

Multiresolution quantum chemistry: Basic theory and initial applications

Robert J. Harrison, George I. Fann, Takeshi Yanai, and Zhengting Gan
Oak Ridge National Laboratory, Oak Ridge, Tennessee 37831

Gregory Beylkin
*Department of Applied Mathematics, University of Colorado at Boulder, Boulder,
Colorado 80309-0526*

(Received 3 March 2004; accepted 19 July 2004)

We describe a multiresolution solver for the all-electron local density approximation Kohn-Sham equations for general polyatomic molecules. The resulting solutions are obtained to a user-specified precision and the computational cost of applying all operators scales linearly with the number of parameters. The construction and use of separated forms for operators (here, the Green's functions for the Poisson and bound-state Helmholtz equations) enable practical computation in three and higher dimensions. Initial applications include the alkali-earth atoms down to strontium and the water and benzene molecules. © 2004 American Institute of Physics. [DOI: 10.1063/1.1791051]

I. INTRODUCTION

We describe a multiresolution solver for the all-electron local density approximation (LDA) Kohn-Sham¹ equations for molecules. The objective of this work is to provide a practical approach for computation on general polyatomic systems without basis set error and with the computational cost of applying all operators scaling linearly with the number of parameters. Besides using a multiresolution approach, a critical step in attaining this objective has been the development of separable representations for kernels of Green's functions.² We describe a prototype, orbital-based implementation with test application to closed-shell systems including the alkali-earth atoms down to strontium, and the water and benzene molecules. We consider that a demonstration of a practical approach for solution of one-electron methods is an essential precursor to direct numerical solution of two- and many-electron problems.

With only a few notable exceptions,^{3,4} mainstream molecular quantum chemistry is performed with the linear combination of atomic orbitals (LCAO) approximation most commonly using atom-centered Gaussian functions. These methods date back 50 years to the work of Roothaan,^{5,6} Hall,⁷ and Boys,⁸ and many of the successes of modern quantum chemistry can be attributed to this framework. By capturing the essence of molecules as interacting atoms, the LCAO approximation provides a very compact representation of molecular orbitals and yields analytically smooth potential energy surfaces. Judicious use of carefully designed atomic orbital (AO) basis sets can yield accurate energy differences (e.g., binding energies) despite large errors in the total energy due to cancellation of the intraatomic error. The use of a Gaussian basis enables efficient computation of the necessary two-electron integrals^{8,9} and makes feasible computations for general polyatomic molecules. In addition, atomic orbitals formed from contractions of primitive Gaussians effectively eliminate the high frequencies in the vicinity of the nucleus. The resulting matrix representation of the kinetic energy operator is better conditioned than that in the

primitive basis. The second-quantized form of quantum theory has been especially successful in formulating post-Hartree-Fock *ab initio* methods. It is developed by projection of the differential equations into a fixed, finite basis, yielding a purely algebraic problem.

However, the LCAO approximation gives rise to severe problems for the accurate treatment of large systems. These problems include the expense of computation with high precision, basis set superposition error,¹⁰ linear dependence problems in high quality calculations, and the fact that AO basis has still to be designed manually for many problems. The second-quantized formulation contributes indirectly to these problems by obscuring the physical interpretation of expressions involving two-electron integrals.

Intuitively, it is clear that a multiresolution approach can address many of these concerns. A systematic approach to multiresolution constructions started with the development of wavelet bases; see Ref. 11 and references therein. For numerical applications, the results in Ref. 12 pointed out a practical approach to reducing the computational cost. One of the results of Ref. 12 was the introduction of the non-standard form (NS form) for representing operators in multiresolution bases. However, the straightforward generalization of the NS form (or for that matter, the standard form) to multiple dimensions is too expensive for practical applications. Our approach is based on using NS form and separated representations of operators which was first used in Ref. 13 and significantly extended in Ref. 2. The basic point of Ref. 2 is that many apparently nonseparable operators are, in fact, separable with a finite but arbitrary precision. Moreover, the number of terms necessary for such representations is remarkably small.

In this paper, we construct and use separable representations of Green's function for the Poisson and bound-state Helmholtz equations. These constructions, combined with multiresolution representations, make our approach practical.

In our approach, we chose to use multiwavelet bases. This selection has been motivated by a number of contradictory requirements for the basis (see Ref. 14). In particular,

we require orthonormality, the interpolating property, and the ability to accommodate boundary conditions while maintaining both accuracy and the order of convergence. It turns out that there are no smooth bases that satisfy all of these conditions. Unexpected positive consequences of using multiwavelets with disjoint supports include a family of derivative operators with analogs of forward and backward differences and a useful connection to the so-called discontinuous finite (or spectral) elements methods.

We note that the Green's functions for the Poisson and bound-state Helmholtz equations can be applied efficiently using the fast multipole method (FMM) (see Ref. 15 and references therein) or using multigrid (see Refs. 16–18). Although there are similarities between the key representations used in the FMM and our separated representations, there are also significant differences. In particular, the separated representations can be used in higher dimensions (for example, in computing six-dimensional integrals) and, after making the approximation, we are still left with analytic expressions for further use. We have not yet made any comparison of speed between these methods.

Previous applications of wavelets and multiresolution analysis to quantum chemistry^{19–28} have almost exclusively employed single-component smooth wavelets and have primarily focused upon periodic systems with pseudopotentials. Prior investigations employing multiwavelets for electronic structure have been limited in scope.^{20,29} Limited applications of wavelets have been made to molecules including all electrons, and the largest prior such calculation that we are aware of is to the oxygen molecule²¹ employing a large unit cell in a periodic code. More extensive application has been made of multigrid approaches^{4,30–32} and, most significantly, the numerical approach of Becke.^{4,33,34}

The features that distinguish the current work from previous related efforts are primarily the use of (1) separated representations of integral operators; (2) multiwavelet bases with disjoint support; (3) the nonstandard form of operators and functions in three dimensions with full local adaptive refinement; and (4) fast application of integral operators to eliminate the iterative solution of differential equations.

In the following, we first present some of the mathematical and numerical methods, and describe in more detail separated representations for operators. A detailed introduction to multiresolution analysis in multiwavelet bases is deferred to an appendix. Subsequently, we briefly describe the Kohn-Sham density functional theory (DFT) equations and present their integral formulation that has several desirable characteristics. Next, we discuss computation of analytic derivatives of the energy with respect to a parameter in the external potential, and describe the iterative solution scheme and our prototype implementation. Finally, we analyze results of calculations on several atoms and molecules.

II. MATHEMATICAL BACKGROUND

The multiresolution constructions employed in this paper are now fairly standard within the mathematical literature (see, e.g., Refs. 11, 14, 35) but since they may not be familiar to most chemists, we provide a nonrigorous description of

the approach in an appendix provided online within EPAPS.³⁶ Here we just summarize the most salient points.

Many objectives of this paper are accomplished, at least in one dimension, by a few central features of the multiresolution representations. However, additional features are necessary to achieve efficient algorithms in higher dimensions.^{2,13}

(1) Multiresolution wavelet and multiwavelet expansions organize functions and operators efficiently in terms of proximity on a given scale and between the length scales.

(2) Simple and efficient algorithms exist to transform between representations at different scales [$O(N)$ decomposition and reconstruction].

(3) There is a simple truncation and adaptive refinement mechanism to maintain the desired accuracy.

(4) A large, physically significant class of differential and integral operators is sparse in wavelet/multiwavelet bases. High-order convergence is achieved for solving partial differential and integral equations.

(5) Multiwavelet bases with disjoint support maintain high-order convergence in the presence of boundary conditions or singularities.

