



HAL
open science

Daubechies Wavelets for Linear Scaling Density Functional Theory

Stephan Mohr, Laura E. Ratcliff, Paul Boulanger, Luigi Genovese, Damien Caliste, Thierry Deutsch, Stefan Goedecker

► **To cite this version:**

Stephan Mohr, Laura E. Ratcliff, Paul Boulanger, Luigi Genovese, Damien Caliste, et al.. Daubechies Wavelets for Linear Scaling Density Functional Theory. *Journal of Chemical Physics*, 2014, 140 (20), pp.204110. 10.1063/1.4871876 . hal-01334186

HAL Id: hal-01334186

<https://hal.science/hal-01334186>

Submitted on 11 Jul 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Daubechies Wavelets for Linear Scaling Density Functional Theory

Stephan Mohr,^{1,2} Laura E. Ratcliff,² Paul Boulanger,^{2,3} Luigi Genovese,² Damien Caliste,² Thierry Deutsch,² and Stefan Goedecker¹

¹*Institut für Physik, Universität Basel, Klingelbergstr. 82, 4056 Basel, Switzerland*

²*Univ. Grenoble Alpes, CEA, INAC-SP2M, F-38000 Grenoble, France*

³*Institut Néel, CNRS and Université Joseph Fourier, B.P. 166, 38042 Grenoble Cedex 09, France*

(Dated: June 4, 2014)

We demonstrate that Daubechies wavelets can be used to construct a minimal set of optimized localized adaptively-contracted basis functions in which the Kohn-Sham orbitals can be represented with an arbitrarily high, controllable precision. Ground state energies and the forces acting on the ions can be calculated in this basis with the same accuracy as if they were calculated directly in a Daubechies wavelets basis, provided that the amplitude of these adaptively-contracted basis functions is sufficiently small on the surface of the localization region, which is guaranteed by the optimization procedure described in this work. This approach reduces the computational costs of DFT calculations, and can be combined with sparse matrix algebra to obtain linear scaling with respect to the number of electrons in the system. Calculations on systems of 10,000 atoms or more thus become feasible in a systematic basis set with moderate computational resources. Further computational savings can be achieved by exploiting the similarity of the adaptively-contracted basis functions for closely related environments, e.g. in geometry optimizations or combined calculations of neutral and charged systems.

I. INTRODUCTION

The Kohn-Sham (KS) formalism of density functional theory (DFT)^{1,2} is one of the most popular electronic structure methods due to its good balance between accuracy and speed. Thanks to the development of new approximations to the exchange correlation functional, this approach now allows many quantities (bond lengths, vibration frequencies, elastic constants, etc.) to be calculated with errors of less than a few percent, which is sufficient for many applications in solid state physics, chemistry, materials science, biology, geology and many other fields. Although the KS approach has some shortcomings – e.g. its inability to accurately describe the HOMO-LUMO separation or many-body (e.g. excitonic) effects, thus reducing its predictive power in the field of optics – it has become the standard for the quantum simulation of matter and also provides a well defined starting point for more accurate methods, such as the *GW* approximation³.

Despite the efforts put forth to increase the efficiency of DFT calculations and the increasing computing power of modern supercomputers, the applicability for standard calculations is limited to systems containing about a thousand atoms, which is small compared to the size of systems of interest in nanoscience. The reason for this is that standard electronic structure programs using systematic basis sets such as plane waves⁴⁻⁶, finite elements⁷ or wavelets⁸ need a number of operations that scales as the number of orbitals, N_{orb} , squared times the number of basis functions, N_{basis} , used to represent them. Since both the number of orbitals and the number of basis functions scale as the number of atoms, the overall cost scales as $\mathcal{O}(N_{orb}^2 N_{basis}) = \mathcal{O}(N_{atom}^3)$. Electronic structure programs that use Gaussians⁹ or atomic orbitals¹⁰ require in a standard implementation a matrix diagonal-

ization which scales as $\mathcal{O}(N_{basis}^3)$.

To circumvent this problem, one can exploit Kohn's nearsightedness principle¹¹⁻¹³, which states that, for systems with a finite gap or for metals at finite temperature, all physical quantities are determined by the local environment. This is a consequence of the exponentially fast decay of the density matrix¹³⁻¹⁹. Therefore, it is theoretically possible to express the KS wavefunctions of a given system in terms of a minimal, localized basis set. In order to get highly accurate results while still keeping the size of the basis relatively small, such a basis has to depend on the local chemical environment. If this basis set were known or could be approximated beforehand, it would lead to a computationally cheap tight-binding like approach^{20,21}. Of course, in practice it is not possible to determine this optimal localized basis set beforehand; instead it has to be built up iteratively during the calculation. This would result in $\mathcal{O}(N_{orb}^3)$ scaling, which is still equivalent to $\mathcal{O}(N_{atom}^3)$, but with a much smaller prefactor than systematic approaches (e.g. plane waves) where the number of basis functions is far greater than the number of orbitals ($N_{basis} \gg N_{orb}$).

However, the use of a strictly localized basis offers yet another possibility. As has been demonstrated during the past twenty years^{22,23}, it is possible to truncate the density matrix and thus transform it into a sparse form by neglecting elements either when they are below a certain threshold, or when they correspond to localized orbitals which are too distant from each other. This reduces the complexity of the algorithm to $\mathcal{O}(N_{orb}) = \mathcal{O}(N_{atom})$ and leads to so-called linear scaling (LS) DFT methods. Even though the exponential decay of the density matrix is as well present for metals at finite temperature, we will – as most $\mathcal{O}(N_{atom})$ approaches – focus on the simpler case of insulators. Methods of this type have been implemented in numerous codes such as ONETEP²⁴,

Conquest²⁵, CP2K²⁶ and SIESTA²⁷. Note, however, that the extent of the truncation impacts the accuracy due to the imposition of an additional constraint on the system, and is therefore left as a freely selectable parameter for the user. This additional constraint also comes at the cost of extra computational steps, so that the prefactor is greater than for standard DFT codes, even for a single iteration in the self-consistency cycle. Furthermore, there can be problems with ill-conditioning when using strictly localized basis sets, which further increases the prefactor. The combination of these two problems means that for small systems the total calculation time is actually greater when one imposes locality, but thanks to the better scaling, there is a crossover point where the new algorithms become more efficient.

Our minimal set of localized adaptively-contracted basis functions, called support functions in the following, is obtained by an environment dependent optimization where the support functions are represented in terms of a fixed underlying wavelet basis set. The term *adaptively-contracted* should not be confused with the terminology of contracted basis functions often used in quantum chemistry, it is simply used to emphasize that there are two levels of basis functions, namely the underlying wavelets basis and the support functions which are built out of them. Because of the environment dependency, the size of this basis set is however for a given accuracy much smaller than the size of typical contracted Gaussian basis sets and we refer to this basis set therefore also as a minimal basis set.

The choice of the underlying basis set is one of the most important aspects impacting the accuracy and efficiency of a linear scaling DFT code. Ideally, it should feature compact support while still being orthogonal, thus allowing for a systematic convergence – properties which are all offered by Daubechies wavelets basis sets²⁸. Furthermore, wavelets have built in multiresolution properties, enabling an adaptive mesh with finer sampling close to the atoms where the most significant part of the orbitals is located; this can be particularly beneficial for inhomogeneous systems. Wavelets also have the distinct advantage that calculations can be performed with all the standard boundary conditions – free, wire, surface or periodic. This also means we can perform calculations on charged and polarized systems using free boundary conditions without the need for a compensating background charge. It is therefore evident that the combination of the above features makes wavelets ideal for a LSDFT code.

This paper is organized as follows. We first give an overview of the method, focussing in particular on the imposition of the localization constraint in Daubechies wavelets. We then discuss the details, highlighting the novel features, following which we consider the calculation of atomic forces. For this latter point, we demonstrate the remarkable result that, thanks to the compact support of Daubechies wavelets, the contribution of the Pulay-like forces, arising from the introduction of the localization regions, can be safely neglected in a typical

calculation. We then present results for a number of systems, illustrating the accuracy of the method for ground state energies and atomic forces. We also demonstrate the improved scaling compared with standard BigDFT, showing that we are able to achieve linear scaling. Finally, we highlight two cases where the minimal basis functions can be reused, resulting in further significant computational savings.

II. MINIMAL ADAPTIVELY-CONTRACTED BASIS

A. Kohn-Sham formalism in a minimal basis set

The standard approach for performing Kohn-Sham DFT calculations is to calculate the Kohn-Sham orbitals $|\Psi_i\rangle$ which satisfy the equation

$$\mathcal{H}_{KS} |\Psi_i\rangle = \varepsilon_i |\Psi_i\rangle, \quad (1)$$

with

$$\mathcal{H}_{KS} = -\frac{1}{2}\nabla^2 + \mathcal{V}_{KS}[\rho] + \mathcal{V}_{PSP}, \quad (2)$$

where $\mathcal{V}_{KS}[\rho]$ contains the Hartree potential – solution to the Poisson equation – and the exchange-correlation potential, while \mathcal{V}_{PSP} contains the potential arising from the pseudopotential and the external potential created by the ions. In the case of BigDFT, these are norm-conserving GTH-HGH²⁹ pseudopotentials and their Krack variants³⁰, possibly with a nonlinear core correction³¹. It is worth noting that the use of pseudopotentials does not only decrease the complexity of the calculation by reducing the number of orbitals and avoiding the need of very high resolution around the nuclei, but also offers the possibility of easily including relativistic effects. Furthermore the calculation of the Hartree and exchange-correlation potentials are done in the same way as in the original version of BigDFT⁸ and are thus not subject to any approximations. For the exchange-correlation part we restrict ourselves to local functionals.

In our approach the KS orbitals are in turn expressed as a linear combination of support functions $|\phi_\alpha\rangle$:

$$|\Psi_i(\mathbf{r})\rangle = \sum_{\alpha} c_i^{\alpha} |\phi_{\alpha}(\mathbf{r})\rangle. \quad (3)$$

The density – which can be obtained from the one-electron orbitals via $\rho(\mathbf{r}) = \sum_i f_i |\Psi_i(\mathbf{r})|^2$, where f_i is the occupation number of orbital i – is given by

$$\rho(\mathbf{r}) = \sum_{\alpha,\beta} \phi_{\alpha}^*(\mathbf{r}) K^{\alpha\beta} \phi_{\beta}(\mathbf{r}), \quad (4)$$

where $K^{\alpha\beta} = \sum_i f_i c_i^{*\alpha} c_i^{\beta}$ is the density kernel. The latter is related to the density matrix formulation of Hernández

and Gillan³², since – as follows from Eq. (3) –

$$\begin{aligned} F(\mathbf{r}, \mathbf{r}') &= \sum_i f_i |\Psi_i(\mathbf{r})\rangle \langle \Psi_i(\mathbf{r}')| \\ &= \sum_{\alpha, \beta} |\phi_\alpha(\mathbf{r})\rangle K^{\alpha\beta} \langle \phi_\beta(\mathbf{r}')|. \end{aligned} \quad (5)$$

Thus the density kernel is the representation of the density matrix in the support function basis. We choose to have real support functions and thus from now on we will neglect the complex notation for this quantity.

