

COMPUTING LOCALIZED REPRESENTATIONS OF THE KOHN–SHAM SUBSPACE VIA RANDOMIZATION AND REFINEMENT*

ANIL DAMLE[†], LIN LIN[‡], AND LEXING YING[§]

Abstract. Localized representation of the Kohn–Sham subspace plays an important role in quantum chemistry and materials science. The recently developed selected columns of the density matrix (SCDM) method [A. Damle, L. Lin, and L. Ying, *J. Chem. Theory Comput.*, 11 (2015), pp. 1463–1469] is a simple and robust procedure for finding a localized representation of a set of Kohn–Sham orbitals from an insulating system. The SCDM method allows the direct construction of a well conditioned (or even orthonormal) and localized basis for the Kohn–Sham subspace. The SCDM algorithm avoids the use of an optimization procedure and does not depend on any adjustable parameters. The most computationally expensive step of the SCDM method is a column pivoted QR factorization that identifies the important columns for constructing the localized basis set. In this paper, we develop a two-stage approximate column selection strategy to find the important columns at much lower computational cost. We demonstrate the effectiveness of this process using the dissociation process of a BH_3NH_3 molecule, an alkane chain, and a supercell with 256 water molecules. Numerical results for the large collection of water molecules show that the two-stage localization procedure can be more than 30 times faster than the original SCDM algorithm and compares favorably with the popular Wannier90 package.

Key words. localization, Kohn–Sham density functional theory, selected columns of the density matrix, column pivoted QR factorization

AMS subject classifications. 65Z05, 65F30

DOI. 10.1137/16M1098589

1. Introduction. Kohn–Sham density functional theory (DFT) [21, 24] is the most widely used electronic structure theory for molecules and systems in condensed phase. The Kohn–Sham orbitals (a.k.a. Kohn–Sham wavefunctions) are eigenfunctions of the Kohn–Sham Hamiltonian and are generally delocalized, i.e., each orbital has significant magnitude across the entire computational domain. Consequently, the information about atomic structure and chemical bonding, which is often localized in real space, may be difficult to interpret directly from Kohn–Sham orbitals. The connection between the delocalized orbitals and localized ones can be established through a *localization* procedure, which has been realized by various numerical methods in the

*Submitted to the journal’s Computational Methods in Science and Engineering section October 12, 2016; accepted for publication (in revised form) May 31, 2017; published electronically December 12, 2017.

<http://www.siam.org/journals/sisc/39-6/M109858.html>

Funding: The first author’s work was partially supported by a National Science Foundation mathematical sciences postdoctoral research fellowship under grant DMS-1606277. The second author’s work was supported by DOE Scientific Discovery through the Advanced Computing (SciDAC) program, the DOE Center for Applied Mathematics for Energy Research Applications (CAMERA) program, and an Alfred P. Sloan fellowship. The third author’s work was supported by the National Science Foundation under awards DMS-1328230 and DMS-1521830 and the U.S. Department of Energy’s Advanced Scientific Computing Research program under award DE-FC02-13ER26134/DE-SC0009409.

[†]Department of Computer Science, Cornell University, Ithaca, NY 14850 (damle@cornell.edu).

[‡]Department of Mathematics, University of California, Berkeley, Berkeley, CA 94720, and Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720 (linlin@math.berkeley.edu).

[§]Department of Mathematics and Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA 94305 (lexing@math.stanford.edu).

literature [15, 30, 29, 18, 10, 14, 34, 2]. The common goal of these methods is to find a set of orbitals that are localized in real space and span the Kohn–Sham invariant subspace, defined as the subspace as that spanned by the Kohn–Sham orbitals. For simplicity, we restrict our discussion to isolated systems and omit the discussion of Brillouin zone sampling in this manuscript.

Mathematically, the problem of finding a localized representation of the Kohn–Sham invariant subspace can be formulated as follows. Assume the collection of Kohn–Sham orbitals are discretized in the real space representation as a tall and skinny matrix $\Psi \in \mathbb{C}^{N \times n_e}$ with orthonormal columns and where $N \gg n_e$. We seek to compute a unitary transformation $Q \in \mathbb{C}^{n_e \times n_e}$ such that the columns of $\Phi = \Psi Q$ are *localized*, i.e., each column of Φ becomes a sparse vector with spatially localized support after truncating entries with relative magnitude smaller than a prescribed threshold. Here, N is the number of grid points in the discrete real space representation of each Kohn–Sham orbital, and n_e is the number of orbitals. In the absence of spin degeneracy, n_e is also the number of electrons in the system.

For a general matrix Ψ it may not be possible to construct such a Q and obtain Φ with the desired structure. However, when Ψ represents a collection of Kohn–Sham orbitals of an insulating system, such localized orbitals generally exist. An important exception is topological insulators with nonvanishing Chern numbers [20, 7]—here we restrict our discussion to systems where localized functions are known to exist. Their construction can be justified physically by the “nearsightedness” principle for electronic matter with a finite so-called HOMO-LUMO gap [23, 36]. The nearsightedness principle can be more rigorously stated as the single particle density matrix being exponentially localized along the off-diagonal direction in its real space representation [4, 23, 6, 9, 8, 33, 25].

The recently developed selected columns of the density matrix (SCDM) procedure [10] provides a simple, accurate, and robust way of constructing localized orbitals. Unlike many existing methods [15, 30, 14, 34, 32], the SCDM method requires no initial guess and does not involve a nonconvex optimization procedure. The core ideas behind it also readily extend to systems with Brillouin zone sampling [11]. The SCDM procedure constructs localized orbitals directly from a column selection procedure implicitly applied to the density matrix. Hence, the locality of the basis is a direct consequence of the locality of the density matrix. The SCDM method can be efficiently performed via a single column pivoted QR (QRCP) factorization. Since efficient implementation of the QRCP is available in both serial and parallel computational environments through the LAPACK [1] and ScaLAPACK [5] libraries, respectively, the SCDM method can be readily adopted by electronic structure software packages.

From a numerical perspective, the computational cost of a QRCP factorization scales as $\mathcal{O}(Nn_e^2)$. The basic form of QRCP [17] is not able to take full advantage of level 3 BLAS operations. Hence for matrices of the same size, the QRCP factorization can still be relatively expensive compared to level 3 BLAS operations such as general matrix-matrix multiplication (GEMM). The computational cost of the single QRCP is not necessarily an issue when the SCDM procedure is used as a postprocessing tool, but it is a potential concern when the SCDM procedure needs to be performed repeatedly. This, for instance, could occur in geometry optimization and molecular dynamics calculations with hybrid exchange–correlation functionals [37, 19], where a localized representation of the Kohn–Sham invariant subspace needs to be constructed in each step to reduce the large computational cost associated with the Fock exchange operator. In fact, [10] demonstrates how our existing SCDM algorithm

may be used to accelerate Hartree–Fock exchange computations. Therefore, here we focus on accelerating the SCDM computation itself.

Practically, any QRCP algorithm may be used within the SCDM procedure. In the serial setting, this includes recently developed methods based on using random projections to accelerate and block up the column selection procedure [28, 13]. In the massively parallel setting one may alternatively use the recently developed communication avoiding rank-revealing QR algorithm [12].

In this paper, we demonstrate that the computational cost of the SCDM procedure can be greatly reduced, to the extent that the column selection procedure is no longer the dominating factor in the SCDM calculation. This is based on the observation that SCDM does not really require the Q and R factors from the QRCP. In fact, only the pivots from the QRCP are needed. More specifically, we develop a two-stage column selection procedure that approximates the behavior of the existing SCDM procedure at a much lower computational cost. Asymptotically, the computational cost is only dominated by two matrix-matrix multiplications of the form ΨQ to construct the localized orbitals, at least one of which is needed in any localization procedure starting from the input Ψ matrix. Notably, the only adjustable parameters we introduce are an oversampling factor and truncation threshold, both of which may be picked without knowledge of the problem’s physical structure.

The approximate column selection procedure consists of two stages. First, we use a randomized procedure to select a set of candidate columns that may be used in an SCDM style localization procedure. The number of candidate columns is only $\mathcal{O}(n_e \log n_e)$ and is much smaller than N . We may use these candidate columns to quickly construct a basis for the subspace that is reasonably localized. In some cases, this fast randomized procedure may provide sufficiently localized columns. Otherwise, we propose a subsequent refinement procedure to improve the quality of the localized orbitals. This is achieved by using a series of QRCP factorizations for matrices of smaller sizes that may be performed concurrently. Numerical results for physical systems obtained from the Quantum ESPRESSO [16] package indicate that the two-stage procedure yields results that are nearly as good as using the columns selected by a full QRCP based SCDM procedure. For large systems, the computational time is reduced by more than one order of magnitude.

