

SPARC: Simulation Package for Ab-initio Real-space Calculations

Qimen Xu^a, Abhiraj Sharma^a, Benjamin Comer^a, Hua Huang^b, Edmond Chow^b, Andrew J. Medford^a, John E. Pask^c, Phanish Suryanarayana^{a,*}

^a*College of Engineering, Georgia Institute of Technology, Atlanta, GA 30332, USA*

^b*College of Computing, Georgia Institute of Technology, Atlanta, GA 30332, USA*

^c*Physics Division, Lawrence Livermore National Laboratory, Livermore, CA 94550, USA*

Abstract

We present SPARC: Simulation Package for Ab-initio Real-space Calculations. SPARC can perform Kohn-Sham density functional theory calculations for isolated systems such as molecules as well as extended systems such as crystals and surfaces, in both static and dynamic settings. It is straightforward to install/use and highly competitive with state-of-the-art planewave codes, demonstrating comparable performance on a small number of processors and increasing advantages as the number of processors grows. Notably, SPARC brings solution times down to a few seconds for systems with $\mathcal{O}(100 - 500)$ atoms on large-scale parallel computers, outperforming planewave counterparts by an order of magnitude and more.

Keywords: Kohn-Sham, Density Functional Theory, Electronic structure, Real-space, Finite-differences

Code Metadata

C1	Current code version	v1.0.0
C2	Permanent link to code/repository used for this code version	https://github.com/SPARC-X/SPARC
C3	Code Ocean compute capsule	
C4	Legal Code License	GNU General Public License v3.0
C5	Code versioning system used	git
C6	Software code languages, tools, and services used	C, MPI, BLAS, LAPACK, ScaLAPACK (optional), MKL (optional)
C7	Compilation requirements, operating environments & dependencies	OS: Unix, Linux, or MacOS
C8	If available Link to developer documentation/manual	https://github.com/SPARC-X/SPARC/tree/master/doc
C9	Support email for questions	phanish.s@gmail.com

Table 1: Code metadata

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

1. Motivation and significance

Over the course of the past few decades, quantum mechanical calculations based on Kohn-Sham density functional theory (DFT) [1, 2] have become a cornerstone of materials research by virtue of the predictive power and fundamental insights they provide. The widespread use of the methodology can be attributed to its generality, simplicity, and high accuracy-to-cost ratio relative to other such *ab initio* approaches [3, 4]. However, while less expensive than wavefunction based methods, the solution of the Kohn-Sham equations remains a formidable task. In particular, the computational cost scales cubically with the number of atoms, severely limiting the range of physical systems accessible to such first principles investigation. These limitations become even more acute in quantum molecular dynamics (QMD) simulations, wherein the equations for the electronic ground state may be solved tens or hundreds of thousands of times to reach time scales relevant to phenomena of interest [3].

The planewave pseudopotential method [5] has been among the most widely used techniques for the solution of the Kohn-Sham problem [6–13]. The underlying Fourier basis is complete, orthonormal, diagonalizes the Laplacian, and provides spectral convergence for smooth problems. As a result, the planewave method is accurate, simple to use since it relies on a single convergence parameter, and highly efficient on moderate computational resources with the use of well optimized Fast Fourier Transforms (FFTs) and efficient preconditioning schemes. However, the Fourier basis restricts the method to periodic boundary conditions, whereby finite systems such as clusters and molecules, as well as semi-infinite systems such as surfaces and nanowires, require the introduction of artificial periodicity with large vacuum regions. This limitation also necessitates the introduction of an unphysical neutralizing background density when treating charged systems in order to avoid Coulomb divergences. Moreover, the global nature of the Fourier basis hampers scalability on parallel computing platforms and complicates the development of linear-scaling methods [14–16], limiting the system sizes and time scales accessible.

The limitations of the planewave method have motivated the development of a number of alternative solution strategies employing systematically improvable, localized representations [17–37]. Among these, perhaps the most mature and widely used to date are the finite-difference methods [38, 39], wherein computational locality is maximized by discretizing all quantities of interest on a uniform real-space grid. As a result, convergence is controlled by a single parameter and both periodic and Dirichlet boundary conditions are naturally accommodated, thus enabling the efficient and accurate treatment of finite, semi-infinite, bulk, and charged systems alike. Moreover, real-space methods are amenable to the development of linear scaling methods, and large-scale parallel computational resources can be efficiently leveraged by virtue of the method’s simplicity, locality, and freedom from communication-intensive transforms such as FFTs [22, 33, 40–42]. With these and other advances, real-space methods have been applied to systems containing thousands of atoms, and have demonstrated substantially reduced solution times compared to established planewave codes in applications to both finite [43] and extended [34] systems.

However, despite the significant advantages afforded by real-space methods, the planewave

method has remained the method of choice in practice for the better part of the past two decades. This is largely due to the ease of use, extensive feature sets, established accuracy/robustness, and straightforward installation of the associated codes, having been in development and production for a longer period of time. Perhaps most importantly, however, planewave codes have typically yielded shorter times to solution using moderate computational resources, as most widely available to researchers in practice [34, 43]. Moreover, even with access to larger-scale machines, real-space codes have not always yielded shorter times to solution, further hindering wider adoption in practice.

In this work, we present an open-source software package for the accurate, efficient, and scalable solution of the Kohn-Sham equations, referred to as SPARC. The package is straightforward to install/use and highly competitive with state-of-the-art planewave codes, demonstrating comparable performance on a small number of processors and order-of-magnitude advantages as the number of processors increases.

2. Software description

The central focus of SPARC is the accurate and efficient solution of the finite-temperature Kohn-Sham equations for the electronic ground state [1, 2, 44]:

$$\left(\mathcal{H}^\sigma \equiv -\frac{1}{2}\nabla^2 + V_{\text{eff}}^\sigma[\rho^\alpha, \rho^\beta; \mathbf{R}] \right) \psi_n^\sigma = \lambda_n^\sigma \psi_n^\sigma, \quad n = 1, 2, \dots, N_s^\sigma, \quad \sigma \in \{\alpha, \beta\}, \quad (1)$$

where the superscript σ denotes the spin, i.e., spin-up or spin-down, \mathcal{H}^σ is the Hamiltonian, ψ_n^σ are the orthonormal orbitals with energies λ_n^σ , V_{eff}^σ is the effective potential, N_s^σ is the number of states, and \mathbf{R} denotes the set of atomic positions. In addition, ρ^σ represents the spin-resolved electron density:

$$\rho^\sigma(\mathbf{x}) = \sum_{n=1}^{N_s^\sigma} g_n^\sigma |\psi_n^\sigma(\mathbf{x})|^2, \quad \sigma \in \{\alpha, \beta\}, \quad \mathbf{x} \in \mathbb{R}^3, \quad (2)$$

where g_n^σ are the orbital occupations, typically given by the Fermi-Dirac distribution. In implementations of the above equations, once a suitable fundamental domain/unit cell has been identified, zero Dirichlet or Bloch-periodic boundary conditions are prescribed on the orbitals along the directions in which the system is finite or extended, respectively.