The density matrix decays exponentially with respect to the distance $|\mathbf{r} - \mathbf{r}'|$ for systems with a finite gap or for metals at finite temperature^{13–19}. In these cases it can therefore be represented by strictly localized basis functions. A natural and exact choice for these would be the maximally localized Wannier functions which have the same exponential decay³³. Of course, these Wannier functions are not known beforehand. Therefore, in our case, the adaptively-contracted basis functions are constructed *in situ* during the self-consistency cycle and are expected to reach a quality similar to that of the exact Wannier functions.

In the formalism we have presented so far, the KS orbitals have to be optimized by minimizing the total energy with respect to the support functions and density kernel. For a self-consistent calculation this is equivalent to minimizing the band structure energy, i.e.

$$E_{BS} = \sum_{\alpha, \beta} K^{\alpha\beta} H_{\alpha\beta}, \quad (6)$$

subject to the orthonormality condition of the KS orbitals,

$$\langle \Psi_i | \Psi_j \rangle = \sum_{\alpha, \beta} c_i^{\alpha*} S_{\alpha\beta} c_j^\beta = \delta_{ij}, \quad (7)$$

where $H_{\alpha\beta} = \langle \phi_\alpha | \mathcal{H}_{KS} | \phi_\beta \rangle$ and $S_{\alpha\beta} = \langle \phi_\alpha | \phi_\beta \rangle$ are the Hamiltonian and overlap matrices of the support functions, respectively. For systems with all occupation numbers being either zero or one, this is equivalent to imposing the idempotency condition on the density kernel $K^{\alpha\beta}$,

$$\sum_{\gamma, \delta} K^{\alpha\gamma} S_{\gamma\delta} K^{\delta\beta} = K^{\alpha\beta}, \quad (8)$$

which can be achieved using the McWeeny purification scheme³⁴ or directly imposing the orthogonality constraint on the coefficients c_i^α

The algorithm therefore consists of two key components: support function and density kernel optimization. The workflow is illustrated in Fig. 1; it consists of a flexible double loop structure, with the outer loop controlling the overall convergence, and two inner loops which optimize the support functions and density kernel, respectively. The first of these inner loops is done non-self-consistently (i.e. with a fixed potential), whereas the second one is done self-consistently.




FIG. 1. Structure of the minimal basis approach: for the basis optimization loop the hybrid scheme can be used instead of trace or energy minimization, and for the kernel optimization loop either direct minimization or the Fermi operator expansion method can be used in place of diagonalization; see Sec IV.




FIG. 2. Least asymmetric Daubechies wavelet family of order $2m = 16$; both the scaling function $\phi(x)$ and wavelet $\psi(x)$ differ from zero only within the interval $[1 - m, m]$.

B. Daubechies wavelets in BigDFT

BigDFT⁸ uses the orthogonal least asymmetric Daubechies²⁸ family of order $2m = 16$, illustrated in Fig. 2. These functions have a compact support and are smooth, which means that they are also localized in Fourier space. This wavelet family is able to exactly represent polynomials up to 8th order. Such a basis is therefore an optimal choice given that we desire at the same time locality and interpolating power. An exhaustive presentation of the use of wavelets in numerical simulations can be found in Ref. 35.

A wavelet basis set is generated by the integer translates of the scaling functions and wavelets, with arguments measured in units of the grid spacing h . In three dimensions, a wavelet basis set can easily be obtained as the tensor product of one-dimensional basis functions, combining wavelets and scaling functions along each coordinate of the Cartesian grid (see *e.g.* Ref. 8).

In a simulation domain, we have three categories of grid points: those which are closest to the atoms (“fine region”) carry one (three-dimensional) scaling function and seven (three-dimensional) wavelets; those which are further away from the atoms (“coarse region”) carry only one scaling function, corresponding to a resolution which is half that of the fine region; and those which are even further away (“empty region”) carry neither scaling functions nor wavelets. The fine region is typically the region where chemical bonding takes place, whereas the coarse region covers the region where the tails of the wavefunctions decay smoothly to zero. We therefore have two resolution levels whilst maintaining a regularly spaced grid in the entire simulation box.

A support function $\phi_\alpha(\mathbf{r})$ can be expanded in this wavelet basis as follows:

$$\begin{aligned} \phi(\mathbf{r}) = & \sum_{i_1, i_2, i_3} s_{i_1, i_2, i_3} \varphi_{i_1, i_2, i_3}(\mathbf{r}) \\ & + \sum_{j_1, j_2, j_3} \sum_{l=1}^7 d_{j_1, j_2, j_3}^{(l)} \psi_{j_1, j_2, j_3}^{(l)}(\mathbf{r}), \end{aligned} \quad (9)$$

where $\varphi_{i_1, i_2, i_3}(\mathbf{r}) = \varphi(x - i_1)\varphi(y - i_2)\varphi(z - i_3)$ is the tensor product of three one-dimensional scaling functions centered at the grid point (i_1, i_2, i_3) , and $\psi_{j_1, j_2, j_3}^{(l)}(\mathbf{r})$ are the seven tensor products containing at least one one-dimensional wavelet centered on the grid point (j_1, j_2, j_3) . The sums over i_1, i_2, i_3 (j_1, j_2, j_3) run over all grid points where scaling functions (wavelets) are centered, i.e. all the points of the coarse (fine) grid. The overall simulation box is chosen to be rectangular and is identical to the one in the standard version of BigDFT; for simplicity the origin is chosen such that there are only positive grid coordinates, i.e. in a corner of the simulation domain.

To determine these regions of different resolution, we construct two spheres around each atom a ; a small one with radius $R_a^f = \lambda^f \cdot r_a^f$ and a large one with radius $R_a^c = \lambda^c \cdot r_a^c$ ($R_a^c > R_a^f$). The values of r_a^f and r_a^c are characteristic for each atom type and are related to the covalent and van der Waals radii, whereas λ^f and λ^c can be specified by the user in order to control the accuracy of the calculation. The fine (coarse) region is then given by the union of all the small (large) spheres, as shown in Fig. 3.

Hence in BigDFT the basis set is controlled by these three user specified parameters. By reducing h and/or increasing λ^c and λ^f the computational degrees of freedom are incremented, leading to a systematic convergence of the energy.

III. LOCALIZATION REGIONS

Thanks to the nearsightedness principle it is possible to define a basis of strictly localized support functions such that the KS orbitals given in terms of this adaptively-contracted basis are *exactly* equivalent to the expression




FIG. 3. A two level adaptive grid for an alkane: the high resolution grid points are shown with bold points while the low resolution grid points are shown with smaller points. Also visible are three localization regions (red, blue and green) with radii of 3.7 Å centered on different atoms in which certain support functions will reside. The coarse grid points shown in yellow do not belong to any of the three localization regions.

based solely on the underlying Daubechies basis. However, as presented so far, the support functions $\phi_\alpha(\mathbf{r})$ of Eq. (9) are expanded over the entire simulation domain. We want them to be strictly localized while still containing various resolution levels, as illustrated by Fig. 3, and so we set to zero all scaling function and wavelet coefficients which lie outside a sphere of radius R_{cut} around the point \mathbf{R}_α on which the support function is centered. In general, these centers \mathbf{R}_α could be anywhere, but we choose them to be centered on an atom a and we thus assume from now on that $\mathbf{R}_\alpha = \mathbf{R}_a$. Consequently we define a localization projector $\mathcal{L}^{(\alpha)}$, which is written in the Daubechies basis space as

$$\mathcal{L}_{i_1, i_2, i_3; j_1, j_2, j_3}^{(\alpha)} = \delta_{i_1 j_1} \delta_{i_2 j_2} \delta_{i_3 j_3} \theta(R_{cut} - |\mathbf{R}_{(i_1, i_2, i_3)} - \mathbf{R}_\alpha|), \quad (10)$$

where θ is the Heaviside function. We use this projector to constrain the function $|\phi_\alpha\rangle$ to be localized throughout the calculation, i.e.

$$|\phi_\alpha\rangle = \mathcal{L}^{(\alpha)} |\phi_\alpha\rangle. \quad (11)$$

Clearly, if $|\phi_\alpha\rangle$ is localized around \mathbf{R}_α and R_{cut} is large enough, $\mathcal{L}^{(\alpha)}$ leaves $|\phi_\alpha\rangle$ unchanged and no approximation is introduced to the KS equation.

It is important to note that the localization constraint of Eq. (11) determines the expression of $\frac{d\phi_\alpha}{d\mathbf{R}_\beta}$ outside the localization region of ϕ_α . Indeed, as the Daubechies basis set is independent of \mathbf{R}_α , differentiating Eq. (11) with respect to \mathbf{R}_β leads to

$$(1 - \mathcal{L}^{(\alpha)}) \left| \frac{d\phi_\alpha}{d\mathbf{R}_\beta} \right\rangle = \delta_{\alpha\beta} \frac{\partial \mathcal{L}^{(\alpha)}}{\partial \mathbf{R}_\alpha} |\phi_\alpha\rangle. \quad (12)$$

This result will be used in Appendix C2 to demonstrate that the Pulay-like forces are negligible for a typical calculation with our approach.

A. Imposing the localization constraint

In what follows, we demonstrate that choosing the support functions to be orthogonal allows for a more straightforward application of the localization constraint. Due to the orthonormality of the KS orbitals we cannot directly minimize the band structure energy (Eq. (6)) with respect to the support functions, rather we have to minimize the following functional:

$$\Omega = \sum_{\alpha,\beta} K^{\alpha\beta} H_{\alpha\beta} - \sum_{i,j} \sum_{\alpha,\beta} \Lambda_{ij} \left(c_i^{\alpha*} c_j^\beta S_{\alpha\beta} - \delta_{ij} \right), \quad (13)$$

with the Lagrange multiplier coefficients Λ_{ij} determined by the relation

$$\sum_{i,j} c_i^{\alpha*} c_j^\beta \Lambda_{ij} = \sum_{\rho,\sigma} K^{\alpha\rho} H_{\rho\sigma} (S^{-1})^{\sigma\beta}, \quad (14)$$

where $(S^{-1})^{\alpha\beta}$ is the inverse overlap matrix. The gradient $|g^\alpha\rangle = \left| \frac{\delta\Omega}{\delta\langle\phi_\alpha|} \right\rangle$ is therefore

$$|g^\alpha\rangle = \sum_{\beta} K^{\alpha\beta} \mathcal{H}_{KS} |\phi_\beta\rangle - \sum_{\beta,\rho,\sigma} K^{\alpha\rho} H_{\sigma\rho} (S^{-1})^{\sigma\beta} |\phi_\beta\rangle \quad (15)$$

