

Interpolative Separable Density Fitting Decomposition for Accelerating Hybrid Density Functional Calculations with Applications to Defects in Silicon

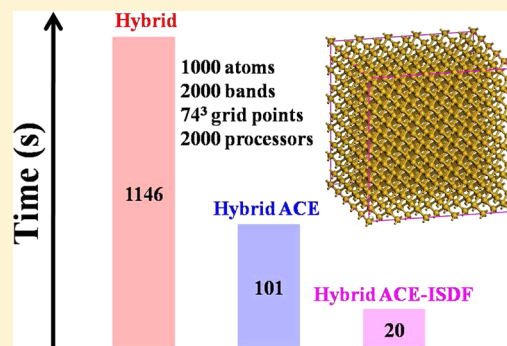
Wei Hu,^{*,†} Lin Lin,^{*,‡,†} and Chao Yang^{*,†}

[†]Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, California 94720, United States

[‡]Department of Mathematics, University of California, Berkeley, California 94720, United States

ABSTRACT: We present a new efficient way to perform hybrid density functional theory (DFT)-based electronic structure calculations. The new method uses an interpolative separable density fitting (ISDF) procedure to construct a set of numerical auxiliary basis vectors and a compact approximation of the matrix consisting of products of occupied orbitals represented in a large basis set such as the plane-wave basis. Such an approximation allows us to reduce the number of Poisson solves from $O(N_e^2)$ to $O(N_e)$ when we apply the exchange operator to occupied orbitals in an iterative method for solving the Kohn–Sham equations, where N_e is the number of electrons in the system to be studied. We show that the ISDF procedure can be carried out in $O(N_e^3)$ operations, with a much smaller preconstant compared to methods used in existing approaches. When

combined with the recently developed adaptively compressed exchange (ACE) operator formalism, which reduces the number of times the exchange operator needs to be updated, the resulting ACE-ISDF method significantly reduces the computational cost associated with the exchange operator by nearly 2 orders of magnitude compared to existing approaches for a large silicon system with 1000 atoms. We demonstrate that the ACE-ISDF method can produce accurate energies and forces for insulating and metallic systems and that it is possible to obtain converged hybrid functional calculation results for a 1000-atom bulk silicon within 10 min on 2000 computational cores. We also show that ACE-ISDF can scale to 8192 computational cores for a 4096-atom bulk silicon system. We use the ACE-ISDF method to geometrically optimize a 1000-atom silicon system with a vacancy defect using the HSE06 functional and computes its electronic structure. We find that the computed energy gap from the HSE06 functional is much closer to the experimental value compared to that produced by semilocal functionals in the DFT calculations.



1. INTRODUCTION

Kohn–Sham density functional theory (KSDF) is the most widely used electronic structure theory in condensed matter physics and quantum chemistry. The fidelity of the results produced by a KSDF calculation often depends on the choice of the exchange and correlation functional.³ Hybrid exchange–correlation functionals, such as B3LYP,⁴ PBE0,⁵ and HSE,^{6,7} are known to be more reliable in producing high fidelity results for a wide range of systems over calculations that make use of local and semilocal exchange–correlation functionals, such as the local density approximation (LDA),^{8–10} the generalized gradient approximation (GGA),^{11–13} and meta-GGA functionals.^{14–16} However, hybrid functionals include a fraction of the Fock exchange operator. Here, $V_X[\{\psi_i\}](\mathbf{r}, \mathbf{r}') = -\sum_{i=1}^n \frac{\psi_i(\mathbf{r})\psi_i(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}$. Applying such an operator to a set of n orbitals $\Psi = [\psi_1(\mathbf{r}), \dots, \psi_n(\mathbf{r})]$, which is often used in an iterative method for solving the Kohn–Sham equations, requires solving $O(N_e^2)$ Poisson-like equations, with effective charges taking the form of $\psi_i(\mathbf{r})\psi_j(\mathbf{r})$ ($1 \leq i, j \leq n$). Here, $n \sim O(N_e)$ and N_e is the number of electrons. This is costly, especially for calculations performed in a

large basis set such as plane-waves and finite elements. In these calculations, an iterative diagonalization procedure is used to solve the KS equations, and the multiplication of V_X with Ψ , which has the complexity of $O(N_e^3)$ with a large preconstant, needs to be performed in each iteration. These multiplications alone often constitute more than 95% of the overall computational time in a conventional approach.

There are two main routes to reducing the computational cost of hybrid functional calculations. The first route is to reduce the cost of multiplying V_X with Ψ . This can be done through efficient parallelization over a large number of processors,^{17–19} or the use of linear scaling methods (i.e., $O(N_e)$ methods).^{20–25} For large systems with a substantial band gap, linear scaling methods use the nearsightedness property²⁶ to construct a sparse approximation to the exchange operator, thereby reducing the cost of computing $V_X\Psi$.²⁷ The second route is to reduce the frequency of computing $V_X\Psi$, and this route is much less explored until recently.^{17,28–31} The adaptively compressed exchange (ACE)

Received: July 27, 2017

Published: September 29, 2017

operator formalism^{30,31} replaces the dense and full rank exchange operator by a low rank operator that is constructed on the fly. The low rank operator is only updated once every few iterations. The ACE operator fully agrees with V_X in the subspace spanned by orbitals in Ψ . The reduced rank of the ACE operator lowers the cost of $V_X\Psi$ calculation without losing accuracy. The ACE formulation is applicable to insulators, semiconductors, and metals. Recently, we have proved that such a low rank compression of the exchange operator is uniquely determined through the ACE formulation and that the updating scheme for the ACE operator converges in the self-consistent field (SCF) iteration both locally and globally for the linearized Hartree–Fock-like equations.³² The ACE formulation can enable hybrid functional calculations in a planewave basis set for more than a thousand atoms³¹ and has recently been integrated¹⁹ into software packages such as Quantum ESPRESSO.³³

In this paper, we develop a new method that combines the strength of both approaches mentioned above to accelerate large-scale hybrid functional calculations. By using an interpolative separable density fitting (ISDF) method first proposed by Lu and Ying in ref 34, we construct a numerical auxiliary basis for $\{\psi_i(\mathbf{r})\psi_j(\mathbf{r})\} (1 \leq i, j \leq n)$ that contains only cN_e basis vectors for a small constant c . As a result, applying V_X to a set of $O(N_e)$ orbitals only requires solving $O(N_e)$ instead of $O(N_e^2)$ Poisson-like equations. Compared to the widely used density fitting and Cholesky decomposition techniques^{35–40} in quantum chemistry, the main feature of the ISDF decomposition is that the fitting coefficient tensor, which is usually written as a three-way tensor, can be analytically separated into the product of two matrices. This is the key to achieving $O(N_e^3)$ scaling and avoiding the $O(N_e^4)$ computational complexity that appears in many other density fitting schemes. The ISDF decomposition is closely related to the recently developed tensor hypercontraction (THC) approach.^{41,42}

The ISDF decomposition replaces $\{\psi_i(\mathbf{r})\psi_j(\mathbf{r})\} (1 \leq i, j \leq n)$ with the product of two matrices. One of the matrices, which can be viewed as a matrix of fitting coefficients, simply consists of $\{\psi_i\psi_j\}$ evaluated at a set of carefully chosen *interpolation points* \hat{r}_μ for $\mu = 1, \dots, N_\mu$ and $N_\mu = cN_e$. The other matrix contains numerical auxiliary basis vectors that we refer to as the *interpolating vectors*. In this paper, these two matrices are determined separately. The matrix containing the fitting coefficients, which depends solely on the choice of interpolation points, is determined first. The matrix containing the interpolating vectors is subsequently obtained through a least-squares fitting procedure. This approach is different from the decomposition proposed in ref 34, where the numerical auxiliary basis and the fitting coefficients are determined simultaneously through a randomized QR factorization with column pivoting (QRCP) applied to $\{\psi_i(\mathbf{r})\psi_j(\mathbf{r})\}$ directly. We find that using randomized QRCP at each SCF iteration is costly and does not speed up hybrid functional calculations. By separating the treatment of the interpolation points from the construction of the matrix containing the interpolating vectors in ISDF, we can use the relatively expensive randomized QRCP procedure to find the interpolation points in advance and only recompute the interpolation vectors whenever $\{\psi_i(\mathbf{r})\psi_j(\mathbf{r})\}$ has been updated using an efficient least-squares procedure that exploits the separable nature of the matrix to be approximated. As a result, we can significantly accelerate hybrid functional calculations using the ISDF decomposition in all but the first SCF iteration.

The ISDF decomposition can be used in the construction of an ACE operator to reduce the number of Poisson solves. In fact, the decomposition itself yields a low-rank approximation of the Fock exchange operator. However, symmetry is not strictly preserved in the ISDF decomposition. The lack of symmetry can introduce numerical stability issues in the convergence of the SCF iteration. We demonstrate how to combine the ISDF decomposition with the ACE formulation in a numerically stable manner by maintaining the symmetry of the compressed exchange operator. The resulting ACE-ISDF method does not rely on the nearsightedness property and is applicable to insulators, semiconductors, and metals. The computational complexity of our new approach is still $O(N_e^3)$, but the preconstant is significantly reduced. The ACE-ISDF method can be efficiently parallelized on high performance supercomputers. Using this technique, we can perform hybrid functional calculations for a bulk silicon system with 1000 atoms in less than 10 wall clock minutes on 2000 computational cores. We find that the cost associated with the exchange operator is reduced by nearly 2 orders of magnitude compared to conventional approaches. Furthermore, this method can also scale to 8192 computational cores for a 4096-atom bulk silicon system.

As an example, we use the ACE-ISDF method to optimize the geometry and compute the electronic structure of a bulk 1000-atom silicon system that contains a single vacancy at the level of the HSE06 hybrid functional calculations.⁷ Our calculation reveals three defect states within the intrinsic energy gap of the silicon. The computed energy gap is much closer to the experimental value compared to GGA functional calculations.

The rest of the paper is organized as follows. In Section 2, we introduce the density fitting approximation in the context of hybrid functional calculations and the interpolative separable density fitting approximation. In Section 3, we develop a new method to efficiently compute the interpolative separable density fitting approximation and to combine with the adaptively compressed exchange formulation. We describe an efficient parallelization strategy in Section 4. The numerical results are given in Section 5, followed by the conclusions Section 6.

