Highly efficient method for Kohn-Sham density functional calculations of 500–10 000 atom systems

M. J. Rayson
Max-Planck-Institut für Eisenforschung, Max-Planck-Straße 1, D-40237 Düsseldorf, Germany

P. R. Briddon
School of Electrical, Electronic and Computer Engineering, Newcastle University, NE1 7RU England, United Kingdom

(Received 18 March 2009; revised manuscript received 28 August 2009; published 4 November 2009)

A method for the solution of the self-consistent Kohn-Sham equations using Gaussian-type orbitals is presented. Accurate relative energies and forces are demonstrated to be achievable at a fraction of the computational expense for large systems. With this approach calculations involving around 1000 atoms can easily be performed with a serial desktop computer and ~10 000 atom systems are within reach of relatively modest parallel computational resources. The method is applicable to arbitrary systems including metals. The approach generates a minimal basis on the fly while retaining the accuracy of the large underpinning basis set. Convergence of energies and forces are given for clusters as well as cubic cells of silicon and aluminum, for which the formation energies of defects are calculated in systems up to 8000 and 4000 atoms, respectively. For these systems the method exhibits linear scaling with the number of atoms in the presently important size range of ~500–3000 atoms.

DOI: 10.1103/PhysRevB.80.205104 PACS number(s): 71.15.Dx

I. INTRODUCTION

The Kohn-Sham formalism[1] of density functional theory[2] (KSDFT) is one of the most widely used techniques for the *ab initio* study of systems containing many electrons and nuclei. Currently, calculations of around 500 atoms (we stress that *any* mention of system size in this work refers to three-dimensional dense packed systems, such as a cubic crystal) can be tackled *routinely* (by which we mean a self-consistent total energy that requires a few hours on 64 processors of a typical modern PC cluster) with a state-of-the-art Gaussian orbital code.[3,4] However, many pressing physical problems exist in the range ~1000–10 000 atoms. In this system size regime updating the density from a given potential $[V(\mathbf{r}) \rightarrow n(\mathbf{r})]$ is by far the time dominant step. Gaussian orbital codes utilizing direct diagonalization require $\sim N^3$ (where N is the number of basis functions) or $\sim N^2 m$ (where m is the number of state vectors) when using iterative methods.[4,5] Therefore, going between 10^3 and 10^4 atoms (using the same number of basis functions per atom) requires at least 1000 times more computational effort assuming no useful locality to be present in the density matrix (DM). All other components of a KSDFT self-consistent iteration in a Gaussian orbital code (such as evaluating the Hamiltonian matrix elements) can be performed in $\mathcal{O}(N_{\text{atom}})$ due to the locality of the basis functions.

The method, developed here, aims to achieve the following: (i) reproduction of accurate forces and relative energy differences with respect to the conventional Gaussian orbital approach, (ii) be valid for arbitrary systems (e.g., including metals and complex systems), and (iii) enable routine calculations for system sizes of around 10^3–10^4 atoms. The general approach will be to generate a small number of contracted functions on the fly (for example, four functions per atom for carbon or silicon using the pseudopotential approximation) that adequately describe the occupied subspace usually requiring at least ~20 basis functions per atom. The time dominant $\sim N^3$ or $\sim N^2 m$ operations in the $[V(\mathbf{r}) \rightarrow n(\mathbf{r})]$ kernel would be reduced by $(\frac{20}{4})^3$ or $(\frac{20}{4})^2$, respectively, or a speed gain of at least 25–125 and a memory saving of a factor of ~25. If a system and problem permitted the use of contracted functions, for example, many properties can be modeled satisfactorily in diamond—a favorable example for contracted basis functions—using 13 functions per atom, in which case the method presented here would be ~30 times quicker in the large system limit. However, as the underpinning basis could be far larger higher accuracy could also be achieved, another important benefit. It must be stressed, if a large high-quality basis was used, for example, 40 primitive functions per atom, the algorithm presented here would be 100–1000 times quicker than the conventional approach; and, remarkably, for large enough systems there would not be a significant change in the calculation time between 20 and 40 primitive functions per atom.

It should be noted that, for materials considered in this work, standard basis sets with four functions per atom, given names of the form STO-nG by the quantum chemistry community, do exist but results obtained using them are unpublishable today due to their very poor accuracy. The ideas presented here enable four functions to be generated employing a generalized and automated contraction scheme which give results comparable in accuracy to calculations using 20 or more functions per atom (for example, 6-311G** in quantum chemistry terminology), but achieving this with time and memory requirements comparable to the STO-3G basis set.

This work is organized as follows: Sec. II outlines the overall approach from a theoretical (rather than computational) point of view and places the method in the context of well-known conventional procedures, Sec. III gives a detailed account of the computational methodology to realize the theory of Sec. II in practice and Sec. IV presents tests of the accuracy of the method and performance results.

II. METHOD

The use of Gaussian orbitals in electronic-structure theory is well documented[6] including various contracted Gaussian basis sets for particular systems.[7] Typically, a basis set is chosen and the calculation performed yielding solutions that minimize the total energy within the space spanned by the basis set. We shall refer to this initial basis, the basis that defines the space within which the solutions can exist, as the "primitive" basis—which in turn spans the primitive space—whether or not this primitive set is formed from single Cartesian Gaussians or linear combinations of them. The primitive basis set will be denoted $\{\phi_i\}[i=1,\ldots,N]$. The Kohn-Sham wave function of level λ is then

$$\psi_\lambda(\mathbf{r}) = \sum_i c_{i\lambda}\phi_i(\mathbf{r}). \tag{1}$$

A. Conventional approach using only the primitive basis

Many approaches have been developed in the past few decades to solve the KSDFT equations using atom-centered functions. Here, we briefly outline the general framework to clearly place later discussions in context.

(1) Build the matrix Hamiltonian in the (primitive) basis $\{\phi_i(\mathbf{r})\}$

$$H_{ij} = \int \phi_i(\mathbf{r})\left(-\frac{1}{2}\nabla^2 + V(\mathbf{r})\right)\phi_j(\mathbf{r}) + \int\int \phi_i(\mathbf{r})V^{\mathrm{nl}}(\mathbf{r},\mathbf{r}')\phi_j(\mathbf{r}')d\mathbf{r}\ d\mathbf{r}', \tag{2}$$

where the nonlocal term [containing $V^{\mathrm{nl}}(\mathbf{r},\mathbf{r}')$] is only required when using norm-conserving pseudopotentials. $V(\mathbf{r})$ is the local potential comprising the local pseudopotential and the Hartree and exchange-correlation potentials. A more detailed account of the linear-scaling Hamiltonian build is given in Sec. III.

(2) Evaluate the overlap matrix,

$$S_{ij} = \int \phi_i(\mathbf{r})\phi_j(\mathbf{r}), \tag{3}$$

and form the generalized eigenproblem

$$\mathbf{Hc} = \mathbf{Sc}\boldsymbol{\Lambda}. \tag{4}$$

(3) Perform the kernel minimization step (this may be direct diagonalization or an iterative minimization[4,5]) and update the density matrix (or the state vectors),

$$b_{ij} = \sum_\lambda^N f(\boldsymbol{\Lambda}_{\lambda\lambda})c_{i\lambda}c_{j\lambda}, \tag{5}$$

and in turn the density,

$$n(\mathbf{r}) = \sum_{ij}^N b_{ij}\phi_i(\mathbf{r})\phi_j(\mathbf{r}), \tag{6}$$

which defines a new $V(\mathbf{r})$. Repeat steps (1)–(3) until self-consistency is obtained.

