

# Robust linear-scaling optimization of compact localized orbitals in density functional theory

Yifei Shi<sup>1</sup> and Rustam Z. Khaliullin<sup>1,\*</sup>

<sup>1</sup>*Department of Chemistry, McGill University, 801 Sherbrooke St. West, Montreal, QC H3A 0B8, Canada*

Locality of compact one-electron orbitals expanded strictly in terms of local subsets of basis functions can be exploited in density functional theory (DFT) to achieve linear growth of computation time with systems size, crucial in large-scale simulations. However, despite advantages of compact orbitals the development of practical orbital-based linear-scaling DFT methods has long been hindered because a compact representation of the electronic ground state is difficult to find in a variational optimization procedure. In this work, we showed that the slow and unstable optimization of compact orbitals originates from nearly-invariant mixing of occupied orbitals with the states that are mostly but not fully localized within the local vector subspace. We proposed a simple and practical method for identifying and removing the problematic nearly-invariant modes using an approximate Hessian and, as a result, developed a robust linear-scaling optimization procedure with a low computational overhead. We demonstrated the new method is highly efficient yet accurate in fixed-nuclei calculations and molecular dynamics simulations for a variety of systems ranging from semiconductors to insulators.

## INTRODUCTION

At present, Kohn-Sham (KS) density functional theory (DFT) is the most popular electronic structure method. Unfortunately, the computational cost of the conventional diagonalization-based KS DFT grows cubically with the number of atoms preventing its application to large systems. To overcome this limitation, substantial efforts have been directed to the development of linearly-scaling (LS) DFT methods [1–3].

In LS DFT methods, the delocalized eigenstates of the effective KS Hamiltonian must be replaced with an alternative set of *local* electron descriptors. Most LS methods explore the natural locality of the one-electron density matrix (DM) [4–10]. They include, among others, the Fermi operator expansion [11, 12], divide-and-conquer [13, 14], and direct DM optimization methods [4, 7, 8]. However, the variational optimization of the DM is very inefficient for accurate DFT calculations which require many basis functions per atom [1, 2, 8, 15, 16]. Therefore, the application of DM-based LS methods have been mostly restricted to minimal-basis tight-binding problems [17–19]. This issue is rectified in optimal-basis DM methods [20–22] that contract large basis sets into a small number of new localized functions and then optimize the DM in the contracted basis. Despite becoming the most popular LS DFT, the efficiency of optimal-basis methods is hampered by the costly optimization of both the contracted orbitals and the DM [23].

One-electron orbitals expanded strictly in subsets of localized basis functions centered on and near a given nuclei represent another type of local electron descriptors, alternative to the DM. Since their introduction [24–26] such localized orbitals have become known under different names including absolutely localized molecular orbitals [27], localized wave functions [28], non-orthogonal

generalized Wannier functions [20], multi-site support functions [21], and non-orthogonal localized molecular orbitals [29]. In this work, we will refer to them as compact localized molecular orbitals (CLMOs) to emphasize that their expansion coefficients are zero for the basis functions outside orbital’s localization subset. Unlike our previous work [16], we avoid using the ALMO acronym [27] because it is commonly used now to refer to compact orbitals restricted to *nonoverlapping* subsets [30–33] - a special case of CLMOs.

From the computational point of view, a direct variation of CLMOs is preferable to the DM for systems with nonvanishing band gap because LS can be achieved with significantly fewer variables. The computational advantages of orbitals-only LS DFT would be especially pronounced in accurate calculations that require many basis functions per atom. CLMOs are also advantageous from the physical point of view because they provide clear, chemically meaningful, transferable description of interactions between atoms and molecules [27, 31, 34–36]. Regrettably, the development of promising orbital-based LS methods has been hindered because of slow and inherently difficult variational optimization of CLMOs [1, 28, 37–40]. In addition to poor convergence of the variational procedure, the straightforward optimization of CLMOs, which cannot be constrained to stay both compact and orthogonal [27], often leads to a “collapsed” electronic state represented by linearly dependent occupied orbitals [28]. The progress of a typical variational optimization of CLMOs in Figure 1 suggests that the collapse, apparent from the vanishing determinant of the CLMO overlap matrix, is closely associated with the convergences problem.

Substantial efforts have been made to mitigate these problems. They include the introduction of an extra set of variational parameters [29, 39, 41] or the energy functionals that can be calculated without inverting the




FIG. 1. The determinant of the CLMO overlap matrix, the energy, and the maximum norm of the energy gradient with respect to the CLMO coefficients in a direct iterative optimization with the conjugate gradient algorithm. Linear system of four hydrogen fluoride molecules interacting through hydrogen bonds is described at the PBE/DZVP level of theory. CLMOs of a molecule are allowed to delocalize only over the two nearest neighbors. The zero energy is set at the energy of the fully delocalized state. For the CLMO state, the energy plateaus even though the gradient does not vanish. For the reference, the maximum norm of the gradient of same optimization using fully delocalized orbitals is also shown.

CLMO overlap matrix [28, 37, 42]. However, despite notable progress [38, 43, 44] the existing CLMO-based DFT algorithms establish convergence by measuring the rate of change of the energy in the course of the self-consistent field (SCF) procedure and require multiple steps to reach the final state. Furthermore, the SCF convergence has not been demonstrated using the norm of the gradient of the energy with respect to the CLMO coefficients – a much stronger convergence criteria that rules out a possibility of a sluggish optimization illustrated in Figure 1.

Importantly, the convergence problem has been reliably solved for weakly-interacting molecular systems. It has been shown that CLMOs can be optimized efficiently only if they are forced to be orthogonal to a set of auxiliary tightly localized nonoverlapping orbitals, precomputed and fixed on their molecules [16, 40, 45]. As will be discussed below, this approach works only when the auxiliary orbitals resemble the final variationally optimal CLMOs and thus cannot be applied to finite-gap systems with strong covalent bonds between atoms.

In this work, a detailed analysis of the origins of the convergence problem enabled us to develop a new linear-scaling CLMO-based DFT method that converges fast for systems of strongly interacting atoms and avoids collapsed states. The proposed method is conceptually simple and does not require any precomputed tightly-localized orbitals nor the optimization of auxiliary variables. These features greatly reduce its computational cost and make it straightforward to implement. Several tests in this work demonstrate the accuracy and efficiency

of the method for systems with finite band gap. It is important to state, however, that this method is not expected to be practical for metals and semimetals.

