

Treecode-accelerated Green iteration for Kohn-Sham density functional theory

Nathan Vaughn^{a,*}, Vikram Gavini^{b,c}, Robert Krasny^a

^a Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, USA

^b Department of Mechanical Engineering, University of Michigan, Ann Arbor, MI 48109, USA

^c Department of Materials Science and Engineering, University of Michigan, Ann Arbor, MI 48109, USA

ARTICLE INFO

Article history:

Available online 5 January 2021

Keywords:

All-electron Kohn-Sham density functional theory

Real-space integral equation

Green iteration

A-priori adaptive mesh refinement

Barycentric Lagrange treecode

GPU acceleration

ABSTRACT

We present a real-space computational method called treecode-accelerated Green Iteration (TAGI) for all-electron Kohn-Sham Density Functional Theory. TAGI is based on a reformulation of the Kohn-Sham equations in which the eigenvalue problem in differential form is converted into a fixed-point problem in integral form by convolution with the modified Helmholtz Green's function. In each self-consistent field (SCF) iteration, the fixed-points are computed by Green Iteration, where the discrete convolution sums are efficiently evaluated by a GPU-accelerated barycentric Lagrange treecode. Other techniques used in TAGI include *a-priori* adaptive mesh refinement, Fejér quadrature, singularity subtraction, gradient-free eigenvalue update, and Anderson mixing to accelerate convergence of the SCF and Green Iterations. Ground state energy computations of several atoms (Li, Be, O) and small molecules (H₂, CO, C₆H₆) demonstrate TAGI's ability to efficiently achieve chemical accuracy.

© 2021 Elsevier Inc. All rights reserved.

1. Introduction

Electronic structure calculations complement materials engineering experiments by predicting properties such as binding energy, inter-atomic forces, magnetization, and doping effects. Density Functional Theory (DFT) [6], which describes a system and its properties by its electron density, has been the workhorse of ground state electronic structure computations. For an N_e -electron system, the Kohn-Sham approach to DFT [7] reduces the $3N_e$ -dimensional problem for the many-body wavefunction to a 3-dimensional problem for the electron density. In particular, the system of N_e interacting electrons is replaced by a fictitious system of N_e non-interacting electrons giving rise to the same electron density. In principle, the Kohn-Sham formulation is exact for the ground state properties of materials systems, but it requires knowledge of the exchange-correlation functional, which is not known explicitly and is modeled in practice. Approximating the exchange-correlation functional is an active area of research [8–10], and better approximations enable Kohn-Sham DFT to more accurately predict ground state materials properties.

Previous related work. There are many options for performing either all-electron or pseudopotential DFT calculations, where, in the latter case, only the valence electrons are computed. Often a basis set is used to represent the wavefunctions and electron density [11]. For periodic systems, the plane-wave basis is widely used for pseudopotential calculations [12–16], and for all-electron calculations that require higher resolution to capture the rapidly oscillating wavefunctions, the aug-

* Corresponding author.

E-mail addresses: njvaughn@umich.edu (N. Vaughn), vikramg@umich.edu (V. Gavini), krasny@umich.edu (R. Krasny).

mented plane wave basis [17] and its variants are employed [18–22]. For non-periodic systems, Gaussian basis sets are widely used in quantum chemistry codes [1,2,23,24] as they afford analytic evaluation of many integral and differential operators. A more recent option is the finite-element basis [25–28], which efficiently treats periodic or non-periodic boundary conditions, and pseudopotential or all-electron systems using higher order finite-elements [29–32].

The previously described methods are based on solving the Kohn-Sham eigenvalue equation, a single-particle Schrödinger-like differential equation. In this work we consider an alternative approach in which the eigenvalue problem in differential form is converted into a fixed-point problem in integral form by convolution with the modified Helmholtz Green's function. While integral equation methods are extensively used for the wave equations arising in classical scattering [3,33–39] and quantum scattering [40–45], these methods have received much less attention for eigenvalue problems corresponding to ground state calculations of the Schrödinger or Kohn-Sham equations. The integral equation approach was first applied by Kalos [46] to solve the Schrödinger equation for 3- and 4-electron systems using Monte Carlo minimization. Later, Zhao et al. [47] used this approach to investigate various 1-electron systems in 3D, where the convolution integrals were computed using the Multi-Level Fast Multipole Method.

In recent work, the integral equation approach was extended to the Hartree-Fock and Kohn-Sham equations, where the electron density was updated in self-consistent field (SCF) iterations, and the fixed-point problem for the wavefunctions and eigenvalues in each SCF was solved by a process called Green Iteration. Harrison et al. [48] implemented Green Iteration for the Kohn-Sham equations in a multiwavelet basis that provides local refinement for each wavefunction, and this is now incorporated in the MADNESS code [49]. The convergence of Green Iteration for the many-body Schrödinger equation was investigated by Mohlenkamp and Young [50,51], who proved that the iteration converges for $N_e = 1$ and $N_e = 2$, provided the interaction potential belongs to the function space $L^2(\mathbb{R}^3) + L^\infty(\mathbb{R}^3)$ and the $L^\infty(\mathbb{R}^3)$ piece can be taken to be arbitrarily small. Khoromskij [52] later extended this proof to Kohn-Sham DFT, where now the electron-electron interaction potential is replaced by the exchange-correlation potential which must satisfy the same function space requirements. Subsequently, Rakhuba and Oseledets [53,54] applied Green Iteration to the Hartree-Fock and Kohn-Sham equations in a Tucker tensor basis that uses low rank approximations of the wavefunctions.

Present work. We present a new integral equation based method called Treecode-Accelerated Green Iteration (TAGI) for all-electron Kohn-Sham DFT calculations. The key features of TAGI that enable accurate and efficient calculations are (1) *a-priori* adaptive mesh refinement, (2) high order quadrature, (3) singularity subtraction for convolution integrals, (4) gradient-free eigenvalue update, (5) Anderson mixing for SCF and Green Iteration, and (6) GPU-accelerated treecode computation of discrete convolution sums.

TAGI is a real-space method in which the fields are represented directly at quadrature points. TAGI uses *a-priori* adaptive mesh refinement to efficiently represent the fields, which vary rapidly near the nuclei but decay smoothly in the far-field. The adaptive refinement scheme results in a set of cuboid cells, which are discretized with Chebyshev points of the first kind, and all integrals are evaluated with the Fejér (“classical” Clenshaw-Curtis) quadrature rule [55,56]. The convolution integrals have singular kernels (Coulomb and Yukawa), which impede the accuracy of the quadrature rule, and we employ singularity subtraction to reduce the error in the quadrature sums. A standard singularity subtraction scheme is used for the Yukawa kernel [57,58] and we developed a modified version for the Coulomb kernel. To further improve accuracy, we use a gradient-free eigenvalue update [48] within Green Iteration to eliminate the error arising from numerical differentiation in the standard gradient eigenvalue update. We analyze the convergence rate of Green Iteration and use a fixed-point acceleration technique to alleviate slow convergence. Finally, the discrete convolution sums are efficiently evaluated using a Barycentric Lagrange Treecode [59] (BLTC), which reduces the computational complexity from $O(N^2)$ to $O(N \log N)$ while introducing a small and controllable approximation error. Furthermore, the BLTC is accelerated on GPUs with OpenACC [60] and across multiple GPUs on a single node with OpenMP. We demonstrate the impact on accuracy and efficiency of each of the previously described features on the carbon monoxide molecule, and then perform ground state energy calculations for several atoms and molecules, demonstrating TAGI's ability to achieve chemical accuracy of 1 mHa/atom.

The paper is organized as follows. Section 2 presents Kohn-Sham DFT, and the standard Self-Consistent Field iteration for computing the ground state density and wavefunctions. Section 3 presents the integral equation formulation we employ and Green Iteration for the resulting fixed-point problem. Section 4 describes the numerical techniques developed in this work to enhance the accuracy of the integral formulation, and demonstrates these ideas on the carbon monoxide molecule. Section 5 investigates the convergence rate of Green Iteration and demonstrates the fixed-point acceleration technique used in TAGI. Section 6 describes the treecode algorithm for computing fast approximations of the convolution integrals and demonstrates the efficiency of the GPU-accelerated implementation used in this work. Section 7 applies TAGI to several atoms and small molecules, achieving chemical accuracy of 1 mHa/atom with respect to reference values. Section 8 provides a summary of our findings, and discusses a path forward for this approach to further improve performance and scale to larger systems.

2. Kohn-Sham density functional theory

The input to Kohn-Sham DFT consists of the positions and atomic numbers of the atoms in the system, and the output consists of the ground-state electron density along with the Kohn-Sham single-electron wavefunctions, from which the desired observables (including ground-state energy and ionic forces) can be computed. The Kohn-Sham equations are

$$\mathcal{H}[\rho]\psi_i(\mathbf{r}) = \varepsilon_i\psi_i(\mathbf{r}), \quad i = 1, 2, \dots, \quad \mathcal{H}[\rho] = -\frac{1}{2}\nabla^2 + V_{\text{eff}}[\rho], \quad (1)$$

where $\mathcal{H}[\rho]$ is the Kohn-Sham Hamiltonian, $\rho = \rho(\mathbf{r})$ is the electron density, ε_i are the Kohn-Sham eigenvalues, and $\psi_i(\mathbf{r})$ are the Kohn-Sham eigenfunctions, also referred to as the Kohn-Sham wavefunctions. Here, we restrict ourselves to a spin-independent formulation on non-periodic systems, but the general ideas presented in this work can be extended to a spin-dependent formulation and periodic geometries in a straightforward manner. The effective Kohn-Sham potential has the form,

$$V_{eff}[\rho](\mathbf{r}) = V_H[\rho](\mathbf{r}) + V_{ext}(\mathbf{r}) + V_{xc}[\rho](\mathbf{r}), \quad (2)$$

where the first two terms are the Hartree potential due to the electron density and the external potential due to the N_A atomic nuclei located at $\mathbf{R}_j$ with charges Z_j , respectively,

$$V_H[\rho](\mathbf{r}) = \int \frac{\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}', \quad V_{ext}(\mathbf{r}) = \sum_{j=1}^{N_A} \frac{-Z_j}{|\mathbf{r} - \mathbf{R}_j|}, \quad (3)$$

and the third term is the exchange-correlation potential $V_{xc}[\rho] = \partial E_{xc}[\rho] / \partial \rho$ depending on the exchange-correlation energy $E_{xc}[\rho]$. The electron density depends on the eigenvalues and wavefunctions,

$$\rho(\mathbf{r}) = 2 \sum_{i=1}^{N_w} f(\varepsilon_i, \mu_F) |\psi_i(\mathbf{r})|^2, \quad f(\varepsilon, \mu_F) = \frac{1}{e^{(\varepsilon - \mu_F)/k_B T} + 1}, \quad (4)$$

where $f(\varepsilon, \mu_F)$ is the fractional occupation computed by Fermi-Dirac statistics [12,61], with Fermi energy μ_F , Boltzmann constant k_B , and temperature T . The Fermi energy μ_F is determined from the constraint on the total number of electrons N_e ,

$$2 \sum_{i=1}^{N_w} f(\varepsilon_i, \mu_F) = N_e. \quad (5)$$

The sums in Eq. (4) and Eq. (5) run over the N_w lowest energy wavefunctions, where N_w is chosen so that the fractional occupation of any higher energy wavefunction is negligible.

The preceding equations constitute a non-linear eigenvalue problem and the standard solution method uses the Self-Consistent Field iteration (SCF) outlined in Algorithm 1. The iteration takes the atomic positions and an initial guess for the electron density as input. The output is the converged electron density and wavefunctions, from which observables are computed. The iteration starts in line 1. In line 2, at the n th step of the iteration, the effective potential $V_{eff}[\rho_{in}^{(n)}]$ is constructed from the input electron density of the current iterate by Eq. (2). In line 3, the eigenvalue problem in Eq. (1) is solved for the eigenpairs (ε_i, ψ_i) . In line 4, the Fermi energy μ_F and fractional occupations $f(\varepsilon_i, \mu_F)$ are computed. In line 5, these quantities are used to compute a new output density $\rho_{out}^{(n)}$ by Eq. (4). In line 6, the scheme checks whether the density has converged to a desired tolerance; if so, then the iteration stops and returns the latest density; otherwise a new input density $\rho_{in}^{(n+1)}$ is constructed by Anderson mixing [62] and the iteration continues. The present work follows this approach, but focuses on the solution of the eigenvalue problem (line 3), which is the most computationally intensive step in the SCF iteration, using treecode-accelerated Green Iteration (TAGI) described below.