The algorithm presented here replaces step (3) above. Given the matrices **H** and **S** we first need to go from the primitive space to the contracted space to perform the diagonalization, then we need to return to the primitive space for the calculation of the density and, optionally, forces.

B. The contracted/filtered basis

The term contracted/filtered will be used to describe the basis construction outlined in the following sections to highlight the similarities (and differences) to conventional contracted basis sets.[7] We define the contracted/filtered basis set $\{\tilde{\phi}_I\}[I=1,\ldots,\tilde{n}]$ in terms of the underlying primitive basis set $\{\phi_i\}[i=1,\ldots,N]$

$$\tilde{\phi}_I(\mathbf{r}) = \sum_i k_{iI}\phi_i(\mathbf{r}); \quad I=1,\ldots,\tilde{n}. \tag{7}$$

Such an approach is conventionally used in Gaussian-type orbital calculations where the matrix of contraction coefficients $\{k_{iI}\}$ are derived from atomic calculations and are frozen throughout a total-energy calculation. In contrast, this work will introduce a procedure where the optimal set $\{k_{iI}\}$ is constructed automatically for the problem in hand without any reference to atomic calculations. This is done by filtering out unwanted high-energy components from the primitive basis to provide a much smaller basis that spans an energy window in the lower eigenspectrum of the Hamiltonian. The approach is based on the filter diagonalization algorithm of Neuhauser[8] and the Fermi operator projection method of Goedecker.[9]

First, the method is presented assuming an orthogonal primitive basis for simplicity before the generalization, and justification of efficacy, for nonorthogonal basis sets is described.

1. Filter diagonalization

Given an ordinary $n\times n$ eigenproblem (where n is very large)

$$\mathbf{Hc} = \mathbf{c}\boldsymbol{\Lambda} \tag{8}$$

imagine trying to calculate some selected eigenvalues in the interval $[\lambda_i,\lambda_f]$. One could form a function of the Hamiltonian $\mathbf{F}(\mathbf{H})$ that suppresses eigenvalue/eigenvector pairs outside the range of interest. This function could, for example, be a Gaussian in form or similar to a Fermi-Dirac function. In this work, as we are interested in the lowest states of the Hamiltonian, we employ a high-temperature Fermi-Dirac function. If we operate on a trial function $|t_k\rangle$ with this matrix function,

$$\mathbf{F}|t_k\rangle = \sum_i |i\rangle f(\lambda_i)\langle i|t_k\rangle = |t_k'\rangle, \tag{9}$$

where $\{|i\rangle,\lambda_i\}$ are the eigenpairs of **F**, then the resulting function $|t'\rangle$ will have been filtered such that its components predominantly lie in the space of the desired eigenvectors. This process is repeated until the space is sufficiently spanned (say by m filtered functions). The subspace of m functions (where $m \ll n$) is then used to construct an $m\times m$ eigenprob-

lem that is diagonalized and the answer obtained.

If the primitive basis is nonorthogonal, resulting in a generalized eigenvalue problem as in Eq. (4) then the filtration step in Eq. (9) simply becomes

$$\mathbf{c}f(\mathbf{\Lambda})\mathbf{c}^T\mathbf{S}|t\rangle = \mathbf{FS}|t\rangle = |t'\rangle. \tag{10}$$

The accuracy of this technique can be increased either by increasing the size of the trial set $\{|t_k\rangle\}$ or by defining a sharper filtration function (a narrower Gaussian or a lower temperature Fermi-Dirac distribution).

2. Imposing localization constraints

Until now the method has been presented without approximation–assuming the relevant subspace can be fully spanned during the filtration step. However, this algorithm would not in itself provide any computational saving in a Gaussian basis as the direct construction of the filtration matrix $\mathbf{F}$ is as expensive as calculating the density. However, a controlled approximation can be introduced. We can impose localization on the construction of $\mathbf{F}$. Therefore, for example, we may construct a filtration matrix for each atom $\mathbf{F}_\alpha[\alpha = 1,\ldots,N_{\text{atom}}]$ using only functions centered on surrounding atoms within a sphere of radius r_{cut}^α (the "filtration radius") centered on the atom (at $\mathbf{R}_\alpha$) in question. The filtration matrix in this work is a Fermi-Dirac DM at a very high temperature chosen for computational convenience and bears *no* relation to the electronic temperature of the calculation. Furthermore, detailed *a priori* knowledge of the chemical potential is not required. The high temperature of the filtration function assures that the functions will be localized. The higher the temperature the more localized the filtered functions are in real space, however, more degrees of freedom are required to span the space filtered by the function. The primitive basis must contain functions that have commensurate localization properties. Ideally, from the point of view of localization—though certainly not in terms of speed and memory efficiency—the basis would be a fine mesh of points in real space. However, we have found that typical, accurate Gaussian primitive basis sets containing a function with an exponent of ~0.1 a.u. to be also well suited to the task, as will be shown in Sec. IV. Once the contracted/filtered basis has been constructed then the subspace eigenproblem—in the space of these basis functions—can be constructed and solved. The computational details of this algorithm will be given in more detail in Sec. III.

Readers may be reminded, at this point, of the well-known divide and conquer algorithm[10,11] where the density (or density matrix) can be built up using local diagonalizations. However, it must be stressed that the method here bears no relation to the divide and conquer method and does not rely on the locality of the density matrix but rather the filtration matrix—and is therefore suitable for metallic systems at low electronic temperature as will be demonstrated in Sec. IV.

3. Subspace and nonorthogonality

When using a nonorthogonal basis set $\{\phi_i\}$, with an overlap matrix $\mathbf{S}$, one may define the dual set

$$\bar{\phi}_i(\mathbf{r}) = \sum_j S_{ij}^{-1}\phi_j, \tag{11}$$

such that

$$\int \bar{\phi}_i(\mathbf{r})\phi_j(\mathbf{r})d\mathbf{r} = \delta_{ij}. \tag{12}$$

When defining the DM, or filtration matrix, these dual complements must be used,[12]

$$F_{ij} = \int\int \bar{\phi}_i(\mathbf{r})F(\mathbf{r},\mathbf{r}')\bar{\phi}_j(\mathbf{r}')d\mathbf{r}\ d\mathbf{r}'. \tag{13}$$

It is known that the upper bound of the decay rate of $\mathbf{S}^{-1}$ goes like $\exp[-C/\kappa(\mathbf{S})]$ where C is a positive constant and $\kappa(\mathbf{S})$ is the condition number of $\mathbf{S}$.[13] Therefore, in general, a nonorthogonal basis set (especially high-quality Gaussian basis sets where κ is large due to approaching overcompleteness) strictly F_{ij} in such a basis does not inherit the locality of $F(\mathbf{r},\mathbf{r}')$. However, there is a different way of viewing this point. Consider a localized function $f(x)$ expanded in a localized basis set $\{\phi_i(x)\}$ with the expansion coefficients $\{x_i\}$ given by the solution of the linear equation

$$\sum_j S_{ij}x_j = y_i, \tag{14}$$

where

$$y_i = \int \phi_i(x)f(x)dx. \tag{15}$$

First, we consider the trivial case where the expansion is exact and given by a local expansion in the coefficient space. In this limit, clearly, if

$$\int \phi_i(x)f(x)dx = 0 \tag{16}$$

it must follow that $x_i=0$. This always follows if the basis is orthogonal (whether the function is exactly representable in the basis or not). Clearly, this result is independent of $\kappa(\mathbf{S})$. Indeed, we could construct $\mathbf{S}$ to be ill-conditioned with an upper bound on the decay of its inverse such that this matrix must always be considered dense. So, although $\mathbf{S}^{-1}$ may be dense the vector $|x\rangle$ can be sparse.

