








Large scale and linear scaling DFT with the CONQUEST code

Cite as: J. Chem. Phys. **152**, 164112 (2020); <https://doi.org/10.1063/5.0005074>

Submitted: 19 February 2020 . Accepted: 03 April 2020 . Published Online: 28 April 2020

Ayako Nakata , Jack S. Baker , Shereif Y. Mujahed , Jack T. L. Poulton, Sergiu Arapan, Jianbo Lin, Zamaan Raza , Sushma Yadav, Lionel Truflandier , Tsuyoshi Miyazaki , and David R. Bowler 

COLLECTIONS

Paper published as part of the special topic on [Electronic Structure Software](#)

Note: This article is part of the JCP Special Topic on Electronic Structure Software.



View Online



Export Citation



CrossMark

ARTICLES YOU MAY BE INTERESTED IN

[Recent developments in the general atomic and molecular electronic structure system](#)

The Journal of Chemical Physics **152**, 154102 (2020); <https://doi.org/10.1063/5.0005188>

[Quantum ESPRESSO toward the exascale](#)

The Journal of Chemical Physics **152**, 154105 (2020); <https://doi.org/10.1063/5.0005082>

[ABINIT: Overview and focus on selected capabilities](#)

The Journal of Chemical Physics **152**, 124102 (2020); <https://doi.org/10.1063/1.5144261>

Lock-in Amplifiers
up to 600 MHz



Watch



Large scale and linear scaling DFT with the CONQUEST code

Cite as: J. Chem. Phys. 152, 164112 (2020); doi: 10.1063/5.0005074

Submitted: 19 February 2020 • Accepted: 3 April 2020 •

Published Online: 28 April 2020



Ayako Nakata,¹ Jack S. Baker,^{2,3} Shereif Y. Mujahed,^{2,3} Jack T. L. Poulton,^{2,3} Sergiu Arapan,^{1,a)} Jianbo Lin,¹ Zamaan Raza,^{1,b)} Sushma Yadav,¹ Lionel Truflandier,⁴ Tsuyoshi Miyazaki,^{1,c)} and David R. Bowler^{1,2,3,d)}

AFFILIATIONS

¹International Centre for Materials Nanoarchitectonics (MANA), National Institute for Materials Science (NIMS), 1-1 Namiki, Tsukuba, Ibaraki 305-0044, Japan

²London Centre for Nanotechnology, University College London, 17-19 Gordon St., London WC1H 0AH, United Kingdom

³Department of Physics & Astronomy, University College London, Gower St., London WC1E 6BT, United Kingdom

⁴Institut des Sciences Moléculaires, Université Bordeaux, 351 Cours de la Libération, 33405 Talence, France

Note: This article is part of the JCP Special Topic on Electronic Structure Software.

^{a)}**Present address:** IT4Innovations, VŠB - Technical University of Ostrava, 17. listopadu 2172/15, 70800 Ostrava-Poruba, Czech Republic.

^{b)}**Present address:** Jülich Centre for Neutron Science (JCNS) at Heinz Maier-Leibnitz Zentrum (MLZ), Lichtenbergstraße 1, 85748 Garching, Germany.

^{c)}**Electronic mail:** MIYAZAKI.Tsuyoshi@nims.go.jp

^{d)}**Author to whom correspondence should be addressed:** david.bowler@ucl.ac.uk

ABSTRACT

We survey the underlying theory behind the large-scale and linear scaling density functional theory code, CONQUEST, which shows excellent parallel scaling and can be applied to thousands of atoms with diagonalization and millions of atoms with linear scaling. We give details of the representation of the density matrix and the approach to finding the electronic ground state and discuss the implementation of molecular dynamics with linear scaling. We give an overview of the performance of the code, focusing in particular on the parallel scaling, and provide examples of recent developments and applications.

Published under license by AIP Publishing. <https://doi.org/10.1063/5.0005074>

I. INTRODUCTION

As computing power has increased and methods have become more sophisticated, computational modeling of materials systems has taken its place alongside experiment and theory in science. Electronic structure methods give insight into bonding and electronic properties of systems, and density functional theory (DFT) is now the *de facto* method used in fields, such as diverse as physics, chemistry, earth sciences, materials science, and biochemistry.¹

However, almost all DFT calculations are focused on a relatively small system size of a few hundred atoms (while feasible, calculations involving more than a thousand atoms are still considered expensive). This size limitation mainly comes from the cost and

scaling of standard DFT implementations (asymptotically, the computer time required scales with the cube of the number of atoms in the system, while the memory scales with the square of the number of atoms).

The use of increasingly large numbers of computational cores is one route to larger scale DFT calculations. Indeed, high-performance computing centers have recently scaled to tens of thousands and hundreds of thousands of cores and exascale computing is scheduled to arrive in the next few years. However, efficient use of systems of this size requires care to ensure that the parallel scaling of any given code remains efficient as the number of cores becomes large. A real-space formulation and the use of local basis functions to represent the Kohn–Sham (KS) eigenstates can help with the

parallel efficiency of the code and enables calculations on larger systems, often with several thousand atoms.^{2–10} Moreover, there is no inherent reason why DFT implementations should scale with the cube of the system size;¹¹ in fact, with local basis functions, DFT can be formulated in a linear scaling approach.^{12,13}

The use of a few hundred atoms in a calculation potentially introduces many approximations and, in some cases, errors. There are systems where larger scale calculations are necessary to model the properties correctly. Features with dimensions in nanometres or more or features with strain fields over nanometres will be seriously restricted by a small calculation as the strain will not be fully relaxed. Dilute alloys can change their properties if the concentration is incorrect (for instance, the metal–insulator transition in doped silicon occurs at around 2×10^{19} dopants per cubic centimetre or one dopant per 50 000 silicon atoms). Calculations on biomolecules often use a small quantum mechanical (QM) region embedded into a larger forcefield (or molecular mechanics, MM) region, and it has been shown¹⁴ that the size and choice of the QM region affect the results strongly, with sizes of at least 500 atoms being required. While accurate forcefields are available and new approaches to the fitting of forcefields are being developed, simulations requiring the electronic structure or bond making and breaking must use DFT or a related approach.

CONQUEST is a large-scale and linear scaling DFT code, which is designed to scale efficiently in parallel and to be applicable to systems with thousands of atoms with full diagonalization for the ground state and to systems with millions of atoms with a linear scaling solution of the ground state. It has recently been released under an open-source MIT license,¹⁵ and in this paper, we describe the implementation of the code and various recent applications. We first describe the approaches to representing the density matrix and finding the electronic ground state and then consider how eigenvectors can be calculated for systems large enough to require linear scaling. We also discuss the calculation of exact exchange (EXX) with linear scaling. We then turn to the movement of atoms considering, in particular, the calculation of stress and the implementation of molecular dynamics (MD), both with reference to linear scaling. After presenting the performance of CONQUEST for various systems, we then illustrate several applications with many thousands of atoms before concluding.

II. METHODS: ELECTRONIC STRUCTURE

A. Pseudopotentials and pseudo-atomic orbitals

The default pseudopotential format used in CONQUEST is the optimized norm-conserving Vanderbilt pseudopotential (ONCVPSP) developed by Hamann¹⁶ from Vanderbilt's approach.¹⁷ This approach has been used to generate two libraries covering most of the periodic table: PseudoDojo¹⁸ and SG15.¹⁹ Both of these libraries give very good values using the Delta comparison to all-electron results,²⁰ with PseudoDojo showing an accuracy comparable to the best PAW (projector augmented wave) results (when using a fully converged plane wave basis set).

This approach to pseudopotential generation is focused on accuracy and, as a consequence, includes semi-core states for many elements beyond the third row of the periodic table, as well as partial core corrections²¹ in many cases. The spacing of the real-space

integration grid required for these pseudopotentials may be finer than that other approaches would give. CONQUEST is also compatible with pseudopotentials in the Troullier–Martins form as generated by SIESTA, which are less stringent and thus less expensive.

As is common for local orbital approaches, we use a neutral atom approach²² where the local part of the pseudopotential and a (neutral) atomic density cancel out each other at large distances from the atom; among other things, this removes the need for an Ewald sum for the ion–ion energies. However, it has been shown²³ that the resulting potential can be deep and require a fine integration grid for convergence. To alleviate this problem, we have implemented the neutral atom projector approach²³ and will report detailed testing in a future publication.

CONQUEST can use pseudo-atomic orbitals (PAOs) as basis functions to represent the support functions (discussed in Sec. II B below), while the valence orbitals also serve to generate the atomic density. We generate the PAOs²⁴ by solving the Schrödinger equation for an isolated, confined atom with a pseudopotential. The basis sets are formed from valence orbitals (with a given number of functions, or zetas, for each angular momentum, each with its own confinement) and polarization functions (typically with $l_v + 1$, where l_v is the angular momentum of the highest occupied state, and a given different number of functions to the valence states).

The confinement is equivalent to a hard wall and can be applied either as a radial cutoff or an energy shift for the orbital, which generates a different radius for each orbital. The default basis sets in CONQUEST are generated either with the same radial cutoffs for all shells or with the same energy shifts. For the energy shifts, we use a tight confinement (a shift of 2 eV) and a loose confinement (a shift of 0.02 eV) to generate two zeta functions, with a third generated when needed using the average radius of the first two. The uniform radial confinement is found as the average of the radii for all valence functions with the large or small energy confinement, with the third radius again found as an average. Semi-core states only use one function with a loose confinement (in this case, the orbital is strongly confined so that even a very small energy shift gives good results).

We have tested our default basis sets [single zeta with polarization (SZP), double zeta with polarization (DZP), and triple zeta with triple polarization (TZTP)] against converged plane wave calculations using the same pseudopotentials.²⁴ We used a wide variety of systems: elemental semiconductors (C, Si, and Ge), oxides (SiO₂ in both α -quartz and stishovite phases, MgO, SrTiO₃, PbTiO₃, and MgSiO₃), non-magnetic bcc Fe, and weakly bonded systems (ice XI and BN sheets). In all cases, we showed that the TZTP basis sets reproduced the converged plane wave results with excellent accuracy: better than 1% in bulk modulus and 0.1% of the lattice constant for a variety of bulk systems (for full details, see Ref. 24).

B. Representing the density matrix

CONQUEST uses the density matrix as the fundamental description of the system being modelled (in contrast to the wavefunctions as is common in many DFT codes). The density matrix is represented in terms of support functions,^{25,26} $\phi_{i\alpha}(\mathbf{r})$, where i indexes the atom and α is the function on the atom, and can be written as

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{i\alpha, j\beta} \phi_{i\alpha}(\mathbf{r}) K_{i\alpha, j\beta} \phi_{j\beta}(\mathbf{r}'). \quad (1)$$

(Note that the density matrix can easily be written in terms of the wavefunctions when these are available or found by linear scaling optimization, as described in Sec. II C. Note also that, while we assume a non-spin-polarized calculation here, the extension to spin polarization is straight-forward and has been implemented in CONQUEST.)

The support functions are local functions that move with the atoms and are strictly localized within a sphere. They are used to form the Hamiltonian and overlap matrices and to represent the density matrix,

$$H_{i\alpha,j\beta} = \int d\mathbf{r} \phi_{i\alpha}(\mathbf{r}) \hat{H} \phi_{j\beta}(\mathbf{r}). \quad (2)$$

The support functions themselves are defined either as primitive PAOs (in a one-to-one mapping) or are represented in terms of one of two basis sets: the PAOs or blip functions.²⁷ We write

$$\phi_{i\alpha}(\mathbf{r}) = \sum_s b_{i\alpha s} \theta_s(\mathbf{r}), \quad (3)$$

where the basis functions $\theta_s(\mathbf{r})$ (either pseudo-atomic orbitals or cubic B-splines) are discussed further below.

1. Multi-site support functions

Since the primitive PAOs are localized around the atoms, we can use them as support functions without any modifications, and a large PAO basis set gives high accuracy, as shown in Sec. II A. However, the computational cost of calculations scales with the cube of the number of support functions. For large-scale calculations, we need to reduce the number of support functions as much as possible. We note that there is a strong link between the basis set chosen and the number of support functions that can be used.²⁸

Multi-site support functions (MSSFs)^{29,30} offer one way to reduce the number of support functions. The MSSFs are constructed as linear combinations of the PAOs not only on the target atom but also on its neighboring atoms, defined by a cutoff radius r_{MS} ,

$$\phi_{i\alpha}(\mathbf{r}) = \sum_{k \in \text{neighbours}} \sum_{\mu \in k} C_{i\alpha,k\mu} \chi_{k\mu}(\mathbf{r}), \quad (4)$$

where $\chi_{k\mu}$ is a PAO on atom k , α is the index of the support functions of atom i , μ is the index of the PAOs of its neighboring atoms k (including i itself), and $C_{i\alpha,k\mu}$ are the linear combination coefficients. Since the MSSFs are no longer atomic orbitals but local molecular orbital (MO)-like functions, the number of MSSFs can be equal to that of a minimal basis.

In CONQUEST, two methods have been implemented to determine the linear-combination coefficients of the MSSFs. One of the methods is the local-filter-diagonalization (LFD) method proposed by Rayson and Briddon,^{5,29,31} and the other is numerical optimisation,³⁰ which will be explained in Sec. II C. In the LFD method, as shown in Eq. (5), the coefficients \mathbf{C} for each atom are found by diagonalizing a subspace Hamiltonian matrix defined by a cluster of radius r_{LFD} centered on the atom constructed with PAOs; the coefficients, \mathbf{C}_{sub} , for the resulting eigenstates, or local molecular orbitals, are projected onto trial support functions, \mathbf{t} , localized on the target atom,

$$\mathbf{C} = \mathbf{C}_{\text{sub}} f(\epsilon_{\text{sub}}) \mathbf{C}_{\text{sub}}^{\text{T}} \mathbf{S}_{\text{sub}} \mathbf{t}, \quad (5)$$




FIG. 1. Difference of total energy per atom (eV) with MSSFs from the full primitive PAO result with respect to the multi-site range r_{MS} . The local diagonalization range r_{LFD} was set to be equal to r_{MS} . The circles and squares correspond to bulk Si and Al. Data taken with permission from A. Nakata, D. R. Bowler, and T. Miyazaki, J. Chem. Theory Comput. **10**, 4813 (2014). Copyright 2014 American Chemical Society.

where \mathbf{S}_{sub} is the overlap matrix and $f(\epsilon_{\text{sub}})$ is a Fermi function with a local Fermi level ϵ_{sub} used to exclude high energy unoccupied local molecular orbitals. Since the MSSFs will depend on the charge density of the system, which will in turn depend on the MSSFs, the linear-combination coefficients need to be determined self-consistently.²⁹ r_{MS} should be equal to or smaller than the subspace region in the LFD method r_{LFD} .

The accuracy of the MSSFs depends on r_{MS} . In Fig. 1, we see that the deviation from the full, unrestricted, primitive PAO result decreases exponentially not only in a gapped system (bulk Si) but also in a metallic system (bulk Al). The number of MSSFs per atom is four, while that of the TZP PAOs is 17 in both Si and Al, giving a fourfold reduction in number and a significant speed-up. An example of the computational time with the MSSFs is demonstrated in Sec. IV A.

2. On-site support functions

When using a linear scaling solver, as described in Sec. II C, we require a sparse approximation to the inverse overlap matrix to act as a metric.³² We have found that multiple zeta basis sets and multi-site support functions are not compatible with our standard linear-scaling method for finding this inverse overlap (Hotelling's method). The reasons for this failure are not yet clear and are under investigation but are most likely to arise from the assumed sparsity pattern of the matrix and the starting value used.³³ As a result, we have been limited in the basis sets that can be used for linear scaling. Blip functions, which will be discussed in Sec. II B 3, offer a route to an accurate linear scaling basis set; however, PAOs are often convenient and efficient, and a restriction to SZ or SZP PAO basis sets is limiting.

We have found recently, however, that an adaptation of the MSSF approach allows linear scaling solution for a sparse inverse overlap matrix while retaining accurate basis sets: on-site support functions (OSSFs). We restrict the PAOs forming the support functions for an atom i to its own PAOs; however, we must be careful to

respect any symmetry of the atomic lattice so that the space spanned by the support functions of the atoms decomposes into complete irreducible representations of the symmetry group.²⁸ The simplest way to ensure that this is respected is to increase the number of support functions such that it encompasses all angular momenta of the PAOs (e.g., for PAOs including $l = 0 \rightarrow 2$, we would need nine SFs, while for PAOs only including $l = 0$ and $l = 2$, we would need six SFs).