## THEORY

### Formalism, notation and main assumptions

In the first step, all atomic nuclei of a system, its electrons, and atom-centered basis set orbitals (AOs) — in our case Gaussian functions — are logically divided into nonoverlapping subsets called *localization centers*, often referred to as *fragments*. For a system with clearly defined molecules, a typical localization center includes all nuclei of a single molecule, associated AOs and electrons. For systems that cannot be partitioned into molecules — the subject of this work — a localization center can be represented by a single nuclei. As a result of the partitioning, each electron acquires a localization-center label. Within Kohn-Sham DFT, electrons are described by molecular orbitals (MOs)  $|\psi_{xi}\rangle$  whose indices now indicate that orbital  $i$  belongs to center  $x$ . Throughout this work, centers are labeled with Latin letters  $x, y, z$  whereas Latin letters  $i, j, k$  label MOs.

In the next step, each atom  $A$  is assigned a predefined element-specific radius  $R_c(A)$  that defines neighbors of each center in an obvious way: two centers are considered neighbors if there are atoms located within a sum of their radii. *Localization domain* for each center is defined as a subset of AOs that are localized on the neighbors of a center  $|\chi_{\bar{x}\mu}\rangle$ . In our notation, index  $\bar{x}$  refers to center  $x$  and its neighbors. Basis set orbitals are denoted with Greek letters  $\mu, \nu, \lambda, \kappa$ .

The basis set orbitals of a localization domain form a subspace in the one-electron vector space and the following projection operator serves as the identity operator on the subspace:

$$\hat{I}_{\bar{x}} = |\chi_{\bar{x}\mu}\rangle S^{\bar{x}\mu, \bar{x}\nu} \langle \chi_{\bar{x}\nu}|, \quad (1)$$

where  $S^{\bar{x}\mu, \bar{x}\nu}$  is a matrix element of the inverse overlap matrix  $S_{\bar{x}\mu, \bar{x}\nu} = \langle \chi_{\bar{x}\mu} | \chi_{\bar{x}\nu} \rangle$ . Note that the conventional tensor notation is used to work with the nonorthogonal orbitals [46]: covariant quantities are denoted with subscripts, contravariant quantities with superscripts, and summation is implied over the same orbital indices but not over the same center and domain indices.

By construction a basis set function may belong to several localization domains or, in other words, domains often overlap.

In the final step, the main approximation of CLMO methods is introduced. It is assumed that the electronic structure of the system can be described accurately by molecular orbitals that are completely localized within domains of their centers

$$|\psi_{xi}\rangle = \hat{I}_{\bar{x}} |\psi_{xi}\rangle \quad (2)$$

Thus this approximation imposes a blocked structure on the matrix of MO coefficients

$$|\psi_{xi}\rangle = |\chi_{\bar{x}\mu}\rangle T^{\bar{x}\mu}_{xi}. \quad (3)$$

The size of blocks is determined entirely by the pre-selected  $R_c$  and does not change with the number of molecules. Therefore, in the limit of large systems, the computational cost of the optimization grows linearly with the number of atoms offering a way of performing LS calculation directly with MOs.

In this work, we consider only spin-unpolarized orbitals evaluated at the  $\Gamma$ -point. The KS DFT energy functional can be written in the conventional way

$$E = 2\text{Tr} \left[ \hat{R} \hat{H} \right] - \frac{1}{2} \iint \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}d\mathbf{r}' + E_{\text{XC}} - \int v_{\text{XC}}(\mathbf{r})\rho(\mathbf{r})d\mathbf{r}, \quad (4)$$

where  $\hat{H}$  is the Kohn-Sham Hamiltonian,  $\hat{R}$  is the projection operator onto the occupied subspace, and  $\rho(\mathbf{r}) = 2|\mathbf{r}|\hat{R}|\mathbf{r}\rangle$  is the electron density.  $\hat{R}$  written to take into account the nonorthogonality of CLMOs

$$\hat{R} = \sum_{x,y} |\psi_{xi}\rangle \sigma^{xi,yj} \langle \psi_{yj}|, \quad (5)$$

where  $\sigma^{xi,yj}$  is a matrix element of the inverse of the CLMO overlap matrix  $\sigma_{zk,xj} = \langle \psi_{zk} | \psi_{xj} \rangle$ .

### Convergence problem and Hessian eigenspectrum

The direct minimization of the energy functional with respect to molecular orbital coefficients is a straightforward reliable low-cost method optimization for fully delocalized orbitals ( $R_c \rightarrow \infty$ ) [47–49] and completely localized orbitals ( $R_c = 0$ ) [16]. The preconditioned conjugate gradient algorithm (PCG) is a particularly efficient energy minimizer that requires iterative evaluation of the energy gradient

$$G_{\bar{x}\mu}^{xi} \equiv \frac{\partial E}{\partial T^{\bar{x}\mu}_{xi}} = 4 \langle \chi_{\bar{x}\mu} | (\hat{I} - \hat{R}) \hat{H} | \psi^{xi} \rangle \quad (6)$$

and only a single inversion of a preconditioner. The latter is typically chosen as an easily-invertible approximation to the exact Hessian. It has been found [16, 48] that, for the two special cases  $R_c \rightarrow \infty$  and  $R_c = 0$ , the preconditioner defined on domain  $\bar{x}$

$$P_{\bar{x}\mu,\bar{x}\nu} = 4 \langle \chi_{\bar{x}\mu} | (\hat{I} - \hat{R})(\hat{I} + \hat{H})(\hat{I} - \hat{R}) | \chi_{\bar{x}\nu} \rangle \quad (7)$$

provides the same rate of convergence of the PCG algorithm as the exact Hessian but at a fraction of the inversion cost. The relation between the preconditioner in Eq. (7) and the exact Hessian is presented in the Supplementary Material.