2.1. Software Architecture

SPARC employs the pseudopotential approximation [5] to facilitate the efficient solution of the Kohn-Sham equations for the whole of the periodic table of elements. In addition, it employs a local real-space formulation of the electrostatics [45, 46], wherein the electrostatic potential — component of V_{eff}^σ that is the sum of ionic and Hartree contributions — is given by the solution of a Poisson problem, with Dirichlet or periodic boundary conditions prescribed along directions in which the system is finite or extended, respectively. In this framework, SPARC performs a uniform real-space discretization of the equations, using a high-order

centered finite-difference approximation for differential operators and the trapezoidal rule for integral operators. The actual code is written in the C language and achieves parallelism through the message passing interface (MPI) [47]. An overview of the SPARC framework for performing Kohn-Sham DFT calculations is shown in Fig. 1.

SPARC can perform single-point calculations, structural relaxations (atom and/or cell), and QMD simulations. For single-point calculations, the electronic ground state is determined for fixed ionic positions and cell dimensions, whereas for structural relaxations, positions and/or cell dimensions are varied to minimize the Kohn-Sham energy using the Hellmann-Feynman atomic forces [34, 43] and/or stress tensor [48]. For QMD, the ionic positions, velocities, and accelerations are evolved by integrating the equations of motion, with or without a thermostat, using the atomic forces. In all cases, the calculations can be either spin-polarized or unpolarized, with various choices of local and semilocal exchange-correlation functionals.

SPARC requires two input files for every calculation: (i) a `.inpt` file containing the options and parameters to be used in the calculation, including the choice of exchange-correlation functional, flag for spin-polarization, type of static/dynamic calculation, ionic temperature in the case of QMD, cell dimensions, boundary condition in each direction, and finite-difference grid specification; and (ii) a `.ion` file containing information on the atomic configuration, including atom types, positions, and paths to corresponding pseudopotential files. Note that, in order to enable detailed control of the simulation, a large number of parameters and options can be specified, as described in the accompanying user guide. However, by virtue of carefully chosen defaults, relatively few parameters typically need be specified in practice. Note also that, since all files are simple, human-readable text, series of simulations are readily scripted. A Python package containing helper functions for generating input files and submitting simulations is also available.

The Kohn-Sham problem for the electronic ground state needs to be solved for every configuration of atoms encountered during the DFT simulation. In SPARC, this is achieved using the self-consistent field (SCF) method [5], which represents a fixed-point iteration with respect to either the electron density or potential. For the first SCF iteration in the simulation, a superposition of isolated atom electron densities is used as the initial guess, whereas for every subsequent atomic configuration encountered, extrapolation based on solutions to previous configurations is employed [49]. The convergence of the SCF iteration is accelerated using the restarted variant [50] of the Periodic Pulay mixing scheme [51] with a real-space preconditioner [52]. In the case of spin-polarized calculations, mixing is performed simultaneously on both components, i.e., on a vector of twice the original length containing both spin-up and spin-down density/potential components.

In each SCF iteration, SPARC performs a partial diagonalization (i.e., eigenvalues and eigenvectors calculated approximately) of the linear eigenproblem using the CheFSI method [53, 54], with multiple Chebyshev filtering steps performed in the first iteration of the simulation [55]. The Hamiltonian-matrix/vector products are performed in matrix-free fashion, using the finite-difference stencil for the Laplacian and the outer product nature of the non-local pseudopotential operator. While doing so, zero-Dirichlet or Bloch-periodic boundary

