SPARC: Accurate and efficient finite-difference formulation and parallel implementation of Density Functional Theory: Extended systems

Swarnava Ghosh^a, Phanish Suryanarayana^{*,a}

^aCollege of Engineering, Georgia Institute of Technology, GA 30332, USA

Abstract

As the second component of SPARC (Simulation Package for Ab-initio Real-space Calculations), we present an accurate and efficient finite-difference formulation and parallel implementation of Density Functional Theory (DFT) for extended systems. Specifically, employing a local formulation of the electrostatics, the Chebyshev polynomial filtered self-consistent field iteration, and a reformulation of the non-local force component, we develop a finite-difference framework wherein both the energy and atomic forces can be efficiently calculated to within desired accuracies in DFT. We demonstrate using a wide variety of materials systems that SPARC achieves high convergence rates in energy and forces with respect to spatial discretization to reference plane-wave result; exponential convergence in energies and forces with respect to vacuum size for slabs and wires; energies and forces that are consistent and display negligible ‘egg-box’ effect; accurate properties of crystals, slabs, and wires; and negligible drift in molecular dynamics simulations. We also demonstrate that the weak and strong scaling behavior of SPARC is similar to well-established and optimized plane-wave implementations for systems consisting up to thousands of electrons, but with a significantly reduced prefactor. Overall, SPARC represents an attractive alternative to plane-wave codes for performing DFT simulations of extended systems.

Key words: Electronic structure, Real-space, Finite-differences, Electrostatics, Atomic forces, Parallel computing

PROGRAM SUMMARY

Manuscript Title: SPARC: Accurate and efficient finite-difference formulation and parallel implementation of Density Functional Theory: Extended systems

Authors: Swarnava Ghosh, Phanish Suryanarayana

Program Title: SPARC

Journal Reference:

Catalogue identifier:

Licensing provisions: GNU General Public License 3 (GPL)

Programming language: C/C++

Computer: Any system with C/C++ compiler

Operating system: Linux

RAM: Problem dependent. Ranges from 80 GB to 800 GB for a system with 2500 electrons.

Number of processors used: 1 to 1000 processors.

Keywords: Electronic structure, Finite-differences, Electrostatics, Atomic forces, Parallel computing

*Corresponding Author (phanish.suryanarayana@ce.gatech.edu)

Classification: 7.3

External routines/libraries: PETSc 3.5.3 (<http://www.mcs.anl.gov/petsc>), MKL 11.2 (<https://software.intel.com/en-us/intel-mkl>), and MVAPICH2 2.1 (<http://mvapich.cse.ohio-state.edu/>).

*Catalogue identifier of previous version:** AFBL_v1_0

*Journal reference of previous version:** Comput. Phys. Comm. 212 (2017) 189

Does the new version supersede the previous version?: Yes

Nature of problem: Calculation of the static and dynamic properties of isolated and extended systems in the framework of Kohn-Sham Density Functional Theory (DFT).

Solution method: High-order finite-difference discretization. Local reformulation of the electrostatics in terms of the electrostatic potential and pseudocharge densities. Application of Bloch-periodic and zero-Dirichlet boundary conditions on the orbitals in the direction of periodicity and vacuum, respectively. Application of periodic and Dirichlet boundary conditions on the electrostatic potential in the direction of periodicity and vacuum, respectively. Integration over the Brillouin zone for extended systems using the Monkhorst-Pack grid. Calculation of the electronic ground-state using the Chebyshev polynomial filtered self-consistent field iteration in conjunction with Anderson based extrapolation/mixing schemes. Reformulation of the non-local component of the force. Geometry optimization using the Polak-Ribiere variant of non-linear conjugate gradients with secant line search. NVE molecular dynamics using the leapfrog method. Parallelization via domain decomposition and over Brillouin zone integration.

Restrictions: Cuboidal domain. System size less than ~ 4000 electrons. Local Density Approximation (LDA). Troullier-Martins pseudopotentials without relativistic or non-linear core corrections.

Reasons for the new version: To enable the study of extended systems like crystals, slabs, and wires using SPARC.

Summary of revisions: Incorporated the ability to study the static and dynamic properties of crystals, slabs, and wires.

Unusual features:

Running time: Problem dependent. Timing results for selected examples provided in the paper.

1. Introduction

Density Functional Theory (DFT) developed by Hohenberg, Kohn, and Sham [1, 2] is a popular ab-initio method for understanding as well as predicting a wide range of materials properties [3, 4, 5, 6]. However, the solution of the DFT problem still remains a formidable task, which severely restricts the size of systems that can be studied. The plane-wave discretization has been a popular choice [7, 8, 9, 10, 11, 12] since it forms a complete and orthonormal set, provides spectral convergence with increasing basis size, enables the efficient evaluation of convolutions through the Fast Fourier Transform (FFT) [13], and is amenable to efficient and effective preconditioning [14, 15]. However, the need for periodicity makes the plane-wave basis unsuitable for the study of non-periodic and localized systems such as clusters, surfaces and wires [16, 17]. Additionally, its non-local nature makes the efficient use of modern large-scale computer architectures particularly challenging [18, 19], and the development of methods that scale linearly with respect to the number of atoms impractical [20, 21].

The aforementioned limitations of the plane-wave basis have motivated the development of various real-space DFT implementations, including finite-differences [22, 23, 24, 25, 26, 27], finite-elements [28, 29, 30,

31, 32, 33], wavelets [34, 35, 36], periodic sinc functions [37], basis splines (B-splines) [38], non-uniform rational B-splines (NURBS) [39], and maximum entropy basis functions [40]. However, despite the success of real-space approaches, plane-wave implementations still remain the method of choice for practical DFT computations. This is mainly due to their superior performance in achieving chemical accuracies on the modest computational resources that are commonly available to researchers [41].

The finite-difference discretization is an attractive option for performing real-space DFT simulations because it generates a standard eigenvalue problem with relatively small spectral width, an attribute that is critical to the eigensolver performance, particularly in the absence of effective real-space preconditioners. In addition, high-order approximations can be easily incorporated, a feature that is essential for efficient electronic structure calculations. However, the lack of an underlying variational structure due to the absence of a basis can result in non-monotonic convergence of the energies and atomic forces. In addition, the reduced accuracy of spatial integrations due to the use of a lower order integration scheme can lead to a pronounced ‘egg-box’ effect [42, 43]—phenomenon arising due to the breaking of the translational symmetry—which can significantly impact the accuracy of structural relaxations and molecular dynamics simulations [44, 45, 46].

In recent work [41], we have developed an accurate and efficient finite-difference formulation and parallel implementation of DFT for isolated systems, which represents the first component of SPARC (Simulation Package for Ab-initio Real-space Calculations). The solution methodology in SPARC includes a local formulation of the electrostatics, the Chebyshev polynomial filtered self-consistent field iteration, and a reformulation of the non-local component of the force. The electrostatic formulation, atomic force calculation, and overall parallel implementation distinguishes SPARC from existing finite-difference DFT packages like PARSEC [47] and OCTOPUS [23]. Notably, for isolated systems, we have found that SPARC significantly outperforms plane-wave codes like ABINIT [9] as well as finite-difference codes like PARSEC and OCTOPUS for achieving the accuracy desired in electronic structure calculations [41].

In this work, we extend the capabilities of SPARC to enable the study of the static and dynamic properties of extended systems like crystals, surfaces, and wires.¹ Through selected examples, we demonstrate that SPARC obtains high rates of convergence in the energy and forces to reference plane-wave results on refining the discretization; exponential convergence in energies and forces with respect to vacuum size for slabs and wires; forces that are consistent with the energy, both being free from any noticeable ‘egg-box’ effect; accurate ground-state properties; and negligible drift in molecular dynamics simulations. Moreover, we show that SPARC displays weak and strong scaling behavior that is similar to well-established plane-wave codes, but with a significantly smaller prefactor.

The remainder of this paper is organized as follows. In Section 2, we provide the mathematical details of DFT for extended systems. In Section 3, we describe its finite-difference formulation and parallel implementation in SPARC, whose accuracy and efficiency is verified through selected examples in Section 4. Finally, we provide concluding remarks in Section 5.

¹Extended systems are infinite in one or more directions. Therefore, the reduction of the calculations to a finite unit-cell require the electrostatics to be carefully reformulated, mainly due to their long-ranged nature. In addition, Bloch-periodic boundary conditions on the orbitals and Brillouin zone integration need to be incorporated, which introduces significant additional complexity into the theoretical, numerical, and computational aspects of DFT formulations and implementations. Also, since the underlying physics of isolated clusters are inherently different from extended systems, a thorough investigation into the accuracy and efficiency of the developed DFT approaches/code for such systems is warranted. This provides the motivation for the current work.

2. Density Functional Theory (DFT)

Consider a unit cell Ω with N atoms and a total of N_e valence electrons. Let the nuclei be positioned at $\mathbf{R} = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_N\}$ and possess valence charges $\{Z_1, Z_2, \dots, Z_N\}$, respectively. Neglecting spin, the system's free energy in Density Functional Theory (DFT) [1, 2] can be written as

$$\mathcal{F}(\Psi, \mathbf{g}, \mathbf{R}) = T_s(\Psi, \mathbf{g}) + E_{xc}(\rho) + K(\Psi, \mathbf{g}, \mathbf{R}) + E_{el}(\rho, \mathbf{R}) - TS(\mathbf{g}), \quad (1)$$

where $\Psi = \{\psi_1, \psi_2, \dots, \psi_{N_s}\}$ is the collection of orbitals with occupations $\mathbf{g} = \{g_1, g_2, \dots, g_{N_s}\}$, ρ is the electron density, and T is the electronic temperature. The electron density itself can be expressed in terms of the orbitals and their occupations as

$$\rho(\mathbf{x}) = 2 \sum_{n=1}^{N_s} \oint_{BZ} g_n(\mathbf{k}) |\psi_n(\mathbf{x}, \mathbf{k})|^2 d\mathbf{k}, \quad (2)$$

where $\mathbf{k}$ denotes the Bloch wavevector, and $\oint_{BZ}$ signifies the volume average over the Brillouin zone. In Eqn. 1, the first term is the electronic kinetic energy, the second term is the exchange-correlation energy, the third term is the non-local pseudopotential energy, the fourth term is the total electrostatic energy, and the final term is the entropic contribution arising from the partial occupations of the orbitals.

Electronic kinetic energy. In Kohn-Sham DFT, the kinetic energy of the non-interacting electrons takes the form

$$T_s(\Psi, \mathbf{g}) = - \sum_{n=1}^{N_s} \oint_{BZ} \int_{\Omega} g_n(\mathbf{k}) \psi_n^*(\mathbf{x}, \mathbf{k}) \nabla^2 \psi_n(\mathbf{x}, \mathbf{k}) d\mathbf{x} d\mathbf{k}. \quad (3)$$

where the superscript $*$ denotes the complex conjugate.

Exchange-correlation energy. The exact form of the exchange-correlation energy is currently unknown. Therefore, a number of approximations have been developed, including the widely used Local Density Approximation (LDA) [2]:

$$E_{xc}(\rho) = \int_{\Omega} \varepsilon_{xc}(\rho(\mathbf{x})) \rho(\mathbf{x}) d\mathbf{x}, \quad (4)$$

where $\varepsilon_{xc}(\rho) = \varepsilon_x(\rho) + \varepsilon_c(\rho)$ is the sum of the exchange and correlation per particle of a uniform electron gas.