Algorithm 1 Self-Consistent Field Iteration (SCF).

input: atomic positions and initial guess for electron density $\rho_{in}^{(0)}$
output: electron density $\rho_{out}^{(n)}$ and Kohn Sham wavefunctions $\psi_i^{(n)}, i = 1, \dots, N_w$
1: for $n = 0, 1, 2, \dots$
2: given $\rho_{in}^{(n)}$, construct effective potential $V_{eff}[\rho_{in}^{(n)}]$ by Eq. (2)
3: using $V_{eff}[\rho_{in}^{(n)}]$, solve eigenvalue problem $\mathcal{H}[\rho_{in}^{(n)}]\psi_i^{(n)} = \varepsilon_i^{(n)}\psi_i^{(n)}, i = 1, \dots, N_w$
4: using $\varepsilon_i^{(n)}$, compute Fermi energy μ_F and fractional occupations $f(\varepsilon_i^{(n)}, \mu_F)$ by Eq. (5)
5: using $f(\varepsilon_i^{(n)}, \mu_F), \psi_i^{(n)}$, construct new density $\rho_{out}^{(n)}$ by Eq. (4)
6: if $\|\rho_{out}^{(n)} - \rho_{in}^{(n)}\|_2 < tol_{scf}$, return $\rho_{out}^{(n)}$
7: else construct new density $\rho_{in}^{(n+1)}$ by Anderson mixing and return to step 2

Having obtained the converged $\varepsilon_i, \psi_i(\mathbf{r}), \rho(\mathbf{r})$, the ground-state energy of the system is

$$E = E_{kin} + E_{xc} + E_H + E_{ext} + E_{ZZ}. \quad (6)$$

In this expression, the first two terms are the kinetic energy and exchange-correlation energy, respectively,

$$E_{kin} = \sum_{i=1}^{N_w} \int \psi_i(\mathbf{r}) \left(-\frac{1}{2} \nabla^2 \right) \psi_i(\mathbf{r}) d\mathbf{r}, \quad E_{xc}[\rho] = \int \varepsilon_{xc}[\rho](\mathbf{r}) \rho(\mathbf{r}) d\mathbf{r}, \quad (7)$$

where $\varepsilon_{xc}[\rho](\mathbf{r})$ is the exchange-correlation energy per electron for the chosen DFT functional, and the remaining three terms are the Hartree energy, external electrostatic energy, and nuclear repulsion energy, respectively,

$$E_H[\rho] = \frac{1}{2} \int V_H[\rho](\mathbf{r}) \rho(\mathbf{r}) d\mathbf{r}, \quad E_{ext}[\rho] = \int V_{ext}(\mathbf{r}) \rho(\mathbf{r}) d\mathbf{r}, \quad E_{ZZ} = \frac{1}{2} \sum_{i,j \neq i} \frac{Z_i Z_j}{|\mathbf{R}_i - \mathbf{R}_j|}. \quad (8)$$

This work employs the Local Density Approximation (LDA) [63,64] for $V_{xc}[\rho]$, $\varepsilon_{xc}[\rho]$ which are computed using the *Libxc* package [65,66].

3. Solution of eigenvalue problem by Green iteration

Several methods are available for solving the eigenvalue problem in each step of the SCF iteration (line 3 in Algorithm 1). Among real-space methods, finite-difference [4,5,67–69] and finite-element [28,31] methods represent the differential operator as a sparse matrix and use iterative techniques to compute the eigenpairs (ε_i, ψ_i) . By contrast, in this work the differential equation is converted into an integral equation by convolution with the modified Helmholtz Green's function [46], and then an iterative technique called Green Iteration is applied to obtain the eigenpairs [48,50,52,53]. We describe these steps below.

Following Kalos [46], the Kohn-Sham equations in Eq. (1) are rewritten in the form

$$\left(\frac{1}{2} \nabla^2 + \varepsilon_i \right) \psi_i = V_{eff}[\rho] \psi_i, \quad (9)$$

where ρ is the electron density for a given SCF iteration. Since the bound state eigenvalues of the Kohn-Sham Hamiltonian are negative, $\varepsilon_i < 0$, Eq. (9) is a modified Helmholtz equation with Green's function,

$$G_{\varepsilon_i}(\mathbf{r}, \mathbf{r}') = -\frac{e^{-\sqrt{-2\varepsilon_i}|\mathbf{r}-\mathbf{r}'|}}{2\pi|\mathbf{r}-\mathbf{r}'|}, \quad (10)$$

where free-space boundary conditions are assumed. Then convolution with Eq. (9) yields the integral form of the Kohn-Sham eigenvalue problem,

$$\psi_i(\mathbf{r}) = \mathcal{G}(\varepsilon_i) \psi_i(\mathbf{r}), \quad i = 1, \dots, N_w, \quad (11)$$

where

$$\mathcal{G}(\varepsilon) \psi(\mathbf{r}) = \int G_{\varepsilon}(\mathbf{r}, \mathbf{r}') V_{eff}[\rho](\mathbf{r}') \psi(\mathbf{r}') d\mathbf{r}', \quad (12)$$

defines a 1-parameter family of linear integral operators.

Note that Eq. (11) can be viewed as a fixed-point problem and this motivates the solution method called Green Iteration described in Algorithm 2. The scheme takes as input the effective potential $V_{eff}[\rho]$ for the current SCF and an initial guess for the eigenpairs $(\varepsilon_i^{(0)}, \psi_i^{(0)})$, $i = 1, \dots, N_w$, and provides the converged eigenpairs (ε_i, ψ_i) as output. Line 1 is the outer loop over wavefunctions and line 2 is the iteration for a given wavefunction. Line 3 applies the integral operator $\mathcal{G}(\varepsilon_i^{(n)})$ to the current wavefunction $\psi_i^{(n)}$. Line 4 updates the eigenvalue; several methods are available and we compare some of them below. Line 5 is the deflation step that orthogonalizes the new wavefunction $\psi_i^{(n+1)}$ against the previously converged wavefunctions, and line 6 normalizes it. Line 7 checks for convergence; if the tolerance is satisfied, then the eigenpair is stored and the process returns to line 1; otherwise the iteration in line 2 continues.

Algorithm 2 Green Iteration.

input: effective potential $V_{eff}[\rho]$ for current SCF

input: initial guess for eigenpairs $(\varepsilon_i^{(0)}, \psi_i^{(0)})$, $i = 1, \dots, N_w$

output: eigenpairs (ε_i, ψ_i) , $i = 1, \dots, N_w$

```

1: for  $i = 1, 2, \dots, N_w$ 
2:   for  $n = 0, 1, 2, \dots$ 
3:     compute  $\psi_i^{(n+1)} = \mathcal{G}(\varepsilon_i^{(n)}) \psi_i^{(n)}$ 
4:     update eigenvalue  $\varepsilon_i^{(n+1)}$ 
5:     orthogonalize  $\psi_i^{(n+1)}$  against previously converged wavefunctions  $\psi_j$ ,  $j < i$ 
6:     normalize  $\psi_i^{(n+1)}$ 
7:     if  $\|\psi_i^{(n+1)} - \psi_i^{(n)}\|_2 < tol_{gi}$  set  $(\varepsilon_i, \psi_i) = (\varepsilon_i^{(n)}, \psi_i^{(n)})$  and return to line 1
8:   else return to line 2 and continue iteration

```

4. Spatial discretization techniques

This section focuses on the spatial discretization techniques used in TAGI. These include the initialization scheme for the electron density and wavefunctions, the quadrature and *a-priori* adaptive mesh refinement techniques, the singularity subtraction schemes used to evaluate the convolution integrals, and the gradient-free approach used to update the eigenvalues. The section concludes by demonstrating the effect of these techniques using the carbon monoxide molecule as an example. Note that Hartree atomic units are used and a table containing the physical and numerical parameters is provided in Appendix A.

4.1. Initial electron density and eigenpairs

Common to many other approaches, TAGI uses a superposition of 1-atom densities to construct the initial guess in the first step of the SCF iteration. In particular, the SCF iteration uses an initial guess for the electron density of the form,

$$\rho^{(0)}(\mathbf{r}) = \sum_{j=1}^{N_a} \rho_j(|\mathbf{r} - \mathbf{R}_j|), \quad (13)$$

where $\rho_j(|\mathbf{r} - \mathbf{R}_j|)$ is a radial 1-atom electron density associated with the j th atom. These 1-atom densities are precomputed by solving a radial version of the Kohn-Sham problem for each atomic species. In addition, Green Iteration requires an initial guess for the eigenpairs, $(\varepsilon_i^{(0)}, \psi_i^{(0)}(\mathbf{r}))$, $i = 1, \dots, N_w$. The number of wavefunctions N_w is determined as follows. Since each wavefunction is occupied up to two electrons, there is a lower bound, $N_w \geq N_e/2$, however there is no sharp upper bound. In practice N_w should be chosen large enough to accommodate all states with significant fractional occupation $f(\varepsilon_i, \mu_F)$. To this end, N_w is initialized to be larger than $N_e/2$, and upon obtaining the eigenpairs, if the fractional occupation of the highest state is negligibly small, then N_w is considered large enough; otherwise, N_w is increased and the process is repeated until the check is satisfied. The initial guess for the eigenpairs depends on whether or not this is the first step in the SCF iteration. In the first step, the wavefunctions are initialized using 1-atom wavefunctions obtained in the radial solve for the initial electron density, multiplied by appropriate spherical harmonics, and the initial eigenvalues are computed by the Rayleigh quotient, $\varepsilon_i^{(0)} = \langle \psi_i^{(0)}, \mathcal{H} \psi_i^{(0)} \rangle$. In subsequent steps of the SCF iteration, the eigenpairs of the previous step are taken as the initial guess.

4.2. Spatial discretization and quadrature schemes

The energy integrals and convolution integrals will be evaluated on a set of cuboid cells representing a bounded computational domain. Using the Hartree energy as an example,

$$E_H = \frac{1}{2} \int V_H(\mathbf{r}) \rho(\mathbf{r}) d\mathbf{r} \approx \frac{1}{2} \sum_{i=1}^{N_c} \int_{C_i} V_H(\mathbf{r}) \rho(\mathbf{r}) d\mathbf{r} \approx \frac{1}{2} \sum_{i=1}^{N_c} \sum_{j=1}^{(p+1)^3} V_H(\mathbf{r}_{ij}) \rho(\mathbf{r}_{ij}) w_{ij}, \quad (14)$$

where N_c is the number of cells, $(p+1)^3$ is the number of quadrature points in each cell, indices i, j refer to quadrature point j in cell i , and w_{ij} are the quadrature weights. The total number of mesh points is denoted by $N_m = (p+1)^3 N_c$. We note that the wavefunctions, electron density, and potential are represented directly on the mesh points. The quadrature scheme uses Chebyshev points of the first kind; on the interval $[-1, 1]$ these are given by

$$x_i = \cos \theta_i, \quad \theta_i = \frac{(2i+1)\pi}{2p+2}, \quad i = 0 : p. \quad (15)$$

A tensor product grid of $(p+1)^3$ Chebyshev points is adapted to each cell; Fig. 1 shows a 2D schematic. Note that the Chebyshev points lie entirely inside the cell and never coincide with a vertex; as explained below this is important because the cells are chosen so that the atoms are located at cell vertices, thereby avoiding the singularity of the nuclear potential. Within each cell, the integrals are evaluated using the Fejér (or “classical” Clenshaw-Curtis) quadrature rule [55,56] with quadrature weights w_{ij} . The $p+1$ point Fejér quadrature rule integrates p th-degree polynomials exactly, so we refer to this as a p th-order quadrature rule.

The cells are defined using an *a-priori* adaptive refinement scheme illustrated in Fig. 2 for a 1-atom example. The goal of the scheme is to produce cells that resolve the regions with significant electron density and wavefunction variation, primarily near the atoms. Level 0 is a large cube surrounding the atoms in the system, with dimensions chosen to ensure that the electron density and wavefunctions are sufficiently small at the boundary. The cube is refined by bisecting it in the three coordinate directions, resulting in eight child cells. Several levels of uniform refinement are performed, and subsequent refinement is done adaptively in the following manner. Given a cell C , we temporarily create the child cells C_i , $i = 1 : 8$, and check the following criterion,




Fig. 1. A tensor product grid of Chebyshev points of the first kind in Eq. (15) with $p = 4$ in a 2D cell.




Fig. 2. Illustration of the refinement scheme for a 1-atom system with the atom located at (•). Four levels of refinement are shown where the final refinement level puts the atom at a cell vertex.

$$\left| \int_C t(\mathbf{r}) d\mathbf{r} - \sum_{i=1}^8 \int_{C_i} t(\mathbf{r}) d\mathbf{r} \right| < tol_m, \quad (16)$$

where $t(\mathbf{r})$ is a test function specified below and tol_m is a user-specified tolerance. The integrals in Eq. (16) are evaluated using the Fejér rule. If Eq. (16) is satisfied, then refinement is not needed and the child cells are discarded; otherwise the child cells are retained and the process continues. Fig. 2 shows the schematic of a possible outcome where the initial cell is refined at level 1, but only the child cell containing the atom is refined at level 2. Once the tolerance is satisfied for every cell, a final refinement step occurs; those cells containing an atom are subdivided so that the atoms lie at cell vertices; this ensures that the Chebyshev grid points never coincide with an atom position and hence the fields (effective potential, wavefunctions, electron density) are smooth on the interior of the cells. If the refinement scheme creates any cells with large aspect ratio, these cells are refined along their longest dimension.

Several options for the refinement test function were considered and we decided to use

$$t(\mathbf{r}) = \sqrt{\rho^{(0)}(\mathbf{r})} V_{ext}(\mathbf{r}), \quad (17)$$

where $\rho^{(0)}(\mathbf{r})$ is the initial electron density in Eq. (13) and $V_{ext}(\mathbf{r})$ is the external potential in Eq. (3). This choice is motivated by several considerations. First, it resembles the function $\psi(\mathbf{r})V_{eff}(\mathbf{r})$ appearing in the integral form of the Kohn-Sham equations (12); this is because near a nucleus, $\sqrt{\rho^{(0)}(\mathbf{r})}$ has the characteristics of an s -orbital atomic wavefunction, capturing the cusp and decay rate, and although $V_{eff}(\mathbf{r})$ is not known, $V_{ext}(\mathbf{r})$ is known and contains the Coulomb singularities that must be resolved. Second, this test function is accessible at the start of the computation and can be evaluated at arbitrary grid points as needed in the refinement scheme, suitable for the *a-priori* adaptive refinement scheme. Fig. 3 shows an example of coarse and fine meshes for the benzene molecule (C_6H_6) obtained using the refinement scheme described above with 4th order quadrature. The molecule lies in the $z = 0$ plane and a truncated portion of the mesh in that plane is shown. The coarse mesh is generated with $tol_m = 1e-4$ and the fine mesh with $tol_m = 3e-6$. The resulting cell density is highest near the twelve nuclei, and the carbon atoms are more highly refined than the hydrogen atoms, as expected since the test function $\sqrt{\rho^{(0)}(\mathbf{r})}V_{ext}(\mathbf{r})$ grows faster at heavier nuclei. Compared to a variety of other refinement schemes we considered, this approach gave the best combination of accuracy and efficiency. Further below we will demonstrate convergence with respect to both the order of the quadrature rule p and the mesh tolerance parameter tol_m .