The remainder of this paper is organized as follows. In section 2 we present both a brief introduction to Kohn–Sham DFT and a summary of the existing SCDM algorithm. Section 3 discusses the new two-stage algorithm we propose and details both the randomized approximate localization stage and the refinement of the column selection. Finally, section 4 demonstrates the effectiveness of the algorithm for various molecules.

2. Preliminaries. For completeness we first provide a brief introduction to Kohn–Sham DFT and the SCDM procedure for finding a localized basis for the Kohn–Sham subspace.

2.1. Kohn–Sham density functional theory. For a given atomic configuration with M atoms at locations $\{R_I\}_{I=1}^M$, Kohn–Sham DFT solves the nonlinear eigenvalue problem

$$(2.1) \quad \begin{aligned} \hat{H}[\hat{\rho}; \{R_I\}] \hat{\psi}_i &= \varepsilon_i \hat{\psi}_i, \\ \hat{\rho}(\mathbf{r}) &= \sum_{i=1}^{n_e} |\hat{\psi}_i(\mathbf{r})|^2, \quad \int \hat{\psi}_i^*(\mathbf{r}) \hat{\psi}_j(\mathbf{r}) \, d\mathbf{r} = \delta_{ij}. \end{aligned}$$

For simplicity we omit the spin degeneracy. The number of electrons is n_e and the eigenvalues $\{\varepsilon_i\}$ are ordered nondecreasingly. The lowest n_e eigenvalues $\{\varepsilon_i\}_{i=1}^{n_e}$ are called the occupied state energies, and $\{\varepsilon_i\}_{j>n_e}$ are called the unoccupied state energies. We assume $\varepsilon_g := \varepsilon_{n_e+1} - \varepsilon_{n_e} > 0$. Here ε_{n_e} is often called the highest occupied molecular orbital (HOMO), ε_{n_e+1} the lowest unoccupied molecular orbital (LUMO), and hence ε_g the HOMO-LUMO gap. For extended systems, if ε_g is uniformly bounded away from zero as the system size increases, the quantum system is an insulating system [27]. The eigenfunctions $\{\hat{\psi}_i\}_{i=1}^{n_e}$ define the electron density $\hat{\rho}(\mathbf{r})$, which in turn defines the Kohn–Sham Hamiltonian

$$(2.2) \quad \hat{H}[\hat{\rho}; \{R_I\}] = -\frac{1}{2}\Delta + \hat{V}_c[\hat{\rho}] + \hat{V}_{\text{xc}}[\hat{\rho}] + V_{\text{ion}}[\{R_I\}].$$

Here Δ is the Laplacian operator for the kinetic energy of electrons,

$$\hat{V}_c[\hat{\rho}](\mathbf{r}) \equiv \int \frac{\hat{\rho}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}'$$

is the Coulomb potential, and \hat{V}_c depends linearly with respect to the electron density $\hat{\rho}$. $\hat{V}_{\text{xc}}[\hat{\rho}]$ depends nonlinearly with respect to $\hat{\rho}$ and characterizes the many body exchange and correlation effect. $V_{\text{ion}}[\{R_I\}]$ is an external potential depending explicitly on the ionic positions and describes the electron-ion interaction potential and is independent of $\hat{\rho}$. Because the eigenvalue problem (2.1) is nonlinear, it is often solved iteratively by a class of algorithms called self-consistent field iterations [27], until (2.1) reaches self-consistency.

In a finite dimensional discretization of (2.2), let N be the number of degrees of freedom. Using a large basis set such as the plane-wave basis set, we have $N = cn_e$ and c is a large constant that is often $10^2 \sim 10^4$. Due to this large constant, we explicitly distinguish N and n_e in the complexity analysis below. For large-scale systems, the cost for storing the Kohn–Sham orbitals is $\mathcal{O}(Nn_e)$, and the cost for computing them is generally $\mathcal{O}(Nn_e^2)$ and scales cubically with respect to n_e . In modern KSDFT calculations the Hartree–Fock exact exchange term is also often taken into account in the form of hybrid functionals [35, 3]. The computational cost for this step not only scales as $\mathcal{O}(Nn_e^2)$ but also has a large preconstant.

When the self-consistent solution of the Kohn–Sham equation is obtained, the existence of a finite HOMO-LUMO gap has important implications for the collective behavior of the occupied Kohn–Sham orbitals $\{\hat{\psi}_i\}_{i=1}^{n_e}$. Since any nondegenerate linear transformation of the set of Kohn–Sham orbitals yields exactly the same physical properties of a system, the physically relevant quantity is the subspace spanned by the Kohn–Sham orbitals $\{\hat{\psi}_i\}_{i=1}^{n_e}$. Various efforts [15, 30, 29, 18, 14, 34] have been made to utilize this fact and to find a set of localized orbitals that form a compressed representation of a Kohn–Sham subspace. In other words, we find a set of functions $\{\hat{\phi}_i\}_{i=1}^{n_e}$ whose span is the same as the span of $\{\hat{\psi}_i\}_{i=1}^{n_e}$. Compared to each Kohn–Sham orbital $\hat{\psi}_i$ which is delocalized in the real space, each compressed orbital $\hat{\phi}_i$ is often localized around an atom or a chemical bond. Hence working with $\hat{\phi}_i$'s can reduce both the storage and the computational cost.

Assume we have access to $\hat{\psi}_j(\mathbf{r})$'s evaluated at a set of discrete grid points $\{\mathbf{r}_i\}_{i=1}^N$. Let $\{\omega_i\}_{i=1}^N$ be a set of positive integration weights associated with the grid points $\{\mathbf{r}_i\}_{i=1}^N$; then the discrete orthonormality condition is given by

$$(2.3) \quad \sum_{i=1}^N \hat{\psi}_j(\mathbf{r}_i) \hat{\psi}_{j'}(\mathbf{r}_i) \omega_i = \delta_{jj'}.$$

Let $\hat{\psi}_j = [\hat{\psi}_j(\mathbf{r}_1), \hat{\psi}_j(\mathbf{r}_2), \dots, \hat{\psi}_j(\mathbf{r}_N)]^T$ be a column vector and $\hat{\Psi} = [\hat{\psi}_1, \dots, \hat{\psi}_{n_e}]$ be a matrix of size $N \times n_e$. We call $\hat{\Psi}$ the *real space representation* of the Kohn–Sham orbitals and define diagonal matrix $W = \text{diag}[\omega_1, \dots, \omega_N]$.

Our method requires the Kohn–Sham orbitals to be represented on a set of real space grid points. This is the case for a plane-wave basis set, as well as other representations such as finite differences, finite elements, and wavelets. For instance, if the Kohn–Sham orbitals are represented using the plane-wave basis functions, their real space representation can be obtained on a uniform grid efficiently with the fast Fourier transform technique and in this case ω_i takes the same constant value for all i . It is in this setting that our method is of particular interest. However, this procedure is also applicable to other basis sets such as Gaussian type orbitals or numerical atomic orbitals when a real space representation of the basis functions is readily available. Therefore, our method is amenable to most electronic structure software packages.

We define $\Psi = W^{\frac{1}{2}} \hat{\Psi}$ such that the discrete orthonormality condition in (2.3) becomes $\Psi^* \Psi = I$, where I is an identity matrix of size n_e . We now seek a compressed basis for the span of Ψ , denoted by the set of vectors $\Phi = [\phi_1, \dots, \phi_{n_e}]$, where each ϕ_i is a sparse vector with spatially localized support after truncating entries with small magnitudes. In such case, ϕ_i is called a localized vector.

2.2. Selected columns of the density matrix. As opposed to widely used procedures such as maximally localized Wannier functions [29], the key difference in the SCDM procedure is that the localized orbitals ϕ_i are obtained directly from columns of the density matrix $P = \Psi \Psi^*$. The aforementioned nearsightedness principle states that for insulating systems, each column of the matrix P is localized. As a result, selecting any linearly independent subset of n_e of them will yield a localized basis for the span of Ψ . However, picking n_e random columns of P may result in a poorly conditioned basis if, for example, there is too much overlap between the selected columns. Therefore, we would like a means for choosing a well-conditioned set of columns, denoted $\mathcal{C} = \{c_1, c_2, \dots, c_{n_e}\}$, to use as the localized basis. Intuitively we expect such a basis to select columns to minimize overlaps with each other when possible.