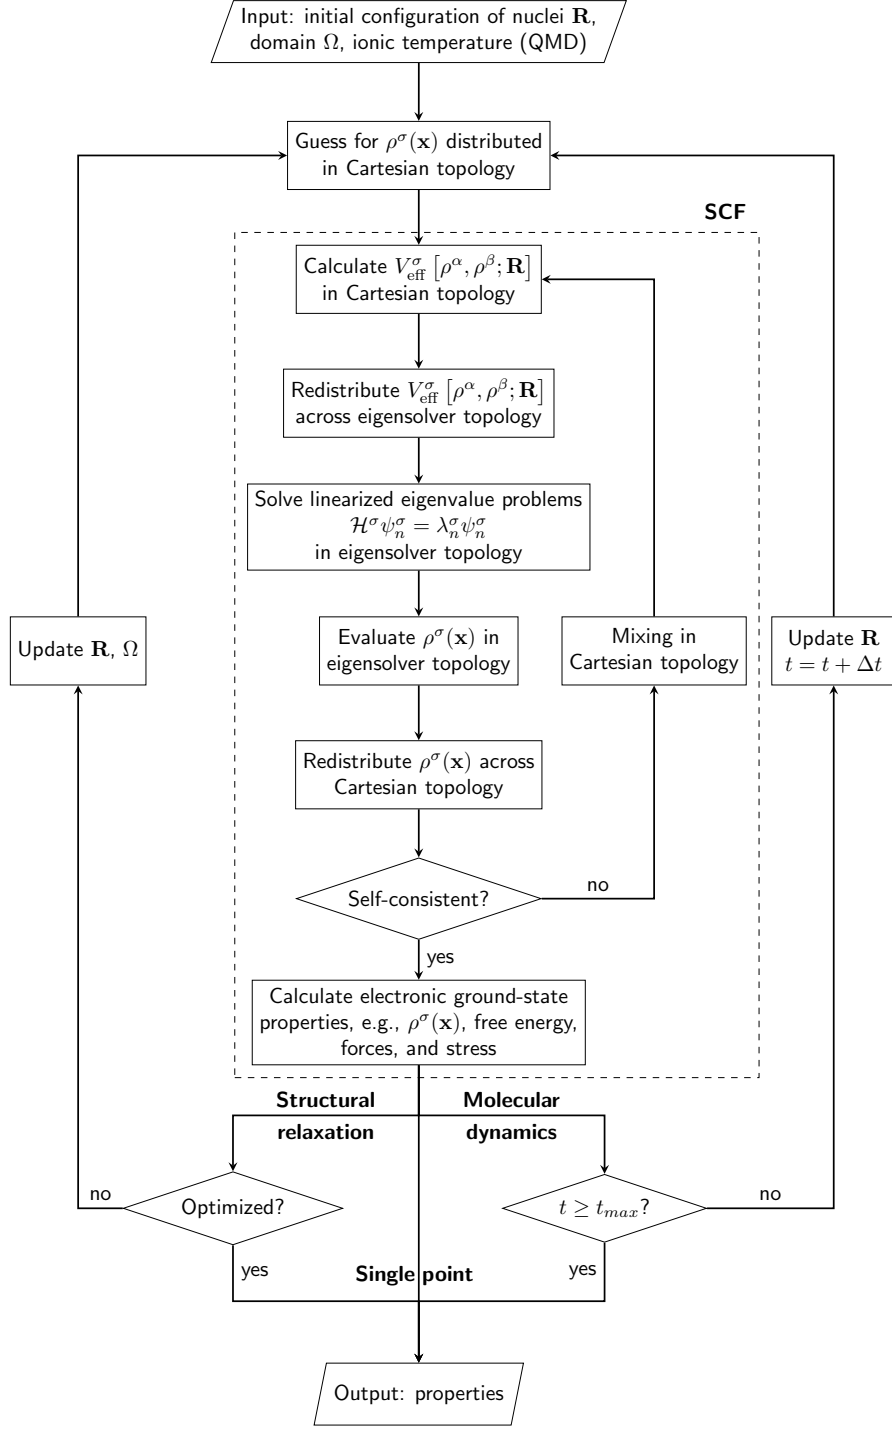


Figure 1: Overview of the SPARC framework for performing Kohn-Sham DFT calculations. The Cartesian topology is formed by embedding a three-dimensional processor grid into the `MPI_COMM_WORLD` communicator. The eigensolver topology is a collection of smaller Cartesian topologies, created by first splitting the `MPI_COMM_WORLD` communicator into multiple spin groups, then splitting each spin group into multiple Brillouin zone integration groups, then splitting each Brillouin zone integration group into multiple band groups, and finally, embedding each band group with a Cartesian topology.

conditions are prescribed on the orbitals along directions in which the system is finite or extended, respectively. In calculating the effective potential, the Poisson problem for the electrostatic potential is solved using the AAR method [56, 57], with Laplacian-vector products again performed in matrix-free fashion using the finite-difference stencil. While doing so, Dirichlet or periodic boundary conditions are prescribed on the electrostatic potential along directions in which the system is finite or extended, respectively. In particular, Dirichlet values are determined using a multipole expansion for isolated systems and a dipole correction for surfaces and nanowires [58, 59].

In SPARC, information pertaining to the overall DFT simulation is written to the `.out` file, including progression of the SCF iteration, electronic ground state energy, maximum atomic force, maximum stress, and various timings. Based on the type of calculation, a `.static`, `.geopt`, or `.aimd` file may also be written. The `.static` file contains information about the single-point calculation, including atom positions, electronic ground state energy, forces, and stress tensor. The `.geopt` file contains information about the structural relaxation, including (i) atom positions, electronic ground state energy, and forces (atomic relaxation), and (ii) cell information and stress tensor (cell relaxation). The `.aimd` file contains information about the QMD simulation, including atom positions, forces, and velocities. To seamlessly continue from a previously stopped simulation, a `.restart` file is written for structural relaxation and QMD calculations. SPARC provides other outputs if specified as well, e.g., a `.eigen` file containing eigenvalues and occupations and `.dens` file containing the charge density.

2.2. Software Functionalities

The current version of SPARC is capable of performing spin-polarized and unpolarized *ab initio* calculations based on Kohn-Sham DFT for isolated systems such as molecules as well as extended systems such as crystals, surfaces, and nanowires, in both static and dynamic settings. Specifically, it can perform single-point calculations for a given atomic configuration, structural relaxations with respect to atom positions and/or cell dimensions [60–63], and NVE/NVT/NVK QMD simulations [64–66]. Available exchange-correlation functionals include various forms of LDA [67, 68] and GGA [69–71]. Types of pseudopotentials employed are ONCV [72] and Troullier-Martin [73], both in `psp8` format [8]. Over the course of simulations, in addition to electronic density and free energy, SPARC can calculate atomic forces, pressure, and the stress tensor for extended systems. The outputs from such DFT calculations can be used to calculate a number of properties, including lattice constant, cohesive energy, polarization, elastic moduli, density of states, electronic band structure, pair distribution function, equations of state, shear viscosity, defect energy, surface energy, absorption energy, equilibrium bond lengths, HOMO-LUMO gap, and dipole moment.

3. Illustrative Examples

We now demonstrate the major functionalities of SPARC through examples representative of physical applications. Specifically, we consider (i) 200-atom NVT QMD simulation for liquid $\text{Al}_{88}\text{Si}_{12}$ alloy at 973 K, with LDA and Γ -point for Brillouin zone integration; (ii)

structural atomic relaxation for a 52-atom system modeling a NH_3 adsorbate on a (110) TiO_2 surface with GGA and 4×4 grid for Brillouin zone integration; (iii) structural cell relaxation for a 102-atom MoS_2 nanotube of diameter 3 nm with GGA and 10 points for Brillouin zone integration; (iv) single-point calculation of a 55-atom icosahedral Co nanoparticle with GGA and spin polarization; and (v) single-point calculations for a 74-system test suite containing isolated systems such as clusters and molecules as well as extended systems such as crystals, surfaces, and nanowires, ranging from 2 to 204 atoms, encompassing 48 different chemical elements and spin-polarized as well as unpolarized calculations.