Numerical algorithms using wavelet bases are similar to other transform methods, e.g., Fourier methods, in that vectors and operators are expanded into a basis and the computations take place in the new system of coordinates. As in all transform methods, availability of a fast algorithm for decomposition and reconstruction is critical. However, unlike the globally defined functions of the Fourier basis, the wavelet/multiwavelet bases have localized support, thus permitting adaptive decomposition of functions and operators and efficient accommodation of boundary conditions. Moreover, due to the vanishing moment property, a large class of operators (which includes those considered in this paper) admit a sparse representation. Using multiwavelets (or discontinuous multiwavelets),³⁵ we achieve a good balance of properties needed for solving partial differential equations (PDEs) and integral equations.¹⁴

Although we have developed a straightforward implementation of multiwavelet bases in dimension 3 for comparison purposes, we use a separable representation of integral kernels (in addition to multiresolution representations) as the main tool to achieve improved performance in three and higher dimensions.²

Separable representation of integral kernels in multiple dimensions

We use a numerical generalization of separation of variables² to avoid the computational cost of the straightforward extension of the multiresolution approach to multiple dimensions. For a given error, ϵ , we represent an operator T with the kernel

$$K(x_{j_1}, x_{j_1}'; x_{j_2}, x_{j_2}'; \dots; x_{j_d}, x_{j_d}')$$

in dimension d as

$$\sum_{l=1}^M s_l K_1^l(x_{j_1}, x_{j_1}') K_2^l(x_{j_2}, x_{j_2}') \cdots K_d^l(x_{j_d}, x_{j_d}'), \quad (1)$$

where s_l is a scalar, $s_1 \gg \dots \gg s_M > 0$, and $K_i^l(x_{j_i}, x_{j'_i})$ are kernels of normalized (norm one) operators T_i^l in dimension $d = 1$. We require the error to be less than ϵ , namely,

$$\left\| T - \sum_{l=1}^M s_l T_1^l \otimes T_2^l \otimes \dots \otimes T_d^l \right\| < \epsilon. \quad (2)$$

We call the scalars s_l separation values and the integer M the separation rank. The smallest M that yields such a representation for a given ϵ is the optimal separation rank. We do not necessarily need the optimal separated representation, but it is essential to obtain M close to the optimal.

In this paper, we consider operators in space of the dimension $d=3$; however we plan to use the same approach for $d=6$ in the future. In dimension $d=2$, the separated representation (1) reduces to a form similar to the singular value decomposition, but with an unusual pairing of indices. Instead of separating the input coordinate $(x_{j'_1}, x_{j'_2})$ from the output coordinate (x_{j_1}, x_{j_2}) , we separate direction $(x_{j_1}, x_{j'_1})$ from direction $(x_{j_2}, x_{j'_2})$.

An immediate question arises as to how to obtain and work with representations (1). In Ref. 2 it was shown that a number of important operators have a small separation rank. In fact, a multidimensional numerical calculus of operators of this form is possible and Ref. 2 provides a brief description of algorithms for this purpose. In this paper we need Eq. (1) to represent two operators, the Poisson kernel and the Green's function for the bound-state Helmholtz operator. Importantly, we combine Eq. (1) with multiresolution multiwavelet representations.

Another question is the number of terms in Eq. (1) since it directly affects the efficiency of algorithms. Theoretically, for operators we are interested in, the separation rank grows as the logarithm of the range of validity of approximation² and, in fact, the actual separation ranks we obtain are mercifully small.

We consider convolutions with kernels

$$G_\mu(x) = \frac{1}{4\pi} \frac{e^{-\mu|x|}}{|x|}, \quad \mu \geq 0, \quad (3)$$

represented in multiwavelet bases. The function G_μ satisfies

$$(-\Delta + \mu^2)G_\mu(x) = \delta(x), \quad x \in \mathbb{R}^3. \quad (4)$$

If $\mu=0$, then $G_0 \equiv G$ in Eq. (3) is the Poisson kernel. For $\mu>0$, G_μ is the kernel of the bound-state Helmholtz operator.

In order to construct the representation of G_μ in multiwavelet bases, we need to compute the integrals

$$r_{ii',jj',kk'}^{n;\mathbf{l}\mathbf{l}'} = \int G_\mu(x-y) \phi_{i,l_1}^n(x_1) \phi_{i',l'_1}^n(y_1) \phi_{j,l_2}^n(x_2) \times \phi_{j',l'_2}^n(y_2) \phi_{k,l_3}^n(x_3) \phi_{k',l'_3}^n(y_3) dx dy, \quad (5)$$

where $x=(x_1, x_2, x_3)$, $y=(y_1, y_2, y_3)$, $\mathbf{l}=(l_1, l_2, l_3)$, and $\mathbf{l}'=(l'_1, l'_2, l'_3)$. The integration in Eq. (5) is over the support of the scaling functions $\phi_{i,l}^n(x) = 2^{n/2} \phi_i(2^n x - l)$, where $i=0, \dots, k-1$, and k is the order of the basis.

Taking advantage of G_μ being a convolution, we obtain

$$r_{ii',jj',kk'}^{n;\mathbf{l}\mathbf{l}'} = r_{ii',jj',kk'}^{n;\mathbf{l}-\mathbf{l}'}$$

and

$$r_{ii',jj',kk'}^{n;\mathbf{l}} = 2^{-3n} \int_{-1}^1 \int_{-1}^1 \int_{-1}^1 G_\mu(2^{-n}(x+\mathbf{l})) \times \Phi_{ii'}(x_1) \Phi_{jj'}(x_2) \Phi_{kk'}(x_3) dx_1 dx_2 dx_3, \quad (6)$$

where

$$\Phi_{ii'}(x) = \int \phi_i(x+y) \phi_{i'}(y) dy, \quad i, i' = 0, \dots, k-1, \quad (7)$$

are cross-correlation functions of the scaling functions.

In order to explain our approach, let us first consider the Poisson kernel, $\mu=0$. Since the Poisson kernel is homogeneous, it is sufficient to set $n=0$ in Eq. (6) and we obtain

$$r_{ii',jj',kk'}^{0;\mathbf{l}} = \int_{-1}^1 \int_{-1}^1 \int_{-1}^1 \frac{1}{|x+\mathbf{l}|} \Phi_{ii'}(x_1) \Phi_{jj'}(x_2) \Phi_{kk'}(x_3) \times dx_1 dx_2 dx_3. \quad (8)$$

The total number of coefficients in Eq. (8), namely, $O(k^6)$ for each shift index \mathbf{l} , is too large for a practical method since then the nominal cost of applying matrices $r_{ii',jj',kk'}^{0;\mathbf{l}}$ is proportional to $O(k^6)$.

However, for any $\epsilon>0$, the integral in Eq. (8) has an approximation with a low separation rank,

$$\left| r_{ii',jj',kk'}^{0;\mathbf{l}} - \sum_{m=1}^M \sigma_m F_{ii'}^{m,l_1} F_{jj'}^{m,l_2} F_{kk'}^{m,l_3} \right| \leq \epsilon, \quad (9)$$

where $M = O(-\ln \epsilon)$. In representation (9) the number of coefficients is Mk^2 for each shift index \mathbf{l} . The separated form of $r_{ii',jj',kk'}^{0;\mathbf{l}}$ in Eq. (9) allows us to apply the operator separately in each direction, with a nominal computational cost of $O(Mk^4)$. Further reduction of the computational cost is carried out by reducing M for each shift index \mathbf{l} . Better than this, the overall computational cost proportional to Mk^3 may be obtained if sparse and/or low-rank representations are used for each of the submatrices which are expected to be of low rank (due to vanishing multipole moments) away from the origin.