The *a-priori* adaptive mesh refinement scheme described above enables TAGI to achieve chemical accuracy for all-electron calculations using only the external potential and initial guess for the electron density, both of which are available at the start of the calculation. Alternatively, it would be interesting to investigate *a-posteriori* adaptive mesh refinement methods, such as those used in the finite element community [70,31], to estimate local errors and improve the efficiency of TAGI's mesh refinement scheme.




Fig. 3. Example of the mesh refinement scheme for the benzene molecule (C_6H_6). 2D slices of the mesh are shown in the plane of the molecule generated with 4th order quadrature in Eq. (17) and (a) $tol_m = 1e-4$, (b) $tol_m = 3e-6$.

4.3. Singularity subtraction

Achieving the necessary accuracy for DFT calculations requires careful treatment of the singular integrals arising in Green Iteration,

$$\psi^{(n+1)}(\mathbf{r}) = - \int V_{eff}(\mathbf{r}') \psi^{(n)}(\mathbf{r}') \frac{e^{-\sqrt{-2\varepsilon^{(n)}}|\mathbf{r}-\mathbf{r}'|}}{2\pi|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}', \quad (18)$$

and the Hartree potential,

$$V_H(\mathbf{r}) = \int \frac{\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}'. \quad (19)$$

The singular $\mathbf{r}' = \mathbf{r}$ term in the quadrature sums is skipped. This error due to skipping the singularity is reduced by weakening the singularities before discretization. For the integral involving the Yukawa kernel in Eq. (18) we implemented a standard singularity subtraction scheme [57,58],

$$\int f(\mathbf{r}') \frac{e^{-k|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' = \int (f(\mathbf{r}') - f(\mathbf{r})) \frac{e^{-k|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' + f(\mathbf{r}) \int \frac{e^{-k|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}'. \quad (20)$$

The second term on the right in Eq. (20) is evaluated analytically,

$$f(\mathbf{r}) \int \frac{e^{-k|\mathbf{r}-\mathbf{r}'|}}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' = \frac{4\pi f(\mathbf{r})}{k^2}, \quad (21)$$

while the singularity in the first term on the right has been weakened, so the quadrature scheme yields a more accurate result. Note however that the exponential decay rate in the Yukawa kernel is $k = \sqrt{-2\varepsilon}$, and a problem arises if $\varepsilon \rightarrow 0$, since in that case the singularity subtraction scheme in Eq. (20) tends to the indeterminate form $\infty - \infty$. This is resolved by introducing a constant gauge shift in the effective potential,

$$V_{eff}(\mathbf{r}) \rightarrow V_{eff}(\mathbf{r}) + V_{shift}. \quad (22)$$

The wavefunctions are unaffected and the eigenvalues simply shift by this amount (the shift is removed before computing energies). Throughout this work we set $V_{shift} = -0.5$, ensuring that the eigenvalues of the occupied states are bounded away from zero.

The scheme described above however does not work for the Hartree potential in Eq. (19) which corresponds $k = 0$. In this case, we employ a modified form of singularity subtraction using a Gaussian function,

$$\int f(\mathbf{r}') \frac{1}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' = \int (f(\mathbf{r}') - f(\mathbf{r}) e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}) \frac{1}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' + f(\mathbf{r}) \int \frac{e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}', \quad (23)$$

where α is a scaling parameter. As before, the second term on the right is evaluated analytically,

$$f(\mathbf{r}) \int \frac{e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}' = 2\pi f(\mathbf{r})\alpha^2, \quad (24)$$

and the singularity in the first term has been weakened. In addition, the Gaussian remains smooth for $\mathbf{r}' \rightarrow \mathbf{r}$, unlike other options, and this ensures the accuracy of the quadrature scheme.

The choice of the scaling factor α is guided by the following considerations. Recall that the first integral on the right in Eq. (23) is computed using a quadrature scheme on a truncated computational domain, and hence the function $f(\mathbf{r})e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}$ should have certain properties. If α is small, then the Gaussian is narrow and the quadrature scheme would struggle to resolve the variation in this function. On the other hand if α is large, then the Gaussian is wide and this would require increasing the size of the computational domain. In more detail, the function $f(\mathbf{r})e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}$ must be sufficiently small when $\mathbf{r}'$ is near the domain boundary to ensure that the effect of the domain truncation is small; there are two cases, (1) when $\mathbf{r}$ lies in the domain interior, then $f(\mathbf{r})$ is not necessarily small, but $e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}$ is small as long as α is not too large, (2) when $\mathbf{r}$ lies near the domain boundary, then $f(\mathbf{r})$ is small while $e^{-|\mathbf{r}-\mathbf{r}'|^2/\alpha^2}$ is bounded. The conclusion is that the Gaussian scaling factor α should not be too small in relation to the spatial discretization and should not be too large in relation to the computational domain size; this work uses domains of size $[-20, 20]^3$ a.u. to $[-30, 30]^3$ a.u. with $\alpha = 1$ a.u., which was determined empirically.

4.4. Gradient-free eigenvalue update

Recall that line 4 in Green Iteration updates the eigenvalue $\varepsilon_i^{(n+1)}$; in this subsection we describe three methods for this purpose. The first method uses the Rayleigh quotient [50],

$$\varepsilon_i^{(n+1)} = \frac{\langle \psi_i^{(n+1)}, \mathcal{H} \psi_i^{(n+1)} \rangle}{\langle \psi_i^{(n+1)}, \psi_i^{(n+1)} \rangle} = \frac{-\frac{1}{2} \langle \psi_i^{(n+1)}, \nabla^2 \psi_i^{(n+1)} \rangle + \langle \psi_i^{(n+1)}, V_{eff} \psi_i^{(n+1)} \rangle}{\langle \psi_i^{(n+1)}, \psi_i^{(n+1)} \rangle}, \quad (25)$$

where $\mathcal{H}$ is the Kohn-Sham differential operator defined in Eq. (1), V_{eff} is the effective potential in the current SCF, and $\psi_i^{(n+1)}$ is the wavefunction computed in line 3 of Green Iteration. The second method applies integration by parts in Eq. (25) to obtain,

$$\varepsilon_i^{(n+1)} = \frac{\frac{1}{2} \langle \nabla \psi_i^{(n+1)}, \nabla \psi_i^{(n+1)} \rangle + \langle \psi_i^{(n+1)}, V_{eff} \psi_i^{(n+1)} \rangle}{\langle \psi_i^{(n+1)}, \psi_i^{(n+1)} \rangle}. \quad (26)$$

In the present framework the gradient $\nabla \psi_i^{(n+1)}$ in Eq. (26) and Laplacian $\nabla^2 \psi_i^{(n+1)}$ in Eq. (25) are computed by spectral differentiation using the values of the wavefunction at the Chebyshev points in each cell [71]. The third method is a gradient-free update suggested by Harrison et al. [48],

$$\varepsilon_i^{(n+1)} = \varepsilon_i^{(n)} - \frac{\langle V_{eff} \psi_i^{(n)}, \psi_i^{(n)} - \psi_i^{(n+1)} \rangle}{\langle \psi_i^{(n+1)}, \psi_i^{(n+1)} \rangle}. \quad (27)$$

In this case, which computes a $\Delta \varepsilon_i$ rather than $\varepsilon_i^{(n+1)}$ itself, the initial guess $\varepsilon_i^{(0)}$ in the first SCF iteration can be given using either Eq. (25) or Eq. (26). The gradient-free eigenvalue update enables the total energy to also be computed in a gradient-free manner using the alternative expression,

$$E = E_{band} - E_H + E_{xc} - \int \rho(\mathbf{r}) V_{xc}[\rho](\mathbf{r}) d\mathbf{r} + E_{ZZ}, \quad (28)$$

where the band energy is the weighted sum of the eigenvalues,

$$E_{band} = 2 \sum_{i=1}^{N_w} f_i(\varepsilon_i, \mu) \varepsilon_i. \quad (29)$$

In contrast to the original expression for the total energy in Eq. (6), the gradient-free expression in Eq. (28) avoids explicitly computing the kinetic energy E_{kin} in Eq. (7) which contains the Laplacian; the kinetic energy is now contained implicitly in the band energy, which is obtained with the gradient-free method. Later below we show that the gradient-free method has the best accuracy of the three approaches described here.

Table 1

Error in the total energy per atom for the carbon monoxide molecule using (a) a fixed mesh, increasing quadrature order p from 4 to 7, and (b) a fixed quadrature order $p = 4$, decreasing mesh refinement parameter tol_m from $3e-6$ to $1e-7$.

(a)				
p	tol_m	Cells	Points	Error (mHa)
4	$3e-7$	5293	661625	1.313
5	$3e-7$	5293	1143288	0.605
6	$3e-7$	5293	1815499	0.311
7	$3e-7$	5293	2710016	0.179

(b)				
p	tol_m	Cells	Points	Error (mHa)
4	$3e-6$	2962	370250	3.946
4	$1e-6$	3676	459500	2.550
4	$3e-7$	5293	661625	1.313
4	$1e-7$	7428	928500	0.674

Table 2

Error in the total energy per atom for the carbon monoxide molecule without using singularity subtraction (column 4) and with using singularity subtraction (column 5).

tol_m	Cells	Points	Error (mHa)	
			Non-SS	SS
$3e-6$	2962	370250	823	3.946
$1e-6$	3676	459500	702	2.550
$3e-7$	5293	661625	558	1.313
$1e-7$	7428	928500	429	0.674

4.5. Accuracy of computational techniques

This subsection documents the accuracy of the previously described computational techniques in the case of the carbon monoxide molecule ($N_A = 2$, $N_e = 14$, $N_w = 8$). The computations use domain $[-20, 20]^3$ a.u., temperature $T = 200$ K, gauge shift $V_{shift} = -0.5$, Green Iteration eigensolve tolerance $tol_{gi} = 1e-7$, SCF tolerance $tol_{scf} = 1e-6$, and Anderson mixing parameter $\beta = 0.5$. Except where specified, the computations use singularity subtraction and the gradient-free eigenvalue update. We report the energy error $|E_{TAGI} - E_{ref}|$, where E_{TAGI} is computed using TAGI and $E_{ref} = -112.47193$ Ha is the reference value converged to $1e-4$ Ha, which was computed using DFT-FE [28,31].

4.5.1. Quadrature rule and mesh refinement scheme

We first demonstrate the effect of the quadrature rule order p and the mesh refinement tolerance tol_m described in section 4.2. To test the effect of the quadrature rule order we generate a mesh using order $p = 4$ and tolerance $tol_m = 3e-7$, and then on this mesh the order p is varied; Table 1a shows that the error is reduced from 1.313 mHa with $p = 4$ to 0.179 mHa with $p = 7$. To test the effect of the mesh refinement tolerance we fix the quadrature order to $p = 4$ and vary the mesh refinement tolerance tol_m ; Table 1b shows that the error is reduced from 3.946 mHa with $tol_m = 3e-6$ to 0.674 mHa with $tol_m = 1e-7$.

4.5.2. Singularity subtraction

Next we demonstrate the effect of the singularity subtraction schemes described in section 4.3. The quadrature order is set to $p = 4$ and a sequence of mesh refinements is performed. The ground state calculation is performed with and without singularity subtraction; in both cases the singular term in the discrete convolution sums is skipped. Table 2 shows that singularity subtraction yields a significant improvement in the accuracy of the total energy, over two orders of magnitude for mesh size $N_m = 661625$ which achieves chemical accuracy.

4.5.3. Eigenvalue updates

Finally, we compare the eigenvalue update methods described in section 4.4, Laplacian update (Eq. (25)), gradient update (Eq. (26)), and gradient-free update (Eq. (27)). The ground state calculation is performed for a sequence of refined meshes. Fig. 4 shows the total energy error per atom versus the number of mesh points N_m for (a) order $p = 4$ and (b) order $p = 6$. We make the following three observations. First, for a given mesh, the gradient-free update is significantly more accurate than the Laplacian and gradient updates. An exception occurs for the coarsest mesh with order $p = 4$, where the gradient update has slightly lower error than the gradient-free update due to a cancellation of errors. In particular, we found that using the gradient update, some of the computed eigenvalues have positive error, while others have negative error, leading to a cancellation of errors in the band energy and total energy, which ceases as the mesh is refined or the quadrature order




Fig. 4. Error in the total energy per atom for the carbon monoxide molecule versus the number of mesh points N_m while using different eigenvalue update methods in Green Iteration for (a) quadrature order $p = 4$ and (b) quadrature order $p = 6$.




Fig. 5. Convergence of the eigenfunction residual $\|\psi_i^{(n+1)} - \psi_i^{(n)}\|_2$ during Green Iteration in the first SCF iteration for the carbon monoxide molecule. The observed residuals (symbols) and the predicted convergence rates r_i (black lines, explained below). (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

is increased. Second, for $p = 4$, the gradient update error saturates around 6 mHa/atom as the mesh is refined. We note that the mesh refinement scheme as described in section 4.2 uses feedback from the gradient-free update, and hence it is not guaranteed to refine cells having large error in the gradient terms. Third, for $p = 6$, the gradient update recovers its convergence as the mesh is refined, although it is still less accurate than the gradient-free update.

