

Fast Exchange with Gaussian Basis Set Using Robust Pseudospectral Method

Sandeep Sharma,* Alec F. White, and Gregory Beylkin

Cite This: <https://doi.org/10.1021/acs.jctc.2c00720>

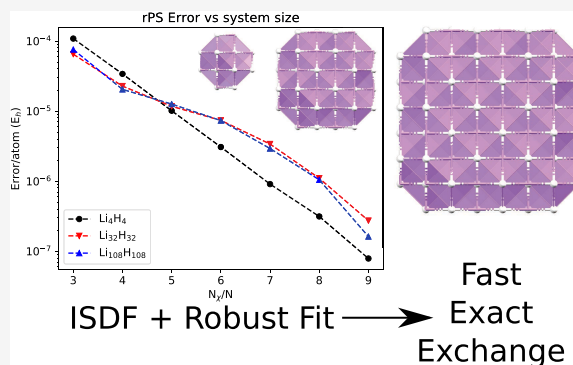
Read Online

ACCESS |

Metrics & More

Article Recommendations

ABSTRACT: In this article, we present an algorithm to efficiently evaluate the exchange matrix in periodic systems when a Gaussian basis set with pseudopotentials is used. The usual algorithm for evaluating exchange matrix scales cubically with the system size because one has to perform $O(N^2)$ fast Fourier transform (FFT). Here, we introduce an algorithm that retains the cubic scaling but reduces the prefactor significantly by eliminating the need to do FFTs during each exchange build. This is accomplished by representing the products of Gaussian basis function using a linear combination of an auxiliary basis the number of which scales linearly with the size of the system. We store the potential due to these auxiliary functions in memory, which allows us to obtain the exchange matrix without the need to do FFT, albeit at the cost of additional memory requirement. Although the basic idea of using auxiliary functions is not new, our algorithm is cheaper due to a combination of three ingredients: (a) we use a robust pseudospectral method that allows us to use a relatively small number of auxiliary basis to obtain high accuracy; (b) we use occ-RI exchange, which eliminates the need to construct the full exchange matrix; and (c) we use the (interpolative separable density fitting) ISDF algorithm to construct these auxiliary basis sets that are used in the robust pseudospectral method. The resulting algorithm is accurate, and we note that the error in the final energy decreases exponentially rapidly with the number of auxiliary functions.



1. INTRODUCTION

The inclusion of exact Hartree–Fock (HF) exchange within the framework of Kohn–Sham (KS) density functional theory (DFT) is critical to the success of DFT for molecular systems, and these “hybrid” functionals are used in almost all modern DFT calculations on molecules. For periodic solids, hybrid functionals can dramatically improve on standard semilocal functionals for a wide variety of properties,^{1–5} but the large computational cost associated with computing the HF exchange is a limiting factor. The evaluation of the nonlocal HF exchange can be computationally demanding in any context, and it is particularly expensive for solids where calculations with hybrid density functionals may be orders of magnitude more expensive than for their purely semilocal counterparts.

Molecular calculations usually use local basis sets where the ratio of basis functions to electrons, N/n , is small, often on the order of 3–10. Although computation of the electron repulsion integrals is naively $O(N^4)$, the locality of the basis implies an asymptotically linear number of significant basis function pairs and quadratic scaling of the classical Coulomb and the HF exchange.⁶ The computation of the Coulomb energy and potential can be further reduced to $O(N)$ by multipole expansion,^{7,8} while the exchange can be computed in linear time by leveraging the sparsity of the density matrix for

nonmetallic systems.^{9–12} This asymptotically linear region is rarely reached in practice, and recent work has focused on reducing the computational cost of practical calculations by tensor factorization. The resolution of the identity (RI) approximation,^{13,14} also called “density fitting,” is the most widely used such method, and efficient approaches for both Coulomb (RI-J)^{15,16} and exchange (RI-K)^{17–20} have been developed. Dunlap introduced a “robust” approximation to the two-electron integrals for which the error in integrals is quadratic in the fitting error for the basis function products.^{21–23} The Cholesky decomposition (CD) is another method that can be used to obtain a decomposition of RI form without the need for optimized auxiliary basis sets.^{24,25} An alternative approach, the pseudospectral (PS) method, is to provide a factorization from a real-space perspective by introducing a basis of grid points and performing one of the integrals analytically.²⁶ The chain-of-spheres exchange

Received: July 11, 2022

(COSX) algorithm²⁷ and related seminumerical exchange algorithms^{28–31} are a particularly general application of the PS method to exchange. The idea of factorizing the integral tensor was taken one step further in the tensor hypercontraction (THC) method of Martinez and co-workers where the four-index two-electron integral tensor is decomposed into a product of five 2-index tensors.^{32–34} The difficulty of efficiently performing the THC tensor decomposition has largely limited its use to correlated methods, but a cubic scaling factorization algorithm was first introduced by Lu and Ying³⁵ and later used to accelerate the computation of exact exchange in periodic^{36,37} and molecular³⁹ calculations under the name “interpolative separable density fitting” (ISDF).

In traditional plane wave DFT calculations, the action of the Coulomb operator on the occupied orbitals can be evaluated by solving n Poisson equations. This scales quadratically in total ($O(nN_g \ln N_g)$ for n electrons and N_g grid points). In contrast, the action of the exchange operator scales cubically ($O(n^2 N_g \ln N_g)$ for n electrons and N_g grid points), a fact that has motivated the development of numerous numerical methods that seek to lessen this cost. Most notable are linear scaling methods^{12,40} for which the sparsity in the exchange operator relies on the system being an insulator.⁴¹ Stochastic density functional theory (sDFT),⁴² including the extension to hybrid functionals,⁴³ can achieve linear scaling without any locality arguments, but controlling the stochastic error results in a very large prefactor. As with molecular systems, the linear regime is usually out-of-reach, and methods that do not improve the scaling, like the adaptively compressed exchange (ACE),⁴⁴ can greatly increase the efficiency in practice. When large super cells are necessary, a screening and/or truncating the coulomb operator in the exchange term can somewhat lessen the cost.⁴⁵ The auxiliary density matrix method⁴⁶ (ADMM) can significantly reduce the cost by approximating the density matrix.⁴⁶

In this work, we present an efficient algorithm for evaluating the exchange matrix in periodic Gaussian-type orbital (GTO) calculations. We begin by providing background information that will help the reader understand the reason for the high cost of exchange matrix evaluation, namely, that one has to perform $O(N^2)$ FFTs. In Section 3, we describe the algorithm which combines several different ideas that eliminate the need for performing FFTs during exchange build and replace them with matrix multiplications. Although the algorithm is still cubic scaling, the prefactor is significantly reduced. In Section 4, we describe the computational details including an efficient parallel implementation. Finally, in Section 5, we show that this algorithm enables computation of the exchange matrix with a cost comparable to the computation of the Coulomb matrix. This allows hybrid DFT to be used for systems where semilocal DFT is feasible.

2. BACKGROUND

In this work, we will solve the Hartree–Fock equations using the self-consistent field (SCF) method, where, for a given set of molecular orbitals, one constructs the Coulomb (\hat{J}) and the exchange (\hat{K}) operators, respectively. The Coulomb and exchange operators are given by the expressions

$$\hat{J}(\mathbf{r}) = \int_{\mathbb{R}^3} \sum_i \frac{|\phi_i(\mathbf{r}')|^2}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'$$

$$\hat{K}(\mathbf{r}, \mathbf{r}') = \sum_i \frac{\phi_i(\mathbf{r})\phi_i(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}$$

where $\phi_i(\mathbf{r})$ are the occupied molecular orbitals. From the equation, one can note that while the Coulomb operator is diagonal, the exchange operator has a rank n . To make progress, one typically introduces a basis set and obtains the Coulomb and exchange matrices J and K , respectively. Although many possible basis functions can be used including Slater-type orbitals,^{47,48} wavelets,^{49–53} numerical basis functions,^{54–57} etc., the two most commonly used ones are the atom-centered Gaussian basis functions^{58–67} and the plane wave basis.^{68–72} Although, in this work, we will use the Gaussian basis functions to represent molecular orbitals, we will also make use of the plane wave basis and their dual basis (the periodized Sinc basis) to simplify certain calculations.^{73–76} We first review the properties of these functions.

2.1. Basis Functions. We begin by introducing periodic basis functions in the unit cell defined by vectors $\mathbf{a}_1, \mathbf{a}_2, \mathbf{a}_3$ such that the volume $\Omega = \mathbf{a}_1 \cdot (\mathbf{a}_2 \times \mathbf{a}_3)$. The reciprocal vectors are denoted by $\mathbf{A}_1, \mathbf{A}_2, \mathbf{A}_3$ so that $\mathbf{a}_i \cdot \mathbf{A}_j = 2\pi\delta_{ij}$.

2.1.1. Plane Wave Basis. The plane wave basis functions ($\xi_{\mathbf{G}}$) are parameterized by the wave vector \mathbf{G} and are given by

$$\xi_{\mathbf{G}}(\mathbf{r}) = \frac{1}{\sqrt{\Omega}} \exp(-i\mathbf{G} \cdot \mathbf{r})$$

It is easy to check that these functions form an orthogonal basis i.e.,

$$\int_{\Omega} \xi_{\mathbf{G}'}(\mathbf{r})^* \xi_{\mathbf{G}}(\mathbf{r}) d\mathbf{r} = \delta_{\mathbf{G}', \mathbf{G}}$$

as long as $\mathbf{G} = n_1\mathbf{A}_1 + n_2\mathbf{A}_2 + n_3\mathbf{A}_3$, where n_i are integers. If in a calculation one retains $n_i = -\frac{N_i}{2}, \dots, \frac{N_i+1}{2}$, where N_i is an integer, $i = 1, 2, 3$, then the total number of plane waves is equal to $N_g = N_1 N_2 N_3$.

2.1.2. Sinc Basis. The dual basis consists of the periodized Sinc functions ($\xi_{\mathbf{R}}$) (referred to as pSinc functions from now on) which are associated with grid points \mathbf{R} . They are obtained by the unitary transformation of the plane wave basis functions

$$\xi_{\mathbf{R}}(\mathbf{r}) = \frac{1}{\sqrt{N_g}} \sum_{\mathbf{G}} \exp(i\mathbf{G} \cdot \mathbf{R}) \xi_{\mathbf{G}}(\mathbf{r}) \quad (1)$$

$$\xi_{\mathbf{G}}(\mathbf{r}) = \frac{1}{\sqrt{N_g}} \sum_{\mathbf{R}} \exp(-i\mathbf{G} \cdot \mathbf{R}) \xi_{\mathbf{R}}(\mathbf{r}) \quad (2)$$

Again it is easy to see that $\xi_{\mathbf{R}}$ are orthogonal if $\mathbf{R} = \frac{n_1\mathbf{a}_1}{N_1} + \frac{n_2\mathbf{a}_2}{N_2} + \frac{n_3\mathbf{a}_3}{N_3}$ are a uniform set of grid points in the unit cell. Since the plane wave and the pSinc bases are related via a unitary transformation, they span the same space.

2.1.3. Atom-Centered Gaussian Basis. We also have the computational basis, the periodized atom-centered Gaussian functions $\phi_{p\mu}(\mathbf{r})$ given by

$$\phi_{p\mu}(\mathbf{r}) = \sum_{\mathbf{T}} \phi_{\mu}(\mathbf{r} + \mathbf{T})$$

where $\phi(\mathbf{r})$ is the atom-centered Gaussian and $\mathbf{T} = n_1\mathbf{a}_1 + n_2\mathbf{a}_2 + n_3\mathbf{a}_3$ are the lattice vectors in the real space. This function is

also known as a Jacobi theta function. Note that $\phi_{p\mu}(\mathbf{r})$ is neither orthogonal nor normalized even if the functions $\phi(\mathbf{r})$ themselves are. The periodized Gaussians can be represented as a linear combination of the plane waves

$$\phi_{p,\mu}(\mathbf{r}) = \sum_{\mathbf{G}} \hat{\phi}_{\mu}(\mathbf{G}) \xi_{\mathbf{G}}(\mathbf{r}) \quad (3)$$

where $\hat{\phi}_{\mu}(\mathbf{G})$ is the Fourier transform of the function $\phi_{\mu}(\mathbf{r})$

$$\hat{\phi}_{\mu}(\mathbf{G}) = \int_{\mathbb{R}^3} \exp(i\mathbf{G} \cdot \mathbf{r}) \phi_{\mu}(\mathbf{r}) d\mathbf{r}$$

The Fourier transform of a Gaussian basis function is also a Gaussian, thus if

$$\phi_{\mu}(\mathbf{r}) = \exp(-\mu|\mathbf{r}-\mathbf{A}_{\mu}|^2)$$

then

$$\hat{\phi}_{\mu}(\mathbf{G}) = \left(\frac{\pi}{\mu}\right)^{3/2} \frac{\exp(-|\mathbf{G}|^2/4\mu)}{\sqrt{\Omega}} \exp(i\mathbf{G} \cdot \mathbf{A}_{\mu}) \quad (4)$$

If the factor μ of the exponent of the Gaussian is small, then a relatively few plane waves N_g are needed to represent it (with a small but finite error). In this work, we will use pseudopotentials so that the Gaussian basis functions will be reasonably flat and a manageable number of plane waves is sufficient to represent them.