Now, if the function is not exactly represented by the basis, in general, the vector $|x\rangle$ is no longer sparse and all coefficients will be used to minimize

$$R(x) = \int \left[f(x) - \sum_i c_i\phi_i(x)\right]^2 dx. \tag{17}$$

However, for the current discussion, the interesting residual is

$$\bar{R}(x) = \int \left[\sum_i c_i\phi_i(x) - \sum_i \bar{c}_i\tilde{\phi}_i(x)dx\right]^2 dx, \tag{18}$$

where the set $\{\tilde{\phi}(x)\}$ is a subset of $\{\phi(x)\}$. Clearly, there is no obvious relationship between $\bar{R}(x)$ and $\kappa(\mathbf{S})$. The above argument is easily extended to bi-linear expansions [see

Eq. (13)] such as the density/filtration matrix.

In fact, as the (Gaussian) basis size is increased it will tend to become overcomplete both locally and globally. We may therefore expect $|\bar{R}(x)|^2$ to reduce as the basis size is increased and $\kappa(\mathbf{S})$ tends to infinity as we are approaching the limit where the function becomes exactly represented locally. This is the opposite of the argument presented in Ref. 12 and shows that the use a nonorthogonal basis cannot be discounted by such reasoning.

4. Choice of trial functions

Ideally the trial functions should be localized both in real space and the space of the Hamiltonian. We cannot know *a priori* which trial functions are maximally localized for both of these spaces, however we can make a reasonable guess. If we consider a free electron gas, as an approximation to our system, then the filtration step in Eq. (9) (where we consider the filtration function to be a zero-temperature Fermi-Dirac distribution with Fermi energy μ) is equivalent to setting to zero Fourier coefficients higher in energy that μ. So, a function localized in both real and Fourier space is a good choice for the trial function. This follows essentially the same argument given in Ref. 9 where this principle is also demonstrated. Therefore, given that our primitive basis is constructed from Gaussians (the ideal choice based on localization in real and Fourier space) we can use trial functions that are essentially unit vectors in the coefficient space of the primitive basis. We therefore choose the broadest Gaussian on each atom and construct trial functions from its different angular momentum components. Therefore, in the current implementation, our trial function is a unit vector in the space of the Gaussian primitives.

C. Forces

The two approximations that enter the calculation when going from the full primitive space to the contracted space are the extent to which the contracted set spans the correct space and the localization constraints imposed on the contraction coefficients. It must be stressed that in the current implementation, in the limit of large systems, the filtration step [calculation of $\{k_{il}\}$ in Eq. (7)] will always be a small fraction of the computational time as it scales linearly with system size whereas the kernel diagonalization scales cubically. However, being able to use a small an $r_{\mathrm{cut}}^{\alpha}$ as possible has the benefit that the size of system where the method becomes advantageous (the *crossover* point) is lower and should the diagonalization step ever be replaced with a low-complexity algorithm this particular crossover point would also be lower.

Strictly speaking, as the contraction coefficients depend on the atomic positions this introduces Pulay-type forces in addition to the usual Pulay forces one obtains from the use of an atom-centered primitive basis. Here, we argue that these additional Pulay forces are not significant; this point is corroborated by the results given in Sec. IV.

Pulay forces are usually associated with the additional forces that arise due to the dependence of a basis on atomic position. Here, we shall use a slightly more general definition and in the interests of clarity we shall define the term *additional forces*. Imagine the calculation of the force on a given atom $\tilde{\mathbf{F}}$ and that this force is in error to the *true* force (in the complete basis set limit) given by $\mathbf{F}$ such that $\mathbf{F}=\tilde{\mathbf{F}}+\mathbf{F}'$. The additional force $\mathbf{F}'$ may come from one of three sources. First, the error could be in the evaluation of the Hamiltonian matrix elements. This additional force is commonly referred to as aliasing. This is far less of an issue in methods employing Gaussian basis functions as the integration grid is independent of the basis size (a property not typical, in practice, of large systematic basis sets), and the evaluation of the Hamiltonian, being linear-scaling, becomes negligible for systems of more than $\sim$100 atoms (see timings given in Ref. 4). In practice, we may arbitrarily reduce the error in our matrix elements while not significantly affecting performance. Second, errors could come from the solution not being fully self-consistent. Without significant errors in the evaluation of the Hamiltonian matrix elements, in a given basis, and that self-consistency is demanded the only other additional force arises from the incomplete nature of the basis. We shall term these forces *incompleteness forces*. In this sense the usual Pulay forces evaluated in atom centered basis codes are an approximation to the incompleteness forces.

With this more general definition in mind we can therefore define incompleteness forces from a given basis with respect to another. Standard Pulay forces are an approximation to the incompleteness forces *from* the primitive basis *with respect to* the complete Hilbert space. The secondary incompleteness forces induced by the contraction procedure are incompleteness forces from the contracted/filtered basis with respect to the primitive basis. We shall term the incompleteness forces between the full and primitive space as primary additional forces and those between the primitive and contracted space as secondary additional forces.

It is easy to conceive of a situation where these secondary additional forces are zero, whereas the primary additional forces are significant. Imagine a self-consistent Hamiltonian $\mathbf{H}^{\mathrm{SCF}}$ constructed in the primitive space. We could take a unitary transformation of the eigenvectors of $\mathbf{H}^{\mathrm{SCF}}$, transform this Hamiltonian in the new basis, rediagonalize and calculate the forces once again and they would be mathematically equivalent to the previous force calculation. However, here we have introduced contraction coefficients that depend on atomic position—yet the secondary Pulay forces are zero by construction. This is precisely the limit that would be obtained by the algorithm presented here when the filtered basis functions of Eq. (7) completely span the occupied subspace. Numerical support showing that acceptable convergence to this limit is achievable using modest filtration radii is given in Sec. IV.

In practice, our numerical results indicate that accurate forces are obtained using much smaller filtration radii than may be expected from consideration of the total-energy convergence. A physical reason for the negligible effect of the secondary additional forces follows from considering the question: where does the majority of the Pulay force come from? One can ask this question a different way; if the primitive Gaussian orbitals were not moved with the atoms what would be the predominant reason for the subspace not span-

ning two similar structures? It seems reasonable to assume that the region around the nucleus would be most affected, therefore it is essential to track the movement of the atoms but not necessarily important to update the contraction coefficients.