Non-local pseudopotential energy. The non-local pseudopotential energy within the Kleinman-Bylander [48] representation takes the form

$$K(\Psi, \mathbf{g}, \mathbf{R}) = 2 \sum_{n=1}^{N_s} \oint_{BZ} g_n(\mathbf{k}) \sum_J \sum_{lm} \gamma_{Jl} \left| \sum_{J'} \int_{\Omega} \chi_{J'lm}^*(\mathbf{x}, \mathbf{R}_{J'}) e^{i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \psi_n(\mathbf{x}, \mathbf{k}) d\mathbf{x} \right|^2 d\mathbf{k}, \quad (5)$$

where the summation index J runs over all atoms in Ω , and the summation index J' runs over the J^{th} atom and its periodic images. In addition, the coefficients γ_{Jl} and projection functions χ_{Jlm} can be expressed as

$$\gamma_{Jl} = \left(\int_{\mathbb{R}^3} \chi_{Jlm}^*(\mathbf{x}, \mathbf{R}_J) u_{Jlm}(\mathbf{x}, \mathbf{R}_J) d\mathbf{x} \right)^{-1}, \quad \chi_{J'lm}(\mathbf{x}, \mathbf{R}_{J'}) = u_{J'lm}(\mathbf{x}, \mathbf{R}_{J'}) (V_{J'l}(\mathbf{x}, \mathbf{R}_{J'}) - V_{J'}(\mathbf{x}, \mathbf{R}_{J'})). \quad (6)$$

Above, u_{Jlm} are the isolated atom pseudowavefunctions, $V_{J'l}$ are the angular momentum dependent pseudopotentials, and $V_{J'}$ are the local components of the pseudopotentials, with l and m signifying the azimuthal and magnetic quantum numbers, respectively.

Electrostatic energy. The total electrostatic energy is the sum of three components:

$$E_{el}(\rho, \mathbf{R}) = \frac{1}{2} \int_{\mathbb{R}^3} \int_{\Omega} \frac{\rho(\mathbf{x})\rho(\mathbf{x}')}{|\mathbf{x} - \mathbf{x}'|} d\mathbf{x} d\mathbf{x}' + \sum_I \int_{\Omega} \rho(\mathbf{x}) V_I(\mathbf{x}, \mathbf{R}_I) d\mathbf{x} + \frac{1}{2} \sum_I \sum_{J \neq I} \frac{Z_I Z_J}{|\mathbf{R}_I - \mathbf{R}_J|}, \quad (7)$$

where the first term—referred to as the Hartree energy—is the classical interaction energy of the electron density, the second term is the interaction energy between the electron density and the nuclei, and the third term is the repulsion energy between the nuclei. The summation index I runs over all atoms in $\mathbb{R}^3$, i.e., all the atoms in Ω and their periodic images.

Electronic entropy. The electronic entropy originating from the partial orbital occupations:

$$S(\mathbf{g}) = -2k_B \sum_{n=1}^{N_s} \int_{BZ} (g_n(\mathbf{k}) \log g_n(\mathbf{k}) + (1 - g_n(\mathbf{k})) \log(1 - g_n(\mathbf{k}))) d\mathbf{k}, \quad (8)$$

where k_B is the Boltzmann constant.

Ground state/Molecular dynamics. The overall ground state in DFT is governed by the variational problem

$$\mathcal{F}_0 = \inf_{\mathbf{R}} \hat{\mathcal{F}}(\mathbf{R}), \quad (9)$$

where

$$\hat{\mathcal{F}}(\mathbf{R}) = \inf_{\Psi, \mathbf{g}} \mathcal{F}(\Psi, \mathbf{g}, \mathbf{R}) \quad s.t. \quad \int_{\Omega} \psi_i^*(\mathbf{x}, \mathbf{k}) \psi_j(\mathbf{x}, \mathbf{k}) d\mathbf{x} = \delta_{ij}, \quad 2 \sum_{n=1}^{N_s} \int_{BZ} g_n(\mathbf{k}) d\mathbf{k} = N_e. \quad (10)$$

In this staggered approach, the electronic ground-state as described by the above equation and the corresponding Hellmann-Feynman forces need to be computed for every configuration of the nuclei encountered during geometry optimization/molecular dynamics.

3. Formulation and implementation

In this section, we describe the real-space formulation and parallel finite-difference implementation of Density Functional Theory (DFT) for extended systems. This represents the second component of SPARC (Simulation Package for Ab-initio Real-space Calculations) [41], a first principles code currently under development for efficient large-scale electronic structure calculations.

Electrostatic reformulation. The non-locality of the electrostatic energy in Eqn. 7 makes its direct real-space evaluation scale as $\mathcal{O}(N^2)$ with respect to the number of atoms. Furthermore, the individual components diverge in extended systems. To overcome this, we adopt a local formulation of the electrostatics [49, 50]:

$$E_{el}(\rho, \mathbf{R}) = \sup_{\phi} \left\{ -\frac{1}{8\pi} \int_{\Omega} |\nabla \phi(\mathbf{x}, \mathbf{R})|^2 d\mathbf{x} + \int_{\Omega} (\rho(\mathbf{x}) + b(\mathbf{x}, \mathbf{R})) \phi(\mathbf{x}, \mathbf{R}) d\mathbf{x} \right\} - E_{self}(\mathbf{R}) + E_c(\mathbf{R}), \quad (11)$$

where ϕ denotes the electrostatic potential. In addition, b represents the total pseudocharge density of the nuclei:

$$b(\mathbf{x}, \mathbf{R}) = \sum_I b_I(\mathbf{x}, \mathbf{R}_I), \quad b_I(\mathbf{x}, \mathbf{R}_I) = -\frac{1}{4\pi} \nabla^2 V_I(\mathbf{x}, \mathbf{R}_I), \quad (12)$$

$$\int_{\Omega} b(\mathbf{x}, \mathbf{R}) d\mathbf{x} = -N_e, \quad \int_{\mathbb{R}^3} b_I(\mathbf{x}, \mathbf{R}_I) d\mathbf{x} = Z_I, \quad (13)$$

where the summation index I runs over all atoms in $\mathbb{R}^3$, and b_I is the pseudocharge density of the I^{th} nucleus that generates the potential V_I . The second to last term in Eqn. 11 represents the self energy associated with the pseudocharge densities:

$$E_{self}(\mathbf{R}) = \frac{1}{2} \sum_I \int_{\Omega} b_I(\mathbf{x}, \mathbf{R}_I) V_I(\mathbf{x}, \mathbf{R}_I) d\mathbf{x}. \quad (14)$$

The last term E_c , whose explicit expression can be found in Appendix A, corrects for the error in the repulsive energy when the pseudocharge densities overlap.

Electronic ground-state. The electronic ground-state for a given position of nuclei is governed by the constrained minimization problem in Eqn. 10. On taking the first variation and utilizing Bloch's theorem [51]:

$$\psi_n(\mathbf{x}, \mathbf{k}) = e^{i\mathbf{k} \cdot \mathbf{x}} u_n(\mathbf{x}, \mathbf{k}), \quad (15)$$

we arrive at

$$\begin{aligned} \left(\mathcal{H} \equiv -\frac{1}{2} \nabla^2 - i\mathbf{k} \cdot \nabla + \frac{1}{2} |\mathbf{k}|^2 + V_{xc} + \phi + e^{-i\mathbf{k} \cdot \mathbf{x}} V_{nl} e^{i\mathbf{k} \cdot \mathbf{x}} \right) u_n &= \lambda_n u_n, \quad n = 1, 2, \dots, N_s, \\ g_n(\mathbf{k}) &= \left(1 + \exp \left(\frac{\lambda_n(\mathbf{k}) - \lambda_f}{k_B T} \right) \right)^{-1}, \quad \text{where } \lambda_f \text{ is s.t. } 2 \sum_{n=1}^{N_s} \int_{BZ} g_n(\mathbf{k}) d\mathbf{k} = N_e, \\ \rho(\mathbf{x}) &= 2 \sum_{n=1}^{N_s} \int_{BZ} g_n(\mathbf{k}) |u_n(\mathbf{x}, \mathbf{k})|^2 d\mathbf{k}, \quad -\frac{1}{4\pi} \nabla^2 \phi(\mathbf{x}, \mathbf{R}) = \rho(\mathbf{x}) + b(\mathbf{x}, \mathbf{R}), \end{aligned} \quad (16)$$

where $i = \sqrt{-1}$, u is a function that is periodic on the unit cell in the directions of periodicity, $\mathcal{H}$ is the Hamiltonian operator, $V_{xc} = \delta E_{xc} / \delta \rho$ is the exchange-correlation potential, λ_f is the Fermi energy, and V_{nl} is the non-local pseudopotential operator:

$$\begin{aligned} V_{nl} f &= \sum_J V_{nl,J} f \\ &= \sum_J \sum_{lm} \gamma_{Jl} \left(\sum_{J'} e^{-i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm} \right) \left(\sum_{J'} \int_{\Omega} \chi_{J'lm}^*(\mathbf{x}, \mathbf{R}_{J'}) e^{i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} f(\mathbf{x}) d\mathbf{x} \right) \end{aligned} \quad (17)$$

Above, the summation index J runs over all atoms in Ω , and the summation index J' runs over the J^{th} atom and its periodic images.

The electronic ground-state is determined by solving the non-linear eigenvalue problem in Eqn. 16 using the Self-Consistent Field (SCF) method [52]. Specifically, a fixed-point iteration is performed with respect to the potential $V_{eff} = V_{xc} + \phi$, which is further accelerated using mixing/extrapolation schemes [53, 54, 55, 56]. In each iteration of the SCF method, the electrostatic potential is calculated by solving the Poisson equation, and the electron density is determined by computing the eigenfunctions of the linearized Hamiltonian. The orthogonality requirement amongst the Kohn-Sham orbitals makes such a procedure scale asymptotically as $\mathcal{O}(N^3)$ with respect to the number of atoms, which severely limits the size of systems that can be studied. To overcome this restrictive scaling, $\mathcal{O}(N)$ approaches [20, 21, 57, 58] will be subsequently developed and implemented into SPARC.

Free energy. The free energy is evaluated using the Harris-Foulkes [59, 60] type functional:

$$\begin{aligned}
\hat{\mathcal{F}}(\mathbf{R}) &= 2 \sum_{n=1}^{N_s} \int_{BZ} g_n(\mathbf{k}) \lambda_n(\mathbf{k}) d\mathbf{k} + \int_{\Omega} \varepsilon_{xc}(\rho(\mathbf{x})) \rho(\mathbf{x}) d\mathbf{x} - \int_{\Omega} V_{xc}(\rho(\mathbf{x})) \rho(\mathbf{x}) d\mathbf{x} \\
&+ \frac{1}{2} \int_{\Omega} (b(\mathbf{x}, \mathbf{R}) - \rho(\mathbf{x})) \phi(\mathbf{x}, \mathbf{R}) d\mathbf{x} - E_{self}(\mathbf{R}) + E_c(\mathbf{R}) \\
&+ 2k_B T \sum_{n=1}^{N_s} \int_{BZ} (g_n(\mathbf{k}) \log g_n(\mathbf{k}) + (1 - g_n(\mathbf{k})) \log(1 - g_n(\mathbf{k}))) d\mathbf{k},
\end{aligned} \tag{18}$$

where E_{self} and E_c are as defined in Eqns. 14 and 46, respectively.

Atomic forces. Once the electronic ground-state has been determined, the atomic forces are calculated using the relation:

$$\begin{aligned}
\mathbf{f}_J &= - \frac{\partial \hat{\mathcal{F}}(\mathbf{R})}{\partial \mathbf{R}_J} \\
&= \sum_{J'} \int_{\Omega} \nabla b_{J'}(\mathbf{x}, \mathbf{R}_{J'}) (\phi(\mathbf{x}, \mathbf{R}) - V_{J'}(\mathbf{x}, \mathbf{R}_{J'})) d\mathbf{x} + \mathbf{f}_{J,c}(\mathbf{R}) \\
&- 4 \sum_{n=1}^{N_s} \int_{BZ} g_n(\mathbf{k}) \sum_{lm} \gamma_{Jl} \Re \left[\left(\sum_{J'} \int_{\Omega} \psi_n^*(\mathbf{x}, \mathbf{k}) e^{-i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}(\mathbf{x}, \mathbf{R}_{J'}) d\mathbf{x} \right) \right. \\
&\quad \times \left. \left(\sum_{J'} \int_{\Omega} \nabla \psi_n(\mathbf{x}, \mathbf{k}) e^{i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^*(\mathbf{x}, \mathbf{R}_{J'}) d\mathbf{x} \right) \right] d\mathbf{k},
\end{aligned} \tag{19}$$

where the summation index J' runs over the J^{th} atom and its periodic images, and $\Re[\cdot]$ denotes the real part of the bracketed expression. The first term is the local component of the force [17], the second term corrects for overlapping pseudocharge densities [61] (Appendix A), and the final term is the non-local component obtained by transferring the derivative on the projectors (with respect to the atomic position) to the orbitals (with respect to space) [62]. This strategy is employed since the orbitals are generally much smoother than the projectors, and therefore more accurate forces can be obtained for a given discretization [63, 41].²

Overview of SPARC. SPARC has been implemented in the framework of the Portable, Extensible Toolkit for scientific computations (PETSc) [65, 66] suite of data structures and routines. The approach adopted for geometry optimization/molecular dynamics is outlined in Fig. 1, whose key components are discussed in detail in the subsections below.