5. Convergence rate of Green iteration

Previously we explained how in each step of the SCF iteration, the Kohn-Sham eigenproblem in Eq. (1) can be converted into a fixed-point problem for the integral operator in Eq. (11) and that the fixed-point problem is solved by Green Iteration. This section examines the convergence rate of Green Iteration; first an example exhibiting slow convergence is presented, then the cause of the problem is identified by reference to power iteration, and finally Anderson mixing is applied to the wavefunctions to accelerate convergence.

5.1. Slow convergence of Green iteration

To illustrate the slow convergence of Green Iteration, we consider the first SCF iteration for the carbon monoxide molecule. Fig. 5 plots the residual of the first seven wavefunctions determined by Green Iteration versus the iteration number. In this case the first two wavefunctions converge rapidly, but the subsequent wavefunctions converge slowly; in particular the 4th wavefunction converges extremely slowly. The result is that Green Iteration requires a total of 1246 iterations to ensure that the first seven wavefunction residuals fall below $1e-8$. This is a tighter tolerance than is used in practice, however it helps illustrate the issue. In the next subsection we examine the cause of this slow convergence, as well as how the predicted rates (black lines) are computed.




Fig. 6. First SCF iteration for the carbon monoxide molecule. The curves $\mu_i(\varepsilon)$ defined by eigenvalue problem Eq. (30) for the integral operator $\mathcal{G}(\varepsilon)$ are plotted versus parameter ε . Intersections with the dashed line $\mu = 1$ yield fixed-points ε_i of Green Iteration. The spectral gap of the integral operator, $\Delta\mu_i = 1 - \mu_{i+1}(\varepsilon_i)$, is indicated at the fixed-points ε_i , (a) $i = 1, 2, 3$, (b) $i = 3, 4, 5$, (c) $i = 5, 6, 7, 8$. Numerical values are given in Table 3.

5.2. Convergence analysis

Recall line 3 of Green Iteration (Algorithm 2), $\psi_i^{(n+1)} = \mathcal{G}(\varepsilon_i^{(n)})\psi_i^{(n)}$, which updates the i th eigenfunction using the operator defined in Eq. (12). The parameter $\varepsilon_i^{(n)}$ changes in each step of the iteration, but as $\varepsilon_i^{(n)} \rightarrow \varepsilon_i$, the scheme converges to power iteration for the operator $\mathcal{G}(\varepsilon_i)$ with deflation against the previously determined eigenfunctions ψ_j , $j < i$ as indicated in line 5 of the algorithm. This suggests that the convergence rate of $\psi_i^{(n)}$ depends on the spectral gap of $\mathcal{G}(\varepsilon_i)$ [72]. To demonstrate this it is useful to define a 1-parameter family of curves $\mu_i(\varepsilon)$ and functions $\phi_i(\varepsilon)$ satisfying the linear eigenvalue equation,

$$\mathcal{G}(\varepsilon)\phi_i(\varepsilon) = \mu_i(\varepsilon)\phi_i(\varepsilon), \quad i = 1, \dots, N_w, \quad (30)$$

subject to the following conditions. For each ε , the eigenvalues are ordered by their magnitude $\mu_1(\varepsilon) \geq \dots \geq \mu_{N_w}(\varepsilon)$. Note that if $\mu_i(\varepsilon) = 1$ for some index i and parameter value ε , then Eq. (30) reduces to the fixed-point problem in Eq. (11), $\psi_i = \mathcal{G}(\varepsilon_i)\psi_i$, in which case we have $\varepsilon = \varepsilon_i$ and $\phi_i(\varepsilon) = \psi_i$ [50,52]. In addition, the usual orthogonality condition, $\phi_i(\varepsilon) \perp \phi_j(\varepsilon)$ for $i \neq j$, is modified to be consistent with the deflation step in Green Iteration; that is, $\psi_i \perp \phi_j(\varepsilon)$ for $i < j$ and $\varepsilon_i < \varepsilon$.

Fig. 6 illustrates this for the first SCF iteration of the carbon monoxide molecule, where the curves $\mu_i(\varepsilon)$ are plotted versus ε for $i = 1 : 8$. Note that for each parameter value ε , the eigenvalues $\mu_i(\varepsilon)$ are computed by power iteration applied to the operator $\mathcal{G}(\varepsilon)$ subject to the modified orthogonality condition stated above. The fixed-points of Green Iteration occur when one of the curves $\mu_i(\varepsilon)$ intersects the line $\mu = 1$, and the plotted curves terminate there because there are no eigenfunctions with $\mu > 1$ due to the orthogonality condition.

Fig. 6 also indicates the spectral gap of the operator $\mathcal{G}(\varepsilon_i)$, defined by $\Delta\mu_i = \mu_i(\varepsilon_i) - \mu_{i+1}(\varepsilon_i) = 1 - \mu_{i+1}(\varepsilon_i)$; due to the continuity of the curves $\mu_i(\varepsilon)$, these are correlated with the spectral gap of the Hamiltonian, defined by $\Delta\varepsilon_i = \varepsilon_{i+1} - \varepsilon_i$; hence both gaps are relatively large for $i = 1, 2$ in Fig. 6a, and relatively small for $i = 3, \dots, 7$ in Fig. 6b,c. Note further that the spectrum of the CO molecule contains a degeneracy; $\varepsilon_5 = \varepsilon_6$, hence ψ_5 and ψ_6 span a degenerate subspace. This degeneracy manifests itself in the spectral analysis in several ways. First, ψ_5 and ψ_6 converge with identical rates in Green Iteration (Fig. 5 parallel purple and black), and second, the $\mu_5(\varepsilon)$ and $\mu_6(\varepsilon)$ curves are identical (Fig. 6c overlapping purple and black). In the case of a degeneracy, the convergence rate of the wavefunctions to the degenerate subspace is governed by the spectral gap to the next distinct eigenvalue. In this example, we define the spectral gaps $\Delta\mu_5$ and $\Delta\mu_6$ with respect to the 7th eigenvalue, $\Delta\mu_5 = 1 - \mu_7(\varepsilon_5)$ and $\Delta\mu_6 = 1 - \mu_7(\varepsilon_6)$. Finally, note that at a fixed-point parameter ε_i , the largest eigenvalue of $\mathcal{G}(\varepsilon_i)$ is $\mu_i = 1$, so the convergence rate of power iteration is $r_i = \mu_{i+1}/\mu_i = (\mu_i - \Delta\mu_i)/\mu_i = 1 - \Delta\mu_i$; hence a large gap $\Delta\mu_i$ leads to rapid convergence of ψ_i and a small gap $\Delta\mu_i$ leads to slow convergence.

Table 3 gives the fixed-points ε_i , spectral gaps $\Delta\varepsilon_i$ and $\Delta\mu_i$, observed and predicted convergence rates r_i , error in the predicted rates, and number of iterations required to achieve the $1e-8$ tolerance in Green Iteration. For each curve in Fig. 5, we observe a rapid initial decay of the residual to roughly $1e-2$ or $1e-3$, followed by geometric convergence down to $1e-8$, $\|\psi_i^{(n+1)} - \psi_i^{(n)}\|_2 \sim r_i^n \|\psi_i^{(1)} - \psi_i^{(0)}\|_2$, where n is the iteration index and r_i is the geometric convergence rate. There are two versions of r_i in Table 3; the observed r_i were obtained from the actual Green Iteration, and the predicted $r_i = \mu_{i+1}/\mu_i = (\mu_i - \Delta\mu_i)/\mu_i = 1 - \Delta\mu_i$ were obtained from the spectrum of the operator in Eq. (30). Reference lines with the predicted rate r_i appear in Fig. 5 adjacent to the corresponding residual curve. In several cases (ψ_1 , ψ_2 , ψ_7), the observed convergence rate is faster than the predicted rate; this is attributed to the iteration not entering the asymptotic power-iteration regime before the tolerance was met. In the slower converging cases (ψ_3 , ψ_4 , ψ_5 , ψ_6), the predicted convergence rates accurately agree with the observed rates, with percent errors 0.115%, 0.082%, 0.053%, and 0.053%. This confirms that the convergence rate of the eigenfunction ψ_i in Green Iteration is controlled by the spectral gap $\Delta\mu_i$ in the integral operator, which in turn is correlated with the spectral gap $\Delta\varepsilon_i$ in the differential operator. Hence, Green Iteration may converge slowly whenever a small spectral gap exists in the Hamiltonian; the next subsection applies a mixing method to overcome this limitation.

Table 3

First SCF iteration for the carbon monoxide molecule. Eigenvalue index, Hamiltonian eigenvalues ε_i , Hamiltonian spectral gap $\Delta\varepsilon_i$, integral operator spectral gap $\Delta\mu_i$, observed convergence rate, predicted convergence rate, accuracy of the prediction, number of iterations for the wavefunction to converge to $1e-8$ tolerance.

index, i	Spectral Gaps			Convergence Rates			Number of Iterations
	ε_i	$\Delta\varepsilon_i$	$\Delta\mu_i$	observed r_i	predicted r_i	% error	
1	-18.870	8.862	0.221	0.460	0.779	69.3	19
2	-10.008	8.822	0.491	0.450	0.509	13.1	20
3	-1.186	0.579	0.132	0.867	0.868	0.115	110
4	-0.607	0.087	0.023	0.9752	0.976	0.082	601
5	-0.520	0.124	0.061	0.9385	0.939	0.053	208
6	-0.520	0.124	0.061	0.9385	0.939	0.053	211
7	-0.396	0.262	0.086	0.819	0.914	11.6	77




Fig. 7. Convergence of the eigenfunction residual $\|\psi_i^{(n+1)} - \psi_i^{(n)}\|_2$ during Green Iteration in the first SCF iterations for the carbon monoxide molecule using wavefunction mixing with mixing parameter $\beta = 0.5$. Wavefunction mixing reduces the total number of iterations from 1246 to 188.

5.3. Wavefunction mixing

While Green Iteration resembles power iteration as noted above, it is a fixed-point iteration and hence is amenable to standard fixed-point acceleration techniques. We define the vector $\mathbf{x} = (\varepsilon, \psi)$, and the inner product between two vectors $\mathbf{x}_1 = (\varepsilon_1, \psi_1)$ and $\mathbf{x}_2 = (\varepsilon_2, \psi_2)$ to be $(\mathbf{x}_1, \mathbf{x}_2) = \varepsilon_1 \varepsilon_2 + \int \psi_1(\mathbf{r}) \psi_2(\mathbf{r}) d\mathbf{r}$. We then use Anderson mixing to update the eigenpairs $(\varepsilon_i^{(n)}, \psi_i^{(n)})$ after each step of Green Iteration, in the same way that the electron density is updated after each step of the SCF iteration. Fig. 7 shows the effect of applying Anderson mixing to the wavefunctions with mixing parameter $\beta = 0.5$, for the same computation as above, the first SCF iteration of the carbon monoxide molecule. The total number of iterations is reduced from 1246 (Green Iteration) to 188 (Green Iteration with wavefunction mixing). $\psi_3 - \psi_6$ still converge the slowest, however they converge significantly faster than without Anderson mixing.

In practice the wavefunction mixing scheme requires a good initial guess to ensure convergence. In the first SCF iteration, to achieve a good initial guess, Green Iteration can be performed without wavefunction mixing until convergence to a user-defined tolerance is achieved, at which point the computed wavefunction is in the basin of attraction of the fixed-point scheme and Anderson wavefunction mixing can be safely applied. In subsequent SCF iterations the initial guess for the eigenpairs tend to be much better and delaying the use of wavefunction mixing is not necessary. Furthermore, the tolerance for Green Iteration tol_{gi} does not have to be the same throughout an SCF iteration. We find that starting with a loose tolerance and gradually tightening it after each step in the SCF is beneficial. The gradual reduction of tol_{gi} increases the number of steps in the SCF for the electron density to converge to tol_{scf} , but it significantly reduces the cost of the first few steps of the SCF iteration, and results in an overall reduction of computation time.

6. Treecode acceleration

The scheme described above requires computing convolutions of the form,

$$u(\mathbf{r}) = \int G(\mathbf{r}, \mathbf{r}') f(\mathbf{r}') d\mathbf{r}', \quad (31)$$

where $G(\mathbf{r}, \mathbf{r}')$ is either the Coulomb kernel in the Hartree potential in Eq. (3), or the Yukawa kernel in the integral operator $\mathcal{G}(\varepsilon)$ in Eq. (12) needed in line 3 of Green Iteration. Upon discretization, the integral in Eq. (31) is approximated by the discrete convolution sum,

$$u_i = \sum_{\substack{j=1 \\ i \neq j}}^{N_m} G(\mathbf{r}_i, \mathbf{r}_j) f_j w_j, \quad i = 1, \dots, N_m, \quad (32)$$

where $u_i \approx u(\mathbf{r}_i)$, $f_j = f(\mathbf{r}_j)$, and w_j are the quadrature weights. Computing u_i by direct summation requires $O(N_m^2)$ operations, and several methods have been developed to reduce the cost including the treecode [73] and fast multipole method [74]. This work employs a recently developed barycentric Lagrange treecode (BLTC) [59] which reduces the operation count to $O(N_m \log N_m)$ using barycentric Lagrange interpolation [75]. We use the BLTC because it has several desirable features; it is efficient in the high accuracy regime demanded by the DFT application, its kernel-independent property facilitates the implementation of the singularity subtraction schemes, the barycentric Lagrange approximation it uses is well suited for GPU acceleration, and the code is publicly available. For clarity of presentation, the singularity subtraction schemes from Eqs. (20) and (23) have been omitted in Eq. (32), but they are easily accommodated in the BLTC and are used in this work. Throughout this section the points $\mathbf{r}_i$ are called target particles, the points $\mathbf{r}_j$ are called source particles, and Eq. (32) expresses the particle-particle interactions. Below we present an overview of the treecode; references can be consulted for more details [59,60,76].