In the general case of finite  $R_c$ , the PCG-based optimization of CLMOs has been investigated thoroughly as a low-cost alternative to LS DM-based DFT methods [16, 28, 29, 37–39, 42, 43]. This approach is expected to be very efficient because both the CLMO coefficient and gradient matrices are small (i.e. number of columns is much smaller than the number of rows) and enforced to be sparse. In addition to this, the inversion of the approximate Hessian in Eq. 7 can be done fast, domain-by-domain. Unfortunately, the PCG algorithm, as well as all other optimization procedures (e.g. Newton-Raphson or DIIS-accelerated diagonalizations) [27], suffer from the aforementioned slow convergence and orbital collapse problems. A closer inspection of the eigenvalues of the preconditioner obtained by solving the generalized eigenproblem for each domain

$$P_{\bar{x}\mu,\bar{x}\nu} A^{\bar{x}\nu}_{xp} = S_{\bar{x}\mu,\bar{x}\lambda} A^{\bar{x}\lambda}_{xp} \Lambda_{xp}, \quad (8)$$

reveals the origin of this and other closely related previously reported problems [1]. The eigenvalues, which approximate the energy curvature along the optimization direction  $|d_{xp}\rangle$  given by the corresponding eigenvector,  $A^{\bar{x}\nu}_{xp} \equiv \langle \chi^{\bar{x}\nu} | d_{xp} \rangle$ , can be divided into three categories according to their magnitudes. Zero eigenvalues in the first category represent optimization directions towards occupied orbitals localized completely within the same domain (Figure 2, case A). As expected the energy is invariant along these occupied-occupied mixing modes. The second category includes large nonzero eigenvalues that represent the optimization in the direction toward unoccupied orbitals (Figure 2, case B) of the domain. These two categories are the only present in the rapidly converging optimization of fully delocalized orbitals ( $R_c \rightarrow \infty$ ) and absolutely localized orbitals ( $R_c = 0$ ). The eigenvalues classified as the third category are extremely small but nonzero (Figure S1 in the Supplementary Material). They appear only for finite  $R_c$  when domains share basis set functions. The optimization along these low-curvature nearly-invariant directions is difficult to converge because various analytical approximations (e.g. approximate Hessian, quadratic linear search) and numerical noise (e.g. finite DFT grids) make calculations imprecise. In other words, optimization with finite  $R_c$  becomes ill-conditioned and the low-curvature modes represents the major barrier to the practical use of promising orbital-based LS DFT methods.

## RESULTS AND DISCUSSION

### Nature of low-curvature modes

What is the physical origin of the low-curvature optimization modes? It has been suggested previously [1] that the sluggish optimization in orbital-based LS DFT




FIG. 2. Illustration of the origin of low-curvature modes in a model vector space spanned by three basis set functions  $|\text{AO}_1\rangle$ ,  $|\text{AO}_2\rangle$ ,  $|\text{AO}_3\rangle$ . The left panel shows  $|\text{CLMO}_1\rangle$  (blue) confined to its domain spanned by  $|\text{AO}_1\rangle$  and  $|\text{AO}_2\rangle$  as well as  $|\text{CLMO}_2\rangle$  (red, three cases) confined to its domain spanned by  $|\text{AO}_2\rangle$  and  $|\text{AO}_3\rangle$ . The right panel shows the behavior of the energy as a function of the position of  $|\text{CLMO}_1\rangle$  — angle  $\alpha$  — for three typical cases. Case C shows that the low-curvature modes arise when  $|\text{CLMO}_2\rangle$  lies almost entirely in the domain of  $|\text{CLMO}_1\rangle$ .

is due to the inexact invariance of the energy with respect to the mixing of occupied orbitals. Although this might indeed present a problem for a series of functionals based on DMs that are not fully idempotent [28, 37, 42] the DM used in this work is constructed with the inverse of the overlap matrix and, therefore, is properly idempotent and exactly invariant to the mixing among occupied states.

We instead suggest that the low-curvature optimization modes represent *hybrid* directions, which are neither occupied nor virtual. From the point of view of the local vector subspace of  $\bar{x}$ , these directions exist because CLMOs of neighbor centers are only partially localized on  $\bar{x}$ . Hybrid directions, originating from CLMOs that are almost but not entirely localized on the domain, are expected to be particularly problematic (Figure 2, case C).

Our hypothesis implies that a hybrid low-curvature mode  $|d_{\bar{x}p}\rangle$  has only small component in the unoccupied subspace of domain  $\bar{x}$ . This component is measured by the residue  $\Delta_{\bar{x}p}$

$$\Delta_{\bar{x}p} \equiv \langle d_{\bar{x}p} | \hat{I}_{\bar{x}} - \hat{R}_{\bar{x}} | d_{\bar{x}p} \rangle, \quad (9)$$

where  $\hat{R}_{\bar{x}}$  is a projector onto the subspace of the occupied CLMOs trimmed with operator  $\hat{T}_{\bar{x}} \approx \hat{I}_{\bar{x}}$  to be fully localized on  $\bar{x}$

$$\hat{R}_{\bar{x}} = \sum_{y,z \in \bar{x}} \hat{T}_{\bar{x}} |\psi_{yi}\rangle \sigma_{\bar{x}}^{yi,zj} \langle \psi_{zj} | \hat{T}_{\bar{x}} \quad (10)$$

and  $\sigma_{\bar{x}}$  is the overlap matrix of the truncated orbitals.

Figure 3 shows that the small-curvature modes indeed lie mostly outside the unoccupied space for a variety of materials (see Supplementary Material for details). It is also interesting to note that the unoccupied fraction of the low-curvature modes increases with the strength of interatomic delocalization (i.e. the degree of covalency of bonding).

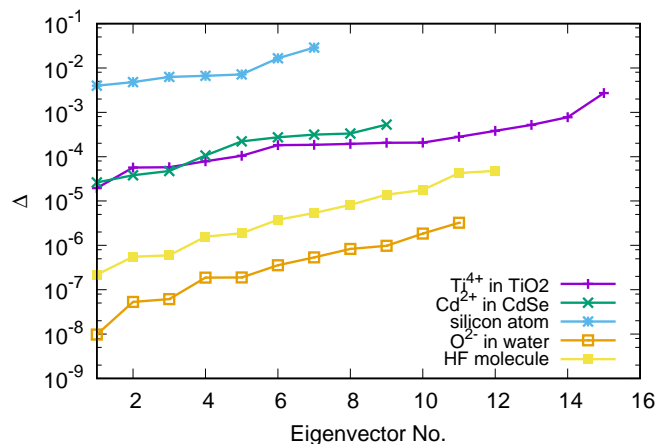


FIG. 3. The norm of the projection of the small-curvature modes on the unoccupied subspace of a domain. Preconditioner eigenvectors with the eigenvalues smaller than 0.02 a.u. are chosen as small-curvature modes. The following localization centers are considered:  $\text{Ti}^{4+}$  ions in the  $\text{TiO}_2$  rutile lattice,  $\text{Cd}^{2+}$  ions in the  $\text{CdSe}$  Wurtzite lattice, Si atom in the diamond silicon lattice,  $\text{O}^{2-}$  in a water tetramer system, hydrogen fluoride molecule in a linear tetramer system. The BLYP/TZV2P level of theory is used for water and PBE/DZVP in all other tests.

Thus the optimization along the low-curvature modes of a domain represents mixing of its occupied CLMO with the occupied orbitals of its neighbors that are not fully localized on the domain. Unsurprisingly, such orbital variations often lead to linear dependencies among occupied orbitals and eventual collapse of the optimization.