²The adopted approach introduces an additional approximation into the calculation of the forces, since the derivative of the projectors can in principle be evaluated exactly, whereas the gradient operator in the finite-difference setting is an approximate one. However, we have found that this is more than offset by the smoothness of the orbitals compared to the projectors. Other approaches to improve the quality of the forces include suitable modification of the pseudopotential [24], double-grid method [64], and high-order spatial integration [43], all of which can possibly be used in conjunction with the current approach to even further reduce the egg-box effect.

Figure 1: Outline of DFT simulations in SPARC for extended systems.

3.1. Finite-difference discretization

Let Ω denote the unit cell—a cuboid aligned with the x , y , and z axes (origin at the center) with sides of length L_1 , L_2 and L_3 , respectively. We discretize Ω using a finite-difference grid with spacings h_1 , h_2 , and h_3 along the x , y , and z axes, respectively, such that $L_1 = n_1 h_1$, $L_2 = n_2 h_2$ and $L_3 = n_3 h_3$ ($n_1, n_2, n_3 \in \mathbb{N}$, $\mathbb{N}$: set of all natural numbers). We designate each finite-difference node using an index of the form (i, j, k) , where $i = 1, 2, \dots, n_1$, $j = 1, 2, \dots, n_2$ and $k = 1, 2, \dots, n_3$. We approximate the Laplacian using the central finite-difference approximation:

$$\nabla_h^2 f|^{(i,j,k)} \approx \sum_{p=0}^{n_o} \left(w_{p,1}(f^{(i+p,j,k)} + f^{(i-p,j,k)}) + w_{p,2}(f^{(i,j+p,k)} + f^{(i,j-p,k)}) + w_{p,3}(f^{(i,j,k+p)} + f^{(i,j,k-p)}) \right), \quad (20)$$

where $f^{(i,j,k)}$ represents the value of the function f at (i, j, k) , and the weights [67, 50]

$$\begin{aligned} w_{0,s} &= -\frac{1}{h_s^2} \sum_{q=1}^{n_o} \frac{1}{q^2}, \quad s = 1, 2, 3, \\ w_{p,s} &= \frac{2(-1)^{p+1}}{h_s^2 p^2} \frac{(n_o!)^2}{(n_o - p)!(n_o + p)!}, \quad p = 1, 2, \dots, n_o, \quad s = 1, 2, 3. \end{aligned} \quad (21)$$

Similarly, we approximate the gradient operator using central finite-differences:

$$\nabla_h f|^{(i,j,k)} \approx \sum_{p=1}^{n_o} \left(\tilde{w}_{p,1}(f^{(i+p,j,k)} - f^{(i-p,j,k)})\hat{\mathbf{e}}_1 + \tilde{w}_{p,2}(f^{(i,j+p,k)} - f^{(i,j-p,k)})\hat{\mathbf{e}}_2 + \tilde{w}_{p,3}(f^{(i,j,k+p)} - f^{(i,j,k-p)})\hat{\mathbf{e}}_3 \right), \quad (22)$$

where $\hat{\mathbf{e}}_1$, $\hat{\mathbf{e}}_2$ and $\hat{\mathbf{e}}_3$ represent unit vectors along the x , y , and z axes, respectively. Further, the weights [67, 50]

$$\tilde{w}_{p,s} = \frac{(-1)^{p+1}}{h_s p} \frac{(n_o!)^2}{(n_o - p)!(n_o + p)!}, \quad p = 1, 2, \dots, n_o, \quad s = 1, 2, 3. \quad (23)$$

We employ the trapezoidal rule for performing spatial integrations [41], i.e.,

$$\int_{\Omega} f(\mathbf{x}) d\mathbf{x} \approx h_1 h_2 h_3 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} f^{(i,j,k)}, \quad (24)$$

using which we approximate the nonlocal pseudopotential operator as

$$\begin{aligned} V_{nl} f|^{(i,j,k)} &= \sum_J V_{nl,J} f|^{(i,j,k)} \\ &\approx h_1 h_2 h_3 \sum_J \sum_{lm} \sum_{p=1}^{n_1} \sum_{q=1}^{n_2} \sum_{r=1}^{n_3} \gamma_{Jl} \left(\sum_{J'} e^{-i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^{(i,j,k)} \right) \left(\sum_{J'} e^{i\mathbf{k} \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^{*(p,q,r)} f^{(p,q,r)} \right), \end{aligned} \quad (25)$$

where the summation index J runs over all atoms in Ω , and the summation index J' runs over the J^{th} atom and its periodic images. We enforce periodic boundary conditions by mapping any index that does not correspond to a node in the finite-difference grid to its periodic image within Ω . We enforce zero Dirichlet boundary conditions by setting $f^{(i,j,k)} = 0$ for any index that does not correspond to a node in the finite-difference grid.

Henceforth, we denote the discrete Hamiltonian matrix as a function of the Bloch wavevector $\mathbf{k}$ by $\mathbf{H}(\mathbf{k}) \in \mathbb{C}^{N_d \times N_d}$, where $N_d = n_1 \times n_2 \times n_3$ is the total number of finite-difference nodes. In addition, we represent the eigenvalues of $\mathbf{H}(\mathbf{k})$ arranged in ascending order by $\lambda_1(\mathbf{k}), \lambda_2(\mathbf{k}), \dots, \lambda_{N_d}(\mathbf{k})$. We store $\mathbf{H}(\mathbf{k})$ and other sparse matrices in compressed row format, and store the discrete $u_n(\mathbf{x}, \mathbf{k})$ as columns of the dense matrix $\mathbf{U}(\mathbf{k}) \in \mathbb{C}^{N_d \times N_s}$. In parallel computations, we partition the domain as $\Omega = \bigcup_{p=1}^{n_p} \Omega_p$, where Ω_p

denotes the domain local to the p^{th} processor, and n_p is the total number of processors. The specific choice of Ω_p corresponds to the PETSc default for structured grids.

3.2. Pseudocharge density generation and self energy calculation

In each structural relaxation/molecular dynamics step, the finite-difference Laplacian is used to assign the pseudocharge densities to the grid [17]:

$$b^{(i,j,k)} = \sum_I b_I^{(i,j,k)}, \quad b_I^{(i,j,k)} = -\frac{1}{4\pi} \nabla_h^2 V_I|^{(i,j,k)}, \quad (26)$$

where the summation index I runs over all atoms in $\mathbb{R}^3$. The discrete form of the pseudocharge density b_J has exponential decay away from $\mathbf{R}_J$ (Appendix B), which allows for its truncation at some suitably chosen

radius $r_{J'}^b$. The corresponding discrete self energy takes the form

$$E_{self}^h = \frac{1}{2} h_1 h_2 h_3 \sum_I \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} b_I^{(i,j,k)} V_I^{(i,j,k)}. \quad (27)$$

In Algorithm 1, we outline the calculation of $b^{(i,j,k)}$ and E_{self}^h , as implemented in SPARC. We use $P_{r^b}^p$ to denote the set of all atoms for which $\Omega_{r_{J'}^b} \cap \Omega_p \neq \emptyset$, where the index J' runs over all atoms in Ω and their periodic images. Here, $\Omega_{r_{J'}^b}$ represents the cuboid with side of lengths $2r_{J,1}^b \approx 2r_{J,2}^b \approx 2r_{J,3}^b (\approx 2r_J^b)$ centered on the J^{th} atom. We have chosen $\Omega_{r_{J'}^b}$ to be a cuboid rather than a sphere due to its simplicity and efficiency within the Euclidean finite-difference discretization. The values of $r_{J,1}^b$, $r_{J,2}^b$, and $r_{J,3}^b$ —integer multiples of the mesh spacings h_1 , h_2 , and h_3 , respectively—are chosen such that the charge constraint in Eqn. 13 is satisfied to within a prespecified tolerance ε_b , i.e.,

$$\left| \frac{h_1 h_2 h_3 \sum_{J'} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} b_{J'}^{(i,j,k)} - Z_J}{Z_J} \right| < \varepsilon_b. \quad (28)$$

For the atom indexed by J' belonging to $P_{r^b}^p$, we determine the overlap region $\Omega_{r_{J'}^b} \cap \Omega_p$, with the subscripts s and e denoting the starting and ending indices, respectively. In this overlap region (and an additional $2n_0$ points in each direction), we interpolate the values of $V_{J'}^{(i,j,k)}$ on to the finite-difference grid using cubic-splines [68]. Next, we determine $b^{(i,j,k)}$ and $E_{self}^{h,p}$ —contribution of the p^{th} processor to the self energy—using Eqns. 26 and 27, respectively. Finally, we calculate the total self energy E_{self}^h by summing the contributions from all the processors.

Algorithm 1: Pseudocharge density generation and self energy calculation

Input: $\mathbf{R}$, V_J , and r_J^b

$b^{(i,j,k)} = 0$, $E_{self}^{h,p} = 0$

for $J' \in P_{r^b}^p$ **do**

 Determine $i_s, i_e, j_s, j_e, k_s, k_e$ of $\Omega_{r_{J'}^b} \cap \Omega_p$

 Determine $V_{J'}^{(i,j,k)} \forall i \in [i_s - n_o, i_e + n_o], j \in [j_s - n_o, j_e + n_o], k \in [k_s - n_o, k_e + n_o]$

$b_{J'}^{(i,j,k)} = -\frac{1}{4\pi} \nabla_h^2 V_{J'}^{(i,j,k)}$, $b^{(i,j,k)} = b^{(i,j,k)} + b_{J'}^{(i,j,k)} \forall i \in [i_s, i_e], j \in [j_s, j_e], k \in [k_s, k_e]$

$E_{self}^{h,p} = E_{self}^{h,p} + \frac{1}{2} h_1 h_2 h_3 b_{J'}^{(i,j,k)} V_{J'}^{(i,j,k)} \forall i \in [i_s, i_e], j \in [j_s, j_e], k \in [k_s, k_e]$

$E_{self}^h = \sum_{p=1}^{n_p} E_{self}^{h,p}$

Output: $b^{(i,j,k)}$ and E_{self}^h

3.3. Brillouin zone integration

The volume averaged integral of any function over the Brillouin zone is approximated as

$$\oint_{BZ} f(\mathbf{k}) d\mathbf{k} \approx \sum_{b=1}^{N_k} w_b f_b, \quad (29)$$

³Note that $r_{J'}^b = r_J^b$, where the index J corresponds to any atom in Ω , and J' corresponds to the J^{th} atom and its periodic images

where $f_b \equiv f(\mathbf{k}_b)$. Here, $\mathbf{k}_b$ and w_b ($b = 1, 2, \dots, N_k$) denote the nodes and weights for integration, respectively. The specific choice of $\mathbf{k}_b$ is commonly referred to as Brillouin zone sampling.