2. PRELIMINARIES

2.1. Density Fitting Approximations in Hybrid Functional Calculations. For simplicity, we consider isolated, gapped systems and omit spin degeneracy. In hybrid functional calculations, the exchange operator is an integral operator defined in terms of the occupied orbitals $\{\varphi_i\}_{i=1}^{N_e}$ with kernel

$$\begin{aligned} V_X[\{\varphi_i\}](\mathbf{r}, \mathbf{r}') &= - \sum_{i=1}^{N_e} \varphi_i(\mathbf{r})\varphi_i(\mathbf{r}')K(\mathbf{r}, \mathbf{r}') \\ &\equiv -P^\rho(\mathbf{r}, \mathbf{r}')K(\mathbf{r}, \mathbf{r}') \end{aligned} \quad (1)$$

where K is either the Coulomb potential $K(\mathbf{r}, \mathbf{r}') = \frac{1}{|\mathbf{r} - \mathbf{r}'|}$ in hybrid functionals such as B3LYP⁴ and PBE0⁵ or the screened Coulomb potential $K(\mathbf{r}, \mathbf{r}') = \frac{\text{erfc}(\mu|\mathbf{r} - \mathbf{r}'|)}{|\mathbf{r} - \mathbf{r}'|}$ in functionals such as HSE.⁶ P^ρ is the density matrix.

When a large basis set such as the planewave basis set is used to discretize the KS equations, it is generally more efficient to perform $V_X[\{\varphi_i\}]\psi_j$, $j = 1, 2, \dots, n$ on the fly in an iterative diagonalization procedure without explicitly constructing or storing $V_X[\{\varphi_i\}]$. In many cases, $n = N_e$, but $\{\varphi_i\}$ and $\{\psi_i\}$ may be different sets of orbitals when used in a self-consistent field (SCF) iteration. It is also possible to have $n > N_e$ when some

unoccupied orbitals are also to be computed. So, we deliberately use different notation to distinguish the occupied orbitals $\{\varphi_i\}$ from generic orbitals $\{\psi_i\}$ that can be either occupied or unoccupied. In order to reach self-consistency for the occupied orbitals $\{\varphi_i\}$, a common practice is to separate the self-consistent field (SCF) iteration into two sets of SCF iterations. In the *inner SCF iteration*, the exchange operation V_X defined by the orbitals $\{\varphi_i\}$ is fixed and the Hamiltonian operator only depends on the density $\rho(\mathbf{r})$. The SCF iteration then proceeds as in KSDFFT calculations with a fixed exchange operator. In the *outer SCF iteration*, the occupied components of the output orbitals $\{\psi_j\}$ can be used as the input orbitals to update the exchange operator. In each inner iteration, the product of $V_X[\{\varphi_i\}]$ and ψ_j need to be evaluated many times using the relation

$$(V_X[\{\varphi_i\}]\psi_j)(\mathbf{r}) = -\sum_{i=1}^{N_e} \varphi_i(\mathbf{r}) \int K(\mathbf{r}, \mathbf{r}') \varphi_i(\mathbf{r}') \psi_j(\mathbf{r}') d\mathbf{r}' \quad (2)$$

The integration in eq 2 is often carried out by solving Poisson-like equations, using, e.g., a fast Fourier transform (FFT) method. The number of equations to be solved is $nN_e \sim O(N_e^2)$. This is typically the most time-consuming component in hybrid functional calculations.

We would like to reduce the number of equations to be solved by exploiting the numerical rank deficiency in the set of right-hand sides $\{\varphi_i(\mathbf{r})\psi_j(\mathbf{r})\}$ in these Poisson-like equations and represent them using a smaller set of linearly independent basis. One possible way to achieve this is through the use of a density fitting method (a.k.a., resolution of identity).^{39,40} In general, given two sets of functions $\{\varphi_i(\mathbf{r})\}_{i=1}^m$, $\{\psi_j(\mathbf{r})\}_{j=1}^n$, a density fitting procedure constructs an auxiliary basis $\{\zeta_\mu\}$, $\mu = 1, 2, \dots, N_\mu$, with $N_\mu \ll mn$, for the set of Hadamard products (i.e., the element-wise product)

$$\{Z_{ij}(\mathbf{r}) := \varphi_i(\mathbf{r})\psi_j(\mathbf{r})\}_{1 \leq i \leq m, 1 \leq j \leq n} \quad (3)$$

so that

$$\varphi_i(\mathbf{r})\psi_j(\mathbf{r}) \approx \sum_{\mu=1}^{N_\mu} \zeta_\mu(\mathbf{r}) C_\mu^{ij} \quad (4)$$

where C_μ^{ij} 's are fitting coefficients. This can be implemented using rank revealing methods such as the singular value decomposition (SVD)⁴³ and the pivoted Cholesky factorization.⁴⁴

In the context of hybrid functional calculations above, we have $m = N_e$ and $n \sim O(N_e)$. Hence, the density fitting procedure compresses $mn \sim O(N_e^2)$ functions into a much smaller set of auxiliary functions $\{\zeta_\mu(\mathbf{r})\}_{\mu=1}^{N_\mu}$. Numerical results indicate that it is often sufficient to choose $N_\mu = cN_e$, where c is a small constant that we refer to as a rank parameter. This parameter determines the computational accuracy of the decomposition (eq 4).

In the standard density fitting procedure, the fitting coefficient tensor $\{C_\mu^{ij}\}$ is treated as a three way tensor and is often obtained through a least-squares fitting procedure. The storage cost of $\{C_\mu^{ij}\}$ is $O(N_e^3)$, and the computational cost of density fitting typically scales as $O(N_e^4)$. When a large basis set such as the plane-wave basis set is used, both the storage and the computational cost can be prohibitively high. As a result, density fitting is rarely used for this type of basis set unless additional locality constraints are enforced.³⁹

2.2. Interpolative Separable Density Fitting Decomposition. In order to reduce the complexity of the density fitting

method, the key is to find a more efficient treatment for the three-way fitting coefficient tensor. This has been achieved by the tensor hypercontraction (THC) method^{41,42} and the interpolative separable density fitting (ISDF) method.³⁴ Both methods use the following compression format:

$$\varphi_i(\mathbf{r})\psi_j(\mathbf{r}) \approx \sum_{\mu=1}^{N_\mu} \zeta_\mu(\mathbf{r}) \varphi_i(\hat{\mathbf{r}}_\mu) \psi_j(\hat{\mathbf{r}}_\mu) \quad (5)$$

Here, $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$ is a subset of real space grid points $\{\mathbf{r}_i\}_{i=1}^{N_g}$ on which the orbitals are evaluated. We refer to $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$ as the interpolation points and $\{\zeta_\mu(\mathbf{r})\}_{\mu=1}^{N_\mu}$ sampled on $\{\mathbf{r}_i\}_{i=1}^{N_g}$ the interpolation vectors. Since the term “interpolative separable” captures clearly the relation between the format (eq 5) and standard density fitting formats, with some abuse of terminology, we also refer to the format (eq 5) as the “ISDF format” or the “ISDF decomposition”.

Comparing eq 4 with eq 5, we see that the ISDF decomposition is a special form of a density fitting decomposition, where the fitting coefficient tensor $C_\mu^{ij} = \varphi_i(\hat{\mathbf{r}}_\mu) \psi_j(\hat{\mathbf{r}}_\mu)$ is given explicitly *without additional computation*. Hence, the storage cost is reduced from $O(N_e^3)$ to $O(N_e^2)$, and the ISDF decomposition is potentially suitable for calculations with a large basis set. The reason why such decomposition can be expected can be understood from the perspective of interpolation. Indeed, if $\{\hat{\mathbf{r}}_\mu\}$ is a set of grid points in the real space, let $\zeta_\mu(\mathbf{r})$ be the Lagrange interpolation function on these grid points satisfying

$$\zeta_\mu(\hat{\mathbf{r}}_{\mu'}) = \begin{cases} 1, & \mu = \mu', \\ 0, & \text{otherwise} \end{cases} \quad (6)$$

then the ISDF decomposition would become sufficiently accurate as one systematically refines the set $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$. In the worst case, all grid points are selected and $N_\mu = N_g$. Since $N_g \sim O(N_e)$, the asymptotic number of interpolation vectors is still smaller than $O(N_e^2)$ even in such scenario. Furthermore, if both $\{\varphi_i(\mathbf{r})\}$ and $\{\psi_j(\mathbf{r})\}$ are sets of sufficiently smooth functions, the number of points N_μ can be expected to be much smaller than N_g .

The THC and ISDF methods differ in terms of the procedure for finding the interpolation points and interpolation vectors and hence the computational complexity. For the THC method, the interpolation points are determined through a quadrature rule, e.g., uniform grid points or grid points from a Gauss-Hermite quadrature rule. However, the focus of THC is not to identify the interpolation vectors but to find an efficient approximation scheme for the four-way Coulomb integral tensor, which can be written in the current context as

$$V_{ijkl} = \int K(\mathbf{r}, \mathbf{r}') \varphi_i(\mathbf{r}) \psi_j(\mathbf{r}) \varphi_k(\mathbf{r}') \psi_l(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \quad (7)$$

Using the ISDF format, we have

$$V_{ijkl} \approx \sum_{\mu, \nu=1}^{N_\mu} \varphi_i(\hat{\mathbf{r}}_\mu) \psi_j(\hat{\mathbf{r}}_\mu) M_{\mu, \nu} \varphi_k(\hat{\mathbf{r}}_\nu) \psi_l(\hat{\mathbf{r}}_\nu) \quad (8)$$

Here, the coefficient matrix is given by the interpolation vectors

$$M_{\mu, \nu} = \int K(\mathbf{r}, \mathbf{r}') \zeta_\mu(\mathbf{r}) \zeta_\nu(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \quad (9)$$

On the other hand, if V_{ijkl} is already known, it is possible to not explicitly evaluate the interpolation vectors but directly obtain $M_{\mu, \nu}$ using the following equation:

$$\begin{aligned} & \sum_{ijkl} \varphi_i(\hat{\mathbf{r}}_\mu) \psi_j(\hat{\mathbf{r}}_\mu) V_{ijkl} \varphi_k(\hat{\mathbf{r}}_\nu) \psi_l(\hat{\mathbf{r}}_\nu) \\ & \approx \sum_{\mu', \nu'=1}^{N_\mu} \left(\sum_i \varphi_i(\hat{\mathbf{r}}_\mu) \varphi_i(\hat{\mathbf{r}}_{\mu'}) \right) \left(\sum_j \psi_j(\hat{\mathbf{r}}_\mu) \psi_j(\hat{\mathbf{r}}_{\mu'}) \right) M_{\mu', \nu'} \left(\sum_k \varphi_k(\hat{\mathbf{r}}_\nu) \varphi_k(\hat{\mathbf{r}}_{\nu'}) \right) \left(\sum_l \psi_l(\hat{\mathbf{r}}_\nu) \psi_l(\hat{\mathbf{r}}_{\nu'}) \right) \end{aligned} \quad (10)$$

Equation 10 can be understood as a least-squares fitting to eq 8. The dominant computational cost comes from the evaluation of the left-hand side, which scales as $O(N_e^5)$ if we do not assume any structure on the four way tensor V_{ijkl} . Hence, the strategy of THC is not suitable for large basis set calculations.