The proposed explanation of the nature of the low-curvature modes is consistent with the fact that, for molecular partitioning, the number of the low-curvature modes is equal to the sum of occupied orbitals on the neighbor fragments. It also explains why the two-stage procedure of Ref. 16 works so well for molecular systems.

In the first stage of this procedure,  $R_c$  is set to zero and the resulting block-diagonal orbitals  $|\psi_{xi}^0\rangle$  are optimized variationally to construct the occupied space projector  $\hat{R}^0$ . This projector is then fixed,  $R_c$  is reset to its original finite value to allow intercenter electron delocalization, and the following trial orbitals are optimized with respect to CLMOs coefficients:

$$|\psi_{xi}\rangle = |\psi_{xi}^0\rangle + \hat{I}_{\bar{x}}(\hat{I} - \hat{R}^0)|\chi_{\bar{x}\mu}\rangle \bar{T}_{xi}^{\bar{x}\mu} \quad (11)$$

For these CLMOs, the gradient and preconditioner

$$\tilde{G}_{\bar{x}\mu}^{xi} \equiv \frac{\partial E}{\partial \bar{T}_{xi}^{\bar{x}\mu}} = \langle \chi_{\bar{x}\mu} | \hat{I} - \hat{R}^0 | \chi_{\bar{x}\nu} \rangle G_{\bar{x}\nu}^{xi} \quad (12)$$

$$\tilde{P}_{\bar{x}\mu, \bar{x}\nu} = \langle \chi_{\bar{x}\mu} | \hat{I} - \hat{R}^0 | \chi_{\bar{x}\lambda} \rangle P_{\bar{x}\lambda, \bar{x}\kappa} \langle \chi_{\bar{x}\kappa} | \hat{I} - \hat{R}^0 | \chi_{\bar{x}\nu} \rangle \quad (13)$$

are related to those in Eqs. (6) and (7) via domain-specific operators

$$\hat{I}_{\bar{x}}(\hat{I} - \hat{R}^0)\hat{I}_{\bar{x}} \quad (14)$$

As shown in Ref. 16, these operators are essential to converging orbital optimization for molecular systems. The present work explains that this is because  $\hat{R}^0$  satisfies two important requirements. First, it is constructed from CLMOs fully localized on their centers. This ensures that each localization domain contains an integer number of occupied states. Second,  $\hat{R}^0$  is already close to the final converged DM for systems of weakly-interacting molecules. This, therefore, guarantees that correct low-curvature optimization modes are removed when in the second optimization stage when  $\hat{I}_{\bar{x}}\hat{R}^0\hat{I}_{\bar{x}}$  is projected out.

Unfortunately, the two-stage approach does not work for systems with strong covalent bonds between localization centers (i.e. atoms). This is illustrated in Figure 4 for the cadmium selenide lattice. The main reason for the deteriorating convergence is the inability of  $\hat{R}^0$ , which is constructed from the block-diagonal CLMOs, to adequately represent the low-curvature modes in the delocalized states encountered in later stages of the optimization.

### Proposed solution to the optimization problem

We propose to solve the convergence problem by detecting the low-curvature modes directly by diagonalizing the approximate Hessian in Eq. (7) and avoiding the optimization along these modes altogether. Although this procedure does not produce fully optimized orbitals they are still expected to provide accurate representation of the ground state. This is because the low-curvature modes are associated mostly with mixing of occupied and nearly-occupied orbitals and, therefore, are expected to be *shallow* optimization directions. That is, they are unlikely to produce a noticeable variational decrease in the energy.

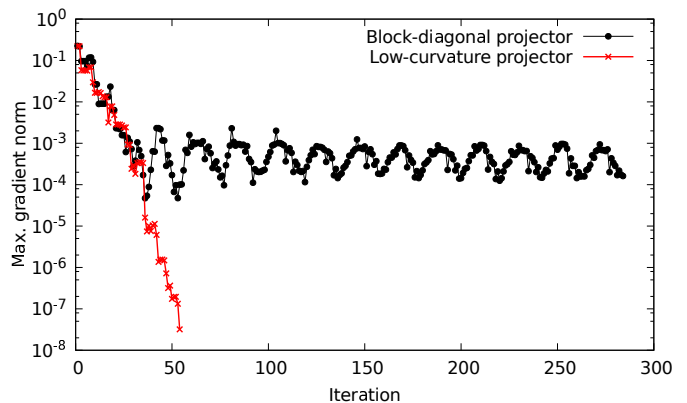


FIG. 4. The maximum norm of the energy gradient for the PBE/DZVP basis [50] optimization of the CLMOs for the hexagonal wurtzite CdSe lattice.  $R_c = 3.5$  Å. Modes with the preconditioner eigenvalues lower than 0.05 a.u. are removed from the optimization. The block-diagonal projector is described in Eq. 14.

To filter out the low-curvature modes we construct the projector onto these modes

$$\hat{L}_{\bar{x}} = |d_{\bar{x}p}\rangle \Theta[\Lambda_{\bar{x}p} - \Lambda_c] \langle d_{\bar{x}p}|, \quad (15)$$

where  $\langle d_{\bar{x}p}|$  are eigenvectors of the approximate Hessian from Eq. (8),  $\Theta$  is the reversed unit step function

$$\Theta[\Lambda_{\bar{x}p} - \Lambda_c] = \begin{cases} 1 & \Lambda_{\bar{x}p} \leq \Lambda_c, \\ 0 & \Lambda_{\bar{x}p} > \Lambda_c, \end{cases} \quad (16)$$

and  $\Lambda_c$  is the curvature threshold below which the optimization mode is classified as a low-curvature mode. The low curvature modes are then projected out from the the gradient and preconditioner in the PCG optimization algorithm

$$\tilde{G}_{\bar{x}\mu}^{xi} = \langle \chi_{\bar{x}\mu} | \hat{I} - \hat{L}_{\bar{x}} | \chi_{\bar{x}\nu} \rangle G_{\bar{x}\nu}^{xi} \quad (17)$$

$$\tilde{P}_{\bar{x}\mu, \bar{x}\nu} = \langle \chi_{\bar{x}\mu} | \hat{I} - \hat{L}_{\bar{x}} | \chi_{\bar{x}\lambda} \rangle P_{\bar{x}\lambda, \bar{x}\kappa} \langle \chi_{\bar{x}\kappa} | \hat{I} - \hat{L}_{\bar{x}} | \chi_{\bar{x}\nu} \rangle. \quad (18)$$

As a result, the PCG optimization procedure neglects the optimization modes with the curvature below pre-selected threshold  $\Lambda_c$ .