We employ the Perdew-Zunger parametrization for LDA [67], PBE variant for GGA [69], and ONCV pseudopotentials [72, 74]. We choose mesh-sizes of 0.55, 0.24, 0.20, 0.35, and 0.18 Bohr for the $\text{Al}_{88}\text{Si}_{12}$, TiO_2 , MoS_2 , Co, and test suite examples, respectively. We present the results in Fig. 2 (illustrations using VESTA [75]) and compare them to established planewave codes Quantum Espresso (QE) [9] and ABINIT [8], as well results from the literature [76, 77]. It is clear that there is excellent agreement between SPARC and established planewave codes. In particular, the average difference in energy from highly converged ABINIT calculations obtained with planewave cutoff 100 Ha for the 74-system test suite is $\sim 2.5 \times 10^{-5}$ Ha/atom, substantially smaller than required in typical applications. Indeed, the agreement is further increased as the discretization is refined in SPARC. Overall, these examples demonstrate the capability of SPARC to obtain highly accurate results for a broad range of system compositions, configurations, and dimensionalities.

4. Impact

Kohn-Sham DFT simulations occupy a large fraction of high-performance computing resources around the world every day [78, 79], a consequence of the unique insights and robust predictions they have been shown to provide. The majority of these calculations are performed using established planewave codes [6–10, 13]. Therefore, any new implementation that is able to consistently outperform these state-of-the-art DFT codes, thereby enabling the ab initio investigation of larger length and time scales than previously accessible, with the accuracy required, stands to have significant and immediate impact. This is particularly true of a code like SPARC that is open-source with minimal dependencies so that it can be easily installed on computers large and small around the world.

Accordingly, we compare the accuracy and efficiency of SPARC to Quantum Espresso (QE) [9], an established state-of-the-art planewave DFT code. We employ the same pseudopotentials [72, 74] and exchange-correlation functionals [67, 69] in both codes. Results and computational parameters for the study, containing a wide range of system sizes, are shown in Fig. 3. It is clear that SPARC demonstrates comparable performance to QE on a small number of processors and increasing advantages as the number of processors grows. In particular, SPARC brings solution times down to a few seconds for systems with $\mathcal{O}(100 - 500)$ atoms on large-scale parallel computers, outperforming QE by more than an order of magnitude. Furthermore, SPARC achieves QMD step times of just over 20 seconds for the largest systems containing more than a thousand atoms, achieved on only 312 cores for Al_{1372} . For such systems and larger, the ability of SPARC to efficiently scale to many thousands of