Fundamentally, as long as the degrees of freedom in the contracted/filtered basis can allow for the correct relative energy with respect to a small structural change with constant contraction coefficients then standard Pulay forces are sufficient.

III. COMPUTATIONAL METHOD

A. Construction of contracted/filtered basis

This section concentrates on the additional computational steps required for the introduction of the contracted/filtered basis. It is assumed that the normal (primitive) Hamiltonian and overlap matrices have been evaluated. Section III D describes an algorithm for the efficient linear-scaling build of the Hamiltonian, which is also used in the primitive basis (or "usual") algorithm. Nevertheless, its speed allows for the advantages of the contracted basis algorithm to be seen for fairly small systems.

We loop over atoms, and construct the minimal (filtered) basis for each atom. The explanation below will concentrate on a single atom centered at $\mathbf{R}_\alpha$.

At this point it is assumed that the Hamiltonian and overlap matrices ($\mathbf{H}$ and $\mathbf{S}$, respectively) formed in the primitive basis are available.

(1) Take a sphere of radius r_{cut}^{α} centered at $\mathbf{R}_\alpha$. Atoms that lie within this sphere contribute their primitive basis functions for the construction of the minimal basis on atom α. Essentially, this procedure defines a set of numbers (of size n_α) in the range $[1,N]$ (or is a subset of the full primitive basis); let us call this set F_α.

(2) Take the n_α rows and columns of $\mathbf{H}$ and $\mathbf{S}$ determined by the set F_α (see Fig. 1).

(3) Solve the $n_\alpha \times n_\alpha$ eigenproblem formed by this set $\overline{\mathbf{H}}\overline{\mathbf{c}}=\overline{\mathbf{S}}\overline{\mathbf{c}}\overline{\mathbf{\Lambda}}$.

(4) Form specific columns of the matrix $\mathbf{f}(\overline{\mathbf{c}})\overline{\mathbf{S}}$ using,

$$\mathbf{f}_{:,\sigma}(\overline{\mathbf{c}}) = \overline{\mathbf{c}}\mathbf{f}(\overline{\mathbf{\Lambda}})\overline{\mathbf{c}}^T\overline{\mathbf{S}}\hat{\mathbf{e}}_\sigma, \tag{19}$$

where $\mathbf{f}(\overline{\mathbf{\Lambda}})$ is the filtration function—at present a high-temperature Fermi function (see Sec. II B)—of the eigenvalue matrix. As $n_\alpha \gg \tilde{n}_\alpha$ (where $\tilde{n}_\alpha$ is the number of contracted/filtered basis functions on atom α), Eq. (19) is most efficiently calculated as a series of matrix-vector multiplications rather than forming the matrix $\overline{\mathbf{c}}\mathbf{f}(\overline{\mathbf{\Lambda}})\overline{\mathbf{c}}^T\overline{\mathbf{S}}$ directly by matrix-matrix multiplications. In any case, in the large system limit this is not the time dominant part of the current algorithm.

(5) Map $\mathbf{f}_{:,\sigma}(\overline{\mathbf{c}})$ (with a space of size n_α) into the corresponding column of $\mathbf{k}$ (with a space of size N). This essentially decompresses the n_α length vector into a sparse representation of a vector of length N with the elements in the appropriate places to refer to the primitive basis.

The time dominant component from the above is step (3). However, the matrices are small–typically with a dimension of $n_\alpha = 500-1000$, which is independent of system size. Therefore, N_{atom} executions of steps 3 and 4 scale linearly with respect to system size with an operation count of $\sim 3N_{\text{atom}}n_\alpha^3$. Furthermore, as each atom is independent this is an extremely efficient operation when performed in parallel.

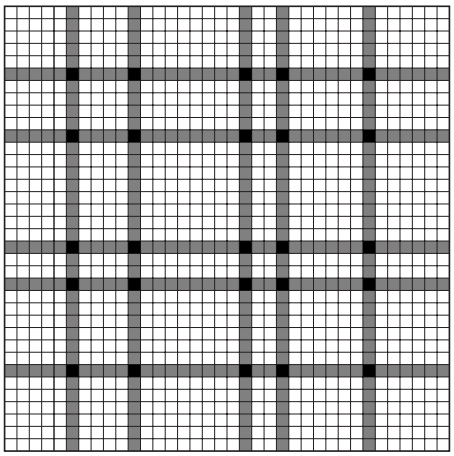

FIG. 1. Schematic of a primitive matrix with $N=36$. The shaded rows and columns are the rows and columns in the set F_α and $n_\alpha = 5$. Here, for example, $F_\alpha=\{6,11,20,23,30\}$. The black elements at the intersection of the rows and columns prescribed by F_α are the elements that form the $n_\alpha \times n_\alpha$ matrices $\overline{\mathbf{H}}$ and $\overline{\mathbf{S}}$.

In our implementation, this filtration step is performed at each stage in the self-consistent field (SCF) cycle thereby giving an optimum basis set at each stage in the calculation. Clearly, the process will also be repeated when atoms are moved, enabling the environment of each atom to be optimally described.

B. Primitive space to subspace transformation

When the above has been performed for every atom (the $\mathbf{k}$ matrix is fully constructed) the subspace eigenproblem can be formed using

$$\begin{aligned}\tilde{H}_{IJ} &= \int \tilde{\phi}_I(\mathbf{r})\hat{H}\tilde{\phi}_J(\mathbf{r})d\mathbf{r} \\ &= \sum_i^N \sum_j^N k_{iI}k_{jJ}\int \phi_i(\mathbf{r})\hat{H}\phi_j(\mathbf{r})d\mathbf{r} \\ &= \sum_i^N \sum_j^N k_{iI}k_{jJ}H_{ij},\end{aligned}$$

$$\tilde{\mathbf{H}} = \mathbf{k}^T\mathbf{H}\mathbf{k} \tag{20}$$

and a similar transformation for the overlap

$$\tilde{\mathbf{S}} = \mathbf{k}^T\mathbf{S}\mathbf{k} \tag{21}$$

giving the subspace eigenproblem

$$\tilde{\mathbf{H}}\tilde{\mathbf{c}} = \tilde{\mathbf{S}}\tilde{\mathbf{c}}\tilde{\mathbf{\Lambda}}. \tag{22}$$

This can then be solved using standard techniques such as direct diagonalization. Note, this subspace eigenproblem is

typically the size of a minimal basis calculation so, for example, carbon or silicon in the pseudopotential approximation would have only four functions per atom. The subspace density matrix $\tilde{b}_{IJ}=\Sigma_\lambda f(\tilde{\Lambda}_{\lambda\lambda})\tilde{\mathbf{c}}_{I\lambda}\tilde{\mathbf{c}}_{J\lambda}$ can now be calculated along with the band-structure energy

$$E_{\mathrm{bs}} = \sum_{IJ}^{\tilde{n}} \tilde{b}_{IJ}\tilde{H}_{IJ}. \tag{23}$$