The procedure of obtaining Eq. (9) consists of three steps. First we consider the integral

$$\frac{1}{r} = \frac{2}{\sqrt{\pi}} \int_{-\infty}^{\infty} e^{-r^2 e^{2s} + s} ds, \quad (10)$$

and discretize it on an interval $[s_{\min}, s_{\max}]$ using the trapezoidal rule.³⁷ The choice of the interval $[s_{\min}, s_{\max}]$ and the number of nodes depend on the desired range $\delta \leq r \leq 1$ and accuracy ϵ . The interval of integration in s must be chosen so that at the end points the integrand and a sufficient number of its derivatives are less than ϵ for all $\delta \leq r \leq 1$. This simple approach yields a number of nodes M that for useful ranges of parameters is proportional to $-\ln \delta$ and to $-\ln \epsilon$. The result, however, is not near optimal, especially near the end of the interval of integration where s is negative.³⁷




FIG. 1. Relative error of approximating $1/r$ via Eq. (11) in the interval $[10^{-5}, 1]$ using an optimal expansion of 52 Gaussians.

We then use an algorithm for constructing a nearly optimal representation using Ref. 38, where the authors extend the approach in Ref. 39. As a result, we find the generalized Gaussian nodes and weights such that

$$\left| \frac{1}{r} - \sum_{m=1}^M \omega_m e^{-p_m r^2} \right| \leq \frac{\epsilon}{r}, \quad (11)$$

where ϵ is the required maximum relative error within the range and where the number of terms, M , is nearly optimal. The optimization procedure in Ref. 38 also allows us to obtain an approximation

$$\left| \frac{1}{r} - \sum_{m=1}^M \omega_m e^{-p_m r^2} \right| \leq \frac{\epsilon}{r^2}, \quad (12)$$

which is sufficient to obtain Eq. (9) and requires fewer terms.

The third step in obtaining Eq. (9) is to substitute the approximation in Eq. (11) or (12) into Eq. (8) and compute one-dimensional integrals. A detailed derivation of these representations will appear separately.⁴⁰

For instance, a 52-term representation was constructed to obtain a relative precision of 10^{-8} in the interval $[10^{-5}, 1]$ using Eq. (11). The representation obtained using the trapezoidal rule in Eq. (10) used 300+ terms, which may be partially reduced through elementary methods.³⁷ The optimal representations was then formed using Ref. 38, yielding 52 terms for the relative accuracy displayed in Fig. 1.

For $\mu > 0$, the kernel G_μ in Eq. (3) is not homogeneous, so that the approximation should be constructed for each scale separately which will yield a compact representation of the matrix elements [see Eq. (5)] rather than of the kernel as a whole. Although this is clearly the most efficient approach, details of which we will present in a sequel, in the prototype code we used an expansion that spans all length scales. To generate the separated representation of G_μ , $\mu \neq 0$, we use instead of Eq. (10) the integral

$$\frac{e^{-\mu r}}{r} = \frac{2}{\sqrt{\pi}} \int_{-\infty}^{\infty} e^{-r^2 e^{2s} - \mu^2 e^{-2s} + s} ds. \quad (13)$$

Due to the superexponential decay at both ends of the range, the trapezoidal rule may be directly applied to this integral to obtain the desired separated representation.³⁷ In our computations with G_μ , we select the range of validity to include $[0.001/(LZ_{\max}), 1.0]$, where Z_{\max} is the maximum nuclear charge and L is the simulation cell size.

We note that for isotropic kernels, it is sufficient to compute the transition matrix elements for positive translations only. Values for negative translations are obtained using symmetries of the basis functions $\phi_i(x) = (-1)^i \phi_i(1-x)$. The one-dimensional (1D) transition matrix elements are computed and a singular value decomposition is performed for these matrices. Away from the singularity, the matrices have very low operator rank.

When applying the operator, a test is performed to see if it is more efficient to use a low-operator-rank form. In 1D, the break-even point is an operator rank less than half the matrix dimension. In 3D, if the transformations are done in order of increasing rank, the break-even ratios are about 3/4, 2/3, and 1/2 for the first, second, and third transformations, respectively.

Efficient application of the separated kernel requires screening based upon both the magnitude of the coefficients of the source function and the coefficients of the nonstandard form of the three-dimensional operator. We currently estimate the l^2 norm of each block of the operator using the power method, namely, the rapidly convergent iteration $\|A\|_2 = \lim_{n \rightarrow \infty} \sqrt{x_n/x_{n-1}}$, where $x_n = A^t A x_{n-1}$. Typically, two to four iterations provide more than one digit precision starting from a random initial guess.

III. KOHN-SHAM EQUATIONS

The nonlinear Kohn-Sham equations [e.g., Eqs. (7.2.7-9) and (7.4.3) in Ref. 1] result from minimization of the DFT energy functional with respect to variation of the occupied orbitals $[\phi_i(r), i=1, \dots, N]$ which define the electron density ρ (here for a closed-shell system),

$$\rho(r) = 2 \sum_{i=1}^N |\phi_i(r)|^2. \quad (14)$$

The occupied orbitals are the lowest N eigenfunctions of the Kohn-Sham operator (also here casually referred to as the Fock operator since the Hartree and Hartree-Fock equations are very similar¹) which implicitly depends upon the orbitals through the density,

$$[-\frac{1}{2}\nabla^2 + V(r)]\phi_i(r) = \epsilon_i \phi_i(r), \quad (15)$$

$$V(r) = V_{\text{ext}}(r) + V_{\text{coul}}(r) + V_{\text{xc}}^{\text{LDA}}(r). \quad (16)$$

In this paper, the external potential V_{ext} includes only the attraction of the electrons to the nuclei,

$$V_{\text{ext}}(r) = - \sum_{\alpha} \frac{Z_{\alpha}}{|r-r_{\alpha}|}. \quad (17)$$

The Coulomb potential $V_{\text{coul}}(r)$ describes the repulsion between electrons,

$$V_{\text{coul}}(r) = \int dr' \frac{\rho(r')}{|r-r'|}. \quad (18)$$

The current work considers only the standard LDA approximation,^{1,41} for which the exchange-correlation potential $V_{\text{xc}}^{\text{LDA}}$ is a scalar function that depends only upon $\rho(r)$. The Hartree-Fock equations are of similar form, but instead of the local exchange correlation potential $V_{\text{xc}}(r)$ include the

nonlocal exchange potential [Eq. (2.5.12) in Ref. 1]. Implementation of gradient-corrected functionals, Hartree–Fock exchange, and hybrid potentials will be discussed in future publications.

The eigenvalues ϵ_i are also referred to as the orbital energies and will be negative for the occupied orbitals. The asymptotic form of the LDA occupied orbitals is obtained by substituting the form $r^\beta \exp -\alpha r$ into Eqs. (15) and collecting terms at long range. The result is, for a neutral system, $\alpha_i = \sqrt{-2\epsilon_i}$ and $\beta_i = 1/\alpha_i - 1$. The nonlocal exchange potential causes the asymptotic behavior of the Hartree–Fock orbitals⁴² to differ from that of the LDA orbitals, but the actual iterative solutions of the LDA and Hartree–Fock equations are very similar.

A. Integral equation formulation

In 1962, Kalos⁴³ used the following Lippmann–Schwinger integral formulation to determine via Monte Carlo sampling the ground state wave function ψ and corresponding eigenvalue E ,

$$\psi = -2G_\mu V\psi, \quad (19)$$

where G_μ is an integral operator with the kernel being the Green's function defined by

$$(-\nabla^2 + \mu^2)G_\mu(r, r') = \delta(r - r') \quad (20)$$

and $\mu = \sqrt{-2E}$. For one particle in three dimensions with free-space boundary conditions,

$$G_\mu(r, r') = (2\pi)^{-3/2} K_{1/2}(\mu|r - r'|)(\mu|r - r'|)^{-1/2} = \frac{e^{-\mu r}}{4\pi r}. \quad (21)$$

To determine the wave function, the integral equation is iterated and the eigenvalue adjusted to conserve the norm of the wave function. In combination with deflation (Sec. VI), it may be used to extract the eigenvectors of the three-dimensional Hamiltonian. Away from the origin, the bound-state Helmholtz Green's function is smooth and decays more rapidly than the Green's function for the Poisson equation. It is therefore very efficiently represented in the multiwavelet basis. This integral formulation of the DFT equations is also commonly used in band structure calculations.⁴⁴ However, it is important to point out that the scattering-state (positive energy) Green's function does not have a sparse representation in wavelet bases since the function is oscillatory at long range and its higher derivatives do not decay rapidly.