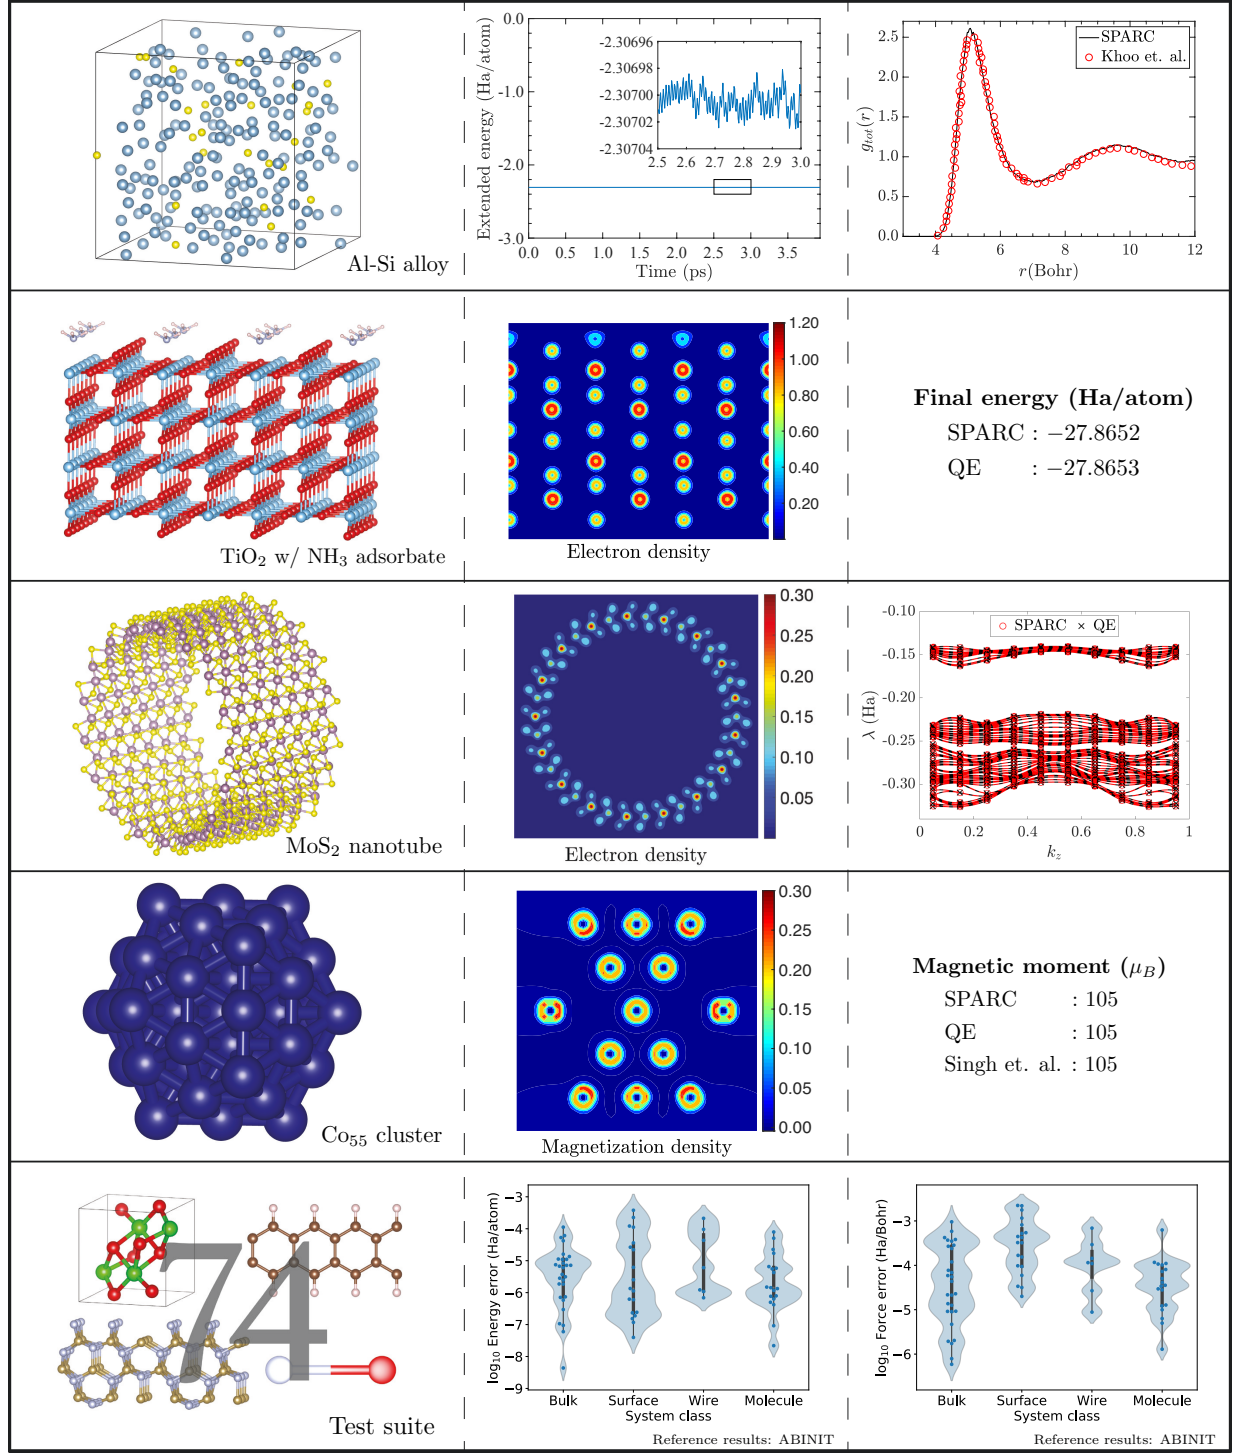


Figure 2: Examples demonstrating the major functionalities of SPARC.

processors and more is currently limited by the subspace diagonalization step performed in each SCF iteration, which due to its cubic scaling and limited parallel scalability takes a

larger fraction of wall time as the system size grows.

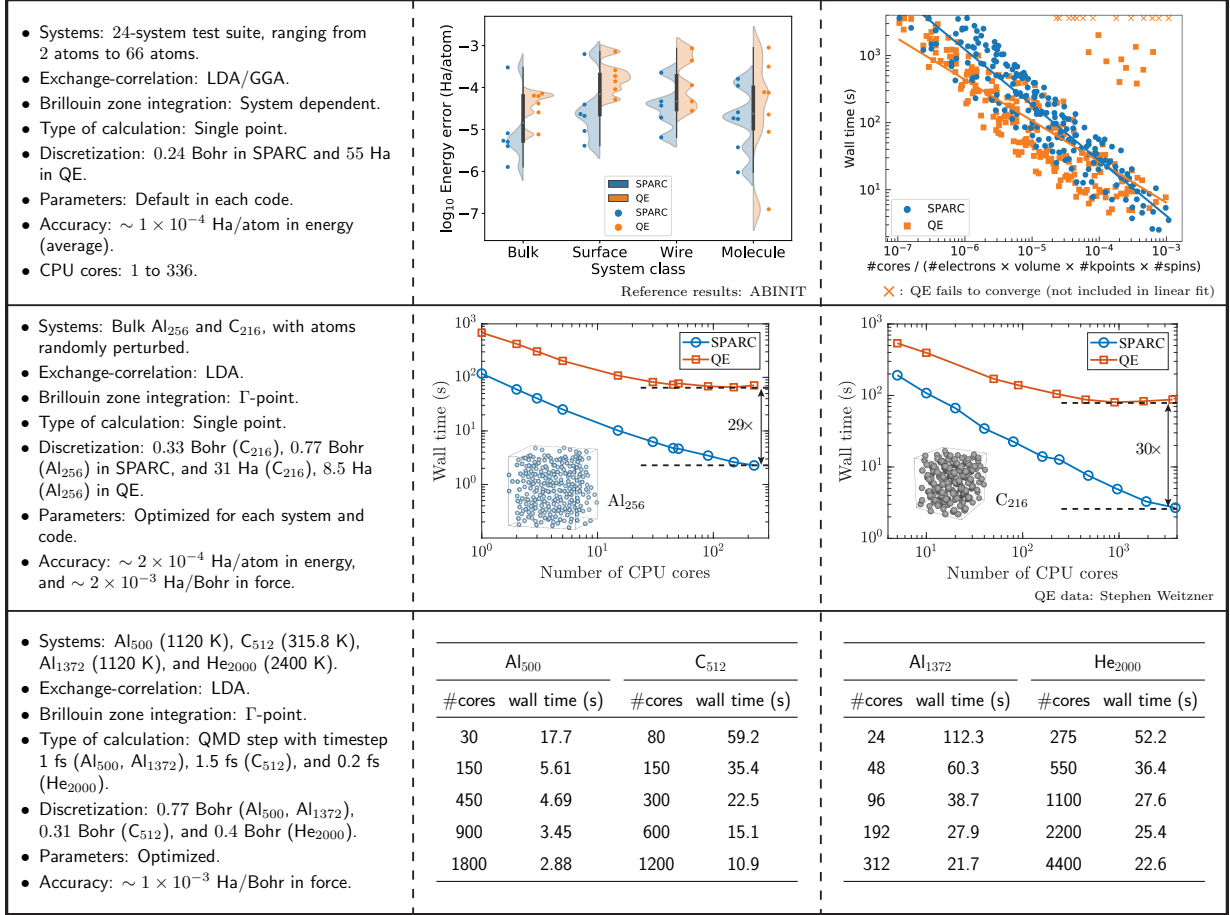


Figure 3: Examples demonstrating the performance of SPARC.

Going forward, we plan to first implement a structure-adapted eigensolver in SPARC to push back the cubic-scaling bottleneck, and then the DDBP method [24] to enable strong scaling of SPARC to still larger numbers of processors, bringing down time to solution still further. The DDBP method will also enable efficient DFT calculations with hybrid functionals and the linear-scaling Spectral Quadrature (SQ) method [42, 80], which will be implemented subsequently. In order to enable the effective use of exascale computing platforms, a parallel engine for SPARC that enables highly efficient distributed memory communication and offloading to GPUs will be completed. Moreover, machine-learning methods will be explored to improve efficiency still further. Along with these developments, we plan to implement cyclic and helical symmetry-adapted DFT formulations that allow for the highly efficient study of associated mechanical deformations as well as systems with such symmetries [81–83]; and a coarse-grained DFT formulation that enables the study of crystal defects at realistic concentrations [84]. Indeed, many of these developments will be accelerated by using the M-SPARC code [85]—same structure, algorithms, input, and output as SPARC—for rapid prototyping.