This is algorithmically accomplished with a QRCP factorization (see, e.g., [17]). More specifically, given a matrix A , a QRCP seeks to compute a permutation matrix Π such that the leading submatrices $(A\Pi)_{1,\dots,k,:}$ for any applicable k are as well conditioned as possible. In particular, if we let \mathcal{C} denote the columns selected by the first n_e columns of Π then $A_{:, \mathcal{C}}$ should be a well-conditioned set of n_e columns of A .

In our setting, this means we would ideally compute an QRCP factorization of the matrix P to identify n_e well-conditioned columns from which we may construct a localized basis. However, this would be highly costly since P is a large matrix of size N . Fortunately, we may equivalently compute the set \mathcal{C} by computing a QRCP factorization of Ψ^* or, in fact, any matrix U with orthogonal columns such that $P = UU^*$. More specifically, we compute

$$(2.4) \quad \Psi^* \Pi = Q \begin{bmatrix} R_1 & R_2 \end{bmatrix},$$

and the first n_e columns of Π encode \mathcal{C} .

The SCDM procedure to construct an orthonormal set of localized basis elements, denoted ϕ_i for $i = 1, \dots, n_e$, and collected as columns of the matrix Φ , is presented in its simplest form in Algorithm 1. In such form the algorithm requires knowledge of the orthogonal factor from the QRCP. However, an alternative description simply requires the column selection \mathcal{C} . We may equivalently write the SCDM algorithm as

Algorithm 1. The SCDM algorithm.

Given: the Kohn–Sham orbitals Ψ

1: Compute a column pivoted QR of Ψ^* , $\Psi^*\Pi = Q [R_1 \ R_2]$

2: Compute $\Phi = \Psi Q$ or, alternatively, $\Phi^* = [R_1 \ R_2] \Pi^*$

Output: a localized basis for the Kohn–Sham subspace Φ

Algorithm 2. An alternative version of the SCDM algorithm.

Given: the Kohn–Sham orbitals Ψ

1: Compute \mathcal{C} associated with a column pivoted QR of Ψ^*

2: Compute the QR factorization $(\Psi_{\mathcal{C},:})^* = QR$

3: Compute $\Phi = \Psi Q$

Output: a localized basis for the Kohn–Sham subspace Φ

in Algorithm 2. Note that in Algorithm 2, the cost of the QR factorization for the matrix $(\Psi_{\mathcal{C},:})^*$ is only $\mathcal{O}(n_e^3)$.

Remark 2.1. There are various equivalent ways to construct the SCDM algorithm. While the simple presentation here differs slightly from the original presentation [10] the two are mathematically equivalent up to a choice of sign for the columns of Φ . The original presentation corresponds to the computation of a QR factorization of $(\Psi_{\mathcal{C},:})^*$ via a Cholesky factorization of $(\Psi_{\mathcal{C},:})(\Psi_{\mathcal{C},:})^*$. QR factorizations are not unique; there is always ambiguity up to a diagonal matrix with entries on the unit circle. However, such an ambiguity does not have any affect on the localization.

This second interpretation allows us to briefly explain why this algorithm constructs localized orbitals. Let D be a diagonal matrix with ± 1 on the diagonal such that DR has positive diagonal entries. This means that

$$Q = (\Psi_{\mathcal{C},:})^* R^{-1}D,$$

where R^{-1} is a Cholesky factor of $[(\Psi_{\mathcal{C},:})(\Psi_{\mathcal{C},:})^*]^{-1}$. Importantly, $(\Psi_{\mathcal{C},:})(\Psi_{\mathcal{C},:})^* = P_{\mathcal{C},\mathcal{C}}$ and, therefore, exhibits off diagonal decay so long as $P_{\mathcal{C},\mathcal{C}}$ is well conditioned. This property is then inherited by R^{-1} [4]. Finally, since $P_{\mathcal{C},:} = \Psi(\Psi_{\mathcal{C},:})^*$

$$\Phi = P_{\mathcal{C},:}R^{-1}D,$$

we may conclude that it is well localized— $P_{\mathcal{C},:}$ is well localized and $R^{-1}D$ does not destroy that locality. Importantly, here we see that all the factorizations we are performing can be thought of as involving submatrices of P .

The overall computational cost of the algorithm is $\mathcal{O}(Nn_e^2)$, and practically the cost is dominated by the single QRCP factorization regardless of the version used. Another key feature of the algorithm, especially for our modifications later, is that because we are effectively working with the spectral projector P , the method performs equivalently if a different orthonormal basis for the range of Ψ is used as input. In physics terminology, the SCDM procedure is gauge-invariant. Last, the key factor in forming a localized basis is the selection of a well-conditioned subset of columns. Small changes to the selected columns, provided they remain nearly as well conditioned, may not significantly impact the overall quality of the basis.

3. The approximate column selection algorithm. When the SCDM procedure is used as a postprocessing tool for a single atomic configuration, the computational cost is usually affordable. In fact in such a situation, the most time-consuming

part of the computation is often the I/O related to the Ψ matrices especially for systems of large sizes. However, when localized orbitals need to be calculated repeatedly inside an electronic structure software package, such as in the context of hybrid functional calculations with geometry optimization or *ab initio* molecular dynamics simulations, the computational cost of SCDM can become relatively large. Here we present an algorithm that significantly accelerates the SCDM procedure.

The core aspect of the SCDM procedure is the column selection procedure. Given a set of appropriate columns the requisite orthogonal transform to construct the SCDM may be computed from the corresponding rows of Ψ , as seen in Algorithm 2. Here we present a two-stage procedure for accelerating this selection of columns and hence the computation of Φ . First, we construct a set of approximately localized orbitals that span the range of Ψ via a randomized method that requires only Ψ and the electron density ρ , though if ρ is not given it may be computed directly from Ψ without increasing the asymptotic computational complexity. We then use this approximately localized basis as the input for a procedure that refines the selection of columns from which the localized basis is ultimately constructed. This is done by using the approximate locality to carefully partition the column selection process into a number of small, local, QRCP factorizations. Each small QRCP may be done in parallel and operates on matrices of much smaller dimension than Ψ .

3.1. Approximate localization. The original SCDM procedure, through the QRCP, examines all N columns of Ψ^* to decide which columns to use to construct Q . However, physical intuition suggests that it is often not necessary to visit all columns to find good pivots. For instance, for a molecular system in vacuum, it is highly unlikely that a pivot comes from a column of the density matrix corresponding to the vacuum space away from the molecule. This inspires us to accelerate the column selection procedure by restricting the candidate columns.

This is accomplished by generating $\mathcal{O}(n_e \log n_e)$ independent and identically distributed random sample columns, using the normalized electron density as the probability distribution function. Indeed, if a column of the density matrix corresponds to the vacuum, then the electron density is very small and hence the probability of picking the column is very low. In statistics this corresponds to leverage score sampling; see, e.g., [26]. This randomized version of the SCDM algorithm is outlined in Algorithm 3.

Algorithm 3. Computing an approximately localized collection of basis vectors.

Given: Kohn–Sham orbitals Ψ , electron density ρ , concentration γ , and failure probability δ

- 1: Sample $(n_e/\gamma) \log n_e/\delta$ elements from $\{1, \dots, N\}$ based on the discrete distribution

$$\Pr(\{j\}) = \rho(j)/n_e$$

and denote this set $\tilde{\mathcal{C}}$

- 2: Compute the column pivoted QR factorization

$$\left(\Psi_{\tilde{\mathcal{C}},:}\right)^* \Pi = QR$$

- 3: Form approximately localized basis $\tilde{\Phi} = \Psi Q$
-

To complete our discussion of Algorithm 3 we must justify the subsampling procedure used to select $\tilde{\mathcal{C}}$. In order to do so we introduce a simple model for the column selection procedure based on the idea that columns “similar” to the ones selected by Algorithm 1 will work well to compute an approximately localized basis. Our underlying assumption is that a set of columns will serve to approximately localize the basis if it contains at least one column in each region where one of the ϕ_i is large. Because the ϕ_i constructed via the SCDM procedure decay exponentially, this is analogous to saying that any column associated with a grid point close enough to the “true” grid point used will suffice. However, by avoiding explicit use of spatial relations of grid points, our algorithm and its parameters are not dependent on the physical geometry.

To codify this postulate, we let $\mathcal{I}_i \subset \{1, \dots, N\}$ be the smallest nonempty set such that

$$(3.1) \quad \sum_{j \in \mathcal{I}_i} |\phi_i(j)|^2 \geq \gamma,$$

where $\gamma \in (0, 1)$. If multiple such sets exist we select the one that maximizes $\sum_{j \in \mathcal{I}_i} |\phi_i(j)|^2$. Now, we may write our assumption more concretely: a column $c_i \in \{1, \dots, N\}$ suffices to approximately construct ϕ_i if it is contained in \mathcal{I}_i . Taking γ to be sufficiently small would enforce adequate sampling to ensure the columns selected by Algorithm 2 are selected to be part of $\tilde{\mathcal{C}}$. However, in practice this is not necessary for the construction of a localized basis and by choosing a larger γ we allow for other columns near the peak of ϕ_i to act as good surrogates.