Beyond providing a simple and rapidly convergent iteration to compute the eigenfunctions, this integral equation is of interest because it does not require the use of derivative operators to determine the wave function. In principle, even the total energy may be computed using only integral operators.

In eigenvalue problems a small perturbation of the matrix bounds only the absolute error of the eigenvalues and, thus, for small eigenvalues the relative precision is worse than it is for the large ones. In using the differential operators we are always concerned with subspaces corresponding to the (relatively) small eigenvalues, whereas for integral opera-

tors (inverses of the differential operators) we are concerned with the large eigenvalues, and, thus, capable of maintaining a better relative precision. Also the integral operator in Eq. (19) can be viewed as a preconditioner to improve the rate of convergence.

Additional improvements are expected by using alternative bases that are being developed by one of us (G.B.) to provide a more effective representation of band limited functions.³⁹

B. Boundary conditions

The derivative operator can employ both periodic and zero Dirichlet boundary conditions.¹⁴ The Green's function [Eq. (21)] used in the integral equation iteration [Eq. (19)] is for free-space boundary conditions, which are that the wave function and its derivative are zero at infinity. It is important to enforce these boundary conditions during the iterative process, otherwise nonphysical solutions may be amplified. To enforce these conditions, we multiply all trial molecular orbitals by a simple mask which is a tensor product of the following function:

$$m(x) = \begin{cases} s(x/\tau), & 0 \leq x \leq \tau \\ 1, & \tau \leq x \leq 1 - \tau \\ s((1-x)/\tau), & 1 - \tau \leq x \leq 1, \end{cases} \quad (22)$$

where

$$s(x) = x^2(3 - 2x).$$

The polynomial $s(x)$ is the first β function which is zero at $x=0$, 1 at $x=1$, and also has zero first derivative at $x=0$, 1. We currently choose $\tau=1/16$.

The above approach to enforce the boundary conditions requires that the box be large enough for the orbitals or wave function to become negligible. If the box is not large enough, it is not possible to converge the equations to a precision greater than the implied truncation of the wave function; the energy is less affected. Although the multiresolution decomposition can efficiently treat a large box, it may be yet more efficient to use a smaller box and either match with an asymptotic form or to include the boundary terms in the integral equation. If the wave function and the bound-state Helmholtz Green's function are substituted into Green's theorem, then the integral iteration is modified as follows:

$$\psi = -2 \int_{\Omega} ds G_\mu(r, s) V(s) \psi(s) + \int_{\partial\Omega} d\sigma [G(r, s) \nabla_s \psi(s) \hat{n} - \psi(s) \nabla_s G(r, s) \cdot \hat{n}], \quad (23)$$

where \hat{n} is an outward unit vector normal to the surface. The first integral is over the finite simulation volume and the second integral is over the surface. The additional terms correspond to single and double layer contributions. We have not yet implemented this approach.

IV. ANALYTIC ENERGY DERIVATIVES

Derivatives of the variational Hartree–Fock or DFT energy with respect to the nuclear coordinates (or any parameter in the external potential) are straightforwardly evaluated since the Hellman–Feynman theorem is obeyed up to the truncation threshold:

$$\frac{\partial E}{\partial q} = \left\langle \frac{\partial V_{\text{ext}}}{\partial q} \right\rangle + O(\epsilon), \quad (24)$$

where V_{ext} is the external potential [usually the sum of the electron–nuclear and nuclear–nuclear potentials, Eq. (17)], and q is a parameter (e.g., a nuclear coordinate). The energy for variational models is quadratic in the error in the wave function due to approximate solution of the equations, and the gradient is linear in this error. However, both the energy and the gradient are linear in the basis truncation error.

V. SMOOTHED NUCLEAR POTENTIAL

If the nuclei are located at dyadic points, the multiwavelet basis can efficiently represent the cusps in orbitals at nuclei—the accuracy and high-order convergence are maintained. This is also the case for the singular nuclear potential, except in the vicinity of the nucleus, where many levels of refinement may be necessary for high (normwise) precision. The ability of the polynomial basis to attain high values at the interval endpoints is beneficial in representing the singularity, and the integrals used to project the potential into the polynomial basis are satisfactorily evaluated by Gauss–Legendre quadrature. Moreover, since the type of singularity is known, it is possible to develop quadratures that take the singularity into account, thus reducing further the number of refinement levels.

If the nuclei are displaced away from dyadic points, the high-order convergence for the wave function breaks down near the nucleus, and additional levels of refinement are necessary both for the orbitals and the potential. Maintaining precision in the wave function near the nucleus is important for computation of accurate analytic gradients and other properties. Also, the Gauss–Legendre quadrature can fail, for instance, if a nucleus coincides with a quadrature point.

Since we have not yet implemented adaptive subdivision of boxes (i.e., division by factors other than 2) we have chosen to smooth or band-limit the nuclear potential. Modifying the potential, so that both the potential and the resulting orbitals are smooth near the nucleus, eliminates unnecessary fine length scales, which improves the efficiency and accuracy of calculations especially with the nuclei at nondyadic points.

The electron–nuclear attraction potential (r^{-1}) has been smoothed replacing it with $u(r/c)/c$ where c is a scalar and the function $u(r)$ is

$$u(r) = \frac{\text{erf } r}{r} + \frac{1}{3\sqrt{\pi}}(e^{-r^2} + 16e^{-4r^2}). \quad (25)$$

This function is displayed in Fig. 2. For $r > 6$, $u(r)$ differs from $1/r$ by less than 64-bit machine precision. The first three moments of the error are zero, i.e.,




FIG. 2. The exact ($1/r$) and smoothed [$u(r)$, Eq. (25)] Coulomb potentials.

$$\int_0^{\infty} dr r^{2+n}(u(r) - r^{-1}) = 0 \quad (26)$$

for $n = 0, 1, 2$. These zero moments ensure that the expectation value of the potential is quite accurate, implying that the error arising from use of the modified potential is mostly second order. Other forms may be preferable, but this has proven satisfactory to date.

The parameter c determines the range of the modification to the potential and thus the size of the error in the total energy. From perturbation theory and empirical tests (solving the one-electron equation to high precision), we determined the following relationship between the total error (ϵ) in the energy and the smoothing parameter c for hydrogenic atoms of charge Z ,

$$c = \left(\frac{0.00435\epsilon}{Z^5} \right)^{1/3}. \quad (27)$$

Tests for the helium isoelectronic Hartree–Fock series showed almost exactly twice this error, consistent with double occupation of the orbital.

In Sec. VIII D, we present results comparing the energy and structure of the oxygen core orbital in water with various values of the smoothing parameter.

VI. DFT SOLUTION SCHEME

The only necessary inputs are the nuclear coordinates and charges, the required final precision, and an initial guess for the orbitals which is currently an STO-3G (Ref. 45) wave function generated with NWChem.⁴⁶ In the near future, we plan to replace the Gaussian LCAO initial guess with a superposition of precomputed numerical or Slater-type atomic orbitals. An appropriate box size (L) for the simulation is most readily determined from an estimate of the energy of the highest occupied molecular orbital (HOMO), the expected asymptotic orbital decay of $\exp(-\sqrt{-2\epsilon_{\text{HOMO}}})$, and the required precision. The smoothing of the nuclear potential (Sec. V) is chosen to match the final required precision in the energy.