SPARC and its variants are currently being used by multiple research groups. Moving forward, the user base is expected to grow, given the current open-source distribution, simplicity of installation and use, high accuracy, and ability to reach larger length and time scales than current state-of-the-art planewave codes. The impact thus stands to be both broad and substantial.

5. Conclusions

SPARC has now become a mature code for performing real-space Kohn-Sham DFT calculations, prompting its open-source release with this publication. Currently, it can perform pseudopotential spin-polarized and unpolarized simulations for isolated systems such as molecules and clusters as well as extended systems such as crystals, surfaces, and nanowires, in both static and dynamic settings. SPARC is not only highly accurate, but also highly competitive with established state-of-the-art planewave codes on modest computational resources, with increasing advantages as the number of processors increases. In particular, SPARC efficiently scales to thousands of processors, bringing solution times for moderate-sized systems consisting of $\mathcal{O}(100 - 500)$ atoms to within a few seconds, making it an attractive choice for QMD simulations in particular. Given its superior scalability, and ability to incorporate attractive features such as linear scaling methods and variety of boundary conditions, SPARC has the potential to enable a number of new and exciting applications in science and engineering that were previously beyond reach.

Conflict of Interest

We wish to confirm that there are no known conflicts of interest associated with this publication and there has been no significant financial support for this work that could have influenced its outcome.

Acknowledgements

This work was supported by grant DE-SC0019410 funded by the U.S. Department of Energy, Office of Science. The work was performed in part under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. Support from the Advanced Simulation & Computing / Physics & Engineering Models program at LLNL is gratefully acknowledged. This research was supported in part through research cyberinfrastructure resources and services provided by PACE at GT, including the Hive cluster (NSF-1828187). Time on the Quartz supercomputer was provided by the Computing Grand Challenge program at LLNL. We thank Donald Hamann for use of and assistance with a development version of the ONCVSP code.

References

- [1] W. Kohn, L. J. Sham, Self-consistent equations including exchange and correlation effects, *Physical Review* 140 (4A) (1965) A1133–A1138.

- [2] P. Hohenberg, W. Kohn, Inhomogeneous electron gas, *Physical Review* 136 (3B) (1964) B864.
- [3] K. Burke, Perspective on density functional theory, *The Journal of Chemical Physics* 136 (2012) 150901.
- [4] A. D. Becke, Perspective: Fifty years of density-functional theory in chemical physics, *The Journal of Chemical Physics* 140 (2014) 18A301.
- [5] R. Martin, *Electronic Structure: Basic theory and practical methods*, Cambridge University Press, 2004.
- [6] G. Kresse, J. Furthmüller, Efficient iterative schemes for ab initio total-energy calculations using a plane-wave basis set, *Physical Review B* 54 (16) (1996) 11169–11186.
- [7] S. J. Clark, M. D. Segall, C. J. Pickard, P. J. Hasnip, M. I. Probert, K. Refson, M. C. Payne, First principles methods using CASTEP, *Zeitschrift für Kristallographie-Crystalline Materials* 220 (5/6) (2005) 567–570.
- [8] 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, First-principles computation of material properties: the ABINIT software project, *Computational Materials Science* 25 (2002) 478–492(15).
- [9] 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, QUANTUM ESPRESSO: a modular and open-source software project for quantum simulations of materials, *Journal of Physics: Condensed Matter* 21 (39) (2009) 395502 (19pp).
- [10] D. Marx, J. Hutter, Ab initio molecular dynamics: Theory and implementation, *Modern methods and algorithms of quantum chemistry* 1 (2000) 301–449.
- [11] S. Ismail-Beigi, T. A. Arias, New algebraic formulation of density functional calculation, *Computer Physics Communications* 128 (1-2) (2000) 1 – 45.
- [12] F. Gygi, Architecture of Qbox: A scalable first-principles molecular dynamics code, *IBM Journal of Research and Development* 52 (1.2) (2008) 137–144.
- [13] M. Valiev, E. Bylaska, N. Govind, K. Kowalski, T. Straatsma, H. V. 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, *Computer Physics Communications* 181 (9) (2010) 1477 – 1489.

- [14] S. Goedecker, Linear scaling electronic structure methods, *Reviews of Modern Physics* 71 (4) (1999) 1085–1123.
- [15] D. R. Bowler, T. Miyazaki, $O(N)$ methods in electronic structure calculations, *Reports on Progress in Physics* 75 (3) (2012) 036503.
- [16] J. Aarons, M. Sarwar, D. Thompsett, C.-K. Skylaris, Perspective: Methods for large-scale density functional calculations on metallic systems, *The Journal of Chemical Physics* 145 (22) (2016) 220901.
- [17] A. D. Becke, Basis-set-free density-functional quantum chemistry, *International Journal of Quantum Chemistry* 36 (S23) (1989) 599–609.
- [18] J. R. Chelikowsky, N. Troullier, Y. Saad, Finite-difference-pseudopotential method: electronic structure calculations without a basis, *Physical Review Letters* 72 (8) (1994) 1240.
- [19] L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, R. Schneider, Daubechies wavelets as a basis set for density functional pseudopotential calculations, *The Journal of Chemical Physics* 129 (1) (2008) 014109.
- [20] A. P. Seitsonen, M. J. Puska, R. M. Nieminen, Real-space electronic-structure calculations: Combination of the finite-difference and conjugate-gradient methods, *Physical Review B* 51 (20) (1995) 14057.
- [21] S. R. White, J. W. Wilkins, M. P. Teter, Finite-element method for electronic structure, *Physical Review B* 39 (9) (1989) 5819.
- [22] J.-I. Iwata, D. Takahashi, A. Oshiyama, T. Boku, K. Shiraishi, S. Okada, K. Yabana, A massively-parallel electronic-structure calculations based on real-space density functional theory, *Journal of Computational Physics* 229 (6) (2010) 2339–2363.
- [23] E. Tsuchida, M. Tsukada, Electronic-structure calculations based on the finite-element method, *Physical Review B* 52 (8) (1995) 5573.
- [24] Q. Xu, P. Suryanarayana, J. E. Pask, Discrete discontinuous basis projection method for large-scale electronic structure calculations, *The Journal of Chemical Physics* 149 (9) (2018) 094104.
- [25] P. Suryanarayana, K. Bhattacharya, M. Ortiz, A mesh-free convex approximation scheme for Kohn-Sham density functional theory, *Journal of Computational Physics* 230 (13) (2011) 5226 – 5238.
- [26] P. Suryanarayana, V. Gavini, T. Blesgen, K. Bhattacharya, M. Ortiz, Non-periodic finite-element formulation of Kohn-Sham density functional theory, *Journal of the Mechanics and Physics of Solids* 58 (2) (2010) 256 – 280.