However, we wish to impose the localization condition $|\phi_\alpha\rangle = \mathcal{L}^{(\alpha)} |\phi_\alpha\rangle$ on the support functions and therefore the functional to be minimized becomes

$$\Omega' = \Omega - \sum_{\alpha} \langle \phi_\alpha | 1 - \mathcal{L}^{(\alpha)} | \ell^\alpha \rangle, \quad (16)$$

where the components of the vector $|\ell^\alpha\rangle$ are the Lagrange multipliers of this locality constraint. The gradient for Ω' , $\left| \frac{\delta\Omega'}{\delta\langle\phi_\alpha|} \right\rangle$, can therefore be written as

$$|g'^\alpha\rangle = |g^\alpha\rangle - (1 - \mathcal{L}^{(\alpha)}) |\ell^\alpha\rangle. \quad (17)$$

Using the stationarity condition $0 = |g'^\alpha\rangle$ and combining with the fact that $1 - \mathcal{L}^{(\alpha)}$ is a projection operator, i.e. $(1 - \mathcal{L}^{(\alpha)})^2 = 1 - \mathcal{L}^{(\alpha)}$, we have

$$(1 - \mathcal{L}^{(\alpha)}) |\ell^\alpha\rangle = (1 - \mathcal{L}^{(\alpha)}) |g^\alpha\rangle. \quad (18)$$

Therefore, using Eq. (17),

$$|g'^\alpha\rangle = \mathcal{L}^{(\alpha)} |g^\alpha\rangle. \quad (19)$$

i.e. the gradient is explicitly localized. This yields the following result for the gradient:

$$|g'^\alpha\rangle = \sum_{\beta,\rho} K^{\alpha\rho} (S^{1/2})_{\rho\beta} \left[\mathcal{L}^{(\alpha)} \mathcal{H}_{KS} \left| \tilde{\phi}_\beta \right\rangle - \sum_{\sigma} \left\langle \tilde{\phi}_\sigma \left| \mathcal{H}_{KS} \left| \tilde{\phi}_\rho \right\rangle \mathcal{L}^{(\alpha)} \left| \tilde{\phi}_\sigma \right\rangle \right] \right]. \quad (20)$$

Here the localized gradient is expressed in terms of the orthogonalized support functions $\left| \tilde{\phi}_\alpha \right\rangle = \sum_{\beta} (S^{-1/2})_{\alpha\beta} |\phi_\beta\rangle$. Requiring the support functions to be orthogonal i.e. $S_{\alpha\beta} = \delta_{\alpha\beta}$ therefore further simplifies the evaluation of the gradient as it no longer becomes necessary to calculate \mathbf{S}^{-1} or $\mathbf{S}^{1/2}$. Moreover, it avoids the need for distinguishing between covariant and contravariant indices³⁶.

B. Localization of the Hamiltonian application

As shown in Ref. 8, the Hamiltonian operator in a Daubechies wavelets basis set is defined by a set of convolution operations, combined with the application of non-local pseudopotential projectors. The nature of these operations is such that $\mathcal{H}_{KS} |\phi_\beta\rangle$ will have a greater extent than $|\phi_\beta\rangle$. We therefore define a second localization operator, $\mathcal{L}'^{(\beta)}$, with a corresponding cutoff radius R'_{cut} , such that R'_{cut} is equal to R_{cut} plus half of the convolution filter length times the grid spacing, which in our case corresponds to an additional eight grid points. When applying the Hamiltonian, we impose

$$\mathcal{H}_{KS} |\phi_\beta\rangle = \mathcal{L}'^{(\beta)} \mathcal{H}_{KS} \mathcal{L}^{(\beta)} |\phi_\beta\rangle. \quad (21)$$

For both the convolution operations and the non-local pseudopotential applications, this procedure guarantees that the Hamiltonian application is exact within the localization region of $\mathcal{L}^{(\beta)}$. However, the values of $\mathcal{H}_{KS} \mathcal{L}^{(\beta)} |\phi_\beta\rangle$ are approximated outside this region due to the semilocal nature of the convolutions and the pseudopotential projectors. This impacts the evaluation of the Hamiltonian matrix $H_{\alpha\beta}$ for all elements whose localization regions do not coincide, and thus also affects the gradient $|g^\alpha\rangle$. Nonetheless, we have verified that further enlargement of R'_{cut} is not needed as it has negligible impact on the accuracy, while adding additional overheads, see Sec.VI.

Apart from these technical details, most of the basic operations are identical to their implementation in the standard BigDFT code⁸ and are therefore not repeated here. The only difference is that these operations are now done only in the localization regions (corresponding to either $\mathcal{L}^{(\beta)}$ or $\mathcal{L}'^{(\beta)}$) and not in the entire computational volume.

IV. SELF-CONSISTENT CYCLE

A. Support function optimization

As an initial guess for the support functions we use atomic orbitals, which are generated by solving the atomic Schrödinger equation and therefore possess long tails which need to be truncated at the borders of the localization regions. If the values at the borders are not negligible, the resulting kink will cause the kinetic energy to become very large due to the definition of the Laplacian operator in a wavelet basis set, and so to assure stability during the optimization procedure, the localization regions would need to be further enlarged. To overcome this problem, even for small localization regions, it is advantageous to decrease the extent of the atomic orbitals before the initial truncation by adding a confining quartic potential centered on each atom, $a_\alpha(\mathbf{r} - \mathbf{R}_\alpha)^4$, to the atomic Schrödinger equation.

For the first few iterations of the outer loop (Fig. 1) we maintain the quartic confining potential of the atomic

input guess. This implies that the total Hamiltonian becomes dependent on the support function, $\mathcal{H}_\alpha = \mathcal{H}_{KS} + a_\alpha(\mathbf{r} - \mathbf{R}_\alpha)^4$, and we can no longer minimize the band structure energy (Eq. (6)) to obtain the support functions. Instead, we choose to minimize the functional

$$\min_{\phi_\alpha} \sum_{\alpha} \langle \phi_\alpha | \mathcal{H}_\alpha | \phi_\alpha \rangle, \quad (22)$$

while applying both orthogonality and localization constraints on the support functions, as discussed in Section III A.

Apart from the improved localization, the use of the confining potential has yet another advantage. The band structure energy (Eq. (6)) is invariant under unitary transformations among the support functions if there are no localization constraints. This corresponds to some zero eigenvalues in the Hessian characterizing the optimization of the support functions. The introduction of a localization constraint violates this invariance and leads to small but non-zero eigenvalues. The condition number, defined as the ratio of the largest and smallest (nonzero) eigenvalue of the Hessian, can thus become very large, potentially turning the optimization into a strongly ill-conditioned problem. On the other hand, if the unitary invariance is heavily violated in Eq. (22) by the introduction of a strong localization potential, the small eigenvalues grow and the condition number improves as a consequence.

After a few iterations of the outer loop, the support functions are sufficiently localized to continue the optimization without a confining potential, i.e. by minimizing the band structure energy. This procedure will lead to highly accurate support functions while still preserving locality. As an alternative it is also possible to define a so-called ‘‘hybrid mode’’ which combines the two categories of support function optimization and thus provides a smoother transition between the two. In this case the target function is given by

$$\Omega^{hy} = \sum_{\alpha} K^{\alpha\alpha} \langle \phi_\alpha | \mathcal{H}_\alpha | \phi_\alpha \rangle + \sum_{\beta \neq \alpha} K^{\alpha\beta} \langle \phi_\alpha | \mathcal{H}_{KS} | \phi_\beta \rangle. \quad (23)$$

In the beginning a strong confinement is used, making this expression similar to the functional of Eq. (22); however the confining potential is reduced throughout the calculation so that towards the end the strength of the confinement is negligible and Eq. (23) reverts to the full energy expression. A prescription for reducing the confinement is presented in Appendix A.

1. Orthogonalization

The Lagrange multiplier formalism conserves the orthonormality of the adaptively-contracted basis only to first order. An additional explicit orthogonalization has to be performed after each update of the adaptively-contracted basis to restore exact orthogonality. This is

done using the Löwdin procedure. The calculation of $\mathbf{S}^{-1/2}$, which is required in this context, can pose a bottleneck. However, as our basis functions are close to orthonormality, the exact calculation can safely be replaced by a first order Taylor approximation. Numerical tests have shown that the error of this approximation is of the same order of magnitude as the inevitable deviation from exact orthonormality which is inherent to our set of support functions due to the strictly enforced locality; consequently the slight non-orthonormality of the adaptively-contracted basis does not significantly increase by the use of the Taylor approximation.

It is important to note that the support functions will only be nearly orthogonal rather than exactly orthogonal, as exact orthogonality is in general not possible for functions exhibiting compact support in a discretized space. This near-orthogonality is in contrast with most other minimal basis implementations which use fully non-orthogonal support functions^{24,25}. The asymptotic decay behavior of the orthogonal and non-orthogonal support functions is identical. However the prefactor differs and leads to a better localization of the non-orthogonal functions³⁷. However, in practice, we have found that the introduction of the orthogonality constraint does not significantly increase the size required for the localization regions, provided that a sufficiently strong confining potential is applied to localize the support functions at the start of the calculation.

In order to counteract the small deviations from idempotency caused by the changing overlap matrix, we purify the density kernel during the support function optimization either by directly orthonormalizing the expansion coefficients of the KS orbitals or by using the McWeeny purification transformation.

2. Gradient and preconditioning

The optimization is done via a direct minimization scheme or with direct inversion of the iterative subspace (DIIS)³⁸, both combined with an efficient preconditioning scheme. The derivation of the gradient of the target function Ω with respect to the support functions involves some subtleties for the trace and hybrid modes since in these cases the Hamiltonian depends explicitly on the support function, leading to an asymmetry of the Lagrange multiplier matrix $\Lambda_{\alpha\beta} = \left\langle \phi_\alpha \left| \sum_{\gamma} S_{\beta\gamma} \frac{\delta\Omega}{\delta\langle\phi_\gamma|} \right. \right\rangle$. In order to correctly derive the gradient expression we follow the same guidelines as Ref. 39; assuming nearly orthogonal orbitals the final result is given by

$$|g_\alpha\rangle = \sum_{\beta} S_{\alpha\beta} \frac{\delta\Omega}{\delta\langle\phi_\beta|} - \frac{1}{2} \sum_{\beta} [\mathbf{S}^{-1} (\mathbf{\Lambda} + \mathbf{\Lambda}^T)]_{\alpha\beta} |\phi_\beta\rangle \quad (24)$$

This is a generalization of the ordinary expression and thus also valid if the Hamiltonian does not explicitly depend on the support function, i.e. for the energy mode.

As discussed in Section III A the gradient is suitably localized once derived, i.e. $|g_\alpha\rangle \leftarrow \mathcal{L}^{(\alpha)}|g_\alpha\rangle$.

To precondition the gradient $|g_\alpha\rangle$ we use the standard kinetic preconditioning scheme⁸. To ensure that the preconditioning does not negatively impact the localization of the gradient, we have found that it is important to add an extra term to account for the confining potential if present. In this case, the expression becomes

$$\left(-\frac{1}{2}\nabla^2 + a_\alpha(\mathbf{r} - \mathbf{R}_\alpha)^4 - \varepsilon_\alpha\right)|g_\alpha^{prec}\rangle = -|g_\alpha\rangle, \quad (25)$$

where ε_α is an approximate value of $\Lambda_{\alpha\alpha}$. This also has the effect of improving the convergence. The inclusion of the confining potential adds only a small overhead as it can be evaluated via convolutions in the same manner as the kinetic energy. Furthermore, the preconditioning equations do not need to be solved with high accuracy, only approximately.