It is also possible to represent the periodized Gaussians as a linear combination of the pSinc basis functions

$$\begin{aligned} \phi_{p\mu}(\mathbf{r}) &= \sum_{\mathbf{R}} \tilde{\phi}_{\mu}(\mathbf{R}) \xi_{\mathbf{R}}(\mathbf{r}) \\ \tilde{\phi}_{\mu}(\mathbf{R}) &= \frac{1}{\sqrt{N_g}} \sum_{\mathbf{G}} \exp(-i\mathbf{G} \cdot \mathbf{R}) \hat{\phi}_{\mu}(\mathbf{G}) \\ \hat{\phi}_{\mu}(\mathbf{G}) &= \frac{1}{\sqrt{N_g}} \sum_{\mathbf{R}} \exp(i\mathbf{G} \cdot \mathbf{R}) \tilde{\phi}_{\mu}(\mathbf{R}) \end{aligned}$$

These equations follow from eqs 3 and 1, respectively. In the following, it will be convenient to switch from the pSinc basis to the plane wave basis using the Fast Fourier transform (FFT), i.e., the coefficients $\tilde{\phi}_{\mu}(\mathbf{R})$ can be calculated from those of $\hat{\phi}_{\mu}(\mathbf{G})$ (and vice versa) rapidly using FFT at a cost of $O(N_g \ln N_g)$.

So far, we have been careful to distinguish between ϕ_{μ} and $\phi_{p\mu}$; however, we will not do so for the rest of the paper where we only use the symbol ϕ_{μ} and it should be understood that it refers to the periodized Gaussian function.

2.2. Integral Evaluation and the Diagonal Approximation. In the following, we evaluate the potential due to a charge density where the charge density is given as a product of two orbitals. If we represent both orbitals $\phi_{\mu}(\mathbf{r})$ and $\phi_{\nu}(\mathbf{r})$ using plane wave basis, we then obtain the charge density $\rho_{\mu\nu}(\mathbf{r})$ as

$$\begin{aligned} \rho_{\mu\nu}(\mathbf{r}) &= \phi_{\mu}(\mathbf{r}) \phi_{\nu}(\mathbf{r}) \\ &= \frac{1}{\sqrt{\Omega}} \sum_{\mathbf{G}} \sum_{\mathbf{G}'}^{N_g} \hat{\phi}_{\mu}(\mathbf{G}) \hat{\phi}_{\nu}(\mathbf{G}') \xi_{\mathbf{G}+\mathbf{G}'}(\mathbf{r}) \\ &= \frac{1}{\sqrt{\Omega}} \sum_{\mathbf{G}} \sum_{\mathbf{G}'}^{2N_g} \hat{\phi}_{\mu}(\mathbf{G}-\mathbf{G}') \hat{\phi}_{\nu}(\mathbf{G}') \xi_{\mathbf{G}}(\mathbf{r}) \\ &\approx \frac{1}{\sqrt{\Omega}} \sum_{\mathbf{G}} \sum_{\mathbf{G}'}^{N_g} \hat{\phi}_{\mu}(\mathbf{G}-\mathbf{G}') \hat{\phi}_{\nu}(\mathbf{G}') \xi_{\mathbf{G}}(\mathbf{r}) \end{aligned}$$

where in going from the second to the third expression, we have changed the variables, and in the last expression, we made an approximation by truncating the summation over \mathbf{G} . By choosing a sufficiently large cutoff N_g , the error due to the approximation can be made negligibly small since we expect the density to be sufficiently smooth so that the contributions to $\rho_{\mu\nu}(\mathbf{G})$ coming from plane waves above the cutoff are small. With this approximation, the final equation starts to look like a convolution and one can show that

$$\begin{aligned} \rho_{\mu\nu}(\mathbf{r}) &\approx \frac{1}{\sqrt{\Omega}} \sum_{\mathbf{G}} \sum_{\mathbf{G}'}^{N_g} \hat{\phi}_{\mu}(\mathbf{G}-\mathbf{G}') \hat{\phi}_{\nu}(\mathbf{G}') \xi_{\mathbf{G}}(\mathbf{r}) \\ &\approx \sqrt{\frac{N_g}{\Omega}} \sum_{\mathbf{R}} \tilde{\phi}_{\mu}(\mathbf{R}) \tilde{\phi}_{\nu}(\mathbf{R}) \xi_{\mathbf{R}}(\mathbf{r}) \end{aligned} \quad (5)$$

The last expression has been called the diagonal approximation and has been extensively used in the work of Steve White in the context of Gausslets^{52,77,78} and before that with Sinc functions.⁷⁹

Using the diagonal approximation described above, we can calculate the two-electron integrals of four Gaussian basis functions as follows

$$\begin{aligned} (\mu\nu|\lambda\sigma) &= \int \int \phi_{\mu}(\mathbf{r}) \phi_{\nu}(\mathbf{r}) \frac{1}{|\mathbf{r}-\mathbf{r}'|} \phi_{\lambda}(\mathbf{r}') \phi_{\sigma}(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \\ &= \frac{N_g}{\Omega} \sum_{\mathbf{R}} \sum_{\mathbf{R}'} \tilde{\phi}_{\mu}(\mathbf{R}) \tilde{\phi}_{\nu}(\mathbf{R}) \tilde{\phi}_{\lambda}(\mathbf{R}') \tilde{\phi}_{\sigma}(\mathbf{R}') \\ &\quad \int \int \xi_{\mathbf{R}}(\mathbf{r}) \frac{1}{|\mathbf{r}-\mathbf{r}'|} \xi_{\mathbf{R}'}(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \\ &= \frac{N_g}{\Omega} \sum_{\mathbf{R}} \sum_{\mathbf{R}'} \tilde{\phi}_{\mu}(\mathbf{R}) \tilde{\phi}_{\nu}(\mathbf{R}) v(\mathbf{R}-\mathbf{R}') \tilde{\phi}_{\lambda}(\mathbf{R}') \tilde{\phi}_{\sigma}(\mathbf{R}') \end{aligned} \quad (6)$$

$$v(\mathbf{R}-\mathbf{R}') = \int \int \xi_{\mathbf{R}}(\mathbf{r}) \frac{1}{|\mathbf{r}-\mathbf{r}'|} \xi_{\mathbf{R}'}(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \quad (7)$$

$$= \frac{1}{N_g} \sum_{\mathbf{G}} \exp(-i\mathbf{G} \cdot \mathbf{R}) \frac{4\pi}{G^2} \exp(i\mathbf{G} \cdot \mathbf{R}') \quad (8)$$

In going from the first to the second expression in eq 6, we have used eq 5, and in going from the second to the third expression, we have used eq 1 and made use of the fact that the Coulomb operator is diagonal and is equal to $\frac{4\pi}{|\mathbf{G}|^2}$ in plane wave basis. It is useful to note that eq 6 takes the form of tensor hypercontraction (THC) and just like in THC one can make use of the integrals in this form to calculate exchange with a cubic scaling cost. As we will show next, the Coulomb matrix

can be evaluated using a linear scaling algorithm because the matrix $v(\mathbf{R} - \mathbf{R}')$ takes a special form shown in eq 7.

2.3. Coulomb Matrix. In the self-consistent field calculation, we construct the coulomb matrix $J_{\mu\nu} = \langle \mu | \hat{J} | \nu \rangle$ for a set of periodized Gaussian functions ϕ_μ and ϕ_ν . Further, we represent the molecular orbitals as a linear combination of the basis functions $\phi_i(\mathbf{r}) = \sum_\lambda C_{\lambda i} \phi_{\lambda}(\mathbf{r})$, where $C_{\lambda i}$ is the matrix of molecular coefficients and we also define the density matrix $D_{\lambda\sigma} = \sum_i C_{\lambda i} C_{\sigma i}$. Using these, the Coulomb matrix can be written as

$$J_{\mu\nu} = \int d\mathbf{r}' \phi_\mu(\mathbf{r}') \phi_\nu(\mathbf{r}') \left(\int d\mathbf{r} \frac{1}{|\mathbf{r} - \mathbf{r}'|} \sum_{\lambda\sigma} D_{\lambda\sigma} \phi_\lambda(\mathbf{r}) \phi_\sigma(\mathbf{r}) \right) \\ = \frac{1}{\Omega} \sum_{\mathbf{R}} (\tilde{\phi}_\mu(\mathbf{R}) \tilde{\phi}_\nu(\mathbf{R})) \left(\sum_{\mathbf{G}} \exp(-i\mathbf{G} \cdot \mathbf{R}) \frac{4\pi}{G^2} \right. \\ \left. \left(\sum_{\mathbf{R}'} \exp(i\mathbf{G} \cdot \mathbf{R}') \left(\sum_{\lambda\sigma} D_{\lambda\sigma} \tilde{\phi}_\lambda(\mathbf{R}') \tilde{\phi}_\sigma(\mathbf{R}') \right) \right) \right)$$

where we have made use of eqs 6 and 7. Also the order in which the brackets are placed specifies the order in which the tensor contractions are performed. Specifically, the steps involved are

1. We first evaluate the density on a grid $\rho(\mathbf{R}') = \sum_{\lambda\sigma} D_{\lambda\sigma} \tilde{\phi}_\lambda(\mathbf{R}') \tilde{\phi}_\sigma(\mathbf{R}')$. This can be performed with a linear cost since for any finite threshold, the atom-centered Gaussian basis functions have a compact support, which implies that for a given λ only an $O(1)$ number of σ have nonzero overlap.
2. Next the density in Fourier space is evaluated $\rho(\mathbf{G}) = \sum_{\mathbf{R}} \exp(i\mathbf{G} \cdot \mathbf{R}') \rho(\mathbf{R}')$, which is done efficiently with $N_g \ln N_g$ cost using the FFT.
3. In the next step, the potential due to this density is evaluated in real space as $V(\mathbf{R}') = \sum_{\mathbf{G}} \exp(-i\mathbf{G} \cdot \mathbf{R}') \frac{4\pi}{G^2} \rho(\mathbf{G})$, which can also be evaluated efficiently using FFT with an $N_g \ln N_g$ cost.
4. Finally, the Coulomb matrix is evaluated as $J_{\mu\nu} = \sum_{\mathbf{R}} \tilde{\phi}_\mu(\mathbf{R}) \tilde{\phi}_\nu(\mathbf{R}) V(\mathbf{R})$, which again is calculated with a linear scaling due to the compact support of the Gaussian basis set.

Thus, the construction of the Coulomb matrix only requires two FFT and step 4 is the dominant cost of the calculation.

2.4. Exchange Matrix. The exchange matrix $\langle \mu | \hat{K} | \nu \rangle$ can be constructed in a number of ways using a quartic scaling algorithm, e.g., using density fitting. However, a cubic scaling algorithm is as follows

$$K_{\mu\nu} = \sum_i \int \phi_\mu(\mathbf{r}') \phi_i(\mathbf{r}) \left(\int \frac{\phi_i(\mathbf{r}') \phi_\nu(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}' \right) d\mathbf{r} \\ = \frac{1}{\Omega} \sum_{\mathbf{R}} (\tilde{\phi}_\mu(\mathbf{R}) \tilde{\phi}_i(\mathbf{R})) \\ \left(\sum_{\mathbf{G}} \exp(-i\mathbf{G} \cdot \mathbf{R}) \frac{4\pi}{G^2} \right. \\ \left. \left(\sum_{\mathbf{R}'} \exp(i\mathbf{G} \cdot \mathbf{R}') \tilde{\phi}_i(\mathbf{R}') \tilde{\phi}_\nu(\mathbf{R}') \right) \right) \quad (9)$$

where the order of tensor contraction is represented by the brackets. The steps involved in constructing the exchange matrix are

1. We first evaluate the product of the molecular orbital $\tilde{\phi}_i(\mathbf{R}')$ and atomic orbital $\tilde{\phi}_\nu(\mathbf{R}')$ on the grid to obtain $\tilde{\rho}_{i\nu}(\mathbf{R}') = \tilde{\phi}_i(\mathbf{R}') \tilde{\phi}_\nu(\mathbf{R}')$, which is an $O(1)$ operation for a given i, ν . The value of this product density is then evaluated in the Fourier space $\tilde{\rho}_{i\nu}(\mathbf{G})$ using FFT.
2. Next the Coulomb kernel $\frac{4\pi}{G^2}$ multiplies $\tilde{\rho}_{i\nu}(\mathbf{G})$ and the inverse FFT is used to evaluate the potential $\tilde{V}_{i\nu}(\mathbf{R}')$ due to $\tilde{\rho}_{i\nu}(\mathbf{R}')$.
3. Finally, the potential is contracted with the product $\tilde{\phi}_\mu(\mathbf{R}) \tilde{\phi}_i(\mathbf{R})$ to obtain the value of $K_{\mu\nu}$.

Thus, one has to perform Nn FFTs and inverse FFTs, one pair each for a given i and ν . The cost of the FFT is $N_g \ln N_g$, which makes the total cost of the exchange evaluations equal to $NnN_g \ln N_g$, where, as explained in Table 1, N , n , and N_g are the number of basis functions, number of electrons (or occupied orbitals), and number of grid points (or plane wave basis), respectively.