We will refer to the new method based on Eqs. (17) and (18) as low-curvature projector (LCP) method.

### Remarks

It is important to make several remarks about the LCP method. Although the gradient and preconditioner equations of the LCP approach resemble those of the two-stage procedure the key difference between them is that  $\hat{L}_{\bar{x}}$  depends on the variational parameters  $T$  while  $\hat{I}_{\bar{x}}\hat{R}^0\hat{I}_{\bar{x}}$

does not. As a consequence, there is no a well-defined trial orbitals and energy functional that correspond to the LCP gradient exactly. However, the LCP gradient in Eq. (17) is an approximation to the gradient of the conventional energy functional with the following trial CLMOs

$$|\psi_{xi}\rangle = \hat{I}_{\bar{x}}(\hat{I} - \hat{R}_{\bar{x}})|\chi_{\bar{x}\mu}\rangle T^{\bar{x}\mu}_{xi} \quad (19)$$

where  $\hat{R}_{\bar{x}}$  is defined in Eq. (10) as the  $T$ -dependent projector onto the subspace of the CLMOs that are truncated to be fully localized on  $\bar{x}$ . The LCP gradient can be obtained from these trial orbitals if the  $T$ -dependence of  $\hat{R}_{\bar{x}}$  is ignored upon differentiation and  $\hat{R}_{\bar{x}}$  is approximated with the low-curvature projector  $\hat{L}_{\bar{x}}$ . As described above the latter approximation makes sense because the low-curvature modes arise from the incomplete localization of the neighbors' CLMOs on  $\bar{x}$ .

It is worth mentioning in passing that finding the exact gradient of the trial CLMOs in Eq. (19) is tremendously difficult both analytically and algorithmically because of the cyclic dependence of  $\hat{R}_{\bar{x}}$  and  $|\psi_{xi}\rangle$  on each other. The LCP approach described here represents a compromise between a computationally efficient and fully variational optimization.

While previous CLMO works [16, 40, 45] dealt with weakly-interacting fragments (e.g. molecules) the LCP approach represents the ultimate atomic partitioning scheme that draws fragment boundaries across strong covalent bonds. Although the LCP optimization is not fully variational it finds the orbitals and energy very close to the ground state if  $\Lambda_c$  is small. Moreover, the LCP approach passes even more stringent test as it appears to produce stable molecular dynamics description of chemical reactions.

It should be noted that any divide-and-conquer approach, including the LCP method, assigns electrons to localization centers *a priori* and, therefore, requires at least a rough idea about the optimal electron distribution. Although the LCP procedure presented here is not applicable to systems with completely unknown bonding properties this issue can be resolved by employing the global optimizer described in Ref. 42 in conjunction with the current approach.

### Implementation

The procedure for finding and projecting out the low-curvature modes is implemented in the CP2K software package [51]. CP2K relies on the mixed Gaussian and plane wave representation of the electronic degrees of freedom [52] and is an ideal platform for the new orbital-based LS method: just a few tightly-localized Gaussian AOs can provide an accurate representation of CLMOs, whereas plane waves ensure a fast LS construction of the KS Hamiltonian for large systems.

All matrix multiplications are performed with the DBCSR library [53] designed for massively-parallel linear-scaling handling of large sparse matrices. A special care is taken to reduce the computational overhead of the optimization procedure for large Gaussian basis sets. To this end, the order of matrix multiplications in Eqs. (17) and (18) is chosen to avoid steps that scale cubically with the size of the Gaussian basis set. The diagonalization of the preconditioner matrices is done independently for each domain.

It is important to note that the construction of the DM requires the inversion of the CLMO overlap matrix. This matrix is small and its size is independent of the size of the basis set. However, it is not confined to individual domains. This inversion is carried out using the iterative Hotelling method [54] that is based entirely on matrix multiplications and becomes LS when the system is large and the CLMO overlap and its inverse are sparse.

Because of the similarity between the new method and the ALMO SCF for molecular systems only minor modifications in the existing ALMO DFT module of CP2K are necessary. The main difference between the algorithms in the two approaches is that the PCG optimization of CLMOs in the LCP method might need more frequent re-evaluation and inversion of the preconditioner.

### Accuracy and efficiency tests

Figure 4 shows that the new optimization procedure converges rapidly for the hexagonal CdSe lattice – a particularly challenging case for LS methods considering the small band gap of this material and strong interactions between atoms. In contrast, it is difficult to achieve convergence for this system using the two-stage method described in Ref. [16]. It is also important to note that the final energy obtained with the new method differs from the energy obtained after 250 iterations of the unconverged two-stage procedure only by 0.03 kJ/mol per Cd atom.

The accuracy of the LCP energies as a function of  $R_c$  is shown in Figure 5 for several challenging systems with strong interaction and significant electron delocalization between fragments: CdSe in the hexagonal wurtzite lattice; liquid water in which each atom is treated as a fragment; silicon in the cubic diamond lattice. Figure 5 demonstrates that the energy converges to the correct ground state energy as  $R_c$  increases. The error of the localization constraints imposed on electrons is small. The insets in Figure 5 also shows that the error of projecting out the low-curvature optimization modes is even smaller than that introduced by finite  $R_c$ . These tests confirm our hypothesis that the low-curvature directions do not produce a noticeable decrease in the energy.

To demonstrate the accuracy of the LCP energies further, we performed a 500 K Monte Carlo (MC) simula-