B. Density kernel optimization

For the optimization of the density kernel we have implemented three schemes: diagonalization, direct minimization and the Fermi operator expansion method (FOE)^{40,41}. Once the kernel has been updated, we recalculate the charge density via Eq (4); the new density is then used to update the potential, with an optional step wherein the density is mixed with the previous one in order to improve convergence. This procedure is repeated until the kernel is converged; in practice, we consider this convergence to have been reached once the mean difference of the density of two consecutive iterations is below a given threshold, i.e. $\Delta\rho < c$.

The direct diagonalization method consists of finding the solution of the generalized eigenproblem for a given Hamiltonian and overlap matrix. Its implementation therefore relies straightforwardly on linear algebra solvers and will not be detailed here.

In the direct minimization approach the band-structure energy is minimized subject to the orthogonality of the support functions. To this end, we express the gradient of the Kohn-Sham orbitals, $|g_i\rangle = \mathcal{H}_{KS}|\Psi_i\rangle - \sum_j \Lambda_{ij}|\Psi_j\rangle$, in terms of the support functions, i.e. $|g_i\rangle = \sum_\alpha d_i^\alpha|\phi_\alpha\rangle$. The d_i^α are obtained by solving

$$\begin{aligned} \sum_\alpha S_{\beta\alpha}d_i^\alpha &= \sum_\alpha H_{\beta\alpha}c_i^\alpha - \sum_j \sum_\alpha S_{\beta\alpha}c_j^\alpha \Lambda_{ji}, \\ \Lambda_{ji} &= \sum_{\gamma,\delta} c_j^{\gamma*} c_i^\delta H_{\gamma\delta}. \end{aligned} \quad (26)$$

The coefficients are optimized using this gradient via steepest descents or DIIS. Once the gradient has converged to the required threshold, the density kernel is calculated from the coefficients and occupancies.

In the Fermi operator expansion method, the density matrix may be defined as a function of the Hamiltonian as $F = f(\mathcal{H})$, where f is the Fermi function. In terms of

the support functions, this corresponds to an expression for the density kernel in terms of the Hamiltonian matrix, i.e. $\mathbf{K} = f(\mathbf{H})$. The central idea of the FOE^{23,40,41} is to find an expression for $f(\mathbf{H})$ which can be efficiently evaluated numerically. One particularly simple possibility is a polynomial expansion; for numerical stability we use Chebyshev polynomials⁴². As will be shown in detail in Appendix B, the density kernel can be constructed using only matrix vector multiplications thanks to the recursion formulae for the Chebyshev matrix polynomials.

1. Suitability of the methods

All three methods for calculating the density kernel (direct diagonalization, direct minimization and FOE) yield the same final result, thus the main differences lie in their performance, where one of the most important points is the performance of the linear algebra. Due to the localized nature of the support functions, the overlap, Hamiltonian and density kernel matrices are in general sparse, with the level and pattern of sparsity depending on the localization radii of the support functions and the dimensionality of the system in question. We can take advantage of this sparsity by storing and using these matrices in compressed form and indeed this is necessary to achieve a fully linear scaling algorithm.

For diagonalization, exploiting the sparsity is very hard due to the lack of efficient parallel solvers for sparse matrices. The method therefore performs badly for large systems due to its cubic scaling, but thanks to the small prefactor it can still be useful for smaller systems consisting of a few hundred atoms.

For direct minimization the situation is better, since both the solution of the linear system of Eq. (26) and the orthonormalization can be approximated using Taylor expansions for \mathbf{S}^{-1} and $\mathbf{S}^{-1/2}$, respectively. It can also be easily parallelized, so that the cubic scaling terms only become problematic for systems containing more than a few thousand atoms, as demonstrated in section VII (see Fig. 7). Indeed, for moderate system sizes, the extra overhead associated with the manipulation of sparse matrices makes dense matrix algebra cheaper, and so many physically interesting systems are already considerably accelerated.

Even though the direct minimization method does not scale linearly in its current implementation, there are some situations where its use is advantageous. For example, the unoccupied states are generally not well represented by the adaptively-contracted basis set following the optimization procedure, and for cases where we have only the density kernel and not the coefficients, it becomes necessary to use another approach to calculate them, such as optimizing a second set of minimal basis functions⁴³, which can be expensive. However, as with the diagonalization approach, the ability to work directly with the coefficients makes it possible to optimize the support functions and coefficients to accurately represent

a few states above the Fermi level at the same time as the occupied states, without significantly impacting the cost.

The FOE approach leads to linear scaling if the sparsity of the Hamiltonian is exploited and if the build up of the density matrix is done within localization regions. These localization regions do not need to be identical to the localization regions employed for the calculation of the adaptively-contracted basis set – in general they will actually be larger. The final result turns out to be relatively insensitive to the choice of size for the density matrix localization regions, above a sensible minimum value. By exploiting this sparsity, we can access system sizes of around ten thousand atoms with a moderate use of parallel resources.

C. Parallelization

Like standard BigDFT we have a multi-level MPI/OpenMP parallelization scheme⁴⁴. The details of the MPI parallelization are presented in Appendix D. In Fig. 4 we show the effective speedup as a function of the number of cores for a large water droplet, keeping the number of OpenMP threads at four. We measured the parallelization up to 3840 cores, by taking as a reference a 160 core run. As can be seen, the effective speedup reaches about 92% of the ideal value at 480 cores, decreasing to 62% for 3840 cores. Fitting the data to Amdahl's law⁴⁵ shows that at least 97% of the code has been parallelized; the true value is higher as the times shown include the communications and are relative to 160 cores rather than a serial run. Fig. 4 also shows a breakdown of the total calculation time into different categories, where we see that the communications start to become limiting for the highest number of cores, demonstrating that this is at the upper bound which is appropriate for this system size. For a larger system, this limit will of course be higher.

V. CALCULATION OF IONIC FORCES

In a self-consistent KS calculation (i.e. when the charge density is derived from the numerical set of wavefunctions), the forces acting on atom a are given by the negative gradient of the band structure energy with respect to the atomic positions \mathbf{R}_a . The Hellmann-Feynman force, given by the expression

$$\mathbf{F}_a^{(HF)} = - \sum_i f_i \left\langle \psi_i \left| \frac{\partial \mathcal{H}}{\partial \mathbf{R}_a} \right| \psi_i \right\rangle, \quad (27)$$

involves only the functional derivative of the Hamiltonian operator. This term is evaluated numerically in the computational setup used to express the ground state energy. As explained in more detail in Appendix C 1, with the cubic version of BigDFT, only the Hellmann-Feynman term




FIG. 4. Parallel scaling for a water droplet containing 960 atoms for between 160 and 3840 cores, with the effective and ideal speedup with respect to 160 cores, including a fit of Amdahl's law [below], and a breakdown of the calculation time for selected number of processors [above]. The categories presented are for communications, linear algebra, convolutions, calculation of the potential and density, and miscellaneous.

contributes to the forces, as the remaining part tends to zero in the limit of small grid spacings.

However, when the KS orbitals are expressed in terms of the support functions, there is an additional contribution which is not captured by the computational setup. As demonstrated in Appendix C 2, it is given by

$$\mathbf{F}_a^{(P)} = -2 \sum_{\alpha, \beta} \text{Re} \left(K^{\alpha\beta} \left\langle \chi_\beta \left| \frac{\partial \mathcal{L}^{(\alpha)}}{\partial \mathbf{R}_a} \right| \phi_\alpha \right\rangle \right), \quad (28)$$

where

$$|\chi_\alpha\rangle = \mathcal{H}_{KS} |\phi_\alpha\rangle - \sum_j \sum_{\rho, \sigma} c_j^\rho \varepsilon_j c_j^{\sigma*} S_{\sigma\alpha} |\phi_\rho\rangle \quad (29)$$

is the residual vector of the support function $|\phi_\alpha\rangle$, which is related to the support function gradient $|g^\alpha\rangle$ (see Eq. (15)). This term can be considered as the equivalent of a Pulay contribution to the ionic forces, arising from the explicit dependence of the localization operators on the atomic positions.

The vector $\frac{\partial \mathcal{L}^{(\alpha)}}{\partial \mathbf{R}_a} |\phi_\alpha\rangle$ only depends on the value of the support function on the borders of the localization regions (Eq. (C12)). Therefore, if the scalar product between the residues $|\chi_\beta\rangle$ and the values of the support functions $|\phi_\alpha\rangle$ at the boundaries of their localization regions is smaller than the norm of the residue itself (quantifying the accuracy of the results), the Pulay term can be safely neglected.

As mentioned in Section IV A, the Laplacian operator in the wavelet basis causes the kinetic energy to be high if the values at the edges are non-negligible. Such a situation is therefore penalized by the energy minimization and so the values at the borders are guaranteed to remain

low. Indeed, we have seen excellent agreement between the Hellmann-Feynman term only and the forces calculated using standard BigDFT, as will be demonstrated in sections VI and VIII A.

The Hellmann-Feynman force is thus the only relevant term even in the adaptively-contracted basis approach and is given by

$$\mathbf{F}_a^{(HF)} = - \sum_{\alpha, \beta} K^{\alpha\beta} \left\langle \phi_\alpha \left| \frac{\partial \mathcal{H}}{\partial \mathbf{R}_a} \right| \phi_\beta \right\rangle. \quad (30)$$

It is identical to the implementation in standard BigDFT⁸ and so the different terms are not repeated here. The only difference is that instead of applying the operator to the wavefunctions, we now apply it to all overlapping support functions. This can be done efficiently since each support function overlaps with only a few neighbors.

VI. ACCURACY

We have applied our minimal basis approach to a number of systems, depicted in Fig. 5, in order to demonstrate both its accuracy and its applicability. All calculations have been done using the local density approximation (LDA) exchange-correlation functional⁴⁶ and HGH pseudopotentials²⁹. However it is worth noting that the sole usage of LDA does not imply a general restriction and other functionals can be used as well; as an example, a PBE⁴⁷ calculation is presented for one system. In addition, we have used free boundary conditions, avoiding the need for the supercell approximation. The values of the wavelet basis parameters for the different systems, as well as the localization radii, were selected in order to achieve accuracies better than 1 meV/atom. This corresponded to values between 0.13 Å and 0.20 Å for h , 5.0 and 7.0 for λ^c and 7.0 and 8.0 for λ^f . Unless otherwise stated, we have used the direct minimization scheme for the density kernel optimization. For hydrogen atoms one basis function was used per atom, whereas for all other elements four basis functions were used per atom, except where otherwise stated.

