Parallel Implementation of Γ -point Pseudopotential Plane-Wave DFT with Exact Exchange

Eric J. Bylaska¹

Environmental Molecular Sciences Laboratory, Pacific Northwest National Laboratory,
P.O. Box 999, Richland, Washington 99352

Kiril Tsemekhman

Department of Chemistry, University of Washington,
Box 351700, Seattle, WA 98195-1700

Scott B. Baden

Department of Computer Science and Engineering, University of California,
San Diego, La Jolla, California 92093

John H. Weare

Department of Chemistry and Biochemistry, University of California,
San Diego, La Jolla, California 92093

Hannes Jonsson

Faculty of Science, VR-II, University of Iceland,
107 Reykjavík, Iceland

¹ Corresponding author: eric.bylaska@pnl.gov

Abstract: Semi-local functionals commonly used in density functional theory (DFT) studies of solids usually fail to reproduce localized states such as trapped holes, polarons, excitons, and solitons. This failure is ascribed to self-interaction which creates a Coulomb barrier to localization. Pragmatic approaches in which the exchange correlation functionals are augmented with small amount of exact exchange (hybrid-DFT, e.g. B3LYP and PBE0) have shown promise in rectifying this type of failure, as well as producing more accurate band gaps and reaction barriers. The evaluation of exact exchange is challenging for large, solid state systems with periodic boundary conditions, especially when plane-wave basis sets are used. We have developed parallel algorithms for implementing exact exchange into pseudopotential plane-wave DFT program and we have implemented them in the NWChem program package. The technique developed can readily be employed in Γ -point plane-wave DFT programs. Furthermore, atomic forces and stresses are straightforward to implement, making it applicable to both confined and extended systems, as well as to Car-Parrinello *ab initio* molecular dynamic simulations. This method has been applied to several systems for which conventional DFT methods do not work well, including calculations for band gaps in oxides and the electronic structure of a charge trapped state in the Fe(II) containing mica, annite.

I. Introduction

A number of failures are known to exist in traditional implementations of plane-wave DFT, such as the underestimation of band gaps, the inability to localize excess spin density, and the underestimation of chemical reaction barriers. These problems are a consequence of having to rely on computationally efficient approximations to the exact exchange-correlation functional (e.g. LDA and GGA) used by plane-wave DFT programs. It is generally agreed upon that the largest error in these approximations is their failure to completely cancel out the orbital self-interaction energies, or in plain terms these approximations result in electrons partially seeing themselves^{1,2}. In the Hartree-Fock approximation, the exchange energy is calculated exactly and no self-interaction is present; however it is by construction a single particle theory and all correlation effects are missing from it. In all practical implementations of DFT the exchange energy is calculated approximately, and the cancellation of the self-interaction is incomplete.

It is not known how much importance to place on this extra interaction. Experience has shown that many of the failures associated with the erroneous self-interaction term can be corrected by approaches in which DFT exchange correlation functionals are augmented with some fraction of exact exchange (hybrid-DFT, e.g. B3LYP^{3,4} and PBE0⁵). These methods have been fairly successful in many problems of molecular chemistry and solid-state physics. For example they were able to yield localized charge states such as polarons, excitons⁶ and solitons, which have posed problems for conventional DFT⁶. They are also known to restore size consistency for systems where conventional DFT functionals are incorrect (e.g. dissociation of H_2^+ ion⁷), and they significantly improve the band gaps⁸⁻¹⁰ and spin structure¹¹ of solids, reaction barriers¹², and NMR shielding constants¹³.

The drawback of hybrid-DFT, however, is that it is significantly more expensive than conventional DFT (e.g. LDA, and GGA); for plane-wave methods hybrid-DFT requires the computation of $O(N_e^2)$ Fast Fourier Transforms (FFT), whereas conventional DFT only needs $O(N_e)$ FFTs, where N_e is the number of electrons. With the advent of new parallel machines, which are 100 to a 1000 times larger than current teraflop machines, these extra computational costs should easily be overcome. Scalable implementations of plane-wave DFT methods began appearing on hundreds of processors in the early to mid 1990s¹⁴⁻¹⁶ and improvements continue to this day¹⁷⁻²⁰. Notably (Gordon Bell Prize) F. Gygi *et al.*²⁰ have scaled a band structure calculation involving $N_e=12,000$ orbitals using 8 Brillouin zone sampling points on 64K nodes of Blue Gene L using the Qbox FPMD code. However, to date parallel algorithms for plane-wave DFT methods have focused only on conventional DFT, and not hybrid-DFT. This has led to a seemingly contradictory result: hybrid-DFT scales worse than conventional DFT on parallel machines²¹, even though the overall computational cost of hybrid-DFT is significantly higher and in principle should scale much better.

In this study, we present a parallel implementation of hybrid-DFT in a pseudopotential plane wave code based on a novel approach to the calculation of the exact exchange term, and illustrate its application to calculations for which conventional DFT does not work well. Our implementation of the hybrid-DFT method is explained in section II. The exact exchange term in this implementation is re-written in terms of maximally localized Wannier orbitals. A critical step in our implementation requires that the integration of the exact exchange integrals be limited to a single unit cell, while at the same time allowing for the treatment of extended systems. Section III presents a strategy for parallelizing the exchange term in plane-wave DFT. The parallel algorithm proposed here is based on the two-dimensional processor grid strategy

proposed by Gygi *et al.* for conventional DFT²⁰. However, our algorithm makes additional use of data replication to overcome the bottleneck associated with exact exchange term. In section IV, the hybrid-DFT method is applied to several systems for which conventional DFT methods do not work well, including calculations for the band gaps of oxides,, and the electronic structure of a charge trapped state in annite. Finally, conclusions are given in section V.