The ISDF method³⁴ uses a different approach to construct the decomposition (eq 5). Instead of using a quadrature rule, it uses a randomized QR factorization with column pivoting (QRCP) procedure⁴⁵ to find the interpolation points for the product pairs $\{\varphi_i(\mathbf{r})\psi_j(\mathbf{r})\}$, which can be potentially much more compact than a set of universal quadrature points. The interpolation vectors can also be deduced from the QRCP decomposition in the same calculation. The computational cost of the ISDF decomposition is only $O(N_e^3)$ and hence is suitable for large basis set calculations. It should be pointed out that the ISDF decomposition yields a decomposition for the Coulomb integral tensor once the coefficient matrix in eq 9 is computed, but the reverse statement is not true. Some of the key steps of the randomized QRCP procedure are discussed in Section 3.

As an example, we consider a water molecule (four occupied bands $N_{\text{band}} = 4$) (Figure 1(a)) in a $10 \text{ \AA} \times 10 \text{ \AA} \times 10 \text{ \AA}$ box, with




Figure 1. (a) Electron density (yellow isosurfaces) and (b) interpolation points (green squares) $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$ ($N_\mu = 8$) selected from the real space grid points $\{\mathbf{r}_i\}_{i=1}^{N_g}$ ($N_g = 66^3$) for a water molecule in a $10 \text{ \AA} \times 10 \text{ \AA} \times 10 \text{ \AA}$ box. The white and red balls denote hydrogen and oxygen atoms, respectively.

the kinetic energy cutoff $E_{\text{cut}} = 60$ hartree. This corresponds to a real space grid of size $66 \times 66 \times 66$ (i.e., $N_g = 66^3$). Figure 1(b) shows how $N_\mu = 8$ interpolation points are distributed in the real space. As we can see, they are closer to the oxygen atom than to hydrogen atoms. This is consistent with the distribution of the electron density. Figure 1(b) also indicates that instead of using a uniform sampling grid, and it may be more advantageous to select the interpolation points adaptively from QRCP, especially for electron densities with an inhomogeneous spatial distribution.

3. ISDF DECOMPOSITION FOR HYBRID FUNCTIONAL CALCULATIONS

Although the ISDF decomposition significantly reduces the number of Poisson-like equations to be solved, the complexity of the randomized QRCP method used to find the interpolation points is still $O(N_e^3)$, which is comparable to the cost of computing $V_X[\{\varphi_i\}]\psi_j$. Hence, at first glance, the ISDF

decomposition may not lead to much efficiency gain in the context of hybrid functional calculations.

Below, we introduce a new method to compute the ISDF decomposition, which separates the treatment of the interpolation vectors and the interpolation points. More specifically, we use the randomized QRCP decomposition only to find the interpolation points and use a least-squares fitting procedure to efficiently find the interpolation vectors without randomization. Although it is more expensive to use the randomized QRCP decomposition, this step only needs to be invoked once in the entire hybrid functional calculation. The calculation of the interpolation vectors still scales as $O(N_e^3)$, but the preconstant is significantly smaller. We find that such choice balances the efficiency and accuracy, especially when the decomposition needs to be repeatedly used such as in a SCF iteration procedure. We also discuss how to combine the ISDF decomposition with the adaptively compressed exchange operator (ACE) formulation in a numerically stable way, so that the overall computational cost can be significantly reduced.

3.1. Finding the Interpolation Vectors. We first discuss how to find the interpolation vectors assuming that the interpolation points $\{\hat{\mathbf{r}}_\mu\}$ are given.

Note that eq 5 can be written as

$$Z - \Theta C = 0 \quad (11)$$

where each column of Z is defined by eq 3 sampled on real space grids $\{\mathbf{r}_i\}_{i=1}^{N_g}$. $\Theta = [\zeta_1, \zeta_2, \dots, \zeta_{N_\mu}]$ contains the interpolating vectors, and the column of C indexed by (i, j) is given by

$$[\varphi_i(\hat{\mathbf{r}}_1)\psi_j(\hat{\mathbf{r}}_1), \dots, \varphi_i(\hat{\mathbf{r}}_\mu)\psi_j(\hat{\mathbf{r}}_\mu), \dots, \varphi_i(\hat{\mathbf{r}}_{N_\mu})\psi_j(\hat{\mathbf{r}}_{N_\mu})]^T$$

Equation 11 is an overdetermined linear system with respect to the interpolation vectors Θ . One possible way to solve the overdetermined system is to impose the Galerkin condition

$$(Z - \Theta C)C^T = 0 \quad (12)$$

It follows that the interpolating vectors can be obtained from

$$\Theta = ZC^T(CC^T)^{-1} \quad (13)$$

Note that the solution given by eq 13 is a least-squares approximation to the solution of eq 5. This is similar to that in the THC method, but the important difference is that the least-squares fitting is applied to the Z matrix, which is the key to reduce the complexity.

It may appear that the matrix–matrix multiplications ZC^T and CC^T take $O(N_e^4)$ operations because the size of Z is $N_g \times (N_\mu n)$ and the size of C is $N_\mu \times (N_\mu n)$. However, both multiplications can be carried out with fewer operations due to the separable structure of Z and C . It follows from the identity

$$\sum_{i,j} \varphi_i \psi_j = \left(\sum_i \varphi_i \right) \left(\sum_j \psi_j \right)$$

that we may rewrite the (k, μ) th element of ZC^T (also denoted by $P_{k\mu}^{\varphi\psi}$) as

$$P_{k\mu}^{\varphi\psi} P^{\varphi\psi} \equiv e_k^T Z C^T e_\mu = P^\varphi(\mathbf{r}_k, \hat{\mathbf{r}}_\mu) P^\psi(\mathbf{r}_k, \hat{\mathbf{r}}_\mu) \quad (14)$$

where $P^\varphi(\mathbf{r}, \hat{\mathbf{r}}_\mu)$ and $P^\psi(\mathbf{r}, \hat{\mathbf{r}}_\mu)$ can be viewed as columns of (quasi) density matrices defined as

$$P^\varphi(\mathbf{r}, \mathbf{r}') = \sum_{i=1}^m \varphi_i(\mathbf{r}) \varphi_i(\mathbf{r}'), \quad P^\psi(\mathbf{r}, \mathbf{r}') = \sum_{j=1}^n \psi_j(\mathbf{r}) \psi_j(\mathbf{r}') \quad (15)$$

Similarly, we can rewrite the (ν, μ) th element of CC^T as

$$e_\nu^T CC^T e_\mu = P^\varphi(\hat{\mathbf{r}}_\nu, \hat{\mathbf{r}}_\mu) P^\psi(\hat{\mathbf{r}}_\nu, \hat{\mathbf{r}}_\mu) \quad (16)$$

Because both P^φ and P^ψ matrices can be evaluated with $O(N_e^3)$ floating point operations, and the multiplications in eqs 14 and 16 are pointwise multiplications (Hadamard products) consuming only $O(N_e^2)$ floating point operations, the computational complexity for computing the interpolation vectors is $O(N_e^3)$.

The ISDF decomposition can be readily used to accelerate the computation in eq 2. Substituting eq 5 into eq 2 yields

$$\begin{aligned} (V_X[\{\varphi_i\}]\psi_j)(\mathbf{r}) &= - \sum_{i=1}^{N_i} (\varphi_i(\mathbf{r}) \int K(\mathbf{r}, \mathbf{r}') \varphi_i(\mathbf{r}') \psi_j(\mathbf{r}') d\mathbf{r}') \\ &\approx - \sum_{i=1}^{N_i} \sum_{\mu=1}^{N_\mu} \varphi_i(\mathbf{r}) \left(\int K(\mathbf{r}, \mathbf{r}') \zeta_\mu(\mathbf{r}') d\mathbf{r}' \right) \\ &\quad \varphi_i(\hat{\mathbf{r}}_\mu) \psi_j(\hat{\mathbf{r}}_\mu) \\ &\equiv - \sum_{\mu=1}^{N_\mu} P^\varphi(\mathbf{r}, \hat{\mathbf{r}}_\mu) V_\mu^\zeta(\mathbf{r}) \psi_j(\hat{\mathbf{r}}_\mu), \end{aligned} \quad (17)$$

where $P^\varphi(\mathbf{r}, \hat{\mathbf{r}}_\mu)$ is given by the quasi density matrix defined in eq 15, and

$$V_\mu^\zeta(\mathbf{r}) \equiv \int K(\mathbf{r}, \mathbf{r}') \zeta_\mu(\mathbf{r}') d\mathbf{r}' \quad (18)$$

can be carried out as the solution to a Poisson-like equation. The ISDF decomposition reduces the total number of Poisson-like equations to be solved from $O(N_e^2)$ to $N_\mu \sim O(N_e)$.

3.2. Finding the Interpolation Points. The problem for finding a suitable set of interpolation points $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$ can be formulated as the following linear algebra problem. Consider the discretized matrix Z of size as an $N_g \times (mn)$ matrix Z and find N_μ rows of Z so that the rest of the rows of Z can be approximated by the linear combination of the selected N_μ rows. This is called an interpolative decomposition,⁴⁵ and a standard method to achieve such a decomposition is the QR factorization with a column pivoting (QRCP) procedure⁴⁵ as

$$Z^T \Pi = QR \quad (19)$$

where Z^T is the transpose of Z , Q is an $mn \times N_g$ matrix that has orthonormal columns, R is an upper triangular matrix, and Π is a permutation matrix chosen so that the magnitude of the diagonal elements of R form a nonincreasing sequence. The magnitude of each diagonal element R indicates how important the corresponding column of the permuted Z^T is and whether the corresponding grid point should be chosen as an interpolation point. The QRCP factorization can be terminated when the $(N_\mu + 1)$ th diagonal element of R becomes less than a predetermined threshold. The leading N_μ columns of the permuted Z^T are considered to be linearly independent numerically. The

corresponding grid points are chosen as the interpolation points. The indices for the chosen interpolation points $\{\hat{\mathbf{r}}_\mu\}$ can be obtained from indices of the nonzero entries of the first N_μ columns of the permutation matrix Π . However, the storage requirement for the matrix Z is $O(N_e^3)$ and the computational cost associated with a standard QRCP procedure is $O(N_e^4)$, which is not so appealing.