Under this assumption, to approximately localize the basis, we must simply ensure that $\tilde{\mathcal{C}}$ contains at least one distinct column in each of the sets $\mathcal{I}_1, \mathcal{I}_2, \dots, \mathcal{I}_{n_e}$, i.e., we need a one-to-one matching between sets and columns. Theorem 3.1 provides an upper bound on the required cardinality of $\tilde{\mathcal{C}}$ to ensure it may be used to approximately localize the basis with high probability. We do require an additional mild assumption that ensures the sets \mathcal{I}_i do not simultaneously overlap significantly and have small support.

THEOREM 3.1. *Let η be the largest constant such that there exist disjoint subsets $\mathcal{I}_i^s \subseteq \mathcal{I}_i$, $i = 1, \dots, n_e$, each satisfying*

$$\sum_{j \in \mathcal{I}_i^s} |\phi_i(j)|^2 \geq \eta \sum_{j \in \mathcal{I}_i} |\phi_i(j)|^2.$$

The set $\tilde{\mathcal{C}}$ constructed by sampling

$$m \geq (n_e/\eta\gamma) \log n_e/\delta$$

elements with replacement from $\{1, \dots, N\}$ based on the discrete distribution

$$\Pr(\{j\}) = \rho(j)/n_e$$

contains a unique element in each \mathcal{I}_i for $i = 1, \dots, n_e$ with probability $1 - \delta$.

Proof. Let \mathcal{F} be the event that $\tilde{\mathcal{C}}$ does not contain a distinct element in one of the sets \mathcal{I}_i . We may write

$$\Pr(\mathcal{F}) \leq \Pr\left(\left\{\tilde{\mathcal{C}} \cap \mathcal{I}_i^s = \emptyset \text{ for some } i\right\}\right)$$

because requiring $\tilde{\mathcal{C}}$ to contain an element in each \mathcal{I}_i^s implies that $\tilde{\mathcal{C}}$ contains a distinct element in each \mathcal{I}_i . Subsequently, by a union bound

$$\Pr(\mathcal{F}) \leq \sum_{i=1}^{n_e} \Pr\left(\{\tilde{\mathcal{C}} \cap \mathcal{I}_i^s = \emptyset\}\right).$$

The event $\{\tilde{\mathcal{C}} \cap \mathcal{I}_i^s = \emptyset\}$ is simply the probability that none of the m samples falls in \mathcal{I}_i^s . Because

$$\rho(j) = \sum_{i=1}^{n_e} |\phi_i(j)|^2,$$

we may find the lower bound for the probability of selecting an element j in \mathcal{I}_i as

$$\Pr(\{j \in \mathcal{I}_i^s\}) \geq \eta\gamma/n_e.$$

Consequently,

$$\sum_{i=1}^{n_e} \Pr\left(\{\tilde{\mathcal{C}} \cap \mathcal{I}_i^s = \emptyset\}\right) \leq n_e(1 - \eta\gamma/n_e)^m$$

and to ensure the probability of missing any of the sets \mathcal{I}_i^s is less than δ we may simply enforce

$$m \geq \frac{n_e}{\eta\gamma} \log n_e/\delta. \quad \square$$

Remark 3.2. The parameter η is necessary in the proof, but algorithmically we simply assume it to be one; alternatively, one could consider choosing $1/(\gamma\eta)$ rather than just $1/\gamma$ as the oversampling factor.

If, for example, we take $\gamma = 1/2$ and $\eta = 1/2$ this bound says $4n_e \log n_e/\delta$ samples suffices for the approximate localization procedure. As expected, when either the failure probability δ or the cardinality of \mathcal{I}_i go to one the required number of samples grows. Furthermore, since $\gamma = \min_i \max_j |\phi_i(j)|^2$ corresponds to each of the sets \mathcal{I}_i containing a single point, that is a lower bound on how small γ can be theoretically. We remark that this theoretical bound may be pessimistic for two reasons. One is its use of the union bound. The other is the introduction of η . The disjoint requirement simplifies the assignment of selected columns to sets for the proof but is a relaxation of what is really needed.

3.2. Accelerating the SCDM procedure using an approximately localized basis. Once the approximate localized orbitals $\tilde{\Phi}$ are obtained, we would like to perform a refinement procedure to further localize the basis. We do this by taking advantage of the locality of the input to the SCDM procedure. In Algorithm 3 we approximate the behavior of the SCDM algorithm by restricting the number of columns of Ψ^* that the QRCP factorization can select from. However, once we have a somewhat localized basis we can efficiently take more columns of Ψ^* into consideration. This allows us to better approximate the original SCDM algorithm and, therefore, construct a more localized basis. We accomplish this with a procedure that resembles the tournament pivoting strategy for computing a QRCP. We first compute a bunch of small QRCP factorizations, each involving columns associated with the support of a subset of the rows of $\tilde{\Phi}^*$. This is computationally feasible because the rows are already somewhat localized. Last, because this procedure generates more candidate

columns than needed we perform one final QRCP to try and get as well-conditioned a set of columns of Ψ^* as possible.

Algorithmically, we first need to select a superset of the ultimately desired columns from which to select the final columns used in the localization. To construct such a superset we consider each approximately localized orbital and how we may refine it. Each orbital, once approximately localized, only exerts influence on, i.e., significantly overlaps with, nearby localized orbitals. Hence, we may refine the selected columns locally. This means that for each orbital we may simply figure out which orbitals it substantially overlaps with and compute a QRCP on just those orbitals (rows) of $\tilde{\Phi}^*$ while simultaneously omitting columns with small norm over those rows. This process will yield a small number of selected columns that we add to a list of potential candidate columns for the final localization. However, because we repeat this process for each localized orbital we might have more than n_e total candidate columns by a small multiplicative factor. Therefore, we perform one final column pivoted QR factorization on these candidate columns to select the final set \mathcal{C} .

In principle, while the localized orbitals get small on large portions of the domain, they are not actually zero. Hence, for any given ϵ we say two orbitals substantially overlap if there is any spatial point where they both have relative magnitude greater than ϵ . We find in practice that choosing ϵ close to the relative value at which one considers an orbital to have vanished suffices, though taking $\epsilon \rightarrow 0$ one recovers the original SCDM algorithm. Importantly, this parameter is completely independent of the geometry of the problem under consideration and thus does not require detailed knowledge of the system to set.

We detail our complete algorithm for computing the orthogonalized SCDM from approximately localized input in Algorithm 4. The only parameter is how small a column is for it to be ignored. At a high level, the goal is simply to generate additional candidate columns by allowing each orbital to interact with its immediate neighbors. Provided that each orbital is only close to a small number of others, and that the approximate localization is sufficient, this procedure may be performed quickly.

Algorithm 4. Refining an approximately localized collection of basis vectors.

- Given: approximately localized Kohn–Sham orbitals $\tilde{\Phi}$ and column tolerance ϵ
- 1: $\mathcal{J}_i = \left\{ j \in \{1, \dots, N\} \mid |\tilde{\phi}_i| > \epsilon \max_k |\tilde{\phi}_i(\mathbf{r}_k)| \right\}$ for $i = 1, \dots, n_e$
 - 2: **for** $i = 1, \dots, n_e$ **do**
 - 3: Set $\mathcal{R}_i = \{j \in \{1, \dots, n_e\} \mid \mathcal{J}_i \cap \mathcal{J}_j \neq \emptyset\}$
 - 4: Set $\mathcal{L}_i = \bigcup_{j \in \mathcal{R}_i} \mathcal{J}_j$
 - 5: Compute a column pivoted QR factorization of $\left(\tilde{\Phi}^* \right)_{\mathcal{R}_i, \mathcal{L}_i}$ and denote the pivot columns \mathcal{C}_i
 - 6: **end for**
 - 7: Set $\tilde{\mathcal{C}} = \cup_i \mathcal{C}_i$
 - 8: Compute the column pivoted QR factorization