3.4. Electron density calculation

On employing the Brillouin zone integration scheme described by Eqn. 29, the electron density takes the form

$$\rho(\mathbf{x}) = 2 \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b g_{nb} |u_{nb}(\mathbf{x})|^2. \quad (30)$$

In each SCF iteration, we calculate the electron density using the Chebyshev-filtered subspace iteration (CheFSI) method [69, 47], wherein we compute approximations to the eigenvectors corresponding to the lowest N_s eigenvalues of $\mathbf{H}_b$, $b = 1, 2, \dots, N_k$. This choice of eigensolver is motivated by the minimal orthogonalization and computer memory costs compared to other alternatives commonly employed in electronic structure calculations, e.g. Locally Optimal Block Preconditioned Conjugate Gradient (LOBPCG) [70]. Additionally, it has been shown to be extremely efficient within SPARC for isolated systems [41].

The implementation of the CheFSI algorithm in SPARC for extended systems consists of three main components. First, we use the rapid growth of Chebyshev polynomials outside the interval $[-1, 1]$ to filter out the unwanted eigencomponents from $\mathbf{U}_b$:

$$\mathbf{U}_{bf} = p_{mb}(\mathbf{H}_b) \mathbf{U}_b, \quad p_{mb}(t) = C_m \left(\frac{t - c_b}{e_b} \right), \quad b = 1, 2, \dots, N_k, \quad (31)$$

where the columns of $\mathbf{U}_{bf}$ represent the filtered $\mathbf{U}_b$, and C_m denotes the Chebyshev polynomial of degree m . Additionally, $e_b = (\lambda_{N_d b} - \lambda_c)/2$ and $c_b = (\lambda_{N_d b} + \lambda_c)/2$, where λ_c is the filter cutoff. Rather than explicitly compute the matrix $p_{mb}(\mathbf{H}_b)$, its product with $\mathbf{U}_b$ is determined using the three term recurrence relation for Chebyshev polynomials.

Next, projecting onto the filtered basis $\mathbf{U}_{bf}$, we arrive at the generalized eigenproblem:

$$\tilde{\mathbf{H}}_b \mathbf{y}_{nb} = \lambda_{nb} \tilde{\mathbf{M}}_b \mathbf{y}_{nb}, \quad n = 1, 2, \dots, N_s, \quad b = 1, 2, \dots, N_k, \quad (32)$$

where λ_{nb} represent approximations to the eigenvalues of $\mathbf{H}_b$, and the dense matrices $\tilde{\mathbf{H}}_b, \tilde{\mathbf{M}}_b \in \mathbb{C}^{N_s \times N_s}$ are determined as

$$\tilde{\mathbf{H}}_b = \mathbf{U}_{bf}^* \mathbf{H}_b \mathbf{U}_{bf}, \quad \tilde{\mathbf{M}}_b = \mathbf{U}_{bf}^* \mathbf{U}_{bf}. \quad (33)$$

After eigendecomposing Eqn. 32 at all the $\mathbf{k}$ -points, the Fermi energy λ_f is calculated by enforcing the constraint on the total number of electrons:

$$2 \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b g_{nb} = N_e, \quad \text{where} \quad g_{nb} = \left(1 + \exp \left(\frac{\lambda_{nb} - \lambda_f}{k_B T} \right) \right)^{-1}. \quad (34)$$

Finally, we perform the subspace rotation

$$\mathbf{U}_b = \mathbf{U}_{bf} \mathbf{Y}_b, \quad b = 1, 2, \dots, N_k, \quad (35)$$

where the columns of the matrix $\mathbf{Y}_b \in \mathbb{C}^{N_s \times N_s}$ contain the eigenvectors $\mathbf{y}_{nb}$. The columns of $\mathbf{U}_b$ so obtained represent approximations to the eigenvectors of $\mathbf{H}_b$. The electron density at the finite-difference grid points is then calculated using the relation

$$\rho^{(i,j,k)} = 2 \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b g_{nb} |u_{nb}^{(i,j,k)}|^2, \quad (36)$$

where the values of $u_{nb}^{(i,j,k)}$ are extracted from the n^{th} column of $\mathbf{U}_b$.

We start the very first SCF iteration of the complete DFT simulation with a randomly generated guess for $\mathbf{U}_b$ ($b = 1, 2, \dots, N_k$), and perform the CheFSI steps multiple times [71] without calculating/updating the electron density. This allows us to obtain a good approximation of the electron density for the second SCF iteration. For every subsequent SCF iteration, we perform the CheFSI steps only once, and use the rotated $\mathbf{U}_b$ from the previous step as the initial guess. Overall, the calculation of the electron density scales as $\mathcal{O}(N_k N_s N_d) + \mathcal{O}(N_k N_s^2 N_d) + \mathcal{O}(N_k N_s^3)$, which makes it $\mathcal{O}(N^3)$ with respect to the number of atoms.

3.5. Free energy calculation

On approximating the spatial integrals using Eqn. 24 and the Brillouin zone integrals using Eqn. 29, we arrive at the discrete form for the free energy at the electronic ground-state:

$$\begin{aligned} \hat{\mathcal{F}}^h = & 2 \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b g_{nb} \lambda_{nb} + h_1 h_2 h_3 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} \left(\varepsilon_{xc}^{(i,j,k)} \rho^{(i,j,k)} - V_{xc}^{(i,j,k)} \rho^{(i,j,k)} + \frac{1}{2} (b^{(i,j,k)} - \rho^{(i,j,k)}) \phi^{(i,j,k)} \right) \\ & - E_{self}^h + E_c^h + 2k_B T \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b (g_{nb} \log g_{nb} + (1 - g_{nb}) \log(1 - g_{nb})) , \end{aligned} \quad (37)$$

where E_{self}^h is the discrete self energy (Eqn. 27), and E_c^h is the discrete repulsive energy correction for overlapping pseudocharges (Eqn. 50). The evaluation of $\hat{\mathcal{F}}^h$ scales as $\mathcal{O}(N_d)$, and therefore $\mathcal{O}(N)$ with respect to the number of atoms.

3.6. Atomic forces calculation

The discrete form of the atomic force presented in Eqn. 19 can be split into three parts:

$$\mathbf{f}_J^h = \mathbf{f}_{J,loc}^h + \mathbf{f}_{J,c}^h + \mathbf{f}_{J,nloc}^h , \quad (38)$$

where $\mathbf{f}_{J,loc}^h$ is the discrete local component, $\mathbf{f}_{J,c}^h$ is the discrete correction corresponding to overlapping pseudocharges, and $\mathbf{f}_{J,nloc}^h$ is the discrete non-local component of the force. Below, we discuss the evaluation of $\mathbf{f}_{J,loc}^h$ and $\mathbf{f}_{J,nloc}^h$ in SPARC for extended systems, with the computation of $\mathbf{f}_{J,c}^h$ proceeding similarly to $\mathbf{f}_{J,loc}^h$.

Local component. On approximating the spatial integral using Eqn. 24, the discrete local component of the force takes the form

$$\mathbf{f}_{J,loc}^h = h_1 h_2 h_3 \sum_{J'} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} \nabla_h b_{J'} \big|^{(i,j,k)} (\phi^{(i,j,k)} - V_{J'}^{(i,j,k)}) , \quad (39)$$

where the summation index J' runs over the J^{th} atom and its periodic images. The calculation of $\mathbf{f}_{J,loc}^h$ is outlined in Algorithm 2, which proceeds as follows. For the atom indexed by J' belonging to $P_{r_b}^p$, we determine the overlap region $\Omega_{r_b}^{J'} \cap \Omega_p$. In this overlap region (and an additional $4n_0$ points in each direction), the values of $V_{J'}^{(i,j,k)}$ are interpolated on to the finite-difference grid using cubic-splines, from which $b_{J'}^{(i,j,k)}$ is calculated using Eqn. 26. Subsequently, $\mathbf{f}_{J,loc}^{h,p}$ —contribution of the p^{th} processor to the

local component of the force—is calculated using Eqn. 39. Finally, the contributions from all processors are summed to obtain $\mathbf{f}_{J,loc}^h$.

Algorithm 2: Calculation of the local component of the atomic force.

Input: $\mathbf{R}$, $\phi^{(i,j,k)}$, V_J , and r_J^b

$\mathbf{f}_{J,loc}^{h,p} = 0$

for $J' \in P_{r_J^b}^p$ **do**

 Determine $i_s, i_e, j_s, j_e, k_s, k_e$ of $\Omega_{r_{J'}^b} \cap \Omega_p$

 Determine $V_{J'}^{(i,j,k)} \forall i \in [i_s - 2n_o, i_e + 2n_o], j \in [j_s - 2n_o, j_e + 2n_o]$, and

$k \in [k_s - 2n_o, k_e + 2n_o]$

$b_{J'}^{(i,j,k)} = -\frac{1}{4\pi} \nabla_h^2 V_{J'}^{(i,j,k)} \forall i \in [i_s - n_o, i_e + n_o], j \in [j_s - n_o, j_e + n_o], k \in [k_s - n_o, k_e + n_o]$

$\mathbf{f}_{J,loc}^{h,p} = \mathbf{f}_{J,loc}^{h,p} + h_1 h_2 h_3 \sum_{i=i_s}^{i_e} \sum_{j=j_s}^{j_e} \sum_{k=k_s}^{k_e} \nabla_h b_{J'}^{(i,j,k)} (\phi^{(i,j,k)} - V_{J'}^{(i,j,k)})$

$\mathbf{f}_{J,loc}^h = \sum_{p=1}^{n_p} \mathbf{f}_{J,loc}^{h,p}$

Output: $\mathbf{f}_{J,loc}^h$

Non-local component. On approximating the spatial integral using Eqn. 24, the discrete non-local component of the force takes the form

$$\mathbf{f}_{J,nloc}^h = -4 \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b g_{nb} \sum_{lm} \gamma_{Jl} \Re [Y_{Jnblm} \mathbf{W}_{Jnblm}] , \quad (40)$$

where

$$Y_{Jnblm} = h_1 h_2 h_3 \sum_{J'} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} \psi_n^{*(i,j,k)} e^{-i\mathbf{k}_b \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^{(i,j,k)} , \quad (41)$$

$$\mathbf{W}_{Jnblm} = h_1 h_2 h_3 \sum_{J'} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} \nabla_h \psi_n^{(i,j,k)} e^{i\mathbf{k}_b \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^{*(i,j,k)} . \quad (42)$$

In Algorithm 3, we outline the calculation of $\mathbf{f}_{J,nloc}^h$. We use $P_{r_J^c}^p$ to denote the set of all atoms for which $\Omega_{r_{J'}^c} \cap \Omega_p \neq \emptyset$, where the index J' runs over all atoms in Ω and their periodic images. In addition, $\Omega_{r_{J'}^c}$ represents the cuboid with side of lengths $2r_{J,1}^c \approx 2r_{J,2}^c \approx 2r_{J,3}^c (\approx 2r_J^c)$ —integer multiples of the spacings h_1, h_2 , and h_3 , respectively—centered on the J'^{th} atom. We have chosen $\Omega_{r_{J'}^c}$ to be a cuboid rather than a sphere due to its simplicity and efficiency within the Euclidean finite-difference discretization. For the atom indexed by J' belonging to $P_{r_J^c}^p$, we determine the overlap region $\Omega_{r_{J'}^c} \cap \Omega_p$, with the subscripts s and e denoting the starting and ending indices, respectively. In this overlap region, we interpolate the radial components of the projectors $\chi_{J'lm}^{(i,j,k)}$ on to the finite-difference grid using cubic-splines. Next, we utilize Eqns. 41 and 42 to determine Y_{Jnblm}^p and $\mathbf{W}_{Jnblm}^p$, respectively, which represent the contributions of the p^{th} processor to Y_{Jnblm} and $\mathbf{W}_{Jnblm}$, respectively. Finally, we sum the contributions from all the processors to obtain Y_{Jnblm} and $\mathbf{W}_{Jnblm}$, which are then used to calculate $\mathbf{f}_{J,nloc}^h$ using Eqn. 40. Overall, the calculation of

the atomic forces scales as $\mathcal{O}(N)$ with respect to the number of atoms.