Table 1. Notation Used in the Rest of the Paper

notation	meaning
n	number of electrons
N	number of atom-centered Gaussian basis functions
N_g	number of grid points or plane wave basis
N_χ	number of grid points used in pseudospectral method
k	number of processors (including MPI and OMP)
k_t	number of threads per node
ϕ_i or with subscripts i, j, k, l, \dots	occupied molecular orbitals
ϕ_μ or with subscripts $\mu, \nu, \lambda, \sigma, \dots$	atom-centered Gaussian basis
χ	pseudospectral fitting functions
\mathbf{R}	location of the grid point (or the center of the periodized Sinc basis)
\mathbf{G}	wavenumber of the plane wave

2.5. Other Steps. For performing the HF calculation, one also needs the kinetic matrix $T_{\mu\nu}$, the nuclear matrix $N_{\mu\nu}$, and the overlap matrix $S_{\mu\nu}$. These matrices are evaluated using standard techniques⁸⁰ and we do not discuss them further. Finally, the Fock matrix $F_{\mu\nu} = T_{\mu\nu} + N_{\mu\nu} + J_{\mu\nu} - K_{\mu\nu}$ is constructed and diagonalized to obtain the molecular coefficients $C_{\lambda i}$ using which the molecular orbitals $\phi_i(\mathbf{r})$ are formed. These molecular orbitals are used to construct the Fock matrix in a self-consistent field cycle. The diagonalization of the Fock matrix to obtain the molecular orbitals scales cubically; however, the prefactor of this step is small so that its cost does not exceed that of Coulomb matrix construction (which scales linearly) for systems with less than 5000 basis functions (for example, see Table 4).

3. SPEEDING UP EXCHANGE EVALUATION

As we mentioned in the previous section, the exchange evaluation scales cubically with the size of the system. In particular, we have to perform Nn FFTs. In practice, the cubic scaling in itself is not a severe limitation (at least for problems in which ≤ 5000 basis functions are utilized) since other steps such as Fock matrix diagonalization scale cubically as well.

However, the high prefactor associated with performing Nn FFTs results in a large overall cost. In principle, one can make the evaluation of the exchange matrix to scale linearly using the fact that $\hat{K}(\mathbf{r}, \mathbf{r}')$ decays exponentially with $|\mathbf{r} - \mathbf{r}'|$ for systems with a band gap or metallic systems at finite temperatures. However, even for insulating systems with a reasonably large band gap, the exponential decay is extremely slow and one does not reach the linear scaling regime until very large system sizes. In this work, we take the point of view that by reducing the prefactor in the cubic scaling algorithm, we can obtain a method for exchange matrix formation that is only slightly more expensive than the formation of the Coulomb matrix. We briefly describe how this is accomplished before going into details in subsequent subsections.

Let us begin by observing that the number of pairs of basis function $(\phi_\mu\phi_\nu)$ scale quadratically with the size of the system. However, eq 5 shows that using the diagonal approximation (which can be made arbitrarily accurate), one can ensure that the products can be represented using a linear combination of pSinc basis, the number of which scales linearly with the size of the system. This fact reduces the scaling of the exchange build to cubic (from the usual quartic), however with a large prefactor. We will retain the cubic scaling but reduce the prefactor by fitting the products of Gaussians using a linear combination of a relatively small number of auxiliary functions. We are able to get away with using a small number of auxiliary functions because we make use of the so-called robust pseudospectral method, which ensures that the error in the exchange matrix is quadratic of the error in the fitting. Below we first show how the auxiliary functions (represented by χ_g) are obtained and then how they are used in the robust pseudospectral method. Finally, we also introduce occ-RI exchange algorithm that further allows us to reduce the cost.

3.1. Interpolative Separable Density Fitting. This algorithm was introduced by Lu and Ying³⁵ for writing down the two-electron integrals in the THC format and later it was adapted for speeding up exchange matrix evaluation in hybrid DFT with plane wave basis.^{36,37} The basis idea is to view the product density

$$\tilde{\rho}_{\mu\nu,\mathbf{R}} = \tilde{\phi}_\mu(\mathbf{R})\tilde{\phi}_\nu(\mathbf{R})$$

as a matrix with N^2 rows corresponding to indices $\mu\nu$ and N_g columns corresponding to the grid points. One then tries to obtain a subset of columns such that all other columns can be written as a linear combination of this subset. Specifically, we will attempt to write

$$\tilde{\rho}_{\mu\nu,\mathbf{R}} = \tilde{\phi}_\mu(\mathbf{R})\tilde{\phi}_\nu(\mathbf{R}) \approx \sum_{\mathbf{R}_g} \tilde{\phi}_\mu(\mathbf{R}_g)\tilde{\phi}_\nu(\mathbf{R}_g)\chi_g(\mathbf{R}) + O(\epsilon) \quad (10)$$

where \mathbf{R}_g is a subset of all grid points that we call interpolation points, $\chi_g(\mathbf{R})$ is a function defined on all grid points, and ϵ is the error incurred due to the approximation. The interpolation grid points \mathbf{R}_g are obtained by performing QR decomposition with column pivoting (QRCP)³⁸ on the matrix $\tilde{\rho}_{\mu\nu,\mathbf{R}}$ which allows one to write it as

$$\tilde{\rho}P = QR$$

where P is a permutation matrix that ensures that the diagonal entries of the matrix R are in the descending order, $R_{11} \geq R_{22} \geq R_{33} \dots$. We then choose a user-defined N_χ set of interpolation points specified by the leading pivot points in P . These provide

an optimal set of interpolation points. The cost of performing this step directly scales as $N^2N_g^2$ (assuming $N^2 > N_g$) and is thus prohibitively expensive. It can be made to scale cubically using a randomized algorithm, and we provide more details in Section 4.3.

Once the interpolation points are selected, we can obtain the function $\chi_g(\mathbf{R})$ using a least square minimization

$$\min_{\chi_{\mathbf{R}_g,\mathbf{R}}} |\tilde{\rho}_{\mu\nu,\mathbf{R}} - \tilde{\rho}_{\mu\nu,\mathbf{R}_g}\chi_{\mathbf{R}_g,\mathbf{R}}|^F$$

where we have written $\chi_g(\mathbf{R})$ as a matrix $\chi_{\mathbf{R}_g,\mathbf{R}}$ using the Einstein notation of summing over repeated indices (which we will continue to do so for the rest of this subsection), and subscript F indices the Frobenius norm. The least squares problem is solved by transforming it into a system of linear equations

$$\begin{aligned} \tilde{\rho}_{\mu\nu,\mathbf{R}_g}\tilde{\rho}_{\mu\nu,\mathbf{R}_g}\chi_{\mathbf{R}_g,\mathbf{R}} &= \tilde{\rho}_{\mu\nu,\mathbf{R}_g}\tilde{\rho}_{\mu\nu,\mathbf{R}} \\ X_{\mathbf{R}_g,\mathbf{R}_g}\chi_{\mathbf{R}_g,\mathbf{R}} &= X_{\mathbf{R}_g,\mathbf{R}} \end{aligned} \quad (11)$$

where $X = \tilde{\rho}^T\tilde{\rho}$. We solve a set of N_g linear equations of size $N_\chi \times N_\chi$ to obtain the fitting functions $\chi_{\mathbf{R}_g,\mathbf{R}}$. A naive evaluation of X will cost $N^2N_gN_g$; however, using the product structure of $\tilde{\rho} = \tilde{\phi}_\mu\tilde{\phi}_\nu$, we reduce the cost to NN_gN_χ

$$X_{\mathbf{R}_g,\mathbf{R}} = (\tilde{\phi}_{\mu,\mathbf{R}_g}\tilde{\phi}_{\nu,\mathbf{R}})^2$$

Thus, by obtaining the interpolation points using QRCP and the fitting functions by solving the least squares problem, we have a way of systematically improving approximation in eq 10. We show in Section 3.2 how to use ISDF to reduce the CPU cost for exchange evaluation.

3.2. Robust Fitting. Substituting eq 10 into eq 6, we obtain a set of approximate integrals

$$\begin{aligned} (\mu\nu|\lambda\sigma) &\approx \frac{1}{\Omega} \sum_{\mathbf{R}_g} \sum_{\mathbf{R}\mathbf{R}'} \tilde{\phi}_\mu(\mathbf{R}_g)\tilde{\phi}_\nu(\mathbf{R}_g) \\ &\quad \left(\chi_g(\mathbf{R}) \sum_{\mathbf{G}} \exp(-i\mathbf{G}\cdot\mathbf{R}) \frac{4\pi}{G^2} \exp(i\mathbf{G}\cdot\mathbf{R}') \right) \tilde{\phi}_\lambda(\mathbf{R}') \\ &\quad \tilde{\phi}_\sigma(\mathbf{R}') + O(\epsilon) \\ &\approx \frac{1}{\Omega} \sum_{\mathbf{R}_g} \sum_{\mathbf{R}\mathbf{R}'} \tilde{\phi}_\mu(\mathbf{R}_g)\tilde{\phi}_\nu(\mathbf{R}_g)V(\mathbf{R}_g,\mathbf{R}')\tilde{\phi}_\lambda(\mathbf{R}')\tilde{\phi}_\sigma(\mathbf{R}') \\ &\quad + O(\epsilon) \end{aligned} \quad (12)$$

where we have defined $V(\mathbf{R}_g,\mathbf{R})$ as the potential due to the function $\chi_g(\mathbf{R})$. However, these integrals are not symmetric with respect to functions on the bra and ket (i.e., $(\mu\nu|\lambda\sigma) \neq (\lambda\sigma|\mu\nu)$), and consequently lead to exchange matrix that is asymmetric. The integrals can be made symmetric, and more importantly, the error can be made $O(\epsilon^2)$ (which is a significant improvement over $O(\epsilon)$ error) using the following modification

$$\begin{aligned}
(\mu\nu|\lambda\sigma) &\approx \frac{1}{\Omega} \sum_{\mathbf{R}_g} \sum_{\mathbf{R}'} \tilde{\phi}_\mu(\mathbf{R}_g) \tilde{\phi}_\nu(\mathbf{R}_g) V(\mathbf{R}_g, \mathbf{R}') \tilde{\phi}_\lambda(\mathbf{R}') \\
&\quad \tilde{\phi}_\sigma(\mathbf{R}') + \frac{1}{\Omega} \sum_{\mathbf{R}} \sum_{\mathbf{R}_g} \tilde{\phi}_\mu(\mathbf{R}) \tilde{\phi}_\nu(\mathbf{R}) V(\mathbf{R}_g, \mathbf{R}) \\
&\quad \tilde{\phi}_\lambda(\mathbf{R}_g) \tilde{\phi}_\sigma(\mathbf{R}_g) - \frac{1}{\Omega} \sum_{\mathbf{R}_g'} \sum_{\mathbf{R}_g} \tilde{\phi}_\mu(\mathbf{R}_g) \\
&\quad \tilde{\phi}_\nu(\mathbf{R}_g) W(\mathbf{R}_g, \mathbf{R}_g') \tilde{\phi}_\lambda(\mathbf{R}_g') \tilde{\phi}_\sigma(\mathbf{R}_g') + O(\epsilon^2)
\end{aligned} \quad (13)$$

where $W(\mathbf{R}_g, \mathbf{R}_g')$ is a symmetric matrix given by

$$W(\mathbf{R}_g, \mathbf{R}_g') = \sum_{\mathbf{R}} V(\mathbf{R}_g, \mathbf{R}) \chi_g(\mathbf{R}) \quad (14)$$

To see where the quadratic error in eq 13 comes from, it is useful to recognize that the two-electron integral is of the form $\sum_{ab} v(a) M(a, b) w(b)$. Now we can approximate $v(a) = \sum_{a_g} v(a_g) \xi_g(a) + \epsilon \delta(a)$ and write a similar expression for $w(b)$, where $\delta(a)$ is a normalized error vector and ϵ is included to signify the magnitude of the error. Using these approximations, one can write

$$\begin{aligned}
\sum_{ab} v(a) M(a, b) w(b) &= \sum_{ab} \left(\sum_{a_g} v(a_g) \xi_g(a) + \epsilon \delta(a) \right) M \\
&\quad (a, b) \left(\sum_{b_g} w(b_g) \xi_g(b) + \epsilon' \delta'(b) \right) \\
&= \sum_{a'b'} v(a_g) \left(\sum_a \xi_g(a) M(a, b) \right) w(b) \\
&\quad + \sum_{a'b'} v(a) \left(\sum_b M(a, b) \xi_g(b) \right) \\
&\quad w(b_g) \\
&\quad + \sum_{a'b'} v(a_g) \\
&\quad \left(\sum_{ab} \xi_g(a) M(a, b) \xi_g(b) \right) w(b) \\
&\quad + \epsilon \epsilon' \sum_{ab} \delta(a) M(a, b) \delta'(b)
\end{aligned}$$

The above expression is exact and in the robust pseudospectral (rPS) technique we include the first three terms and exclude the final term which is quadratic in the error ϵ . The approximation made in eq 13 is known as the robust pseudospectral (rPS) and was introduced recently.⁸¹ If one only includes the first line of the equation, then we get the pseudospectral (PS) method, and if one only uses the last line of the equation, then we get the tensor hypercontraction (THC) method. Recently, it was emphasized by Valeev et al.⁸¹ that the quadratic error can be obtained using rPS. In fact, this approach has been previously used in other contexts including robust density fitting^{21,22} and even in quantum Monte Carlo for the evaluation of the reduced density matrices.^{82–85}