$$\left(\tilde{\Phi}_{\tilde{\mathcal{C}}:,} \right)^* \Pi = QR$$

- 9: Form the localized basis $\Phi = \tilde{\Phi}Q$
-

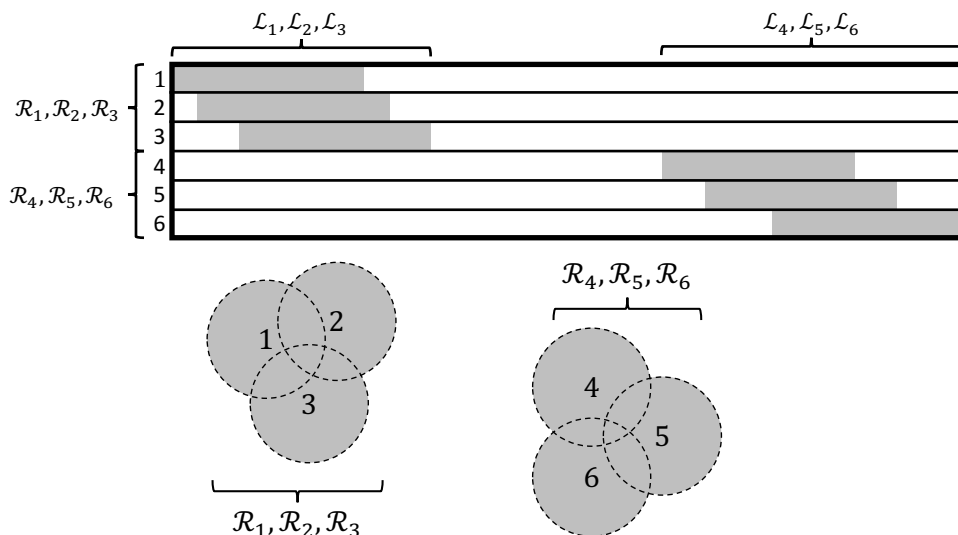


FIG. 1. Matrix (top) and physical domain (bottom) associated with two collections of approximately localized orbitals whose support (lightly shaded region) is disjoint after truncation.




FIG. 2. Matrix (top) and physical domain (bottom) associated with a chain of approximately localized orbitals whose support (lightly shaded region) is connected after truncation.

To illustrate the behavior of this algorithm, we sketch the behavior of Algorithm 4 in two cases. In one case, after the approximate localization, we have two sets of orbitals whose support sets after truncation are disjoint. This is shown in Figure 1. Here, simply computing two independent QRCP factorizations is actually equivalent to computing the QRCP of the entire matrix. As we see Algorithm 4 partitions the orbitals into two sets and then only considers the columns with significant norm over the orbital set. In the second case, illustrated in Figure 2, we have a chain of orbitals

whose support set after truncation forms a connected region in the spatial domain. In this situation we do not actually replicate the computation of a QRCP of the whole matrix, but rather for each orbital we compute a local QRCP ignoring interactions with distant orbitals. More specifically, any column mostly supported on a given orbital will be minimally affected by orthogonalization against columns associated with distant orbitals. Therefore, we may simply ignore those orthogonalization steps and still closely match the column selection procedure.

Remark 3.3. In Algorithm 4 it may sometimes be the case that several of the \mathcal{R}_i are identical (e.g., as in Figure 1). In this case, as a matter of practical performance optimization, one may simply skip instances of the loop over i that would be computing a QRCP of the exact same matrix as a prior instance of the loop. Similarly, if there is a set $\mathcal{S} \subset \{1, \dots, n_e\}$ such that $(\cup_{i \in \mathcal{S}} \mathcal{R}_i) \cap (\cup_{i \in \mathcal{S}^c} \mathcal{R}_i) = \emptyset$ and $|\cup_{i \in \mathcal{S}} \mathcal{R}_i|$ is small, we may combine the instances of the loop corresponding to $i \in \mathcal{S}$ into a single small QRCP by simply using $\cup_{i \in \mathcal{S}} \mathcal{R}_i$ as the set of rows. This avoids redundant work if a small collection of orbitals just interact among themselves and are disjoint from all others.

3.3. Computational complexity. The computational cost of Algorithms 3 and 4 is formally limited by the cost of computing a single GEMM ΨQ at cost $\mathcal{O}(n_e^2 N)$. However, this single BLAS 3 operation must appear in any localization procedure starting with the Ψ matrix as an input. This GEMM operation is also highly optimized in both sequential and parallel computational environments. Furthermore, at large scale one can approximate the product with $\mathcal{O}(n_e N)$ complexity using sparse linear algebra, since the support of the result is sparse and one may infer the support based on the column indices from which Q is built. For these reasons, we let T_{mult} represent the cost of computing ΨQ .

Using this notation, the computational cost of Algorithm 3 is

$$\mathcal{O}(n_e N \log n_e + n_e^3 \log n_e) + T_{mult}.$$

More specifically, the random selection of columns costs $\mathcal{O}(N n_e \log n_e)$ and the subsequent QR factorization costs $\mathcal{O}(n_e^3 \log n_e)$. If we assume that the support of the approximately localized basis used as input for Algorithm 4 and the number of nearby orbital centers are bounded by some constant independent of N and n_e , which is reasonable for models where the molecule is growing and the discretization quality remains constant, the computational cost is

$$\mathcal{O}(n_e N + n_e^3) + T_{mult}.$$

Under these assumptions each of the, at most, n_e small QR factorizations has constant cost. However, the procedure for finding \mathcal{J}_i introduces a dependency on N . While the computational cost of the refinement procedure and randomized algorithms is broadly similar, the randomized algorithm is significantly faster in practice because drawing the random samples is cheaper than the support computations needed in the refinement algorithm.

Last, our algorithms are memory efficient. In fact, for practical purposes their memory cost is bounded by the cost of storing Ψ plus a few work arrays of size N . Besides the storage for Ψ , all of the matrices we operate on cost at most $\mathcal{O}(n_e^3 \log n_e)$ to store and may be discarded after use. Furthermore, the QR factorizations and matrix multiplication may be done in place excepting a few work arrays of length N .

4. Numerical examples. To demonstrate the performance of our method we use three examples that capture the different facets of our algorithm. The first example is the dissociation of a BH_3NH_3 molecule, and the second example is the alkane chain. We select these two examples not because they are computationally expensive but because they clearly demonstrate that our approximate localization algorithm is effective in two very different regimes. In particular, the effectiveness of the algorithm is independent of whether localized orbitals form one single group or multiple disconnected groups. Our third example is a large supercell with 256 water molecules. We demonstrate the performance gains over the existing SCDM method and provide a comparison with Wannier90 [31].

In all of the examples here we assume we have access to the electron density ρ from the electronic structure calculation and therefore exclude its computation from the timings of our randomized method. We use $3n_e \log n_e$ samples in the randomized algorithm, corresponding to $\gamma = 1/3$, for all of the experiments. Furthermore, to more clearly illustrate the advantages of our method, we separately report timings for computing the orthogonal transform Q that localizes the orbitals and subsequent computation of the localized orbitals by a single matrix product. Here we only consider the orthogonalized SCDM, as discussed in this paper. For the refinement algorithm we set the relative truncation threshold at $\epsilon = 5 \times 10^{-2}$ and we observe that this is sufficient for our new algorithm to closely match the results of the existing algorithm. Finally, in all of the experiments here we use 2.5×10^{-2} as the relative truncation threshold of the localized orbitals when counting the fraction of entries that are nonzero. This measure of “locality” has the advantage of not depending on the geometry of the physical system. However, for completeness we also provide spread computations in the final example. Prior work validates the expected exponential decay of the orbitals by varying the truncation threshold [10].

All numerical results shown were run on a quad-socket Intel Xeon E5-4640 processor clocked at 2.4 GHz with 1.5 TB of RAM and our algorithms were implemented in MATLAB R2015a. Our code is sequential and the only multithreading is in the LAPACK and BLAS routines called by MATLAB. The storage cost of all the algorithms presented here is $\mathcal{O}(n_e N)$, which is also the cost to store Ψ in memory. In our largest example a copy of Ψ cost roughly 60 GB to store. The Kohn–Sham orbitals were computed using Quantum ESPRESSO [16], and VMD [22] was used for plotting orbitals and molecules in the alkane and water examples.

4.1. BH_3NH_3 . First, we demonstrate the performance of the approximate column selection method for the dissociation process of a BH_3NH_3 molecule. The main purpose of this example is to demonstrate that our approximate localization algorithm is equally effective when localized orbitals form disconnected groups or a single group. Figure 3 shows the localized orbitals for three atomic configurations, where the distance between B and N atoms is 1.18, 3.09, and 4.96 Bohr, respectively, and also shows the locality for orbitals computed by Algorithms 1, 3, and 4 for each of the three atomic configurations. Here we plot an isosurface of the orbitals at a value of 2.5×10^{-2} .