6.1. Source clusters and target batches

The treecode starts by dividing the source particles into a hierarchical tree of source clusters, where the root cluster is the minimal bounding box enclosing the computational domain. The root is divided into child clusters by bisection in each dimension, and the child clusters are recursively subdivided until they contain fewer than N_L particles; these are the leaves of the tree. After division each cluster is shrunk to the minimal bounding box containing its particles. Typically a cluster is divided into eight children, but if shrinking would cause the aspect ratio to be greater than $\sqrt{2}$, the cluster is instead divided into either two or four children. Note that the source clusters in the treecode are rectangular boxes, and in general they are different than the cells in the adaptive mesh. For efficiency purposes as explained below, the target particles are also organized into a set of localized batches containing fewer than N_B particles, and then the particle-particle interactions are organized into batch-cluster interactions between the target particles in a batch and the source particles in a cluster. In this work we set $N_B = N_L$, and since the target particles and source particles correspond to the same set (the N_m quadrature points), the target batches are equivalent to the leaf source clusters in the tree.

6.2. Particle-cluster approximation by barycentric Lagrange interpolation

Note that the sum in Eq. (32) can be rewritten as

$$u_i = \sum_{\substack{j=1 \\ i \neq j}}^{N_m} G(\mathbf{r}_i, \mathbf{r}_j) f_j w_j = \sum_C u(\mathbf{r}_i, C), \quad (33)$$

where the second sum is taken over a set of source clusters C , and

$$u(\mathbf{r}_i, C) = \sum_{\mathbf{r}_j \in C} G(\mathbf{r}_i, \mathbf{r}_j) f_j w_j \quad (34)$$

is the interaction between a target particle $\mathbf{r}_i$ and a source cluster $C = \{\mathbf{r}_j\}$. Following [59], the particle-cluster interaction can be approximated by 3D polynomial interpolation of degree n ,

$$u(\mathbf{r}_i, C) \approx \sum_{k_1=0}^n \sum_{k_2=0}^n \sum_{k_3=0}^n G(\mathbf{r}_i, \mathbf{s}_{\mathbf{k}}) \hat{f}_{\mathbf{k}}, \quad \hat{f}_{\mathbf{k}} = \sum_{\mathbf{r}_j \in C} L_{k_1}(r_{j1}) L_{k_2}(r_{j2}) L_{k_3}(r_{j3}) f_j w_j, \quad (35)$$

where $\mathbf{r}_j = (r_{j1}, r_{j2}, r_{j3})$ is a source particle, $\mathbf{s}_{\mathbf{k}} = (s_{k1}, s_{k2}, s_{k3})$ is a tensor product grid of interpolation points, $L_k(t)$ are the 1D Lagrange interpolating polynomials, and $\hat{f}_{\mathbf{k}}$ are weights associated with the approximation. It is important to note that the approximation in Eq. (35) has the same direct sum structure as the exact interaction in Eq. (34); in one case the target particle $\mathbf{r}_i$ interacts with the source particles $\mathbf{r}_j$ and in the other it interacts with the interpolation points $\mathbf{s}_{\mathbf{k}}$; however in both cases the necessary kernel evaluations are independent of each other and can be efficiently computed in parallel on a GPU; this is in contrast to other fast summation schemes based on analytic series expansions, such as the Taylor treecode [76], where the approximations are recursive, which limits performance on the GPU. A further point is that the approximation weights $\hat{f}_{\mathbf{k}}$ are independent of the target particle $\mathbf{r}_i$, so they can be precomputed and reused for different targets. Next we examine the decision of when to use the approximation in Eq. (35), then return to the choice of interpolation points $\mathbf{s}_{\mathbf{k}}$ and the structure of the interpolating polynomial $L_k(t)$.




Fig. 8. 2D diagram showing a target batch interacting with a source cluster. The target batch of radius r_B is a distance R from the source cluster of radius r_C . If $(r_B + r_C)/R < \theta$ and $(n + 1)^3 < N_S$, this interaction will be approximated using Eq. (35).




Fig. 9. Example of a cluster C in 2D, (a) source particles $\mathbf{r}_j$ in C (these are quadrature points in the adaptive mesh, here defined with order $p = 3$), (b) Chebyshev grid of interpolation points $\mathbf{s}_k$ (here defined with degree $n = 4$, these represent the cluster through the particle-cluster approximation in Eq. (35)).

In this work the particle-cluster interactions $u(\mathbf{r}_i, C)$ are organized into batch-cluster interactions. For a given batch of target particles, the decision on whether to apply the approximation in Eq. (35) is controlled by the multipole acceptance criterion (MAC) which in this work has the form,

$$\frac{r_B + r_C}{R} < \theta, \quad (n + 1)^3 < N_S, \quad (36)$$

where r_B is the target batch radius, r_C is the source cluster radius, R is the distance between the target batch center and source cluster center, n is the interpolation degree, and N_S is the number of source particles in the cluster. The first part of the MAC ensures the approximation's accuracy and is diagrammed in Fig. 8, while the second part ensures its efficiency.

For the interpolation points, the BLTC uses Chebyshev points of the 2nd kind,

$$s_k = \cos \theta_k, \quad \theta_k = \pi k/n, \quad k = 0, \dots, n, \quad (37)$$

which are defined on the interval $[-1, 1]$ and are linearly mapped to clusters located elsewhere. In addition the BLTC uses the barycentric form of the Lagrange polynomials [75],

$$L_k(t) = \frac{\frac{b_k}{t - s_k}}{\sum_{k'=0}^n \frac{b_{k'}}{t - s_{k'}}}, \quad k = 0, \dots, n, \quad (38)$$

where due to the scale-invariance of this form [77], the barycentric weights are

$$b_k = (-1)^k \delta_k, \quad \delta_k = \begin{cases} 1, & k = 1, \dots, n-1, \\ 1/2, & k = 0, n. \end{cases} \quad (39)$$

Fig. 9 shows a 2D example of a cluster C comprised of seven quadrature cells; Fig. 9a shows the source particles $\mathbf{r}_j$ in C (these are quadrature points in the adaptive mesh, here defined with order $p = 3$), and Fig. 9b shows the Chebyshev grid of interpolation points $\mathbf{s}_k$ in C (here defined with degree $n = 4$, these represent the cluster through the particle-cluster approximation in Eq. (35)). The treecode has two options for computing particle-cluster interactions; the direct sum in Eq. (34) requires $O(N_S)$ operations, where N_S is the number of source particles in C , while the approximation in Eq. (35) requires $O(n^3)$ operations for interpolation of degree n ; hence the approximation is more efficient when N_S is large and n is small, as in the example in Fig. 9.




Fig. 10. Comparison of the treecode approximation error to the underlying discretization error for the carbon monoxide molecule. The mesh contains $N_m = 661625$ points and results in a discretization error of $|E_{ref} - E_{ds}|/N_A = 6.56e-4$ Ha/atom (red dashed line). Solid curves and symbols show the treecode approximation error $|E_{ds} - E_{tc}|/N_A$ for treecode MAC parameter $0.35 \leq \theta \leq 0.8$ and interpolation degree $n = 6, 8, 10$.

6.3. Treecode algorithm

The treecode is described in Algorithm 3. The input consists of the quadrature points $\mathbf{r}_i$, weights f_i, w_i , and the treecode parameters including MAC θ , interpolation degree n , maximum leaf size N_L , maximum batch size N_B . Line 1 builds the tree of source clusters and the set of target batches. Line 2 computes the approximation weights $\hat{f}_{\mathbf{k}}$ for each source cluster. Line 3 computes the batch-cluster interactions for each target batch via the recursive function COMPUTEPOTENTIAL; if the MAC in Eq. (36) is satisfied, then the approximation is computed with Eq. (35); if the MAC fails because $(r_B + r_C)/R \geq \theta$, then there are two options; if the cluster is a leaf, then the batch interacts directly with the cluster by Eq. (34), otherwise, the cluster has children and the batch interacts with each child; if the MAC fails because $(n+1)^3 \geq N_S$, then the batch interacts directly with the cluster by Eq. (34). The treecode operation count is $O(N_m \log N_m)$, where the factor N_m is the number of target particles and the factor $\log N_m$ is the number of levels in the tree.

Algorithm 3 Treecode.

input: quadrature points $\mathbf{r}_i$ and weights $f_i, w_i, i = 1, \dots, N_m$
input: treecode MAC parameter θ , interpolation degree n , maximum leaf size N_L
output: approximate potential $u_i, i = 1, \dots, N_m$

- 1: build tree of source clusters $\{C\}$ and set of target batches $\{B\}$
- 2: for each source cluster, compute approximation weights $\hat{f}_{\mathbf{k}}$ in Eq. (35)
- 3: for each target batch, COMPUTEPOTENTIAL(B , root_cluster)
- 4:
- 5: **function** COMPUTEPOTENTIAL($Batch$, $Cluster$)
- 6: if MAC is satisfied, compute batch-cluster approximation by Eq. (35)
- 7: else if $(r_B + r_C)/R \geq \theta$
- 8: if $Cluster$ is a leaf, compute batch-cluster interaction by direct sum in Eq. (34)
- 9: else for each $Child$ of $Cluster$, COMPUTEPOTENTIAL($Batch$, $Child$)
- 10: else if $(n+1)^3 \geq N_S$
- 11: compute batch-cluster interaction by direct sum in Eq. (34)
- 12: **end function**

6.4. Treecode accuracy

The sum in Eq. (32) is a discretization of the convolution integral in Eq. (31), and it is important to ensure that the treecode approximation error is less than discretization error. We document the accuracy of the treecode for the carbon monoxide molecule with domain $[-20, 20]^3$ a.u., quadrature order $p = 4$, mesh refinement tolerance $tol_m = 3e-7$, SCF tolerance $tol_{scf} = 1e-5$, and Green Iteration tolerance $tol_{gi} = 1e-6$. In this case the number of mesh points is $N_m = 661625$. We compute the ground-state energy with and without the treecode, and record the discretization error $|E_{ref} - E_{ds}|/N_A$, and treecode approximation error $|E_{ds} - E_{tc}|/N_A$, where E_{ref} is the reference energy computed by DFT-FE, and E_{ds}, E_{tc} are computed by the present method using direct summation and the treecode, respectively. Throughout this work the source clusters and target batch size parameters are set to $N_L = N_B = 2000$. Fig. 10 shows that the discretization error of the present method is $|E_{ref} - E_{ds}|/N_A = 6.56e-4$ Ha/atom (red dashed line), while the treecode approximation error is much smaller. This confirms that in this range of parameter values, the numerical errors introduced by the treecode do not upset the accuracy of the discretization.

Table 4

Treecode accuracy and acceleration for the carbon monoxide molecule using quadrature order $p = 4$ and mesh refinement parameter tol_m , giving mesh size N_m , Hartree energy E_H (Ha) in Eq. (8) for electron density in first SCF iteration, ds error (discretization error, computed using $tol_m = 1e-8$ as reference), tc error (treecode error, $|E_H(ds) - E_H(tc)|$ using MAC $\theta = 0.7$ and interpolation degree $n = 8$). Run time (s) for the direct sum (ds) and treecode (tc) and treecode speedup (ds/tc) on a 6-core CPU and a single GPU.

tol_m	N_m	E_H (Ha)	ds error	tc error	6-core CPU time (s)			GPU time (s)		
					ds	tc	ds/tc	ds	tc	ds/tc
1e-3	141000	74.88640	6.20e-4	5.31e-7	70.42	26.24	2.68	0.39	0.47	0.82
1e-4	184750	74.87690	8.88e-3	1.20e-6	150.13	41.40	3.63	0.67	0.68	0.97
1e-5	249500	74.88668	9.00e-4	8.85e-6	216.17	80.32	2.69	1.13	1.13	1.00
1e-6	459500	74.88551	2.70e-4	3.68e-5	638.23	196.11	3.25	3.57	2.72	1.31
1e-7	928500	74.88574	4.00e-5	6.33e-6	2509.2	486.54	5.16	13.70	6.49	2.11
1e-8	2224375	74.88578	na	2.57e-6	15239.6	1373.90	11.09	78.31	17.27	4.54

6.5. Treecode efficiency on a 6-core CPU and single GPU

This section documents the BLTC efficiency on a 6-core 2.6 GHz Intel Core i7 processor and a single NVIDIA Titan V GPU. The treecode was programmed in C, with OpenMP directives for parallelizing over multiple CPU cores and OpenACC directives for running on the GPU. As noted above, the GPU implementation takes advantage of the fact that the particle-cluster interaction in Eq. (34) and the approximation in Eq. (35) both have the same direct sum structure involving independent kernel evaluations. The GPU processes the particle-particle interactions between the target batch and source cluster in parallel without thread divergence; this is because the MAC applies uniformly to all particles in a given target batch [60]. In practice, interaction lists are precomputed for each target batch to identify the source clusters interacting with the batch.

The performance of the BLTC on both platforms is demonstrated by computing the Hartree energy E_H in Eq. (8) for the carbon monoxide molecule using the electron density from the first SCF iteration. Results are shown in Table 4 using direct summation and the treecode, for quadrature order $p = 4$ and mesh refinement parameter tol_m between $1e-3$ and $1e-8$ yielding the indicated mesh size N_m . The direct sum energy values in the 3rd column converge as the mesh is refined, and the 4th column shows the corresponding discretization error using the value $E_H = 74.88578$ obtained with $tol_m = 1e-8$ as the reference. The 5th column records the treecode approximation error for MAC $\theta = 0.7$ and degree $n = 8$; this is the difference between the value of E_H computed by direct summation (column 3) and the value computed by the treecode (not shown). The results show that the treecode approximation error is well below the discretization error and within chemical accuracy.