Algorithm 3: Calculation of the non-local component of the atomic force

Input: $\mathbf{R}$, $\psi_n^{(i,j,k)}$, γ_{Jl} , χ_{Jlm} , and r_J^c
 $Y_{Jnblm}^p = 0$, $\mathbf{W}_{Jnblm}^p = 0$
for $J' \in P_{r_J^c}^p$ **do**
 Determine starting and ending indices $i_s, i_e, j_s, j_e, k_s, k_e$ for $\Omega_{r_{J'}}^c \cap \Omega_p$
 Determine $\chi_{J'lm}^{(i,j,k)} \forall i \in [i_s, i_e], j \in [j_s, j_e], k \in [k_s, k_e]$
 $Y_{Jnblm}^p = Y_{Jnblm}^p + h_1 h_2 h_3 \psi_n^{*(i,j,k)} e^{-i\mathbf{k}_b \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^{(i,j,k)} \forall i \in [i_s, i_e], j \in [j_s, j_e], k \in [k_s, k_e]$
 $\mathbf{W}_{Jnblm}^p = \mathbf{W}_{Jnblm}^p + h_1 h_2 h_3 \nabla_h \psi_n^{(i,j,k)} e^{i\mathbf{k}_b \cdot (\mathbf{R}_J - \mathbf{R}_{J'})} \chi_{J'lm}^{*(i,j,k)} \forall i \in [i_s, i_e], j \in [j_s, j_e], k \in [k_s, k_e]$
 $Y_{Jnblm} = \sum_{p=1}^{n_p} Y_{Jnblm}^p$, $\mathbf{W}_{Jnblm} = \sum_{p=1}^{n_p} \mathbf{W}_{Jnblm}^p$
 $\mathbf{f}_{J,nloc}^h = -4 \sum_{n=1}^{N_s} \sum_{b=1}^{N_k} w_b g_{nb} \sum_{lm} \gamma_{Jl} \Re[Y_{Jnblm} \mathbf{W}_{Jnblm}]$
Output: $\mathbf{f}_{J,nloc}^h$

4. Examples and Results

In this section, we verify the accuracy and efficiency of SPARC for performing DFT calculations of extended systems. In all the examples, we utilize a twelfth-order accurate finite-difference discretization ($n_0 = 6$), the Perdew-Wang parametrization [72] of the correlation energy calculated by Ceperley-Alder [73], norm-conserving Troullier-Martins pseudopotentials [74], and a smearing of $k_B T = 0.01$ Ha. The cutoff radii for the non-local projectors and the local component of the pseudopotentials are listed in Appendix B. We utilize the Monkhorst-Pack [75] grid for performing integrations over the Brillouin zone. Unless specified otherwise, the simulations correspond to $\mathbf{k} = [0.0 \ 0.0 \ 0.0]$ (Γ -point). We use the notation h to denote the mesh-size when a common spacing is employed in all directions, i.e., $h_1 = h_2 = h_3 = h$.

We solve the linear system corresponding to the Poisson problem using the block-Jacobi preconditioned [76] minimal residual method (MINRES) [77]. Within the CheFSI approach, we utilize a polynomial of degree $m = 20$ for Chebyshev filtering; the Lanczos method [78] for calculating the extremal eigenvalues of the Hamiltonian $\mathbf{H}_b$; and LAPACK's [79] implementation of the QR algorithm [80] for solving the subspace eigenproblem in Eqn. 32. We calculate the Fermi energy using Brent's method [81], and use Anderson extrapolation [82] with relaxation parameter of 0.3 and mixing history of 7 for accelerating the convergence of the SCF method. We employ the Polak-Ribiere variant of non-linear conjugate gradients with a secant line search [83] for geometry optimization. Finally, we use the leapfrog integration scheme [84] for performing NVE molecular dynamics.

All the results presented here are converged to within the 'chemical accuracy' of 0.001 Ha/atom in the energy and 0.001 Ha/Bohr in the forces.⁴ Wherever applicable, the results obtained by SPARC are compared to the well-established plane-wave code ABINIT [9, 85, 86]. The error in energy is defined to be the difference in the magnitude, and the error in forces is defined to be the maximum difference in any component. All simulations are performed on a computer cluster consisting of 16 nodes with the following configuration: Altus 1804i Server - 4P Interlagos Node, Quad AMD Opteron 6276, 16C, 2.3 GHz, 128GB,

⁴In this section, we will use the term 'chemical accuracy' to denote the convergence of energy and atomic forces—with respect to the fully converged DFT results—to within 0.001 Ha/atom and 0.001 Ha/Bohr, respectively.

DDR3-1333 ECC, 80GB SSD, MLC, 2.5" HCA, Mellanox ConnectX 2, 1-port QSFP, QDR, memfree, CentOS, Version 5, and connected through InfiniBand cable.

4.1. Convergence with discretization

First, we verify convergence of the energy and atomic forces computed by SPARC with respect to spatial discretization. For this study, we choose three examples: $2 \times 2 \times 2$ unit cells of lithium hydride with lattice constant of 7.37 Bohr and corner lithium atom perturbed by [0.57 0.43 0.37] Bohr; $2 \times 2 \times 2$ unit cells of silicon with lattice constant of 10.68 Bohr and corner atom perturbed by [0.93 0.50 0.20] Bohr; and $2 \times 2 \times 2$ unit cells of gold with lattice constant of 8.0 Bohr, $1 \times 1 \times 2$ Brillouin zone integration, and a face-centered atom perturbed by [0.85 0.50 0.30] Bohr. All errors are defined with respect to ABINIT, wherein we employ plane-wave cutoffs of 46, 40, and 46 Ha for lithium hydride, silicon, and gold, respectively. This results in energy and forces that are converged to within 5×10^{-6} Ha/atom and 5×10^{-6} Ha/Bohr, respectively.

It is clear from Fig. 2—plots of the error in energy and atomic forces with respect to mesh-size—that there is systematic convergence to the reference plane-wave result. On performing a fit to the data, we obtain average convergence rates of approximately $\mathcal{O}(h^{10})$ in the energy and $\mathcal{O}(h^9)$ in the forces. In doing so, the chemical accuracy desired in electronic structure calculations is readily attained. These results demonstrate that SPARC is able to obtain high convergence rates in both the energy and forces, which contributes to its accuracy and efficiency. Moreover, the energies and forces in SPARC converge at comparable rates, without the need for additional measures such as double-grid [64] or high-order integration [43] techniques.⁵

Figure 2: Convergence of the energy and atomic forces with respect to mesh size to reference planewave result for the lithium hydride, silicon, and gold systems. The straight lines represent linear fits to the data.

4.2. Convergence with vacuum size for slabs and wires

Next, we verify the convergence of the energy and atomic forces computed by SPARC with respect to vacuum size for slabs and wires. As representative examples, we choose a $1 \times 1 \times 5$ aluminum slab with lattice constant 7.78 Bohr; and a silicon nanowire of lattice constant 10.16 Bohr, radius 15 Bohr and length

⁵Since the calculation of the forces involves taking derivatives of the energy with respect to the atomic positions, they are susceptible to larger egg-box effects, which contributes to a typically slower convergence of the forces relative to the energy [43].

10.16 Bohr. The slab has vacuum in the z direction, and the wire has vacuum in the y and z directions. We utilize mesh-spacings of $\{h_1, h_2, h_3\} = \{0.5985, 0.5985, 0.6\}$ Bohr and $\{h_1, h_2, h_3\} = \{0.406, 0.4, 0.4\}$ Bohr for the slab and wire, respectively.

In Fig. 3, we present convergence of the energy and atomic forces with vacuum size for SPARC and ABINIT, wherein the results obtained for a vacuum of 18 Bohr are used as reference. We observe that SPARC achieves exponential convergence in both the energy and forces to accuracies well below those desired in DFT calculations. In particular, even a vacuum of around 7 Bohr is sufficient to obtain chemical accuracy in both energy and forces.⁶ However, in ABINIT, there is a stagnation of the error, which can be attributed to the spurious interactions between images resulting from the need for periodic boundary conditions.

Figure 3: Convergence of energy and atomic forces with respect to vacuum size for the aluminum slab and silicon wire.

4.3. Bulk properties

We now verify the ability of SPARC to accurately calculate material bulk properties. We select silicon—8-atom unit cell with $4 \times 4 \times 4$ Brillouin zone integration—as the representative example. In SPARC, we use a mesh-size of $h = 0.407$ Bohr. We compare the results with ABINIT, wherein we choose a plane-wave energy cutoff of 40 Ha, which results in energies that are converged to within 5×10^{-6} Ha/atom. In Fig. 4, we plot the energy so computed by SPARC and ABINIT as a function of the lattice constant. We observe that there is very good agreement between SPARC and ABINIT, with the curves being practically indistinguishable. From a cubic fit to the data, we find that the predicted equilibrium lattice constant, energy, and bulk modulus are in agreement to within 0.003 Bohr, 1×10^{-5} Ha/atom, and 0.18 GPa, respectively.

⁶Indeed, the convergence with vacuum size can be further accelerated in SPARC (particularly for systems with large dipole moments) by incorporating more accurate boundary conditions for slabs and wires [16].

Figure 4: Variation of energy with lattice constant for bulk silicon.

Next, we compare the band structure plot at the equilibrium lattice constants determined above, i.e., 10.157 Bohr for SPARC and 10.160 Bohr for ABINIT. Specifically, we choose the $L - \Gamma - X - \Gamma$ circuit, whose coordinates in terms of the reciprocal lattice vectors are $[-0.5 \ 0.5 \ 0.5]$, $[0.0 \ 0.0 \ 0.0]$, $[1.0 \ 0.0 \ 0.0]$, and $[1.0 \ 1.0 \ 1.0]$, respectively. We discretize the $L - \Gamma$, $\Gamma - X$, and $X - \Gamma$ line segments into 10, 12, and 17 divisions, respectively. At each resulting k -point, we determine the band structure (at the electronic ground-state) in SPARC by repeating the CheFSI steps until convergence. In Fig. 5, we present the band structure plot so computed by SPARC and compare it with that calculated by ABINIT. It is clear that there is very good agreement, with the curves being nearly identical. In particular, the HOMO eigenvalue, LUMO eigenvalue, and bandgap are in agreement to within 7×10^{-5} Ha, 1×10^{-5} Ha, and 6×10^{-5} Ha, respectively.

Figure 5: Band structure plot for bulk silicon.

4.4. Slab and wire properties

Here, we verify ability of SPARC to accurately calculate the properties of slabs and wires. For this purpose, we choose the aluminum slab ($1 \times 1 \times 1$ Brillouin zone integration) and silicon wire ($4 \times 1 \times 1$ Brillouin zone integration) described in Section 4.2 as representative examples. In SPARC, we utilize a vacuum of 6.95 Bohr and 5.59 Bohr for the slab and wire, respectively. In ABINIT, the corresponding values are 12.55 Bohr and 10 Bohr, respectively. The energy and forces so computed in ABINIT are accurate to

within 1×10^{-4} Ha/atom and 1×10^{-4} Ha/Bohr, respectively.⁷

First, we calculate the unrelaxed surface energy of the aluminum slab with cross-sectional area A :

$$\gamma = N \left(\frac{\hat{\mathcal{F}}_{slab} - \hat{\mathcal{F}}_{bulk}}{2A} \right), \quad (43)$$

where $\hat{\mathcal{F}}$ is the free energy/atom of the slab, and $\hat{\mathcal{F}}_{bulk}$ is the free energy/atom of bulk aluminum. The surface energies so calculated by SPARC and ABINIT are in agreement to within 0.001 J/m². In addition, the maximum difference in the atomic forces is within 3×10^{-4} Ha/Bohr. In Fig. 6a, we present the contours of the electron density computed by SPARC on the $y = -3.89$ Bohr plane of the slab.

Figure 6: Electron density contours for the aluminum slab and silicon wire.

Next, we determine the density of states (DOS) of the silicon wire using the relation

$$D(\mathcal{E}) = 2 \sum_{n=1}^{N_s} \oint_{BZ} \delta(\lambda_n(\mathbf{k}) - \mathcal{E}) d\mathbf{k}, \quad (44)$$

where δ is the Dirac delta function. In Fig. 7, we present the DOS computed by SPARC and ABINIT. It is clear that there is excellent agreement between SPARC and ABINIT, with the curves being practically indistinguishable. In addition, the energy and atomic forces are in agreement to within 7×10^{-5} Ha/atom and 3×10^{-4} Ha/Bohr, respectively. In Fig. 6b, we present the contours of the electron density computed by SPARC on the $x = 0$ Bohr plane of the wire.