We find that Algorithm 4 automatically identifies that the localized orbitals should be treated as two disconnected groups for the dissociated configuration, and as one single group for the bonded configuration. We see that in all cases, the randomized method works quite well on its own. Furthermore, after applying Algorithm 4 to the output of the randomized method, the sparsity of the orbitals is nearly indistinguishable from that of the original SCDM algorithm. In these three scenarios

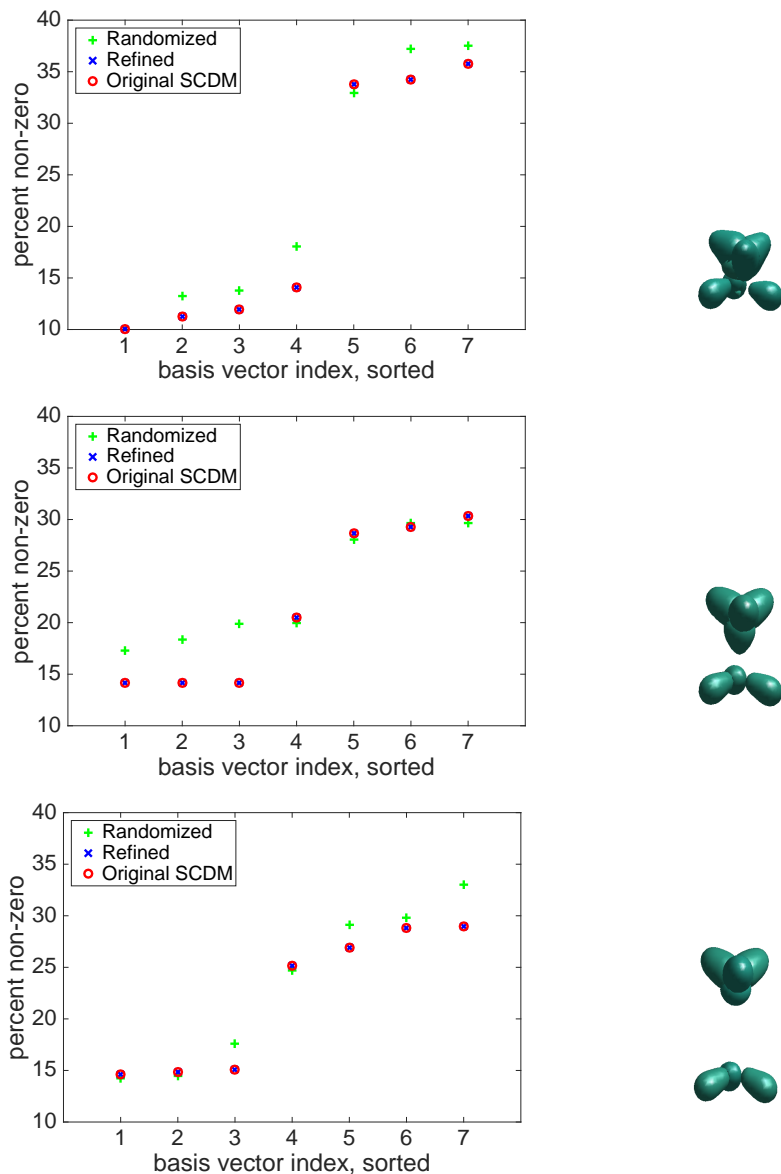


FIG. 3. Sparsity (left) of localized orbitals computed by Algorithms 1, 3, and 4 based on fraction of nonzero entries after truncation and orbital isosurfaces (right) at 2.5×10^{-2} generated by Algorithm 4 when using the output of Algorithm 3 as input. Three different configurations moving from the bonded configuration (top) to the dissociated configuration (bottom).

the condition number of $P_{:,c}$, equivalently $(\Psi_{C,:})^*$, is never larger than three. Here we let $\mathcal{C} \subset \tilde{\mathcal{C}}$ denote the final set of columns we have selected via the last QRCP in Algorithm 4 as it corresponds to the columns of the density matrix from which we ultimately build the localized basis.

4.2. Alkane chain. Our second example is the alkane chain (atomic configuration shown in Figure 4). Similar to the BH_3NH_3 , this example is not computationally




FIG. 4. Atomic configuration of the alkane chain.




FIG. 5. Histogram of localized orbitals for the alkane chain example computed by three different algorithms based on fraction of nonzero entries after truncation: (top) output of the randomized algorithm, (middle) output of the refinement algorithm applied to the output of the randomized algorithm, and (bottom) output of the original SCDM algorithm.

expensive but confirms that the approximate localization algorithm is effective even when all localized orbitals form one large connected group. We demonstrate that the refinement process still achieves the desired goal. In this example $N = 820, 125$ and $n_e = 100$.

Figure 5 shows histograms of the fraction of nonzero entries after truncation for the randomized method, the refinement procedure applied to the output of the randomized method, and our original algorithm. We observe that the randomized method actually serves to localize the orbitals rather well. However, the output is




FIG. 6. The same (as determined by picking those maximally correlated) orbital as generated by the randomized algorithm (left), the refinement procedure (middle), and the original SCDM algorithm (right). Here an isosurface value of 5×10^{-3} was used and the colors delineate positive and negative regions.

TABLE 1

Runtime for localization algorithms as applied to an alkane chain.

Operation	Time (s)
Matrix-matrix multiplication ΨQ	0.2713
Randomized version, Algorithm 3	0.0261
Refinement step, Algorithm 4	0.7176
Total cost of our two-stage algorithm	1.0149
Original algorithm, Algorithm 1	6.9412

clearly not as good as that produced by the original SCDM algorithm. However, once the refinement algorithm is applied we see that, while not identical, the locality of the localized orbitals basically matches that of the localized orbitals generated by Algorithm 1. This is further illustrated in Figure 6, which shows isosurfaces for a localized orbital generated by each of the three methods. Once again the columns of the density matrix we ultimately select are very well conditioned; in fact the condition number is less than two.

Table 1 illustrates the computational cost of our new algorithms as compared to the original version. The refinement algorithm is over nine times faster than the original algorithm, though it cannot be performed without first approximately localizing the orbitals. Hence, we also provide the time of getting the orthogonal transform out of Algorithm 4, which is the sum of the preceding three lines in the table. Even taking the whole pipeline into account we see a speedup of about a factor of close to seven.

4.3. Water molecules. We now consider a three-dimensional system consisting of 256 water molecules (part of the atomic configuration shown in Figure 7). In this example, $N = 7,381,125$ and $n_e = 1024$.

Figure 8 shows histograms of the fraction of nonzero entries after truncation for the randomized method, the refinement procedure applied to the output of the randomized method, and the original SCDM algorithm when applied to the 256 water molecule system. As before, the randomized method actually serves to localize the orbitals rather well. However, there is still visible difference between the result of the randomized algorithm and the original SCDM algorithm. Application of the refinement algorithm achieves a set of localized orbitals that broadly match the quality of the ones computed by the original SCDM algorithm. Similar to the previous two examples $P_{:,c}$ is very well conditioned; its condition number is once again less than two.

Table 2 illustrates the computational cost of our new algorithms as compared to Algorithm 1. As before, the randomized algorithm for computing is very fast and in this case much faster than the matrix-matrix multiplication required to construct the localized orbitals themselves. This makes the algorithm particularly attractive in practice when ρ is given and, as in many electronic structure codes, the application




FIG. 7. Part of the atomic configuration of 256 water molecules.

of Q to Ψ may be effectively parallelized. Furthermore, the complete procedure for getting orthogonal transform out of Algorithm 4, the sum of the preceding three lines in the table, is more than 30 times faster than Algorithm 1.

Finally, we use this example to provide a comparison with the popular Wannier90 [31] software package and further demonstrate the quality of orbitals computed by our algorithms. While we have previously been looking at sparsity after truncation to evaluate the quality of the orbitals, we now also evaluate them based on the spread criteria Wannier90 tries to minimize. Loosely speaking this corresponds to the sum of the variance of each orbital [30]. All of the spreads here were computed by Wannier90 to ensure the same quantity was being measured in each case. Importantly, this is the quantity Wannier90 seeks to minimize and we therefore do not expect to do better under this metric. For example, given our localized orbitals as input, Wannier90 should always be able to at least slightly decrease the objective function value. For this comparison, we used the random initial guess option in Wannier90 and the default convergence tolerance of 10^{-10} . Notably, the random choice in Wannier90 does not correspond to a random gauge; rather an initial guess is constructed by randomly placing Gaussian functions in space.