C. Subspace to primitive space transformation

Once the density matrix is constructed in the subspace ($\tilde{\mathbf{b}}$) one can transform back to the primitive space using

$$\begin{aligned} n(\mathbf{r}) &= \sum_{IJ}^{\tilde{n}} \tilde{b}_{IJ}\tilde{\phi}_I(\mathbf{r})\tilde{\phi}_J(\mathbf{r}) \\ &= \sum_{IJ}^{\tilde{n}} \sum_{ij}^{N} k_{iI}k_{jJ}\tilde{b}_{IJ}\phi_i(\mathbf{r})\phi_j(\mathbf{r}) \\ &= \sum_{ij}^{N} b_{ij}\phi_i(\mathbf{r})\phi_j(\mathbf{r}). \end{aligned} \tag{24}$$

Equation (24) demonstrates that only elements of $\tilde{b}_{IJ}$ that are finite in $\tilde{\mathbf{S}}$ are required to construct elements of b_{ij} that are finite in $\mathbf{S}$. Therefore the primitive $\mathbf{b}$ can be calculated thus,

$$b_{ij} = \sum_{IJ}^{\tilde{n}} k_{iI}k_{jJ}\tilde{b}_{IJ} \tag{25}$$

$$\mathbf{b} = \mathbf{k}\tilde{\mathbf{b}}\mathbf{k}^T, \tag{26}$$

where $\tilde{b}_{IJ}$ is calculated and stored in same sparse format as $\tilde{\mathbf{S}}$, other elements if $\tilde{b}_{IJ}$ are ignored entirely. Once back in the primitive space the calculation proceeds as normal (e.g., updating the potential and Hamiltonian and calculation of forces). The subspace energy weighted density matrix $\tilde{w}_{ij}$ can be constructed in a similar fashion as $\tilde{\mathbf{b}}$ and w_{ij} generated using the same transformation in Eq. (26). The matrix multiplications in Eqs. (20), (21), and (26) are sparse due to the localization approximation outlined in Sec. II B 2 allowing the primitive space $\rightarrow$ subspace and subspace $\rightarrow$ primitive space transformations to be performed in a time that scales linearly with systems size. We employ a blocked compressed sparse column storage distribution and our optimized matrix products obtain approximately 25% level 3 BLAS performance. Only the diagonalization of the subspace eigenproblem (22) and the construction of $\tilde{\mathbf{b}}$ do not scale linearly with respect to the system size.

D. Implementation of efficient, linear-scaling Hamiltonian build

In this section we consider building the Hamiltonian matrix in a basis of Gaussian orbitals. The method presented here uses a hierarchical Fourier representation of the potential coupled with the locality of the primitive basis in real and Fourier space, resulting in a very efficient scheme. This is clearly also of importance if we wish to treat systems of 10 000 atoms or more.

Evaluation of the Hamiltonian matrix is of course essentially a problem of quadrature

$$H_{ij} = \int \phi_i(\mathbf{r})\hat{H}\phi_j(\mathbf{r})d\mathbf{r},$$

where the functions $\phi_i(\mathbf{r})$ are either Cartesian Gaussian functions

$$\phi_i(\mathbf{r}) = (x-R_{ix})^{n_x}(y-R_{iy})^{n_y}(z-R_{iz})^{n_z}\exp[-a_i(\mathbf{r}-\mathbf{R}_i)^2]$$

or the Bloch sums of these needed when periodic systems are being considered (the usual scenario in which the code is applied)

$$\phi_{i\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{L}} \phi_i(\mathbf{r}-\mathbf{L})e^{i\mathbf{k}\cdot\mathbf{L}},$$

where the vectors $\mathbf{L}$ are the lattice vectors of the system.

The simplest term in the Hamiltonian is the kinetic energy matrix which can be found using standard formulas.[14] The external potential in this work is provided by a pseudopotential[15] that consists of short ranged terms that can be easily be evaluated analytically (overlap terms weighted by powers of r); a long ranged local term (which may be treated in the same manner as the many body term below) and finally short ranged nonlocal terms that can be evaluated a sums of products of overlap matrices.

The final part of the Hamiltonian is the many body part, written in the local-density approximation or generalized gradient approximation (GGA) as a functional of charge density and terms involving its gradient. The charge density is trivially evaluated on a real-space grid (in linear scaling time) and the Hartree potential found using Fourier transforms. The exchange-correlation potential can also be evaluated point by point on the grid in the standard way; in the GGA this can be done either directly or using the approach of White and Bird.[16] We are thus left with a potential expressed in real space and the matrix elements of this

$$V_{ij} = \int \phi_i(\mathbf{r})\phi_j(\mathbf{r})V(\mathbf{r})d\mathbf{r}$$

can be readily found by summation over grid points on hierarchical real-space grids containing a set of Fourier curtailed representations of the potential in real space as is now described.

For purposes of discussion, we can write the potential as a Fourier series

$$V(\mathbf{r}) = \sum_{\mathbf{G}} V(\mathbf{G})e^{i\mathbf{G}\cdot\mathbf{r}} \tag{27}$$

so that our integral becomes

$$V_{ij} = \Omega\left(\frac{a_i + a_j}{\pi}\right)^{3/2} S_{ij} \sum_{\mathbf{G}} \phi_{ij}^*(\mathbf{G}) V(\mathbf{G}), \tag{28}$$

where $\phi_{ij}(\mathbf{G})$ is the Fourier transform of $\phi_{ij}(\mathbf{r}) = \phi_i \phi_j$, a_i is the decay exponent of $\phi_i(\mathbf{r})$ and Ω is the volume of the unit cell.

This sum is rapidly convergent as

$$\phi_{ij}^*(\mathbf{G}) \sim e^{-G^2/4(a_i+a_j)}$$

becoming converged when $|\mathbf{G}| = G_{ij}^{\text{cut}}$, where

$$G_{ij}^{\text{cut}} = \sqrt{4(a_i + a_j)(\tau - \sigma)}, \tag{29}$$

where $S_{ij} \sim e^{-\sigma}$ and τ determines the accuracy in that the truncation error is $e^{-\tau}$. An accurate and efficient linearly scaling evaluation of Eq. (28) is obtained by recasting the sum into real space

$$\sum_{\mathbf{G}} \phi_{ij}^*(\mathbf{G}) V(\mathbf{G}) = \frac{1}{N} \sum_{\alpha} \phi_{ij}(\mathbf{r}_\alpha) V_\alpha, \tag{30}$$

where N is the number of points in the real-space grid and V_α is *not* $V(\mathbf{r}_\alpha)$, but rather the real-space representation of its truncated Fourier expansion

$$V_\alpha = \sum_{|\mathbf{G}| < G_{ij}^{\text{cut}}} V(\mathbf{G}) e^{i\mathbf{G}\cdot\mathbf{r}_\alpha},$$

where G_{ij}^{cut} is determined from Eq. (29). Higher values of G_{ij}^{cut} require finer resolution grids in real space but sums are evaluated over correspondingly smaller spatial extents due to the decay of ϕ_{ij}. Therefore, the *number* of real-space integration points for *all* ϕ_{ij} is the same. This is a crucial point, leading to an immense increase in speed.