For a sequence of thresholds, $\epsilon = 10^{-3}, 10^{-5}, \dots$ down to the required final precision, we select the basis to be an odd order (k) of multiwavelets such that $\epsilon = 10^{2-k}$. This empirically seems the most efficient choice, but this conclusion is implementation dependent. If the order of wavelets is too low, then the functions are refined more deeply and the integral kernels decay less rapidly. If it is too high, then unne-

essary work is done processing matrices that are currently treated as dense and/or full rank. Odd orders are more efficient due to the superior error in derivatives.¹⁴

At each threshold (ϵ), we project the previous solution (or the initial guess) into the current basis and also compute the nuclear potential and mask (Sec. III B). Then, the integral equation (Sec. III A) is iteratively solved to a maximum residual norm in any orbital of $\max(\epsilon, 3 \times 10^{-5})$. It is neither useful nor feasible to solve to more than the available precision, so at each level we solve the equations to a residual norm of not less than ϵ . Also, if we are only interested in the energy, there is no need to determine the wave function to high precision. Therefore, for an energy accuracy of 10^{-5} the orbitals are determined to a precision of just 3×10^{-4} , though some results reported in this paper were computed to a much higher precision for benchmark purposes.

With just one molecular orbital, iteration of the integral equation converges satisfactorily. Following Kalos,⁴³ we note that with an exact eigenfunction as input, an incorrect eigenvalue in the integral equation results in the norm of the wave function not being conserved. If ψ is the exact wave function and ϵ is an approximate eigenvalue, we wish to compute a correction Δ to the eigenvalue. Starting from the integral equation, denoting the kinetic energy operator as T ,

$$\psi = -(T - \epsilon - \Delta)^{-1}(V\psi), \quad (28)$$

we expand the operator in a Taylor series and obtain

$$\psi = -(T - \epsilon)^{-1}(V\psi) - \Delta(T - \epsilon)^{-2}(V\psi) + O(\Delta^2). \quad (29)$$

Left projection with $V\psi$ and rearrangement yield the following update for the energy:

$$\Delta = -\frac{\langle V\psi | \psi - \tilde{\psi} \rangle}{\|\tilde{\psi}\|^2}, \quad (30)$$

where

$$\tilde{\psi} = -(T - \epsilon)^{-1}(V\psi). \quad (31)$$

Since the changing norm of the wave function has been absorbed into the energy update, the wave function correction (δ) is written as follows:

$$\delta = \frac{\tilde{\psi}}{\|\tilde{\psi}\|} - \psi. \quad (32)$$

The convergence of this iteration is empirically tested in Sec. VIII B and a related variational expression is described in Ref. 44.

For many-electron systems, we must extract multiple eigenpairs from the Fock operator. Straightforward iteration of the integral equation does not work because all roots will collapse to the lowest root unless the initial guesses are very close to the correct solutions. Two modifications are necessary. First, we use deflation to recast the integral equation for each orbital as a ground-state problem. Let P_i denote a projector onto the space of the eigenfunctions of lower energy than orbital i . At convergence, the i th occupied orbital ϕ_i will be the lowest energy solution of

$$(1 - P_i)H(1 - P_i)\phi_i = \epsilon_i\phi_i, \quad (33)$$

which may be rearranged as

$$(H - P_iH(1 - P_i) - (1 - P_i)HP_i + P_iHP_i)\phi_i = \epsilon_i\phi_i. \quad (34)$$

Since $P_i\phi_i = 0$, only the first two terms on the left-hand side are nonzero. The second term may be included in the potential, thereby incorporating the effect of deflation into iteration of the integral equation. However, we note that if, prior to each iteration, the Hamiltonian or Fock matrix is diagonalized in the space of occupied orbitals, then the second term is also zero and the unmodified integral equation may be used. The second modification is to orthogonalize the updated orbitals in order of increasing energy.

The coupled, nonlinear integral equations are solved using a Krylov-subspace accelerated inexact Newton method⁴⁷ which is similar in spirit to direct inversion in iterative subspace.⁴⁸ The total residual is formed by concatenating the residuals in each orbital [Eq. (31)] and eigenvalue [Eq. (30)].

Each iteration involves (1) optionally computing and diagonalizing the Fock matrix in the space of occupied orbitals (the nonlinear equation solver requires a consistent definition of the variables, so it is necessary to keep track of the maximum overlap of the input and output orbitals and to keep phases consistent); (2) computing the density as the sum of the square of the molecular orbitals; (3) computation of the Coulomb potential by convolution with the separated form of the Poisson kernel; (4) computation of the exchange-correlation potential as a local function of the density; (5) iteration of the integral equation once for each of the orbitals, which requires multiplication by the potential, and then convolution with the separated form of the Helmholtz kernel; (6) update of the eigenvalues and orbitals via the subspace solver; and (7) orthogonalize the resulting orbitals.

The iterative solution proceeds in a mostly out-of-core fashion, so that at any instant only a few numerical functions are in memory, with this number not depending upon the number of electrons in the system.

VII. PROTOTYPE 3D IMPLEMENTATION

Our initial implementation uses Python⁴⁹ for high-level control and C/C++/Fortran for computationally intensive operations including matrix transformations, quadratures, and the innermost loops. At the highest level, we have defined a Function class that includes methods for evaluation, compression, reconstruction, addition, multiplication by a function or scalar, differentiation, application of the Laplacian, and other operations. The operator overloading capabilities of Python provide great expressivity and enable very compact programs. For instance, if `psi` is an instance of the Function class representing an orbital, and similarly V represents the potential, then the following statement applies the Fock operator to the orbital:

$$H_{\text{psi}} = -0.5 * (\text{DelSq} * \text{psi}) + V * \text{psi}.$$

`DelSq` is an empty class that is never instantiated. If an instance of Function is multiplied on the right by `DelSq`, the function's Laplacian method is invoked. Evaluation of a function at a point with the natural semantics `psi(x, y, z)` is accomplished by overloading the function call operator.

Internally, a compressed function is represented using the full sparsity of the translations at each level, but currently not exploiting the sparsity or low rank within a block of coefficients. The blocks of difference coefficients of a function in compressed form are stored in a directory (or associative array or hash table), d , which has an entry for each significant level. Each of these entries, $d[n]$, is a directory indexed by the translations. Only the significant (above threshold) translations are stored. Finally, $d[n][lx,ly,lz]$ is a Tensor object that stores the multiwavelet coefficients d_l^n . The sum coefficients are similarly represented in a reconstructed function. The Python Tensor class is implemented in C++. SWIG⁵⁰ is used for facile integration of Python with C, C++, and Fortran. Persistent and efficient storage of functions on disk is straightforwardly accomplished by providing a Python dictionary-like interface to a directory on disk.

Addition of a compressed function is straightforward—the sum coefficients at level zero and the difference coefficients at all levels are simply added. Addition of a reconstructed function may be similarly accomplished by adding the sum coefficients at all levels, noting that the result may have scaling function coefficients at multiple length scales which must be correctly accounted for during a subsequent compression. The Function class provides an operation that performs $af+bg$, where a and b are scalars and f and g are functions. Multiplication of two reconstructed functions is performed as described in Refs. 14 and 51 by reconstructing the functions on the locally finest level, optionally recurring down one level to preserve the accuracy of the approximation, tabulating the functions on the quadrature grid, multiplying the values, projecting back into the scaling function basis, and finally compressing and truncating. Application of a local function [i.e., $g(f(r))$] is similarly accomplished. Evaluation of a compressed function at a point is performed by recurring down the tree, accumulating the sum coefficients until they are reconstructed at the finest level in the box containing the evaluation point. The scaling functions (Legendre polynomials) are then evaluated at the point of interest and contracted with the coefficients. If many points are to be evaluated, it is more efficient to first reconstruct the function in the scaling function basis at the locally finest level.