This approach bears some similarity to the polarized atomic orbital method^{34,35} although that method imposes no restrictions on the number of functions and uses a different approach to find the orbital coefficients. In our approach, we use the LFD method described in Sec. II B 1 using a trial vector, which is extended to include the polarization orbitals. We find that the resulting support functions can be inverted efficiently (interestingly, it is often more efficient than a simple SZP PAO basis set).

When using OSSF with linear scaling, we are still investigating the most efficient approach for finding the ground state; this involves optimizing the density matrix, the OSSF coefficients, and the charge density. Introducing self-consistency between the OSSF coefficients and the charge density is straightforward, but in a naive loop would add considerably to the computational time. Optimizing the energy with respect to the OSSF coefficients is also straightforward, but the most efficient approach (i.e., when to update which parts of the optimization) requires further research.

The basis sets found using OSSFs are significantly better than the simple SZP PAO basis set, as shown in Fig. 2. Here, we see that, for a slightly disturbed eight atom bulk silicon cell, as the LFD range is progressively increased, the rate of convergence of the density matrix optimization improves. The quality of the resulting approximate inverse overlap matrix is also improved, and the energy and forces on the atoms are significantly better with the OSSF basis sets. We show results for different basis sets in Table I: primitive PAOs (SZP, DZP, and TZTP), MSSF for different ranges, and OSSFs. The MSSF and OSSF calculations do not update the SF coefficients after finding self-consistency; for MSSFs, the LFD range is set to 15 bohrs




FIG. 2. Convergence of linear scaling density matrix optimization for different basis sets: SZP and OSSF with different LFD ranges. The system considered is an eight atom bulk silicon cell, slightly disturbed from the perfect crystal structure.

TABLE I. Comparison of different basis sets for an eight atom bulk silicon cell slightly disturbed from the perfect crystal structure. Primitive basis sets have 9, 13, and 27 support functions, respectively; MSSFs have four support functions and a LFD range of 15 bohrs and OSSFs have nine support functions. MSSFs and OSSFs are not updated after the initial calculation of the coefficients. The range on the density matrix for the $\mathcal{O}(N)$ calculations was 16 bohrs.

Basis	Energy (hartree)	Force (hartree/bohr)	Time relative to TZTP
SZP	−33.714	−0.001 70	0.12
DZP	−33.819	−0.001 51	0.23
TZTP	−33.838	−0.001 45	1.00
MSSF 5 bohrs	−33.800	−0.001 54	0.68
MSSF 8 bohrs	−33.821	−0.001 43	0.73
MSSF 10 bohrs	−33.818	−0.001 43	0.73
MSSF 12 bohrs	−33.828	−0.001 44	0.81
OSSF 8 bohrs	−33.625	−0.001 62	0.63
OSSF 10 bohrs	−33.813	−0.001 42	0.63
OSSF 12 bohrs	−33.812	−0.001 42	0.60
OSSF 15 bohrs	−33.820	−0.001 42	0.86
OSSF $\mathcal{O}(N)$ 8 bohrs	−33.605	−0.001 74	4.20
OSSF $\mathcal{O}(N)$ 10 bohrs	−33.784	−0.001 54	3.68
OSSF $\mathcal{O}(N)$ 12 bohrs	−33.782	−0.001 54	2.87
OSSF $\mathcal{O}(N)$ 15 bohrs	−33.792	−0.001 54	2.47

throughout. Note that the energy is not variational with respect to the radius in this simple process, although with optimization (as described in Sec. II C 3), it will be variational.

As seen in Fig. 2, it is evident that the OSSFs give a significant improvement to the performance of the linear scaling solver, and from Table I, we can see that they are comparable to the MSSFs for the accuracy of forces and the timing. Most notably, we see that with the OSSF found with the LFD radius set to 15 bohrs, the linear scaling solver is only 2 times slower than diagonalization with the full TZTP basis set and 8 times slower than the DZP basis set. This performance difference is expected for such a small system where linear scaling solvers are less efficient than diagonalization and related solvers. However, it suggests that the choice of basis functions is important in implementing linear scaling. We note that the quality of both OSSF and MSSF basis sets would be improved by the optimization of the coefficients, as described in Sec. II C 3, but even with these simple approximations, good performance is achieved.

3. Blip functions

While PAOs are a convenient basis set, they do not permit systematic convergence of the energy with respect to the basis: while adding extra basis functions will increase the size of the variational space and lower the energy, there are two parameters that offer different degrees of freedom (maximum angular momentum and the number of radial functions, or zetas, per angular momentum channel), and there are no guarantees of how adding to each parameter will change the energy.

The blip functions,²⁷ which are piecewise continuous cubic splines defined on a cubic grid that moves with the atoms, are a basis set that can be systematically converged. The blip grid spacing corresponds directly to a plane wave cutoff, allowing the basis set to be improved systematically (of course, the support functions are confined within a radius, but it has been shown that the total energy converges variationally and rapidly with this radius^{27,36}). The most efficient procedure for initialization and optimization of the blip coefficients along with the charge density and for linear scaling approaches to finding the density matrix is the subject of on-going research.

C. Solving for the ground state

The ground state electronic structure in CONQUEST is defined by three related quantities: the support functions, the density matrix, and self-consistency between the charge density and the Kohn–Sham potential. The self-consistency procedure is a standard part of DFT and related codes,³⁷ and we implement the Pulay approach,³⁸ which works efficiently.

While the overall search for the ground state could be considered as an optimization in a space formed by both the support function coefficients and the density matrix elements, it is easier to consider how the density matrix is found for a given set of support functions and then to discuss methods for optimizing the energy with respect to the support function coefficients. We consider first the two approaches to solving for the density matrix: exact diagonalisation,³⁹ which scales cubically but makes no approximations, and linear scaling, which imposes a range on the density matrix.

1. Density matrix: Exact diagonalization

We perform diagonalization of the Hamiltonian using ScaLAPACK⁴⁰ and are also investigating the use of ELPA⁴¹ (which uses the same interface and may scale better to large numbers of processes).

Since we apply periodic boundary conditions to our simulation cell, the Brillouin zone must be sampled appropriately; we have implemented the Monkhorst–Pack⁴² sampling method as a default approach to Brillouin zone sampling, but any arbitrary set of k-points can be used. (At present, we do not account for the symmetry of the simulation cell beyond time-reversal symmetry, as the code is designed for large-scale simulations that are unlikely to show significant symmetries.)

The Kohn–Sham eigenstates are represented in terms of the support functions, with the density matrix found as

$$|\psi_{n\mathbf{k}}\rangle = \sum_{i\alpha} c_{i\alpha}^{n\mathbf{k}} |\phi_{i\alpha}\rangle, \quad (6)$$

$$K_{i\alpha,j\beta} = \sum_{n\mathbf{k}} f_{n\mathbf{k}} w_{\mathbf{k}} c_{i\alpha}^{n\mathbf{k}} (c_{j\beta}^{n\mathbf{k}})^*, \quad (7)$$

where the weight of each k-point is given as $w_{\mathbf{k}}$ and the occupancy of the eigenstate is $f_{n\mathbf{k}}$ (which is found using a simple Fermi–Dirac distribution or the Methfessel–Paxton⁴³ approach). The diagonalization at each k-point can be assigned to a sub-group of processes, enabling a calculation using many k-points to be sped up significantly.

2. Density matrix: Linear scaling

To achieve linear scaling in computational time with the system size, we restrict the range of the density matrix (the range of a matrix $A_{i\alpha,j\beta} = 0$ is defined in terms of the distance between atoms i and j , $R_{ij} = |\mathbf{R}_i - \mathbf{R}_j|$, and is restricted by setting matrix elements to zero when R_{ij} is greater than a cutoff distance R_c), and optimize the band energy, $E_{\text{band}} = 2 \text{Tr}[KH]$, with respect to the density matrix elements. When this approach is coupled with strictly local basis functions, all matrices are sparse and all matrix operations scale linearly with the system size.

During the optimization, we must constrain the electron number (a relatively straightforward task²⁵) and also the idempotency of the density matrix (a much more complex task in a variational context). We follow the LNV (Li, Nunes, Vanderbilt) approach where we write the density matrix K in terms of an *auxiliary* density matrix L using the McWeeny transform,

$$K = 3LSL - 2LSLSL. \quad (8)$$

This drives the density matrix, K , toward idempotency (strictly, it is driven toward *weak* idempotency, where the eigenvalues lie between zero and one, but may not be exactly zero and one). If, as above, we write $E = \text{Tr}[KH]$, then we can use the gradient $\partial E / \partial L_{i\alpha,j\beta}$ to minimize the energy with respect to the density matrix, and the density matrix K will be driven toward idempotency as the minimization proceeds. Here, the range of K is the same as the range of the Hamiltonian (which is naturally sparse); it is L whose range is restricted and this range sets the accuracy of the calculation.

The initial density matrix is generated from the Hamiltonian using an iterative procedure based on a generalization of the McWeeny transform.^{46,47} We use an approximate, sparse inverse overlap matrix as the metric for the optimization, found using the iterative Hotelling method. As discussed in Sec. II B, this has certain consequences for the basis sets that can be used, but both OSSF and blip basis sets are promising for accurate, linear scaling calculations.

3. Optimizing support functions

As mentioned in Sec. II B, we can construct support functions by taking linear combinations of PAOs (MSSFs or OSSFs) or blips. The linear combination coefficients can be optimized numerically by minimizing the DFT total energy with respect to the coefficients.³⁰ For MSSFs and OSSFs, the coefficients obtained by the LFD method generally form good initial values for the numerical optimization. The initial blip coefficients are found as a best fit to PAOs. In this section, we demonstrate optimization of SF coefficients for the MSSFs, although the formalism is identical for the other approaches. We note that these optimization processes are liable to ill-conditioning, which can be mitigated.⁴⁸

Figure 3 shows the energy–volume (E – V) curves of bulk Si³⁰ calculated with simple LFD (filled symbols) and optimization of the MSSF coefficients (open symbols). The number of MSSFs per atom is four, while that of the primitive TZDP (3s3p2d) PAOs is 22. Table II summarizes the lattice constant a_0 obtained by fitting the E – V curves with the Birch–Murnaghan equation. The results are improved, i.e., becoming closer to the results of the primitive PAOs,




FIG. 3. Energy–volume curves of bulk Si demonstrating the effect of optimizing MSSF coefficients after the initial LFD process. Symbols correspond to the calculated energies by primitive TZDP PAOs and MSSFs with multi-site ranges (r_{MS}) 17.0 bohrs, 8.0 bohrs, and 5.0 bohrs using the LFD method (filled symbols) and numerical optimization (open symbols). The local filter diagonalization range r_{LFD} was set to be equal to r_{MS} . Adapted with permission from A. Nakata, D. R. Bowler, and T. Miyazaki, Phys. Chem. Chem. Phys. 17, 31427 (2015). Copyright 2015 PCCP Owner Societies.

by the numerical optimization in all cases. When r_{MS} is large, e.g., 17.0 bohrs, since the MSSFs found with LFD give accurate results, the change from numerical optimization is small. When r_{MS} is 8.0 bohrs, the difference with and without the numerical optimization is significantly larger, but both LFD and the numerical optimization show reasonable accuracy. On the other hand, when r_{MS} is as small as 5.0 bohrs, the result with the LFD method is not accurate, with a 1% deviation from the full TZDP result, but we find significant improvement of the accuracy from numerical optimization, reducing the percentage deviation to 0.2%.

TABLE II. Lattice constants a_0 of bulk Si calculated with MSSFs with ranges r_{MS} of 5.0 bohrs, 8.0 bohrs, and 17.0 bohrs and percent deviations (%D) from a_0 calculated with the primitive TZDP PAOs.

r_{MS}	a_0		%D	
	LFD	Opt	LFD	Opt
5.0	5.447	5.406	1.0	0.2
8.0	5.403	5.400	0.2	0.1
17.0	5.393	5.395	0.0	0.0
TZDP	5.395		...	

D. Electronic structure for large systems

Linear scaling, or $\mathcal{O}(N)$, calculations that work with the density matrix implicitly integrate over energy and produce only the sum of the occupied eigenvalues and not any of the Kohn–Sham eigenstates of the system. However, we often want to know individual eigenstates to analyze the electronic structure of the system, although generally within a relatively small energy range. These can be found efficiently from the converged ground-state Hamiltonian by using the Sakurai–Sugiura (SS) method.⁴⁹ The SS method^{50,51} is an efficient interior eigenproblem solver for large sparse matrices using contour integrals in the complex plane, which provides the eigenvalues and eigenvectors in a finite eigenvalue range with high parallel efficiency. We use the SS method as it is much more scalable in parallel than other approaches, such as shift-and-invert Lanczos.⁵² We first optimize the electronic Hamiltonian with the $\mathcal{O}(N)$ method in CONQUEST and then obtain the eigenstates of the Hamiltonian in a finite energy window with a one-shot SS calculation. Here, we demonstrate the usefulness of the combination of the $\mathcal{O}(N)$ and SS methods by showing two examples.

The first example is the energy-specific electron-density distribution in a hut-shaped Ge cluster on the Si(001) surface consisting of 23 737 atoms (the physical system is described in more detail in Sec. V A). Figure 4 shows the electron density distribution in the energy range $[-0.01 \text{ eV}; +0.02 \text{ eV}]$ around the Fermi level, obtained by calculating the Kohn–Sham eigenvectors in this range with the SS method. The calculation for the eigenvalues and eigenvectors required 146 s using 64 nodes of the K supercomputer. We also calculated the eigenstates in the same energy range for a larger Ge/Si(001) system, consisting of 194 573 atoms, in 2399 s using 6400 nodes. (Note that the times quoted are just for the SS eigensolutions, which are performed as post-processing calculations using the output from CONQUEST.)

The second example is the density of states (DOS) of a DNA system in water, which consists of 3439 atoms. The DOS calculated with MSSFs (see Sec. II B 1) (4774 functions) and primitive PAOs (27 883 functions) and their difference are shown in Fig. 5. The DOS from the MSSFs is very close to the full primitive PAO DOS in the occupied states and in the unoccupied states near the Fermi level, while the DOS in the unoccupied states far from Fermi level are quite different. This is because the MSSFs are determined by optimizing the occupied states with a small number of support functions, and the accuracy of the MSSFs for unoccupied states often becomes poor. To improve this poor description, we can use the SS method. First, we optimize the electronic density of the target system using MSSFs, and we re-construct the electronic Hamiltonian using the primitive PAOs with the optimized density. Then, we use the SS method to




FIG. 4. Electronic density distributions (blue) of the Ge hut clusters (light gray) on the Si(001) (dark gray) (totally 23 737 atoms) around the Fermi level. Reprinted with permission from Nakata *et al.*, J. Chem. Theory Comput. 13, 4146 (2017). Copyright 2017 American Chemical Society.




FIG. 5. Density of states of hydrated DNA obtained with (a) multi-site support functions (MSSFs) with multi-site range 8.0 bohrs (blue) and DZP PAOs using MSSF charge density (red) and (b) DZP PAOs using DZP SCF charge density (black). The difference of (a) MSSF (blue) and DZP with the MSSF density (red) from (b) DZP is also shown in (c). The eigenstates in (a) and (b) were obtained by using the SS method. Reprinted with permission from Nakata *et al.*, J. Chem. Theory Comput. **13**, 4146 (2017). Copyright 2017 American Chemical Society.

obtain the eigenstates. Thus, we can obtain the DOS even in unoccupied states far from the Fermi level quite accurately, as shown in Fig. 5.