⁷As shown in Section 4.2, there is stagnation in the convergence of the energy and atomic forces with vacuum size in ABINIT, which limits its accuracy.

Figure 7: Density of states for the silicon wire.

4.5. Geometry optimization

We now verify the capacity of SPARC to perform accurate geometry optimizations. To do so, we first check the consistency of the atomic forces with the energy. As representative examples, we select single unit cells of lithium and aluminum with lattice constants of 6.59 Bohr and 8 Bohr, respectively, and discretize them using mesh sizes of $h = 0.549$ Bohr and 0.615 Bohr, respectively. In Fig. 8, we plot the variation in energy and force when the body centered lithium atom is displaced along the body diagonal, and the corner aluminum atom is displaced along the body diagonal. Specifically, in Fig. 8a, we plot the computed energy and its curve fit using cubic splines. In Fig. 8b, we plot the computed atomic force and the derivative of the cubic spline fit to the energy. The evident agreement demonstrates that the computed energy and atomic forces are indeed consistent. Moreover, there is no noticeable ‘egg-box’ effect [87]—a phenomenon arising due to the breaking of the translational symmetry—at meshes required for obtaining chemical accuracies.⁸

⁸This effect can further be reduced by using the double grid method [64], high-order spatial integration [43], or suitably modifying the pseudopotential [23, 24]

(a) Computed energy and its cubic spline fit

(b) Computed force and the derivative of the cubic spline fit to the energy

Figure 8: Variation in the energy and atomic force as a function of atomic displacement. In lithium, the body centered lithium atom is displaced along the body diagonal. In aluminum, the corner atom is displaced along the body diagonal.

Next, we determine the overall ground-state for two systems: $3 \times 3 \times 3$ unit cells of lithium and $2 \times 2 \times 2$ unit cells of aluminum, both with a vacancy. We compare the results with ABINIT, wherein we use a plane wave cutoff of 30 Ha, which results in energy and forces that are converged to within 5×10^{-6} Ha/atom and 5×10^{-6} Ha/Bohr, respectively. We calculate the vacancy formation energy $\mathcal{E}_{vf}$ using the relation [88]

$$\mathcal{E}_{vf} = \mathcal{F}_0 \left(N-1, 1, \frac{N-1}{N} \Omega \right) - \left(\frac{N-1}{N} \right) \mathcal{F}_0(N, 0, \Omega), \quad (45)$$

where $\mathcal{F}_0(N, n_v, \Omega)$ denotes the energy of the system with N occupied lattice sites and n_v vacancies. The vacancy formation energy so computed by SPARC and ABINIT is in agreement to within 7×10^{-4} Ha. In addition, the fully relaxed atomic positions differ by no more than 1.6×10^{-3} Bohr. On refining the mesh, the agreement between SPARC and ABINIT further improves. For example, consider a mesh-size of $h = 0.284$ Bohr and $h = 0.390$ Bohr for lithium and aluminum, respectively. The vacancy formation energy so computed by SPARC and ABINIT are in agreement to within 9×10^{-6} Ha, and the maximum difference in the final atomic positions is 1.3×10^{-4} Bohr. The contours of electron density on the mid-plane of these systems are plotted in Fig. 9.

(a) $3 \times 3 \times 3$ unit cells of lithium with a vacancy

(b) $2 \times 2 \times 2$ unit cells of aluminum with a vacancy

Figure 9: Mid-plane electron density contours.

4.6. Molecular dynamics

We now study the capability of SPARC to perform accurate Born-Oppenheimer molecular dynamics calculations [89]. Specifically, we perform a 1 ps NVE molecular dynamics simulation for the aluminum system consisting of $5 \times 5 \times 5$ FCC unit cells (500 atoms). We utilize a mesh-size of $h = 0.667$ Bohr, an initial ionic temperature of $T = 3157.75$ K, and a time step of 1 fs. We assign identical initial velocities to all the atoms and set the initial accelerations to be zero. At every molecular dynamics step, we set the electronic temperature (smearing) to be equal to the ionic temperature.

In Fig. 10, we plot the variation in the energy (free energy, kinetic energy, and total energy) and temperature as a function of time. The plots start at 40 fs, which corresponds to the time required to achieve statistical equilibrium. During the simulation, the mean and standard deviation of the total energy is -2.07842 Ha/atom and 1.6×10^{-4} Ha/atom, respectively. The drift in total energy as obtained from the linear fit to the data is 1×10^{-4} Ha/atom-ps.⁹ Overall, the lack of any significant drift verifies that there are no systematic errors in SPARC, which further confirms the accuracy of the energy and atomic forces in SPARC.

⁹We have found that the drift in the total energy further reduces as the spatial discretization is refined.

Figure 10: The variation of total energy, free energy, kinetic energy and temperature during the NVE molecular dynamics simulation for $5 \times 5 \times 5$ FCC unit cells of aluminum at an initial temperature of $T = 3157.75$ K.

4.7. Scaling and Performance

Having verified the accuracy of SPARC relative to ABINIT in previous subsections, we now compare their efficiency¹⁰. As representative systems, we choose $n \times n \times n$ ($n \in \{3, 4, 5, 6, 7\}$) unit cells of aluminum with a vacancy. In SPARC, we employ a mesh-size of $h = 0.778$ Bohr and Chebyshev polynomial filter of degree 15. In ABINIT, we use a plane-wave energy cutoff of 9 Ha. We choose all the other parameters in both codes so as to achieve an overall accuracy of 5×10^{-4} Ha/atom and 5×10^{-4} Ha/Bohr in the energy and forces, respectively. The times reported here include the calculation of the electronic ground-state as well as the atomic forces, i.e., geometry optimization is not performed.

First, we compare the strong scaling of SPARC and ABINIT for $6 \times 6 \times 6$ FCC unit cells of aluminum with a vacancy. We utilize 4, 8, 64, 144, 480, and 576 cores for performing the simulation with ABINIT, which it suggests is optimal in the range of 1 to 1000 cores. For SPARC, we select 4, 8, 27, 128, 384, and

¹⁰A comparison of the accuracy and efficiency of SPARC with other finite-difference DFT codes is presented in Appendix D.

512 cores. In Fig. 11a, we plot the wall time taken by SPARC and ABINIT as a function of the number of processors. We observe that both SPARC and ABINIT display similar trends with respect to strong scaling, with curves being close to parallel and no further reduction in wall time observed after approximately 600 cores. However, the prefactors of SPARC are significantly smaller, by up to a factor of 4.

Next, we compare the weak scaling of SPARC with ABINIT for $3 \times 3 \times 3$, $4 \times 4 \times 4$, $5 \times 5 \times 5$, $6 \times 6 \times 6$, and $7 \times 7 \times 7$ unit cells of aluminum, each with a vacancy. The number of electrons in these systems range from 321 to 4116. For both SPARC and ABINIT, we fix the number of electrons per core to be approximately 96, and choose at most 4 cores from every compute node. In Figure 11b, we present the results so obtained for the variation in total CPU time versus the number of electrons. We determine the scaling for SPARC and ABINIT to be $\mathcal{O}(N_e^{2.64})$ and $\mathcal{O}(N_e^{2.99})$ respectively. The prefactor for SPARC is again noticeably smaller, with speedups over ABINIT ranging from factors of 2.2 to 6.

Figure 11: Strong and weak scaling for SPARC and ABINIT. The system utilized for strong scaling is $6 \times 6 \times 6$ FCC unit cells of aluminum with a vacancy. The systems employed for weak scaling are $3 \times 3 \times 3$, $4 \times 4 \times 4$, $5 \times 5 \times 5$, $6 \times 6 \times 6$, and $7 \times 7 \times 7$ unit cells of aluminum, each with a vacancy.

Finally, we compare the minimum wall time achievable by SPARC and ABINIT for the aforementioned systems. We restrict the maximum number of electrons per core to 96. In SPARC, we choose the number of cores as multiples of 32, whereas we select the number of cores and parallelization scheme in ABINIT that it suggests are optimal. In Table 1, we present the results for the minimum wall time so achieved. We observe that SPARC outperforms ABINIT by factors larger than 2.5 for all the systems considered. In particular, SPARC requires a factor of approximately 5 less wall time for the largest system. Overall, we conclude that SPARC is a very efficient DFT formulation and implementation that is very competitive with currently existing highly optimized plane-wave codes for extended systems.¹¹

¹¹In previous work [41], we have already shown that SPARC is highly competitive with plane-wave codes for isolated systems.

System	SPARC	ABINIT
$3 \times 3 \times 3$	0.43 (64)	1.77 (188)
$4 \times 4 \times 4$	1.69 (96)	5.67 (320)
$5 \times 5 \times 5$	6.99 (256)	17.42 (396)
$6 \times 6 \times 6$	20.98 (640)	113.87 (480)
$7 \times 7 \times 7$	79.15 (704)	398.77 (795)

Table 1: Minimum wall time in minutes for $n \times n \times n$ ($n \in \{3, 4, 5, 6, 7\}$) FCC unit cells of aluminum with a vacancy. The number in brackets denotes the number of cores on which the minimum wall time is achieved.

5. Concluding Remarks

In this work, we have extended the capabilities of SPARC (Simulation Package for Ab-initio Real-space Calculations)—finite-difference formulation and parallel implementation of Density Functional Theory (DFT)—to enable the study of extended systems like crystals, slabs, and wires. Specifically, utilizing the Chebyshev polynomial filtered self-consistent field iteration in combination with the reformulation of the electrostatics and the non-local atomic force component, we have developed a framework that enables the efficient evaluation of energies and atomic forces to within the desired accuracies in DFT while employing the finite-difference representation.

Using a wide variety of materials systems, we have demonstrated that SPARC obtains systematic and high rates of convergence in the energy and forces with mesh-size to reference plane-wave results; exponential convergence with vacuum size for slabs and wires; energy and forces that are consistent and have negligible ‘egg-box’ effect; accurate ground-state properties; and negligible drift in molecular dynamics simulations. We have also shown that SPARC displays weak and strong scaling that is similar to well-established and optimized plane-wave codes for systems having up to thousands of electrons, but with a noticeably smaller prefactor. Additionally, we have established that SPARC significantly outperforms other finite-difference DFT packages.

We conclude by noting that there is scope for significant improvement of the current SPARC implementation. In particular, enabling the parallelization over bands is expected to enhance the performance of SPARC, especially in the context of strong scaling. Further, incorporating scalable diagonalization techniques for the solution of the subspace eigenproblem in parallel will enable the study of significantly larger systems. These improvements along with additional optimization of code are expected to further improve the efficiency of SPARC. Finally, removing the dependency on external packages (i.e., PETSc and MKL) is expected to significantly improve the portability and usability of SPARC, making it another worthy subject for future work.

Acknowledgements

The authors gratefully acknowledge the support of National Science Foundation under Grant Number 1333500. The authors also gratefully acknowledge the valuable comments and suggestions of the anonymous referee.

References

- [1] P. Hohenberg, W. Kohn, Physical Review 136 (1964) B864–B871.
- [2] W. Kohn, L. J. Sham, Physical Review 140 (1965) A1133–A1138.