Differentiation is currently performed using a periodic central difference approximation as described in Ref. 14. A zero boundary condition is enforced by embedding the solution volume in zeroes (i.e., references to coefficients outside the solution volume are treated as zero). A negative, self-adjoint approximation to the Laplacian is formed also as described in Ref. 14. The derivatives, for multiwavelets of odd order k , are accurate to order k ;¹⁴ for even order multiwavelets, the order of the error in the derivatives is one less. Therefore, odd multiwavelets are preferred.

The Function class is able to export a description of the current basis (i.e., a list of the significant translations at each level) and also can restrict to a specified basis. This functionality is useful during iterative processes, such as diagonalization. The class is also able to output a description of the

adaptive mesh and a 3D tabulation of the function in formats suitable for visualization by Opendx.⁵²

Multiwavelets of any order are supported. The two-scale coefficients are generated using Alpert's algorithm⁵³ in extended precision floating point numbers in Python (e.g., 156-bit arithmetic is used to generate the two-scale coefficients for order 10). The extended precision is necessary only to generate the coefficients which are stored for subsequent use. Standard double-precision arithmetic is used for all other operations.

VIII. INITIAL APPLICATIONS

A. Helium isoelectronic sequence

To verify that the results did not depend upon the size of the simulation cell, and to demonstrate the ability of the multiresolution representation to accommodate multiple length scales, we performed high-accuracy Hartree–Fock calculations upon the helium isoelectronic sequence He, Be²⁺, Mg¹⁰⁺, and Ca²⁰⁺ using simulation cell of sizes of 20, 40, and 80 a.u.

To attain independence of the cell size, the compression of the nuclear potential must take into account its dependence upon the cell size. The molecular coordinates are in the cube $[-L/2, L/2]$ (in atomic units) and these are mapped to the unit cube $[0,1]$ for the computation. The transformation scales lengths by L , so the nuclear potential due to a nucleus of charge Z at the origin becomes $V(x) = -Z/(L|x-1/2|)$ with $0 \leq x \leq 1$. In order to maintain precision independent of L , one can either compress $LV(x)$ with the standard precision (ϵ) and subsequently scale by $1/L$, or compress $V(x)$ with precision ϵ/L . We currently do the latter. Similarly, as noted above (Sec. II), the lower limit for the separated form of the bound-state Helmholtz operator must also scale correctly with the cell size.

With these modifications, the energies are found to be independent of the cell size and agree to at least seven decimal places in atomic units (the requested precision) with the results of Thakkar.⁵⁴

The ability to maintain precision independent of the cell size suggests that calculations on large molecules containing atoms at least into the third period will maintain the desired precision. Doubling the simulation cell size adds an additional level to the simulation or, in general, an additional $4^3 - 2^3 = 48$ boxes at the highest level, independent of the number of boxes or length scales already present. In the special case of a single atom at the origin, doubling the box size at most adds eight additional boxes around the origin at the finest level. Thus, the calculations in larger boxes are only slightly more expensive.

B. Hydrogen molecular ion and molecule

Initial molecular calculations were performed upon H₂⁺ and H₂ for which accurate Hartree–Fock results are available in the literature.^{20,55,56} The results for H₂ (bond length 1.4 bohrs, box size $L = 89.6$ bohrs) are given in Table I. In contrast to the recent three-dimensional, mixed-basis results of Pahl and Handy,⁵⁵ no extrapolation was necessary and the best result is accurate to about 10^{-10} a.u.

TABLE I. Hartree–Fock energy of the hydrogen molecule ($r=1.4$ bohrs) computed with various order wavelets (k). The truncation threshold in each calculation was 10^{2-k} and the nuclear potential smoothing parameter was chosen [Eq. (27) and Sec. VI] to yield an energy accurate to at least $1E-10$. For comparison, the best variational energy we are aware of is $-1.133\ 629\ 571\ 456$ due to Mitin (Ref. 56) using a very large Gaussian basis including off-center functions.

k	Energy
5	$-1.133\ 556\ 788\ 8$
7	$-1.133\ 629\ 435\ 3$
9	$-1.133\ 629\ 569\ 8$
11	$-1.133\ 629\ 571\ 3$
13	$-1.133\ 629\ 571\ 4$

The H_2 molecule (bond length 1.4 bohrs) was also used to examine the convergence behavior for iteration of the integral equation with and without use of the subspace information, and with and without the multiscale solution. All calculations started from the STO-3G orbital generated at a bond length of 2 bohrs. Simple iteration of the integral equation, updating the eigenvalue according to Eq. (30), using a multiwavelet of order $k=9$ and a truncation threshold of 10^{-7} , converged to a residual norm of 10^{-6} in ten iterations. Use of the iterative subspace information⁴⁷ to accelerate convergence reduces the number of iterations. For many-electron systems, this is essential for reliable convergence. Repeating these two calculations with the multiscale solver (Sec. VI) requires a few more iterations overall, but requires only one iteration at the most accurate, and most expensive, threshold. Calculations at successive resolutions are approximately twofold to fourfold more expensive.

C. LDA calculations on atoms

We have implemented the local density approximation¹ (LDA; the Dirac–Slater exchange potential with the VWN-5 correlation potential⁴¹) for closed shell systems. In order to verify the implementation and to explore possible issues with calculations on many-electron atoms, we performed calculations on the neutral atoms He, Be, Mg, Ca, and Sr. The results agree with the atomic data from the NIST database,⁵⁷ which are reported to six decimal places in atomic units.

Figure 3 displays a radial plot of the strontium s orbitals. The correct asymptotic decay (Sec. III) is observed for each orbital until the truncation threshold is encountered. Previously, as described in Sec. VI, for all solution thresholds we employed diagonalization within the occupied space to incorporate the effects of deflation. Since the integral and differential forms of the LDA equations are only consistent up to the truncation threshold, diagonalization inevitably mixed the orbitals, resulting in less satisfactory asymptotic forms. The LDA equations were still being solved to the desired precision and the energy was correct. However, since the diagonalization does not significantly accelerate convergence once the eigenfunctions are identified to low precision, and to avoid mixing the final eigenfunctions of the integral equation, we presently only diagonalize in the occupied subspace with the initial solution threshold. This was how displayed orbitals were obtained.




FIG. 3. Plot of the numerical strontium s orbitals. The orbitals demonstrate the correct asymptotic behavior until the truncation threshold is reached and accurately span about 11 orders of magnitude in value.

D. LDA calculations on polyatomic molecules

LDA calculations were performed upon water and benzene near their equilibrium geometries. In Table II, we report the geometry and other simulation parameters. In Table III, we report the corresponding energies and other information for each threshold used in the iterative solution. The best energies should be accurate to 10^{-7} hartree for both systems, with each orbital being determined with residual norm of 3×10^{-5} . These results demonstrate that high precision is attainable for general polyatomic systems.

The chosen geometry and simulation cell size (Table II) for both water and benzene place the nuclei at dyadic points at some level of refinement (level 1 for oxygen and level 7 for hydrogen). Beyond this level of adaptive refinement, the nuclei will always be at grid nodes. An important advantage of the multiwavelet basis (Sec. II) is that it can maintain high-order convergence if singular points (e.g., the cusps in the wave function) are located at nodes. It is possible to adaptively refine the grid by unequally dividing boxes (rather than exactly in two), which would enable the nuclei to sit always at grid nodes regardless of the molecular geometry.

TABLE II. Geometries and simulation cell size (both in atomic units) for LDA calculations on water and benzene.