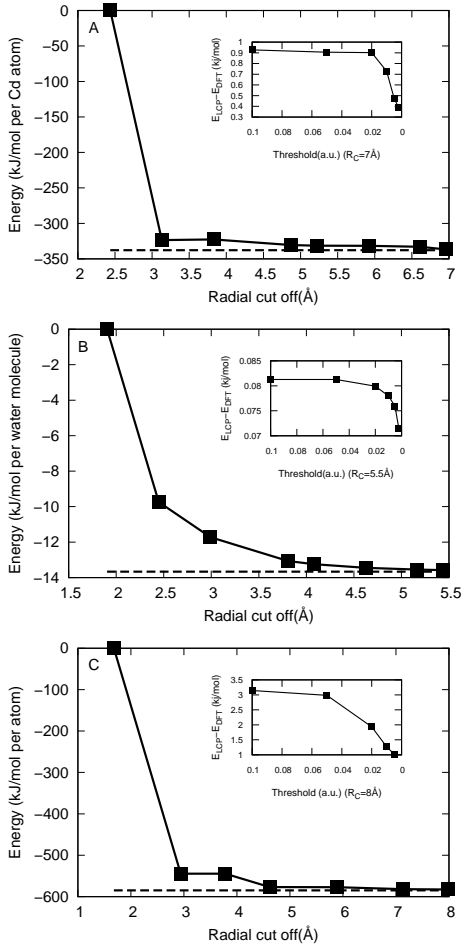


FIG. 5. Energy from the low-curvature projector and the dependence on the localization radii  $R_c$  for: A. PBE/DZVP calculation for the hexagonal wurtzite phase of CdSe, B. BLYP/TZV2P calculation for atomically-partitioned liquid water, C. PBE/DZVP calculation for the cubic diamond phase of silicon. The convergence criteria is  $\|\tilde{G}\|_{\max} < 10^{-5}$  a.u. The zero energy is set at the energy of the fully optimized block-diagonal CLMOs ( $R_c = 0$ ). Dashed line shows the reference energy of conventional DFT calculation ( $E_{\text{DFT}}$ ) without localization constrains ( $R_c \rightarrow \infty$ ). Inset: LCP energy dependence on the low-curvature threshold  $\Lambda_c$  at a fixed  $R_c$ . For the inset, the energy of the unconstrained orbitals is chosen as zero.

tion of silicon in the cubic diamond lattice with a defect created by replacing two neighboring silicon atoms with two carbon atoms. Figure 6 shows that the silicon-silicon radial distribution function (RDF) at equilibrium almost perfectly reproduces that calculated using conventional DFT methods for fully delocalized electrons.

Molecular dynamics (MD) simulations present an even more stringent test on the accuracy of the LCP method. While minor energy errors are not crucial in fixed-nuclei calculations, geometry optimizations, and Monte-Carlo sampling they tend to accumulate in MD trajectories leading to non-physical sampling and eventual failure of




FIG. 6. Silicon-silicon RDF calculated for the cubic diamond phase of silicon with carbon atoms introduced to create a point defect. XC/basis simulations are performed at  $T = 500$  K with the conventional orbital-transformation [48] method for the delocalized electrons and the LCP method with  $R_c = 6$  Å and  $\Lambda_c = 0.02$  a.u. The system includes 62 silicon atoms and two carbon atoms.

simulations. Accurate calculation of forces is therefore crucial for molecular dynamics simulations, the stability of which are judged by the accumulated drift in the conserved quantity (e.g. total energy of the system). Although CLMOs obtained with the LCP method are not strictly variational we chose to neglect this error and invoke the Hellmann–Feynman theorem [55] in the calculation of atomic forces. This procedure is then used to perform an LCP-based MD simulation of a protonated water nanocluster containing 62 water molecules and two protons. In the course of a three-picosecond MD simulation, the two protons hop around breaking and forming covalent bonds with water molecules. To reproduce motion of protons around the nanodroplet each atom was treated as a localization center. Setting  $R_c = 4$  Å and  $\Lambda_c = 0.005$  a.u. produces sufficiently optimized CLMOs to ensure the stability of the MD simulation. Figure 7 shows that the drift in the conserved quantity is acceptably small compared to the magnitude of the fluctuations in the potential energy. The applicability of the LCP method in MD simulations will be explored further especially in conjunction with the modified Langevin integrator method [56].

To test the computational efficiency of the newly designed method, we compared its performance to the orbital transformation (OT) method [48, 58] — a well-optimized low-prefactor cubic scaling DFT method for conventional fully delocalized orbitals. The calculation is done for the hexagonal CdSe system of different sizes.  $R_c$  is set to 3.5 Å, which is large enough to accurately describe the electronic structure of the system (Figure 5) and  $\Lambda_c = 0.02$  a.u. Figure 8 demonstrates that the LCP method is asymptotically LS. The LS regime is achieved when the MO overlap matrix and its inverse are sparse.

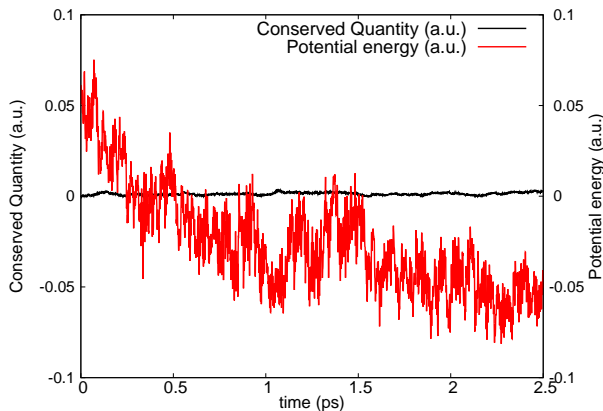


FIG. 7. The potential energy and the conserved quantity in a LCP-based MD simulation of a protonated water nanocluster  $(\text{H}^+)_{2}(\text{H}_2\text{O})_{62}$ . The temperature of the NVT simulations is set to 298 K and controlled by a canonical velocity re-scaling thermostat [57] with the coupling time constant of 50 fs. The potential energy and conserved quantity are shifted for clarity.

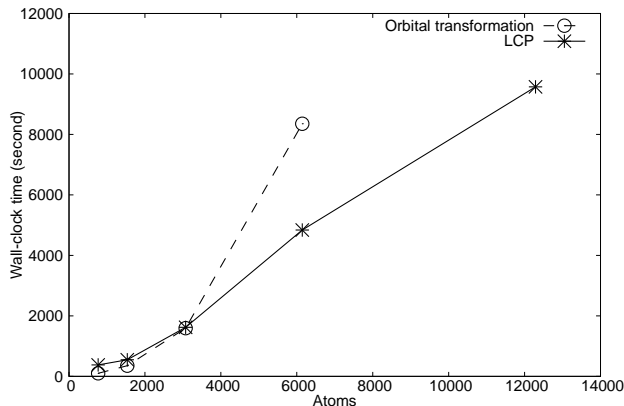


FIG. 8. CLMO optimization timing benchmark for the hexagonal CdSe lattice. The PBE/DZVP calculation with  $R_c = 3.5 \text{ \AA}$  and  $\Lambda_c = 0.02 \text{ a.u.}$  is done on 256 compute cores.

For a dense 3D system like CdSe it takes on the order of  $\sim 3000$  atoms.