For the optimal evaluation of the sum in Eq. (30), representations of V_α are required on grids of different resolution for each term. This is clearly not feasible and so we use a fixed number of grids setup with resolutions varying roughly linearly between one point to the finest needed, and each term V_{ij} is evaluated on the nearest grid with higher resolution. Basis functions with higher angular momentum are trivially dealt with by inserting the Cartesian prefactors into the sum over grid points.

IV. PERFORMANCE AND RESULTS

As mentioned in Sec. I there are a number of different measures of convergence one can use to compare the conventional approach to the method outlined here, namely: (i) absolute energies, (ii) relative energies, and (iii) forces. Therefore, detailed results will be presented to reflect the convergence of each of these quantities. Broadly speaking two types of result are given. Results for smallish unit cells of up to 1000 atoms are first considered as such systems can be treated using conventional methods and this enables the accuracy of the filtration scheme to be monitored as a function of parameters such as the filtration radius. Second, results for much larger systems, up to 8000 atoms are given to illustrate the utility of the approach for such huge unit cells.

Uncontracted Cartesian Gaussian functions

TABLE I. Errors in the absolute and relative energies of two isomers of the molecule $C_{100}H_{200}Cl_2$ with a 5.6×10^{-3} a.u. (0.15 eV) energy difference.

Filtration radius (a.u.)	Absolute error energy A (a.u.)	Absolute error energy B (a.u.)	Relative error energy (a.u.)
3.5	4.7×10^{-1}	4.4×10^{-1}	2.4×10^{-2}
6.5	1.9×10^{-2}	1.6×10^{-2}	2.2×10^{-3}
9.5	3.4×10^{-3}	3.0×10^{-3}	3.8×10^{-4}

$$\phi_i(\mathbf{r}) = (x - R_{ix})^{n_x} (y - R_{iy})^{n_y} (z - R_{iz})^{n_z} \exp[-a_i(\mathbf{r} - \mathbf{R}_i)^2]$$

are used to form our primitive set. For each atom, we typically use four different exponents a_i and we multiply each Gaussian function $\exp[-a_i(\mathbf{r}-\mathbf{R}_i)^2]$ by the Cartesian prefactors, including all combinations of n_x, n_y, and n_z such that $n_x + n_y + n_z \leq \ell$. This produces four functions for $\ell = 1$ and ten functions for $\ell = 2$.

In our notation, exponents are arranged from lowest to highest (most diffuse Gaussian first) and the standard nomenclature is used to define the angular momentum. So, for example, a *ddpp* basis has four exponents with $10+10+4+4 = 28$ functions. Such a basis set applied to an atom such as carbon or silicon would be considered large for a routine quantum chemistry application. The *ddpp* basis sets used include diffuse functions. Their exponents are 0.16145, 0.46343, 1.31473, and 3.75324 for silicon and 0.1595905, 0.5947590, 1.2420759, and 2.7551283 for aluminum.

All results for periodic systems have been calculated for the Γ point unless otherwise mentioned. This is because this first implementation of the filtration algorithm has been written in real arithmetic. This means that results can only be compared to experiment for the largest unit cells—certainly this can only be done above 1000 atoms for silicon, and for a metal such as aluminum even 4000 atoms is not fully converged. This does not affect our timings or convergence rates in other respects. Most comparisons are made with a conventional reference calculation, performed with the same primitive basis set and not employing the filtration scheme.

Attention will be paid to the timings of the various components of the algorithm in the analysis that follows. We wish to make it clear that although many of the molecules and solids possess a high symmetry, no use has been made of this fact when performing the calculation.

Results for three types of systems are considered: molecules; a defect in a semiconductor and a defect in a metal. In all cases below the "temperature" of the Fermi-type filtration function was set to $kT = 0.1$ Ha. The positioning of the filtration Fermi level was set to be approximately equivalent to the eigenvalue of the lowest eigenvalue with an occupation lower than 10^{-4}.

A. Clusters

The first system to be considered in an alkane chain, $C_{100}H_{202}$ in which two hydrogen atoms are substituted by chlorine atoms to give $C_{100}Cl_2H_{200}$. The first test will be to

TABLE II. Timings, energy error, and force error for $C_{50}H_{102}$ alkane. Calculations performed on a single 2.2 GHz Opteron CPU.

Filtration radius (a.u.)	Filtration time (secs)	Subspace DM time (secs)	Abs. Error energy (a.u.)	Max error Force (a.u.)
∞	N/A	43.49	0	0
3.5	0.55	1.15	1.7×10^{-1}	5.2×10^{-3}
4.5	3.56	1.89	4.9×10^{-2}	2.6×10^{-3}
5.5	13.83	3.05	8.4×10^{-3}	8.9×10^{-4}
6.5	17.66	3.43	3.7×10^{-3}	1.9×10^{-4}
7.5	21.97	3.62	2.0×10^{-3}	8.6×10^{-5}

look at the energy difference between two isomers one of which has both chlorine atoms on attached to the same carbon; the other has them attached to neighboring carbon atoms. This energy difference is 5.6 mHa (0.15 eV) as calculated in our basis set using a conventional calculation.

Table I presents the deviations from a reference calculation, both for the absolute energies and the energy difference. It is seen that the absolute energy is converged to 10^{-5} Ha per atom for a cut-off radius $r_c=9.5$. A further challenge is to make sure that the difference of two energies is also reproduced. The final column of I shows that the 5.6 mHa (0.15 eV) difference is reproduced to ~0.01 eV.

Table II displays the accuracy of the forces obtained for a 152 atom $C_{50}H_{102}$ alkane with all positions displaced slightly from equilibrium. Again, fairly modest filtration radii yield acceptable levels of convergence—certainly good enough to permit structural relaxation. This is a demonstration of the validity of the arguments presented in Sec. II C. Also presented in the table are the relative times for performing the filtration step and the subspace kernel. It is clear that for such a small 152 atom system the time taken by the filtration process dominates the run.

B. Ideal vacancy in silicon

The system considered in this section is the ideal vacancy created by removing a single atom from cubic silicon unit cells containing 512, 1000, 1728, 2744, 4096, 5832, and 8000 atoms. All calculations presented in this section use the *ddpp* basis described above and were performed on 3.0 GHz Intel Xeon processors with Infiniband interconnects.

Table III compares the formation energy of the ideal vacancy, obtained from calculations on the 1000 atom unit cell, for different filtration radii with the answer using the full *ddpp* basis. As the subspace was filtered from the *ddpp* primitive basis this result defines the *exact* result within the context of this comparison. It is seen that the formation energy converges rapidly to the correct answer. Compared to the full primitive basis calculation, the magnitude of the error in the forces of the atoms surrounding the vacancy was 1.4×10^{-4} a.u. (~7 meV/Å) for a filtration radius of 12 a.u.

TABLE III. Convergence of the formation energy of the ideal vacancy in silicon with respect to filtration radius. Calculations were performed in 1000 atom cubic unit cells.

Filtration radius (a.u.)	5	8	10	12	14	*ddpp*
Formation energy (eV)	4.23	4.24	4.21	4.19	4.19	4.18

TABLE IV. Formation energy of the ideal vacancy in silicon from Γ-point calculations over a range of cubic-cell sizes. The filtration radius was fixed at 12 a.u. and the primitive basis was *ddpp*.