The remainder of Table 4 records computation times on the 6-core CPU and GPU for direct summation (ds) and the treecode (tc). Averaging over the six runs in Table 4, direct summation runs 192 times faster on the GPU than on the 6-core CPU, while the treecode runs 70 times faster. On both platforms the treecode is faster than direct summation, and the speedup (ds/tc) increases as the mesh is refined; this is consistent with $O(N_m^2)$ scaling for direct summation and $O(N_m \log N_m)$ scaling for the treecode. In particular, for the largest mesh size with approximately 2.2 million mesh points, the treecode computation time on the GPU is less than 18 s, which is about 4.5 times faster than direct summation.

6.6. BLTC parallel efficiency on a single GPU node with 1, 2 or 4 GPUs

This subsection documents the parallel efficiency of the BLTC on a single GPU node running with 1, 2 or 4 GPUs, using OpenMP to parallelize across GPUs with one thread assigned to each GPU. The test system has 10 million particles randomly located in a cube interacting via the Coulomb kernel. The work is divided into two stages; stage 1 encompasses the precomputing tasks in lines 1-2 of Algorithm 3 and stage 2 encompasses the batch-cluster computing in line 3. Fig. 11 shows the parallel efficiency of each stage and the entire computation as the number of GPUs increases from 1 to 4. The precompute stage scales less efficiently than the compute phase, due to some serial computation embedded in these tasks (85% on 2 GPUs, 63% on 4 GPUs), but this accounts for only a small fraction of the total computation time. The compute stage has close to ideal scaling (98% on 2 GPUs, 94% on 4 GPUs); moreover this stage accounts for a large fraction of the total computation time and therefore the treecode achieves 90% efficiency for the entire computation on 4 GPUs.

7. Ground state energy computations for atoms and molecules

The ground-state energy of several atoms (Li, Be, O) and small molecules (H_2 , CO, C_6H_6) was computed using treecode-accelerated Green Iteration (TAGI) with the LDA exchange-correlation functional [63,64]. For each system the SCF iteration continued until the density residual fell below $tol_{scf} = 1e-4$. tol_{gi} was set to $3e-3$ for the first step in the SCF, then gradually reduced to $1e-5$ over the next four steps. The TAGI discretization parameters (quadrature order p , mesh refinement parameter tol_m) and treecode parameters (degree n , MAC θ) were chosen to ensure chemical accuracy of 1 mHa/atom in the computed ground-state energy. The computations were performed on a single node, where the treecode (written in C with OpenMP+OpenACC) was run on the four GPUs and the remainder of the code (written in Python) was run in serial on one CPU core.




Fig. 11. Parallel efficiency of BLTC on a single GPU node for 10 million particles interacting via Coulomb kernel. Treecode MAC parameter $\theta = 0.7$ and interpolation degree $n = 7$ yield treecode approximation error $2.31\text{e-}06$ (L_2 error with respect to direct sum). The results show computation time (s) and ideal scaling time (s) using 1, 2 and 4 GPUs for (a) stage 1 (precompute), (b) stage 2 (compute), and (c) total time. For comparison, the direct sum time on 4 GPUs is 1668 s.

Table 5

Ground-state energy computations of atoms and small molecules using TAGI with errors computed with respect to reference values E_{ref} computed using DFT-FE [31]. TAGI discretization parameters (quadrature order p , mesh refinement parameter tol_m , mesh size N_m), treecode parameters (degree n , MAC θ); error (Ha/atom) and total wall clock computation time (s) on a single node with 4 GPUs.

system	N_e	p	tol_m	N_m	n	θ	E_{TAGI}	error (Ha/atom)	time (s)
Li	3	4	$7\text{e-}6$	232000	6	0.8	-7.334051	$4.59\text{e-}4$	35.75
Be	4	3	$1\text{e-}5$	179712	6	0.8	-14.445658	$5.32\text{e-}4$	29.35
O	8	4	$3\text{e-}7$	421000	6	0.8	-74.469668	$-3.38\text{e-}4$	143.09
H ₂	2	3	$1\text{e-}3$	51200	6	0.8	-1.13584	$9.05\text{e-}4$	6.67
CO	14	4	$3\text{e-}7$	661625	7	0.7	-112.473372	$-7.21\text{e-}4$	932.03
C ₆ H ₆	42	4	$3\text{e-}6$	1464500	8	0.6	-230.193158	$-3.64\text{e-}4$	13947.89




Fig. 12. 2D slices of the adaptively refined mesh and computed electron density ρ for (a) carbon monoxide molecule (CO), and (b) benzene molecule (C₆H₆). The slices are taken at $z = 0$ and show $[-15, 15]^2$ a.u. in xy -plane.

Table 5 presents the parameters and results for each system. The numerical parameters (p, tol_m, n, θ) are chosen to ensure chemical accuracy in the energy, and the heavier carbon and oxygen atoms require somewhat higher numerical resolution than the lighter hydrogen, lithium, and beryllium atoms. In particular, a larger mesh size N_m requires slightly tighter treecode parameters (increasing degree from $n = 6$ to $n = 7, 8$, decreasing MAC from $\theta = 0.8$ to $\theta = 0.7, 0.6$). Column 9 records the error in the ground-state energy computed by TAGI with respect to reference energies computed to 0.1 mHa/atom accuracy with DFT-FE [31], showing that TAGI achieves chemical accuracy. Column 10 records the total wall clock computation time (s). The benzene molecule (C₆H₆, $N_e = 42$) is the largest system considered; the computation used approximately 1.5 million mesh points and required less than 4 hours of wall clock time.

Fig. 12 show 2D slices of the adaptively refined mesh and computed electron density for the carbon monoxide molecule 12a and benzene molecule 12b. Both molecules are located in the $z=0$ plane; for the CO molecule, the carbon atom is at $(-1.06581, 0, 0)$ and the oxygen atom is at $(+1.06581, 0, 0)$; for the benzene molecule, the six carbon atoms are at $(\mp 0.682781, \pm 2.548170, 0)$, $(\pm 2.548230, \mp 0.682767, 0)$, $(\pm 1.86544, \pm 1.86541, 0)$, and the six hydrogen atoms are at $(\mp 1.21247, \pm 4.52502, 0)$, $(\pm 4.52548, \mp 1.21333, 0)$, $(\pm 3.31252, \pm 3.31351, 0)$. The slices are taken in the $z = 0$ plane and show the region $[-15, 15]$ a.u. in the xy -plane. The adaptive mesh successfully captures the variation in magnitude of the electron density.

8. Conclusions

We presented a real-space method for all-electron Kohn-Sham DFT computations called Treecode-Accelerated Green Iteration (TAGI). TAGI is based on a reformulation of the Kohn-Sham equations in which the eigenvalue problem for the energies and wavefunctions (ε_i, ψ_i) in differential form is recast as a fixed-point problem in integral form by convolution with the bound state Helmholtz Green's function [46]. In each step of the SCF iteration the fixed-points are computed by Green Iteration, where the convolution integrals are discretized on an adaptively refined mesh and the discrete convolution sums are efficiently evaluated using a GPU-accelerated treecode.

TAGI relies on several key techniques to achieve chemical accuracy and computational efficiency. First, the Fejér quadrature rule and *a-priori* adaptive mesh refinement based on integration of a test function are used to compute integrals and represent the fields. Second, singularity subtraction is applied to evaluate convolution integrals having singular kernels; in particular a standard scheme is used for the Yukawa kernel in the Green Iteration convolutions, and we developed a new scheme for the Coulomb kernel in the Hartree potential convolutions since the standard scheme is inapplicable in that case. Third, a gradient-free method is employed to update the eigenvalues in Green Iteration. Fourth, the fixed-point iteration for the wavefunctions and eigenvalues in Green Iteration is accelerated using Anderson mixing. Fifth, the discrete convolution sums are computed efficiently using a barycentric Lagrange treecode (BLTC), which reduces the operation count from $O(N_m^2)$ to $O(N_m \log N_m)$, where N_m is the number of mesh points. The GPU implementation of the BLTC is facilitated by the properties of barycentric Lagrange interpolation including its scale-invariance and the fact that the particle-cluster approximation in Eq. (35) consists of independent kernel evaluations which can be evaluated concurrently [60].

We demonstrated the effect of these techniques on the carbon monoxide molecule. First, we investigated the quadrature rule and *a-priori* adaptive mesh refinement scheme, showing that the ground-state energy of the CO molecule is computed to within chemical accuracy of 1 mHa/atom using roughly 600,000 quadrature points and 4th order quadrature. Second, we demonstrated the effect of the singularity subtraction schemes; the ground-state computation was performed with and without singularity subtraction on a sequence of progressively refined meshes, and we observed a 100-fold reduction in error using singularity subtraction. Third, we compared three methods for updating the eigenvalues in Green Iteration, 1) Laplacian update using the Rayleigh quotient of the Hamiltonian differential operator, 2) gradient update using integration by parts to reduce the order of the operator, and 3) gradient-free update; on a wide range of meshes the gradient-free update yields a 10-fold improvement in accuracy over the gradient update, and a 100-fold improvement over the Laplacian update. Fourth, we investigated the convergence of Green Iteration in the first SCF iteration for the CO molecule, showing that the spectral gap in the Hamiltonian eigenvalues controls the rate of convergence of the eigenfunctions in Green Iteration; in particular, a small gap $\Delta\varepsilon_i = |\varepsilon_{i+1} - \varepsilon_i|$ implies slow convergence of the eigenfunction ψ_i ; the convergence rates were predicted and then verified computationally; finally we showed that in the case of slow convergence, Green Iteration can be accelerated by applying Anderson mixing to the eigenpairs, yielding a 6-fold reduction in the number of iterations in this example. Fifth, we demonstrated the treecode's ability to rapidly compute accurate approximations of the discrete convolution sums; in ground-state computations for the CO molecule, we showed that the treecode approximation error can be driven below the discretization error; we then demonstrated the speedup of the treecode over direct summation on both a 6-core CPU (parallelized with OpenMP) and a GPU (parallelized with OpenACC), achieving an $11\times$ speedup on the CPU and a $4.5\times$ speedup on the GPU; finally we observed a $70\times$ speedup of the treecode running on the GPU in comparison with the CPU. A primary advantage of TAGI is that the discrete convolutions are arithmetically intensive and well suited for GPU acceleration. This is significant as GPUs become increasingly dominant in HPC systems (GPUs account for 95% of available floating point operations on a Summit node at Oak Ridge National Lab [78]).

We then performed TAGI computations for several atoms and small molecules on a single node with 4 GPUs, and verified the accuracy of the ground-state energy with respect to reference values. The results demonstrate the chemical accuracy, systematic convergence, and computational efficiency afforded by TAGI on these benchmark systems.

Future work. The purpose of this work was to demonstrate the feasibility of TAGI as a real-space integral equation based approach for all-electron Kohn-Sham DFT calculations. At present TAGI is less efficient than more mature highly optimized state-of-the-art codes; for example, on a single compute node, DFT-FE computes the ground state energy of benzene approximately $20\times$ faster than TAGI. However there are several opportunities to improve the efficiency of this initial implementation of TAGI as discussed below.

First, TAGI currently employs an *a priori* mesh refinement scheme and it may be beneficial to investigate instead an *a posteriori* adaptive mesh refinement technique as used in finite element methods [70,31]. Additionally it would be interesting to replace the current Cartesian mesh cells by curvilinear cells adapted to the molecular geometry [79,80].

Second, the present form of Green Iteration explicitly computes the eigenvectors in each step of the SCF iteration. However it has been shown that computing an orthonormal basis for the eigenspace by Chebyshev filtered subspace iteration can reduce the computational cost in comparison with explicit eigensolves [81,82]; for example, a $10\times$ speedup over previous methods was reported using this technique [28,83]. A subspace iteration approach could potentially be applied as well in the integral equation framework as an alternative to the present implementation of Green Iteration.

Third, as a real-space method TAGI naturally lends itself to domain decomposition, which will enable the use of parallelized orthogonalization routines, which can also be GPU-accelerated, and a distributed memory implementation of BLTC for

parallelized integration [60]. Additionally, a recent version of the treecode based on barycentric Hermite interpolation [84] has been shown to be more efficient than barycentric Lagrange interpolation in the high accuracy regime demanded by the DFT application, and this could also improve TAGI's efficiency.

Fourth, TAGI can be extended to treat a wider variety of materials systems. In particular, periodic systems such as bulk materials or crystal structures can be accommodated using periodic treecodes [85,86]. Furthermore, we are developing a pseudopotential version of TAGI using Optimized Norm-Conserving Vanderbilt (ONCV) pseudopotentials; in this approach the singular Coulomb potential is replaced by a regularized pseudopotential, alleviating the need for highly refined meshes near the nuclei and enabling efficient calculations of larger materials systems. Preliminary results extending TAGI to ONCV calculations and using an MPI parallelized domain decomposition can be found in the reference [87].

CRedit authorship contribution statement

Nathan Vaughn: Code developer, Methodology, Data curator, First draft writer. **Vikram Gavini:** Supervisor, Conceptualization, Methodology, Draft editor. **Robert Krasny:** Supervisor, Conceptualization, Methodology, Draft editor.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgement

We thank the reviewers for carefully reading the original submission and providing useful suggestions. This work was supported by National Science Foundation grant DMS-1819094, and the Michigan Institute for Computational Discovery and Engineering (MICDE) and Mcubed program at the University of Michigan. RK thanks Lunmei Huang and Li Wang for early discussions on the Green's function approach to DFT. NV thanks Bikash Kanungo and Sambit Das for providing reference values and the single-atom radial data. Computational resources and services were provided by Advanced Research Computing-Technology Services (ARC-TS) at the University of Michigan.