The key point is that one can evaluate the potential due to the functions $\chi_g(\mathbf{R})$ only at the beginning of the calculation and then during subsequent evaluations of the exchange matrix one can avoid having to do FFT. Below we show the order in

which the tensor contraction can be performed to obtain the exchange matrix

$$K_{\mu\nu} = \frac{1}{\Omega} \sum_{\mathbf{R}_g} \tilde{\phi}_\mu(\mathbf{R}_g) \left(\sum_{\mathbf{R}'} V(\mathbf{R}_g, \mathbf{R}') \left(\sum_i \left(\sum_{\sigma} C_{\sigma i} \tilde{\phi}_\sigma(\mathbf{R}_g) \right) \left(\sum_{\lambda} C_{\lambda i} \tilde{\phi}_\lambda(\mathbf{R}') \right) \right) \tilde{\phi}_\nu(\mathbf{R}') \right) \quad (15)$$

where we have only focused on the first term of eq 13 (other terms can be treated in a similar way). More explicitly, the steps involved are

1. First contraction over λ and σ indices is carried out to obtain molecular orbitals at all grid points ($\tilde{\phi}_i(\mathbf{R}')$) and interpolation grid points ($\tilde{\phi}_i(\mathbf{R}_g)$), respectively. The cost of both these calculations is $O(nN_g)$ because of the locality of the atomic orbitals.
2. Next the contraction over i is carried out to obtain the matrix $\rho(\mathbf{R}_g, \mathbf{R}') = \sum_i \tilde{\phi}_i(\mathbf{R}') \tilde{\phi}_i(\mathbf{R}_g)$ and the cost of this contraction is $O(nN_g N_\chi)$.
3. Next we element-wise multiply V and ρ matrices to obtain $V_2(\mathbf{R}_g, \mathbf{R}') = V(\mathbf{R}_g, \mathbf{R}') \rho(\mathbf{R}_g, \mathbf{R}')$ which costs $O(N_\chi N_g)$.
4. The matrix $V_2(\mathbf{R}_g, \mathbf{R}')$ is contracted with $\tilde{\phi}_\mu(\mathbf{R}_g)$ to obtain a new matrix $M_\mu(\mathbf{R}') = \sum_{\mathbf{R}_g} V_2(\mathbf{R}_g, \mathbf{R}') \tilde{\phi}_\mu(\mathbf{R}_g)$ at a cost of $O(NN_g N_\chi)$.
5. Finally, we contract over \mathbf{R}' to obtain exchange matrix $K_{\mu\nu} = \sum_{\mathbf{R}'} M_\mu(\mathbf{R}') \tilde{\phi}_\nu(\mathbf{R}')$ at a cost of $O(N_g N_\chi^2)$.

Given that $N_g > N_\chi > N > n$ we can see that several steps of the algorithm are cubic scaling with step 4 being the most expensive. Note that no FFT evaluations are involved, and in Section 5, we will show that this leads to significant speed-up of the calculation.

3.3. occ-RI. If one is only interested in calculating the total energy or the energy of the occupied orbitals, then it is easy to show that only the rectangular part of the Fock matrix $F_{i\mu}$ is needed, where i are the labels of the occupied molecular orbitals and μ are the atomic orbitals. This was first introduced by Manzer et al.²⁰ to reduce the cost of construction of the exchange matrix. Using $F_{i\mu}$ instead of the full $F_{\nu\mu}$ matrix during the SCF cycle does not deteriorate the rate of convergence and the only drawback is that the virtual orbital energies are not evaluated correctly. However, this shortcoming can be overcome at the very end of the SCF cycle by evaluating the full matrix one time. This idea is also at the heart of the efficiency of the ACE method proposed by Lin,⁴⁴ with the difference being that the Fock matrix is not explicitly constructed but is written as a sum of outer product of n vectors. This trick can readily be used with the integrals given in eq 13, and we obtain

$$\begin{aligned}
K_{\mu\nu} = & \sum_{\mathbf{R}'} \left(\sum_{\mathbf{R}_g} \tilde{\phi}_j(\mathbf{R}_g) V(\mathbf{R}_g, \mathbf{R}') \left(\sum_i \left(\sum_{\sigma} C_{\sigma i} \tilde{\phi}_{\sigma}(\mathbf{R}_g) \right) \right. \right. \\
& \left. \left. \left(\sum_{\lambda} C_{\lambda i} \tilde{\phi}_{\lambda}(\mathbf{R}') \right) \right) \right) \tilde{\phi}_{\nu}(\mathbf{R}') \\
& + \sum_{\mathbf{R}_g} \left(\sum_{\mathbf{R}} \tilde{\phi}_j(\mathbf{R}) V(\mathbf{R}_g, \mathbf{R}) \left(\sum_i \left(\sum_{\sigma} C_{\sigma i} \tilde{\phi}_{\sigma}(\mathbf{R}) \right) \right. \right. \\
& \left. \left. \left(\sum_{\lambda} C_{\lambda i} \tilde{\phi}_{\lambda}(\mathbf{R}_g) \right) \right) \right) \tilde{\phi}_{\nu}(\mathbf{R}_g) \\
& - \sum_{\mathbf{R}'_g} \left(\sum_{\mathbf{R}_g} \tilde{\phi}_j(\mathbf{R}_g) W(\mathbf{R}_g, \mathbf{R}'_g) \left(\sum_i \left(\sum_{\sigma} C_{\sigma i} \tilde{\phi}_{\sigma}(\mathbf{R}_g) \right) \right. \right. \\
& \left. \left. \left(\sum_{\lambda} C_{\lambda i} \tilde{\phi}_{\lambda}(\mathbf{R}'_g) \right) \right) \right) \tilde{\phi}_{\nu}(\mathbf{R}'_g)
\end{aligned} \quad (16)$$

The cost of performing contractions in the first line is $O(nN_{\chi}N_g)$, of second line is $O(nN_{\chi}N_g)$, and third line is $O(nN_{\chi}^2)$ which are all lower than the leading cost of $O(NN_gN_{\chi})$ of evaluating the exchange matrix without occ-RI. However, as we will see in the next section, when performing the calculation in parallel, the cost of the second step shown above is either $O(nN_{\chi}N_g/k) + O(nN_{\chi}N)$ or $O(NN_gN_{\chi}/k)$, where k is the number of processors (assuming we do not want to incur additional communication overhead) and the choice between the two is not always obvious. In fact, if one goes with the second choice, then occ-RI provides no advantage.

4. COMPUTATIONAL DETAILS

Having outlined the basic algorithm for fast evaluation of exchange we will outline the computational details of the program paying particular attention to the way in which it is parallelized by making use of both the message passing interface (MPI) and OpenMP (OMP) together. The summary of the memory and CPU cost of the various steps of the algorithm are displayed in Table 2, and the details are presented in the section specified in the fourth column.

4.1. Voronoi Partitioning for Each Atom. We begin by partitioning the grid points $\{\mathbf{R}\}$ into disjoint set of points κ_I , one set for each nucleus I , such that all points in the set are closer to atom I (or its periodic images) than any other atom (or its periodic images). The algorithm for doing Voronoi

partitioning for periodic unit cells is standard.⁸⁶ The sets are distributed in a round-robin manner between different processors and roughly an equal number of grid points end up on each processor.

After the partitioning is performed, each processor is used to evaluate the value of periodized Gaussian basis on the grid point belonging to κ_I associated with it. These values, up to a user-defined threshold ϵ_1 (usually 10^{-8}), are stored in memory. Thus, we end up getting a different matrix $\tilde{\phi}_{\mu}^I(\mathbf{R})$ for each atom I , which specifies the values of only those functions μ that have a non-negligible value on a subset of grid points \mathbf{R} in the Voronoi partition κ_I . The memory cost of storing this matrix is $O(1)$ because only those μ are included that have a value above ϵ_1 for at least one grid point in κ_I and due to the local nature of Gaussians, this only happens for a number of functions that are asymptotically system-size-independent. The number of grid points \mathbf{R} in κ_I are also system-size-independent. Thus, the total memory requirement for storing the Gaussian basis set is linear in system size and it is equally divided among the different processors, each processor stores $O(N/k)$ amount of memory. Although in principle one should be able to calculate the matrix $\tilde{\phi}_{\mu}^I(\mathbf{R})$ (the value of all Gaussian basis functions at all grid points) at an $O(N)$ cost by evaluating the values of the Gaussian basis functions in real space (each Gaussian has compact support up to a finite threshold). However, in our algorithm, we use an $O(N^2)$ algorithm, whereby we evaluate the value of the basis functions in the reciprocal space (recall that the Fourier transform of a Gaussian is also a Gaussian), and then we use FFT to evaluate the functions at all grid points. This procedure is more expensive but one avoids the need for lattice summations (needed for at least small unit cells) and is more convenient to implement. This calculation is split up nearly evenly among the various processors leading to an asymptotic $O(NN_g/k)$ cost per processor. The cost of this step is small enough compared to the rest of the steps that even for large systems, it does not become a bottleneck.

4.2. Parallel Coulomb Matrix Formation. With the Gaussians stored in memory, the three steps of the construction of the Coulomb matrix are evaluated in parallel on each processor. The density matrix $D_{\mu\nu}$ is replicated on each processor, and first, the density $\rho(\mathbf{R}') = \sum_{\lambda\sigma} D_{\lambda\sigma} \tilde{\phi}_{\lambda}(\mathbf{R}') \tilde{\phi}_{\sigma}(\mathbf{R}')$ is calculated on each processor separately only for the grid points associated with it. The current algorithm is naturally linear scaling because only those basis λ, σ are contracted that have a non-negligible contribution to the same Voronoi partition κ_I . The cumulative density is then evaluated by summing up the contributions to density coming from all processors using the MPI command MPI_Allreduce, after which the potential is evaluated by calling FFT two times. The calls to FFT are only parallelized using OMP locally on each processor. The cost of these FFT calculations is negligible and constitutes a very small fraction of the overall cost, and thus not carrying it out in parallel does not cause significant overhead. Finally, after the potential is obtained, the third step in which the potential is contracted with the Gaussians on the local grid to obtain the Coulomb operator $J_{\mu\nu} = \sum_{\mathbf{R}} \tilde{\phi}_{\mu}(\mathbf{R}) \tilde{\phi}_{\nu}(\mathbf{R}) V(\mathbf{R})$ is carried out on each processor and then reduced together. Thus, one calculates those elements of the Coulomb operator $J_{\mu\nu}$ for which both the indices μ, ν have non-negligible values on the Voronoi partitions associated with a given processor. However, the entire Coulomb matrix is replicated on each processor. Thus, the CPU cost of the algorithm is $O(N_g/k)$ and the memory cost is $O(N^2)$.

Table 2. Memory and CPU Cost for Constructing the Various Tensors in the Algorithm^a

steps	memory	CPU	section
$\phi_{\mu}(\mathbf{R})$	$O(N_g/k)$	$O(NN_g/k)$	4.1
$J_{\mu\nu}$	$O(N^2)$	$O(N_g/k)$	4.2
$\chi_g(\mathbf{R})$	$O(N_{\chi}N_g/k)$	$O(N_{\chi}^3/k_I)$	4.3
$V(\mathbf{R}_g, \mathbf{R}')$	$O(N_{\chi}N_g/k)$	$O((N_{\chi}N_g \ln N_g)/k)$	4.3
$W(\mathbf{R}_g, \mathbf{R}'_g)$	$O(N_{\chi}^2/k)$	$O(N_{\chi}^2N_g/k)$	4.3
$K_{\mu\nu}$	$O(N^2)$	$O(nN_{\chi}N_g/k) + O(nN_{\chi}N)$	4.4

^aAll of the symbols are defined in Table 1. For details, see the text in the sections pointed out in the last column.

The Voronoi partitioning allows us to screen the overlap between Gaussians efficiently and results in a linear scaling algorithm that is distributed over different processors almost ideally.