Number of atoms	512	1000	1728	2744	4096	5832	8000
Formation energy (eV)	4.12	4.19	4.22	4.23	4.24	4.24	4.24

Having demonstrated that a cut-off radius of 12 a.u. provides acceptable convergence, Table IV then shows the formation energy as a function of unit-cell size for this radius. The results demonstrate the stability of the underpinning Gaussian orbital method and the ability of the approach to maintain a consistently high level of accuracy over a large range of system sizes. In passing, the results also illustrate the large errors (~0.1 eV) incurred as a consequence of either using only the Γ point or defect-defect interactions in small unit cells containing only ~500 atoms. Table V gives a further check on the convergence of the formation energy with respect to filtration radius, this time for the larger 2744 atom system. The results follow the same trends as seen in the smaller 1000 atom cell discussed above. Clearly a 2744 atom unit cell is of a size that a conventional calculation is very demanding, and as such is not included in the table.

Moving the focus from accuracy to timing, Fig. 2 shows timings for a range of system sizes, again for the ideal vacancy in silicon. The timings pertain to the converged cut-off radius (12 a.u.). The number of processes used for each system is scaled in a manner approximately proportional to the number of atoms (16 processors for 1000 atoms; 64 for 4096 and so on), which would result in horizontal lines for an algorithm scaling linearly both with respect to the number of atoms and parallel speed-up. It is immediately clear that the algorithm is very fast, with only 200 s per SCF iteration even for large systems up to 2000–3000 atoms. The filtration step exhibits the expected excellent scaling properties and completely dominates the time for smaller systems (up to 3000 atoms).

The construction of the subspace DM is dominated by a Householder reduction (ScaLAPACK routine pdsytrd), which not only scales cubically with the number of atoms but also has less than perfect parallel efficiency. Between approximately 500 and 3000 atoms, however, linear scaling (in terms of both operation count and parallel efficiency) is

TABLE V. Convergence of the formation energy of the ideal vacancy in silicon with respect to the filtration radius. Calculations were performed using 2744 and 2743 atom cubic cells.

Filtration radius (a.u.)	8	10	12	14
Formation energy (eV)	4.28	4.25	4.23	4.23

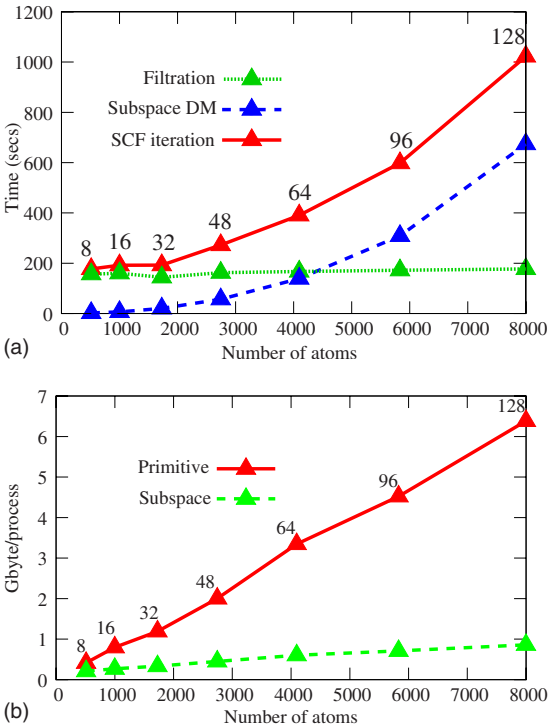

FIG. 2. (Color online) Upper: timings of the time dominant contributions to the SCF iteration for the ideal vacancy in silicon in a range of cubic unit cells. The numbers above the line give the number of cores used. Lower: a comparison of dominant (eigenproblem) memory consumption. The solid line is for the conventional algorithm and the dashed line relates to the subspace method presented in this work.

achieved and as a result only very moderate facilities are needed to perform these runs. The cubically scaling operations become time dominant between 6000 and 7000 atoms. It should be noted that the sparse matrix-matrix multiplications that are needed to transform between the primitive space and subspace [Eqs. (20) and (26)] are insignificant in terms of their contribution to the timings at all considered system sizes.

The lower panel in Fig. 2 displays the memory required, both for the conventional algorithm and the subspace kernel. It is immediately clear that the subspace approach produces a significant saving in memory over the conventional approach. The dashed line includes memory for storage of the sparse matrices $\mathbf{k},\mathbf{H},\mathbf{Hk},\tilde{\mathbf{b}}\mathbf{k}^T$ and the sub Hamiltonian $\tilde{\mathbf{H}}$. This memory scales linearly with system size and so should remain approximately constant on the graph. However, for convenience we have replicated the matrix $\mathbf{k}$ to simplify our implementation of the parallel sparse matrix computations. $\mathbf{k}$ is the smallest of the matrices involved, but its replication means it accounts for 336 Mbyte/process of the 699 Mbyte/process required for the 8000 atom system. A better implementation would reduce this to ~3 Mbyte/process. Thus, it is clearly seen that the filtration algorithm delivers an order of magnitude saving in memory as well as the increase in speed discussed above.

TABLE VI. Relative energy of relaxed structures with respect to the relaxed $I_3:T_d$ system in eV. The primitive basis was *ddpp* and results for filtration radii of 10 and 12 a.u. are shown.

Defect	10 a.u.	12 a.u.	*ddpp*
$I_3:C_{3v}$	0.478	0.479	0.476
$I_3:D_{2d}$	0.673	0.668	0.666
$I_3:W$	1.293	1.293	1.291
$I_3:C_2$	1.658	1.657	1.637

C. Interstitials in silicon

To further investigate the accuracy of the forces, and the efficacy of their use for structural relaxation, we have considered the structures and relative energies of tri-interstitial I_3 structures in silicon. These are both important technologically, with interstitials playing an important role in the processing required for the fabrication of electronic devices, but are also suitable for testing purposes as the bonding between the interstitial atoms and between them and the surrounding network is very different in the different possible structures. For example, some structures have all atoms fully fourfold coordinated, others possess dangling bonds. Many bond angles are greatly distorted from the tetrahedral angle, with angles varying between 60° and 120°.

The precise structures considered will be identified here according to their point group symmetry. The lowest energy form considered here is the compact structure with T_d symmetry, in which an atom in the ideal lattice is replaced by a tetrahedral structural structure of four silicon atoms. Two higher energy metastable structures considered have all atoms fourfold coordinated (albeit with very distorted bond angles) and have symmetries C_{3v} and D_{2d}.[17,18] A further candidate structure[19] is composed of split interstitials, has C_2 symmetry and a higher energy on account of having dangling bonds. The final structure considered is that previously ascribed to the W center.[17] The basis sets and calculational parameters are the same as those employed for the vacancy in the previous section. It should be noted that we are not claiming that these are the only, or even the lowest energy,

TABLE VII. Root mean square (rms) and maximum (max) errors in relaxed structures of various native defects in silicon. A filtration radius of 12 a.u. was used. The superscripts r, b, and θ refer to the errors in positions, bond lengths, and bond angles, respectively.