- [27] C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, M. C. Payne, Introducing ONETEP: Linear-scaling density functional simulations on parallel computers, *The Journal of Chemical Physics* 122 (8) (2005) 084119.
- [28] D. R. Bowler, R. Choudhury, M. J. Gillan, T. Miyazaki, Recent progress with large-scale ab initio calculations: the CONQUEST code, *physica status solidi (b)* 243 (5) (2006) 989–1000.
- [29] 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, *Computer Physics Communications* 246 (2020) 106853.
- [30] A. Castro, H. Appel, M. Oliveira, C. A. Rozzi, X. Andrade, F. Lorenzen, M. A. L. Marques, E. K. U. Gross, A. Rubio, octopus: a tool for the application of time-dependent density functional theory, *Physica Status Solidi B-Basic Solid State Physics* 243 (11) (2006) 2465–2488.
- [31] E. L. Briggs, D. J. Sullivan, J. Bernholc, Real-space multigrid-based approach to large-scale electronic structure calculations, *Physical Review B* 54 (1996) 14362–14375.
- [32] J.-L. Fattebert, Finite difference schemes and block Rayleigh quotient iteration for electronic structure calculations on composite grids, *Journal of Computational Physics* 149 (1) (1999) 75 – 94.
- [33] F. Shimojo, R. K. Kalia, A. Nakano, P. Vashishta, Linear-scaling density-functional-theory calculations of electronic structure based on real-space grids: design, analysis, and scalability test of parallel algorithms, *Computer Physics Communications* 140 (3) (2001) 303 – 314.
- [34] S. Ghosh, P. Suryanarayana, SPARC: Accurate and efficient finite-difference formulation and parallel implementation of density functional theory: Extended systems, *Computer Physics Communications* 216 (2017) 109–125.
- [35] T. A. Arias, Multiresolution analysis of electronic structure: semicardinal and wavelet bases, *Reviews of Modern Physics* 71 (1) (1999) 267–311.
- [36] J. E. Pask, P. A. Sterne, Finite element methods in ab initio electronic structure calculations, *Modelling and Simulation in Materials Science and Engineering* 13 (2005) R71–R96.
- [37] L. Lin, J. Lu, L. Ying, E. Weinan, Adaptive local basis set for Kohn–Sham density functional theory in a discontinuous Galerkin framework i: Total energy calculation, *Journal of Computational Physics* 231 (4) (2012) 2140–2154.
- [38] T. L. Beck, Real-space mesh techniques in density-functional theory, *Reviews of Modern Physics* 72 (4) (2000) 1041–1080.

- [39] Y. Saad, J. R. Chelikowsky, S. M. Shontz, Numerical methods for electronic structure calculations of materials, *SIAM Review* 52 (1) (2010) 3–54.
- [40] Y. Hasegawa, J.-I. Iwata, M. Tsuji, D. Takahashi, A. Oshiyama, K. Minami, T. Boku, F. Shoji, A. Uno, M. Kurokawa, et al., First-principles calculations of electron states of a silicon nanowire with 100,000 atoms on the k computer, in: *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, ACM, 2011, p. 1.
- [41] D. Osei-Kuffuor, J.-L. Fattebert, Accurate and scalable $O(N)$ algorithm for first-principles molecular-dynamics computations on large parallel computers, *Physical Review Letters* 112 (4) (2014) 046401.
- [42] P. Suryanarayana, P. P. Pratapa, A. Sharma, J. E. Pask, SQDFT: Spectral Quadrature method for large-scale parallel $O(N)$ Kohn–Sham calculations at high temperature, *Computer Physics Communications* 224 (2018) 288–298.
- [43] S. Ghosh, P. Suryanarayana, SPARC: Accurate and efficient finite-difference formulation and parallel implementation of density functional theory: Isolated clusters, *Computer Physics Communications* 212 (2017) 189–204.
- [44] N. D. Mermin, Thermal properties of the inhomogeneous electron gas, *Physical Review* 137 (1965) A1441–A1443.
- [45] P. Suryanarayana, D. Phanish, Augmented Lagrangian formulation of orbital-free density functional theory, *Journal of Computational Physics* 275 (0) (2014) 524–538.
- [46] S. Ghosh, P. Suryanarayana, Higher-order finite-difference formulation of periodic orbital-free density functional theory, *Journal of Computational Physics* 307 (2016) 634–652.
- [47] W. Gropp, E. Lusk, A. Skjellum, *Using MPI: portable parallel programming with the message-passing interface*, Vol. 1, MIT press, 1999.
- [48] A. Sharma, P. Suryanarayana, On the calculation of the stress tensor in real-space kohn-sham density functional theory, *The Journal of Chemical Physics* 149 (19) (2018) 194104.
- [49] D. Alfe, Ab initio molecular dynamics, a simple algorithm for charge extrapolation, *Computer Physics Communications* 118 (1) (1999) 31–33.
- [50] P. P. Pratapa, P. Suryanarayana, Restarted Pulay mixing for efficient and robust acceleration of fixed-point iterations, *Chemical Physics Letters* 635 (2015) 69–74.
- [51] A. S. Banerjee, P. Suryanarayana, J. E. Pask, Periodic Pulay method for robust and efficient convergence acceleration of self-consistent field iterations, *Chemical Physics Letters* 647 (2016) 31–35.