E. Exact exchange

Exact exchange (EXX) correction to the original Kohn–Sham formulation of DFT, leading to the class of *hybrid* exchange–correlation functionals, has become very popular, since in the vast majority of cases, it improves the overall reliability of the DFT predictions. Depending on the implementation—mainly basis set and boundary conditions—orbital dependence is introduced to the KS-DFT formalism via the EXX energy standard expression,

$$E_x = -\frac{1}{4} \int d\mathbf{r} d\mathbf{r}' \frac{\rho(\mathbf{r}, \mathbf{r}') \rho(\mathbf{r}', \mathbf{r})}{|\mathbf{r} - \mathbf{r}'|} \\ = -\sum_{n,m} \int d\mathbf{r} d\mathbf{r}' \frac{\psi_n^*(\mathbf{r}) \psi_n(\mathbf{r}') \psi_m(\mathbf{r}) \psi_m(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}, \quad (9)$$

where $\{\psi_n\}$ is the set of N occupied KS states, and can bring the computational cost to a prohibitive level more rapidly than the pure local-density approximation (LDA)/GGA (generalized gradient approximation) DFT when increasing the system size. Within the framework of CONQUEST, where the density matrix of Eq. (9) is expanded onto a set of M localized and real basis functions, the exchange energy reads $E_x = -\text{Tr}\{KX\}$, with the exchange matrix (X) elements given by

$$X_{i\alpha,j\beta} = \sum_{k\mu,l\nu} \int d\mathbf{r} d\mathbf{r}' \frac{\rho_{i\alpha,k\mu}(\mathbf{r}) K_{k\mu,l\nu} \rho_{l\nu,j\beta}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} \\ = \sum_{k\mu,l\nu} \int d\mathbf{r} d\mathbf{r}' \frac{\phi_{i\alpha}(\mathbf{r}) \phi_{k\mu}(\mathbf{r}) K_{k\mu,l\nu} \phi_{l\nu}(\mathbf{r}') \phi_{j\beta}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}. \quad (10)$$

As a result, calculation of X requires to evaluate at most M^4 electron repulsion integrals (ERIs) defined by the integrand of Eq. (10). The first equality in the equation above outlines the fact that evaluating an ERI is similar to computing the Hartree energy, with in

place of the full electronic density localized pair densities ($\rho_{i\alpha,k\mu}$, $\rho_{l\nu,j\beta}$) coupled by the density matrix elements $K_{k\mu,l\nu}$. Consequently, ERI calculation can be performed by solving a Poisson equation into a predefined local cell. Note that, contrary to the Hartree energy, the solution of this equation should be free of periodic boundary conditions.

When dealing with a numerical basis set, such as the PAOs, several options to compute the ERIs are available, with, for instance, the semi-analytic solution given by Toyoda and Ozaki^{53,54} combining the fast-spherical Bessel transform for the radial integration and a more traditional analytic method for the spherical harmonic part. Another approach is based on the experience of Gaussian-type orbital (GTO) ERI solvers.⁵⁵ In that case, the PAO-ERIs are transformed into a set of contracted GTO-ERIs, which are then calculated analytically.^{56,57} Instead, we use a route that circumvents the calculation of the ERIs and works for any smooth finite-range functions, which is particularly well suited for $\mathcal{O}(N)$ approaches based on the pseudopotential approximation. The key part is to perform the sum over the index $l\nu$ before solving for the Coulomb potential of the pair densities; this simple re-ordering increases the efficiency of the procedure markedly. For this, we introduce the contraction functions, $\Phi_{k\mu}(\mathbf{r}')$, as

$$\Phi_{k\mu}(\mathbf{r}') = \sum_{l\nu} K_{k\mu,l\nu} \phi_{l\nu}(\mathbf{r}'). \quad (11)$$

It should be noted that the domain over which these functions are defined requires some care.⁵⁸ The sum over $l\nu$ needs only to include those support functions $\phi_{l\nu}$ overlapping with $\phi_{j\beta}$, as $\Phi_{k\mu}$ will be multiplied by this function. Contracted densities are then defined as

$$\tilde{\rho}_{k\mu,j\beta}(\mathbf{r}') = \Phi_{k\mu}(\mathbf{r}') \phi_{j\beta}(\mathbf{r}'), \quad (12)$$

and the resulting Coulomb potential,

$$\tilde{v}_{k\mu,j\beta}(\mathbf{r}) = \int d\mathbf{r}' \frac{\tilde{\rho}_{k\mu,j\beta}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}, \quad (13)$$

is calculated by solving Poisson's equation. Once the potential has been found, a further contraction over $k\mu$ is performed to create

$$\Omega_{j\beta}(\mathbf{r}) = \sum_k \tilde{v}_{k\mu,j\beta}(\mathbf{r}) \phi_{k\mu}(\mathbf{r}), \quad (14)$$

where, again, the sum over support functions $k\mu$ needs only to include those functions that overlap with functions $i\alpha$. The exchange matrix elements are then calculated by numerical integration,

$$X_{i\alpha,j\beta} = \int d\mathbf{r} \phi_{i\alpha}(\mathbf{r}) \Omega_{j\beta}(\mathbf{r}). \quad (15)$$

The process—from Eqs. (11)–(15)—by which the EXX is calculated in CONQUEST will be referred to as the contraction reduction integral (CRI). The set of function $\Omega_{j\beta}$ is effectively defined by the density matrix range and the need for $j\beta$ to overlap with atoms lv , which naturally controls the number of functions entering in the sums of Eqs. (12) and (14). Note that the calculation time can be reduced by imposing a range condition (R_X) on the exchange matrix. This is related to the sparsity property¹¹ of $\rho(\mathbf{r}, \mathbf{r}')$ and the truncation of all the operators involved in the Hamiltonian.²⁵

Practical tests on the efficiency of this approach algorithm were carried out on a set of isolated water clusters (H_2O)_{*n*} ($n \leq 20$) with fused cubes structures.⁵⁹ Calculations of the exchange energy were performed after the KS density matrix had been converged using the standard self-consistent-field (SCF) method. As a result, the timings presented below for exact exchange (EXX) energy can be compared to a single SCF cycle, as found in the hybrid-DFT calculation. For this demonstration, SZP PAO orbitals have been used for hydrogen and oxygen with cutoff radii of 4.7 a.u. and 3.8 a.u., respectively. We emphasize that the main conclusions of this work can be easily extended to more flexible basis sets as long as the support functions are localized. The central processing unit (CPU) times used for the computation of EXX are reported in Fig. 6 as a function of the number of atoms using (i) the explicit evaluation of the full set of ERI, (ii) the CRI approach, and (iii) the CRI approach with partial storage of the PAO on the grids. Comparing the formal scalings obtained for




FIG. 6. Comparison of CPU times necessary to compute EXX in isolated water clusters as a function of number atoms (N) using the explicit ERI calculation and the CRI method. Ideal N^4 and N^3 scalings are given by plain lines.

the CRI methods against the full ERI approach, it becomes clear that the former reduces the quartic scaling to cubic with respect to the size of the system.

At this point, we should emphasize that exchange energy values obtained with the three schemes are fully identical, their accuracies being only dependent on the Poisson solver used to evaluate the pair potential in Eq. (13). Among the various numerical methods, one can choose to evaluate the Coulomb potential in reciprocal or real space. Whereas the former is the most appropriate for periodic neutral systems—when the positively charged nuclei compensate exactly the electronic charge density—it becomes less reliable for isolated and/or charged systems.⁶⁰ Several schemes have been developed to tackle this problem.^{61–64} Alternatives based on the discrete variable representation (DVR) of Eq. (13), which avoids the direct resolution of the Poisson equation, have been proposed.⁶⁵ The density is generally expanded in a direct product of one-dimensional localized real-space basis functions^{65–67} as, for instance, interpolating scaling functions (ISF). After extended comparisons between the DVR-ISF developed by Genovese *et al.*^{68,69} and corrected FFT-based schemes,^{70–72} we found that systematic convergence of the ERI is obtained with a better accuracy and at a lower cost using the real space Poisson solver.

As shown in Fig. 7, if a finite range R_X is introduced within the CRI algorithm, the CPU time can be significantly reduced, allowing linear scaling to be achieved for clusters with more than 36 atoms (with $R_X = 7.0$ au). Computational resources further decrease with shorter EXX ranges along with faster onset of the linear-scaling regime. The EXX accuracy with respect to the range of the exchange matrix is shown in Fig. 8 for the cluster $(\text{H}_2\text{O})_{20}$ presenting the “boxkite” structure. After somewhat erratic behavior at low values, it is found that an accuracy below 0.5 mhartree is reached for $R_X \geq 8$. Even though the non-local nature of the EXX interaction may need some special care when introducing a cutoff radius on X elements, it is reasonable to believe that the CRI implementation, along with a judicious choice of convergence parameters, is opening the way to exact exchange calculations on 100 000+ atoms with CONQUEST for a fair efficiency/accuracy ratio.




FIG. 7. Variation of the CPU time with respect to the range R_X (in a.u.) for the calculation of EXX in isolated water clusters using the CRI method.




FIG. 8. Convergence of the EXX energy with respect to the exchange range R_X for the cluster $(\text{H}_2\text{O})_{20}$. Error is given with respect to the exact calculation.

III. METHODS: MOVING ATOMS

A. Forces and stresses

Forces have been available in CONQUEST for some time, as described elsewhere,^{73,74} with the force being the exact differential of the energy, including Pulay forces where appropriate.

Calculation of the stress tensor has been recently implemented within the current release of CONQUEST. The definition of the stress tensor is standard,²

$$\sigma_{\alpha\beta} = \frac{\partial E}{\partial \epsilon_{\alpha\beta}} = \frac{\partial E}{\partial r_\alpha} r_\beta, \quad (16)$$

where α and β are Cartesian directions indices, and the second equality holds for *most* contributions to the stress. In this case, the first term is the force, so most contributions to the stress tensor can be calculated at the same time the forces are calculated. There are a few exceptions to this, but they are easily evaluated.^{2,75}

The original formulation of stress within DFT is traced back to the pioneering work of Nielsen and Martin,^{76,77} where a formulation for the stress was expressed for the first time in the framework

of the local-density approximation (LDA) and later derived in more detail.^{78,79} We have chosen to omit the factor of $\frac{1}{\Omega}$ in Eq. (16) since it *averages* the total stress over the macroscopic simulation cell and in a case where the volume Ω is not well defined would give spurious results. Note that *pressure*, as calculated at present, uses the volume of the simulation cell for the purpose of conversion, and if there is vacuum in any direction, the pressure should not be considered accurate. For this reason, CONQUEST internally uses values of stress to optimize simulation cells.

Stress is an extremely useful quantity: it is used to optimize simulation cell parameters, although this requires care to converge both the integration grid spacing and the numbers of k-points. Additionally, it is used in the NPT ensemble for molecular dynamics.

Our implementation of stress is valid for both exact diagonalization and linear scaling solvers. However, we have found that the stress converges extremely slowly with respect to density matrix truncation. Figure 9 shows the convergence of force (i.e., energy differences), total energy, and stress with density matrix truncation for three different elemental semiconductors with very different gaps: carbon, silicon, and germanium. Calculations were performed on the diamond structure (with a small perturbation in the case of the force calculation) at the optimal lattice parameter found using exact diagonalization, with an integration grid spacing of 0.1 bohr radii and an $8 \times 8 \times 8$ Γ -centered Monkhorst-Pack grid. To aid comparison between exact diagonalization and linear scaling calculations, we used the simplest basis set, i.e., single zeta, although this does not change the final results significantly. The plots show the difference between the $\mathcal{O}(N)$ and full diagonalization results. The full diagonalization results for the stresses were all less than 0.001 hartree (less than 0.1 GPa when converted to a pressure). For the forces, the full diagonalization results were 0.036 hartree/bohr for C, 0.016 hartree/bohr for Si, and 0.014 hartree/bohr for Ge. The total energies were -47.891 hartree for C, -33.611 hartree for Si, and -39.589 hartree for Ge.

The spatial decay of the density matrix is not analytically described for complex materials but can be shown to decay approximately exponentially with gap,^{80–82}

$$\rho(\mathbf{r}, \mathbf{r}') \propto \exp(-\gamma |\mathbf{r} - \mathbf{r}'|). \quad (17)$$

We can see in Fig. 9 that the rates of convergence of the different materials with density matrix truncation decrease with the decrease




FIG. 9. Plots of difference between the exact diagonalization result and $\mathcal{O}(N)$ result for (a) force, (b) total energy, and (c) stress (not normalized by the simulation cell volume) for carbon (cross symbols), silicon (plus symbols), and germanium (diamond symbols).

in gap size, as expected. It is notable that the initial errors are largest for the stresses and that significant differences in the stress remain even at very large density matrix ranges. We will investigate this fully in a future publication, but we are confident that this comes from the implicit dependence of energy on the density matrix truncation range, which should be included in a stress calculation as it will change as the unit cell is changed; however, an analytic form for this stress is not available.

B. Structure optimization

Structural optimization can be performed using a variety of standard approaches: the L-BFGS algorithm for atomic optimization, conjugate gradients for atomic and simulation cell optimization, and quenched molecular dynamics (both in a simple form and using the FIRE algorithm⁸³). We note that some form of preconditioning will become increasingly important as system sizes increase, and we are planning to implement some recently proposed preconditioners.^{84,85}

C. Molecular dynamics

Since the calculated forces are accurate and we can treat large systems, it is reasonable to expect that we can perform reliable molecular dynamics of large complex systems using CONQUEST. Unfortunately, it is not so easy to realize reliable MD simulations with the linear-scaling DFT technique or with MSSF. We have two key issues here. First, the calculation time for each MD step should be small enough to reach a meaningful simulation time. Second, density matrix should be sufficiently accurate to produce reliable MD simulations. During structure optimization, we can refine the accuracy step by step, without significant penalty. In many cases, we need a rather high accuracy only in the later stages of structure optimization. On the other hand, for MD simulations, we need to calculate the density matrix accurately at every step to ensure that the correct trajectory is followed. The accuracy of the density matrix depends on the tolerance to which it is optimized. Here, the optimized quantities are the auxiliary density matrix L in the linear-scaling calculations and PAO coefficients of the support functions in the MSSF method. Hereafter, we focus on the linear-scaling calculations.

For efficiency, we need a good initial guess of the L matrix at each MD step, and the simplest way, which should be efficient, is to use the L matrix optimized at the previous step. However, as is well known, this breaks the time-reversibility of the dynamics, resulting in a “drift” in the constant of motion over time.⁸⁶ Figure 10 shows the Born–Oppenheimer total energy (E_{BO}), defined as the sum of the ionic kinetic energy T and the DFT total energy V_{BO} , for linear-scaling MD simulations of a 64-atom silicon crystalline system with different tolerances on the optimization of the L matrix. The simulations are performed with the velocity-Verlet integrator with a time step of 0.5 fs in a microcanonical (NVE) ensemble with initial velocities set so that the system temperature is 300 K. The symbols in the figure show the time evolution of E_{BO} , which should be constant in reliable NVE-MD simulations. The results show that we need a very strict tolerance for stable MD simulations. Note that, if we use McWeeny initialization at every MD step (shown by the solid line in Fig. 10), E_{BO} is almost constant even if we use a rough




FIG. 10. Time evolution of the Born–Oppenheimer total energy (E_{BO}) obtained by McWeeny initialization at every step (solid line) and by reusing the L matrix from the previous step for different tolerances (symbols). Symbols indicate tolerances of 1.6×10^{-5} (circles), 1.6×10^{-7} (squares), and 1.6×10^{-9} (triangles). Reprinted with permission from M. Arita, D. R. Bowler, and T. Miyazaki, *J. Chem. Theory Comput.* **10**, 5419 (2014). Copyright 2014 American Chemical Society.

tolerance. However, this leads to a high computational cost at each iteration.