## CONCLUSIONS

This work presents a new linear-scaling DFT method based entirely on the conceptually simple variational optimization of localized molecular orbitals. Unlike many existing LS methods, the method does not perform the typical concomitant variation of the density matrix and thus has low computational cost. The notorious convergence problem that has plagued orbital-based LS DFT methods is resolved by avoiding the optimization of orbitals along low-curvature modes – the directions associated with tiny eigenvalues of the electronic Hessian. An efficient procedure is designed to identify the low-

curvature modes on-the-fly without the computationally costly construction and diagonalization of the Hessian. It is also shown that the low-curvature modes — an unavoidable repercussion of imposing localization constraints — can be safely neglected because they are associated with the mixing of nearly occupied states and do not produce a noticeable variational decrease in the energy.

The new methodology, which is expected to be applicable to systems with nonvanishing band gap, is tested on a variety of materials including rutile phase of titania, hexagonal phase of cadmium selenide, and cubic diamond phase of silicon with and without defects. These tests demonstrate that the method is accurate and efficient even when localization centers are represented by single atoms and there are strong covalent interaction between the centers. Furthermore, preliminary tests on protonated water nanoclusters suggest that the atomic partitioning does not present problems for the new method and it is sufficiently robust to enable stable molecular dynamics simulation of bond-breaking and formation processes.

The developed LS DFT method is expected to have a significant impact on computational modeling of complex systems. Due to its low computational overhead, the method will enable DFT calculations on previously inaccessible length scales making completely new chemical phenomena amenable to simulations.

## ACKNOWLEDGMENTS

The research was funded by the Natural Sciences and Engineering Research Council of Canada through the Discovery Grant (RGPIN-2016-05059). The authors are grateful to Compute Canada for computer time.

\* rustam.khaliullin@mcgill.ca

- [1] Goedecker, S. Linear scaling electronic structure methods. *Rev. Mod. Phys.* **71**, 1085 (1999).
- [2] Bowler, D. & Miyazaki, T. Methods in electronic structure calculations. *Rep. Prog. Phys.* **75**, 036503 (2012).
- [3] Zaleśny, R., Papadopoulos, M. G., Mezey, P. G. & Leszczynski, J. *Linear-Scaling Techniques in Computational Chemistry and Physics: Methods and Applications*, vol. 13 (Springer Science & Business Media, 2011).
- [4] Li, X.-P., Nunes, R. & Vanderbilt, D. Density-matrix electronic-structure method with linear system-size scaling. *Phys. Rev. B* **47**, 10891 (1993).
- [5] Lee, T.-S., York, D. M. & Yang, W. Linear-scaling semiempirical quantum calculations for macromolecules. *J. Chem. Phys.* **105**, 2744–2750 (1996).
- [6] Li, X. S., Millam, J. M., Scuseria, G. E., Frisch, M. J. & Schlegel, H. B. Density matrix search using direct inversion in the iterative subspace as a linear scaling al-



- ternative to diagonalization in electronic structure calculations. *J. Chem. Phys.* **119**, 7651 (2003).
- [7] Shao, Y., Saravanan, C., Head-Gordon, M. & White, C. A. Curvy steps for density matrix-based energy minimization: Application to large-scale self-consistent-field calculations. *J. Chem. Phys.* **118**, 6144 (2003).
- [8] VandeVondele, J., Borstnik, U. & Hutter, J. Linear scaling self-consistent field calculations with millions of atoms in the condensed phase. *J. Chem. Theory Comput.* **8**, 3565–3573 (2012).
- [9] Kussmann, J., Beer, M. & Ochsenfeld, C. Linear-scaling self-consistent field methods for large molecules. *Wiley Interdisciplinary Reviews: Computational Molecular Science* **3**, 614–636 (2013).
- [10] Aarons, J., Sarwar, M., Thompsett, D. & Skylaris, C.-K. Perspective: Methods for large-scale density functional calculations on metallic systems. *J. Chem. Phys.* **145**, 220901 (2016).
- [11] Goedecker, S. & Colombo, L. Efficient linear scaling algorithm for tight-binding molecular dynamics. *Phys. Rev. Lett.* **73**, 122 (1994).
- [12] Goedecker, S. & Teter, M. Tight-binding electronic-structure calculations and tight-binding molecular dynamics with localized orbitals. *Phys. Rev. B* **51**, 9455 (1995).
- [13] Yang, W. Direct calculation of electron density in density-functional theory. *Phys. Rev. Lett.* **66**, 1438 (1991).
- [14] Yang, W. A local projection method for the linear combination of atomic orbital implementation of density-functional theory. *J. Chem. Phys.* **94**, 1208–1214 (1991).
- [15] Arita, M., Bowler, D. R. & Miyazaki, T. Stable and efficient linear scaling first-principles molecular dynamics for 10000+ atoms. *J. Chem. Theory Comput.* **10**, 5419–5425 (2014).
- [16] Khaliullin, R. Z., VandeVondele, J. & Hutter, J. Efficient linear-scaling density functional theory for molecular systems. *J. Chem. Theory Comput.* **9**, 4421–4427 (2013).
- [17] Richters, D. & Kühne, T. D. Self-consistent field theory based molecular dynamics with linear system-size scaling. *J. Chem. Phys.* **140**, 134109 (2014).
- [18] Goringe, C., Bowler, D. & Hernandez, E. Tight-binding modelling of materials. *Rep. Prog. Phys.* **60**, 1447 (1997).
- [19] Ratcliff, L., Conduit, G., Hine, N. & Haynes, P. Band structure interpolation using optimized local orbitals from linear-scaling density functional theory. *Phys. Rev. B* **98**, 125123 (2018).
- [20] Skylaris, C.-K., Haynes, P. D., Mostofi, A. A. & Payne, M. C. Introducing onetep: Linear-scaling density functional simulations on parallel computers. *J. Chem. Phys.* **122**, 084119 (2005).
- [21] Nakata, A., Bowler, D. R. & Miyazaki, T. Optimized multi-site local orbitals in the large-scale dft program conquest. *Phys. Chem. Chem. Phys.* **17**, 31427–31433 (2015).
- [22] Mohr, S. *et al.* Accurate and efficient linear scaling dft calculations with universal applicability. *Phys. Chem. Chem. Phys.* **17**, 31360–31370 (2015).
- [23] Mostofi, A. A., Haynes, P. D., Skylaris, C.-K. & Payne, M. C. Preconditioned iterative minimization for linear-scaling electronic structure calculations. *J. Chem. Phys.* **119**, 8842 (2003).
- [24] Matsuoka, O. Expansion methods for adams–gilbert equations. i. modified adams–gilbert equation and common and fluctuating basis sets. *J. Chem. Phys.* **66**, 1245–1254 (1977).
- [25] Stoll, H. & Preuß, H. On the direct calculation of localized hf orbitals in molecule clusters, layers and solids. *Theor. Chem. Acc.* **46**, 11–21 (1977).
- [26] Mehler, E. L. Self-consistent, nonorthogonal group function approximation for polyatomic systems. i. closed shells. *J. Chem. Phys.* **67**, 2728–2739 (1977).
- [27] Stoll, H., Wagenblast, G. & Preuß, H. On the use of local basis sets for localized molecular orbitals. *Theor. Chem. Acc.* **57**, 169–178 (1980).
- [28] Ordejón, P., Drabold, D. A., Martin, R. M. & Grumbach, M. P. Linear system-size scaling methods for electronic-structure calculations. *Phys. Rev. B* **51**, 1456 (1995).
- [29] Burger, S. K. & Yang, W. Linear-scaling quantum calculations using non-orthogonal localized molecular orbitals. *J. Phys.: Condens. Matter* **20**, 294209 (2008).
- [30] Khaliullin, R. Z., Head-Gordon, M. & Bell, A. T. An efficient self-consistent field method for large systems of weakly interacting components. *J. Chem. Phys.* **124**, 204105 (2006).
- [31] Khaliullin, R. Z., Cobar, E. A., Lochan, R. C., Bell, A. T. & Head-Gordon, M. Unravelling the origin of intermolecular interactions using absolutely localized molecular orbitals. *J. Phys. Chem. A* **111**, 8753–8765 (2007).
- [32] Khaliullin, R. Z., Bell, A. T. & Head-Gordon, M. Analysis of charge transfer effects in molecular complexes based on absolutely localized molecular orbitals. *J. Chem. Phys.* **128**, 184112 (2008).
- [33] Horn, P. R., Sundstrom, E. J., Baker, T. A. & Head-Gordon, M. Unrestricted absolutely localized molecular orbitals for energy decomposition analysis: Theory and applications to intermolecular interactions involving radicals. *J. Chem. Phys.* **138**, 134119 (2013).
- [34] Mo, Y., Gao, J. & Peyerimhoff, S. D. Energy decomposition analysis of intermolecular interactions using a block-localized wave function approach. *J. Chem. Phys.* **112**, 5530–5538 (2000).
- [35] Khaliullin, R. Z. & Kühne, T. D. Microscopic properties of liquid water from combined ab initio molecular dynamics and energy decomposition studies. *Phys. Chem. Chem. Phys.* **15**, 15746–15766 (2013).
- [36] Mo, Y. The block-localized wavefunction (blw) perspective of chemical bonding. *The Chemical Bond: Fundamental Aspects of Chemical Bonding* 199–232 (2014).
- [37] Mauri, F., Galli, G. & Car, R. Orbital formulation for electronic-structure calculations with linear system-size scaling. *Phys. Rev. B* **47**, 9973 (1993).
- [38] Fattebert, J.-L. & Gygi, F. Linear scaling first-principles molecular dynamics with controlled accuracy. *Comput. Phys. Commun.* **162**, 24–36 (2004).
- [39] Peng, L., Gu, F. L. & Yang, W. Effective preconditioning for ab initio ground state energy minimization with non-orthogonal localized molecular orbitals. *Phys. Chem. Chem. Phys.* **15**, 15518–15527 (2013).
- [40] Tsuchida, E. Ab initio molecular dynamics simulations with linear scaling: application to liquid ethanol. *J. Phys.: Condens. Matter* **20**, 294212 (2008).
- [41] Yang, W. Absolute-energy-minimum principles for linear-scaling electronic-structure calculations. *Phys. Rev. B* **56**, 9294 (1997).
- [42] Kim, J., Mauri, F. & Galli, G. Total-energy global optimizations using nonorthogonal localized orbitals. *Phys. Rev. B* **52**, 1640 (1995).