Appendix A. Symbols

Variables, fields, and operators	Symbol
Electron density, Eq. (4)	ρ
Effective potential, Eq. (2)	$V_{eff}[\rho]$
Hamiltonian operator, Eq. (1)	$\mathcal{H}[\rho]$
Hamiltonian eigenpairs, Eq. (1)	(ε_i, ψ_i)
Green Function Integral operator, Eq. (12)	$\mathcal{G}(\varepsilon)$
Integral eigenpairs, Eq. (30)	(μ_i, ϕ_i)
Physical parameters and constants	Symbol
Number of atoms, Eq. (3)	N_A
Atomic positions, Eq. (3)	$\mathbf{R}_j$
Nuclear charges, Eq. (3)	Z_j
Number of electrons, Eq. (5)	N_e
Boltzmann constant, Eq. (4)	k_B
Temperature, Eq. (4)	T
Fermi Energy, Eq. (5)	μ_F
Numerical Parameters	Symbol
Number of wavefunctions, Eq. (5)	N_w
Mesh refinement parameter, Eq. (16)	tol_m
Number of cells, Eq. (14)	N_c
Quadrature order, Eq. (14)	p
Number of quadrature points, Eq. (14)	N_m
SCF convergence tolerance, Algorithm (1)	tol_{scf}
Green Iteration convergence tolerance, Algorithm (2)	tol_{gi}
Anderson mixing parameter	β
Gaussian singularity subtraction parameter, Eq. (23)	α
Gauge shift, Eq. (22)	V_{shift}
Treecode MAC parameter, Eq. (36)	θ
Treecode interpolation degree, Eq. (35)	n
Treecode maximum batch size	N_B
Treecode maximum leaf cluster size	N_L
Number of source particles in cluster	N_S

References

- [1] M. Challacombe, E. Schwegler, J. Almlöf, Fast assembly of the Coulomb matrix: a quantum chemical tree code, *J. Chem. Phys.* 104 (1996) 4685–4698, <https://doi.org/10.1063/1.471163>.
- [2] M. Challacombe, E. Schwegler, Linear scaling computation of the Fock matrix, *J. Chem. Phys.* 106 (1997) 5526–5536, <https://doi.org/10.1063/1.473575>.
- [3] W.C. Chew, E. Michielssen, J.M. Song, J.M. Jin, *Fast and Efficient Algorithms in Computational Electromagnetics*, Artech House, 2001.
- [4] O. Cohen, L. Kronik, A. Brandt, Locally refined multigrid solution of the all-electron Kohn-Sham equation, *J. Chem. Theory Comput.* 9 (2013) 4744–4760, <https://doi.org/10.1021/ct400479u>.
- [5] A. Natan, A. Benjamini, D. Naveh, L. Kronik, M.L. Tiago, S.P. Beckman, J.R. Chelikowsky, Real-space pseudopotential method for first principles calculations of general periodic and partially periodic systems, *Phys. Rev. B* 78 (2008) 075109, <https://doi.org/10.1103/PhysRevB.78.075109>.
- [6] P. Hohenberg, W. Kohn, Inhomogeneous electron gas, *Phys. Rev.* 136 (3B) (1964) B864–B871, <https://doi.org/10.1103/PhysRev.136.B864>.
- [7] W. Kohn, L.J. Sham, Self-consistent equations including exchange and correlation effects, *Phys. Rev.* 140 (4A) (1965) A1133–A1138, <https://doi.org/10.1103/PhysRev.140.A1133>.
- [8] K. Burke, Perspective on density functional theory, *J. Chem. Phys.* 136 (15) (2012) 150901, <https://doi.org/10.1063/1.4704546>.
- [9] R.O. Jones, Density functional theory: its origins, rise to prominence, and future, *Rev. Mod. Phys.* 87 (2015) 897–923, <https://doi.org/10.1103/RevModPhys.87.897>.
- [10] N. Mardirossian, M. Head-Gordon, Thirty years of density functional theory in computational chemistry: an overview and extensive assessment of 200 density functionals, *Mol. Phys.* 115 (19) (2017) 2315–2372, <https://doi.org/10.1080/00268976.2017.1333644>.
- [11] L. Lin, J. Lu, L. Ying, Numerical methods for Kohn-Sham density functional theory, *Acta Numer.* 28 (2019) 405–539, <https://doi.org/10.1017/S0962492919000047>.
- [12] G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set, *Phys. Rev. B* 54 (16) (1996) 11169–11186, <https://doi.org/10.1103/PhysRevB.54.11169>.
- [13] X. Gonze, J.-M. Beuken, R. Caracas, F. Detraux, M. Fuchs, G.-M. Rignanese, L. Sindic, M. Verstraete, G. Zerah, F. Jollet, M. Torrent, A. Roy, M. Mikami, P. Ghosez, J.-Y. Raty, D. Allan, First-principles computation of material properties: the ABINIT software project, *Comput. Mater. Sci.* 25 (3) (2002) 478–492, [https://doi.org/10.1016/S0927-0256\(02\)00325-7](https://doi.org/10.1016/S0927-0256(02)00325-7).
- [14] M.D. Segall, P.J.D. Lindan, M.J. Probert, C.J. Pickard, P.J. Hasnip, S.J. Clark, M.C. Payne, First-principles simulation: ideas, illustrations and the CASTEP code, *J. Phys. Condens. Matter* 14 (11) (2002) 2717–2744, <https://doi.org/10.1088/0953-8984/14/11/301>.
- [15] P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G.L. Chiarotti, M. Cococcioni, I. Dabo, A. Dal Corso, S. de Gironcoli, S. Fabris, G. Fratesi, R. Gebauer, U. Gerstmann, C. Gougousis, A. Kokalj, M. Lazzeri, L. Martin-Samos, N. Marzari, F. Mauri, R. Mazzarello, S. Paolini, A. Pasquarello, L. Paulatto, C. Sbraccia, S. Scandolo, G. Sclauzero, A.P. Seitsonen, A. Smogunov, P. Umari, R.M. Wentzcovitch, Quantum ESPRESSO: a modular and open-source software project for quantum simulations of materials, *J. Phys. Condens. Matter* 21 (39) (2009) 395502, <https://doi.org/10.1088/0953-8984/21/39/395502>.
- [16] P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. Buongiorno Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A. Dal Corso, S. de Gironcoli, P. Delugas, R.A. DiStasio, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo, R. Gebauer, U. Gerstmann, F. Giustino, T. Gorni, J. Jia, M. Kawamura, H.-Y. Ko, A. Kokalj, E. Küçükbenli, M. Lazzeri, M. Marsili, N. Marzari, F. Mauri, N.L. Nguyen, H.-V. Nguyen, A. Otero-de-la Roza, L. Paulatto, S. Poncè, D. Rocca, R. Sabatini, B. Santra, M. Schlipf, A.P. Seitsonen, A. Smogunov, I. Timrov, T. Thonhauser, P. Umari, N. Vast, X. Wu, S. Baroni, Advanced capabilities for materials modelling with quantum ESPRESSO, *J. Phys. Condens. Matter* 29 (46) (2017) 465901, <https://doi.org/10.1088/1361-648X/aa8f79>.
- [17] T. Loucks, J.C. Slater, Augmented plane wave method: a guide to performing electronic structure calculations, in: *Phys. Today*, vol. 20, 1967, pp. 92–93.
- [18] O.K. Andersen, Linear methods in band theory, *Phys. Rev. B* 12 (8) (1975) 3060–3083, <https://doi.org/10.1103/PhysRevB.12.3060>.
- [19] E. Wimmer, H. Krakauer, M. Weinert, A.J. Freeman, Full-potential self-consistent linearized-augmented-plane-wave method for calculating the electronic structure of molecules and surfaces: O₂ molecule, *Phys. Rev. B* 24 (2) (1981) 864–875, <https://doi.org/10.1103/PhysRevB.24.864>.
- [20] M. Weinert, E. Wimmer, A.J. Freeman, Total-energy all-electron density functional method for bulk solids and surfaces, *Phys. Rev. B* 26 (8) (1982) 4571–4578, <https://doi.org/10.1103/PhysRevB.26.4571>.
- [21] E. Sjöstedt, L. Nordström, D. Singh, An alternative way of linearizing the augmented plane-wave method, *Solid State Commun.* 114 (1) (2000) 15–20, [https://doi.org/10.1016/S0038-1098\(99\)00577-3](https://doi.org/10.1016/S0038-1098(99)00577-3).
- [22] G.K.H. Madsen, P. Blaha, K. Schwarz, E. Sjöstedt, L. Nordström, Efficient linearization of the augmented plane-wave method, *Phys. Rev. B* 64 (19) (2001) 195134, <https://doi.org/10.1103/PhysRevB.64.195134>.
- [23] M.J. Frisch, G.W. Trucks, H.B. Schlegel, G.E. Scuseria, M.A. Robb, J.R. Cheeseman, G. Scalmani, V. Barone, G.A. Petersson, H. Nakatsuji, X. Li, M. Caricato, A.V. Marenich, J. Bloino, B.G. Janesko, R. Gomperts, B. Mennucci, H.P. Hratchian, J.V. Ortiz, A.F. Izmaylov, J.L. Sonnenberg, D. Williams-Young, F. Ding, F. Lipparini, F. Egidi, J. Goings, B. Peng, A. Petrone, T. Henderson, D. Ranasinghe, V.G. Zakrzewski, J. Gao, N. Rega, G. Zheng, W. Liang, M. Hada, M. Ehara, K. Toyota, R. Fukuda, J. Hasegawa, M. Ishida, T. Nakajima, Y. Honda, O. Kitao, H. Nakai, T. Vreven, K. Throssell, J.A. Montgomery Jr., J.E. Peralta, F. Ogliaro, M.J. Bearpark, J.J. Heyd, J.N. Brothers, K.N. Kudin, V.N. Staroverov, T.A. Keith, R. Kobayashi, J. Normand, K. Raghavachari, A.P. Rendell, J.C. Burant, S.S. Iyengar, J. Tomasi, M. Cossi, J.M. Millam, M. Klene, C. Adamo, R. Cammi, J.W. Ochterski, R.L. Martin, K. Morokuma, O. Farkas, J.B. Foresman, D.J. Fox, *Gaussian 16 Revision B.01*, Gaussian Inc. Wallingford CT 2016.
- [24] M. Valiev, E. Bylaska, N. Govind, K. Kowalski, T. Straatsma, H. Van Dam, D. Wang, J. Nieplocha, E. Apra, T. Windus, W. de Jong, NWChem: a comprehensive and scalable open-source solution for large scale molecular simulations, *Comput. Phys. Commun.* 181 (9) (2010) 1477–1489, <https://doi.org/10.1016/j.cpc.2010.04.018>.
- [25] E. Tsuchida, M. Tsukada, Large-scale electronic-structure calculations based on the adaptive finite-element method, *J. Phys. Soc. Jpn.* 67 (11) (1998) 3844–3858, <https://doi.org/10.1143/JPSJ.67.3844>.
- [26] J.E. Pask, B.M. Klein, C.Y. Fong, P.A. Sterne, Real-space local polynomial basis for solid-state electronic-structure calculations: a finite-element approach, *Phys. Rev. B* 59 (19) (1999) 12352–12358, <https://doi.org/10.1103/PhysRevB.59.12352>.
- [27] L. Lehtovaara, V. Havu, M. Puska, All-electron density functional theory and time-dependent density functional theory with high-order finite elements, *J. Chem. Phys.* 131 (5) (2009) 1–10, <https://doi.org/10.1063/1.3176508>.
- [28] P. Motamarri, M. Nowak, K. Leiter, J. Knap, V. Gavini, Higher-order adaptive finite-element methods for Kohn-Sham density functional theory, *J. Comput. Phys.* 253 (2013) 308–343, <https://doi.org/10.1016/j.jcp.2013.06.042>.
- [29] P. Motamarri, V. Gavini, Subquadratic-scaling subspace projection method for large-scale Kohn-Sham density functional theory calculations using spectral finite-element discretization, *Phys. Rev. B* 90 (11) (2014) 115127, <https://doi.org/10.1103/PhysRevB.90.115127>.
- [30] B. Kanungo, V. Gavini, Large-scale all-electron density functional theory calculations using an enriched finite-element basis, *Phys. Rev. B* 95 (3) (2017) 035112, <https://doi.org/10.1103/PhysRevB.95.035112>.
- [31] P. Motamarri, S. Das, S. Rudraraju, K. Ghosh, D. Davydov, V. Gavini, DFT-FE – a massively parallel adaptive finite-element code for large-scale density functional theory calculations, *Comput. Phys. Commun.* 246 (2020) 106853, <https://doi.org/10.1016/j.cpc.2019.07.016>.
- [32] S. Das, P. Motamarri, V. Gavini, B. Turcksin, Y.W. Li, B. Leback, Fast, scalable and accurate finite-element based ab initio calculations using mixed precision computing: 46 PFLOPS simulation of a metallic dislocation system, in: *Int. Conf. High Perform. Comput. Networking, Storage Anal. SC, IEEE Computer Society, New York, NY, USA, 2019*, pp. 1–11.