To solve this problem, CONQUEST uses the XLBOMD method^{86–88} with the density matrix minimization (DMM) method. The extended Lagrangian used in CONQUEST is⁸⁹

$$\mathcal{L}^{\text{XBO}}(\mathbf{X}, \dot{\mathbf{X}}, \mathbf{R}, \dot{\mathbf{R}}) = \mathcal{L}^{\text{BO}}(\mathbf{R}, \dot{\mathbf{R}}) + \frac{1}{2}\mu \text{Tr}[\dot{\mathbf{X}}^2] - \frac{1}{2}\mu\omega^2 \text{Tr}[(\mathbf{L}\mathbf{S} - \mathbf{X})^2], \quad (18)$$

where \mathbf{S} is the overlap matrix and \mathbf{X} is a sparse matrix associated with $\mathbf{L}\mathbf{S}$ rather than \mathbf{L} to maintain the orthogonal metric. μ is the fictitious electronic mass and ω is the curvature of the electronic harmonic potential. If we take the limit $\mu \rightarrow 0$, \mathcal{L}^{XBO} becomes \mathcal{L}^{BO} and we have equations of motion for nuclear positions and \mathbf{X} , and for \mathbf{X} ,

$$\ddot{\mathbf{X}} = \omega^2(\mathbf{L}\mathbf{S} - \mathbf{X}). \quad (19)$$

If we apply the Verlet scheme to calculate \mathbf{X} , we have

$$\mathbf{X}(t + \delta t) = 2\mathbf{X}(t) - \mathbf{X}(t - \delta t) + \delta t^2 \omega^2[\mathbf{L}(t)\mathbf{S}(t) - \mathbf{X}(t)],$$

i.e., the trajectory of $\mathbf{X}(t)$ is time-reversible and evolves in a harmonic potential centered on the ground-state density $\mathbf{L}(t)\mathbf{S}(t)$. The matrix $\mathbf{X}\mathbf{S}^{-1}$ is then used as the initial guess for the L -matrix.

If we use this method, the total energy E_{BO} is stable and the MD trajectories do not strongly depend on the tolerance or the range R_L in the $\mathcal{O}(N)$ calculations.⁸⁹ Figure 11 shows the variation of the total energy E_{BO} with the simulation time for different values of R_L . The fluctuations in the energy are smaller for larger

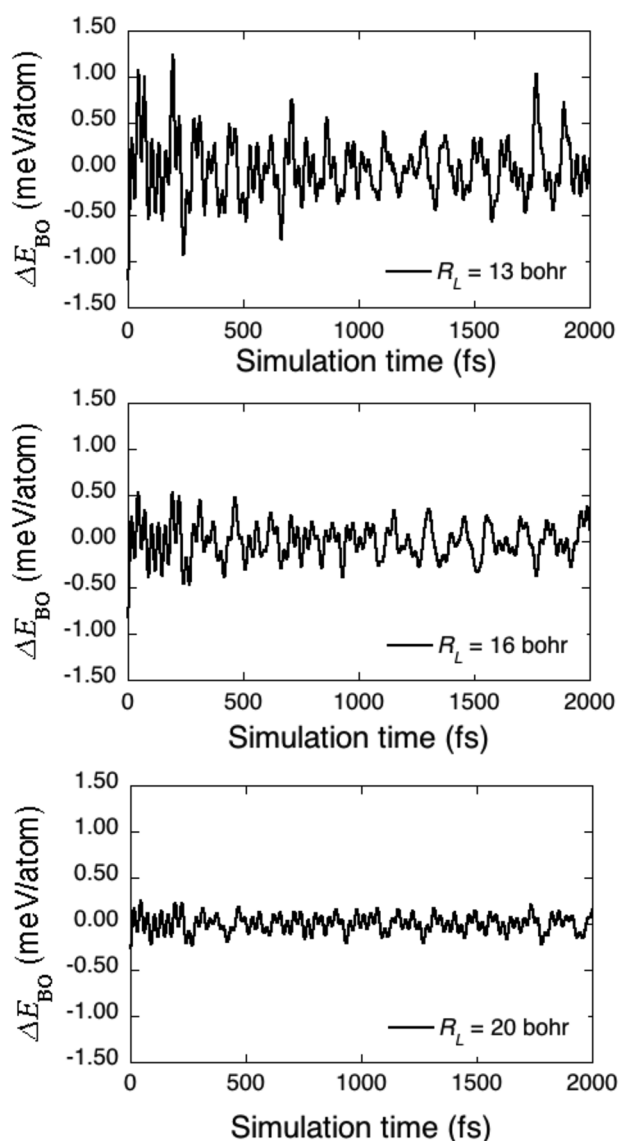


FIG. 11. Variation of E_{BO} during the NVE MD simulation of crystalline Si, with $R_L = 13$ bohrs (top), 16 bohrs (middle), and 20 bohrs (bottom). Reprinted with permission from M. Arita, D. R. Bowler, and T. Miyazaki, *J. Chem. Theory Comput.* **10**, 5419 (2014). Copyright 2014 American Chemical Society.

R_L , but even with $R_L = 13$ bohrs, the energy drift in E_{BO} is very small, meaning that the MD simulation is stable.

In practice, the \mathbf{X} -matrix sometimes moves away from the harmonic center over time, increasing the number of SCF iterations required to reach the ground state over the course of a simulation. To remove this instability, the dissipative term, $a \sum_{m=0}^M c_m \mathbf{X}(t - m\delta t)$, is included.⁹⁰ In principle, this dissipation term may break the time-reversible symmetry, but it is made to have a minimal effect and it is found that the MD simulations with the term is stable.

Using this XLBOMD + DMM method, we can also treat the canonical ensemble and perform constant temperature (NVT) MD

simulations, for example, using the Nosé–Hoover chain (NHC) method.⁹¹ The detailed explanation of the integration scheme used for the canonical ensemble is provided in Subsection 2 a of the Appendix. Figure 12 shows the time evolution of the temperature, the constant of motion for the NHC method, and the DFT

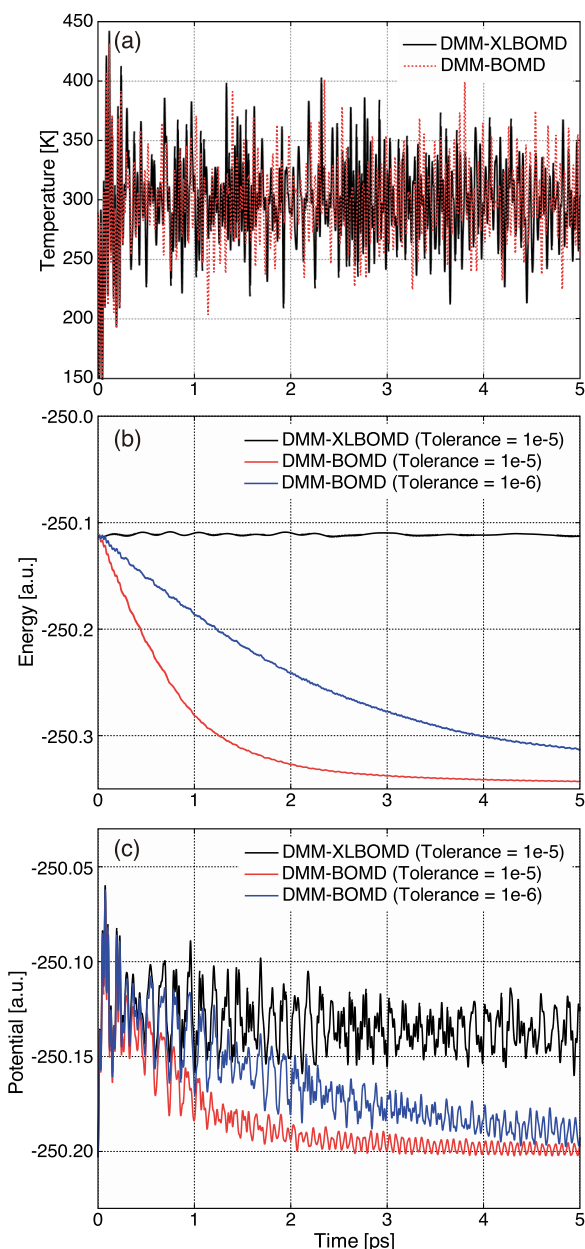


FIG. 12. Time evolution of the MD simulations of 64-atom Si crystalline systems in the canonical ensemble for (a) temperature, (b) constant of motion for the NHC method, and (c) the potential energy (V_{BO}). The results with the XLBOMD method and those without using the XLBOMD method, referred to as DMM-BOMD, are compared. Reproduced with permission from Hirakawa *et al.*, *J. Phys.: Condens. Matter* **29**, 405901 (2017). Copyright 2017 IOP Publishing.

potential energy V_{BO} in NVT-MD simulations of the same crystalline 64-atom silicon system at 300 K with and without the XLBOMD method. We find that temperatures are stable and close to 300 K in both simulations. However, we again observe the problem of drift in the constant of motion when we do not use the XLBOMD method, while there are no such problems in the XLBOMD + DMM simulations. More importantly, the profile of V_{BO} is completely different between the usual DMM and XLBOMD + DMM MD simulations. CONQUEST can also perform NVT simulations using the SVR (stochastic velocity rescaling) thermostat,⁹² which is extremely efficient and provides excellent conservation of the constant of motion, as described in Subsection 2 d of the Appendix.

Since the stress tensors can be also calculated using CONQUEST with the DMM method, as shown in Sec. III A, it is also possible to include the degrees of freedom of the unit cell for NPT simulations with a given pressure using the Parrinello–Rahman equations of motion.⁹³ CONQUEST uses the Martyna–Tobias–Tuckerman–Klein modification,⁹⁴ coupling the constant pressure equations of motion to a Nosé–Hoover chain thermostat to recover the NPT ensemble. The integration scheme used in the NPT ensemble is also explained in Subsection 2 b of the Appendix.

This scheme is tested on a bulk crystalline silicon system containing 1000 atoms and the $\mathcal{O}(N)$ method for finding the electronic ground state, as shown in Fig. 13. A minimal basis set (SZ) was employed together with a grid cutoff of 100 hartree and the PBE (Perdew, Burke, Ernzerhof) exchange-correlation functional. The extended-Lagrangian scheme described above was used, with a velocity Verlet integrator for the X matrix and fifth order dissipation. The system was equilibrated using a Berendsen-type weak coupling thermostat and barostat at a temperature of 300 K and a pressure of 0.1 GPa. The cell volume was allowed to vary but

constrained to be cubic. An integration time step of 0.5 fs was used, with a 5th-order Yoshida–Suzuki integration scheme and thermostat and barostat coupling time periods of 15 fs and 160 fs, respectively. An *ad hoc* drag was applied to the barostat, reducing the velocities of the cell and its Nosé–Hoover thermostats by 5% for each time step. This was found to improve the stability, preventing the amplification of “ringing” of the barostat, with a minimal impact on energy conservation.

It can be seen that in order to achieve good energy conservation, the L-tolerance should be lower than 10^{-5} , with a significant drift in the conserved quantity occurring at looser tolerances; without the XLBOMD scheme, the tolerance required would be much tighter. We note that the NPT integrator is considerably more sensitive to the time step due to coupling between the thermostat and barostat degrees of freedom and that in this case, a time step of 1.0 fs also resulted in a significant energy drift, although we are seeking to alleviate this sensitivity.

IV. PERFORMANCE

Here, we demonstrate the performance of CONQUEST showing two examples, one for the MSSFs with diagonalization and another for the $\mathcal{O}(N)$ calculations.

A. Performance of MSSF

A recent study on the graphene/Rh(111) interface¹⁰ showcases both the accuracy and efficiency of the MSSFs. This study used large basis sets of PAOs contracted to a minimal size using the MSSF formalism, i.e., 15 and 22 PAOs of rhodium and carbon atoms are contracted to 6 and 4 MSSFs, respectively. In Ref. 10, it was




FIG. 13. NPT molecular dynamics on the 1000 atom bulk silicon system. The three lines demonstrate the effect of varying the tolerance applied to the optimization of the energy with respect to the density matrix during the $\mathcal{O}(N)$ solution (“L-tolerance”).




FIG. 14. Density of states of m14 graphene/Rh(111) calculated with plane-waves (black), PAOs (green), and MSSFs (blue). The red line in the lower panel represents the difference between the DOS calculated with PAOs and MSSFs. Reproduced with permission from Romero-Muñoz *et al.*, J. Phys.: Condens. Matter **30**, 505901 (2018). Copyright IOP Publishing.

demonstrated that the PAOs and MSSFs show comparable accuracy with plane-waves for the electronic and atomic structures of graphene/Rh(111), as shown in Fig. 14. The accuracy of PAOs has been further investigated in Ref. 24.

This study then demonstrates the great reduction in computational effort by using MSSFs. Table III shows the computational times of a SCF step for the graphene/Rh(111) systems consisting of 1544 and 3088 atoms (shown in Fig. 15). MSSFs clearly require more computational time for matrix construction than the PAOs, which comes from the calculations of the linear combination coefficients, as explained in Sec. II. On the other hand, the time to diagonalize the electronic Hamiltonian is reduced significantly by using MSSFs because the diagonalization time scales cubically with the number of support functions. For the 1544-atom system, the total time, i.e., the summation time of matrix construction and diagonalization, is reduced by a factor of ≈ 3 from 1256.9 s to 439.6 s. For the 3088-atom




FIG. 15. Atomic structure of the graphene/Rh(111) system (3088 atoms).

systems, when using 108 processes, the total time is reduced by a factor of ≈ 18 from 37 803.5 s to 2156.3 s, which indicates that the use of MSSFs becomes more efficient as systems become larger. Comparing the time for the matrix construction for the 1544 atoms with 432 MPI processes and that for the 3088 atoms with 864 MPI processes, i.e., when both the system size and the number of processes are doubled, the times are very close to each other, which indicates that the construction of the MSSFs is $\mathcal{O}(N)$ and parallelized ideally.

B. Performance of $\mathcal{O}(N)$ calculations on massively parallel computers

The performance of CONQUEST on the Japanese Fujitsu-made K-computer is of real significance.⁹⁵ This computer once topped the TOP500 list⁹⁶ (June and November 2011) and 8 years later *still featured* on the list in 20th place (November 2019) due to its impressive peak performance of 11 280.4 TFLOPS from its 705 024 physical cores. CONQUEST was found to display almost ideal parallel efficiency, as shown in Fig. 16(c), utilizing up to 200 000 physical cores⁹⁵ on systems up to 2×10^6 atoms.⁶ (At present, there is no dynamic fault-tolerance built in to CONQUEST to account for failure of nodes during a run; however, the frequency with which restart files are written can be controlled at a fine-grained level, which makes recovery from a crash easy.) Using crystalline silicon systems as a benchmark, it was demonstrated that in the $\mathcal{O}(N)$ mode of operation that both strong scaling (the wall time for a fixed number of atoms, increasing the physical core count) and weak scaling (the wall time for a fixed number of atoms/physical core, increasing the number of atoms) perform very well. Specifically, for strong scaling, it is found that performance is good should the number of atoms/core be ≥ 4 , but for weak scaling, the performance is close to *perfect* for any given number of atoms per core all the way up to 2 000 000 atoms.⁶ Strong scaling has also been tested on the UK national supercomputer ARCHER, a Cray XC30 MPP system [Fig. 16(a)]. This also demonstrates the high efficiency of the code until about 5 atoms/core. Going to fewer atoms/core than this starts to significantly impact the performance of CONQUEST; for this particular test, more than 50 atoms/core was feasible, but for more stringent tests, it would require large amounts of memory. When testing the scalability of the $\mathcal{O}(N)$ algorithm itself [Fig. 16(b)], we see that we achieve near-perfect linear scaling with the system size even in the range of 2560–24 565 atoms.

V. APPLICATIONS

There are a multitude of physical systems to which CONQUEST has been applied. Studies using both exact diagonalization (with and without the use of MSSFs) and the $\mathcal{O}(N)$ mode of operation have

TABLE III. Computational time for the self-consistent-field calculation step with PAOs and MSSFs for graphene/Rh(111) performed on the supercomputer SGI ICE X in NIMS (data from Ref. 10).

	1544 atoms		3088 atoms		
	PAO	MSSF	PAO	MSSF	MSSF
No. of MPI process	432	432	108	108	864
Time (s)					
Matrix construction	64.3	400.4	155.7	1455.4	405.9
Diagonalization	1192.5	39.2	37 647.7	700.8	165.9
Total ^a	1256.9	439.6	37 803.5	2156.3	571.8

^aSummation of matrix construction and diagonalization.