The key idea used in ref 34 to lower the cost of QRCP is to use a random matrix to subsample columns of the matrix Z to form a smaller matrix \tilde{Z} of size $N_g \times \tilde{N}_\mu$, where \tilde{N}_μ is only slightly larger than N_μ . It can be shown that under some mild assumptions the reduction in the number of columns in a randomly subsampled Z does not have much impact the quality of the interpolation points $\{\hat{\mathbf{r}}_\mu\}$. However, we do not present the theoretical analysis here but merely describe the algorithmic ingredients. We refer readers to refs 46 and 47 for more detailed analysis of randomized sampling methods.

There are a number of ways to subsample columns of Z . Instead of using the subsampled Fourier transform as in ref 34, here we choose two orthogonalized Gaussian matrices G^φ, G^ψ of size $m \times p$ and $n \times p$, respectively, where $p \sim O(\sqrt{N_\mu})$ is chosen to satisfy $\tilde{N}_\mu = p^2$, and use them to construct a set of subsampled products defined by

$$\tilde{Z}_{\alpha\beta}(\mathbf{r}) = \left(\sum_{i=1}^m \varphi_i(\mathbf{r}) G_{i\alpha}^\varphi \right) \left(\sum_{j=1}^n \psi_j(\mathbf{r}) G_{j\beta}^\psi \right), \quad 1 \leq \alpha, \beta \leq p \quad (20)$$

The corresponding discretized matrix \tilde{Z} is of size $N_g \times \tilde{N}_\mu$. Applying the QRCP procedure to \tilde{Z} yields

$$\tilde{Z}^T \Pi = QR \quad (21)$$

where the interpolation points $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$ are given by the first N_μ columns of the permutation matrix Π . Since the random matrices G^φ and G^ψ are only applied to $\{\varphi_i\}$ and $\{\psi_j\}$, respectively, the storage cost for \tilde{Z} is $O(N_e^2)$, and the computational cost for generating \tilde{Z} , which is dominated by the cost of matrix–matrix multiplications is $O(N_e^{2.5})$. The reduced matrix size allows the computational cost of the QRCP procedure to be reduced to $O(N_e^3)$. Since the QRCP algorithm has been implemented in standard linear algebra software packages such as LAPACK and ScaLAPACK,⁴⁸ the implementation and parallelization of ISDF is relatively straightforward.

Our numerical results indicate that the cost of the randomized QRCP method can be comparable to that of computing $V_X[\{\varphi_i\}]\psi_j$. However, while the interpolation vectors depends sensitively on the shape of the input orbitals $\{\varphi_i\}$ and $\{\psi_j\}$ and need to be recomputed whenever they are updated, the interpolation points are much less sensitive to small changes of the orbitals. This is because the significance of the interpolation points is only to indicate which columns of Z^T are important. In practice, we find it sufficient to determine the interpolation points at the beginning of the hybrid functional calculations starting from a set of KS orbitals from a converged GGA calculation and to use such interpolation points throughout the SCF iterations.

3.3. Combining with the Adaptively Compressed Exchange Operator Formulation. The ISDF decomposition can be combined with the recently developed adaptively compressed exchange operator (ACE) formulation^{30,31} to further reduce the cost of hybrid functional KSDFT calculations.

In the ACE formulation, the operator V_X is replaced by a rank- n operator V_X^{ACE} that satisfies

$$V_X \psi_j = V_X^{\text{ACE}} \psi_j$$

for $j = 1, 2, \dots, n$, where n is much less than the total number of grid points N_g . The operator V_X^{ACE} can be written in the form

$$V_X^{\text{ACE}}(\mathbf{r}, \mathbf{r}') = - \sum_{k=1}^n \xi_k(\mathbf{r}) \xi_k(\mathbf{r}') \quad (22)$$

for some vectors $\{\xi_k\}$. The application of V_X^{ACE} to a set of orbitals resembles the application of a nonlocal pseudopotential operator. This does not require any Poisson solves and is much cheaper than applying V_X to these orbitals in an iterative diagonalization procedure used to update the set of occupied orbitals $\{\varphi_i\}$ in the SCF iteration. However, the construction of V_X^{ACE} , which must be performed in each (outer) SCF iteration, still requires applying V_X to ψ_j to produce

$$W_j(\mathbf{r}) = (V_X[\{\varphi_i\}]\psi_j)(\mathbf{r}) \quad j = 1, \dots, n \quad (23)$$

The basis vectors $\{\xi_k\}$ that appear in eq 22 are obtained from W_j via

$$\xi_k(\mathbf{r}) = \sum_{j=1}^n W_j(\mathbf{r})(L^{-T})_{jk} \quad (24)$$

where L is the lower triangular Cholesky factor of the matrix M . The (i, j) th element of M is given by

$$M_{ij} = \int \psi_i(\mathbf{r}) W_j(\mathbf{r}) d\mathbf{r} \quad (25)$$

Because M is a symmetric negative definite matrix of size n , a Cholesky factorization can be used to decompose $-M$ as $-M = LL^T$, where L is unit lower triangular. Since eq 23 is performed only in each outer iteration, which is less frequent than applying V_X to ψ_j in every step of the diagonalization procedure, the use of ACE significantly reduces the cost of hybrid functional calculation, without requiring any approximation to the computation of the W_j 's.

The ISDF decomposition can be readily used to accelerate the computation of W_j 's in eq 23. However, straightforward computation using in eq 25 may result in an M matrix that is not symmetric, let alone being negative definite. To see this, we combine eq 17 and eq 25 to obtain

$$M_{ij} = \sum_{\mu=1}^{N_\mu} \left(\int \psi_i(\mathbf{r}) P^\varphi(\mathbf{r}, \mathbf{r}_\mu) V_\mu^\zeta(\mathbf{r}) d\mathbf{r} \right) \psi_j(\mathbf{r}_\mu)$$

which may be different from

$$M_{ji} = \sum_{\mu=1}^{N_\mu} \left(\int \psi_j(\mathbf{r}) P^\varphi(\mathbf{r}, \mathbf{r}_\mu) V_\mu^\zeta(\mathbf{r}) d\mathbf{r} \right) \psi_i(\mathbf{r}_\mu)$$

The lack of symmetry may result in numerical stability problems in the subsequent Cholesky factorization of M . In order to overcome this problem, we can apply the ISDF decomposition in a symmetric fashion as follows. Note that

$$M_{ij} = \sum_{l=1}^{N_\zeta} \int \psi_i(\mathbf{r}) \varphi_l(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \varphi_l(\mathbf{r}') \psi_j(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \quad (26)$$

Hence, we can use ISDF to expand both the $\varphi_l(\mathbf{r})\psi_i(\mathbf{r})$ and $\varphi_l(\mathbf{r}')\psi_j(\mathbf{r}')$ pairs in terms of the interpolating vector $\{\zeta_\mu\}$ to obtain

$$\begin{aligned} M_{ij} &\approx \sum_{l=1}^{N_\zeta} \sum_{\mu,\nu=1}^{N_\mu} \left(\int \zeta_\mu(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \zeta_\nu(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \right) \varphi_l(\mathbf{r}_\mu) \varphi_l(\mathbf{r}_\nu) \psi_i(\mathbf{r}_\mu) \psi_j(\mathbf{r}_\nu) \\ &= \sum_{\mu,\nu=1}^{N_\mu} \left(\int \zeta_\mu(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \zeta_\nu(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \right) P^\varphi(\mathbf{r}_\mu, \mathbf{r}_\nu) \psi_i(\mathbf{r}_\mu) \psi_j(\mathbf{r}_\nu) \\ &= \sum_{\mu,\nu=1}^{N_\mu} \psi_i(\mathbf{r}_\mu) \tilde{M}_{\mu\nu} \psi_j(\mathbf{r}_\nu), \end{aligned} \quad (27)$$

where

$$\tilde{M}_{\mu\nu} = \left(\int \zeta_\mu(\mathbf{r}) K(\mathbf{r}, \mathbf{r}') \zeta_\nu(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \right) P^\varphi(\mathbf{r}_\mu, \mathbf{r}_\nu) \quad (28)$$

Since both the first and second factors on the right-hand side of eq 28 are symmetric, $\tilde{M}_{\mu\nu}$ is clearly symmetric, which guarantees the symmetry and the definiteness of M defined in eq 27.

In the symmetric formulation given by eq 27, M is automatically a symmetric negative definite matrix. Therefore, the Cholesky factorization of M yields the ACE operator according to eq 22. Hence we refer to the combined method using ACE and ISDF as the ACE-ISDF method.

4. PARALLEL IMPLEMENTATION

In this section, we demonstrate an efficient parallel implementation of the ACE-ISDF method for hybrid functional calculations in a planewave basis set.

We denote the matrix of the discretized orbitals by $\Phi = [\varphi_1, \dots, \varphi_{N_\zeta}]$, and $\Psi = [\psi_1, \dots, \psi_n]$. When P_e processors are used, these orbitals are stored using the 1D column cyclic partition as shown in Figure 2(a), so that the application of the Hamiltonian operator (excluding the exchange part) to the orbitals Ψ can be easily parallelized. In particular, the Laplacian operator can be

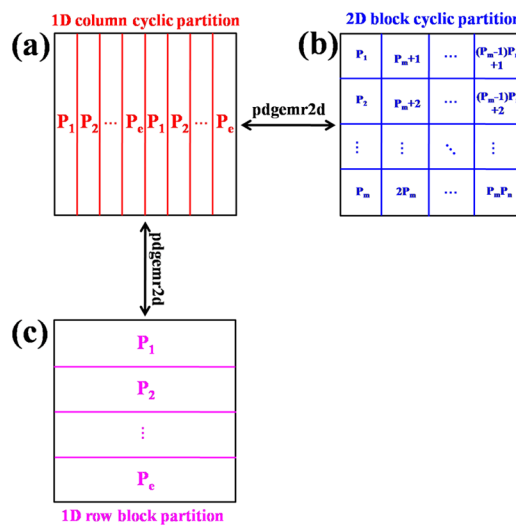


Figure 2. Three different types of data partition for the matrix used in the ACE-ISDF formulation for hybrid density functional calculations: (a) 1D column cyclic partition ($1 \times P_e$ MPI processor grid), (b) 2D block cyclic partition ($P_m \times P_n$ MPI processor grid), and (c) 1D row block partition ($P_e \times 1$ MPI processor grid). P_e is total computational cores used in the ACE-ISDF formulation and $P_m \times P_n = P_e$.