A. Benchmark systems

We demonstrate excellent agreement with the traditional cubic scaling method for both energy and forces – of the order of 1 meV/atom for the energy and a few meV/Å for the forces, as shown in Tab. I. We also demonstrate systematic convergence of the total energy and forces with respect to localization radius for a molecule of C₆₀, where the largest localization regions are close to the total system size, as depicted in Fig. 6. For all these systems the level of accuracy achieved for the forces using the Hellmann-Feynman term only is of the same order as that of the cubic code.

B. Silicon defect energy

In order to demonstrate the accuracy of our method for a practical application we calculated the energy of a vacancy defect in a hydrogen terminated silicon cluster containing 291 atoms, shown in Fig. 5(f). As shown in Tab. II, the difference in defect energy between the cubic reference calculation and the linear version is 129 meV using 4 support functions per Si atom and 1 support function per hydrogen atom. Even more accurate results can be achieved by increasing the number of support functions per Si atom to 9, which reduces the error to 12 meV. Increasing the localization radii does not further improve the accuracy as the result is already within the noise level. To achieve these results, support functions were optimized using the hybrid mode and the density kernel was optimized using the FOE approach with a cutoff of 7.94 Å for the kernel construction.

C. Consistency of energies and forces

Following the discussion in Section V, we have calculated the average value of the support functions on the borders of their localization regions for various systems and found this to be at least three orders of magnitude smaller than the norm of the support function residue (defined in Eq. (C1)). This is in line with our expectations, as discussed in Section V, and implies that the Pulay terms should be negligible compared to the error introduced to the Hellmann-Feynman term due to the localization constraint. Indeed, this agrees with the calculated forces for the systems presented thus far. To further verify that the Pulay term can be neglected and to quantify the different sources of errors, we have also checked that the calculated forces are consistent with the energy, i.e. that they correspond to its negative derivative. To this end, initial and final configurations $\mathbf{R}^{(a)}$ and $\mathbf{R}^{(b)}$ of a given system were chosen, where \mathbf{R} represents the atomic positions. Small steps $\Delta \mathbf{R}$ were then taken between $\mathbf{R}^{(a)}$ to $\mathbf{R}^{(b)}$. If the forces \mathbf{F} are correctly evaluated we should have

$$\Delta E = \int_a^b \mathbf{F}(\mathbf{r}) \cdot d\mathbf{r} \approx \sum_{\mu=a}^b \mathbf{F}(\mathbf{R}_\mu) \cdot \Delta \mathbf{R}_\mu, \quad (31)$$

where μ labels the intermediate steps between configurations $\mathbf{R}^{(a)}$ and $\mathbf{R}^{(b)}$. This approximation can be compared with the exact value obtained by directly calculating the energy differences, i.e. $E(\mathbf{R}^{(b)}) - E(\mathbf{R}^{(a)})$. These two values should agree with each other up to the noise level of the calculation.

To analyze the different terms contributing to the noise in the forces, we use a combination of the hybrid and FOE methods with five progressive setups which give an estimate of the magnitude of the various error sources:

1. Using the cubic scaling scheme where all orbitals can extend over the full simulation cell.




FIG. 5. The different systems studied, where gray denotes carbon, white hydrogen, red oxygen, gold nitrogen, bronze boron, green silicon and blue sulfur. The values used for the support function localization radii are also given.

	Num. atoms	Energy (eV)			Forces (eV/ \AA)		
		Min. Basis	Cubic	(Min. - Cub.)/atom	Min. Basis	Cubic	Av. (Min. - Cub.)
Cinchonidine LDA	44	$-4.273 \cdot 10^3$	$-4.273 \cdot 10^3$	$2.0 \cdot 10^{-3}$	$2.734 \cdot 10^{-2}$	$2.772 \cdot 10^{-2}$	$4.3 \cdot 10^{-3}$
Cinchonidine PBE	44	$-4.274 \cdot 10^3$	$-4.274 \cdot 10^3$	$3.0 \cdot 10^{-4}$	$2.581 \cdot 10^{-2}$	$2.610 \cdot 10^{-2}$	$2.2 \cdot 10^{-3}$
Boron cluster	80	$-6.141 \cdot 10^3$	$-6.141 \cdot 10^3$	$2.4 \cdot 10^{-3}$	$1.305 \cdot 10^{-2}$	$1.316 \cdot 10^{-2}$	$4.2 \cdot 10^{-3}$
Alkane	257	$-1.592 \cdot 10^4$	$-1.592 \cdot 10^4$	$3.7 \cdot 10^{-4}$	$3.760 \cdot 10^{-2}$	$3.767 \cdot 10^{-2}$	$1.0 \cdot 10^{-3}$
Water	450	$-7.011 \cdot 10^4$	$-7.012 \cdot 10^4$	$1.7 \cdot 10^{-3}$	$3.730 \cdot 10^{-2}$	$3.730 \cdot 10^{-2}$	$2.5 \cdot 10^{-3}$

TABLE I. Energy and force differences between the minimal basis approach and the standard cubic version of the code. For the forces, the root mean squared force is given for each approach, as well as the average difference between each component.

2. Using the minimal basis approach but without localization constraints or confining potential.
3. Applying a confining potential but no localization constraints.
4. Using a finite localization radius of 7.94\AA for the density kernel, but not for the support functions.
5. Applying in addition strong localization radii of 4.76\AA to the support functions.

This test was done for a 92 atom alkane – despite the relatively small system size the introduction of finite cutoff radii for the support functions and the density kernel

construction has a strong effect since for chain-like structures the volume of the localization region is only a small fraction of the total computational volume. The results are shown in Tab. III, for a step size of $\Delta \mathbf{R}_\mu = 0.003\text{\AA}$. As expected, without the application of the localization constraint, the errors for the minimal basis calculations are of the same order of magnitude as the reference cubic calculation. This is also the case when a finite cutoff radius for the construction of the density kernel is introduced. Once a finite localization is imposed on the support functions the discrepancy between the energy difference and the force integral increases by an order of magnitude, however it remains small, agreeing with our




FIG. 6. Convergence of the total energy and forces with respect to the localization radius as compared to the cubic values for a molecule of C_{60} . For large localization regions the accuracy of the forces is sufficient for high accuracy geometry optimizations.

	pristine eV	vacancy eV	Δ eV	$\Delta - \Delta_{\text{cubic}}$ meV
cubic	-20674.223	-20563.056	111.167	-
4/1	-20667.556	-20556.518	111.038	129
9/1	-20672.856	-20561.701	111.155	12

TABLE II. Total energies and energy differences for a silicon cluster (Fig. 5(f)) with and without a vacancy defect. The first column shows the number of support functions per Si and H atom, respectively, and the last column shows the difference in defect energy between the cubic and linear versions.

previous observations about the Pulay forces for large enough localization radii.

Setup	Sup. func.	Conf. pot.	$I^{\alpha\beta}$ cutoff	ϕ_e cutoff	ΔE	$\int \mathbf{F}(\mathbf{r}) \cdot d\mathbf{r}$	diff.
1	✗	✗	✗	✗	5.666961	5.667190	$-2.3 \cdot 10^{-4}$
2	✓	✗	✗	✗	5.666966	5.666999	$-3.3 \cdot 10^{-5}$
3	✓	✓	✗	✗	5.666958	5.667024	$-6.5 \cdot 10^{-5}$
4	✓	✓	✓	✗	5.667239	5.667024	$2.1 \cdot 10^{-4}$
5	✓	✓	✓	✓	5.669992	5.673043	$-3.1 \cdot 10^{-3}$

TABLE III. Force calculations for the five setups described in the text, with all quantities given in eV.

VII. SCALING AND CROSSOVER POINT

We have applied the minimal basis method to alkanes, applying both the direct minimization and FOE approaches for the kernel optimization. The time taken per iteration is compared with the traditional cubic-scaling version in Fig. 7. The number of iterations required to reach convergence is similar for the cubic and minimal basis approaches and is approximately constant across system sizes, so that the total time taken shows similar




FIG. 7. Comparison between the time taken per iteration and memory usage for the cubic scaling and minimal basis approaches using both direct minimization (Dmin) and FOE for increasing length alkanes, where the time is for the wavefunction optimization, neglecting the input guess. The coefficients are shown for the corresponding cubic polynomials. A fixed number of 301 MPI tasks and 8 OpenMP threads was used.

behaviour. The results clearly demonstrate the improved scaling of the method, with a crossover point for the total time at around 150 atoms. This will of course be system dependent – the chain like nature of the alkanes makes them a particularly favorable system for the minimal basis approach. We also plot cubic polynomials for the timing data; whilst the cubic scaling approach only has a very small cubic term, both this and the quadratic term are noticeably reduced for both the FOE and direct minimization approaches. Indeed, the FOE method is predominantly linear scaling, compared to direct minimization which has larger quadratic and cubic terms, mainly due to the linear algebra, as expected.

The minimal basis approach also gives considerable savings in memory; for the above example the memory requirements for the cubic version still prohibit calculations on systems bigger than around 2000 atoms for the chosen number of processors, whereas for the minimal basis method the memory requirements allow calculations of up to 4000 atoms using direct minimization, and nearly 8000 atoms with FOE.

To take full advantage of the improvements made to BigDFT, it is not enough for the time taken per iteration to scale favorably with respect to system size, it is also necessary for the number of iterations needed to reach convergence not to increase with system size. We have demonstrated such behavior for increasing sized randomly generated non-equilibrium water droplets, as shown in Fig. 8. The number of iterations required to reach a good level of agreement with the cubic scaling version of the code remains approximately constant, with the fluctuations due to the random noise in the bond




FIG. 8. Convergence behavior for water droplets, showing the convergence with respect to outer loop iteration number for the system containing 60 atoms [above] and the energy difference with the converged cubic value after certain numbers of iterations [below]. The number of iterations is indicated by the color used; circles denote iterations with a confining potential and squares without.

lengths of the water molecules. Furthermore, the energy converges rapidly to a value very close to that obtained with the cubic code, as illustrated by the upper panel. We have also observed similar convergence behavior for other systems, including alkanes, as mentioned above.

VIII. FLEXIBILITY OF THE MINIMAL BASIS FORMALISM

A. Geometry optimization

As a further test of the quality of the forces and as a demonstration of the flexibility of the minimal basis formalism, we have performed a geometry optimization for a segment of a SiC nanotube containing 288 atoms, depicted in Fig. 5(g). Here we can take advantage of the minimal basis formalism by reusing the optimized support functions from the previous geometry step as an improved input guess, moving them with the atoms using an interpolation scheme to account for atomic displacements which are not multiples of the grid spacing h . This has the effect of reducing the number of iterations required to converge the support functions for each new geometry. In fact, for cases where the atoms have only moved a small amount, they will hardly need optimizing at all and so substantial savings can be made. A similar procedure also exists for the cubic version, but the minimal basis approach can profit much more because of the direct relation between the support function centers and the atomic positions.