- [3] R. O. Jones, O. Gunnarsson, *Reviews of Modern Physics* 61 (1989) 689.
- [4] T. Ziegler, *Chemical Reviews* 91 (1991) 651–667.
- [5] W. Kohn, A. D. Becke, R. G. Parr, *The Journal of Physical Chemistry* 100 (1996) 12974–12980.
- [6] R. O. Jones, *Reviews of modern physics* 87 (2015) 897.
- [7] G. Kresse, J. Furthmüller, *Physical Review B* 54 (1996) 11169–11186.
- [8] M. D. Segall, P. J. D. Lindan, M. J. Probert, C. J. Pickard, P. J. Hasnip, S. J. Clark, M. C. Payne, *Journal of Physics: Condensed Matter* 14 (2002) 2717–2744.
- [9] 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. C. Allan, *Computational Materials Science* 25 (2002) 478–492(15).
- [10] 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. Gougoussis, 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, *Journal of Physics: Condensed Matter* 21 (2009) 395502 (19pp).
- [11] S. Ismail-Beigi, T. A. Arias, *Computer Physics Communications* 128 (2000) 1 – 45.
- [12] F. Gygi, *IBM Journal of Research and Development* 52 (2008) 137–144.
- [13] J. Cooley, J. Tukey, *Mathematics of Computation* 19 (1965) 297.
- [14] M. C. Payne, M. P. Teter, D. C. Allan, T. Arias, J. Joannopoulos, *Reviews of Modern Physics* 64 (1992) 1045–1097.
- [15] J. Hutter, H. P. Lüthi, M. Parrinello, *Computational Materials Science* 2 (1994) 244–248.
- [16] A. Natan, A. Benjamini, D. Naveh, L. Kronik, M. L. Tiago, S. P. Beckman, J. R. Chelikowsky, *Physical Review B* 78 (2008) 075109.
- [17] P. Suryanarayana, K. Bhattacharya, M. Ortiz, *Journal of the Mechanics and Physics of Solids* 61 (2013) 38 – 60.
- [18] F. Bottin, S. Leroux, A. Knyazev, G. Zerah, *Computational Materials Science* 42 (2008) 329–336.
- [19] M. E. Tuckerman, D. Yarne, S. O. Samuelson, A. L. Hughes, G. J. Martyna, *Computer Physics Communications* 128 (2000) 333–376.
- [20] S. Goedecker, *Rev. Mod. Phys.* 71 (1999) 1085–1123.
- [21] D. R. Bowler, T. Miyazaki, *Reports on Progress in Physics* 75 (2012) 036503.
- [22] J. R. Chelikowsky, N. Troullier, Y. Saad, *Physical review letters* 72 (1994) 1240.

- [23] A. Castro, H. Appel, M. Oliveira, C. A. Rozzi, X. Andrade, F. Lorenzen, M. A. L. Marques, E. K. U. Gross, A. Rubio, *Physica Status Solidi B-Basic Solid State Physics* 243 (2006) 2465–2488.
- [24] E. Briggs, D. Sullivan, J. Bernholc, *Physical Review B* 54 (1996) 14362.
- [25] J.-L. Fattebert, *Journal of Computational Physics* 149 (1999) 75–94.
- [26] F. Shimojo, R. K. Kalia, A. Nakano, P. Vashishta, *Computer Physics Communications* 140 (2001) 303–314.
- [27] J.-I. Iwata, D. Takahashi, A. Oshiyama, T. Boku, K. Shiraishi, S. Okada, K. Yabana, *Journal of Computational Physics* 229 (2010) 2339 – 2363.
- [28] J. E. Pask, B. M. Klein, C. Y. Fong, P. A. Sterne, *Physical Review B* 59 (1999) 12352–12358.
- [29] S. R. White, J. W. Wilkins, M. P. Teter, *Physical Review B* 39 (1989) 5819.
- [30] E. Tsuchida, M. Tsukada, *Physical Review B* 52 (1995) 5573.
- [31] P. Suryanarayana, V. Gavini, T. Blesgen, K. Bhattacharya, M. Ortiz, *Journal of the Mechanics and Physics of Solids* 58 (2010) 256 – 280.
- [32] P. Motamarri, M. Iyer, J. Knap, V. Gavini, *Journal of Computational Physics* 231 (2012) 6596–6621.
- [33] J. Fang, X. Gao, A. Zhou, *Journal of Computational Physics* 231 (2012) 3166–3180.
- [34] T. A. Arias, *Reviews of Modern Physics* 71 (1999) 267.
- [35] K. Cho, T. Arias, J. Joannopoulos, P. K. Lam, *Physical Review Letters* 71 (1993) 1808.
- [36] L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, et al., *The Journal of chemical physics* 129 (2008) 014109.
- [37] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, M. C. Payne, *The Journal of Chemical Physics* 122 (2005) 084119.
- [38] D. R. Bowler, R. Choudhury, M. J. Gillan, T. Miyazaki, *physica status solidi (b)* 243 (2006) 989–1000.
- [39] A. Masud, R. Kannan, *Computer Methods in Applied Mechanics and Engineering* 241 (2012) 112–127.
- [40] P. Suryanarayana, K. Bhattacharya, M. Ortiz, *Journal of Computational Physics* 230 (2011) 5226 – 5238.
- [41] S. Ghosh, P. Suryanarayana, *Computer Physics Communications* 212 (2017) 189-204.
- [42] T. Ono, M. Heide, N. Atodiresei, P. Baumeister, S. Tsukamoto, S. Blügel, *Physical Review B* 82 (2010) 205115.
- [43] N. S. Bobbitt, G. Schofield, C. Lena, J. R. Chelikowsky, *Phys. Chem. Chem. Phys.* 17 (2015) 31542–31549.

- [44] X. Andrade, D. Strubbe, U. De Giovannini, A. H. Larsen, M. J. Oliveira, J. Alberdi-Rodriguez, A. Varas, I. Theophilou, N. Helbig, M. J. Verstraete, et al., *Physical Chemistry Chemical Physics* 17 (2015) 31371–31396.
- [45] P. Li, X. Liu, M. Chen, P. Lin, X. Ren, L. Lin, C. Yang, L. He, *Computational Materials Science* 112, Part B (2016) 503 – 517. *Computational Materials Science in China*.
- [46] E. Artacho, *Periodic Linear Combination of Atomic Orbitals and Order-N Methods*, Wiley Online Library, pp. 77–92.
- [47] Y. Zhou, Y. Saad, M. L. Tiago, J. R. Chelikowsky, *Physical Review E* 74 (2006) 066704.
- [48] L. Kleinman, D. Bylander, *Physical Review Letters* 48 (1982) 1425.
- [49] J. E. Pask, P. A. Sterne, *Phys. Rev. B* 71 (2005) 113101.
- [50] S. Ghosh, P. Suryanarayana, *Journal of Computational Physics* 307 (2016) 634 – 652.
- [51] F. Bloch, *Zeitschrift für physik* 52 (1929) 555–600.
- [52] J. C. Slater, *The self-consistent field for molecules and solids*, volume 4, McGraw-Hill New York, 1974.
- [53] H.-r. Fang, Y. Saad, *Numerical Linear Algebra with Applications* 16 (2009) 197–221.
- [54] L. Lin, C. Yang, *SIAM Journal on Scientific Computing* 35 (2013) S277–S298.
- [55] P. P. Pratapa, P. Suryanarayana, *Chemical Physics Letters* 635 (2015) 69–74.
- [56] A. S. Banerjee, P. Suryanarayana, J. E. Pask, *Chemical Physics Letters* 647 (2016) 31 – 35.
- [57] P. Suryanarayana, *Chemical Physics Letters* 584 (2013) 182–187.
- [58] P. Suryanarayana, *Chemical Physics Letters* 555 (2013) 291 – 295.
- [59] J. Harris, *Physical Review B* 31 (1985) 1770.
- [60] W. M. C. Foulkes, R. Haydock, *Physical review B* 39 (1989) 12520.
- [61] P. Suryanarayana, D. Phanish, *Journal of Computational Physics* 275 (2014) 524 – 538.
- [62] K. Hirose, T. Ono, Y. Fujimoto, S. Tsukamoto, *First-Principles Calculations in Real-Space Formalism*, Imperial College Press, 2005.
- [63] P. P. Pratapa, P. Suryanarayana, J. E. Pask, *Computer Physics Communications* 200 (2016) 96–107.
- [64] T. Ono, K. Hirose, *Phys. Rev. Lett.* 82 (1999) 5016–5019.
- [65] S. Balay, J. Brown, , K. Buschelman, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, H. Zhang, *PETSc Users Manual*, Technical Report ANL-95/11 - Revision 3.4, Argonne National Laboratory, 2013.
- [66] S. Balay, W. D. Gropp, L. C. McInnes, B. F. Smith, in: E. Arge, A. M. Bruaset, H. P. Langtangen (Eds.), *Modern Software Tools in Scientific Computing*, Birkhäuser Press, 1997, pp. 163–202.

- [67] D. A. Mazziotti, Chemical physics letters 299 (1999) 473–480.
- [68] J. H. Ahlberg, E. N. Nilson, J. L. Walsh, Mathematics in Science and Engineering, New York: Academic Press, 1967 1 (1967).
- [69] Y. Zhou, Y. Saad, M. L. Tiago, J. R. Chelikowsky, Journal of Computational Physics 219 (2006) 172–184.
- [70] A. V. Knyazev, SIAM journal on scientific computing 23 (2001) 517–541.
- [71] Y. Zhou, J. R. Chelikowsky, Y. Saad, Journal of Computational Physics 274 (2014) 770–782.
- [72] J. P. Perdew, Y. Wang, Physical Review B 45 (1992) 13244.
- [73] D. M. Ceperley, B. J. Alder, Phys. Rev. Lett. 45 (1980) 566–569.
- [74] N. Troullier, J. L. Martins, Physical Review B 43 (1991) 1993–2006.
- [75] H. J. Monkhorst, J. D. Pack, Physical Review B 13 (1976) 5188.
- [76] G. H. Golub, C. F. Van Loan, Matrix computations, volume 3, JHU Press, 2012.
- [77] C. C. Paige, M. A. Saunders, SIAM journal on numerical analysis 12 (1975) 617–629.
- [78] C. Lanczos, An iteration method for the solution of the eigenvalue problem of linear differential and integral operators, United States Governm. Press Office, 1950.
- [79] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, D. Sorensen, LAPACK Users’ Guide, Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.
- [80] D. S. Watkins, Fundamentals of matrix computations, volume 64, John Wiley & Sons, 2004.
- [81] W. H. Press, Numerical recipes 3rd edition: The art of scientific computing, Cambridge university press, 2007.
- [82] D. G. Anderson, Journal of the ACM (JACM) 12 (1965) 547–560.
- [83] J. R. Shewchuk, An introduction to the conjugate gradient method without the agonizing pain, 1994.
- [84] D. C. Rapaport, The Art of Molecular Dynamics Simulation, Cambridge University Press, New York, NY, USA, 2nd edition, 2004.
- [85] X. Gonze, B. Amadon, P.-M. Anglade, J.-M. Beuken, F. Bottin, P. Boulanger, F. Bruneval, D. Caliste, R. Caracas, M. Cote, et al., Computer Physics Communications 180 (2009) 2582–2615.
- [86] X. Gonze, Zeitschrift für Kristallographie 220 (2005) 558–562.
- [87] V. Brázdová, D. R. Bowler, Atomistic computer simulations: a practical guide, John Wiley & Sons, 2013.
- [88] M. Gillan, Journal of Physics: Condensed Matter 1 (1989) 689.

- [89] D. Marx, J. Hutter, Modern methods and algorithms of quantum chemistry 1 (2000) 301–449.
- [90] J. E. Pask, N. Sukumar, S. E. Mousavi, International Journal for Multiscale Computational Engineering 10 (2012) 83–99.