Table 3 shows the spreads and time to solution for all of the algorithms. First, we observe that the spread using our two-stage procedure almost exactly matches that of the output of Algorithm 1 while being over an order of magnitude faster. Second, the output from our algorithms, even the randomized one, is close to the local minimum found by Wannier90. Figure 9 directly compares the spreads of our localized orbitals from the two-stage algorithm and those at a local minimum found by Wannier90. As expected, there are discrepancies, but we are generally within approximately 0.05 \AA^2 per orbital. This aligns with our expectations based on the objective function difference. For comparison, we computed a localized basis for an isolated H_2O molecule. Using Algorithm 1 yielded spreads of 0.46, 0.48, 0.56, and 0.57 \AA^2 . Subsequently running Wannier90 to convergence yielded spreads of 0.47, 0.48, 0.55, and 0.55 \AA^2 , which matched the results when starting with a Wannier90 generated random initial guess.

Admittedly, the time to solution for Wannier90 may depend strongly on the initial guess, e.g., in this experiment the initial spread for Wannier90 was 5339 \AA^2

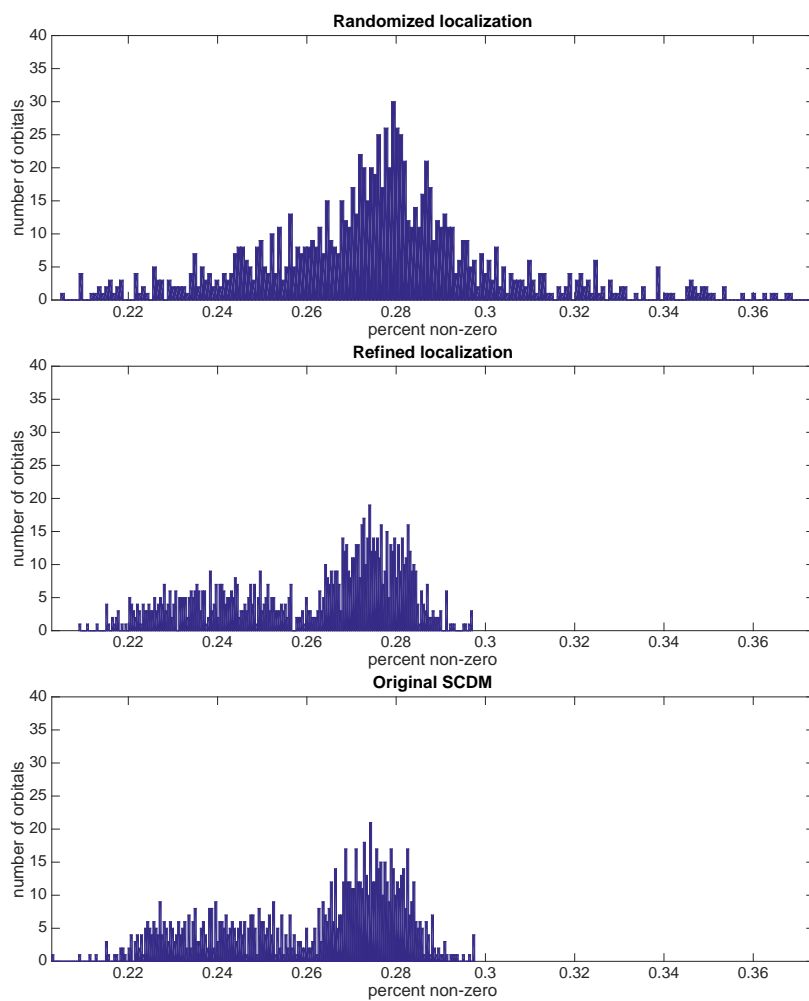


FIG. 8. Histogram of localized orbitals for 256 water molecules computed by three different algorithms based on fraction of nonzero entries after truncation: (top) output of the randomized algorithm, (middle) output of the refinement algorithm applied to the output of the randomized algorithm, and (bottom) output of the original SCDM algorithm.

TABLE 2

Runtime for localization algorithms as applied for 256 water molecules.

Operation	Time (s)
Matrix-matrix multiplication ΨQ	47.831
Randomized version, Algorithm 3	14.024
Refinement step, Algorithm 4	78.361
Total cost of our two-stage algorithm	140.22
Original algorithm, Algorithm 1	4496.1

and convergence took 20 iterations. However, a poor initial guess could result in Wannier90 failing to converge, converging to a worse local minimum, or taking longer to converge. Our algorithms are direct and have no such dependence on an initial guess. In this experiment, each iteration of Wannier90 took roughly 85 seconds. So

TABLE 3
Final spread of orbitals for 256 water molecules and time to solution.

Algorithm	Spread (\AA^2)	Time to solution (s)
Algorithm 1	589.91	4496.1
Randomized version, Algorithm 3	636.60	14.024
Two-stage procedure using Algorithms 3 and 4	589.97	140.22
Wannier90	550.20	1715.2

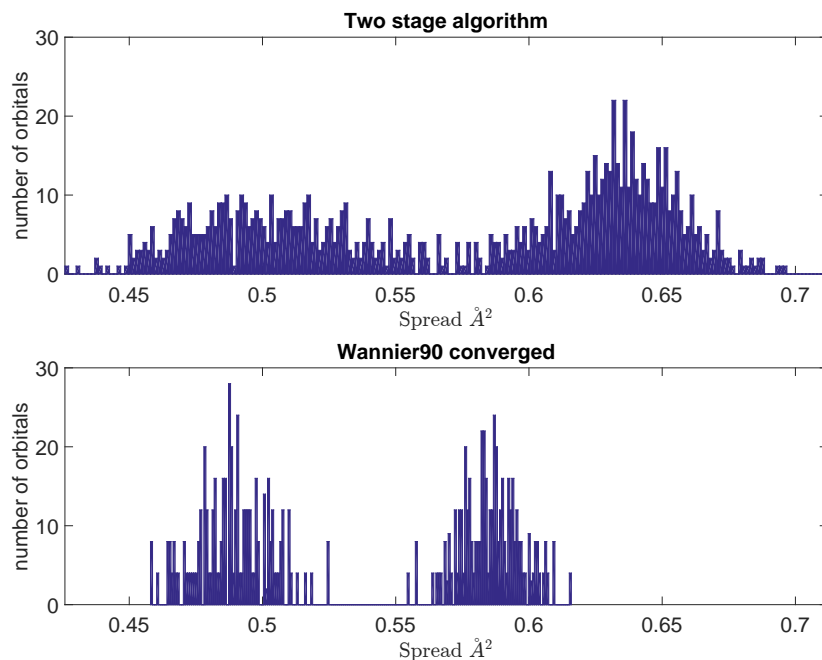


FIG. 9. Histogram of localized orbitals for 256 water molecules computed by two different algorithms based on the Wannier90 spread functional: (top) output of the refinement algorithm applied to the output of the randomized algorithm, and (bottom) converged local minimum for Wannier90.

even two iterations cost more than our two-stage algorithm and we have omitted the cost of computing the overlap matrices based on Ψ , which Wannier90 requires as input and we computed using Quantum ESPRESSO on 256 processors. Collectively, these results make our algorithm an attractive alternative to Wannier90. If a local minimum of the objective is desired, the output from Algorithm 3 or 4 may be used as an algorithmically constructed initial guess for Wannier90.

Remark 4.1. Seeding Wannier90 with our two-stage algorithm convergence took 16 iterations and we appear to arrive at an equivalent local minima. For this system, it would appear the area around the local minimum is quite flat and the default absolute convergence criteria is rather stringent. In fact, after one iteration starting with our algorithm Wannier90 yields localized functions whose spread is within 0.4% of the converged value. To get to within 1% starting from the random initial guess takes five iterations.

5. Conclusion. We have presented a two-stage algorithm to accelerate the computation of the SCDM for finding localized representation of the Kohn–Sham invariant

subspace. We first utilize an algorithm based on random sampling to approximately localize the basis and then perform a subsequent refinement step. This method can achieve computational gains of over an order of magnitude for systems of relatively large sizes. Furthermore, the orbitals computed are qualitatively and quantitatively similar to those generated by the original SCDM algorithm. Last, for large systems we observe that our algorithm may provide an attractive alternative to Wannier90 and, at the very least, may provide a simple method for the construction of a good initial guess.

Rapid computation of a localized basis allows for its use within various computations in electronic structure calculation where a localized basis may need to be computed repeatedly. This includes computations such as molecular dynamics and time-dependent DFT with hybrid exchange-correlation functions. Finally, the ideas inherent to the SCDM procedure have potential applicability in problems outside of Kohn–Sham DFT because the structural and behavioral properties we exploit are not necessarily unique to the problem. Our new algorithms would admit faster computation in such contexts as well.