We compared the convergence behavior and time taken for the minimal basis approach both with and without




FIG. 9. Geometry optimization for a segment of a SiC nanotube for the cubic and minimal basis approaches where the support functions are regenerated from atomic orbitals at each geometry step (“reset”) and where they are reformatted for reuse at the next iteration (“reformat”). The time taken for each step, cumulative time, force convergence and average distance from the final structure optimized using standard BigDFT are plotted for each step of the geometry optimization.

reusing the support functions at each geometry step with that for the standard cubic approach, for which the results are shown in Fig. 9. It is clear that the Hellmann-Feynman forces are sufficiently accurate to optimize the structure to the required level – in this case forces of below 10^{-2} eV/Å are readily achieved. For this system size we are below the crossover point, such that when the support functions are reset at each geometry step the time taken per step is greater than that for the cubic approach. However, the reuse of support functions results in a significant reduction in the number of steps required to fully converge the support functions, and so the total time is less than that required for the cubic approach. This means that the crossover point will be reduced for geometry optimizations or molecular dynamics calculations, opening up further possibilities for the highly accurate study of dynamics of large systems.

B. Charged systems

As previously mentioned, the ability to use free boundary conditions is essential for charged systems. This has enabled us to perform calculations of isolated segments of ladder polythiophene (LPT) (Fig. 5(h)), initially in a neutral state and then adding a charge of plus or minus two electrons. The support functions from the neutral case are also well suited to the charged system so that only kernel optimizations are required, which can reduce the computational cost by an order of magnitude.

In Tab. IV we compare the agreement between the minimal basis approach and the standard cubic approach for

Q	ΔE_Q	$ \Delta(E_Q - E_0) $
0	opt. 176	–
-2	opt. 147	28
	unopt. 292	117
+2	opt. 128	47
	unopt. 200	24

TABLE IV. Energy differences between the standard cubic version and the minimal basis approach for absolute energies (ΔE_Q) and energy differences with the neutral system ($|\Delta(E_Q - E_0)|$) for 63 atom segments of LPT with a net charge Q . Results are shown both for fully optimized support functions (“opt”) and for support functions reused from the neutral calculation (“unopt”). All results are in meV.

a system containing 63 atoms. We demonstrate an agreement of the order of 100 meV for the energy differences for both the fully optimized set of support functions and the reuse of the support functions from the neutral system. For the negatively charged calculations we have also confirmed that this level of accuracy is maintained up to 300 atoms, beyond which size the cost of calculations with the cubic version of BigDFT increases significantly.

In order to converge the results obtained with the minimal basis to a good level of accuracy, we used 9 support functions for carbon and sulfur and 1 per hydrogen. For charged systems we have found that the direct minimization method is more stable, as it allows us to update the coefficients in smaller steps before updating the kernel and therefore density, rather than fully converging them before each update.

We expect such support function reuse to be generally applicable for systems where the addition of a charge only results in a perturbation of the electronic structure. However it may be necessary to optimize a few unoccupied states (using direct minimization or diagonalization) in order to ensure that the adaptively-contracted basis is sufficiently accurate for negatively charged systems.

IX. CONCLUSION

We have presented a self-consistent minimal basis approach within BigDFT which leads to a reduced scaling behavior with system size and allows the treatment of larger systems than can be treated with the cubic version; for very large systems linear scaling is clearly visible. The use of a small set of nearly orthogonal adaptively-contracted basis functions which are optimized *in situ* in the underlying wavelet basis set gives rise to sparse matrices of relatively small size. For the optimization of these so-called support functions we use a confining potential which on the one hand helps to keep the support functions strictly localized, and on the other hand helps to alleviate the notorious ill-conditioning which is typical of linear scaling approaches.

The standard cubic scaling version of BigDFT has been

previously demonstrated to give highly accurate results and so we use this as a standard of comparison for our method. We have demonstrated for a number of different systems excellent agreement with the cubic version for both energy and forces. In particular, we have demonstrated that it is not necessary to include Pulay-like correction terms to the atomic forces, thanks to the nature of the Laplacian operator in the wavelet basis which ensures the support functions remain negligible on the borders of the localization regions. In addition, we have shown consistent convergence behavior across a range of system sizes. From the viewpoint of scaling with the number of atoms we have demonstrated linear scaling for the FOE method where the linear algebra has been written to exploit the sparsity of the matrices.

Finally, we have highlighted some of the advantages of using localized support functions expressed in a wavelet basis set. These include the ability to further accelerate geometry optimizations by reusing the support functions from the previous geometry, and the possibility of achieving a good level of accuracy for a charged calculation by reusing the support functions from a neutral calculation.

By directly working in the basis of the support functions, we can therefore reduce the number of degrees of freedom needed to express the KS operators for a targeted accuracy. Aside from reducing the computational overhead, this flexible approach paves the way for future developments, where the adaptively-contracted basis functions can be reused in other situations, including for example constrained DFT calculations of large systems. Work is ongoing in this direction.

The authors would like to acknowledge funding from the European project MMM@HPC (RI-261594), the CEA-NANOSCIENCE BigPOL project, the ANR projects SAMSON (ANR-AA08-COSI-015) and NEWCASTLE, and the Swiss CSCS grants s142 and h01. CPU time and assistance were provided by CSCS, IDRIS, Oak Ridge National Laboratory and Argonne National Laboratory.

Appendix A: Prescription for reducing the confinement

To derive a prescription for reducing the confinement, it is assumed that the change in the target function between successive iterations of the minimization procedure can be approximated to first order by

$$\Delta\Omega'_{(n)} = \sum_{\alpha} \langle g_{(n)}^{\alpha} | \Delta\phi_{(n)}^{\alpha} \rangle, \quad (\text{A1})$$

where $|\Delta\phi_{(n)}^{\alpha}\rangle$ is the change in support function between iterations n and $n+1$, i.e. $|\Delta\phi_{(n)}^{\alpha}\rangle = |\phi_{(n+1)}^{\alpha}\rangle - |\phi_{(n)}^{\alpha}\rangle$, and $|g_{(n)}^{\alpha}\rangle$ is the gradient of the target function with respect to the support function at iteration n . Due to the influence of the confinement and the localization regions, the gradient of the support functions and thus $\Delta\Omega'_{(n)}$ will not go

down to zero. However, the actual change in the target function, $\Delta\Omega_{(n)}$, will at some point go to zero, meaning that further optimization becomes impossible for the localization region and confining potential currently used. In this case, the only way to further minimize the target function is to decrease the confining potential. Therefore at each step of the minimization the ratio between the actual and estimated decreases in the target function is determined:

$$\kappa = \frac{\Delta\Omega_{(n)}}{\Delta\Omega'_{(n)}}. \quad (\text{A2})$$

This value is then used to update the confinement prefactor, a_α , at the start of the following support function optimization loop, via

$$a_\alpha^{new} = \kappa a_\alpha^{old}. \quad (\text{A3})$$

If κ is of the order of one, this implies there is still some scope for optimizing the support functions using the current confining potential and it should not be updated. If, on the other hand, κ is much smaller, it will hardly be possible to further improve the support functions and so the magnitude of the confining potential should be decreased. In this way one gets a smooth transformation from the hybrid expression to the energy expression.

Appendix B: Fermi operator expansion

In the FOE method the density kernel is given as a sum of Chebyshev polynomials. Since these polynomials are only defined in the interval $[-1, 1]$, it is necessary to shift and scale the Hamiltonian such that its eigenvalue spectrum lies within this interval. If ε_{min} and ε_{max} are the smallest and largest eigenvalues that would result from diagonalizing the Hamiltonian matrix according to $\mathbf{H}\mathbf{c}_i = \varepsilon_i\mathbf{S}\mathbf{c}_i$, then the scaled Hamiltonian, $\tilde{\mathbf{H}}$, has to be built using

$$\tilde{\mathbf{H}} = \sigma(\mathbf{H} - \tau\mathbf{S}), \quad (\text{B1})$$

with

$$\sigma = \frac{2}{\varepsilon_{max} - \varepsilon_{min}}, \quad \tau = \frac{\varepsilon_{min} + \varepsilon_{max}}{2}. \quad (\text{B2})$$

Now the density kernel can be calculated according to

$$\mathbf{K}' \approx \mathbf{p}(\tilde{\mathbf{H}}') = \frac{c_0}{2}\mathbf{I} + \sum_{i=1}^{n_{pl}} c_i \mathbf{T}^i(\tilde{\mathbf{H}}') \quad (\text{B3})$$

with

$$\tilde{\mathbf{H}}' = \mathbf{S}^{-1/2}\tilde{\mathbf{H}}\mathbf{S}^{-1/2}, \quad (\text{B4})$$

where \mathbf{I} is the identity matrix, \mathbf{T}^i the Chebyshev polynomial of order i , \mathbf{S} the support function overlap matrix and $\mathbf{S}^{-1/2}$ is calculated using a first order Taylor expansion.

To determine the expansion coefficients c_i , one has to recall that the density matrix of Eq. (5) is a projection operator onto the occupied subspace of the KS orbitals:

$$\langle \psi_i | F | \psi_j \rangle = f(\varepsilon_j) \delta_{ij}. \quad (\text{B5})$$

Since F and \mathcal{H} have the same eigenfunctions, one can express the polynomial $p(\mathcal{H})$ in the same way, leading to

$$\langle \psi_i | p(\mathcal{H}) | \psi_j \rangle = p(\varepsilon_j) \delta_{ij} \quad (\text{B6})$$

with

$$p(\varepsilon) = \frac{c_0}{2} + \sum_{i=1}^{n_{pl}} c_i T^i(\varepsilon). \quad (\text{B7})$$

By comparing Eqs. (B5) and (B6) it becomes clear that the polynomial expansion $p(\varepsilon)$ has to approximate the Fermi function $f(\varepsilon)$ in the interval $[-1, 1]$. Thus the coefficients c_i are simply given by the expansion of the Fermi function in terms of the Chebyshev polynomials. The time for this step is negligible compared to the other operations related to the FOE. However, in practice it turns out that it is advantageous to replace the Fermi function by

$$f(\varepsilon) = \frac{1}{2} \left[1 - \operatorname{erf} \left(\frac{\varepsilon - \mu}{\Delta\varepsilon} \right) \right], \quad (\text{B8})$$

since it approaches the limits 0 and 1 faster as one goes away from the chemical potential. $\Delta\varepsilon$ is typically a fraction of the band gap.

The last step is to evaluate the Chebyshev polynomials and to build the density kernel. If the l th column of the Chebyshev matrix \mathbf{T} is denoted by \mathbf{t}_l , then these vectors fulfill the recursion relation

$$\begin{aligned} \mathbf{t}_l^0 &= \mathbf{e}_l, \\ \mathbf{t}_l^1 &= \tilde{\mathbf{H}}' \mathbf{e}_l, \\ \mathbf{t}_l^{j+1} &= 2\tilde{\mathbf{H}}' \mathbf{t}_l^j - \mathbf{t}_l^{j-1}, \end{aligned} \quad (\text{B9})$$

where \mathbf{e}_l is the l th column of the identity matrix. The l th column of the density kernel, denoted by \mathbf{k}_l , is then given by the linear combination of all the columns \mathbf{t}_l according to Eq. (B3), i.e.

$$\mathbf{k}_l' = \frac{c_0}{2} \mathbf{t}_l^0 + \sum_{i=1}^{n_{pl}} c_i \mathbf{t}_l^i. \quad (\text{B10})$$

This demonstrates that the density kernel can be constructed using only matrix vector multiplications.

