. Pérez-Jordá and W. Yang, *J. Chem. Phys.* **104**, 8003 (1996).
- ⁶⁴A. D. Becke, *J. Chem. Phys.* **88**, 2547 (1988).
- ⁶⁵T. H. Donning, Jr. *J. Chem. Phys.* **90**, 1007 (1989).
- ⁶⁶M. A. Watson, P. Safek, P. Macak, M. Jaszuński, and T. Helgaker, *Chem. Eur. J.* (in press).