4.3. ISDF. Next we evaluate the interpolation points and the fitting functions for use in exchange build. The input to the program is a number c that specifies the number of interpolation points that will be used as a factor of the number of basis functions e.g., $c = 2$ means that twice as many interpolation points and fitting functions as the number of basis functions will be used. When the number of fitting functions becomes equal to N_g , one gets the exact result. Calculating all interpolation points for the entire system scales as $O(N^2 N_g^2)$ which is prohibitive. So as a first step, a randomized algorithm^{87,88} is used to reduce the cost. In this algorithm, one first constructs two random matrices G^1 and G^2 with orthogonal columns of size $N \times p$, where $p = \sqrt{N_\chi} + m$, where we round up $\sqrt{N_\chi}$ and m is a small integer usually less than 5. Using these matrices, we first construct the randomized density matrix

$$\tilde{\rho}_{mn,\mathbf{R}} = \left(\sum_{\mu} G_{\mu,m}^1 \phi_{\mu}(\mathbf{R}) \right) \left(\sum_{\nu} G_{\nu,m}^2 \phi_{\nu}(\mathbf{R}) \right)$$

where the size of the cumulative index $mn \sim N_\chi$, i.e., the number of fitting bases that will be used. The QRCP decomposition is then performed on this much smaller matrix $\tilde{\rho}_{mn,\mathbf{R}}$ instead of the full matrix $\tilde{\rho}_{\mu\nu,\mathbf{R}}$ which reduces the scaling to $O(N_\chi^2 N_g)$. The step is still fairly slow and without additional simplification can dominate the overall cost of the calculation even for small systems. Previously Dong et al.³⁷ have proposed to use a method based on centroid Voronoi tessellation (CVT) with a weighted K-Mean algorithm to reduce the cost to $O(N_\chi N_g)$. Here we follow a different approach. We first calculate a small subset of fitting points for each Voronoi partitioning κ_I (described in Section 4.1) for each atom I . If the number of atom-centered basis functions on an atom I is N_I , then we obtain a set of $\sim cN_I + 10$ fitting points from each partition using the same randomized QRCP as described above. Each of these calculations is extremely fast and independent of the size of the system because both the number of grid points in κ_I and the number of basis functions N_I are small and system-size-independent. In the second step, all of these fitting points are accumulated to form a set $\{\mathbf{R}_h\}$ where the number of points is of similar number as N_χ but much smaller than N_g . Now we try to find N_χ interpolation points from all of these $\{\mathbf{R}_h\}$ grid points again using the same randomized algorithm described above, only this time we end up with a much smaller matrix

$$\tilde{\rho}_{mn,\mathbf{R}_h} = \left(\sum_{\mu} G_{\mu,m}^1 \phi_{\mu}(\mathbf{R}_h) \right) \left(\sum_{\nu} G_{\nu,m}^2 \phi_{\nu}(\mathbf{R}_h) \right)$$

on which one can readily perform QRCP decomposition at a cost of $O(N_\chi^3)$. The QRCP is parallelized using OMP and thus the CPU cost is $O(N_\chi^3/k_t)$, where k_t is the number of threads per node. The construction of the interpolation points is no longer the dominant cost of the overall calculation unless one goes to very large system sizes (see Table 4). In our current implementation, the QRCP is only parallelized using OMP and not using MPI+OMP. In a future publication, we will use

Scalapack⁸⁹ to further reduce the scaling from $O(N_\chi^3/k_t)$ to $O(N_\chi^3/k)$.

Once the interpolation points have been obtained, one can now calculate the fitting function $\chi_g(\mathbf{R})$ by solving eq 11. It is worth noting that it contains N_g equations which can be solved in parallel to obtain $\chi_g(\mathbf{R})$ with only a subset of \mathbf{R} that are associated with the processor. Thus, the CPU cost of this step is $O(N_\chi^3/k_t) + O(NN_\chi N_g/k)$. The first term arises because of the need to perform Cholesky decomposition of the matrix $X_{\mathbf{R}_g, \mathbf{R}_g}$ and the term arises because of the need to solve eq 11. Following this, one needs to obtain $W(\mathbf{R}_g, \mathbf{R}_g')$ using eq 14, and to do that, one has to first obtain the potential $V(\mathbf{R}_g, \mathbf{R})$ due to each function $\chi_g(\mathbf{R})$ (see eq 12). If we view $\chi_g(\mathbf{R})$ as a matrix of size $N_\chi \times N_g$, then after solving the linear equations, we obtain a matrix of size roughly $N_\chi \times (N_g/k)$ on each processor where we have k processors, i.e., each processor contains a subset of columns of the entire matrix. To obtain the potential, we first use MPI_Alltoall to obtain matrices of size $(N_\chi/k) \times N_g$, i.e., each processor now only contains a subset of rows of the entire matrix. FFT is performed in parallel on these rows to obtain the potential $V(\mathbf{R}_g, \mathbf{R})$ where each processor again only retains a submatrix of size $(N_\chi/k) \times N_g$ and this results in a CPU cost of $O(N_\chi N_g \ln(N_g)/k)$. After this, a call to MPI_Alltoall is used to now distribute the matrix $V(\mathbf{R}_g, \mathbf{R})$ with a column-wise split such that each processor ends up with a submatrix of size $N_\chi \times (N_g/k)$. Finally, having access to $V(\mathbf{R}_g, \mathbf{R})$ and $\chi_g(\mathbf{R})$ one can evaluate the inner product defined in eq 14 to obtain $W(\mathbf{R}_g, \mathbf{R}_g')$, which can again be distributed over all processors leading to a cost of $O(N_\chi^2 N_g/k)$. Finally, in all of these calculations, the matrices $\chi_g(\mathbf{R})$, $V(\mathbf{R}_g, \mathbf{R})$, $W(\mathbf{R}_g, \mathbf{R}_g')$ are all distributed evenly among the processors and thus require a memory of $O(N_\chi N_g/k)$, $O(N_\chi N_g/k)$, $O(N_\chi^2/k)$, respectively.

4.4. Parallel Exchange Matrix Formation. As mentioned in the text below eq 16, when one uses the occ-RI, one has two choices on how to perform contraction in the second line of the equation. One can first calculate $M_j(\mathbf{R}_g) = \sum_{\mathbf{R}} \tilde{\phi}_j(\mathbf{R}) V(\mathbf{R}_g, \mathbf{R})$ followed by $K_{j\nu} = \sum_{\mathbf{R}_g} M_j(\mathbf{R}_g) \tilde{\phi}_\nu(\mathbf{R}_g)$, or one can reverse the order of these two contractions. In the first case, only the first step can be parallelized and the second step has to be performed on each processor and this leads to the computational cost $O(nN_\chi N_g/k) + O(nN_\chi N)$, while when the reversed order is used, the computational cost becomes $O(NN_\chi N_g/k)$. In practice, the first algorithm tends to be faster unless the number of processors is very large because the ratio N_g/n can be on the order of 10^3 or more.

4.5. Other Considerations. The nuclear matrix takes the form

$$\begin{aligned} N_{\mu\nu} &= \int \frac{-Z_I}{|\mathbf{R}_I - \mathbf{r}|} \rho_{\mu\nu}(\mathbf{r}) d\mathbf{r} \\ &= \frac{-Z_I}{\Omega} \sum_{\mathbf{G}} \frac{4\pi}{G^2} \exp(i\mathbf{G} \cdot \mathbf{R}_I) \rho_{\mu\nu}(\mathbf{G}) \end{aligned}$$

where $\rho_{\mu\nu}(\mathbf{G})$ are the Fourier components of the density $\rho(\mathbf{r}) = \phi_\mu(\mathbf{r})\phi_\nu(\mathbf{r})$, and here, we have represented the nucleus as a delta distribution of charge Z_I at position \mathbf{R}_I . The summation is truncated at N_g and the largest error is incurred when both functions μ, ν are sharp. The expression for the two-electron integral for the same sharp functions is given as

$$(\mu\nu|\mu\nu) = \frac{1}{\Omega} \sum_{\mathbf{G}} \frac{4\pi}{G^2} |\rho_{\mu\nu}(\mathbf{G})|^2$$

where again the error is due to the use of a finite N_g . However, if the nuclear charge is truly assumed to be distributed as a delta function, then it is clear that the error incurred in the two-electron integrals is quadratic of that of the one incurred while calculating nuclear integrals. This trend is somewhat tempered due to the fact that pseudopotentials used in this work are not quite delta functions, but nonetheless, the nuclear integrals still show the largest error. To reduce the error, we use an effectively larger N_g while calculating the nuclear integral, which has to be done only once and use a sparser grid for the rest of the calculation (including for the Coulomb and exchange matrix formation).

In periodic systems, one often encounters the problem of linear dependencies of the Gaussian basis sets. This can be understood by looking at eq 4; if the value of μ is smaller than $0.3/L^2$, where L is the length of the super cell, then one can confirm that $\hat{\phi}_{\mu}(\mathbf{G})$ is effectively a constant. For small unit cells, several Gaussian functions in a standard basis set can become constants and thus are linearly dependent. However, when one tries to reach the thermodynamic limit by increasing L (the dimension of the super cell), then the Gaussians are no longer going to be linearly dependent and the number of linearly dependent functions do not increase linearly with the number of super cells in the system (of course the basis functions corresponding to the γ point of the primitive unit cell are still present in the super cell which are constants). Thus, in the large super cell limit, we do not expect the linear dependency problem to be significantly worse than in a large relatively uniform molecule or a cluster. Nevertheless, it is true that using Gaussian basis functions, it is difficult to obtain results in the basis set limit because it is difficult to systematically improve the basis set without running into linear dependencies (which also happens for molecular systems). One way to overcome this is to use a mixed plane wave Gaussian basis set. In our algorithm, it is relatively straightforward to do so such that sharp features are described by the Gaussian basis function and diffuse features are described by plane wave basis functions. We will pursue this line of work in the future.

It is well known that the exchange energy has an integrable singularity that disappears in the infinite systems size limit. However, when one is using finite-sized super cells, the singularity persists (note the presence of $1/G^2$ in the denominator of eq 8) and one needs ways of regularizing it. Several approaches have been proposed^{90,91} to do so including using truncated coulomb⁹² or the minimum image convention.^{59,93} Here, we have used a very simple approach whereby we remove the $\mathbf{G} = 0$ term in the exchange matrix evaluation and include a correction term $nM/2$, where M is the Madelung constant of the super cell. It is worth pointing out that one can readily use the truncated Coulomb kernel in exchange evaluation using our algorithm.

5. RESULTS

In this section, we present calculations using two benchmark systems Li–H solid ($(\text{Li}_4\text{H}_4)_n$) and diamond with unit cell $(\text{C}_8)_n$. For both these systems, we make use of the Goedecker–Teter–Hutter (GTH) pseudopotentials^{94,95} and GTH-DZVP and GTH-TZVP basis sets.⁷⁶ As pointed in the previous

section, the nuclear integrals are evaluated using a large N_g cutoff (in fact part of the calculation is performed in real space). For the calculation of the Coulomb and exchange matrix, $N_g = 35^3$ is used for the conventional unit cell. Numerical experiments showed that this grid is sufficient to deliver an error of less than 0.5 mHartree accuracy in Li–H system and 2 μ Hartree in diamond. In this section, we first look at the relative accuracy of a single-step randomized ISDF calculation versus Voronoi-based one that we have proposed here. Second, we compare the accuracy of the rPS and THC approximations for obtaining the exchange energy and accuracy of the rPS as we change the basis set and the size of the system. Finally, we present data on the cost of the calculations for Li–H of various system sizes, with the largest one being $\text{Li}_{256}\text{H}_{256}$ that contains 4864 basis functions.

5.1. Accuracy of Approximate ISDF. In our algorithm, we do not perform a single large ISDF calculation, instead we first perform a single small ISDF calculation for each atom and then do a single ISDF calculation to select the most important pivot points out of the collection of all of the pivots points obtained from the individual calculations (let us call it v-ISDF for Voronoi-ISDF). Table 3 shows the error incurred in Li_4H_4

Table 3. Error in mE_h for Li_4H_4 and $\text{C}_8(\text{Diamond})$ as N_{χ} Is Increased with respect to Results Obtained with Exact Exchange^a

$c = N_{\chi}/N$	Li_4H_4		$\text{C}_8(\text{diamond})$	
	ISDF	v-ISDF	ISDF	v-ISDF
3	0.86	0.67	44.15	39.15
4	0.18	0.23	6.02	6.39
5	0.03	0.08	1.15	1.41
6	0.01	0.03	0.17	0.29

^aIt is noteworthy that the errors of randomized ISDF (labeled ISDF) and Voronoi-ISDF (the algorithm proposed in this work) are quite similar. The differences between the two are slightly larger than the random errors one expects to see due to the use of randomized algorithm.

and C_8 systems when the usual randomized ISDF is used as opposed to v-ISDF. We note that the errors are quite similar pointing to the fact that the collection of pivot points that we obtain from individual ISDF calculations contains the important pivot points. One can thus devise other algorithms to select the most important points out of this small subset, including for example the K-Means algorithm presented previously.³⁷ We will work on these aspects in a future publication, particularly because the single ISDF calculation needed can become expensive for large system sizes as shown in Table 4.