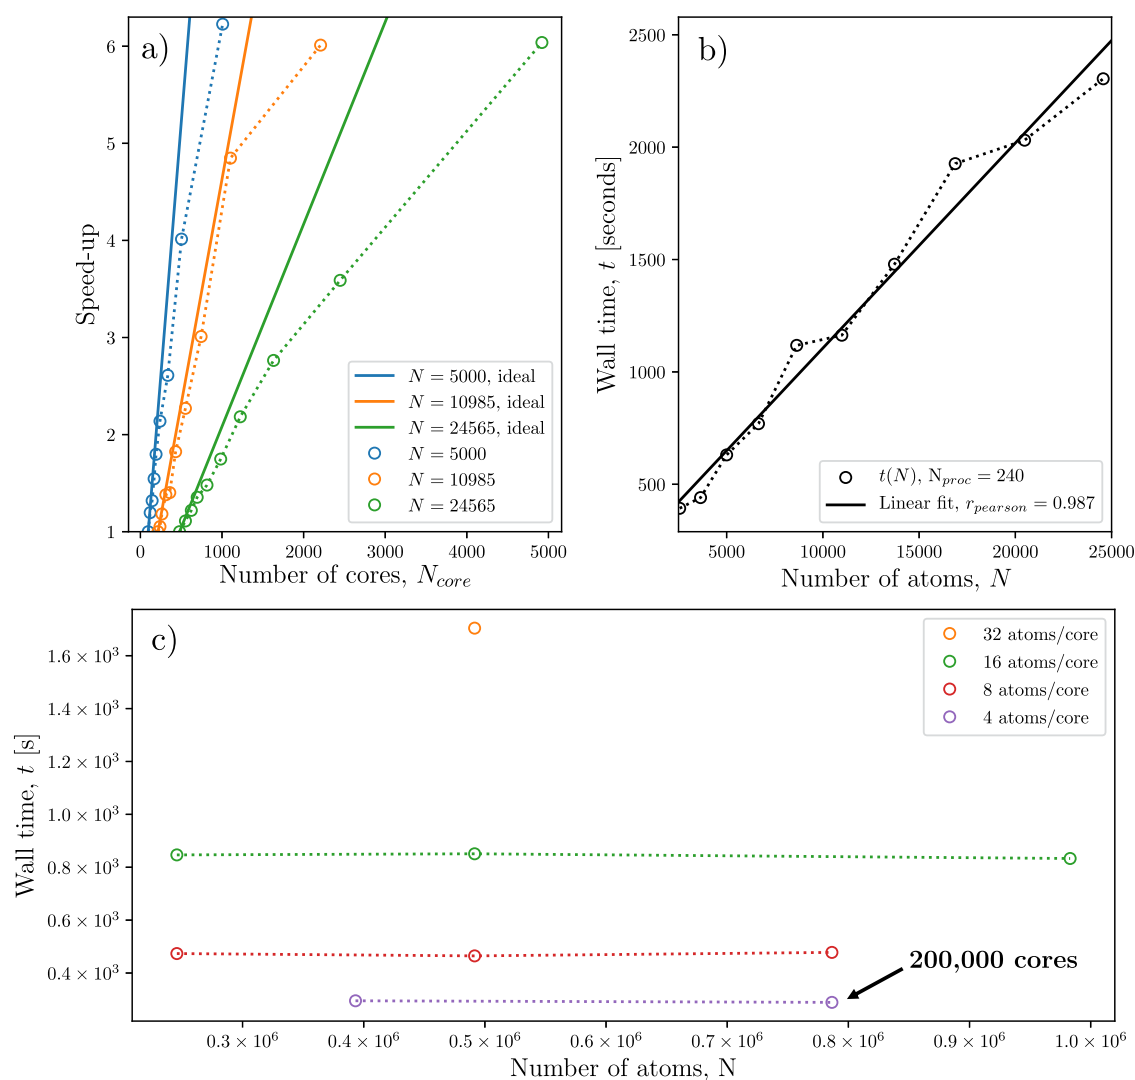


FIG. 16. Scaling performance for CONQUEST on the K-computer and ARCHER. (a) Strong scaling on the UK national supercomputer, ARCHER, up to 4920 processors (from 50 atoms/core up to 5 atoms/core). Calculations are performed on bulk PbTiO_3 with an L_{range} of $14a_0$ and a SZ basis. (b) Demonstration of the scaling of the $\mathcal{O}(N)$ algorithm on ARCHER for the same material system as (a). (c) Weak scaling on the K-computer up to 1 000 000 atoms for bulk Si.

all been exploited in large-scale structural relaxations and molecular dynamics. In the solid state, the code has been used to study the properties of nanowires,^{97,98} Ge hut clusters on Si (001) surfaces,⁹⁹ charge transport properties,¹⁰⁰ interfaces between graphene with metals,¹⁰ and ferroelectric domain morphologies in perovskite oxide heterostructures. The code has also been applied to complex biological systems including hydrated DNA^{101,102} and gramicidin-A.¹⁰³ It is the purpose of this section to outline some of these studies and to suggest areas that the code could find new applications.

A. Nanoscale Ge/Si systems

One of the most important targets for the large-scale DFT study is nano-structured semiconductors. Among them, Ge/Si systems

have many attractive properties as a candidate for next-generation devices. Heteroepitaxy and strained growth in Ge/Si systems can be used as important techniques to control the structures and to explore new favorable properties.

CONQUEST was first applied to study the stability of Ge three-dimensional islands on the Si substrate, called hut clusters, made of four equivalent Ge(105) facets. Experimentally, this 3D structure appears when the coverage of Ge atoms becomes large after the formation of a two-dimensional (2D) structure with defects.¹⁰⁴ Here, CONQUEST calculations were performed with LDA and non-self-consistent mode using a minimal basis set (SZ), whose accuracy was thoroughly investigated for Ge/Si systems.¹⁰⁵ The stability of the 3D structure in the heteroepitaxy systems is usually determined by the competition between the energy gain of the strain relief by the 3D

structure and the energy loss due to the increase in the surface area due to the formation of facets in the 3D structures. However, the Ge/Si(001) system has the unique property that the strained Ge(105) facet is more stable than the strained Ge(001) surface,^{106,107} even accounting for the increase in surface area. Thus, in order to clarify the stability of the 3D structure, it is necessary to include the effects of the edges between the facets and the finite area of the actual facets. For this, we need to treat the actual size of the hut clusters with a Si substrate. Standard DFT methods cannot treat the 3D structure with a size similar to experiments, but it is possible using CONQUEST with structure optimization. In the early study with CONQUEST, the total energies of systems having the same coverage of Ge atoms were compared between the 2D and 3D structures and it was found that the 3D structure becomes more stable when the coverage of Ge is larger than 2.7 monolayers.⁹⁹ This is close to the minimum coverage showing the transition from 2D to 3D growth in experiments, supporting the high accuracy of the present DFT method.

Further studies considered the stability of a single Ge dimer adsorbed at various sites on the facets.¹⁰⁸ This study aimed to clarify the initial process during the formation of a new facet layer. Experimentally, it has been reported that elongated hut clusters tend to grow, under certain growth conditions, by increasing the length of the longer side while keeping the width (shorter side) unchanged [see Fig. 17(a)¹⁰⁹]. The detailed mechanism underlying the growth of new facet layers is extremely difficult to obtain from experiments, since the complete facet is formed rapidly. We expect large-scale DFT calculations to play a significant role in clarifying these processes. By performing structure optimization for more than 100 different sites for the adsorption site of a single Ge dimer, as shown in Fig. 17(c), it was suggested that the top or the edges of the facets are the most preferable sites and higher positions are more stable than lower ones. This kind of study is now possible with CONQUEST using a parallel supercomputer. The largest system in this study contains about 200 000 atoms, whose structure is shown in Fig. 17(b).

Together with the study of double and triple dimer adsorptions, it was concluded that the new layer of the facet is very likely to grow from top to bottom.

Recently, CONQUEST was also applied to study Si/Ge and Ge/Si core-shell nanowires using the SZP basis set with self-consistency. Semiconductor nanowires are promising candidates for the next-generation vertical-type transistors¹¹¹ and have been extensively studied both experimentally and theoretically. The core-shell type nanowires have many interesting and attractive properties^{112–114} for next-generation electronics. All of these properties, however, will depend strongly on the size of the core and shell. Using CONQUEST with the $O(N)$ method, strain distributions were calculated for nanowires with different sizes, shown in Fig. 18(a) up to experimentally accessible sizes. These are hexagonal Si/Ge core-shell nanowires along the $\langle 110 \rangle$ direction and with numbers of atoms ranging from 612 to 2404. The strain distributions in the core region of these nanowires are shown in Fig. 18(b). We can see that the strain is distributed anisotropically, depending on the direction of the bonds, and that large variations of strains exist in the interface and surface regions.

The structure of a more circular Si/Ge core-shell nanowire was also investigated and its band structure was calculated with the Sakurai-Sugiura (SS) method explained in Sec. II D using the optimized structure and the self-consistent charge density obtained by $O(N)$ calculations. The occupied eigenstates near the Fermi level were also calculated and are shown in Fig. 18(c). We can clearly see that the distribution is anisotropic and localized in the Ge-shell region. The effect of arsenic doping and its dependence on the doping sites in the Si nanowires were also recently reported⁹⁸ using a rather high quality basis set (TZTP), with the MSSF method.

B. PbTiO₃ films on SrTiO₃ substrates

Studies of the perovskite oxides can also make good use of large-scale electronic structure calculations. CONQUEST can be used to




FIG. 17. Linear-scaling DFT study of the Ge 3D structure on the Si substrate using CONQUEST. (a) Experimental observation showing that Ge hut clusters grow, under certain conditions, by increasing the length of the longer side while keeping the width (shorter side) unchanged. (b) The optimized structure of the largest structural model for the Ge hut cluster on the Si substrate, which contains about 200 000 atoms. (c) Adsorption energy map of single Ge dimers adsorbed on the {105} facets. Disks show the position of a given dimer on a facet projected on the x-y plane. Dimers are labeled according to their height. Adsorption energy of dimers increases from red to green. top: energy values of single Ge dimers on small (circles) and large (rectangles) facets. (a) Reprinted with permission from M. R. McKay, J. A. Venables, and J. Drucker, *Phys. Rev. Lett.* **101**, 216104 (2008). Copyright 2008 American Physical Society. [(b) and (c)] Reproduced with permission from T. Miyazaki, *ECS Trans.* **86**, 269–279 (2018). Copyright 2018 IOP Publishing.¹¹⁰

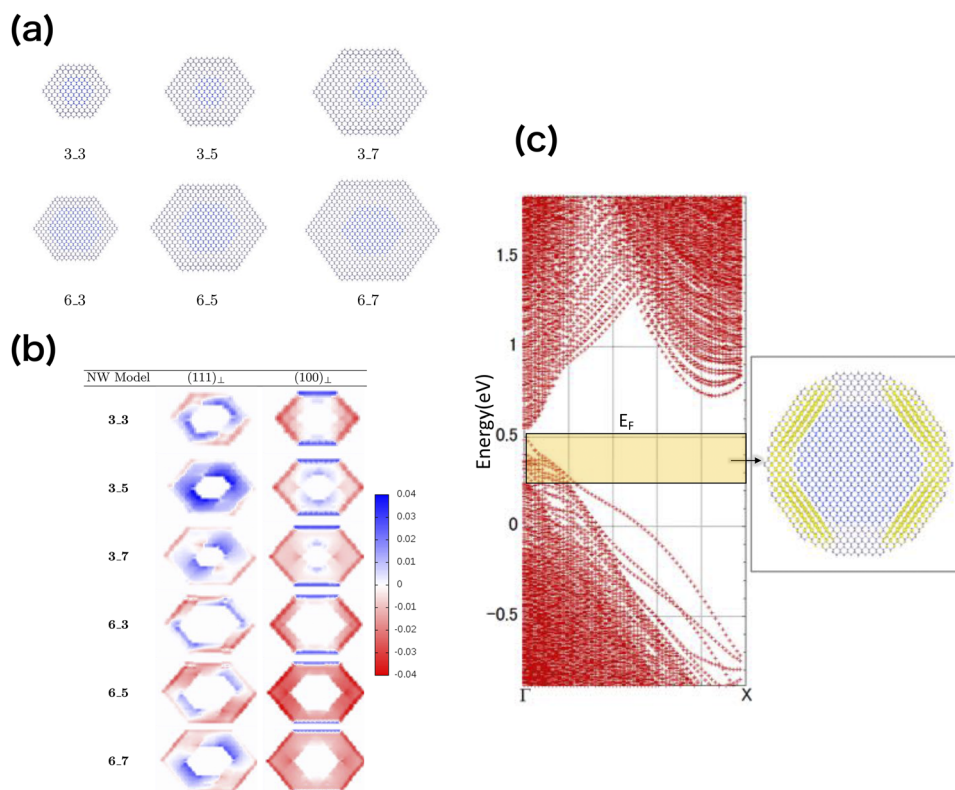


FIG. 18. (a) Structural models used in the study of Si/Ge nanowires along the $\langle 110 \rangle$ direction labeled C_S, where the index C represents the number of layers in the core and S represents the surface. Shell thickness increases from left to right and core thickness from top to bottom. (b) Average bond strain map for the cross section of the Ge shell of the SiGe-NWs. Maps for the bonds along different directions are shown, with extension illustrated in blue and compression in red. (c) Band structure of the circular Si/Ge core-shell nanowire and the charge density constructed from the occupied orbitals near the Fermi level (in the range shown in the band structure). (a) and (b) are reproduced with permission from O'Rourke *et al.*, J. Phys.: Condens. Matter **30**, 465303 (2018). Copyright 2018 IOP Publishing. (c) is reproduced with permission from T. Miyazaki, ECS Trans. **86**, 269–279 (2018). Copyright 2018 IOP Publishing.

study large supercells of technologically relevant piezoelectric alloys such as $\text{PbZr}_x\text{Ti}_{1-x}\text{O}_3$, where approximations designed to circumvent the need for large supercell calculations (such as the virtual crystal approximation) are unable to quantify local structural distortions.¹¹⁵ The study of ferroelectric domains in thin films is another problem requiring large-scale electronic structure calculations and

accurate structural relaxations. Using the MSSF method and a large basis set of PAOs (DZDP), the nature of the ferroelectric flux closure domains in thin PbTiO_3 films on SrTiO_3 substrates was revealed. Using the initial geometry displayed in Fig. 19(a), we were able to relax the system to a stringent 0.01 eV/Å force tolerance using quenched molecular dynamics. The force reduction for the first 50

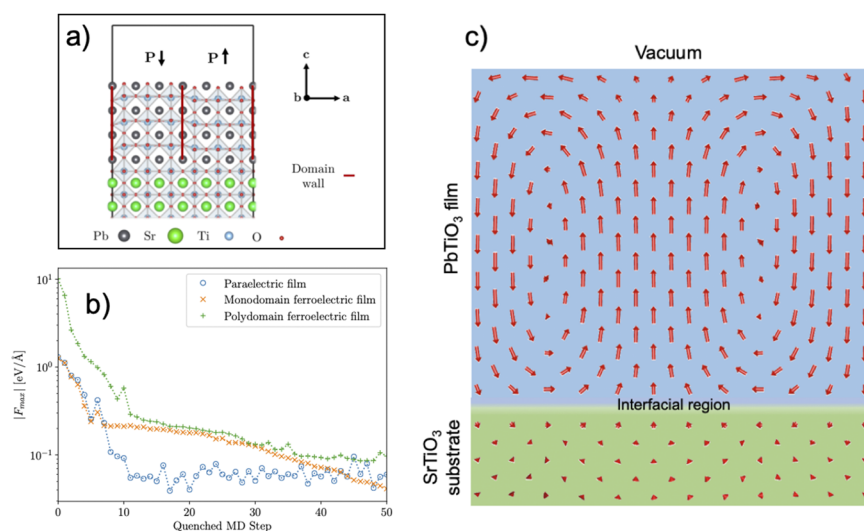


FIG. 19. The results of structural relaxation calculations with CONQUEST. (a) The initial geometry used to study ferroelectric flux closure domains in PbTiO_3 films on SrTiO_3 substrates. This example shows a three unit cell deep film with a domain period of six unit cells. (b) The evolution the magnitude of the maximum force on any atom for the first 50 quenched molecular dynamics steps for three different film configurations. (c) The local polarization vector field of a nine unit cell deep PbTiO_3 film (2088 atoms) on a SrTiO_3 substrate.