Since the correct value of the Fermi energy is initially unknown, this procedure has to be repeated until the correct value has been found, so that $\operatorname{Tr}(\mathbf{K}')$ is equal to the number of electrons in the system. Finally the kernel \mathbf{K} is given by $\mathbf{S}^{-1/2}\mathbf{K}'\mathbf{S}^{-1/2}$ and the band-structure energy can then be calculated by reversing the scaling and shifting operations:

$$E_{BS} = \frac{\operatorname{Tr}(\mathbf{K}\tilde{\mathbf{H}})}{\sigma} + \tau \operatorname{Tr}(\mathbf{K}\mathbf{S}). \quad (\text{B11})$$

Appendix C: Pulay forces

1. The traditional cubic approach

Numerically, the set of $|\Psi_i\rangle$ is expressed in a finite basis set. This means that the action of \mathcal{H}_{KS} can in principle lie *outside* the span of the $|\Psi_i\rangle$. Let us suppose that the KS Hamiltonian and orbitals are expressed in a basis set which is complete enough to describe them within a targeted accuracy Δ . For the Daubechies basis in the traditional BigDFT approach, this happens when the grid spacing h is such as to describe the PSP and orbital oscillations, and the radii $\lambda^{c,f}$ such as to contain the decreasing tails of the wavefunctions. This situation indeed corresponds to the traditional setup of a BigDFT run. We can therefore define a residual function

$$|\chi_i\rangle = \mathcal{H}_{KS} |\Psi_i\rangle - \epsilon_i |\Psi_i\rangle, \quad (\text{C1})$$

which is of course zero when the *numerical* KS orbital is the *exact* KS orbital. By definition $\langle \Psi_j | \chi_i \rangle = 0 \forall i, j$. The norm of this vector, once *projected* in the basis set used to express $|\Psi_i\rangle$, is often used as a convergence criterion for the ground state energy.

Even though the basis set is finite, the orthogonality of the KS orbitals holds exactly, implying $\text{Re} \left(\left\langle \Psi_i \left| \frac{d\Psi_i}{d\mathbf{R}_a} \right. \right\rangle \right) = 0$. It is thus easy to show that the *numerical* atomic forces are defined as follows:

$$-\frac{dE_{BS}}{d\mathbf{R}_a} = -\sum_i \langle \Psi_i | \frac{\partial \mathcal{H}_{KS}}{\partial \mathbf{R}_a} | \Psi_i \rangle - 2 \sum_i \text{Re} \left(\left\langle \chi_i \left| \frac{d\Psi_i}{d\mathbf{R}_a} \right. \right\rangle \right), \quad (\text{C2})$$

where the first term of the right hand side of the above equation is the Hellmann-Feynman contribution to the forces. The norm of $|\chi_i\rangle$ (Eq. (C1)) can be reduced within the *same* basis set to meet the targeted accuracy Δ . Therefore the projection of $\left| \frac{d\Psi_i}{d\mathbf{R}_a} \right\rangle$ onto the basis set used for the calculation can be safely neglected as it is associated with the same numerical precision. Consequently, the atomic forces can be evaluated by the Hellmann-Feynman term only as the remaining part is proportional to Δ .

2. The minimal basis approach

As mentioned in the main text, when the KS orbitals are expressed in terms of the support functions, an addi-

tional Pulay-like term should in principle be taken into account. To demonstrate this, we define – in analogy to Eq. (C1) – the support function residue $|\chi_\alpha\rangle$, which becomes, using the identity $\tilde{H}_{\rho\sigma} = \sum_j c_j^\rho \varepsilon_j c_j^{\sigma*}$,

$$\begin{aligned} |\chi_\alpha\rangle &= \mathcal{H}_{KS} |\phi_\alpha\rangle - \left(\sum_{\rho,\sigma} |\phi_\rho\rangle \tilde{H}_{\rho\sigma} \langle \phi_\sigma | \right) |\phi_\alpha\rangle \\ &= \mathcal{H}_{KS} |\phi_\alpha\rangle - \sum_j \sum_{\rho,\sigma} c_j^\rho \varepsilon_j c_j^{\sigma*} S_{\sigma\alpha} |\phi_\rho\rangle. \end{aligned} \quad (\text{C3})$$

Next, inserting the definition of χ_i (Eq. (C1)) into the non-Hellmann-Feynman contribution of Eq. (C2) and using the relation $\text{Re} \left(\left\langle \Psi_i \left| \frac{d\Psi_i}{d\mathbf{R}_a} \right. \right\rangle \right) = 0$ one obtains

$$\mathbf{F}_a - \mathbf{F}_a^{(HF)} = -2 \sum_i \text{Re} \left(\left\langle \Psi_i \left| \mathcal{H}_{KS} \left| \frac{d\Psi_i}{d\mathbf{R}_a} \right. \right. \right\rangle \right). \quad (\text{C4})$$

Expanding the KS orbitals in terms of the support functions, using the relation $H_{\alpha\beta} = \sum_j \sum_{\rho,\sigma} \varepsilon_j c_j^\rho c_j^{\sigma*} S_{\alpha\rho} S_{\sigma\beta}$ and the orthonormality of the KS orbitals, we can write

$$\begin{aligned} \mathbf{F}_a - \mathbf{F}_a^{(HF)} &= -2 \sum_i \sum_{\alpha,\beta} \text{Re} \left(c_i^{\alpha*} c_i^\beta \left\langle \phi_\alpha \left| \mathcal{H}_{KS} \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right) \\ &\quad - 2 \sum_i \sum_{\alpha,\beta} \text{Re} \left(c_i^{\alpha*} \frac{dc_i^\beta}{d\mathbf{R}_a} \langle \phi_\alpha | \mathcal{H}_{KS} | \phi_\beta \rangle \right) \\ &= -2 \sum_i \sum_{\alpha,\beta} \text{Re} \left(c_i^{\alpha*} c_i^\beta \left\langle \phi_\alpha \left| \mathcal{H}_{KS} \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right) \\ &\quad - 2 \sum_i \sum_{\beta,\sigma} \text{Re} \left(\frac{dc_i^\beta}{d\mathbf{R}_a} \varepsilon_i c_i^{\sigma*} S_{\sigma\beta} \right). \end{aligned} \quad (\text{C5})$$

From the orthonormality of the KS orbitals one can derive the relation

$$2 \sum_{\alpha,\beta} \text{Re} \left(\frac{dc_i^\alpha}{d\mathbf{R}_a} c_i^{\beta*} S_{\alpha\beta} \right) = - \sum_{\alpha,\beta} c_i^{\alpha*} c_i^\beta \frac{dS_{\alpha\beta}}{d\mathbf{R}_a}. \quad (\text{C6})$$

Inserting this into Eq. (C5) yields

$$\begin{aligned} \mathbf{F}_a - \mathbf{F}_a^{(HF)} &= -2 \sum_i \sum_{\alpha,\beta} \text{Re} \left(c_i^{\alpha*} c_i^\beta \left\langle \phi_\alpha \left| \mathcal{H}_{KS} \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right) \\ &\quad + \sum_i \sum_{\beta,\sigma} \text{Re} \left(c_i^\beta c_i^{\sigma*} \frac{dS_{\sigma\beta}}{d\mathbf{R}_a} \varepsilon_i \right). \end{aligned} \quad (\text{C7})$$

Again using the KS orthonormality condition, we can write

$$\begin{aligned}
\mathbf{F}_a - \mathbf{F}_a^{(HF)} &= -2 \sum_i \sum_{\alpha,\beta} \text{Re} \left(c_i^{\alpha*} c_i^\beta \left\langle \phi_\alpha \left| \mathcal{H}_{KS} \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right) + \sum_{i,j} \sum_{\alpha,\beta,\rho,\sigma} \text{Re} \left(c_i^{\alpha*} c_i^\beta c_j^{\sigma*} \varepsilon_i c_j^\rho S_{\alpha\rho} \frac{dS_{\sigma\beta}}{d\mathbf{R}_a} \right) \\
&= -2 \sum_{\alpha,\beta} \text{Re} \left(K^{\beta\alpha} \left\langle \phi_\alpha \left| \mathcal{H}_{KS} \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right) + 2 \sum_j \sum_{\alpha,\beta,\rho,\sigma} \text{Re} \left(K^{\beta\alpha} c_j^{\sigma*} \varepsilon_j c_j^\rho S_{\alpha\rho} \left\langle \phi_\sigma \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right), \tag{C8}
\end{aligned}$$

which becomes in terms of the support function residue of Eq.(C3)

$$\mathbf{F}_a - \mathbf{F}_a^{(HF)} = -2 \sum_{\alpha,\beta} \text{Re} \left(K^{\beta\alpha} \left\langle \chi_\alpha \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right). \tag{C9}$$

This result contains Eq. (C2) when no localization projectors are applied to the support function. Therefore the only term of the forces which cannot be captured within the localization regions is the part which is projected outside. The extra Pulay term due to the localization constraint is therefore

$$\mathbf{F}_a^{(P)} = -2 \sum_{\alpha,\beta} \text{Re} \left(K^{\beta\alpha} \left\langle \chi_\alpha \left| (1 - \mathcal{L}^{(\beta)}) \left| \frac{d\phi_\beta}{d\mathbf{R}_a} \right. \right. \right\rangle \right). \tag{C10}$$

Using Eq. (12), we can show

$$\mathbf{F}_a^{(P)} = -2 \sum_{\alpha,\beta} \text{Re} \left(K^{\beta\alpha} \left\langle \chi_\alpha \left| \frac{\partial \mathcal{L}^{(\beta)}}{\partial \mathbf{R}_a} \left| \phi_\beta \right. \right. \right\rangle \right). \tag{C11}$$

When the localization regions are atom-centered, the derivative of the projector $\mathcal{L}^{(a)}$ (as defined in Eq. (11)) can be evaluated analytically in the underlying basis set and is given by

$$\begin{aligned}
\frac{\partial \mathcal{L}^{(a)}}{\partial \mathbf{R}_\beta} &= \delta_{\alpha\beta} \delta_{i_1 j_1} \delta_{i_2 j_2} \delta_{i_3 j_3} \frac{\mathbf{R}_{(i_1, i_2, i_3)} - \mathbf{R}_\alpha}{R_{cut}} \\
&\quad \times \delta(R_{cut} - |\mathbf{R}_{(i_1, i_2, i_3)} - \mathbf{R}_\alpha|) \tag{C12}
\end{aligned}$$

This demonstrates that the Pulay term is only associated with the value of the support functions at the border of the localization regions.