applied through the use of a sequential fast Fourier transformation (FFT) library. Moreover, the application of the local and nonlocal pseudopotentials in the real space representation is also rather straightforward. The application of the ACE operator to Ψ involves two matrix–matrix multiplication operations and can be done most efficiently by using a row-based partition shown in Figure 2(c) (see, e.g., ref 31 for more details). However, in the ACE-ISDF procedure, a 2D block cyclic partition shown in Figure 2(b) is the most efficient data distribution scheme for performing a number of dense linear algebra operations such as QRCP implemented in the ScaLAPACK software package.⁴⁸ The conversion among different data storage formats is performed using the pdgemr2d subroutine in the ScaLAPACK software package.⁴⁸

More specifically, we describe the various quantities in the ACE-ISDF method and the corresponding storage formats in Figure 3. Starting from Φ and Ψ distributed in the column cyclic

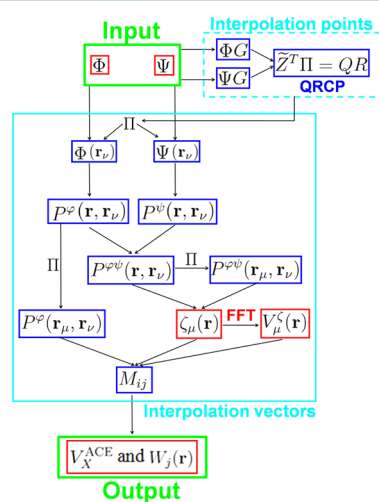


Figure 3. Flowchart of the ACE-ISDF formulation for constructing the ACE exchange operator in PWDFT. Red and blue boxes, respectively, represent 1D column cyclic partition and 2D block cyclic partition for the matrices used in the ACE formulation. Once the interpolation points $\{\mathbf{r}_\mu\}_{\mu=1}^{N_\mu}$ are obtained from QRCP at the first outer SCF iteration, the interpolation vectors $\{\zeta_\mu(\mathbf{r})\}_{\mu=1}^{N_\mu}$ can be updated directly from the input vectors at other outer SCF iterations.

partition, we first transform these matrices into the 2D block cyclic partition to generate the \tilde{Z} matrix. We then perform the QRCP procedure to obtain the permutation matrix Π . The interpolation points $\{\hat{\mathbf{r}}_\mu\}_{\mu=1}^{N_\mu}$ are retrieved from the permutation matrix Π . This is a small vector of size N_μ and are shared among all processors.

In order to construct the interpolation vectors, we distribute the columns of the quasi density matrices $P^\phi(\mathbf{r}_\nu, \hat{\mathbf{r}}_\mu)$ and $P^\psi(\mathbf{r}_\nu, \hat{\mathbf{r}}_\mu)$ in a 2D block cyclic fashion so that the matrix ZC^T defined in eq 14 can be evaluated in parallel via local Hadamard multiplications. The matrix CC^T that appears in eq 13 can be obtained by simply subsampling rows of ZC^T .

The resulting discretized interpolation vectors $\Omega = [\zeta_1, \dots, \zeta_{N_\mu}]$ can be obtained by calling a ScaLAPACK linear equation solver. The 2D block cyclically distributed solution is redistributed and partitioned by a 1D column partition for computing the Coulomb-like potential for the interpolation vectors $V^\zeta = [V_1^\zeta, \dots, V_{N_\mu}^\zeta]$ as in eq 18. $\{V^\zeta\}$ are then converted back to the 2D block cyclic distribution pattern. Finally, the $\{W_j\}$ in eq 23 can be

computed using matrix–matrix multiplication in the 2D partition as in eq 17 and then converted to 1D row partition.

In order to implement the symmetric formulation for the M matrix in eq 27 as required by ACE, we form the matrix \tilde{M} in eq 28 in parallel within a 2D block cyclic distribution scheme, and the M matrix in eq 27 can be obtained by two parallel matrix–matrix multiplication calls. Finally, we perform a parallel Cholesky factorization of M on the 2D block cyclic grid, and the ACE vectors $\Xi = [\xi_1, \dots, \xi_n]$ are partitioned by rows on a 1D processor grid. This gives the V_X^{ACE} implicitly and can be readily used in subsequent iterations.

5. NUMERICAL RESULTS

We demonstrate the performance of the ACE-ISDF method using the DGDFT (Discontinuous Galerkin Density Functional Theory) software package.^{49–53} DGDFT is a massively parallel electronic structure software package designed for large scale DFT calculations involving up to tens of thousands of atoms. It includes a self-contained module called PWDFT for performing plane-wave-based electronic structure calculations (mostly for benchmarking and validation purposes). We implemented the ACE-ISDF method in PWDFT. We use Message Passing Interface (MPI) to handle data communication and the Hartwigsen–Goedecker–Hutter (HGH) norm-conserving pseudopotential.⁵⁴ All calculations use the HSE06 functional.⁷ All calculations are carried out on the Edison systems at the National Energy Research Scientific Computing Center (NERSC). Each node consists of two Intel “Ivy Bridge” processors with 24 cores in total and 64 gigabyte (GB) of memory. Our implementation only uses MPI. The number of cores is equal to the number of MPI ranks used in the simulation.

In this section, we demonstrate the accuracy of the ACE-ISDF method for accelerating hybrid functional calculations using a bulk silicon system Si_{216} and a disordered system $\text{Al}_{176}\text{Si}_{24}$ as shown in Figure 4(a) and (b), respectively. The Si_{216} system is semiconducting with an energy gap of $E_{\text{gap}} = 1.45$ eV, and the $\text{Al}_{176}\text{Si}_{24}$ system is metallic with $E_{\text{gap}} < 0.1$ eV. The density of states of the two systems are shown in Figure 4(c) and (d), respectively. All systems are closed shell systems, and the number




Figure 4. Atomic structures of (a) Si_{216} and (b) $\text{Al}_{176}\text{Si}_{24}$. The yellow and pink balls denote silicon and aluminum atoms, respectively. The total densities of states (DOS) of (c) semiconducting Si_{216} and (d) metallic $\text{Al}_{176}\text{Si}_{24}$. The Fermi levels are marked by green dotted lines.

Table 1. Accuracy of Hybrid Functional Calculations (HSE06) Obtained by ACE-ISDF Method as a Function of Rank Parameter c for Si_{216} and $\text{Al}_{176}\text{Si}_{24}$ ^a

ACE-ISDF: Semiconducting Si_{216} ($N_{\text{band}} = 432$)						
c	E_{VBM}	E_{CBM}	E_{gap}	ΔE_{HF}	ΔE	ΔF
4.0	6.5303	8.4367	1.9064	3.53×10^{-03}	4.13×10^{-03}	1.49×10^{-03}
5.0	6.5923	8.3652	1.7729	2.37×10^{-03}	2.75×10^{-03}	1.30×10^{-03}
6.0	6.6786	8.2535	1.5749	7.34×10^{-04}	1.06×10^{-03}	1.00×10^{-03}
7.0	6.6893	8.1554	1.4661	3.56×10^{-04}	4.61×10^{-04}	6.61×10^{-04}
8.0	6.6763	8.1341	1.4578	1.16×10^{-04}	1.64×10^{-04}	3.05×10^{-04}
9.0	6.6652	8.1164	1.4512	4.74×10^{-05}	7.50×10^{-05}	1.48×10^{-04}
10.0	6.6565	8.1085	1.4520	2.33×10^{-05}	4.11×10^{-05}	1.09×10^{-04}
12.0	6.6487	8.1001	1.4514	7.27×10^{-06}	1.54×10^{-05}	5.88×10^{-05}
16.0	6.6467	8.0959	1.4492	1.57×10^{-06}	2.92×10^{-06}	1.67×10^{-05}
20.0	6.6466	8.0942	1.4476	5.40×10^{-07}	7.87×10^{-07}	5.54×10^{-06}
ACE	6.6467	8.0934	1.4466	0.00×10^{-00}	0.00×10^{-00}	0.00×10^{-00}
ACE-ISDF: Metallic $\text{Al}_{176}\text{Si}_{24}$ ($N_{\text{band}} = 312$)						
c	E_{VBM}	E_{CBM}	E_{gap}	ΔE_{HF}	ΔE	ΔF
4.0	7.8907	7.9963	0.1056	7.20×10^{-03}	8.06×10^{-03}	8.96×10^{-03}
5.0	7.8173	7.9103	0.0930	3.46×10^{-03}	3.76×10^{-03}	3.96×10^{-03}
6.0	7.7810	7.8833	0.1023	1.33×10^{-03}	1.69×10^{-03}	2.32×10^{-03}
7.0	7.7805	7.8742	0.0937	5.97×10^{-04}	6.41×10^{-04}	1.60×10^{-03}
8.0	7.7717	7.8710	0.0993	1.90×10^{-04}	2.03×10^{-04}	5.55×10^{-04}
9.0	7.7719	7.8710	0.0991	6.92×10^{-05}	7.44×10^{-05}	3.10×10^{-04}
10.0	7.7713	7.8699	0.0986	3.20×10^{-05}	3.55×10^{-05}	1.53×10^{-04}
12.0	7.7712	7.8695	0.0983	1.16×10^{-05}	1.38×10^{-05}	9.23×10^{-05}
16.0	7.7704	7.8698	0.0994	3.26×10^{-06}	4.43×10^{-06}	4.34×10^{-05}
20.0	7.7703	7.8695	0.0992	1.27×10^{-06}	1.93×10^{-06}	2.18×10^{-05}
ACE	7.7701	7.8695	0.0994	0.00×10^{-00}	0.00×10^{-00}	0.00×10^{-00}

^aThe unit for VBM (E_{VBM}), CBM (E_{CBM}), and the energy gap E_{gap} is eV. The unit for the error in the Hartree–Fock exchange energy ΔE_{HF} and the total energy ΔE is the Hartree/atom, and the unit for the error in atomic forces ΔF is Hartree/Bohr. We use results from the ACE-enabled hybrid functional calculations as the reference.

of occupied bands is $N_{\text{band}} = N_e/2$. We include two unoccupied bands for computing the energy gap in the systems. We show the parallel scalability of our implementation using seven bulk silicon systems with 64 to 4096 atoms.³¹ Finally, we use the ACE-ISDF method in a hybrid DFT calculation to study the electronic structure of vacancy defect in a silicon supercell that contains 1000 Si atoms.