Appendix

A. Electrostatic correction for overlapping pseudocharge densities

In ab-initio calculations, even when the pseudopotential approximation is employed, the repulsive energy is still calculated by treating the nuclei as point charges. Since the electrostatic formulation in this work does not make this distinction, there is disagreement with convention when the pseudocharge densities overlap. The correction which reestablishes agreement can be written as [61]

$$E_c(\mathbf{R}) = \frac{1}{2} \int_{\Omega} \left(\tilde{b}(\mathbf{x}, \mathbf{R}) + b(\mathbf{x}, \mathbf{R}) \right) V_c(\mathbf{x}, \mathbf{R}) d\mathbf{x} + \frac{1}{2} \sum_I \int_{\Omega} b_I(\mathbf{x}, \mathbf{R}_I) V_I(\mathbf{x}, \mathbf{R}_I) d\mathbf{x} - \frac{1}{2} \sum_I \int_{\Omega} \tilde{b}_I(\mathbf{x}, \mathbf{R}_I) \tilde{V}_I(\mathbf{x}, \mathbf{R}_I) d\mathbf{x}, \quad (46)$$

where the summation index I runs over all atoms in $\mathbb{R}^3$, and

$$V_c(\mathbf{x}, \mathbf{R}) = \sum_I \left(\tilde{V}_I(\mathbf{x}, \mathbf{R}_I) - V_I(\mathbf{x}, \mathbf{R}_I) \right). \quad (47)$$

In addition, $\tilde{b}$ denotes the reference pseudocharge density, and $\tilde{b}_I$ represents the spherically symmetric and compactly supported reference charge density of the I^{th} nucleus that generates the potential $\tilde{V}_I$, i.e.,

$$\tilde{b}(\mathbf{x}, \mathbf{R}) = \sum_I \tilde{b}_I(\mathbf{x}, \mathbf{R}_I), \quad \tilde{b}_I(\mathbf{x}, \mathbf{R}_I) = -\frac{1}{4\pi} \nabla^2 \tilde{V}_I(\mathbf{x}, \mathbf{R}_I), \quad (48)$$

$$\int_{\mathbb{R}^3} \tilde{b}(\mathbf{x}, \mathbf{R}) d\mathbf{x} = -N_e, \quad \int_{\mathbb{R}^3} \tilde{b}_I(\mathbf{x}, \mathbf{R}_I) d\mathbf{x} = Z_I. \quad (49)$$

The discrete form of the repulsive energy correction in Eqn. 46 takes the form

$$E_c^h = \frac{1}{2} h_1 h_2 h_3 \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} \left((\tilde{b}^{(i,j,k)} + b^{(i,j,k)}) V_c^{(i,j,k)} + \sum_I b_I^{(i,j,k)} V_I^{(i,j,k)} - \sum_I \tilde{b}_I^{(i,j,k)} \tilde{V}_I^{(i,j,k)} \right), \quad (50)$$

where the integrals have been approximated using the trapezoidal rule in Eqn. 24.

The correction to the atomic forces can be written as [61]

$$\begin{aligned} \mathbf{f}_{J,c}(\mathbf{R}) = & \frac{1}{2} \sum_{J'} \int_{\Omega} \left[\nabla \tilde{b}_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \left(V_c(\mathbf{x}, \mathbf{R}) - \tilde{V}_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \right) + \nabla b_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \left(V_c(\mathbf{x}, \mathbf{R}) + V_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \right) \right. \\ & + \left(\nabla \tilde{V}_{J'}(\mathbf{x}, \mathbf{R}_{J'}) - \nabla V_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \right) \left(\tilde{b}(\mathbf{x}, \mathbf{R}) + b(\mathbf{x}, \mathbf{R}) \right) + b_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \nabla V_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \\ & \left. - \tilde{b}_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \nabla \tilde{V}_{J'}(\mathbf{x}, \mathbf{R}_{J'}) \right] d\mathbf{x}, \end{aligned} \quad (51)$$

whose discrete form:

$$\begin{aligned}
\mathbf{f}_{J,c}^h = & \frac{1}{2}h_1h_2h_3 \sum_{J'} \sum_{i=1}^{n_1} \sum_{j=1}^{n_2} \sum_{k=1}^{n_3} \left(\nabla_h \tilde{b}_{J'} |^{(i,j,k)} \left(V_c^{(i,j,k)} - \tilde{V}_{J'}^{(i,j,k)} \right) + \nabla_h b_{J'} |^{(i,j,k)} \left(V_c^{(i,j,k)} + V_{J'}^{(i,j,k)} \right) \right. \\
& + \left. \nabla_h (\tilde{V}_{J'}^{(i,j,k)} - V_{J'}^{(i,j,k)}) |^{(i,j,k)} \left(\tilde{b}^{(i,j,k)} + b^{(i,j,k)} \right) + b_{J'}^{(i,j,k)} \nabla_h V_{J'} |^{(i,j,k)} - \tilde{b}_{J'}^{(i,j,k)} \nabla_h \tilde{V}_{J'} |^{(i,j,k)} \right).
\end{aligned} \tag{52}$$

As before, the summation J' runs over the J^{th} atom and its periodic images. For the reference potential, we choose the potential that has been previously employed for generating neutralizing densities in all-electron calculations [90].

B. Pseudopotential parameters

The cutoff radii r_J^c employed in this work for the different angular momentum components of the Troullier-Martins pseudopotential are listed in Table 2. The $l = 0$ component is chosen to be local in all the calculations.

Atom type	Radial cutoff (Bohr)		
	$l = 0$	$l = 1$	$l = 2$
H	1.25	—	—
Li	2.40	2.40	—
Al	2.60	2.60	—
Si	1.80	1.80	1.80
Au	2.60	2.60	2.60

Table 2: Cutoff radii for non-local projectors within the Troullier-Martins pseudopotential.

C. Discrete pseudocharge density properties

The continuous pseudocharge density of the J^{th} atom has compact support in a sphere of radius r_J^c —cutoff radius for the local pseudopotential—centered at that atom. However, the corresponding discrete pseudocharge density has infinite extent due to the use of the finite-difference approximated Laplacian. In Fig. 12, for a mesh-size $h = 0.5$ Bohr, we plot the normalized error in the net enclosed charge as a function of the pseudocharge radius r_J^b . It is clear that there is exponential decay, which allows for truncation at some finite radius without significant loss of accuracy.

Figure 12: Normalized error in the net enclosed charge as a function of pseudocharge radius for $h = 0.5$ Bohr.

In Fig. 13, we plot the truncation radius r_J^b required to achieve the accuracy of $\varepsilon_b = 10^{-8}$ (Eqn. 28) as a function of mesh-size h . We observe that r_J^b becomes smaller as the discretization is refined, with $r_J^b \rightarrow r_J^c$ as $h \rightarrow 0$. The slight non-monotonicity observed at some places in the curves arises due to the fact that r_J^b is chosen to be a multiple of h within SPARC.

Figure 13: Variation of pseudocharge radius as a function of mesh size.

D. Comparison of SPARC with other finite-difference DFT codes

We now study the accuracy and efficiency of SPARC relative to PARSEC [22] and OCTOPUS [23]—two well-established DFT codes that employ the finite-difference discretization. In all three codes, we choose the same pseudopotential, details of which are provided in Appendix B. In addition, we choose the same number of states and values for the parameters in each code. For the eigensolver, we employ the default in each code, i.e., CheFSI in SPARC and PARSEC, and Conjugate Gradients in OCTOPUS.¹² First, we plot in Fig. 14 the convergence of the energy and atomic forces as a function of mesh-size for the system

¹²Other choices of eigensolver in OCTOPUS (e.g. RMM-DIIS) do not significantly change the nature of the results presented here.

consisting of $3 \times 3 \times 3$ FCC unit cells of aluminum with a vacancy. The error has been defined with respect to ABINIT,¹³ wherein we employ a plane-wave cutoff of 32 Ha, which results in energy and forces that are converged to within 1×10^{-6} Ha/atom and 1×10^{-6} Ha/Bohr, respectively. We observe that the prefactors associated with the convergence of the energy in OCTOPUS and the forces in PARSEC are significantly larger than those for the other two codes. Notably, SPARC demonstrates the best convergence in terms of energy and forces amongst the three finite-difference packages.

Figure 14: Convergence of the energy and atomic forces with respect to mesh size to reference planewave result for $3 \times 3 \times 3$ FCC unit cells of aluminum with a vacancy. The error in energy for PARSEC and OCTOPUS has been defined with respect to the results obtained for a mesh-size of $h = 0.201$ Bohr (See footnote 13)

Next, we study the weak and strong scaling of SPARC relative to PARSEC and OCTOPUS. Specifically, we perform the weak and strong scaling tests described in Section 4.7, with a mesh of $h = 0.778$ Bohr employed in all three codes. We present the results obtained in Fig. 15, wherein the time taken for the first SCF iteration has been excluded [41]. In the strong scaling test, the minimum wall time achieved by SPARC is smaller by factors of 12.2, 93, and 5.8 relative to PARSEC, OCTOPUS, and ABINIT, respectively.¹⁴ In the weak scaling test, the increase in CPU time with number of electrons for SPARC, PARSEC, OCTOPUS, and ABINIT is $\mathcal{O}(N_e^{2.64})$, $\mathcal{O}(N_e^{3.90})$, $\mathcal{O}(N_e^{3.98})$ and $\mathcal{O}(N_e^{2.99})$ respectively. It is clear that SPARC is able to outperform PARSEC because of significantly higher efficiency in strong scaling, and is able to outperform OCTOPUS because of significantly smaller prefactor and superior strong scaling.¹⁵

¹³For crystals, the energy computed by PARSEC and OCTOPUS is found to converge to an answer that is significantly different from ABINIT. Therefore, the error in energy for PARSEC and OCTOPUS has been defined with respect to the results obtained for a mesh-size of $h = 0.201$ Bohr.

¹⁴It is worth noting that the local electrostatic reformulation in SPARC significantly outperforms conventional Fourier based approaches for moderate to large systems, especially in the context of parallel computing. For example, in the strong scaling test, the time taken for the electrostatics in SPARC is a factor of 2.7 smaller compared to OCTOPUS on 4 cores. This speedup increases to 62.6 for the calculation on 384 cores.

¹⁵It is worth noting that unlike isolated systems where all three codes take the same number of SCF iterations [41], for extended systems OCTOPUS takes significantly more SCF iterations compared to SPARC/PARSEC. For example, for $4 \times 4 \times 4$ unit cells of aluminum with a vacancy, SPARC/PARSEC takes 19 iterations, whereas OCTOPUS takes 58 SCF iterations. Similarly, for $5 \times 5 \times 5$ unit cells of aluminum with a vacancy, SPARC/PARSEC takes 26 iterations, whereas OCTOPUS takes 79 iterations.

Figure 15: Strong and weak scaling for SPARC, PARSEC, OCTOPUS, and ABINIT. The system utilized for strong scaling is $6 \times 6 \times 6$ FCC unit cells of aluminum with a vacancy. The systems employed for weak scaling are $3 \times 3 \times 3$, $4 \times 4 \times 4$, $5 \times 5 \times 5$, $6 \times 6 \times 6$, and $7 \times 7 \times 7$ unit cells of aluminum, each with a vacancy.

Finally, we compare the minimum wall time—excluding the time for the first SCF iteration—that can be achieved by SPARC, PARSEC, OCTOPUS, and ABINIT for the aforementioned systems. From the results presented in Table 3, we find that SPARC demonstrates speedup by up to factors of 12.2, 72.3, and 5.55 compared to PARSEC, OCTOPUS, and ABINIT, respectively. Overall, these results demonstrate that SPARC is an efficient DFT formulation and implementation that is not only highly competitive with well-established plane-wave codes, but also significantly outperforms well-established finite-difference codes.

System	SPARC $h = 0.778$ Bohr	PARSEC $h = 0.778$ Bohr	OCTOPUS $h = 0.778$ Bohr	ABINIT $E_{cut} = 9$ Ha
$3 \times 3 \times 3$	0.36 (64)	1.58 (64)	3.64 (8)	1.59 (180)
$4 \times 4 \times 4$	1.47 (96)	6.88 (128)	25.33 (64)	5.27 (320)
$5 \times 5 \times 5$	6.39 (256)	43.8 (256)	465.24 (256)	16.51 (396)
$6 \times 6 \times 6$	19.06 (640)	233 (512)	1212.75 (384)	108.77 (480)
$7 \times 7 \times 7$	71.80 (704)	705 (512)	—	388.67 (795)

Table 3: Minimum wall time in minutes for $n \times n \times n$ ($n \in \{3, 4, 5, 6, 7\}$) FCC unit cells of aluminum with a vacancy. The number in brackets represents the number of cores on which the minimum wall time is achieved.