FIG. 20. (a) Structure model for the ion channel gramicidin A embedded in DMPC lipid bilayers sandwiched with bulk water regions. (b) Left: comparison of 1MAG (in blue) and 1JNO (in red) experimental structures. Middle and right: optimized (red/light) and initial (black/dark) structures of the isolated gA molecule starting from 1MAG (middle) or 1JNO (right) models. (c) Change of the residual during the density matrix minimization (DMM) step in the $\mathcal{O}(N)$ calculations of the gA system shown in (a). In this calculation, the search direction in the minimization was reset at every 40 iterations. (a) and (b) are reproduced with permission from Todorović *et al.*, J. R. Soc. Interface **10**, 20130547 (2013). Copyright Royal Society (UK).

steps is shown in Fig. 19(b). Figure 19(c) shows the local polarization vector field of a nine unit cell deep film. Such a field is calculated using the relaxed structure, the deviation in displacement from high symmetry sites, and the Born effective charge tensors.¹¹⁶

C. Biological systems

Complex biological systems are one of the most important targets for large scale DFT simulations.¹⁴ CONQUEST has already been applied to several biological systems, such as hydrated DNA,^{101,102} dihydrofolate reductase (DHFR),¹¹⁷ and the gramicidin A (gA) ion channel¹⁰³ systems. In the study of the gA system, the optimized structure of the isolated gA molecule, shown in Fig. 20(b), was first calculated for the two previously reported structural models, 1MAG and 1JNO. The electronic structure of the gA molecule was also

analyzed, and it was concluded that the side chains of gA do not affect the electrostatic potential in the pore of gA. This kind of study of the isolated gA molecule cannot explain the selectivity of the ion permeation in the gA system, and it suggests the importance of simulating the system in the channel environment. We should treat the gA molecule in lipid bilayers sandwiched by bulk water regions, as shown in Fig. 20(a). Using CONQUEST, we can perform stable self-consistent DFT calculations of such a complex system made of 17 102 atoms, having a rather irregular charge distribution. Figure 20(c) shows that the density matrix minimization, for a SZP basis set, of this system is robust. It is important to note that, for stability in the self-consistency process, we need to update the charge density as well as the density matrix at each step in the calculation. More detailed information about this large-scale DFT study on the gA system will be reported in the future.




FIG. 21. Snapshot structures for (a) Ge/Si core-shell and (b) Si/Ge core-shell nanowires at 3000 K and (c) melting SiO₂ at 3000 K and 10 GPa. (d) Hydrated DNA system (details of the simulation are found in Ref. 102).

D. Large-scale MD simulations with CONQUEST

In Sec. V A, we introduced energetically stable structures of the perfect epitaxial models for Si/Ge core-shell nanowires. However, it is also important to investigate defects or Si-Ge intermixing at the interface for the actual nanowires. In addition, we sometimes need to clarify the thermodynamic stability or the dynamical processes. In such cases, molecular dynamics simulations based on DFT (DFT-MD) are useful and important. Using the XLBOMD + DMM method, explained in Sec. III C, we are now able to do practical and reliable self-consistent DFT-MD simulations of very large systems. The MD simulations in this section used SZ (NWs, SiO₂) and SZP (hydrated DNA) basis sets.

For perfect, epitaxial Si/Ge core-shell nanowires, DFT-MD simulations of nanowires containing 4788 atoms, whose diameter is 10.4 nm (Si core is 7.2 nm and the thickness of the Ge shell part is 1.6 nm), at 900 K were recently performed. The DFT-MD simulations confirmed that the structure is stable at least up to 10 ps. This does not guarantee that the perfect epitaxial model is more stable than other structures containing defects or intermixing, but it indicates that the model is at least a meta-stable structure. We also performed DFT-MD simulations of Si/Ge and Ge/Si core-shell nanowires at 3000 K, whose snapshot structures are shown in Figs. 21(a) and 21(b). We observed that the Ge region melted first in both cases. As linear-scaling DFT-MD simulations on such large systems are now practical, we expect that they can be used to explore possible structures of various types of defects or intermixing effects at the Si/Ge interfaces by a local heating technique. Such study is now in progress. Furthermore, as we explained in Sec. III C, we can now perform DFT-MD simulations at a constant high temperature and a given high pressure. Structural properties of melting SiO₂ [Fig. 21(c)] are now being investigated using CONQUEST.

Of course, complex biomolecules, such as DNA in water, are also an important target for large-scale DFT-MD studies using CONQUEST, with a snapshot shown in Fig. 21(d). It is noteworthy that free energy calculations based on the blue moon ensemble method are now available with CONQUEST.¹¹⁸ We expect that a variety of dynamical processes or enzyme reactions in biological systems will be studied with CONQUEST in the future.

VI. CONCLUSIONS

We have summarized the principles behind the implementation of the CONQUEST code, which enables it to address large scale DFT simulations up to around 10 000 atoms with exact diagonalization, and significantly larger systems, at least up to millions of atoms using linear scaling DFT. We showed how support functions can be represented in three ways, leading to a powerful approach for representing the density matrix. We also gave details on approaches to find the electronic structure of large systems, even with linear scaling, and indicated how hybrid DFT methods can be extended to extremely large systems.

We gave details of atomic movement, particularly molecular dynamics, and how the implementation and performance are affected by the use of linear scaling methods. We demonstrated that accurate, linear scaling MD is feasible with reasonable computational time for standard ensembles (NVE, NVT, and NPT, although care is needed with the calculation of stress and linear scaling). We

then showed the performance of the MSSF approach, and how it opens up the possibility of exact diagonalization simulations with many thousands of atoms. We also investigated the parallel performance of the code, both with MSSF and linear scaling, finding excellent performance including perfect scaling for certain approaches. We ended by giving examples of applications of the code on systems with sizes ranging from hundreds of atoms to hundreds of thousands of atoms.

While large scale DFT calculations are challenging, in terms of the preparation of the system, the computing resources required, and the analysis of large data sets, it is clear that they are also now feasible for the majority of users. It is expected that the size of most DFT calculations will grow from a few hundred atoms to many thousands, enabling greater accuracy, and new systems will be addressed.

ACKNOWLEDGMENTS

The authors would like to express their gratitude to all the developers and beta testers of CONQUEST who have contributed to the code over the years. In particular, Professor Mike Gillan initiated and oversaw the development and implementation for many years, and his input was invaluable. We also acknowledge, in alphabetical order, Michiaki Arita, Veronika Brazdova, Marius Buerkle, Rathin Choudhury, Chris Goringe, Eduardo Hernandez, Teruo Hirakawa, Chathurangi Kumarasinghe, Conn O'Rourke, Takao Otsuka, Alex Sena, Kane Shenton, Teppei Suzuki, Umberto Terranova, Milica Todorovic, Lianheng Tong, and Antonio Torralba.

This work was supported by the World Premier International Research Centre Initiative (WPI Initiative) on Materials Nanoarchitectonics (MANA), "Exploratory Challenge on Post-K computer" by MEXT, and JSPS Grant-in-Aid for Scientific Research (Grant Nos. 18H01143, 17H05224, and 15H01052). L.T. was supported by the BBSRC (Grant No. BB/H024217/1), "Linear Scaling Density Functional Theory for Biochemistry."

Calculations were performed on the Numerical Materials Simulator at NIMS, the supercomputer HA8000 system at Kyushu University, and by using the computational resources of the K computer provided by the RIKEN Advanced Institute for Computational Science through the HPCI System Research project (Project Nos. hp160129, hp170264, hp180175, hp180226, and hp190096).

The authors are grateful for computational support from the UK Materials and Molecular Modeling Hub, which is partially funded by EPSRC (Grant No. EP/P020194), for which access was obtained via the UKCP consortium and funded by EPSRC (Grant Ref. No. EP/P022561/1). We also acknowledge computational support from the UK national high-performance computing service, ARCHER, for which access was obtained via the UKCP consortium and funded by EPSRC (Grant Ref. No. EP/K013564/1).

The data that support the findings of this study are available from the corresponding authors upon reasonable request.

APPENDIX: FURTHER DETAILS OF MD IMPLEMENTATION

Here, we explain the details of the integration scheme used in the molecular dynamics, since it is important for the actual implementation, and those related to the stability of the molecular dynamics.

1. Microcanonical ensemble

The microcanonical ensemble is generated simply by solving Hamilton's equation of motion for the Hamiltonian

$$\mathcal{H} = \sum_{i=1}^N \frac{\mathbf{p}_i^2}{2m_i} + U(\mathbf{r}_i), \quad (\text{A1})$$

resulting in the following equations of motion:

$$\dot{\mathbf{r}}_i = \frac{\mathbf{p}_i}{m_i}, \quad (\text{A2})$$

$$\dot{\mathbf{p}}_i = -\frac{\partial U(\mathbf{r}_i)}{\partial \mathbf{r}_i} = \mathbf{F}_i. \quad (\text{A3})$$

These equations are integrated using the velocity Verlet algorithm.

2. Non-Hamiltonian molecular dynamics

Hamiltonian dynamics describe systems that are isolated from their surroundings, and in order to generate the canonical and isobaric-isothermal ensembles, the system must be coupled to an external bath (heat for the former and heat and stress in the case of the latter). In the extended system approach, a set of *non-Hamiltonian* equations of motion including degrees of freedom for a thermostat and/or barostat are posited and shown to generate the correct statistical ensemble *post hoc*.

a. Canonical (NVT) ensemble

The Nosé–Hoover Hamiltonian^{119,120} for the canonical ensemble can be written as

$$\mathcal{H} = \sum_i \frac{1}{2} m_i s^2 \dot{\mathbf{r}}_i^2 + U(\mathbf{r}_i) + \frac{1}{2} Q \dot{s}^2 - (n_f + 1) k_B T \ln s, \quad (\text{A4})$$

where \mathbf{r}_i and $\dot{\mathbf{r}}_i$ are the position and velocity of particle i , respectively, U is the potential energy (in this case, the DFT total energy), s is a dimensionless quantity that can be interpreted *post hoc* as a time step scaling factor, Q is the fictitious mass of the heat bath, and n_f is the number of ionic degrees of freedom. Hamilton's equations of motion can then be solved to generate the Nosé–Hoover equations of motion. However, Martyna *et al.* demonstrated that this method does not generate an ergodic trajectory and proposed an alternative formulation with a chain of M coupled heat thermostats of mass Q_k , each with “position” η_k and conjugate momentum p_{η_k} ,¹²¹ resulting in the following equations of motion:

$$\dot{\mathbf{r}}_i = \frac{\mathbf{p}_i}{m_i}, \quad (\text{A5})$$

$$\dot{\mathbf{p}}_i = -\frac{\partial U(\mathbf{r})}{\partial \mathbf{r}_i} - \frac{p_{\eta_1}}{Q_1} \mathbf{p}_i, \quad (\text{A6})$$

$$\dot{\eta}_k = \frac{p_{\eta_k}}{Q_k}, \quad (\text{A7})$$

$$\dot{p}_{\eta_1} = \left(\sum_{i=1}^N \frac{\mathbf{p}_i^2}{m_i} - n_f k_B T \right) - \frac{p_{\eta_2}}{Q_{\eta_2}} p_{\eta_1}, \quad (\text{A8})$$

$$\dot{p}_{\eta_k} = \left(\frac{p_{\eta_{k-1}}^2}{Q_{k-1}} - k_B T \right) - \frac{p_{\eta_{k+1}}}{Q_{k+1}} p_{\eta_k}, \quad (\text{A9})$$

$$\dot{p}_{\eta_M} = \left(\frac{p_{\eta_{M-1}}^2}{Q_{M-1}} - k_B T \right). \quad (\text{A10})$$

These equations are integrated by constructing an appropriate Liouvillian and translated into an algorithm via the Trotter–Suzuki expansion, as described by Hirakawa *et al.*⁹¹

b. Isobaric-isothermal (NPT) ensemble

The Parrinello–Rahman equations of motion⁹³ extend the constant volume equations of motion to include the degrees of freedom of the unit cell via the extended system approach. CONQUEST uses the Martyna–Tobias–Tuckerman–Klein modification,⁹⁴ coupling the constant pressure equations of motion to a Nosé–Hoover chain thermostat to recover the NPT ensemble. For a cell unconstrained unit cell, the equation of motion are

$$\dot{\mathbf{r}}_i = \frac{\mathbf{p}_i}{m_i} + \frac{\mathbf{p}_g}{W_g} \mathbf{r}_i, \quad (\text{A11})$$

$$\dot{\mathbf{p}}_i = \mathbf{F}_i - \frac{\mathbf{p}_g}{W_g} \mathbf{p}_i - \left(\frac{1}{N_f} \right) \frac{\text{Tr}[\mathbf{p}_g]}{W_g} \mathbf{p}_i - \frac{p_\xi}{Q} \mathbf{p}_i, \quad (\text{A12})$$

$$\dot{\mathbf{h}} = \frac{\mathbf{p}_g \mathbf{h}}{W_g}, \quad (\text{A13})$$

$$\dot{\mathbf{p}}_g = V(\mathbf{P}_{\text{int}} - \mathbf{I} P_{\text{ext}}) + \left[\frac{1}{N_f} \sum_{i=1}^N \frac{\mathbf{p}_i^2}{m_i} \right] \mathbf{I} - \frac{p_\xi}{Q} \mathbf{p}_g, \quad (\text{A14})$$

$$\dot{\xi} = \frac{p_\xi}{Q}, \quad (\text{A15})$$

$$\dot{p}_g = \sum_{i=1}^N \frac{\mathbf{p}_i^2}{m_i} + \frac{1}{W_g} \text{Tr}[\mathbf{p}_g^T \mathbf{p}_g] - (N_f + d^2) kT, \quad (\text{A16})$$

where \mathbf{r}_i , \mathbf{p}_i , and m_i are the position, momentum, and mass of particle i , respectively, ξ , p_ξ , and Q are the position, momentum, and mass of the thermostat, respectively, and \mathbf{h} , \mathbf{p}_g and W_g are the matrix of lattice vectors, their velocities, and the barostat mass, respectively. For simplicity, only a single Nosé–Hoover thermostat is included, but in CONQUEST, a Nosé–Hoover chain is used. The Liouvillian is constructed, and the integrator is constructed using the splitting of Shinoda *et al.*¹²²

$$iL = iL_r + iL_h + iL_v + iL_{\text{bath}}, \quad (\text{A17})$$

which can be further decomposed,

$$iL_{\text{bath}} = iL_{\text{box}} + iL_{\text{particles}}, \quad (\text{A18})$$

$$iL_{\text{box}} = iL_{v_{\text{box}}} + iL_\xi + iL_{v_{\xi_1}} + iL_{v_{\xi_k}} + iL_{v_{\xi_M}}, \quad (\text{A19})$$

$$iL_{\text{particles}} = iL_{v_{\text{part}}} + iL_\xi + iL_{v_{\xi_1}} + iL_{v_{\xi_k}} + iL_{v_{\xi_M}}. \quad (\text{A20})$$

Then, using Liouville's theorem, we have

$$iL_r = \sum_{i=1}^N [\mathbf{v}_i + \mathbf{v}_g \mathbf{r}_i] \cdot \nabla_{\mathbf{r}_i}, \quad (\text{A21})$$

$$iL_h = \sum_{\alpha,\beta} \mathbf{v}_{g,\alpha\beta} \mathbf{h}_{\alpha\beta} \frac{\partial}{\partial \mathbf{h}_{\alpha\beta}}, \quad (\text{A22})$$

$$iL_v = \sum_{i=1}^N \left(\frac{\mathbf{F}_i}{m_i} \right) \cdot \nabla_{\mathbf{v}_i}, \quad (\text{A23})$$

$$\begin{aligned} iL_{\text{bath}} &= iL_{\text{vpart}} + iL_{\text{vbox}} + iL_{\xi} + iL_{v_{\xi_1}} + iL_{v_{\xi_k}} + iL_{v_{\xi_M}} \\ &= \sum_{i=1}^N \left[- \left\{ \mathbf{v}_g + \frac{1}{N_f} \text{Tr}(\mathbf{v}_g) + v_{\xi_1} \right\} \mathbf{v}_i \right] \nabla_{\mathbf{v}_i} \\ &\quad + \sum_{\alpha,\beta} \left[\frac{F_{\text{box}}}{W} - v_{\xi_1} \mathbf{v}_{g,\alpha\beta} \right] \frac{\partial}{\partial \mathbf{v}_{g,\alpha\beta}} + \sum_{k=1}^M v_{\xi_k} \frac{\partial}{\partial v_{\xi_k}} \\ &\quad + \left[\frac{F_{\text{NHC}_1}}{Q_1} - v_{\xi_1} v_{\xi_2} \right] \frac{\partial}{\partial v_{\xi_1}} \\ &\quad + \sum_{k=2}^M \left[\frac{1}{Q_k} (Q_{k-1} v_{\xi_{k-1}}^2 - k T_{\text{ext}}) - v_{\xi_k} v_{\xi_{k+1}} \right] \frac{\partial}{\partial v_{\xi_k}} \\ &\quad + \left[\frac{1}{Q_M} (Q_{M-1} v_{\xi_{M-1}}^2 - k T_{\text{ext}}) \right] \frac{\partial}{\partial v_{\xi_M}}. \end{aligned} \quad (\text{A24})$$