5.1. Accuracy. In the previous work,^{30,31} we demonstrated that the ACE formulation can significantly accelerate hybrid functional calculations without loss of accuracy. Hence, the results from the ACE calculation is used as the baseline for comparison in assessing the accuracy of the ACE-ISDF method. Table 1 shows the convergence of the ACE-ISDF method as a function of the rank parameter c for the Si_{216} and $\text{Al}_{176}\text{Si}_{24}$ systems. We measure the accuracy in terms of the valence band maximum (VBM) energy level, the conduction band minimum (CBM) energy levels, the energy gap, the Hartree–Fock (HF) exchange energy, and the total energy as well as the atomic force. The energy cutoff E_{cut} is set to 20 hartree. We define the errors in the HF energy, the total energy and the atomic force respectively as

$$\Delta E_{\text{HF}} = (E_{\text{HF}}^{\text{ACE-ISDF}} - E_{\text{HF}}^{\text{ACE}})/N_A$$

$$\Delta E = (E^{\text{ACE-ISDF}} - E^{\text{ACE}})/N_A$$

$$\Delta F = \max_I |F_I^{\text{ACE-ISDF}} - F_I^{\text{ACE}}|$$

Here, N_A is the total number of atoms, and I is the atom index.

Our calculations show that the ACE-ISDF method can produce highly accurate results with a moderate rank parameter c . (Recall that the rank of the ISDF approximation is $N_\mu = cN_{\text{band}}$.) The accuracy of the approximation can be improved systematically by increasing the rank parameter c . When a relatively small c value (e.g., $c = 6.0$) is used, the error in the total energy of both the Si_{216} and the $\text{Al}_{176}\text{Si}_{24}$ systems is already below that required to reach the chemical accuracy of 1 kcal/mol (1.6 Hartree/atom).⁵⁵ For Si_{216} , the errors in the HF energy, the total energy and the atomic force systematically decrease from $O(10^{-3})$ to $O(10^{-6} \sim 10^{-7})$ when c is adjusted from 6.0 to 20.0. We note that the total energy convergence with respect to the rank parameter is similar between Si_{216} and $\text{Al}_{176}\text{Si}_{24}$. The fact that the rank parameter c is independent of the band gap makes ISDF more attractive than linear scaling methods^{56–59} whose accuracy is controlled by the level of truncation in the density matrix, which in turn depends strongly on the band gap.

5.2. Efficiency. We demonstrate the efficiency of the ACE-ISDF method by showing its performance in a hybrid DFT calculation for a bulk silicon system with 1000 atoms ($N_{\text{band}} = 2000$) on 2000 cores. In each outer iteration, the cost of hybrid functional calculations consists of the cost for constructing the ACE operator (with or without ISDF) and the amount of work performed in the inner SCF iterations.

Table 2 shows the wall clock time spent in major components of the ACE-ISDF and ACE calculations. The main cost for constructing the ACE operator without using the ISDF decomposition is in the solution of $O(N_e^2)$ Poisson-like equations via FFTs. For this silicon system, the number of Poisson-like

Table 2. Wall Clock Time (in s) Spent in Components of ACE-ISDF- and ACE-Enabled Hybrid DFT Calculations Related to the Exchange Operator for Si₁₀₀₀ on 2000 Edison Cores at Different E_{cut} Levels^a

Si ₁₀₀₀		ACE-ISDF			ACE
E_{cut}	N_g	IP	IV	Other (FFT)	FFT
10	74 ³	50.22	11.13	8.56 (0.28)	101.10
20	104 ³	105.95	24.52	20.52 (1.17)	148.73
30	128 ³	222.36	40.67	32.88 (1.31)	302.98
40	148 ³	454.42	63.56	54.95 (3.08)	807.31

^aThe corresponding number of real space grid points used to represent the wavefunction is labeled by N_g . We use the rank parameter $c = 6.0$ in the ACE-ISDF calculation.

equations to be solved in each outer iteration is as large as $N_{\text{bands}}^2 = 4,000,000$. To show the detailed cost of constructing the ACE operator using the ISDF method, we report the timing measurements for selecting the interpolation points (IP), computing the interpolation vectors (IV) and other linear algebra operations and FFTs (labeled by “Other”). The IP selection is performed once only in the first outer SCF iteration. Computing the IVs constitutes a major part of the cost in the construction of the ACE operator in subsequent outer SCF iterations. The time spent in solving Poisson equations using FFTs, which we list in the parentheses next to time spent in the remaining parts of the ACE-ISDF calculation for comparison, is negligibly small. The reason that this cost is so small is that the use of ISDF significantly reduces the number of Poisson equations to be solved from $N_{\text{bands}}^2 = 4,000,000$ to $N_{\mu} = 12,000$ (when the rank parameter c is set to 6.)

When E_{cut} is set to 10 hartree and c is set to 6.0, the IP and IV computations take 50.22 and 11.13 s, respectively, in ACE-ISDF. The total amount of time spent in the construction of the ACE operator via ISDF in the first SCF iteration, which is 70 s (50.22 + 11.13 + 8.56), is already lower than that used to construct the ACE operator without ISDF (a procedure dominated by solving a larger number of Poisson-like equations via FFTs), which is roughly 101.10 s. In each subsequent outer SCF iteration, a total of 19.86 s (11.13 + 8.56) is used to construct the ACE operator via ISDF. This is already comparable to the time of one SCF iteration in GGA calculations, which is 17.89 s.

Note that for some complex systems, more inner SCF iterations might be required in each outer SCF iteration to reach convergence. For example, for the disordered Al₁₇₆Si₂₄ system, we need to use 14 inner SCF iterations per outer SCF iteration. As a result, the cost difference between ACE and ACE-ISDF is magnified, and ACE-ISDF is even more advantageous in these situation.

To illustrate the reduction of cost the ACE-ISDF scheme has achieved, we report that the average time spent in the construction and application of the exchange operator per outer SCF iteration in the conventional hybrid functional calculations is 1146.36 s, which is nearly 2 orders of magnitude higher than that used in ACE-ISDF. The large cost mainly comes from the fact that conventional hybrid functional calculations require solving $O(N_c^2)$ Poisson-like equations using FFTs in each step of an iterative diagonalization procedure (e.g., PPCG) when the exchange operator is applied to a set of $O(N_c)$ orbitals. For this system, on average 17 such operations need to be performed during each outer iteration. The ACE formulation reduces the cost by only requiring the exchange operator to be applied once

per outer iteration, and the ACE-ISDF method further reduces the cost for this single application of the exchange operator.

As we discussed earlier, since IP calculation is significantly more expensive, it is important to perform such a calculation only once in order to make ACE-ISDF efficient. We found that the Hartree–Fock exchange energy E_{HF} obtained from using a fixed set of IPs throughout the SCF iteration differs only slightly from that obtained by recalculating IPs in each outer SCF iteration (dynamic IP) as shown in Table 3 for Si₂₁₆ and Al₁₇₆Si₂₄. We can

Table 3. Comparison among Hartree–Fock Exchange Energies Computed by ACE Only and by ACE-ISDF with a Fixed or Changing Set of IPs^a

Systems	fixed IP in ACE-ISDF	dynamic IP in ACE-ISDF	ACE
Si ₂₁₆	−45.3225	−45.3250	−45.4835
Al ₁₇₆ Si ₂₄	−25.9112	−25.9093	−26.1745

^aThe rank parameter c used in ACE-ISDF is set to 6.0 for the Si₂₁₆ and Al₁₇₆Si₂₄ systems.

clearly see from the table that the difference between the first and the second columns is much smaller than the difference between the second and third columns. This shows that using fixed IP is well justified. Hence, in practice, we only use the QRCP decomposition in the first outer SCF iteration to select IPs, whereas the update of the basis vectors $\{\zeta_{\mu}(\mathbf{r})\}_{\mu=1}^{N_{\mu}}$ is performed in each subsequent outer SCF iterations.

Table 2 also shows that as E_{cut} is increased from 10 to 40 hartree the time spent in constructing the ACE operator increases from 101.10 to 807.31 s. The time spent in identifying the interpolation points via QRCP also increases from 50.22 to 454.42 s. However, the time used to compute the interpolation vectors only increases from 11.13 to 63.56 s. Given that the interpolating points only need to be selected once, the performance gain achieved by ISDF is more notable at a larger E_{cut} value. This feature is particularly attractive for calculations that requiring large kinetic energy cutoff, such as those involving transition metal oxides.

5.3. Parallel Scalability. To illustrate the strong parallel scalability of the ACE-ISDF method for large-scale hybrid DFT calculations, we report the change of the wallclock time in one outer SCF iteration with respect to the number of cores for the Si₁₀₀₀ system. The measured wallclock time includes time spent in one inner SCF iteration and in the ACE-ISDF operator construction. We also report the weak scaling of ACE-ISDF by showing the variation of the wallclock time with respect to the system size for a calculation that uses 8192 cores. These results are given in Figure 5 which shows that our implementation of the ACE-ISDF method scales nearly perfectly up to 2000 cores for the Si₁₀₀₀ system. It also shows ACE-ISDF scales well with respect to the system size (up to 4096 atoms) on 8192 cores.

5.4. Application to Vacancy Defect in Silicon. Silicon is one of the most important materials in industry due to its remarkable properties and a wide range of applications in electronics. However, the presence of defects can significantly affect these properties. Furthermore, defects also can be extremely useful for designing innovative electronic devices.^{60,61}

Therefore, accurate description of the electronic structures of silicon defects⁶² is required to examine their effects on the electronic devices.

It is well known that DFT calculations based on LDA and GGA functionals are not reliable in predicting electronic structures of nanosystems.⁶³ In particular, such DFT calculations




Figure 5. (a) The change of wallclock time in one outer SCF iteration with respect to the number of cores for the Si_{1000} system (strong scaling). (b) The change of wallclock time with respect to system size (weak scaling) on 8192 cores. The black dotted lines represent the ideal scaling.

tend to underestimate the energy gap of semiconductors. For example, GGA calculations that use the PBE functional¹³ give an energy gap of 0.69 eV in silicon, which is much smaller than that in bulk silicon measured in the experiments (1.17 eV).⁶⁴ The use of hybrid functionals can mitigate this type of error.

The defect concentration in silicons used in experimental studies is about 10^{18} cm^{-3} . To faithfully represent experimental conditions and avoid the nonphysical interactions between a defect and its images introduced by periodic boundary conditions, a large unit cell containing thousands of silicon atoms is required to in a computational study. Systems of this size is beyond the capability of existing plane-wave DFT software when a hybrid functional is used in the calculation. In this section, we employ the ACE-ISDF method implemented in PWDFDT to calculate the energy levels of a vacancy defect in an 1000-atom silicon system using the HSE06 hybrid functional. The defect concentration is about $5 \times 10^{19} \text{ cm}^{-3}$ in this case, which is close to concentration used in experimental studies. Our calculation based on the HSE06 functional yields an energy gap of 1.28 eV, which is very close to the experimental value of 1.17 eV. We also find that a calculation that uses a smaller unit cell that contains 512 atoms yields an energy gap of 1.32 eV,³¹ and hence, the size effect plays an important role.