5.2. Accuracy of rPS vs THC. Figure 1 shows the error as a function of c when rPS and THC are used for the two systems. It is clear that the error incurred by rPS is significantly lower than THC consistent with the fact that the error in the former is expected to be square of that of the latter. Two trends stand out in the left two subfigure. First, it appears that for the same value of c the error in the larger TZ basis set is smaller than that in the smaller DZ basis set and this trend is observed for both systems. Interestingly, when one plots the error versus N_{χ} rather than c , the errors in the two curves nearly overlap and again this trend is observed for both systems (see rightmost subfigure). Given that the cost of formation of the exchange matrix is $O(nN_{\chi}N_g/k) + O(nN_{\chi}N)$, we expect the

Table 4. Comparison of the Cost of Various Steps of the Algorithm Relative to the Cost of the Coulomb Matrix Construction as the Size of the Lithium Hydride System Is Increased^a

system	N	K	QRCP ^b	$\chi_g(\mathbf{R})$	$W(\mathbf{R}_g, \mathbf{R}'_g)/(\text{MPI})$	diagonalize ^b
Li ₄ H ₄	76	1.2	1.0	0.4	1.3/(0.0)	0.0
Li ₃₂ H ₃₂	608	1.0	0.3	0.3	0.7/(0.0)	0.0
Li ₁₀₈ H ₁₀₈	2052	1.1	0.6	0.4	0.6/(0.3)	0.0
Li ₁₄₄ H ₁₄₄	2736	1.4	0.7	0.5	0.7/(0.3)	0.0
Li ₁₉₂ H ₁₉₂	3648	1.6	1.8	0.5	1.9/(1.4)	0.3
Li ₂₅₆ H ₂₅₆	4864	1.9	3.9	0.6	2.9/(1.9)	0.6

^aAll calculations are Γ point calculations and do not make use of k -point symmetry. For Li₄H₄, the various parameters are $n = 16$, $N = 76$, $N_\chi = 304$, and $N_g = 42875$, and for other systems, each of these parameters increase in proportion to the size of the system. Column “K” refers to the cost of exchange build, column “QRCP” refers to the cost of performing of obtaining the interpolation points, column “ $\chi_g(\mathbf{R})$ ” refers to the cost of obtaining the interpolation functions by solving the least square problem, and column $W(\mathbf{R}_g, \mathbf{R}'_g)$ refers to the cost of building the matrix $W(\mathbf{R}_g, \mathbf{R}'_g)$ in eq 13. All Hartree–Fock calculations converged with eight iterations, and the CPU cost is relative to the cost for constructing the Coulomb matrix 8 times. In column 6, the number in brackets represents the time spent in just the MPI calls. ^bThis is parallelized using only OMP and not MPI.

cost to increase linearly with the increase in the size of the basis set N (assuming N_χ and N_g remain constant). The second trend one can observe is that the error decreases extremely rapidly for the Li₄H₄ compared to the C₈(diamond). To understand this fact, in Figure 2 we have plotted the diagonal elements of the \mathbf{R} matrix that is obtained by performing the QR decomposition of the $\tilde{\rho}_{mn, \mathbf{R}_h}$ matrix. These diagonal elements reflect the importance of the various pivot points and explain why one has to include a much larger number of pivot points for C₈(diamond) to obtain comparable accuracy. The results show that some amount of experimentation is needed to decide on the appropriate number of c needed for a desired accuracy. However, the results in Figure 1 show that for both systems, the results converge roughly exponentially fast with the increase in c .

5.3. Accuracy with System Size. In Figure 3, we have plotted the energy/atom as a function of c for three systems Li₄H₄, Li₃₂H₃₂, and Li₁₀₈H₁₀₈. The errors/atom are quite




Figure 2. Diagonal elements of the \mathbf{R} matrix obtained by performing the QR decomposition of the $\tilde{\rho}_{mn, \mathbf{R}_h}$ matrix. The diagonal elements decay much faster for Li₄H₄ compared to the C₈(diamond).




Figure 3. Intensive error (Error/atom) is approximately independent of the size of the system for the same value of c .




Figure 1. (Left and middle) Error incurred due to rPS and THC approximations in Li₄H₄ and C₈(diamond) systems, respectively. In both systems, one sees that the accuracy of rPS is significantly higher than that of THC, and it can be explained by the quadratic error, we expect to get in the former as opposed to the linear error of the latter. A noteworthy observation is that for the same value of c , the error is smaller with the larger basis set (TZ) rather than the smaller basis set (DZ). In fact, when one plots the error versus N_χ rather than $c = N_\chi/N$ (as is done in the rightmost figure), one sees that the curves for the two basis sets nearly overlap, which shows that the error one makes is largely independent of the size of the basis set for a fixed number of fitting functions. Another observation is that the error is highly system-dependent (see text and Figure 2).

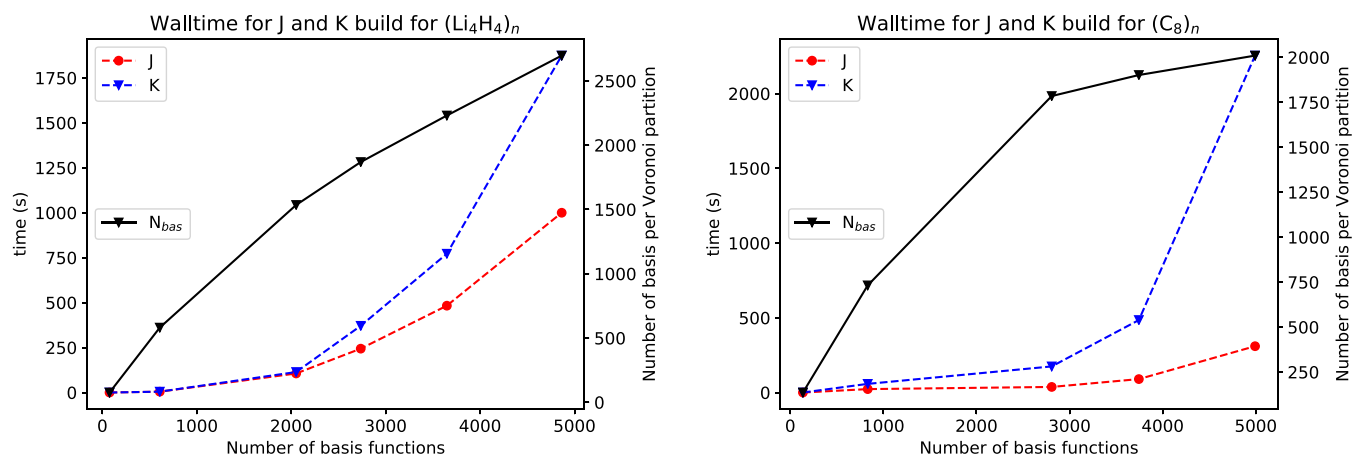


Figure 4. Wall time needed for a single build of the Coulomb and exchange matrix on nodes containing two Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz processors (red and blue curves). The systems used to obtain this figure are $(\text{Li}_4\text{H}_4)_n$ (lithium hydride) and $(\text{C}_8)_n$ (diamond), where $n = 1, 4, 8, 27, 36, 48$, and 64 for the lithium hydride system and $n = 1, 4, 8, 36$, and 48 for diamond. For lithium hydride, a $c = 4$ was used, and for diamond, $c = 6$ was used. These values of c are chosen to ensure that sub mE_h errors are obtained for the $n = 1$ calculation. The timings were normalized for that of 1 node; for example, the largest calculation for lithium hydride contained 64 unit cells and the figures show a timing of 1000 s for Coulomb build; however, the wall time was $1000/16 = 62.5$ s on 16 nodes. The black line shows the number of basis functions with a value $>10^{-8}$ on the Voronoi partition. Note that this number is not yet a constant, which leads to the super linear scaling of the Coulomb matrix construction. The relative cost of exchange is higher in diamond because a higher value of $c = 6$ as opposed to $c = 4$ for lithium hydride is used. Note that the cost of exchange increases linearly with c while the cost of the Coulomb matrix construction is independent of it.

similar for the three systems and in all cases decrease exponentially with c . The figure indicates that one can experiment on a smaller system to estimate the value of c needed for a desired accuracy and then the calculation can be scaled to a larger system size.

5.4. Cost of the Calculations. Figure 4 shows the cost of a single Coulomb and exchange build for lithium hydride and diamond for various system sizes. The scaling of the Coulomb matrix appears to be nonlinear mainly because the system contains highly diffuse functions that do not decay rapidly enough. For example, the black line in both the graphs shows that the number of basis functions that have a non-negligible value on the grid points of the Voronoi partition increases with the size of the system. Asymptotically, this number becomes a constant, but because of the presence of extremely diffuse functions, this does not happen even for the largest system that contains more than 1000 electrons. In our current implementation, the grid spacing is determined by the sharpest Gaussian, while the number of basis functions with non-negligible value is determined by the most diffuse function. This naturally leads to a suboptimal performance in the Coulomb build. In the software packages CP2K⁶⁰ and PySCF,^{96,97} a technique known as multigrid (different from the multigrid approach used in the solution of partial differential equation) is used, whereby several grids are utilized from sparse to dense with sparse grid used to represent diffuse basis functions and the dense for sharp basis function. This allows one to reach the linear scaling regime rapidly. Another option in this context is to use (and further develop) the Gaussian basis set recently introduced by Ye and Berkelbach⁹⁸ that has fewer diffuse functions, which leads to fewer linear dependency problems and as a side benefit would allow one to reach the linear scaling regime more quickly without having to implement multigrid approaches. Finally, one can also replace diffuse Gaussians with plane waves which will not only allow us to more systematically increase the accuracy of the calculation but will also help reduce the overall cost by reaching the linear

scaling regime sooner. We plan to pursue these lines of work in a future publication.

Table 4 shows the timings for the various steps of the calculation for a series of Li–H solids of increasing super cell size. All calculations are performed using computational nodes containing two Intel(R) Xeon(R) CPU E5-2680 v3 @ 2.50GHz processors with a total number of threads equaling 24. In all calculations, a value of $c = 4$ was used. From these calculations, it is clear that the cost of performing QRCP to obtain the interpolation points increases quite rapidly with the system size. This is particularly true because the QRCP is not parallelized over the number of processors and only OMP is used. We are optimistic that using Scalapack, one can reduce the cost of this step. Besides this step, the cost of evaluating $W(\mathbf{R}_g, \mathbf{R}'_g)$ also increases rather steeply. A significant part of the cost is incurred due to the calls to the MPI_All-to-all function which is given in the bracket of column 6. The relative cost of evaluation of the exchange matrix never increases beyond 2 compared to the Coulomb matrix even for large systems. However, for even larger systems, this relative cost will keep increasing because the scaling of the two steps is not the same. Number of nodes needed increases quadratically with the size of the system because the memory cost of storing $V(\mathbf{R}_g, \mathbf{R})$ is $N_x N_g$, for example, 16 nodes were needed to perform the largest calculation in Table 4.

6. CONCLUSIONS

In this work, we have presented an algorithm for reducing the cost of evaluation of the exchange matrix such that it is only slightly more expensive than the evaluation of the Coulomb matrix. We have done so by reducing the prefactor without changing the scaling. Currently, the most expensive step of the algorithm is obtaining the interpolation points using QRCP which is cubic scaling with a fairly large overhead. The QRCP cost can be decreased in two ways. First, one can parallelize the current algorithm using Scalapack to effectively make use of both the MPI and OMP, while currently we only make use of OMP. Second, we can reduce the computational cost of the

algorithm as follows. In the first step of the current algorithm, we obtain a set of interpolation points from each atom, and in the second step, we perform a large QRCP calculation to identify the most important points in this set. The second step is by far the dominant cost and this can be replaced by selecting points using a criterion different from the QRCP method. For example, the centroid Voronoi tessellation (CVT) with K-mean algorithm³⁷ can be used to replace the last step. We expect this to be an effective approach because CVT is only used for selecting a subset of points that are already close to optimal.

In addition, the algorithm can be extended in several ways. First, we can use a mixture of plane wave and Gaussian basis function by removing diffuse Gaussians and replacing them with a set of plane wave basis. By increasing the number of plane wave basis we expect to be able to reach basis set limit when calculating energy differences (e.g., atomization energies). It remains to be seen how effective rPS algorithm remains at selecting pivot points to represent the product density of this mixed basis set. Also the number of basis functions will increase rapidly with the threshold and one will most likely have to resort to direct diagonalization. Second, our algorithm has similarities with the discontinuous Galerkin method,^{86,99} and it is possible to obtain a set of Galerkin basis that are localized to a given Voronoi partition. This will have a significant advantage that the fitting functions will also be perfectly localized to an atomic domain and thus the cubic scaling of the QRCP step and the evaluation of the $W(\mathbf{R}_g, \mathbf{R}_g')$ will be reduced to linear and quadratic scaling, respectively. Third, we can implement k -point symmetry which is expected to reduce the cost of the exchange evaluation to $O(n_k \ln(n_k))$, where n_k is the number of k -points.¹⁰¹ Fourth, if one wants to avoid the use of pseudo-potential, then sharp Gaussians are needed. To include such Gaussians one can split the solution of Poisson's equation between real space and reciprocal space,⁸⁰ with real-space calculations requiring the evaluation of mixed-Gaussians-plane wave integrals which we have recently developed.¹⁰⁰ Alternatively, one has to use an irregular grid, where wavelet basis with multiresolution analysis is an attractive approach. Several wavelet bases such as interpolating wavelets, Coiflets, and Gausslets allow one to use the diagonal approximation that is needed for our algorithm to work. Finally, it is possible to use the rPS integrals in correlated calculations for periodic and molecular systems.

AUTHOR INFORMATION

Corresponding Author

Sandeep Sharma – Department of Chemistry, University of Colorado Boulder, Boulder, Colorado 80309, United States;
orcid.org/0000-0002-6598-8887; Email: sanshar@gmail.com

Authors

Alec F. White – Quantum Simulation Technologies, Inc., Boston, Massachusetts 02135, United States

Gregory Beylkin – Department of Applied Mathematics, University of Colorado, Boulder, Colorado 80309, United States

Complete contact information is available at:
<https://pubs.acs.org/10.1021/acs.jctc.2c00720>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

A.F.W. and S.S. were partly supported by the DOE grant DE-SC0022385. S.S. was also partly supported by NSF Career award CHE-2145209. The authors thank Toru Shiozaki and Daniel Moberg for helpful discussions.