In this instance, we use M Nosé–Hoover heat baths. The equations of motion can then be expanded via the Trotter–Suzuki identity and directly translated into an algorithm,

$$e^{iL\Delta t} = e^{iL_{\text{bath}} \frac{\Delta t}{2}} e^{iL_v \frac{\Delta t}{2}} e^{iL_h \frac{\Delta t}{2}} e^{iL_r \Delta t} e^{iL_h \frac{\Delta t}{2}} e^{iL_v \frac{\Delta t}{2}} e^{iL_{\text{bath}} \frac{\Delta t}{2}}. \quad (\text{A25})$$

This integrator is tested on a bulk crystalline silicon system, as explained in Sec. III C.

c. Weak-coupling thermostat and barostat

The Berendsen weak coupling method¹²⁴ involves global coupling to a pressure and/or heat bath via a Langevin-type equation of motion with a global friction constant. In the case of the thermostat, the ionic velocities are rescaled by a factor λ , which is scaled toward the target temperature T_0 by the coupling frequency $1/\tau_T$,

$$\lambda = \left[1 + \frac{\Delta t}{\tau_T} \left(\frac{T_0}{T} - 1 \right) \right]^{\frac{1}{2}}. \quad (\text{A26})$$

Similarly, for the barostat, the cell is rescaled by the matrix μ , which is scaled toward a target pressure tensor \mathbf{P}_0 by the pressure coupling frequency $1/\tau_P$ and the estimated bulk modulus β ,

$$\mu = \mathbf{I} - \frac{\beta \Delta t}{3 \tau_P} (\mathbf{P}_0 - \mathbf{P}). \quad (\text{A27})$$

While trivial to implement, weak coupling will not generate the correct canonical or isobaric–isothermal velocity distribution, and the thermostat has the pathological effect of systematically transferring energy to the most slowly changing degrees of freedom (the “flying ice cube” effect). However, it may be useful for equilibration.

d. Stochastic velocity rescaling

Stochastic velocity rescaling¹²³ is essentially a modification of the weak coupling method that does not suffer from the flying ice cube effect. A correctly constructed random force is added to enforce the correct NVT (or NPT) phase space distribution. The kinetic energy is rescaled such that the change in kinetic energy between ionic steps is

$$dK = (\bar{K} - K) \frac{dt}{\tau} + 2 \sqrt{\frac{K \bar{K}}{N_f}} \frac{dW}{\sqrt{\tau}}, \quad (\text{A28})$$

where \bar{K} is the target kinetic energy (i.e., heat bath temperature), dt is the time step, τ is the time scale of the thermostat, N_f is the number of degrees of freedom, and dW is a Wiener process. In practice, the particle velocities are rescaled by a factor α , defined as

$$\begin{aligned} \alpha^2 &= e^{-\Delta t/\tau} + \frac{\bar{K}}{N_f K} \left(1 - e^{-\Delta t/\tau} \right) \left(R_1^2 + \sum_{i=2}^{N_f} R_i^2 \right) \\ &\quad + 2e^{-\Delta t/2\tau} \sqrt{\frac{\bar{K}}{N_f K}} \left(1 - e^{-\Delta t/\tau} \right) R_1, \end{aligned} \quad (\text{A29})$$

where R_i is a set of N_f normally distributed random numbers with unitary variance. This thermostat can be used in NPT dynamics⁹² by barostatting the system via the Parrinello–Rahman method, but with additional R_i 's for the cell degrees of freedom, thermostating the cell velocities as well as the ionic velocities.

REFERENCES

- R. O. Jones, “Density functional theory: Its origins, rise to prominence, and future,” *Rev. Mod. Phys.* **87**, 897–923 (2015).
- J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, “The SIESTA method for ab initio order- N materials simulation,” *J. Phys.: Condens. Matter* **14**, 2745–2779 (2002); [arXiv:cond-mat/0111138v1](#) [cond-mat].
- D. R. Bowler, T. Miyazaki, and M. J. Gillan, “Recent progress in linear scaling *ab initio* electronic structure techniques,” *J. Phys.: Condens. Matter* **14**, 2781–2798 (2002).
- C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, “Introducing ONETEP: Linear-scaling density functional simulations on parallel computers,” *J. Chem. Phys.* **122**, 084119 (2005).
- M. J. Rayson and P. R. Briddon, “Highly efficient method for Kohn–Sham density functional calculations of 500–10 000 atom systems,” *Phys. Rev. B* **80**, 205104 (2009).
- D. R. Bowler and T. Miyazaki, “Calculations for millions of atoms with density functional theory: Linear scaling shows its potential,” *J. Phys.: Condens. Matter* **22**, 074207 (2010).
- J. VandeVondele, U. Borštnik, and J. Hutter, “Linear scaling self-consistent field calculations with millions of atoms in the condensed phase,” *J. Chem. Theory Comput.* **8**, 3565–3573 (2012).
- S. Mohr, L. E. Ratcliff, L. Genovese, D. Caliste, P. Boulanger, S. Goedecker, and T. Deutsch, “Accurate and efficient linear scaling DFT calculations with universal applicability,” *Phys. Chem. Chem. Phys.* **17**, 31360–31370 (2015).
- T. Ozaki, M. Fukuda, and G. Jiang, “Efficient $O(N)$ divide-conquer method with localized single-particle natural orbitals,” *Phys. Rev. B* **98**, 245137 (2018).
- C. Romero-Muñoz, A. Nakata, P. Pou, D. R. Bowler, T. Miyazaki, and R. Pérez, “High-accuracy large-scale DFT calculations using localized orbitals in complex electronic systems: The case of graphene–metal interfaces,” *J. Phys.: Condens. Matter* **30**, 505901 (2018).

- ¹¹W. Kohn, "Density functional and density matrix method scaling linearly with the number of atoms," *Phys. Rev. Lett.* **76**, 3168–3171 (1996).
- ¹²S. Goedecker, "Linear scaling electronic structure methods," *Rev. Mod. Phys.* **71**, 1085 (1999).
- ¹³D. R. Bowler and T. Miyazaki, " $\mathcal{O}(n)$ methods in electronic structure calculations," *Rep. Prog. Phys.* **75**, 036503 (2012).
- ¹⁴D. J. Cole and N. D. M. Hine, "Applications of large-scale density functional theory in biology," *J. Phys.: Condens. Matter* **28**, 393001 (2016).
- ¹⁵D. R. Bowler, T. Miyazaki, A. Nakata, L. Truflandier, M. Arita, J. S. Baker, V. Brazdova, R. Choudhury, S. Y. Mujahed, J. T. Poulton, Z. Raza, A. Sena, U. Terranova, L. Tong, A. Torralba, I. J. Bush, C. M. Goringe, E. H. Hernandez, and M. J. Gillan, CONQUEST: Large-scale DFT calculations, 2020, <https://github.com/OrderN/CONQUEST-release>.
- ¹⁶D. R. Hamann, "Optimized norm-conserving Vanderbilt pseudopotentials," *Phys. Rev. B* **88**, 085117 (2013).
- ¹⁷D. Vanderbilt, "Optimally smooth norm-conserving pseudopotentials," *Phys. Rev. B* **32**, 8412–8415 (1985).
- ¹⁸M. J. van Setten, M. Giantomassi, E. Bousquet, M. J. Verstraete, D. R. Hamann, X. Gonze, and G.-M. Rignanese, "The PseudoDojo: Training and grading a 85 element optimized norm-conserving pseudopotential table," *Comput. Phys. Commun.* **226**, 39–54 (2018).
- ¹⁹M. Schlipf and F. Gygi, "Optimization algorithm for the generation of ONCV pseudopotentials," *Comput. Phys. Commun.* **196**, 36–44 (2015).
- ²⁰K. Lejaeghere, G. Bihlmayer, T. Björkman, P. Blaha, S. Blugel, V. Blum, D. Caliste, I. E. Castelli, S. J. Clark, A. Dal Corso, S. de Gironcoli, T. Deutsch, J. K. Dewhurst, I. Di Marco, C. Draxl, M. Dufak, O. Eriksson, J. A. Flores-Livas, K. F. Garrity, L. Genovese, P. Giannozzi, M. Giantomassi, S. Goedecker, X. Gonze, O. Granas, E. K. U. Gross, A. Gulans, F. Gygi, D. R. Hamann, P. J. Hasnip, N. A. W. Holzwarth, D. Iuşan, D. B. Jochym, F. Jollet, D. Jones, G. Kresse, K. Koepnik, E. Kucukbenli, Y. O. Kvashnin, I. L. M. Locht, S. Lubeck, M. Marsman, N. Marzari, U. Nitzsche, L. Nordstrom, T. Ozaki, L. Paulatto, C. J. Pickard, W. Poelmans, M. I. J. Probert, K. Refson, M. Richter, G.-M. Rignanese, S. Saha, M. Scheffler, M. Schlipf, K. Schwarz, S. Sharma, F. Tavazza, P. Thunstrom, A. Tkatchenko, M. Torrent, D. Vanderbilt, M. J. van Setten, V. Van Speybroeck, J. M. Wills, J. R. Yates, G.-X. Zhang, and S. Cottenier, "Reproducibility in density functional theory calculations of solids," *Science* **351**, aad3000 (2016).
- ²¹S. G. Louie, S. Froyen, and M. L. Cohen, "Nonlinear ionic pseudopotentials in spin-density-functional calculations," *Phys. Rev. B* **26**, 1738–1742 (1982).
- ²²O. F. Sankey and D. J. Niklewski, "*Ab initio* multicenter tight-binding model for molecular-dynamics simulations and other applications in covalent systems," *Phys. Rev. B* **40**, 3979–3995 (1989).
- ²³T. Ozaki and H. Kino, "Efficient projector expansion for the *ab initio* LCAO method," *Phys. Rev. B* **72**, 045121 (2005).
- ²⁴D. R. Bowler, J. S. Baker, J. T. L. Poulton, S. Y. Mujahed, J. Lin, S. Yadav, Z. Raza, and T. Miyazaki, "Highly accurate local basis sets for large-scale DFT calculations in conquest," *Jpn. J. Appl. Phys., Part 1* **58**, 100503 (2019).
- ²⁵E. Hernández, M. J. Gillan, and C. M. Goringe, "Linear-scaling density-functional-theory technique: The density-matrix approach," *Phys. Rev. B* **53**, 7147–7157 (1996).
- ²⁶E. Hernández and M. J. Gillan, "Self-consistent first-principles technique with linear scaling," *Phys. Rev. B* **51**, 10157–10160 (1995).
- ²⁷E. Hernández, M. J. Gillan, and C. M. Goringe, "Basis functions for linear-scaling first-principles calculations," *Phys. Rev. B* **55**, 13485–13493 (1997).
- ²⁸A. S. Torralba, M. Todorović, V. Brázdová, R. Choudhury, T. Miyazaki, M. J. Gillan, and D. R. Bowler, "Pseudo-atomic orbitals as basis sets for the $\mathcal{O}(N)$ DFT code CONQUEST," *J. Phys.: Condens. Matter* **20**, 294206 (2008).
- ²⁹A. Nakata, D. R. Bowler, and T. Miyazaki, "Efficient calculations with multisite local orbitals in a large-scale DFT code CONQUEST," *J. Chem. Theory Comput.* **10**, 4813 (2014).
- ³⁰A. Nakata, D. R. Bowler, and T. Miyazaki, "Optimized multi-site local orbitals in the large-scale DFT program CONQUEST," *Phys. Chem. Chem. Phys.* **17**, 31427 (2015).
- ³¹M. J. Rayson, "Rapid filtration algorithm to construct a minimal basis on the fly from a primitive Gaussian basis," *Comput. Phys. Commun.* **181**, 1051–1056 (2010).
- ³²C. A. White, P. Maslen, M. S. Lee, and M. Head-Gordon, "The tensor properties of energy gradients within a non-orthogonal basis," *Chem. Phys. Lett.* **276**, 133–138 (1997).
- ³³We note that the overlap matrix can be inverted without problem when using exact diagonalisation.
- ³⁴M. S. Lee and M. Head-Gordon, "Polarized atomic orbitals for self-consistent field electronic structure calculations," *J. Chem. Phys.* **107**, 9085 (1997).
- ³⁵G. Berghold, M. Parrinello, and J. Hutter, "Polarized atomic orbitals for linear scaling methods," *J. Chem. Phys.* **116**, 1800–1810 (2002).
- ³⁶D. R. Bowler, I. J. Bush, and M. J. Gillan, "Practical methods for *ab initio* calculations on thousands of atoms," *Int. J. Quantum Chem.* **77**, 831 (2000).
- ³⁷N. D. Woods, M. C. Payne, and P. J. Hasnip, "Computing the self-consistent field in Kohn–Sham density functional theory," *J. Phys.: Condens. Matter* **31**, 453001 (2019).
- ³⁸P. Pulay, "Convergence acceleration of iterative sequences. The case of SCF iteration," *Chem. Phys. Lett.* **73**, 393 (1980).
- ³⁹As is common in the field, we use the term "exact diagonalisation" throughout the paper to indicate that the eigenstates and density matrix are found without search or approximation.
- ⁴⁰L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley, *SciLAPACK Users' Guide*, Society for Industrial and Applied Mathematics, Philadelphia, PA, 1997.
- ⁴¹A. Marek, V. Blum, R. Johanni, V. Havu, B. Lang, T. Auckenthaler, A. Heinecke, H.-J. Bungartz, and H. Lederer, "The ELPA library: Scalable parallel eigenvalue solutions for electronic structure theory and computational science," *J. Phys.: Condens. Matter* **26**, 213201 (2014).
- ⁴²H. J. Monkhorst and J. D. Pack, "Special points for Brillouin-zone integrations," *Phys. Rev. B* **13**, 5188–5192 (1976).
- ⁴³M. Methfessel and A. T. Paxton, "High-precision sampling for Brillouin-zone integration in metals," *Phys. Rev. B* **40**, 3616 (1989).
- ⁴⁴X.-P. Li, R. W. Nunes, and D. Vanderbilt, "Density-matrix electronic-structure method with linear system-size scaling," *Phys. Rev. B* **47**, 10891–10894 (1993).
- ⁴⁵R. W. Nunes and D. Vanderbilt, "Generalization of the density-matrix method to a nonorthogonal basis," *Phys. Rev. B* **50**, 17611–17614 (1994).
- ⁴⁶A. H. R. Palser and D. E. Manolopoulos, "Canonical purification of the density matrix in electronic-structure theory," *Phys. Rev. B* **58**, 12704–12711 (1998).
- ⁴⁷D. R. Bowler and M. J. Gillan, "Density matrices in $\mathcal{O}(N)$ electronic structure calculations: Theory and applications," *Comput. Phys. Commun.* **120**, 95–108 (1999).
- ⁴⁸D. R. Bowler and M. J. Gillan, "Length-scale ill conditioning in linear-scaling DFT," *Comput. Phys. Commun.* **112**, 103–111 (1998).
- ⁴⁹A. Nakata, Y. Futamura, T. Sakurai, D. R. Bowler, and T. Miyazaki, "Efficient calculation of electronic structure using $\mathcal{O}(N)$ density functional theory," *J. Chem. Theory Comput.* **13**, 4146 (2017).
- ⁵⁰T. Sakurai and H. Sugiura, "A projection method for generalized eigenvalue problems using numerical integration," *J. Comput. Appl. Math.* **159**, 119–128 (2003).
- ⁵¹T. Sakurai, Y. Futamura, and H. Tadano, "Efficient parameter estimation and implementation of a contour integral-based eigensolver," *J. Algorithms Comput. Technol.* **7**, 249–269 (2013).
- ⁵²I. Yamazaki, H. Tadano, T. Sakurai, and T. Ikegami, "Performance comparison of parallel eigensolvers based on a contour integral method and a Lanczos method," *Parallel Comput.* **39**, 280–290 (2013).
- ⁵³M. Toyoda and T. Ozaki, "Numerical evaluation of electron repulsion integrals for pseudoatomic orbitals and their derivatives," *J. Chem. Phys.* **130**, 124114 (2009).
- ⁵⁴M. Toyoda and T. Ozaki, "LIBERI: Library for numerical evaluation of electron-repulsion integrals," *Comput. Phys. Commun.* **181**, 1455–1463 (2010).
- ⁵⁵J. T. Fermann and E. F. Valeev, "Libint: Machine-generated library for efficient evaluation of molecular integrals over Gaussians," freely available at <http://libint.valeev.net/> or one of the authors, 2003.
- ⁵⁶H. Shang, Z. Li, and J. Yang, "Implementation of screened hybrid density functional for periodic systems with numerical atomic orbitals: Basis function fitting and integral screening," *J. Chem. Phys.* **135**, 034110 (2011).