Figure 6 shows the electronic structure of the vacancy defect in the Si_{1000} system computed with two different exchange–correlation functionals (GGA-PBE and HSE06). We fully relax the structures, respectively, with the GGA and HSE06 exchange–correlation functionals by using the steepest descent algorithm with the Barzilai–Borwein line search method⁶⁵ to optimize the geometry of the atomic configuration. The optimal unit cell lattice constants we obtained are 5.46 and 5.45 Å, respectively, for the GGA-PBE and HSE06 exchange–correlation functional-based calculations. These values are close to those reported in previous theoretical studies.⁶⁶

Our DFT calculations show that hybrid functional HSE06 calculations accurately describe the VBM and CBM energy levels of the silicon and the defect energy levels introduced by the vacancy defect. Furthermore, we can clearly see that the vacancy defect introduces three defect states into the intrinsic energy gap (1.28 eV) of silicon. These include a single defect state a_1 and a doubly degenerate state e , 0.14 and 1.04 eV above the VBM energy of silicon, respectively. The GGA-based DFT calculations fail to accurately predict these defect states.




Figure 6. Comparison of VBM, CBM, and defect energy levels of a 1000-atom silicon system that contains a vacancy defect. Two different types of exchange–correlation functionals (GGA-PBE and HSE06) are used for geometry optimization and for computing the electronic structure. All the energy levels are referenced to the VBM energy, which is set to zero. The black arrows are used to mark the occupied energy levels.

6. CONCLUSION

In this paper, we demonstrate that the interpolative separable density fitting (ISDF) decomposition can be used to reduce the cost hybrid functional calculations for large systems in a plane-wave DFT code. The reduction in cost results from the construction of a set of cN_c numerical auxiliary basis vectors, where c is a modest constant. Using these auxiliary basis vectors instead of N_c^2 products of the occupied orbitals, we only need to solve $O(N_c)$ Poisson-like equations instead of $O(N_c^2)$ equations to apply an approximate exchange operator to a set of occupied orbitals. The accuracy of the approximation depends entirely on the rank parameter c , and we find that the choice of c is insensitive with respect to the band gap.

The ISDF decomposition can be performed in $O(N_c^3)$ operations. The interpolation points are chosen by a randomized QR factorization with column pivoting (QRCP). It is relatively expensive compared to other parts of the ISDF calculation. However, this procedure only needs to be carried out once for all during the first outer SCF iteration in a hybrid functional calculation. The interpolation vectors can be computed via a least-squares fitting procedure that makes use of the separable nature of the functions to be fitted. The complexity of this step still scales as $O(N_c^3)$ but with a significantly smaller preconstant compared to the cost of applying the uncompressed exchange operator or the cost of QRCP. We are currently also exploring other methods for selecting interpolation points that avoid the use of the QRCP procedure, especially in the context of geometry optimization and *ab initio* molecular dynamics simulation.

For a moderate choice of rank parameter, the error in the total energy per atom and the force can be kept under 10^{-3} Hartree/atom and 10^{-3} Hartree/Bohr, respectively, for both semiconducting and metallic systems. Meanwhile the computational time can be reduced by up to an order of magnitude for applying the exchange operator once to all Kohn–Sham orbitals. We demonstrated that the ISDF decomposition can be combined with the adaptively compressed exchange operator (ACE) formulation to reduce the cost of ACE operator construction. The resulting ACE-ISDF method exhibits excellent parallel scalability on high performance computers and significantly reduces the time required to perform hybrid functional calculations by nearly 2 orders of magnitude. In particular, the

time spent in the ACE-ISDF-enabled hybrid functional calculation is only marginally higher than that spent in DFT calculations that use local and semilocal functionals.

However, we also find that hybrid functional calculations often require more iterations to converge compared to GGA calculations. One main reason is the two-level SCF iteration structure in hybrid functional calculations, which may be inefficient especially in the context of *ab initio* molecular dynamics simulation. Further reduction of the number of SCF iterations may close the final gap between hybrid functional calculations and calculations with local and semilocal functionals and thus opens the door to the accurate simulation of a vast range of nanomaterials using hybrid functionals beyond reach today.

AUTHOR INFORMATION

Corresponding Authors

*E-mail: whu@lbl.gov (W.H.).

*E-mail: linlin@math.berkeley.edu (L.L.).

*E-mail: cyang@lbl.gov (C.Y.).

ORCID

Wei Hu: 0000-0001-9629-2121

Lin Lin: 0000-0001-6860-9566

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

This work was partly supported by the Scientific Discovery through Advanced Computing (SciDAC) program funded by U.S. Department of Energy, Office of Science, Advanced Scientific Computing Research and Basic Energy Sciences (W.H., L.L., and C.Y.), the National Science Foundation under Grant DMS-1652330 (L.L.), and the Center for Applied Mathematics for Energy Research Applications (CAMERA) (L.L. and C.Y.). The authors thank the National Energy Research Scientific Computing (NERSC) Center and the Berkeley Research Computing program at the University of California, Berkeley, for making computational resources available to them. We thank Francois Gygi, Jianfeng Lu, Meiyue Shao, and Lexing Ying for helpful discussions.

REFERENCES

- (1) Hohenberg, P.; Kohn, W. Inhomogeneous Electron Gas. *Phys. Rev.* **1964**, *136*, B864.
- (2) Kohn, W.; Sham, L. J. Self-Consistent Equations Including Exchange and Correlation Effects. *Phys. Rev.* **1965**, *140*, A1133.
- (3) Perdew, J. P.; Schmidt, K. Jacobs Ladder of Density Functional Approximations for the Exchange-Correlation Energy. *AIP Conf. Proc.* **2000**, *577*, 1.
- (4) Becke, A. D. A New Mixing of Hartree-Fock and Local Density Functional Theories. *J. Chem. Phys.* **1993**, *98*, 1372.
- (5) Perdew, J. P.; Ernzerhof, M.; Burke, K. Rationale for Mixing Exact Exchange with Density Functional Approximations. *J. Chem. Phys.* **1996**, *105*, 9982.
- (6) Heyd, J.; Scuseria, G. E.; Ernzerhof, M. Hybrid Functionals based on a Screened Coulomb Potential. *J. Chem. Phys.* **2003**, *118*, 8207.
- (7) Heyd, J.; Scuseria, G. E.; Ernzerhof, M. Erratum: "Hybrid functionals based on a screened Coulomb potential" [*J. Chem. Phys.* **118**, 8207 (2003)]. *J. Chem. Phys.* **2006**, *124*, 219906.
- (8) Ceperley, D. M.; Alder, B. J. Ground State of the Electron Gas by a Stochastic Method. *Phys. Rev. Lett.* **1980**, *45*, 566.
- (9) Perdew, J. P.; Zunger, A. Self-Interaction Correction to Density-Functional Approximations for Many-Electron Systems. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1981**, *23*, 5048.

(10) Goedecker, S.; Teter, M.; Hutter, J. Separable Dual-Space Gaussian Pseudopotentials. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1996**, *54*, 1703.

(11) Becke, A. D. Density-Functional Exchange-Energy Approximation with Correct Asymptotic Behavior. *Phys. Rev. A: At., Mol., Opt. Phys.* **1988**, *38*, 3098.

(12) Lee, C.; Yang, W.; Parr, R. G. Development of the Colle-Salvetti Correlation-Energy Formula into a Functional of the Electron Density. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1988**, *37*, 785.

(13) Perdew, J. P.; Burke, K.; Ernzerhof, M. Generalized Gradient Approximation Made Simple. *Phys. Rev. Lett.* **1996**, *77*, 3865.

(14) Tao, J.; Perdew, J. P.; Staroverov, V. N.; Scuseria, G. E. Climbing the Density Functional Ladder: Nonempirical Meta-Generalized Gradient Approximation Designed for Molecules and Solids. *Phys. Rev. Lett.* **2003**, *91*, 146401.

(15) Sun, J.; Ruzsinszky, A.; Perdew, J. P. Strongly Constrained and Appropriately Normed Semilocal Density Functional. *Phys. Rev. Lett.* **2015**, *115*, 036402.

(16) Sun, J.; Perdew, J. P.; Ruzsinszky, A. Semilocal Density Functional Obeying a Strongly Tightened Bound for Exchange. *Proc. Natl. Acad. Sci. U. S. A.* **2015**, *112*, 685–689.

(17) Duchemin, I.; Gygi, F. A scalable and accurate algorithm for the computation of Hartree-Fock exchange. *Comput. Phys. Commun.* **2010**, *181*, 855–860.

(18) Bylaska, E. J.; Tsemekhman, K.; Baden, S. B.; Weare, J. H.; Jonsson, H. Parallel implementation of Γ -point pseudopotential plane-wave DFT with exact exchange. *J. Comput. Chem.* **2011**, *32*, 54–69.

(19) Barnes, T. A.; Kurth, T.; Carrier, P.; Wichmann, N.; Prendergast, D.; Kent, P. R. C.; Deslippe, J. Improved Treatment of Exact Exchange in Quantum ESPRESSO. *Comput. Phys. Commun.* **2017**, *214*, 52–58.

(20) Goedecker, S. Linear Scaling Electronic Structure Methods. *Rev. Mod. Phys.* **1999**, *71*, 1085.

(21) Bowler, D. R.; Miyazaki, T. $O(N)$ methods in Electronic Structure Calculations. *Rep. Prog. Phys.* **2012**, *75*, 036503.

(22) Guidon, M.; Hutter, J.; VandeVondele, J. Auxiliary Density Matrix Methods for Hartree-Fock Exchange Calculations. *J. Chem. Theory Comput.* **2010**, *6*, 2348–2364.

(23) DiStasio, R. A., Jr.; Santra, B.; Li, Z.; Wu, X.; Car, R. The Individual and Collective effects of Exact Exchange and Dispersion Interactions on the Ab Initio Structure of Liquid Water. *J. Chem. Phys.* **2014**, *141*, 084502.

(24) Dawson, W.; Gygi, F. Performance and Accuracy of Recursive Subspace Bisection for Hybrid DFT Calculations in Inhomogeneous Systems. *J. Chem. Theory Comput.* **2015**, *11*, 4655–4663.

(25) Damle, A.; Lin, L.; Ying, L. Compressed Representation of Kohn-Sham Orbitals via Selected Columns of the Density Matrix. *J. Chem. Theory Comput.* **2015**, *11*, 1463–1469.