Molecule	Atom	x	y	z
Water	$L=36.8$			
	O	0.0	0.0	0.0
	H	1.4375	0.0	1.15
	H	-1.4375	0.0	1.15
Benzene	$L=45.0$			
	C	0.0	$\pm 2.614\ 746\ 093\ 75$	0.0
	C	$\pm 2.263\ 183\ 593\ 75$	$\pm 1.318\ 359\ 375$	0.0
	H	0.0	$\pm 4.658\ 203\ 125$	0.0
	H	$\pm 4.020\ 996\ 093\ 75$	$\pm 2.329\ 101\ 562\ 5$	0.0

TABLE III. Energies for calculations on water (at dyadic and nondyadic geometries) and benzene at various thresholds (the wavelet order being adjusted as described in the text). The maximum size of any orbital is given in Mwords (i.e., the number of coefficients stored). All calculations were performed with an energy precision of 10^{-7} hartree.

Molecule	Threshold	Energy	Max. orbital size (Mwords)
Water (dyadic)	10^{-3}	-75.910 585 6	0.11
	10^{-5}	-75.913 532 5	0.58
	10^{-7}	-75.913 555 7	5.1
	10^{-9}	-75.913 556 3	50
Water (nondyadic)	10^{-3}	-75.911 011 9	0.11
	10^{-5}	-75.913 547 4	0.57
	10^{-7}	-75.913 555 9	5.2
	10^{-9}	-75.913 556 3	51
Benzene	10^{-3}	-230.184 463 0	0.26
	10^{-5}	-230.201 642 4	1.5
	10^{-7}	-230.201 702 8	12
	10^{-9}	-230.201 704 8	156

Since we have not yet implemented this, there is the possibility of the energy or other aspects of the simulation not being translationally invariant.

To investigate translational invariance of the energy, we repeated the water calculation, displacing the molecule away from the origin by $1/9.9$ ($=0.101\ 01\dots$) a.u., thereby ensuring that the nuclei are not at dyadic points at any level of refinement. The results are also included in Table III. The energy is observed to be translationally invariant within a precision controlled by the truncation threshold and convergence of the orbitals; this is to be expected since it is guaranteed by the multiresolution algorithm. Also, the orbitals do not differ significantly in size from those computed at the dyadic point. This is perhaps a consequence of two things: First, an additional few levels of refinement in the immediate vicinity of the nuclei adds relatively few new coefficients to the overall simulation. Second, we have constructed the smoothed nuclear potential (Sec. V) so that the resulting orbitals are smooth (Gaussian) near the nucleus.

To illustrate the effect of the smoothed nuclear potential (Sec. 5) upon the orbitals, Fig. 4 displays a slice through the




FIG. 4. Slices through the water oxygen 1s orbital for values of the nuclear potential smoothing parameter (Sec. V) $c=0.05$, $0.016\ 67$, $0.005\ 56$, $0.001\ 85$, and 10^{-6} . The corresponding energies are $-76.919\ 073$, $-76.914\ 719$, $-76.913\ 599$, $-76.913\ 530$, and $-76.913\ 533$, respectively. The smoother orbitals correspond to larger values of the smoothing parameter, and the lines for the smallest two values are indistinguishable at the resolution of the plot.

water oxygen 1s orbital for several values of the smoothing parameter (c in Sec. V). The molecular geometry was as given in Table II, which places the nuclei at dyadic points. The displayed orbitals were computed by self-consistently solving the LDA equations with the smoothed potential with a finest truncation threshold of 10^{-5} . The resulting orbitals are smooth at the nucleus until the effect of smoothing falls below the truncation threshold. At this point there is an actual cusp that can be exactly represented with the nucleus at a dyadic point.

Finally, to reflect more routine chemical computations, the calculations on water and benzene were repeated using C_{2v} and D_{2h} symmetries, respectively, and with a total energy precision of 0.05 kcal/mol. The total computational times, starting from an STO-3G initial guess and using a single 2.4 GHz Pentium-IV processor, are 91 and 850 s, respectively. With our current prototype code these calculations are expensive, but we anticipate substantial performance improvements in future versions. We note that the ratio of time for the two calculations ($850/91=9.3$) is close to that expected from quadratic scaling upon the number of occupied orbitals [$(21/5)^2/2=8.8$].

IX. CONCLUSIONS

We have formulated and demonstrated the fully adaptive, multiresolution solution, with guaranteed precision, of the all-electron density functional equations for general polyatomic molecules. The most significant development of this work is the application of efficient, accurate, low-separation rank representations of integral operators which enable practical computation with multiwavelet bases in three dimensions and lay the foundation for computation in even higher dimensions. Adaptive local refinement (which is different for each orbital) is combined with the nonstandard form of operators for efficient computation with a low memory requirement since all required operators are Toeplitz. The multiwavelet bases with disjoint support provide high-order convergence and can maintain this order even in the presence of singularities and boundaries located at dyadic points. These bases also enable efficient computation on modern cached-based processors since many computations are phrased as small (often low-rank) matrix-matrix operations. The fast (i.e., with a computational cost scaling linearly with respect to the number of coefficients and logarithmically with the simulation volume and required precision) application of integral operators replaces the iterative solution of poorly conditioned differential equations with the solution of well conditioned integral equations.

This paper has focused on practical details of the fast algorithms, but it is important to note that underlying this apparent complexity is a standard, well-understood orthogonal basis set—Legendre (or interpolating) polynomials. As a consequence, most standard quantum mechanical methods and interpretations are immediately applicable, and most physical operators fit automatically into the $O(N)$ framework with the computation of simple matrix elements over Legendre polynomials with standard numerical quadrature. Interpretation of densities and orbitals in terms of atomic

populations can be readily accomplished by projection into a minimum atomic basis.⁵⁸

The quantum chemical algorithms are written at a very high level in terms of operators and functions. This should be compared with the conventional approach in atomic bases which manipulates explicitly indexed, sparse, multidimensional arrays of matrix elements (one- and two-electron integrals). This has led to the very rapid and compact implementation of energies, analytic derivatives, and linear response theory for excited states of LDA, Hartree–Fock, gradient-corrected, and hybrid functionals. These developments will be discussed in future publications.

The performance of the current prototype code is satisfactory in that it reliably delivers the requested precision, all of the operations scale correctly (by construction) with the system size, and the execution time is significantly lower than that required to obtain similar precision in the total energy from conventional methods. Currently, the time (91 s) to determine the energy of a single water molecule to 10^{-5} hartree is expected to be significantly less than a LCAO calculations yielding the same precision, and the numerical calculations using canonical orbitals are expected to have a predominantly quadratic scaling that approaches linearity if the canonical orbitals are localized and sparse linear-algebra techniques are adopted. This is much better than the expected scaling of conventional methods in high-quality bases. However, chemistry is not primarily concerned with the precision of total energies, but with that of energy differences. Relatively small atom-centered bases can compute accurate energy differences if the large, primarily intra-atomic errors cancel. Incorporating some of these advantages into the multiresolution approach is a topic of future research. We anticipate that a more refined implementation can increase the speed of the current code about threefold. Faster than this will require new algorithms for application of the integral operators and improved basis sets that reduce the amount of oversampling implicit in the use of polynomials.

Finally, as stated at the beginning of this paper, we regard this solution of one-electron models as the essential first step in addressing the more significant basis and computational problems that arise in many-body theories. The development of low-separation rank representations² is central to this effort.

ACKNOWLEDGMENTS

R.J.H., T.Y., and G.F. were funded by the Scientific Discovery through Advanced Computing (SciDAC) program of the U.S. Department of Energy, the division of Basic Energy Science, Office of Science, under Contract No. DE-AC05-00OR22725 with Oak Ridge National Laboratory. G.F. was partially supported by the Office of Advanced Scientific Computing Research, Program in Mathematics, Information and CS through the Scientific Application Prototype Program of SciDAC. Funding for Z.G. was provided by ORNL Laboratory Directed Research and Development Funds. The work of G.B. was supported in part by University of Virginia subcontract under Grant No. MDA972-00-1-0016 and NSF/ITR Grant Nos. ACI-0082982 and DMS-0219326. This research was performed in part using the resources of the National

Energy Scientific Computing Center which is supported by the Office of Energy Research of the U.S. Department of Energy under Contract No. DE-AC03-76SF0098 and the Center for Computational Sciences at Oak Ridge National Laboratory under Contract No. DE-AC05-00OR22725.