	Δ^r (mÅ)		Δ^b (mÅ)		Δ^θ (deg)	
Defect	rms	max	rms	max	rms	max
T_d	0.2	1.0	0.5	2.6	0.04	0.1
C_{3v}	0.2	0.9	0.5	0.5	0.07	0.1
D_{2d}	0.7	3.5	0.7	1.6	0.10	0.2
W	0.4	1.0	0.6	1.1	0.08	0.1
C_2	0.2	1.6	0.5	1.6	0.07	0.1

TABLE VIII. Convergence of formation energy of the ideal vacancy in Al (Γ point only) with respect to the filtration radius using a *ddpp* basis. Calculations were performed using 500 atom cubic unit cells. Fermi smearing with $kT=0.02$ eV was used in all these calculations.

Filtration radius (a.u.)	6	8	10	12	*ddpp*
Formation energy (eV)	0.33	0.36	0.41	0.40	0.40

structures for the I_3 defect—they are only chosen as candidates to give a more demanding test for the filtered basis set. These structures were all relaxed using the full primitive basis (the conventional algorithm), and filtered basis sets with cut-off radii of 10 and 12 a.u. These optimizations were performed in a 515 atom unit cell (512 atom cubic supercell plus the I_3 defect) chosen to make it possible to perform the same calculation with the conventional algorithm within the *ddpp* basis described above.

The resulting energies relative to the ground state structure are given in Table VI. It is seen that the relative energies are given for the 12 a.u. cutoff are correct to a typical accuracy of 1 meV, perfectly acceptable for this type of calculation—certainly such a small difference is one or two orders of magnitude smaller than the changes currently regarded as acceptable if a different pseudopotential or exchange correlation functional is used or if another parameter (for example, cell size, number of special k points) is changed.

In comparing the resulting structures, we have produced maximum absolute and root mean square (rms) differences in position $|\mathbf{R}-\mathbf{R}'|$ between the relaxed positions and the nearest-neighbor bond lengths and associated bond angles obtained with and without filtration. It is seen from Table VII that over all structures the maximum rms error in position is $<10^{-3}$ Å with a similar result for bond length and an rms error for bond angle being about 0.1°. We consider these very small differences and the associated results in the previous section provide numerical support to the assertions made in Sec. II C.

D. Ideal vacancy in aluminum

Finally, to demonstrate the method works equally well for metallic systems we consider the ideal vacancy in aluminum. This defect is created by removing a single atom from cubic unit cells containing 500 and 4000 atoms. Again, we first focus on accuracy, and Table VIII gives the calculated formation energy in a 500 atom cubic cell using only the Γ point. Clearly the use of Γ is not appropriate physically, but a direct comparison with a conventional calculation requires a small unit cell.

TABLE IX. Convergence of the formation energy of the ideal vacancy in aluminum with respect to the filtration radius. Calculations were performed using 4000 atoms cubic unit cells. Fermi smearing with $kT=0.02$ eV was used in all these calculations.

Filtration radius (a.u.)	8	10	12
Formation energy (eV)	0.59	0.61	0.60

It can be seen that the filtration radius needed to achieve a convergence in the formation energy of ~0.01 eV is both still rather small and similar to that used in silicon. Table IX shows the same set of data in a 4000 atom cell (again, a conventional calculation employing the full primitive basis calculation which would have been too computationally demanding to include).

The large difference in the values produced by the 500 and 4000 atom systems is a consequence of us using the Γ point. Nevertheless, Table IX shows that convergence with respect to the filtration radius is the same in a metal as in a semiconductor.

As a final demonstration of the power of the approach a 14^3 cell of aluminum (10 976 atoms) was modeled. The initial total energy on 256 cores (Intel Xeon 3.0 GHz) took only 138 min. To our knowledge, this is, to date, the largest calculation of its type for a cubic dense-packed metallic system. As far as we know, in fact, the only example in the literature to date of a calculation approaching this system size, using an algorithm capable of also treating metals, was presented in Refs. 20 and 21 where timings for large hydrogen terminated silicon clusters were presented.

V. CONCLUSION

A method based on the filter diagonalization algorithm with locality constraints has been implemented within the AIMPRO Gaussian orbital electronic structure code. Relative energies and forces of defects in both semiconducting and metallic systems were reproduced accurately. The method is capable of performing calculations of around 1000–2000 atoms in serial on a standard desktop computer and systems up to ~10 000 atoms can be tackled with modest computational resources. Due to the weak dependence of primitive basis size on performance and memory this also allows for far more accurate calculations to be performed. An extremely efficient linear scaling algorithm for the Hamiltonian build in a Gaussian basis was also presented.

ACKNOWLEDGMENTS

We thank J. P. Goss for providing us with specimen structures for the defects in silicon. M.J.R. also gratefully acknowledges the support of the Alexander von Humboldt Foundation.

[1] W. Kohn and L. J. Sham, Phys. Rev. **140**, A1133 (1965).
[2] P. Hohenberg and W. Kohn, Phys. Rev. **136**, B864 (1964).
[3] R. Jones and P. R. Briddon, Semicond. Semimetals **51A**, 288 (1998).
[4] M. J. Rayson and P. R. Briddon, Comput. Phys. Commun. **178**, 128 (2008).
[5] J. VandeVondele and J. Hutter, J. Chem. Phys. **118**, 4365 (2003).
[6] P. R. Briddon and R. Jones, Phys. Status Solidi B **217**, 131 (2000).
[7] J. P. Goss, M. J. Shaw, and P. R. Briddon, Top. Appl. Phys. **104**, 69 (2007).
[8] D. Neuhauser, J. Chem. Phys. **100**, 5076 (1994).
[9] S. Goedecker, J. Comput. Phys. **118**, 261 (1995).
[10] W. T. Yang, Phys. Rev. Lett. **66**, 1438 (1991).
[11] W. T. Yang and T. S. Lee, J. Chem. Phys. **103**, 5674 (1995).
[12] S. Goedecker, Rev. Mod. Phys. **71**, 1085 (1999).
[13] P. E. Maslen, C. Ochsenfeld, C. A. White, M. S. Lee, and M. Head-Gordon, J. Phys. Chem. A **102**, 2215 (1998).
[14] S. Obara and A. Saika, J. Chem. Phys. **84**, 3963 (1986).
[15] C. Hartwigsen, S. Goedecker, and J. Hutter, Phys. Rev. B **58**, 3641 (1998).
[16] J. A. White and D. M. Bird, Phys. Rev. B **50**, 4954 (1994).
[17] B. J. Coomer, J. P. Goss, R. Jones, S. Öberg, and P. R. Briddon, Physica B **273-274**, 505 (1999).
[18] M. Gharaibeh, S. K. Estreicher, and P. E. Fedders, Physica B **273-274**, 532 (1999).
[19] D. A. Richie, J. Kim, S. A. Barr, K. R. A. Hazzard, R. Hennig, and J. W. Wilkins, Phys. Rev. Lett. **92**, 045501 (2004).
[20] Y. Zhou, Y. Saad, M. L. Tiago, and J. R. Chelikowsky, Phys. Rev. E **74**, 066704 (2006).
[21] J. R. Chelikowsky, A. T. Zayak, T. Chan, M. L. Tiago, Y. Zhou, and Y. Saad, J. Phys.: Condens. Matter **21**, 064207 (2009).