REFERENCES

- (1) 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*, 1–8.
- (2) Kümmel, S.; Kronik, L. Orbital-dependent density functionals: Theory and applications. *Rev. Mod. Phys.* **2008**, *80*, 3–60.
- (3) Finazzi, E.; Di Valentin, C.; Pacchioni, G.; Selloni, A. Excess electron states in reduced bulk anatase TiO₂: Comparison of standard GGA, GGA + U, and hybrid DFT calculations. *J. Chem. Phys.* **2008**, *129*, No. 154113.
- (4) Xiao, H.; Tahir-Kheli, J.; Goddard, W. A. Accurate band gaps for semiconductors from density functional theory. *J. Phys. Chem. Lett.* **2011**, *2*, 212–217.
- (5) Basera, P.; Saini, S.; Arora, E.; Singh, A.; Kumar, M.; Bhattacharya, S. Stability of non-metal dopants to tune the photo-absorption of TiO₂ at realistic temperatures and oxygen partial pressures: A hybrid DFT study. *Sci. Rep.* **2019**, *9*, No. 11427.
- (6) Almlöf, J.; Faegri, K.; Korsell, K. Principles for a direct SCF approach to LCAO-MO ab-initio calculations. *J. Comput. Chem.* **1982**, *3*, 385–399.
- (7) White, C. A.; Johnson, B. G.; Gill, P. M. W.; Head-Gordon, M. The continuous fast multipole method. *Chem. Phys. Lett.* **1994**, *230*, 8–16.
- (8) Kudin, K. N.; Scuseria, G. E. *A fast multipole method for periodic systems with arbitrary unit cell geometries*, Technical report, **1998**.
- (9) Challacombe, M.; Schwegler, E. Linear scaling computation of the Fock matrix. *J. Chem. Phys.* **1997**, *106*, 5526.
- (10) 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.
- (11) Ko, H.-Y.; Jia, J.; Santra, B.; Wu, X.; Car, R.; DiStasio, R. A., Jr. Enabling Large-Scale Condensed-Phase Hybrid Density Functional Theory Based Ab Initio Molecular Dynamics. 1. Theory, Algorithm, and Performance. *J. Chem. Theory Comput.* **2020**, *16*, 3757–3785.
- (12) Goedecker, S. Linear scaling electronic structure methods. *Rev. Mod. Phys.* **1999**, *71*, 1085–1123.
- (13) Früchtl, H. A.; Kendall, R. A.; Harrison, R. J.; Dyall, K. G. An implementation of RI-SCF on parallel computers. *Int. J. Quantum Chem.* **1998**, *64*, 63–69.
- (14) Weigend, Florian. A fully direct RI-HF algorithm: Implementation, optimised auxiliary basis sets, demonstration of accuracy and efficiency. *Phys. Chem. Chem. Phys.* **2002**, *4*, 4285–4291.
- (15) Sierka, M.; Hogekamp, A.; Ahlrichs, R. Fast evaluation of the Coulomb potential for electron densities using multipole accelerated resolution of identity approximation. *J. Chem. Phys.* **2003**, *118*, 9136.
- (16) Sodt, A.; Subotnik, J. E.; Head-Gordon, M. Linear scaling density fitting. *J. Chem. Phys.* **2006**, *125*, No. 194109.
- (17) Polly, R.; Werner, H. J.; Manby, F. R.; Knowles, P. J. Fast Hartree-Fock theory using local density fitting approximations. *Mol. Phys.* **2004**, *102*, 2311–2321.
- (18) Alex, S.; Head-Gordon, M. Hartree-Fock exchange computed using the atomic resolution of the identity approximation. *J. Chem. Phys.* **2008**, *128*, No. 104106.
- (19) Manzer, S. F.; Epifanovsky, E.; Head-Gordon, M. Efficient implementation of the pair atomic resolution of the identity approximation for exact exchange for hybrid and range-separated density functionals. *J. Chem. Theory Comput.* **2015**, *11*, 518–527.
- (20) 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*, No. 024113.

- (21) Dunlap, B. I. Robust variational fitting: Gaspar's variational exchange can accurately be treated analytically. *J. Mol. Struct.:THEO-CHEM* **2000**, 501–502, 221–228.
- (22) Dunlap, B. I. Robust and variational fitting: Removing the four-center integrals from center stage in quantum chemistry. *J. Mol. Struct.:THEO-CHEM* **2000**, 529, 37–40.
- (23) Hollman, D. S.; Schaefer, H. F.; Valeev, E. F. Fast construction of the exchange operator in an atom-centred basis with concentric atomic density fitting. *Mol. Phys.* **2017**, 115, 2065–2076.
- (24) Beebe, N. H. F.; Linderberg, J. Simplifications in the generation and transformation of two-electron integrals in molecular calculations. *Int. J. Quantum Chem.* **1977**, 12, 683–705.
- (25) Aquilante, F.; Lindh, R.; Pedersen, T. B. Unbiased auxiliary basis sets for accurate two-electron integral approximations. *J. Chem. Phys.* **2007**, 127, No. 114107.
- (26) Friesner, R. A. Solution of self-consistent field electronic structure equations by a pseudospectral method. *Chem. Phys. Lett.* **1985**, 116, 39–43.
- (27) Neese, F.; Wennmohs, F.; Hansen, A.; Becker, U. Efficient, approximate and parallel Hartree-Fock and hybrid DFT calculations. A 'chain-of-spheres' algorithm for the Hartree-Fock exchange. *Chem. Phys.* **2009**, 356, 98–109.
- (28) Laqua, H.; Thompson, T. H.; Kussmann, J.; Ochsenfeld, C. Highly efficient, linear-scaling seminumerical exact-exchange method for graphic processing units. *J. Chem. Theory Comput.* **2020**, 16, 1456.
- (29) Laqua, H.; Kussmann, J.; Ochsenfeld, C. Efficient and linear-scaling seminumerical method for local hybrid density functionals. *J. Chem. Theory Comput.* **2018**, 14, 3451.
- (30) Liu, F.; Kong, J. Efficient computation of exchange energy density with Gaussian basis functions. *J. Chem. Theory Comput.* **2017**, 13, 2571.
- (31) Hilke, B.; Kaup, M. Efficient self-consistent implementation of local hybrid functional. *J. Chem. Theory Comput.* **2015**, 11, 1540.
- (32) Hohenstein, E. G.; Parrish, R. M.; Sherrill, C. D.; Martínez, T. J. Communication: Tensor hypercontraction. III. Least-squares tensor hypercontraction for the determination of correlated wavefunctions. *J. Chem. Phys.* **2012**, 137, No. 221101.
- (33) Parrish, R. M.; Hohenstein, E. G.; Martínez, T. J.; Sherrill, C. D. Tensor hypercontraction. II. Least-squares renormalization. *J. Chem. Phys.* **2012**, 137, No. 224106.
- (34) Hohenstein, E. G.; Parrish, R. M.; Martínez, T. J. Tensor hypercontraction density fitting. I. Quartic scaling second- and third-order Møller-Plesset perturbation theory. *J. Chem. Phys.* **2012**, 137, 1085.
- (35) 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.
- (36) Hu, W.; Lin, L.; Yang, C. Interpolative Separable Density Fitting Decomposition for Accelerating Hybrid Density Functional Calculations with Applications to Defects in Silicon. *J. Chem. Theory Comput.* **2017**, 13, 5420–5431.
- (37) Dong, K.; Hu, W.; Lin, L. Interpolative Separable Density Fitting through Centroidal Voronoi Tessellation with Applications to Hybrid Functional Electronic Structure Calculations. *J. Chem. Theory Comput.* **2018**, 14, 1311–1320.
- (38) Engler, H. The behavior of the QR-factorization algorithm with column pivoting. *Appl. Math. Lett.* **1997**, 10, 7–11.
- (39) Lee, J.; Lin, L.; Head-Gordon, M. Systematically Improvable Tensor Hypercontraction: Interpolative Separable Density-Fitting for Molecules Applied to Exact Exchange, Second- and Third-Order Møller-Plesset Perturbation Theory. *J. Chem. Theory Comput.* **2020**, 16, 243–263.
- (40) Bowler, D. R.; Miyazaki, T. O(N) methods in electronic structure calculations. *Rep. Prog. Phys.* **2012**, 75, No. 036503.
- (41) Wu, X.; Selloni, A.; Car, R. Order- N implementation of exact exchange in extended insulating systems. *Phys. Rev. B* **2009**, 79, No. 085102.
- (42) Baer, R.; Neuhauser, D.; Rabani, E. Self-averaging stochastic kohn-sham density-functional theory. *Phys. Rev. Lett.* **2013**, 111, No. 106402.
- (43) Neuhauser, D.; Rabani, E.; Cytter, Y.; Baer, R. Stochastic Optimally Tuned Range-Separated Hybrid Density Functional Theory. *J. Phys. Chem. A* **2016**, 120, 3071–3078.
- (44) Lin, L. Adaptively Compressed Exchange Operator. *J. Chem. Theory Comput.* **2016**, 12, 2242–2249.
- (45) Todorova, T.; Seitsonen, A. P.; Hutter, J.; Kuo, I.-F. W.; Mundy, C. J. Molecular dynamics simulation of liquid water: Hybrid density functionals. *J. Phys. Chem. B* **2006**, 110, 3685–3691.
- (46) Guidon, M.; Hutter, J.; VandeVondele, J. Auxiliary Density Matrix Methods for Hartree-Fock Exchange Calculations. *J. Chem. Theory Comput.* **2010**, 6, 2348–2364.
- (47) Cohen, A. J.; Handy, N. C. Density functional generalized gradient calculations using Slater basis sets. *J. Chem. Phys.* **2002**, 117, 1470.
- (48) Chong, D. P.; Van Lenthe, E.; Van Gisbergen, S.; Baerends, E. J. Even-tempered slater-type orbitals revisited: From hydrogen to krypton. *J. Comput. Chem.* **2004**, 25, 1030–1036.
- (49) Goedecker, S. Wavelets and their application for the solution of partial differential equations in physics. *Max-Planck Institute for Solid State Research*, **2009**.
- (50) Yanai, T.; Fann, G. I.; Beylkin, G.; Harrison, R. J. Multiresolution quantum chemistry in multiwavelet bases: excited states from time-dependent Hartree-Fock and density functional theory via linear response. *Phys. Chem. Chem. Phys.* **2015**, 17, 31405–31416.
- (51) Genovese, L.; Neelov, A.; Goedecker, S.; Deutsch, T.; Ghasemi, S. A.; Willand, A.; Caliste, D.; Zilberberg, O.; Rayson, M.; Bergman, A.; Schneider, R. Daubechies wavelets as a basis set for density functional pseudopotential calculations. *J. Chem. Phys.* **2008**, 129, No. 014109.
- (52) White, S. R.; Miles Stoudenmire, E. Multisliced gausslet basis sets for electronic structure. *Phys. Rev. B* **2019**, 99, No. 081110.
- (53) Brakstad, A.; Wind, P.; Jensen, S. R.; Frediani, L.; Hopmann, K. H. Multiwavelets applied to metal-ligand interactions: Energies free from basis set errors. *J. Chem. Phys.* **2021**, 154, No. 214302.
- (54) Blum, V.; Gehrke, R.; Hanke, F.; Havu, P.; Havu, V.; Ren, X.; Reuter, K.; Scheffler, M. Ab initio molecular simulations with numeric atom-centered orbitals. *Comput. Phys. Commun.* **2009**, 180, 2175–2196.
- (55) Andrade, X.; Strubbe, U.; De Giovannini, U.; Larsen, A.; Oliveira, M. J. T.; Alberdi-Rodriguez, J.; Varas, A.; Theophilou, I.; Helbig, N.; Verstraete, M. J.; Stella, L.; Nogueira, F.; Aspuru-Guzik, A.; Castro, A.; Marques, M. A. L.; Rubio, A. Real-space grids and the Octopus code as tools for the development of new simulation approaches for electronic systems. *Phys. Chem. Chem. Phys.* **2015**, 17, 31371–31396.
- (56) Baye, D.; Dohet-Eraly, J. Confined helium on Lagrange meshes. *Phys. Chem. Chem. Phys.* **2015**, 17, 31417–31426.
- (57) Enkovaara, J.; Rostgaard, C.; Mortensen, J. J.; Chen, J.; Dulak, M.; Ferrighi, L.; Gavnholt, J.; Glinsvad, C.; Haikola, V.; Hansen, H. A.; Kristoffersen, H. H.; Kuisma, M.; Larsen, A. H.; Lehtovaara, L.; Ljungberg, M.; Lopez-Acevedo, O.; Moses, P. G.; Ojanen, J.; Olsen, T.; Petzold, V.; Romero, N. A.; Stausholm-Møller, J.; Strange, M.; Tritsaris, G. A.; Vanin, M.; Walter, M.; Hammer, B.; Häkkinen, H.; Madsen, G. K. H.; Nieminen, R. M.; Nørskov, J. K.; Puska, M.; Rantala, T. T.; Schiøtz, J.; Thygesen, K. S.; Jacobsen, K. W. Electronic structure calculations with GPAW: a real-space implementation of the projector augmented-wave method. *J. Phys.: Condens. Matter* **2010**, 22, No. 253202.
- (58) Ye, H. Z.; Berkelbach, T. C. Fast periodic Gaussian density fitting by range separation. *J. Chem. Phys.* **2021**, 154, No. 131104.
- (59) Irmeler, A.; Burow, A. M.; Pauly, F. Robust Periodic Fock Exchange with Atom-Centered Gaussian Basis Sets. *J. Chem. Theory Comput.* **2018**, 14, 4567–4580.
- (60) Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khaliullin, R. Z.; Schütt, O.