- [43] Fatttebert, J.-L. & Gygi, F. Linear-scaling first-principles molecular dynamics with plane-waves accuracy. *Phys. Rev. B* **73**, 115124 (2006).
- [44] Osei-Kuffuor, D. & Fatttebert, J.-L. Accurate and scalable  $\mathcal{O}(n)$  algorithm for first-principles molecular-dynamics computations on large parallel computers. *Phys. Rev. Lett.* **112**, 046401 (2014).
- [45] Tsuchida, E. Augmented orbital minimization method for linear scaling electronic structure calculations. *J. Phys. Soc. Jpn.* **76**, 034708 (2007).
- [46] Head-Gordon, M., Maslen, P. E. & White, C. A. A tensor formulation of many-electron theory in a nonorthogonal single-particle basis. *J. Chem. Phys.* **108**, 616–625 (1998).
- [47] Galli, G. & Parrinello, M. Large scale electronic structure calculations. *Phys. Rev. Lett.* **69**, 3547 (1992).
- [48] VandeVondele, J. & Hutter, J. An efficient orbital transformation method for electronic structure calculations. *J. Chem. Phys.* **118**, 4365–4369 (2003).
- [49] Van Voorhis, T. & Head-Gordon, M. A geometric approach to direct minimization. *Mol. Phys.* **100**, 1713–1721 (2002).
- [50] VandeVondele, J. & Hutter, J. Gaussian basis sets for accurate calculations on molecular systems in gas and condensed phases. *J. Chem. Phys.* **127**, 114105 (2007).
- [51] Cp2k open source molecular dynamics. URL <http://www.cp2k.org>.
- [52] VandeVondele, J. *et al.* Quickstep: Fast and accurate density functional calculations using a mixed gaussian and plane waves approach. *Comput. Phys. Commun.* **167**, 103–128 (2005).
- [53] Borštnik, U., VandeVondele, J., Weber, V. & Hutter, J. Sparse matrix multiplication: The distributed block-compressed sparse row library. *Parallel Comput.* **40**, 47–58 (2014).
- [54] Hotelling, H. Some new methods in matrix calculation. *Ann. Math. Stat.* **14**, 1–34 (1943).
- [55] Feynman, R. P. Forces in molecules. *Phys. Rev.* **56**, 340 (1939).
- [56] Scheiber, H., Shi, Y. & Khaliullin, R. Z. Communication: Compact orbitals enable low-cost linear-scaling ab initio molecular dynamics for weakly-interacting systems. *J. Chem. Phys.* **148**, 231103 (2018).
- [57] Bussi, G., Donadio, D. & Parrinello, M. Canonical sampling through velocity rescaling. *J. Chem. Phys.* **126**, 014101 (2007).
- [58] Weber, V., VandeVondele, J., Hutter, J. & Niklasson, A. M. Direct energy functional minimization under orthogonality constraints. *J. Chem. Phys.* **128**, 084113 (2008).