- [52] S. Kumar, Q. Xu, P. Suryanarayana, On preconditioning the self-consistent field iteration in real-space density functional theory, *Chemical Physics Letters* 739 (2020) 136983.
- [53] Y. Zhou, Y. Saad, M. L. Tiago, J. R. Chelikowsky, Self-consistent-field calculations using Chebyshev-filtered subspace iteration, *Journal of Computational Physics* 219 (1) (2006) 172–184.
- [54] Y. Zhou, Y. Saad, M. L. Tiago, J. R. Chelikowsky, Parallel self-consistent-field calculations via Chebyshev-filtered subspace acceleration, *Physical Review E* 74 (6) (2006) 066704.
- [55] Y. Zhou, J. R. Chelikowsky, Y. Saad, Chebyshev-filtered subspace iteration method free of sparse diagonalization for solving the Kohn–Sham equation, *Journal of Computational Physics* 274 (2014) 770–782.
- [56] P. P. Pratapa, P. Suryanarayana, J. E. Pask, Anderson acceleration of the Jacobi iterative method: An efficient alternative to Krylov methods for large, sparse linear systems, *Journal of Computational Physics* 306 (2016) 43–54.
- [57] P. Suryanarayana, P. P. Pratapa, J. E. Pask, Alternating Anderson–Richardson method: An efficient alternative to preconditioned Krylov methods for large, sparse linear systems, *Computer Physics Communications* 234 (2019) 278–285.
- [58] W. R. Burdick, Y. Saad, L. Kronik, I. Vasiliev, M. Jain, J. R. Chelikowsky, Parallel implementation of time-dependent density functional theory, *Computer Physics Communications* 156 (1) (2003) 22–42.
- [59] 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, *Physical Review B* 78 (7) (2008) 075109.
- [60] J. R. Shewchuk, An introduction to the conjugate gradient method without the agonizing pain, 1994.
- [61] J. Nocedal, Updating quasi-newton matrices with limited storage, *Mathematics of Computation* 35 (1980) 773.
- [62] E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, P. Gumbsch, Structural relaxation made simple, *Physical Review Letters* 97 (17) (2006) 170201.
- [63] W. H. Press, Numerical recipes 3rd edition: The art of scientific computing, Cambridge University Press, 2007.
- [64] M. P. Allen, D. J. Tildesley, *Computer Simulation of Liquids*, Oxford University Press, 1987.

- [65] W. C. Swope, H. C. Andersen, P. H. Berens, K. R. Wilson, A computer simulation method for the calculation of equilibrium constants for the formation of physical clusters of molecules: Application to small water clusters, *The Journal of Chemical Physics* 76 (1982) 637.
- [66] P. Minari, G. J. Martyna, M. E. Tuckerman, Algorithms and novel applications based on the isokinetic ensemble. i. biophysical and path integral molecular dynamics, *The Journal of Chemical Physics* 118 (2003) 2510.
- [67] J. P. Perdew, A. Zunger, Self-interaction correction to density-functional approximations for many-electron systems, *Physical Review B* 23 (1981) 5048–5079.
- [68] J. P. Perdew, Y. Wang, Accurate and simple analytic representation of the electron-gas correlation energy, *Physical Review B* 45 (23) (1992) 13244–13249.
- [69] J. P. Perdew, K. Burke, M. Ernzerhof, Generalized gradient approximation made simple, *Physical Review Letters* 77 (18) (1996) 3865.
- [70] B. Hammer, L. B. Hansen, J. K. Norskov, Improved adsorption energetics within density-functional theory using revised Perdew-Burke-Ernzerhof functionals, *Physical Review B* 59 (1999) 7413.
- [71] J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou, K. Burke, Restoring the density-gradient expansion for exchange in solids and surfaces, *Physical Review Letters* 100 (13) (2008) 136406.
- [72] D. R. Hamann, Optimized norm-conserving Vanderbilt pseudopotentials, *Physical Review B* 88 (8) (2013) 085117.
- [73] N. Troullier, J. L. Martins, Efficient pseudopotentials for plane-wave calculations, *Physical Review B* 43 (3) (1991) 1993–2006.
- [74] M. Schlipf, F. Gygi, Optimization algorithm for the generation of oncv pseudopotentials, *Computer Physics Communications* 196 (2015) 36–44.
- [75] K. Momma, F. Izumi, VESTA 3 for three-dimensional visualization of crystal, volumetric and morphology data, *Journal of applied crystallography* 44 (6) (2011) 1272–1276.
- [76] K. H. Khoo, T.-L. Chan, M. Kim, J. R. Chelikowsky, Ab initio molecular dynamics simulations of molten $\text{Al}_{1-x}\text{Si}_x$ alloys, *Physical Review B* 84 (2011) 214203.
- [77] R. Singh, P. Kroll, Structural, electronic, and magnetic properties of 13-, 55-, and 147-atom clusters of Fe, Co, and Ni: A spin-polarized density functional study, *Physical Review B* 78 (2008) 245404.
- [78] B. Austin, W. Bhimji, T. Butler, J. Deslippe, 2014 NERSC workload analysis, http://portal.nersc.gov/project/mpccc/baustin/NERSC_2014_Workload_Analysis_v1.1.pdf.

- [79] L. J. Vernon, IC Application Performance Team analysis, as part of IC Knights Special Project, Tech. Rep. LANL, 2015.
- [80] P. Suryanarayana, On spectral quadrature for linear-scaling density functional theory, *Chemical Physics Letters* 584 (2013) 182–187.
- [81] A. S. Banerjee, P. Suryanarayana, Cyclic density functional theory: A route to the first principles simulation of bending in nanostructures, *Journal of the Mechanics and Physics of Solids* 96 (2016) 605–631.
- [82] S. Ghosh, A. S. Banerjee, P. Suryanarayana, Symmetry-adapted real-space density functional theory for cylindrical geometries: Application to large group-IV nanotubes, *Physical Review B* 100 (12) (2019) 125143.
- [83] S. Kumar, P. Suryanarayana, Bending moduli for thirty-two select atomic monolayers from first principles, *arXiv preprint arXiv:2003.10936*.
- [84] P. Suryanarayana, K. Bhattacharya, M. Ortiz, Coarse-graining Kohn-Sham density functional theory, *Journal of the Mechanics and Physics of Solids* 61 (1) (2013) 38–60.
- [85] Q. Xu, A. Sharma, P. Suryanarayana, M-SPARC: Matlab-Simulation Package for Ab-initio Real-space Calculations, *SoftwareX* 11 (2020) 100423.