- [33] V. Rokhlin, Rapid solution of integral equations of scattering theory in two dimensions, *J. Comput. Phys.* 86 (2) (1990) 414–439, [https://doi.org/10.1016/0021-9991\(90\)90107-C](https://doi.org/10.1016/0021-9991(90)90107-C).
- [34] R. Jorgenson, R. Mittra, Efficient calculation of the free-space periodic Green's function, *IEEE Trans. Antennas Propag.* 38 (5) (1990) 633–642, <https://doi.org/10.1109/8.53491>.
- [35] R. Coifman, V. Rokhlin, S. Wandzura, The fast multipole method for the wave equation: a pedestrian prescription, *IEEE Antennas Propag. Mag.* 35 (3) (1993) 7–12, <https://doi.org/10.1109/74.250128>.
- [36] L.N. Medgyesi-Mitschang, J.M. Putnam, M.B. Gedera, Generalized method of moments for three-dimensional penetrable scatterers, *J. Opt. Soc. Am. A* 11 (4) (1994) 1383, <https://doi.org/10.1364/JOSAA.11.001383>.
- [37] E. Bleszynski, M. Bleszynski, T. Jaroszewicz, AIM: adaptive integral method for solving large-scale electromagnetic scattering and radiation problems, *Radio Sci.* 31 (5) (1996) 1225–1251, <https://doi.org/10.1029/96RS02504>.
- [38] J.L. Volakis, K. Sertel, *Integral Equation Methods for Electromagnetics*, Scitech, 2011.
- [39] M.M. Botha, Solving the volume integral equations of electromagnetic scattering, *J. Comput. Phys.* 218 (1) (2006) 141–158, <https://doi.org/10.1016/j.jcp.2006.02.004>.
- [40] L.D. Faddeyev, B. Seckler, The inverse problem in the quantum theory of scattering, *J. Math. Phys.* 4 (1) (1963) 72–104, <https://doi.org/10.1063/1.1703891>.
- [41] B.R. Johnson, D. Secrest, The solution of the nonrelativistic quantum scattering problem without exchange, *J. Math. Phys.* 7 (12) (1966) 2187–2195, <https://doi.org/10.1063/1.1704905>.
- [42] E. Alt, P. Grassberger, W. Sandhas, Reduction of the three-particle collision problem to multi-channel two-particle Lippmann-Schwinger equations, *Nucl. Phys. B* 2 (2) (1967) 167–180, [https://doi.org/10.1016/0550-3213\(67\)90016-8](https://doi.org/10.1016/0550-3213(67)90016-8).
- [43] R.I. Masel, R.P. Merrill, W.H. Miller, Quantum scattering from a sinusoidal hard wall: atomic diffraction from solid surfaces, *Phys. Rev. B* 12 (12) (1975) 5545–5551, <https://doi.org/10.1103/PhysRevB.12.5545>.
- [44] S.K. Adhikari, Quantum scattering in two dimensions, *Am. J. Phys.* 54 (4) (1986) 362–367, <https://doi.org/10.1119/1.14623>.
- [45] K.T. Hecht, Operator form of scattering Green's function and the integral equation for the scattering problem, in: *Scattering*, Springer New York, New York, NY, 2000, pp. 477–480.
- [46] M.H. Kalos, Monte Carlo, Calculations of the ground state of three- and four-body nuclei, *Phys. Rev.* 128 (4) (1962) 1791–1795, <https://doi.org/10.1103/PhysRev.128.1791>.
- [47] Z. Zhao, N. Kovvali, W. Lin, C.-H. Ahn, L. Couchman, L. Carin, Volumetric fast multipole method for modeling Schrödinger's equation, *J. Comput. Phys.* 224 (2) (2007) 941–955, <https://doi.org/10.1016/j.jcp.2006.11.003>.
- [48] R.J. Harrison, G.I. Fann, T. Yanai, Z. Gan, G. Beylkin, Multiresolution quantum chemistry: basic theory and initial applications, *J. Chem. Phys.* 121 (23) (2004) 11587–11598, <https://doi.org/10.1063/1.1791051>.
- [49] R.J. Harrison, G. Beylkin, F.A. Bischoff, J.A. Calvin, G.I. Fann, J. Fosso-Tande, D. Galindo, J.R. Hammond, R. Hartman-Baker, J.C. Hill, J. Jia, J.S. Kottmann, M.J. Ou, J. Pei, L.E. Ratcliff, M.G. Reuter, A.C. Richie-Halford, N.A. Romero, H. Sekino, W.A. Shelton, B.E. Sundahl, W.S. Thornton, E.F. Valeev, Á. Vázquez-Mayagoitia, N. Vence, T. Yanai, Y. Yokoi, *Madness: A Multiresolution, Adaptive Numerical Environment for Scientific Simulation*, *SIAM J. Sci. Comput.*, vol. 38, Society for Industrial and Applied Mathematics Publications, 2016, pp. S123–S142, [arXiv:1507.01888](https://arxiv.org/abs/1507.01888).
- [50] M.J. Mohlenkamp, T. Young, Convergence of Green iterations for Schrödinger equations, *Recent Adv. Comput. Sci.* (2008) 201–208, https://doi.org/10.1142/9789812792389_0010.
- [51] M.J. Mohlenkamp, Function space requirements for the single-electron functions within the multiparticle Schrödinger equation, *J. Math. Phys.* 54 (6) (2013) 062105, <https://doi.org/10.1063/1.4811396>.
- [52] B.N. Khoromskij, On tensor approximation of Green iterations for Kohn-Sham equations, *Comput. Vis. Sci.* 11 (4–6) (2008) 259–271, <https://doi.org/10.1007/s00791-008-0097-x>.
- [53] M.V. Rakhuba, I.V. Oseledets, Fast multidimensional convolution in low-rank tensor formats via cross approximation, *SIAM J. Sci. Comput.* 37 (2) (2015) A565–A582, <https://doi.org/10.1137/140958529>.
- [54] M.V. Rakhuba, I.V. Oseledets, Grid-based electronic structure calculations: the tensor decomposition approach, *J. Comput. Phys.* 312 (2016) 19–30, <https://doi.org/10.1016/j.jcp.2016.02.023>.
- [55] L. Fejér, Mechanische Quadraturen mit positiven Cotesschen Zahlen, *Math. Z.* 37 (1) (1933) 287–309, <https://doi.org/10.1007/BF01474575>.
- [56] L.N. Trefethen, Is Gauss quadrature better than Clenshaw-Curtis?, *SIAM Rev.* 50 (1) (2008) 67–87, <https://doi.org/10.1137/060659831>.
- [57] P. Rabinowitz, Approximate methods of higher analysis. L.V. Kantorovich and V.I. Krylov, translated from the third Russian edition by Curtis D. Benster. Interscience, New York, 1959, *Science* 134 (3487) (1961) 1358, <https://doi.org/10.1126/science.134.3487.1358>.
- [58] P.M. Anselone, Singularity subtraction in the numerical solution of integral equations, *J. Aust. Math. Soc. Ser. B, Appl. Math.* 22 (4) (1981) 408–418, <https://doi.org/10.1017/S0334270000002757>.
- [59] L. Wang, R. Krasny, S. Tlupova, A kernel-independent treecode based on barycentric Lagrange interpolation, *Commun. Comput. Phys.* 28 (2020) 1415–1436.
- [60] N. Vaughn, L. Wilson, R. Krasny, A GPU-accelerated Barycentric Lagrange Treecode, in: *2020 IEEE International Parallel and Distributed Processing Symposium Workshops, IPDPSW, New Orleans, LA, USA, 2020*, pp. 701–710.
- [61] S. Goedecker, Linear scaling electronic structure methods, *Rev. Mod. Phys.* 71 (4) (1999) 1085–1123, <https://doi.org/10.1103/RevModPhys.71.1085>.
- [62] D.G. Anderson, Iterative procedures for nonlinear integral equations, *J. ACM* 12 (4) (1965) 547–560, <https://doi.org/10.1145/321296.321305>.
- [63] D.M. Ceperley, B.J. Alder, Ground state of the electron gas by a stochastic method, *Phys. Rev. Lett.* 45 (7) (1980) 566–569, <https://doi.org/10.1103/PhysRevLett.45.566>.
- [64] J.P. Perdew, A. Zunger, Self-interaction correction to density-functional approximations for many-electron systems, *Phys. Rev. B* 23 (10) (1981) 5048–5079, <https://doi.org/10.1103/PhysRevB.23.5048>.
- [65] M.A.L. Marques, M.J.T. Oliveira, T. Burnus, Libxc: a library of exchange and correlation functionals for density functional theory, *Comput. Phys. Commun.* 183 (10) (2012) 2272–2281, <https://doi.org/10.1016/j.cpc.2012.05.007>.
- [66] S. Lehtola, C. Steigemann, M.J. Oliveira, M.A. Marques, Recent developments in libxc — a comprehensive library of functionals for density functional theory, *SoftwareX* 7 (2018) 1–5, <https://doi.org/10.1016/j.softx.2017.11.002>.
- [67] J.R. Chelikowsky, N. Troullier, Y. Saad, Finite-difference-pseudopotential method: electronic structure calculations without a basis, *Phys. Rev. Lett.* 72 (8) (1994) 1240–1243, <https://doi.org/10.1103/PhysRevLett.72.1240>.
- [68] J. Bernholc, M. Hodak, W. Lu, Recent developments and applications of the real-space multigrid method, *J. Phys. Condens. Matter* 20 (29) (2008) 294205, <https://doi.org/10.1088/0953-8984/20/29/294205>.
- [69] Y. Saad, J.R. Chelikowsky, S.M. Shontz, Numerical methods for electronic structure calculations of materials, *SIAM Rev.* 52 (1) (2010) 3–54, <https://doi.org/10.1137/060651653>.
- [70] H. Chen, X. Dai, X. Gong, L. He, A. Zhou, Adaptive finite element approximations for Kohn–Sham models, *Multiscale Model. Simul.* 12 (4) (2014) 1828–1869, <https://doi.org/10.1137/130916096>.
- [71] L.N. Trefethen, *Spectral Methods in MATLAB*, SIAM, 2008.
- [72] L.N. Trefethen, D.I. Bau, *Numerical Linear Algebra*, SIAM Society for Industrial and Applied Mathematics, 2000.

- [73] J. Barnes, P. Hut, A hierarchical $O(N \log N)$ force-calculation algorithm, *Nature* 324 (6096) (1986) 446–449, <https://doi.org/10.1038/324446a0>.
- [74] L. Greengard, V. Rokhlin, A fast algorithm for particle simulations, *J. Comput. Phys.* 73 (2) (1987) 325–348, [https://doi.org/10.1016/0021-9991\(87\)90140-9](https://doi.org/10.1016/0021-9991(87)90140-9).
- [75] J.-P. Berrut, L.N. Trefethen, Barycentric Lagrange interpolation, *SIAM Rev.* 46 (3) (2004) 501–517, <https://doi.org/10.1137/S0036144502417715>.
- [76] P. Li, H. Johnston, R. Krasny, A Cartesian treecode for screened Coulomb interactions, *J. Comput. Phys.* 228 (10) (2009) 3858–3868, <https://doi.org/10.1016/j.jcp.2009.02.022>.
- [77] H.E. Salzer, Lagrangian interpolation at the Chebyshev points $x_{n,v} \equiv \cos(v\pi/n)$, $v = 0(1)n$; some unnoted advantages, *Comput. J.* 15 (1972) 156–159.
- [78] I. Buck, Reaching the summit: accelerated computing powering world's fastest supercomputer, <https://blogs.nvidia.com/blog/2018/06/08/worlds-fastest-exascale-ai-supercomputer-summit/>, 2018.
- [79] E. Tsuchida, M. Tsukada, Adaptive finite-element method for electronic-structure calculations, *Phys. Rev. B* 54 (1996) 7602–7605, <https://doi.org/10.1103/PhysRevB.54.7602>.
- [80] E. Tsuchida, M. Tsukada, Large-scale electronic-structure calculations based on the adaptive finite-element method, *J. Phys. Soc. Jpn.* 67 (11) (1998) 3844–3858, <https://doi.org/10.1143/JPSJ.67.3844>.
- [81] Y. Zhou, Y. Saad, M.L. Tiago, J.R. Chelikowsky, Self-consistent-field calculations using Chebyshev-filtered subspace iteration, *J. Comput. Phys.* 219 (1) (2006) 172–184, <https://doi.org/10.1016/j.jcp.2006.03.017>.
- [82] Y. Zhou, Y. Saad, M.L. Tiago, J.R. Chelikowsky, Parallel self-consistent-field calculations via Chebyshev-filtered subspace acceleration, *Phys. Rev. E* 74 (2006) 066704, <https://doi.org/10.1103/PhysRevE.74.066704>.
- [83] Y. Zhou, J.R. Chelikowsky, Y. Saad, Chebyshev-filtered subspace iteration method free of sparse diagonalization for solving the Kohn-Sham equation, *J. Comput. Phys.* 274 (2014) 770–782, <https://doi.org/10.1016/j.jcp.2014.06.056>.
- [84] R. Krasny, L. Wang, A treecode based on barycentric Hermite interpolation for electrostatic particle interactions, *Comput. Math. Biophys.* 7 (2019) 73–84, <https://doi.org/10.1515/cmb-2019-0006>.
- [85] Z.-H. Duan, R. Krasny, A treecode algorithm for computing Ewald summation of dipolar systems, in: *Proceedings of the 2003 ACM Symposium on Applied Computing, SAC '03*, Association for Computing Machinery, New York, NY, USA, 2003, pp. 172–177.
- [86] H.A. Boateng, Periodic Coulomb tree method: an alternative to parallel particle mesh Ewald, *J. Chem. Theory Comput.* 16 (1) (2020) 7–17, <https://doi.org/10.1021/acs.jctc.9b00648>.
- [87] N. Vaughn, GPU Accelerated Barycentric Treecodes and their Application to Kohn-Sham Density Functional Theory, Ph.D. thesis, University of Michigan, 2020.