- ⁵⁷X. Qin, H. Shang, H. Xiang, Z. Li, and J. Yang, "HONPAS: A linear scaling open-source solution for large system simulations," *Int. J. Quantum Chem.* **115**, 647–655 (2015).
- ⁵⁸L. Truflandier, T. Miyazaki, and D. R. Bowler, "Linear-scaling implementation of exact exchange using localized numerical orbitals and contraction reduction integrals," [arXiv:1112.5989v2](https://arxiv.org/abs/1112.5989v2) (2012).
- ⁵⁹D. J. Wales and M. P. Hodges, "Global minima of water clusters (H₂O)_n, $n \leq 21$, described by an empirical potential," *Chem. Phys. Lett.* **286**, 65–72 (1998).
- ⁶⁰A. Castro, A. Rubio, and M. J. Stott, "Solution of Poisson's equation for finite systems using plane-wave methods," *Can. J. Phys.* **81**, 1151–1164 (2003).
- ⁶¹M. R. Jarvis, I. D. White, R. W. Godby, and M. C. Payne, *Phys. Rev. B* **56**, 14972–14978 (1997).
- ⁶²G. J. Martyna and M. E. Tuckerman, *J. Chem. Phys.* **110**, 2810 (1999).
- ⁶³C. A. Rozzi, D. Varsano, A. Marini, E. K. U. Gross, and A. Rubio, *Phys. Rev. B* **73**, 205119 (2006).
- ⁶⁴I. Dabo, B. Kozinsky, N. E. Singh-Miller, and N. Marzari, *Phys. Rev. B* **77**, 115139 (2008).
- ⁶⁵H.-S. Lee and M. E. Tuckerman, *J. Chem. Phys.* **129**, 224108 (2008).
- ⁶⁶M. A. Watson and K. Hirao, *J. Chem. Phys.* **129**, 184107 (2008).
- ⁶⁷K. Varga, Z. Zhang, and S. T. Pantelides, *Phys. Rev. Lett.* **93**, 176403 (2004).
- ⁶⁸L. Genovese, T. Deutsch, A. Neelov, S. Goedecker, and G. Beylkin, *J. Chem. Phys.* **125**, 074105 (2006).
- ⁶⁹L. Genovese, T. Deutsch, and S. Goedecker, *J. Chem. Phys.* **127**, 054704 (2007).
- ⁷⁰G. Onida, L. Reining, R. W. Godby, R. Del Sole, and W. Andreoni, *Phys. Rev. Lett.* **75**, 818–821 (1995).
- ⁷¹P. E. Blöchl, *J. Chem. Phys.* **103**, 7422 (1995).
- ⁷²P. A. Schultz, *Phys. Rev. B* **60**, 1551–1554 (1999).
- ⁷³T. Miyazaki, D. R. Bowler, R. Choudhury, and M. J. Gillan, "Atomic force algorithms in density functional theory electronic-structure techniques based on local orbitals," *J. Chem. Phys.* **121**, 6186–6194 (2004).
- ⁷⁴A. S. Torralba, D. R. Bowler, T. Miyazaki, and M. J. Gillan, "Non-self-consistent density-functional theory exchange-correlation forces for GGA functionals," *J. Chem. Theory Comput.* **5**, 1499 (2009).
- ⁷⁵F. Knuth, C. Carbogno, V. Atalla, V. Blum, and M. Scheffler, "All-electron formalism for total energy strain derivatives and stress tensor components for numeric atom-centered orbitals," *Comput. Phys. Commun.* **190**, 33–50 (2015).
- ⁷⁶O. H. Nielsen and R. M. Martin, "First-principles calculation of stress," *Phys. Rev. Lett.* **50**, 697–700 (1983).
- ⁷⁷O. H. Nielsen and R. M. Martin, "Stresses in semiconductors: *Ab initio* calculations on Si, Ge, and GaAs," *Phys. Rev. B* **32**, 3792–3805 (1985).
- ⁷⁸O. H. Nielsen and R. M. Martin, "Quantum-mechanical theory of stress and force," *Phys. Rev. B* **32**, 3780–3791 (1985).
- ⁷⁹O. H. Nielsen and R. M. Martin, "Erratum: Quantum-mechanical theory of stress and force," *Phys. Rev. B* **35**, 9308 (1987).
- ⁸⁰S. Ismail-Beigi and T. A. Arias, "Locality of the density matrix in metals, semiconductors, and insulators," *Phys. Rev. Lett.* **82**, 2127–2130 (1999).
- ⁸¹L. He and D. Vanderbilt, "Exponential decay properties of Wannier functions and related quantities," *Phys. Rev. Lett.* **86**, 5341–5344 (2001).
- ⁸²S. N. Taraskin, D. A. Drabold, and S. R. Elliott, "Spatial decay of the single-particle density matrix in insulators: Analytic results in two and three dimensions," *Phys. Rev. Lett.* **88**, 196405 (2002).
- ⁸³E. Bitzek, P. Koskinen, F. Gähler, M. Moseler, and P. Gumbsch, "Structural relaxation made simple," *Phys. Rev. Lett.* **97**, 170201 (2006).
- ⁸⁴D. Packwood, J. Kermode, L. Mones, N. Bernstein, J. Woolley, N. Gould, C. Ortner, and G. Csányi, "A universal preconditioner for simulating condensed phase materials," *J. Chem. Phys.* **144**, 164109 (2016).
- ⁸⁵M. V. Fernández-Serra, E. Artacho, and J. M. Soler, "Model hessian for accelerating first-principles structure optimizations," *Phys. Rev. B* **67**, 100101 (2003).
- ⁸⁶A. M. N. Niklasson, "Extended Born-Oppenheimer molecular dynamics," *Phys. Rev. Lett.* **100**, 123004 (2008).
- ⁸⁷A. M. N. Niklasson, C. J. Tymczak, and M. Challacombe, "Time-reversible Born-Oppenheimer molecular dynamics," *Phys. Rev. Lett.* **97**, 123001 (2006).
- ⁸⁸A. M. N. Niklasson, "Next generation extended Lagrangian first principles molecular dynamics," *J. Chem. Phys.* **147**, 054103 (2017).
- ⁸⁹M. Arita, D. R. Bowler, and T. Miyazaki, "Stable and efficient linear scaling first-principles molecular dynamics for 10000+ atoms," *J. Chem. Theory Comput.* **10**, 5419 (2014).
- ⁹⁰A. M. N. Niklasson, P. Steneteg, A. Odell, N. Bock, M. Challacombe, C. J. Tymczak, E. Holmström, G. Zheng, and V. Weber, "Extended Lagrangian Born-Oppenheimer molecular dynamics with dissipation," *J. Chem. Phys.* **130**, 214109 (2009).
- ⁹¹T. Hirakawa, T. Suzuki, D. R. Bowler, and T. Miyazaki, "Canonical-ensemble extended Lagrangian Born-Oppenheimer molecular dynamics for the linear scaling density functional theory," *J. Phys.: Condens. Matter* **29**, 405901 (2017).
- ⁹²G. Bussi, T. Zykova-Timan, and M. Parrinello, "Isothermal-isobaric molecular dynamics using stochastic velocity rescaling," *J. Chem. Phys.* **130**, 074101 (2009).
- ⁹³M. Parrinello and A. Rahman, "Polymorphic transitions in single crystals: A new molecular dynamics method," *J. Appl. Phys.* **52**, 7182 (1981).
- ⁹⁴G. J. Martyna, M. E. Tuckerman, D. J. Tobias, and M. L. Klein, "Explicit reversible integrators for extended systems dynamics," *Mol. Phys.* **87**, 1117 (1996).
- ⁹⁵M. Arita, S. Arapan, D. R. Bowler, and T. Miyazaki, "Large-scale DFT simulations with a linear-scaling DFT code CONQUEST on K-computer," *J. Adv. Simul. Sci. Eng.* **1**, 87–97 (2014).
- ⁹⁶H. Meuer, E. Strohmaier, J. Dongarra, H. Simon, and M. Meuer, Top 500 list, 2012.
- ⁹⁷C. O'Rourke, S. Y. Mujahed, C. Kumarasinghe, T. Miyazaki, and D. R. Bowler, "Structural properties of silicon-germanium and germanium-silicon core-shell nanowires," *J. Phys.: Condens. Matter* **30**, 465303 (2018).
- ⁹⁸C. Kumarasinghe and D. R. Bowler, "DFT study of undoped and As-doped Si nanowires approaching the bulk limit," *J. Phys.: Condens. Matter* **32**, 035304 (2019).
- ⁹⁹T. Miyazaki, D. R. Bowler, M. J. Gillan, and T. Ohno, "The energetics of hut-cluster self-assembly in Ge/Si(001) from linear-scaling DFT calculations," *J. Phys. Soc. Jpn.* **77**, 123706 (2008).
- ¹⁰⁰Y. Li, M. Buerkle, G. Li, A. Rostamian, H. Wang, Z. Wang, D. R. Bowler, T. Miyazaki, L. Xiang, Y. Asai, G. Zhou, and N. Tao, "Gate controlling of quantum interference and direct observation of anti-resonances in single molecule charge transport," *Nat. Mater.* **18**, 357–363 (2019).
- ¹⁰¹T. Otsuka, T. Miyazaki, T. Ohno, D. R. Bowler, and M. J. Gillan, "Accuracy of order- N density-functional theory calculations on DNA systems using CONQUEST," *J. Phys.: Condens. Matter* **20**, 294201 (2008).
- ¹⁰²T. Otsuka, M. Taiji, D. R. Bowler, and T. Miyazaki, "Linear-scaling first-principles molecular dynamics of complex biological systems with the Conquest code," *Jpn. J. Appl. Phys., Part 1* **55**, 1102B1 (2016).
- ¹⁰³M. Todorović, D. R. Bowler, M. J. Gillan, and T. Miyazaki, "Density-functional theory study of gramicidin A ion channel geometry and electronic properties," *J. R. Soc. Interface* **10**, 20130547 (2013).
- ¹⁰⁴Y.-W. Mo, D. E. Savage, B. S. Swartzentruber, and M. G. Lagally, "Kinetic pathway in Stranski-Krastanov growth of Ge on Si(001)," *Phys. Rev. Lett.* **65**, 1020–1023 (1990).
- ¹⁰⁵T. Miyazaki, D. R. Bowler, R. Choudhury, and M. J. Gillan, "Density functional calculations for Ge(105): Local basis sets and $O(N)$ methods," *Phys. Rev. B* **76**, 115327 (2007).
- ¹⁰⁶T. Hashimoto, Y. Morikawa, and K. Terakura, "Stability and electronic structure of Ge(105)1×2: A first-principles theoretical study," *Surf. Sci.* **576**, 61–66 (2005).
- ¹⁰⁷P. Raiteri, D. B. Migas, L. Miglio, A. Rastelli, and H. von Känel, "Critical role of the surface reconstruction in the thermodynamic stability of {105}Ge pyramids on Si(001)," *Phys. Rev. Lett.* **88**, 256103 (2002).
- ¹⁰⁸S. Arapan, D. R. Bowler, and T. Miyazaki, "A linear scaling DFT study of the growth of a new 105 facet layer on a Ge hut cluster," [arXiv:1510.00526](https://arxiv.org/abs/1510.00526) (2015).
- ¹⁰⁹M. R. McKay, J. A. Venables, and J. Drucker, "Kinetically suppressed Ostwald ripening of Ge/Si (100) hut clusters," *Phys. Rev. Lett.* **101**, 216104 (2008).
- ¹¹⁰T. Miyazaki, "Invited large-scale DFT study of Ge/Si 3D nanoislands and core-shell nanowires," *ECS Trans.* **86**, 269–279 (2018).

- ¹¹¹C. Thelander, P. Agarwal, S. Brongersma, J. Eymery, L. F. Feiner, A. Forchel, M. Scheffler, W. Riess, B. J. Ohlsson, U. Gösele, and L. Samuelson, "Nanowire-based one-dimensional electronics," *Mater. Today* **9**, 28–35 (2006).
- ¹¹²W. Lu, J. Xiang, B. P. Timko, Y. Wu, and C. M. Lieber, "One-dimensional hole gas in germanium/silicon nanowire heterostructures," *Proc. Natl. Acad. Sci. U. S. A.* **102**, 10046 (2005).
- ¹¹³J. Xiang, W. Lu, Y. Hu, Y. Wu, H. Yan, and C. M. Lieber, "Ge/Si nanowire heterostructures as high-performance field-effect transistors," *Nature* **441**, 489–493 (2006).
- ¹¹⁴N. Fukata, M. Mitome, T. Sekiguchi, Y. Bando, M. Kirkham, J.-I. Hong, Z. L. Wang, and R. L. Snyder, "Characterization of impurity doping and stress in Si/Ge and Ge/Si core-shell nanowires," *ACS Nano* **6**, 8887–8895 (2012).
- ¹¹⁵J. S. Baker and D. R. Bowler, "First-principles soft-mode lattice dynamics of $\text{PbZr}_{0.5}\text{Ti}_{0.5}\text{O}_3$ and shortcomings of the virtual crystal approximation," *Phys. Rev. B* **100**, 224305 (2019).
- ¹¹⁶B. Meyer and D. Vanderbilt, "*Ab initio* study of ferroelectric domain walls in PbTiO_3 ," *Phys. Rev. B* **65**, 104111 (2002).
- ¹¹⁷M. J. Gillan, D. R. Bowler, A. S. Torralba, and T. Miyazaki, "Order- N first-principles calculations with the conquest code," *Comput. Phys. Commun.* **177**, 14–18 (2007), part of special issue: Proceedings of the Conference on Computational Physics 2006.
- ¹¹⁸T. Hirakawa, D. R. Bowler, T. Miyazaki, Y. Morikawa, and L. A. Truflandier, "Blue moon ensemble simulation of aquation free energy profiles applied to mono and bifunctional platinum anticancer drugs," [arXiv:2003.01418](https://arxiv.org/abs/2003.01418) (2020).
- ¹¹⁹S. Nosé, "A unified formulation of the constant temperature molecular dynamics methods," *J. Chem. Phys.* **81**, 511 (1984).
- ¹²⁰W. G. Hoover, "Canonical dynamics: Equilibrium phase-space distributions," *Phys. Rev. A* **31**, 1695 (1985).
- ¹²¹G. J. Martyna, M. L. Klein, and M. Tuckerman, "Nosé–Hoover chains: The canonical ensemble via continuous dynamics," *J. Chem. Phys.* **97**, 2635 (1992).
- ¹²²W. Shinoda, M. Shiga, and M. Mikami, "Rapid estimation of elastic constants by molecular dynamics simulation under constant stress," *Phys. Rev. B* **69**, 134103 (2004).
- ¹²³G. Bussi, D. Donadio, and M. Parrinello, "Canonical sampling through velocity rescaling," *J. Chem. Phys.* **126**, 014101 (2007).
- ¹²⁴H. J. C. Berendsen, J. P. M. Postma, W. F. van Gunsteren, A. Dinola, and J. R. Haak, "Molecular dynamics with coupling to an external bath," *J. Chem. Phys.* **81**, 3684 (1984).