(26) Kohn, W. Density Functional and Density Matrix Method Scaling Linearly with the Number of Atoms. *Phys. Rev. Lett.* **1996**, *76*, 3168.

(27) Rebolini, E.; Izsák, R.; Reine, S. S.; Helgaker, T.; Pedersen, T. B. Comparison of Three Efficient Approximate Exact-Exchange Algorithms: The Chain-of-Spheres Algorithm, Pair-Atomic Resolution-of-the-Identity Method, and Auxiliary Density Matrix Method. *J. Chem. Theory Comput.* **2016**, *12*, 3514–3522.

(28) Manzer, S.; Horn, P. R.; Mardirossian, N.; Head-Gordon, M. Fast, Accurate Evaluation of Exact Exchange: The Occ-RI-K Algorithm. *J. Chem. Phys.* **2015**, *143*, 024113.

(29) Boffi, N. M.; Jain, M.; Natan, A. Efficient Computation of the Hartree-Fock Exchange in Real-Space with Projection Operators. *J. Chem. Theory Comput.* **2016**, *12*, 3614–3622.

(30) Lin, L. Adaptively Compressed Exchange Operator. *J. Chem. Theory Comput.* **2016**, *12*, 2242–2249.

(31) Hu, W.; Lin, L.; Banerjee, A. S.; Vecharynski, E.; Yang, C. Adaptively Compressed Exchange Operator for Large-Scale Hybrid Density Functional Calculations with Applications to the Adsorption of Water on Silicene. *J. Chem. Theory Comput.* **2017**, *13*, 1188–1198.

(32) Lin, L.; Lindsey, M. Convergence of adaptive compression methods for Hartree-Fock-like equations. *arXiv:1703.05441*, 2017.

- (33) Giannozzi, P.; Baroni, S.; Bonini, N.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Chiarotti, G. L.; Cococcioni, M.; Dabo, I.; Dal Corso, A. D.; de Gironcoli, S.; Fabris, S.; Fratesi, G.; Gebauer, R.; Gerstmann, U.; Gougoussis, C.; Kokalj, A.; Lazzeri, M.; Martin-Samos, L.; Marzari, N.; Mauri, F.; Mazzarello, R.; Paolini, S.; Pasquarello, A.; Paulatto, L.; Sbraccia, C.; Scandolo, S.; Sclauzero, G.; Seitsonen, A. P.; Smogunov, A.; Umari, P.; Wentzcovitch, R. M. QUANTUM ESPRESSO: A Modular and Open-Source Software Project for Quantum Simulations of Materials. *J. Phys.: Condens. Matter* **2009**, *21*, 395502.
- (34) Lu, J.; Ying, L. Compression of the Electron Repulsion Integral Tensor in Tensor Hypercontraction Format with Cubic Scaling Cost. *J. Comput. Phys.* **2015**, *302*, 329–335.
- (35) Beebe, N.; Linderberg, J. Simplifications in the generation and transformation of two-electron integrals in molecular calculations. *Int. J. Quantum Chem.* **1977**, *12*, 683–705.
- (36) Koch, H.; Sánchez de Merás, A.; Pedersen, T. B. Reduced scaling in electronic structure calculations using Cholesky decompositions. *J. Chem. Phys.* **2003**, *118*, 9481–9484.
- (37) Aquilante, F.; Pedersen, T. B.; Lindh, R. Low-cost evaluation of the exchange Fock matrix from Cholesky and density fitting representations of the electron repulsion integrals. *J. Chem. Phys.* **2007**, *126*, 194106.
- (38) Manzer, S.; Horn, P. R.; Mardirossian, N.; Head-Gordon, M. Fast, accurate evaluation of exact exchange: The occ-RI-K algorithm. *J. Chem. Phys.* **2015**, *143*, 024113.
- (39) Ren, X.; Rinke, P.; Blum, V.; Wieferink, J.; Tkatchenko, A.; Sanfilippo, A.; Reuter, K.; Scheffler, M. Resolution-of-Identity Approach to Hartree-Fock, Hybrid Density Functionals, RPA, MP2 and GW with Numeric Atom-Centered Orbital Basis Functions. *New J. Phys.* **2012**, *14*, 053020.
- (40) Weigend, F. A fully direct RI-HF algorithm: Implementation, optimized auxiliary basis sets, demonstration of accuracy and efficiency. *Phys. Chem. Chem. Phys.* **2002**, *4*, 4285–4291.
- (41) Parrish, R. M.; Hohenstein, E. G.; Martínez, T. J.; Sherrill, C. D. Tensor hypercontraction. II. Least-squares renormalization. *J. Chem. Phys.* **2012**, *137*, 224106.
- (42) Parrish, R. M.; Hohenstein, E. G.; Martínez, T. J.; Sherrill, C. D. Discrete variable representation in electronic structure theory: Quadrature grids for least-squares tensor hypercontraction. *J. Chem. Phys.* **2013**, *138*, 194107.
- (43) Golub, G. H.; Reinsch, C. Singular Value Decomposition and Least Squares Solutions. *Numerische Mathematik* **1970**, *14*, 403–420.
- (44) Harbrecht, H.; Peters, M.; Schneider, R. On the Low-Rank Approximation by the Pivoted Cholesky Decomposition. *Applied Numerical Mathematics* **2012**, *62*, 428–440.
- (45) Chan, T. F.; Hansen, P. C. Some Applications of the Rank Revealing QR Factorization. *SIAM J. Sci. Statist. Comput.* **1992**, *13*, 727–741.
- (46) Liberty, E.; Woolfe, F.; Martinsson, P.-G.; Rokhlin, V.; Tygert, M. Randomized Algorithms for the Low-Rank Approximation of Matrices. *Proc. Natl. Acad. Sci. U. S. A.* **2007**, *104*, 20167–20172.
- (47) Halko, N.; Martinsson, P. G.; Tropp, J. A. Finding Structure with Randomness: Probabilistic Algorithms for Constructing Approximate Matrix Decompositions. *SIAM Rev.* **2011**, *53*, 217–288.
- (48) Auckenthaler, T.; Blum, V.; Bungartz, H. J.; Huckle, T.; Johanni, R.; Krämer, L.; Lang, B.; Lederer, H.; Willems, P. R. Parallel Solution of Partial Symmetric Eigenvalue Problems from Electronic Structure Calculations. *Parallel Comput.* **2011**, *37*, 783–794.
- (49) Lin, L.; Lu, J.; Ying, L.; E, W. Adaptive Local Basis Set for Kohn-Sham Density Functional Theory in a Discontinuous Galerkin Framework I: Total Energy Calculation. *J. Comput. Phys.* **2012**, *231*, 2140–2154.
- (50) Hu, W.; Lin, L.; Yang, C. DGDFT: A Massively Parallel Method for Large Scale Density Functional Theory Calculations. *J. Chem. Phys.* **2015**, *143*, 124110.
- (51) Hu, W.; Lin, L.; Yang, C. Edge Reconstruction in Armchair Phosphorene Nanoribbons Revealed by Discontinuous Galerkin Density Functional Theory. *Phys. Chem. Chem. Phys.* **2015**, *17*, 31397–31404.
- (52) Banerjee, A. S.; Lin, L.; Hu, W.; Yang, C.; Pask, J. E. Chebyshev Polynomial Filtered Subspace Iteration in the Discontinuous Galerkin Method for Large-Scale Electronic Structure Calculations. *J. Chem. Phys.* **2016**, *145*, 154101.
- (53) Zhang, G.; Lin, L.; Hu, W.; Yang, C.; Pask, J. E. Adaptive Local Basis Set for Kohn-Sham Density Functional Theory in a Discontinuous Galerkin Framework II: Force, Vibration, and Molecular Dynamics Calculations. *J. Comput. Phys.* **2017**, *335*, 426–443.
- (54) Hartwigsen, C.; Goedecker, S.; Hutter, J. Relativistic Separable Dual-Space Gaussian Pseudopotentials from H to Rn. *Phys. Rev. B: Condens. Matter Mater. Phys.* **1998**, *58*, 3641.
- (55) Pople, J. A. Nobel Lecture: Quantum Chemical Models. *Rev. Mod. Phys.* **1999**, *71*, 1267.
- (56) Schwegler, E.; Challacombe, M. Linear Scaling Computation of the Hartree-Fock Exchange Matrix. *J. Chem. Phys.* **1996**, *105*, 2726.
- (57) Ochsenfeld, C.; White, C. A.; Head-Gordon, M. Linear and Sublinear Scaling Formation of Hartree-Fock-Type Exchange Matrices. *J. Chem. Phys.* **1998**, *109*, 1663.
- (58) Shang, H.; Li, Z.; Yang, J. Implementation of Screened Hybrid Density Functional for Periodic Systems with Numerical Atomic Orbitals: Basis Function Fitting and Integral Screening. *J. Chem. Phys.* **2011**, *135*, 034110.
- (59) Shang, H.; Li, Z.; Yang, J. Implementation of Exact Exchange with Numerical Atomic Orbitals. *J. Phys. Chem. A* **2010**, *114*, 1039–1043.
- (60) Newman, R. C. Defects in Silicon. *Rep. Prog. Phys.* **1982**, *45*, 1163.
- (61) Watkins, G. D. Defects in Silicon. *Mater. Sci. Semicond. Process.* **2000**, *3*, 227.
- (62) Freysoldt, C.; Grabowski, B.; Hickel, T.; Neugebauer, J.; Kresse, G.; Janotti, A.; Van de Walle, C. G. First-Principles Calculations for Point Defects in Solids. *Rev. Mod. Phys.* **2014**, *86*, 253–305.
- (63) Heyd, J.; Peralta, J. E.; Scuseria, G. E.; Martin, R. L. Energy Band Gaps and Lattice Parameters Evaluated with the Heyd-Scuseria-Ernzerhof Screened Hybrid Functional. *J. Chem. Phys.* **2005**, *123*, 174101.
- (64) Ledoux, G.; Gong, J.; Huisken, F.; Guillois, O.; Reynaud, C. Photoluminescence of Size-Separated Silicon Nanocrystals: Confirmation of Quantum Confinement. *Appl. Phys. Lett.* **2002**, *80*, 4834.
- (65) Barzilai, J.; Borwein, J. M. Two-point step size gradient methods. *IMA J. Numer. Anal.* **1988**, *8*, 141–148.
- (66) Lucero, M. J.; Henderson, T. M.; Scuseria, G. E. Improved Semiconductor Lattice Parameters and Gaps from a Middle-Range Screened Hybrid Exchange Functional. *J. Phys.: Condens. Matter* **2012**, *24*, 145504.