Appendix D: Parallelization

It is a natural choice to divide the support functions between MPI tasks so that each one handles only a subset of support functions. For some operations these can be treated independently but for others, such as the calculation of scalar products between overlapping support functions needed to build the overlap and Hamiltonian matrices, communication of support functions between MPI tasks is required. One could directly exchange in a point-to-point fashion the parts of the support functions which overlap with each other, so that the scalar products can be calculated locally on each task. Although conceptually straightforward, this has severe drawbacks. Firstly the amount of data being communicated is tremendous

since the support functions generally have quite a notable overlap. This also results in a very poor ratio between computation and communication – in the extreme case where each task handles only one support function, each communicated element is only used for one operation. Secondly there can be enormous load imbalancing for free boundary conditions as support functions in the center of the system usually have more neighboring support functions than those near the edges. Finally, the data is split into a large number of small messages, which could result in a large overhead due to the latency of the network.

We therefore use a different approach, which requires a so-called “transposed” rather than “direct” arrangement of data. In this layout the simulation cell is partitioned among MPI tasks and the support functions are distributed to the various tasks such that each one can calculate a partial overlap matrix for a given region of the cell. Each task therefore has to receive those parts of all support functions which extend into its region. The partial matrices are then summed to build the full overlap matrix using MPI_Allreduce. This partitioning of the cell is done such that the load balancing among the MPI tasks is optimal, which in general does not correspond to a naive uniform distribution of the simulation cell. To determine the optimal layout a weight is assigned to each grid point, given by m^2 , where m is the number of support functions touching it (if symmetry can be exploited the weight should rather be $\frac{m(m+1)}{2}$); the total weight (i.e. the sum of all partial weights) is then divided among all MPI tasks as evenly as possible. In Fig. 10 this procedure is illustrated with a toy example, where in the upper part the support functions and their overlaps are shown and in the lower part the resulting direct and transposed (both naive and optimal) data layouts are given.

In addition to the better load balancing this approach has the advantage that considerably less data has to be communicated – since the transposed layout is just a redistribution of the standard layout, the total amount of data that is communicated is equal to the total size of all support functions, whereas in the point-to-point approach, the same data is often sent to multiple processes. Furthermore, the communication can be done more efficiently: after some local rearrangement of the data for each MPI task, it can be communicated with a single MPI call (MPI_Alltoallv) – in practice there are two calls since the coarse and fine parts are handled separately. After the data has been received some local rearrangement is again required to reach the correct layout. These three steps – local rearrangement, communication and further local rearrangement – are illustrated in Fig. 11. Due




FIG. 10. An example depicting four support functions (Roman numerals) in a system consisting of four grid points (Arabic numerals) [above]. The support functions are constructed such that each extends over two grid points. The various data layouts are also illustrated [below]: the direct layout where each MPI task has all the data for certain support functions [left]; the naive transposed layout where each MPI task has all the data for some grid points [centre]; and the optimized transposed layout which is similar to the previous case, but with an optimal load balancing [right].




FIG. 11. Illustration of the transposition process for the system shown in Fig. 10. In step a, the data is rearranged locally on each MPI task, after which it is communicated using a single collective call (`MPI_Alltoallv`), as shown in step b. Finally in step c it is again rearranged locally to reach the final layout.

to the latency of the network, two MPI calls will likely be more efficient than the very large number of small messages that have to be sent for the point-to-point approach.

For the calculation of the charge density, which is formally identical to the calculation of scalar products, a similar approach is used. Since these two operations are the most important ones from the viewpoint of communication and parallelization, this results in an excellent scaling with respect to the number of cores.

- ¹ P. Hohenberg and W. Kohn, *Phys. Rev.* **136**, B864 (1964)
- ² W. Kohn and L. J. Sham, *Phys. Rev.* **140**, A1133 (1965)
- ³ L. Hedin and S. Lundqvist, in *Solid State Physics*, Vol. 23, edited by F. Seitz, D. Turnbull, and H. Ehrenreich (Academic Press, New York, 1969) pp. 1–181
- ⁴ X. Gonze, B. Amadon, P.-M. Anglade, J.-M. Beuken, F. Bottin, P. Boulanger, F. Bruneval, D. Caliste, R. Caracas, M. Côté, T. Deutsch, L. Genovese, P. Ghosez, M. Giantomassi, S. Goedecker, D. Hamann, P. Hermet, F. Jollet, G. Jomard, S. Leroux, M. Mancini, S. Mazevet, M. Oliveira, G. Onida, Y. Pouillon, T. Rangel, G.-M. Rignanese, D. Sangalli, R. Shaltaf, M. Torrent, M. Verstraete, G. Zerah, and J. Zwanziger, *Comput. Phys. Commun.* **180**, 2582 (2009)
- ⁵ P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, A. D. 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, *J. Phys.: Condens. Matter* **21**, 395502 (2009)
- ⁶ M. D. Segall, P. J. D. Lindan, M. J. Probert, C. J. Pickard, P. J. Hasnip, S. J. Clark, and M. C. Payne, *J. Phys.: Condens. Matter* **14**, 2717 (2002)
- ⁷ J. E. Pask and P. A. Sterne, *Modell. Simul. Mater. Sci. Eng.* **13**, R71 (2005)
- ⁸ L. Genovese, A. Neelov, S. Goedecker, T. Deutsch, S. A. Ghasemi, A. Willand, D. Caliste, O. Zilberberg, M. Rayson, A. Bergman, and R. Schneider, *J. Chem. Phys.* **129**, 014109 (2008)
- ⁹ M. Valiev, E. Bylaska, N. Govind, K. Kowalski, T. Straatsma, H. V. Dam, D. Wang, J. Nieplocha, E. Apra, T. Windus, and W. de Jong, *Comput. Phys. Commun.* **181**, 1477 (2010)
- ¹⁰ V. Blum, R. Gehrke, F. Hanke, P. Havu, V. Havu, X. Ren, K. Reuter, and M. Scheffler, *Comput. Phys. Commun.* **180**, 2175 (2009)
- ¹¹ W. Kohn, *Phys. Rev.* **133**, A171 (1964)
- ¹² W. Kohn, *Phys. Rev. Lett.* **76**, 3168 (1996)
- ¹³ S. Goedecker, *Phys. Rev. B* **58**, 3501 (1998)
- ¹⁴ J. D. Cloizeaux, *Phys. Rev.* **135**, A685 (1964)
- ¹⁵ J. D. Cloizeaux, *Phys. Rev.* **135**, A698 (1964)
- ¹⁶ W. Kohn, *Phys. Rev.* **115**, 809 (1959)
- ¹⁷ R. Baer and M. Head-Gordon, *Phys. Rev. Lett.* **79**, 3962 (1997)
- ¹⁸ S. Ismail-Beigi and T. A. Arias, *Phys. Rev. Lett.* **82**, 2127 (1999)
- ¹⁹ L. He and D. Vanderbilt, *Phys. Rev. Lett.* **86**, 5341 (2001)
- ²⁰ M. Elstner, D. Porezag, G. Jungnickel, J. Elsner, M. Haugk, T. Frauenheim, S. Suhai, and G. Seifert, *Phys. Rev. B* **58**, 7260 (1998)
- ²¹ B. Aradi, B. Hourahine, and T. Frauenheim, *J. Phys. Chem. A* **111**, 5678 (2007)
- ²² D. R. Bowler and T. Miyazaki, *Rep. on Prog. in Phys.* **75**, 036503 (2012)
- ²³ S. Goedecker, *Rev. Mod. Phys.* **71**, 1085 (1999)

- ²⁴ C.-K. Skylaris, P. D. Haynes, A. A. Mostofi, and M. C. Payne, *J. Chem. Phys.* **122**, 084119 (2005)
- ²⁵ D. R. Bowler and T. Miyazaki, *J. Phys.: Condens. Matter* **22**, 074207 (2010)
- ²⁶ J. VandeVondele, M. Krack, F. Mohamed, M. Parrinello, T. Chassaing, and J. Hutter, *Comput. Phys. Commun.* **167**, 103 (2005)
- ²⁷ J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, *J. Phys.: Condens. Matter* **14**, 2745 (2002)
- ²⁸ I. Daubechies, *Ten Lectures on Wavelets* (SIAM, 1992)
- ²⁹ C. Hartwigsen, S. Goedecker, and J. Hutter, *Phys. Rev. B* **58**, 3641 (1998)
- ³⁰ M. Krack, *Theor. Chem. Acc.* **114**, 145 (2005)
- ³¹ A. Willand, Y. O. Kvashnin, L. Genovese, A. Vázquez-Mayagoitia, A. K. Deb, A. Sadeghi, T. Deutsch, and S. Goedecker, *J. Chem. Phys.* **138**, 104109 (2013)
- ³² E. Hernández and M. J. Gillan, *Phys. Rev. B* **51**, 10157 (1995)
- ³³ N. Marzari and D. Vanderbilt, *Phys. Rev. B* **56**, 12847 (1997)
- ³⁴ R. McWeeny, *Rev. Mod. Phys.* **32**, 335 (1960)
- ³⁵ S. Goedecker, *Wavelets and Their Application: For the Solution of Partial Differential Equations in Physics* (Presses polytechniques et universitaires romandes, 1998)
- ³⁶ A. Ruiz-Serrano and C.-K. Skylaris, *J. Chem. Phys.* **139**, 164110 (2013)
- ³⁷ P. W. Anderson, *Phys. Rev. Lett.* **21**, 13 (1968)
- ³⁸ P. Pulay, *Chem. Phys. Lett.* **73**, 393 (1980)
- ³⁹ S. Goedecker and C. J. Umrigar, *Phys. Rev. A* **55**, 1765 (1997)
- ⁴⁰ S. Goedecker and L. Colombo, *Phys. Rev. Lett.* **73**, 122 (1994)
- ⁴¹ S. Goedecker and M. Teter, *Phys. Rev. B* **51**, 9455 (1995)
- ⁴² W. H. Press, S. A. Teukolsky, W. T. Vetterling, and B. P. Flannery, *Numerical Recipes 3rd Edition: The Art of Scientific Computing*, 3rd ed. (Cambridge University Press, New York, NY, USA, 2007)
- ⁴³ L. E. Ratcliff, N. D. M. Hine, and P. D. Haynes, *Phys. Rev. B* **84**, 165131 (2011)
- ⁴⁴ L. Genovese, B. Videau, M. Ospici, T. Deutsch, S. Goedecker, and J.-F. Méhaut, *C. R. Méc.* **339**, 149 (2011)
- ⁴⁵ G. M. Amdahl, *IEEE Solid-State Circuits Newsletter* **12**, 19 (2007)
- ⁴⁶ D. M. Ceperley and B. J. Alder, *Phys. Rev. Lett.* **45**, 566 (1980)
- ⁴⁷ J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996), ISSN 0031-9007
- ⁴⁸ S. De, A. Willand, M. Amsler, P. Pochet, L. Genovese, and S. Goedecker, *Phys. Rev. Lett.* **106**, 225502 (2011)
- ⁴⁹ P. Pochet, L. Genovese, D. Caliste, I. Rousseau, S. Goedecker, and T. Deutsch, *Phys. Rev. B* **82**, 035431 (2010)