- Schiffmann, F.; Golze, D.; Wilhelm, J.; Chulkov, S.; Bani-Hashemian, M. H.; Weber, V.; Borštnik, U.; Taillefumier, M.; Jakobovits, A. S.; Lazzaro, A.; Pabst, H.; Müller, T.; Schade, R.; Guidon, M.; Andermatt, S.; Holmberg, N.; Schenter, G. K.; Hehn, A.; Bussy, A.; Belleflamme, F.; Tabacchi, G.; Glöß, A.; Lass, M.; Bethune, I.; Mundy, C. J.; Plessl, C.; Watkins, M.; VandeVondele, J.; Krack, M.; Hutter, J. CP2K: An electronic structure and molecular dynamics software package -Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **2020**, *152*, No. 194103.
- (61) Paier, J.; Diaconu, C. V.; Scuseria, G. E.; et al. Accurate Hartree-Fock energy of extended systems using large Gaussian basis sets. *Phys. Rev. B* **2009**, *80*, 1–8.
- (62) Causa, M.; Dovesi, R.; Orlando, R.; Pisani, Cesare.; Saunders, V. R. Treatment of the exchange interactions in Hartree-Fock LCAO calculations of periodic systems. *J. Phys. Chem. A* **1988**, *92*, 909–913.
- (63) Causá, M.; Dovesi, R.; Orlando, R.; Pisani, C.; Saunders, V. R. Treatment of the exchange interactions in Hartree-Fock linear combination of atomic orbital calculations of periodic systems. *J. Phys. Chem. B* **1988**, *92*, 909–913.
- (64) Kudin, K. N.; Scuseria, G. E. Revisiting infinite lattice sums with the periodic fast multipole method. *J. Chem. Phys.* **2004**, *121*, 2886–2890.
- (65) Kudin, K. N.; Scuseria, G. E. Linear-scaling density-functional theory with Gaussian orbitals and periodic boundary conditions: Efficient evaluation of energy and forces via the fast multipole method. *Phys. Rev. B* **2000**, *61*, 16440–16453.
- (66) Sun, Q. Exact exchange matrix of periodic hartree-fock theory for all-electron simulations. 2020, arXiv:2012.07929. arXiv.org e-Print archive. <https://arxiv.org/abs/2012.07929>.
- (67) Lee, J.; Morales, M. A.; Malone, F. D. A phaseless auxiliary-field quantum Monte Carlo perspective on the uniform electron gas at finite temperatures: Issues, observations, and benchmark study. *J. Chem. Phys.* **2021**, *154*, No. 064109.
- (68) Kresse, G.; Furthmüller, J. effective iterative schemes for ab initio total-energy calculations using a plane-wave basis set. *Phys. Rev. B* **1996**, *54*, No. 11169.
- (69) Kresse, G.; Furthmüller, J. Efficiency of ab-initio total energy calculations for metals and semiconductors using a plane-wave basis set. *Comput. Mater. Sci.* **1996**, *6*, 15–50.
- (70) Giannozzi, P.; Andreussi, O.; Brumme, T.; Bunau, O.; Buongiorno Nardelli, M.; Calandra, M.; Car, R.; Cavazzoni, C.; Ceresoli, D.; Cococcioni, M.; Colonna, N.; Carnimeo, L.; Dal Corso, A.; de Gironcoli, S.; Delugas, P.; DiStasio, R. A.; Ferretti, A.; Floris, A.; Fratesi, G.; Fugallo, G.; Gebauer, R.; Gerstmann, U.; Giustino, F.; Gorni, T.; Jia, J.; Kawamura, M.; Ko, H.-Y.; Kokalj, A.; Küçükbenli, E.; Lazzeri, M.; Marsili, M.; Marzari, N.; Mauri, F.; Nguyen, N. L.; Nguyen, H.-V.; de-la Roza, A. Otero.; Paulatto, L.; Poncé, S.; Rocca, D.; Sabatini, R.; Santra, B.; Schlipf, M.; Seitsonen, A. P.; Smogunov, A.; Timrov, I.; Thonhauser, T.; Umari, P.; Vast, N.; Wu, X.; Baroni, S. Advanced capabilities for materials modelling with quantum ESPRESSO. *J. Phys.: Condens. Matter* **2017**, *29*, No. 465901.
- (71) Gonze, X.; Beuken, J.-M.; Caracas, R.; Detraux, F.; Fuchs, M.; Rignanese, G.-M.; Sindic, L.; Verstraete, M.; Zerah, G.; Jollet, F.; Torrent, M.; Roy, A.; Mikami, M.; Ghosez, Ph.; Raty, J.-Y.; Allan, D. C. First-principles computation of material properties: the abinit software project. *Comput. Mater. Sci.* **2002**, *25*, 478–492.
- (72) Holzwarth, N. A. W.; Matthews, G. E.; Dunning, R. B.; Tackett, A. R.; Zeng, Y. Comparison of the projector augmented-wave, pseudopotential, and linearized augmented-plane-wave formalisms for density-functional calculations of solids. *Phys. Rev. B* **1997**, *55*, 2005–2017.
- (73) Füsti-Molnár, L.; Pulay, P. The Fourier transform coulomb method: Efficient and accurate calculation of the coulomb operator in a gaussian basis. *J. Chem. Phys.* **2002**, *117*, 7827–7835.
- (74) Füsti-Molnár, L.; Pulay, P. Accurate molecular integrals and energies using combined plane wave and gaussian basis sets in molecular electronic structure theory. *J. Chem. Phys.* **2002**, *116*, 7795–7805.
- (75) Sun, Q.; Berkelbach, T. C.; McClain, J. D.; Chan, G. K. L. Gaussian and plane-wave mixed density fitting for periodic systems. *J. Chem. Phys.* **2017**, *147*, No. 164119.
- (76) Vandevondele, J.; Krack, M.; Mohamed, F.; Parrinello, M.; Chassaing, T.; Hutter, J. Quickstep: Fast and accurate density functional calculations using a mixed Gaussian and plane waves approach. *Comput. Phys. Commun.* **2005**, *167*, 103–128.
- (77) Qiu, Y.; White, S. R. Hybrid gausslet/Gaussian basis sets. *J. Chem. Phys.* **2021**, *155*, No. 184107.
- (78) White, S. R. Hybrid grid/basis set discretizations of the Schrödinger equation. *J. Chem. Phys.* **2017**, *147*, No. 244102.
- (79) Jones, J. R.; Rouet, F. H.; Lawler, K. V.; Vecharynski, E.; Ibrahim, K. Z.; Williams, S.; Abeln, B.; Yang, C.; McCurdy, W.; Haxton, D. J.; Li, X. S.; Rescigno, T. N. An efficient basis set representation for calculating electrons in molecules. *Mol. Phys.* **2016**, *114*, 2014–2028.
- (80) Sharma, S.; Beylkin, G. Efficient evaluation of two-center gaussian integrals in periodic systems. *J. Chem. Theory Comput.* **2021**, *17*, 3916–3922.
- (81) Pierce, K.; Rishi, V.; Valeev, E. F. Robust Approximation of Tensor Networks: Application to Grid-Free Tensor Factorization of the Coulomb Interaction. *J. Chem. Theory Comput.* **2021**, *17*, 2217–2230.
- (82) Mahajan, A.; Lee, J.; Sharma, S. Selected configuration interaction wave functions in phaseless auxiliary field quantum monte carlo. *J. Chem. Phys.* **2022**, *156*, No. 174111.
- (83) Whitlock, P. A.; Ceperley, D. M.; Chester, G. V.; Kalos, M. H. Properties of liquid and solid he 4. *Phys. Rev. B* **1979**, *19*, 5598.
- (84) Ceperley, D. M.; Kalos, M. H. Quantum Many-Body Problems. In *Monte Carlo Methods in Statistical Physics*; Springer, 1986; pp 145–194.
- (85) Rothstein, S. M. A survey on pure sampling in quantum monte carlo methods. *Can. J. Chem.* **2013**, *91*, 505–510.
- (86) Faulstich, F. M.; Wu, X.; Lin, L. Discontinuous galerkin method with voronoi partitioning for quantum simulation of chemistry. 2020, arXiv:2011.00367. arXiv.org e-Print archive. <https://arxiv.org/abs/2011.00367>.
- (87) 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.
- (88) Liberty, E.; Woolfe, F.; Martinsson, P.-G.; Rokhlin, V.; Tytgert, M. Randomized algorithms for the low-rank approximation of matrices. *Proc. Natl. Acad. Sci. U.S.A.* **2007**, *104*, 20167–20172.
- (89) Blackford, L. S.; Choi, J.; Cleary, A.; D’Azevedo, E.; Demmel, J.; Dhillon, I.; Dongarra, J.; Hammarling, S.; Henry, G.; Petitet, A.; Stanley, K.; Walker, D.; Whaley, R. C. *ScaLAPACK Users’ Guide*; Society for Industrial and Applied Mathematics: Philadelphia, PA, 1997.
- (90) Sundararaman, R.; Arias, T. A. Regularization of the Coulomb singularity in exact exchange by Wigner-Seitz truncated interactions: Towards chemical accuracy in nontrivial systems. *Phys. Rev. B* **2013**, *87*, No. 165122.
- (91) Spencer, J.; Alavi, A. Efficient calculation of the exact exchange energy in periodic systems using a truncated Coulomb potential. *Phys. Rev. B* **2008**, *77*, No. 193110.
- (92) Guidon, M.; Hutter, J.; VandeVondele, J. Robust periodic Hartree-Fock exchange for large-scale simulations using Gaussian basis sets. *J. Chem. Theory Comput.* **2009**, *5*, 3010–3021.
- (93) Tymczak, C. J.; Weber, V. T.; Schwegler, E.; Challacombe, M. Linear scaling computation of the Fock matrix. VIII. Periodic boundaries for exact exchange at the Gamma point. *J. Chem. Phys.* **2005**, *122*, No. 124105.
- (94) Hartwigsen, C.; Goedecker, S.; Hutter, J. Relativistic separable dual-space gaussian pseudopotentials from h to rn. *Phys. Rev. B* **1998**, *58*, 3641–3662.
- (95) Goedecker, S.; Teter, M.; Hutter, J. Separable dual-space gaussian pseudopotentials. *Phys. Rev. B* **1996**, *54*, 1703–1710.
- (96) Sun, Q.; Zhang, X.; Banerjee, S.; Bao, P.; Barbry, M.; Blunt, N. S.; Bogdanov, N. A.; Booth, G. H.; Chen, J.; Cui, Z.-H.; Eriksen, J. J.;

Gao, Y.; Guo, S.; Hermann, J.; Hermes, M. R.; Koh, K.; Koval, P.; Lehtola, S.; Li, Z.; Liu, J.; Mardirossian, N.; McClain, J. D.; Motta, M.; Mussard, B.; Pham, H. Q.; Pulkin, A.; Purwanto, W.; Robinson, P. J.; Ronca, E.; Sayfutyarova, E. R.; Scheurer, M.; Schurkus, H. F.; Smith, J. E. T.; Sun, C.; Sun, S.-N.; Upadhyay, S.; Wagner, L. K.; Wang, X.; White, A.; Whitfield, J. D.; Williamson, M. J.; Wouters, S.; Yang, J.; Yu, J. M.; Zhu, T.; Berkelbach, T. C.; Sharma, S.; Sokolov, A. Y.; Chan, G. K.-L. Recent developments in the PySCF program package. *J. Chem. Phys.* **2020**, *153*, No. 024109.

(97) Sun, Q.; Berkelbach, T. C.; Blunt, N. S.; Booth, G. H.; Guo, S.; Li, Z.; Liu, J.; McClain, J. D.; Sayfutyarova, E. R.; Sharma, S.; Wouters, S.; Kin-Lic Chan, G. PySCF: the Python-based simulations of chemistry framework. *WIREs Comput. Mol. Sci.* **2018**, *8*, No. e1304.

(98) Ye, H.-Z.; Berkelbach, T. C. Correlation-Consistent Gaussian Basis Sets for Solids Made Simple. *J. Chem. Theory Comput.* **2022**, *18*, 1595.

(99) McClean, J. R.; Faulstich, F. M.; Zhu, Q.; O’Gorman, B.; Qiu, Y.; White, S. R.; Babbush, R.; Lin, L. Discontinuous Galerkin discretization for quantum simulation of chemistry. *New J. Phys.* **2020**, *22*, No. 093015.

(100) Beylkin, G.; Sharma, S. A fast algorithm for computing the boys function. *J. Chem. Phys.* **2021**, *155*, No. 174117.

(101) Wu, K.; Qin, X.; Hu, W.; Yang, J. Low-Rank Approximation Accelerated Plane-Wave Hybrid Functional Calculations with k-point Sampling. *J. Chem. Theory Comput.* **2022**, *18*, 206.