- ¹R. G. Parr and W. Yang, *Density-Functional Theory of Atoms and Molecules*, The International Series of Monographs on Chemistry Vol. 16 (Oxford University Press, New York, 1989).
- ²G. Beylkin and M. J. Mohlenkamp, Proc. Natl. Acad. Sci. U.S.A. **99**, 10246 (2002); See <http://www.pnas.org/cgi/content/abstract/112329799v1>
- ³A. D. Becke and R. M. Dickson, J. Chem. Phys. **89**, 2993 (1988).
- ⁴A. D. Becke and R. M. Dickson, J. Chem. Phys. **92**, 3610 (1990).
- ⁵C. Roothaan, Rev. Mod. Phys. **23**, 69 (1951).
- ⁶C. Roothaan, Rev. Mod. Phys. **32**, 179 (1960).
- ⁷G. Hall, Proc. R. Soc. London **A230**, 541 (1951).
- ⁸S. Boys, Proc. R. Soc. London **A200**, 542 (1950).
- ⁹T. Helgaker and P. Taylor, *Modern Electronic Structure Theory* (World Scientific, Singapore, 1995), pp. 725–856.
- ¹⁰S. F. Boys and F. Bernardi, Mol. Phys. **19**, 553 (1970).
- ¹¹S. Jaffard, Y. Meyer, and R. D. Ryan, *Wavelets: Tools for Science & Technology* [Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2001], revised edition.
- ¹²G. Beylkin, R. Coifman, and V. Rokhlin, Commun. Pure Appl. Math. **44**, 141 (1991); Yale University Technical Report No. YALEU/DCS/RR-696, 1989 (unpublished).
- ¹³G. Beylkin and R. Cramer, SIAM J. Sci. Comput. (USA) **24**, 81 (2002).
- ¹⁴B. Alpert, G. Beylkin, D. Gines, and L. Vozovoi, J. Comput. Phys. **182**, 149 (2002); See <ftp://amath.colorado.edu/pub/wavelets/papers/mwa.pdf>
- ¹⁵H. Cheng, L. Greengard, and V. Rokhlin, J. Comput. Phys. **155**, 468 (1999).
- ¹⁶U. Trottenberg, C. Oosterlee, and A. Schüller, *Multigrid* (Academic, New York, 2001).
- ¹⁷T. Beck, Int. J. Quantum Chem. **65**, 477 (1997).
- ¹⁸J. Bernholc, E. Briggs, D. Sullivan, C. Brabec, M. Nardelli, K. Rapcewicz, C. Roland, and M. Wensell, Int. J. Quantum Chem. **65**, 531 (1997).
- ¹⁹I. P. Daykov, T. A. Arias, and T. D. Engeness, Phys. Rev. Lett. **90**, 216402 (2003).
- ²⁰A. Maloney, J. Kinsey, and B. Johnson, J. Chem. Phys. **117**, 3548 (2002).
- ²¹T. Engeness and T. Arias, Phys. Rev. B **65**, 165106 (2002).
- ²²T. Arias, Rev. Mod. Phys. **71**, 267 (1999).
- ²³S. Goedecker and O. Ivanov, Phys. Rev. B **59**, 7270 (1999).
- ²⁴R. Lippert, T. Arias, and A. Edelman, J. Comput. Phys. **140**, 278 (1998).
- ²⁵S. Goedecker and O. Ivanov, Solid State Commun. **105**, 665 (1998).
- ²⁶P. Fischer and M. Defranceschi, SIAM (Soc. Ind. Appl. Math.) J. Numer. Anal. **35**, 1 (1998).
- ²⁷M. Brewster, G. Fann, and Z. Yang, J. Math. Chem. **22**, 117 (1997).
- ²⁸K. Cho, T. Arias, J. Joannopoulos, and P. Lam, Phys. Rev. Lett. **71**, 1808 (1993).
- ²⁹A. Niklasson, C. Tymczak, and H. Roder, Phys. Rev. B **66**, 155120 (2002).
- ³⁰T. Beck, Rev. Mod. Phys. **72**, 1041 (2000).
- ³¹J. Wang and T. Beck, J. Chem. Phys. **112**, 9223 (2000).
- ³²J. Fattbert and J. Bernholc, Phys. Rev. B **62**, 1713 (2000).
- ³³R. M. Dickson and A. D. Becke, J. Chem. Phys. **99**, 3898 (1993).
- ³⁴R. M. Dickson and A. D. Becke, J. Phys. Chem. **100**, 16105 (1996).
- ³⁵B. Alpert, SIAM J. Math. Anal. **24**, 246 (1993).
- ³⁶See EPAPS Document No. E-JCPA6-121-311438 for Appendix to Multiresolution Quantum Chemistry: basic theory and initial applications. A direct link to this document may be found in the online article's HTML reference section. The document may also be reached via the EPAPS homepage (<http://www.aip.org/pubservs/epaps.html>) or from <ftp.aip.org> in the directory /epaps/. See the EPAPS homepage for more information.
- ³⁷R. J. Harrison, G. I. Fann, T. Yanai, and G. Beylkin, *Lecture Notes in Computer Science* (Springer, Heidelberg, 2003), Vol. 2660, pp. 103–110.
- ³⁸G. Beylkin and L. Monzón (unpublished).
- ³⁹G. Beylkin and L. Monzón, Appl. Comput. Harmon. Anal. **12**, 332 (2002).
- ⁴⁰G. Beylkin, G. Fann, R. Harrison, and R. Cramer (unpublished).
- ⁴¹S. H. Vosko, L. Wilk, and M. Nusair, Can. J. Phys. **58**, 1200 (1980).
- ⁴²N. C. Handy, M. T. Marron, and H. J. Silverst, Phys. Rev. **180**, 45 (1969).
- ⁴³M. H. Kalos, Phys. Rev. **128**, 1791 (1962).
- ⁴⁴A. Gonis and W. Butler, *Multiple Scattering in Solids* (Springer, Berlin, 1999).

⁴⁵W. Hehre, R. Stewart, and J. Pople, *J. Chem. Phys.* **51**, 2657 (1969).

⁴⁶R. Kendall, E. Apra, D. Bernholdt *et al.*, *Comput. Phys. Commun.* **128**, 260 (2000).

⁴⁷R. J. Harrison, *J. Comput. Chem.* **25**, 328 (2004).

⁴⁸P. Pulay, *Chem. Phys. Lett.* **73**, 393 (1980).

⁴⁹See <http://www.python.org>

⁵⁰See <http://www.swig.org>

⁵¹G. Beylkin and J. M. Keiser, *J. Comput. Phys.* **132**, 233 (1997).

⁵²See <http://www.opendx.org>

⁵³B. Alpert, Ph.D. thesis, Yale University, 1990.

⁵⁴T. Koga, M. Omura, H. Teruya, and A. Thakkar, *J. Phys. B* **28**, 3113 (1995).

⁵⁵F. A. Pahl and N. C. Handy, *Mol. Phys.* **100**, 3199 (2002).

⁵⁶A. Mitin, *Phys. Rev. A* **6201**, 010501 (2000).

⁵⁷See <http://physics.nist.gov/PhysRefData/DFTdata/contents.html>

⁵⁸B. Bursten and R. Fenske, *J. Chem. Phys.* **67**, 3138 (1977).