Acknowledgments. The authors thank Stanford University and the Stanford Research Computing Center for providing computational resources and support that have contributed to these research results. The authors also thank the anonymous referees for their many helpful suggestions.

REFERENCES

- [1] E. ANDERSON, Z. BAI, C. BISCHOF, S. BLACKFORD, J. DEMMEL, J. DONGARRA, J. DU CROZ, A. GREENBAUM, S. HAMMARLING, A. MCKENNEY, AND D. SORENSEN, *LAPACK Users' Guide*, 3rd ed., SIAM, Philadelphia, 1999.
- [2] F. AQUILANTE, T. B. PEDERSEN, A. S. DE MERÁS, AND H. KOCH, *Fast noniterative orbital localization for large molecules*, *J. Chem. Phys.*, 125 (2006), 174101.
- [3] A. D. BECKE, *Density functional thermochemistry. III. The role of exact exchange*, *J. Chem. Phys.*, 98 (1993), pp. 5648–5652.
- [4] M. BENZI, P. BOITO, AND N. RAZOUK, *Decay properties of spectral projectors with applications to electronic structure*, *SIAM Rev.*, 55 (2013), pp. 3–64.
- [5] L. S. BLACKFORD, J. CHOI, A. CLEARY, E. D'AZEVEDO, J. DEMMEL, I. DHILLON, J. DONGARRA, S. HAMMARLING, G. HENRY, A. PETITET, K. STANLEY, D. WALKER, AND R. C. WHALEY, *ScaLAPACK Users' Guide*, SIAM, Philadelphia, 1997.
- [6] E. I. BLOUNT, *Formalism of Band Theory*, *Solid State Phys.* 13, Academic Press, New York, 1962, pp. 305–373.
- [7] C. BROUDER, G. PANATI, M. CALANDRA, C. MOURougANE, AND N. MARZARI, *Exponential localization of Wannier functions in insulators*, *Phys. Rev. Lett.*, 98 (2007), 046402.
- [8] J. CLOIZEAUX, *Analytical properties of n-dimensional energy bands and Wannier functions*, *Phys. Rev.*, 135 (1964), pp. A698–A707.
- [9] J. CLOIZEAUX, *Energy bands and projection operators in a crystal: Analytic and asymptotic properties*, *Phys. Rev.*, 135 (1964), pp. A685–A697.
- [10] A. DAMLE, L. LIN, AND L. YING, *Compressed representation of Kohn–Sham orbitals via selected columns of the density matrix*, *J. Chem. Theory Comput.*, 11 (2015), pp. 1463–1469.
- [11] A. DAMLE, L. LIN, AND L. YING, *SCDM-K: Localized orbitals for solids via selected columns of the density matrix*, *J. Comput. Phys.*, 334 (2017), pp. 1–15.
- [12] J. DEMMEL, L. GRIGORI, M. GU, AND H. XIANG, *Communication Avoiding Rank Revealing QR Factorization with Column Pivoting*, Tech. Report UCB/EECS-2013-46, EECS Department, University of California, Berkeley, 2013.
- [13] J. A. DUERSCH AND M. GU, *Randomized QR with column pivoting*, *SIAM J. Sci. Comput.*, 39 (2017), pp. C263–C291.
- [14] W. E, T. LI, AND J. LU, *Localized bases of eigensubspaces and operator compression*, *Proc. Natl. Acad. Sci. USA*, 107 (2010), pp. 1273–1278.
- [15] J. M. FOSTER AND S. F. BOYS, *Canonical configurational interaction procedure*, *Rev. Modern Phys.*, 32 (1960), pp. 300–302.

- [16] P. GIANNOZZI, S. BARONI, N. BONINI, M. CALANDRA, R. CAR, C. CAVAZZONI, D. CERESOLI, G. L. CHIAROTTI, M. COCCIONI, I. DABO, A. DAL CORSO, S. DE GIRONCOLI, S. FABRIS, G. FRATESI, R. GEBAUER, U. GERSTMANN, C. GOUGOUSSIS, A. KOKALJ, M. LAZZERI, L. MARTIN-SAMOS, N. MARZARI, F. MAURI, R. MAZZARELLO, S. PAOLINI, A. PASQUARELLO, L. PAULATTO, C. SBRACCIA, S. SCANDOLO, G. SCLAUZERO, A. P. SEITSONEN, A. SMOGUNOV, P. UMARI, AND R. M. WENTZCOVITCH, *Quantum ESPRESSO: A modular and open-source software project for quantum simulations of materials*, J. Phys. Condens. Matter, 21 (2009), 395502.
- [17] G. H. GOLUB AND C. F. VAN LOAN, *Matrix Computations*, 3rd ed., Johns Hopkins University Press, Baltimore, 1996.
- [18] F. GYGI, *Compact representations of Kohn–Sham invariant subspaces*, Phys. Rev. Lett., 102 (2009), 166406.
- [19] F. GYGI AND I. DUCHEMIN, *Efficient computation of Hartree–Fock exchange using recursive subspace bisection*, J. Chem. Theory Comput., 9 (2012), pp. 582–587.
- [20] M. Z. HASAN AND C. L. KANE, *Colloquium: Topological insulators*, Rev. Modern Phys., 82 (2010), 3045.
- [21] P. HOHENBERG AND W. KOHN, *Inhomogeneous electron gas*, Phys. Rev., 136 (1964), pp. B864–B871.
- [22] W. HUMPHREY, A. DALKE, AND K. SCHULTEN, *VMD—Visual molecular dynamics*, J. Molecular Graphics, 14 (1996), pp. 33–38.
- [23] W. KOHN, *Density functional and density matrix method scaling linearly with the number of atoms*, Phys. Rev. Lett., 76 (1996), pp. 3168–3171.
- [24] W. KOHN AND L. SHAM, *Self-consistent equations including exchange and correlation effects*, Phys. Rev., 140 (1965), pp. A1133–A1138.
- [25] L. LIN AND J. LU, *Sharp Decay Estimates of Discretized Green’s Functions for Schrödinger Type Operators*, arXiv:1511.07957, 2015.
- [26] M. W. MAHONEY AND P. DRINEAS, *Cur matrix decompositions for improved data analysis*, Proc. Nat. Acad. Sci. USA, 106 (2009), pp. 697–702.
- [27] R. MARTIN, *Electronic Structure: Basic Theory and Practical Methods*, Cambridge University Press, New York, 2004.
- [28] P. G. MARTINSSON, G. ORTÍ, N. HEAVNER, AND R. VAN DE GELIN, *Householder QR factorization with randomization for column pivoting (HQRRP)*, SIAM J. Sci. Comput., 39 (2017), pp. C96–C115.
- [29] N. MARZARI, A. A. MOSTOFI, J. R. YATES, I. SOUZA, AND D. VANDERBILT, *Maximally localized Wannier functions: Theory and applications*, Rev. Modern Phys., 84 (2012), pp. 1419–1475.
- [30] N. MARZARI AND D. VANDERBILT, *Maximally localized generalized Wannier functions for composite energy bands*, Phys. Rev. B, 56 (1997), pp. 12847–12865.
- [31] A. A. MOSTOFI, J. R. YATES, Y. LEE, I. SOUZA, D. VANDERBILT, AND N. MARZARI, *Wannier90: A tool for obtaining maximally-localised Wannier functions*, Comput. Phys. Comm., 178 (2008), pp. 685–699.
- [32] J. I. MUSTAFA, S. COH, M. L. COHEN, AND S. G. LOUIE, *Automated construction of maximally localized Wannier functions: Optimized projection functions method*, Phys. Rev. B, 92 (2015), 165134.
- [33] G. NENCIU, *Existence of the exponentially localised Wannier functions*, Comm. Math. Phys., 91 (1983), pp. 81–85.
- [34] V. OZOLIŅŠ, R. LAI, R. CAFLISCH, AND S. OSHER, *Compressed modes for variational problems in mathematics and physics*, Proc. Natl. Acad. Sci. USA, 110 (2013), pp. 18368–18373.
- [35] J. P. PERDEW, M. ERNZERHOF, AND K. BURKE, *Rationale for mixing exact exchange with density functional approximations*, J. Chem. Phys., 105 (1996), pp. 9982–9985.
- [36] E. PRODAN AND W. KOHN, *Nearsightedness of electronic matter*, Proc. Natl. Acad. Sci. USA, 102 (2005), pp. 11635–11638.
- [37] X. WU, A. SELLONI, AND R. CAR, *Order- N implementation of exact exchange in extended insulating systems*, Phys. Rev. B, 79 (2009), 085102.