II. Hybrid-DFT Formalism Based on Maximally Localized Wannier Orbitals

The DFT energy and one-electron Hamiltonian (or gradient) with hybrid-DFT functionals for a set of $N_{occ}^{\uparrow} + N_{occ}^{\downarrow}$ unrestricted orbitals, $\{\psi_{\sigma,n}\}$ may be written as^{1,22}

$$E_{Hybrid-DFT} = \sum_{\sigma=\uparrow,\downarrow} \sum_{n=1}^{N_{occ}^{\sigma}} \int \psi_{\sigma,n}^*(\mathbf{r}) \left(-\frac{1}{2} \nabla^2 + V_l(\mathbf{r}) + \hat{V}_{NL} \right) \psi_{\sigma,n}(\mathbf{r}) d\mathbf{r} \\ + \frac{1}{2} \iint \frac{\rho(\mathbf{r})\rho(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r}d\mathbf{r}' + \alpha E_{x-exact}[\{\psi\}] + \beta E_x[\rho_{\uparrow}, \rho_{\downarrow}] + E_c[\rho_{\uparrow}, \rho_{\downarrow}] + E_{ion-ion} \quad (1)$$

$$H\psi_{\sigma,n}(\mathbf{r}) = \frac{\delta E_{Hybrid-DFT}}{\delta \psi_{\sigma,n}^*(\mathbf{r})} \\ = \left(-\frac{1}{2} \nabla^2 + V_l(\mathbf{r}) + \hat{V}_{NL} + V_H[\rho] + \beta V_x^{\sigma}[\rho_{\uparrow}, \rho_{\downarrow}] + V_c^{\sigma}[\rho_{\uparrow}, \rho_{\downarrow}] \right) \psi_{\sigma,n}(\mathbf{r}) \\ - \alpha \sum_{m=1}^{N_{occ}^{\sigma}} K_{nm}^{\sigma}(\mathbf{r}) \psi_{\sigma,m}(\mathbf{r}) \quad (2)$$

where α and β are constants. The total density $\rho = \rho_{\uparrow} + \rho_{\downarrow}$ is the sum of the spin densities given by

$$\rho_{\uparrow}(\mathbf{r}) = \sum_{n=1}^{N_{occ}^{\uparrow}} |\psi_{\uparrow,n}(\mathbf{r})|^2 \quad \text{and} \quad \rho_{\downarrow}(\mathbf{r}) = \sum_{n=1}^{N_{occ}^{\downarrow}} |\psi_{\downarrow,n}(\mathbf{r})|^2 \quad (3)$$

The local and non-local pseudopotentials, V_l and $\hat{V}_{NL}$ represent the electron-ion interaction, the double integral represents the Coulomb interaction between the electrons, and $E_{ion-ion}$ is the Coulomb interaction between the ions. $E_x[\rho_\uparrow, \rho_\downarrow]$ is the local exchange functional, $E_c[\rho_\uparrow, \rho_\downarrow]$ is the local correlation functional, $E_{x-exact}[\{\psi\}]$ is the exact exchange energy given by the standard expression²³

$$E_{x-exact}[\{\psi\}] = -\frac{1}{2} \sum_{\sigma=\uparrow, \downarrow} \sum_{n=1}^{N_{occ}^\sigma} \iint \frac{\rho_{nm}^\sigma(\mathbf{r}) \rho_{mn}^\sigma(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \quad (4)$$

And the exchange kernel $K_{nm}^\sigma(\mathbf{r})$ is

$$K_{nm}^\sigma(\mathbf{r}) = \int \frac{\rho_{mn}^\sigma(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r}' \quad (5)$$

where the overlap densities are given by

$$\rho_{nm}^\sigma(\mathbf{r}) = \psi_{\sigma,n}^*(\mathbf{r}) \psi_{\sigma,m}(\mathbf{r}) \quad (6)$$

The implementation of the exchange term in terms of plane-waves is straightforward for finite systems. The integrals

$$\iint \frac{\rho_{nm}^\sigma(\mathbf{r}) \rho_{mn}^\sigma(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \Rightarrow \iint_{\Omega} \frac{\rho_{nm}^\sigma(\mathbf{r}) \rho_{mn}^\sigma(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \quad (7)$$

are restricted to the volume Ω of one unit cell using the aperiodic or free-space boundary conditions procedure given in Appendix A. For extended systems, traditional implementations are done in terms of plane-wave expanded Bloch states²⁴

$$\psi_{\sigma,n\mathbf{k}}(\mathbf{r}) = \frac{e^{i\mathbf{k}\cdot\mathbf{r}}}{\sqrt{\Omega}} \sum_{\mathbf{G}} \psi_{\sigma,n\mathbf{k}}(\mathbf{G}) e^{i\mathbf{G}\cdot\mathbf{r}} \quad (8)$$

Substituting these states into Eq. 4 results in the following formula

$$E_{x-exact}[\{\psi\}] = -\frac{1}{2\Omega} \sum_{\sigma=\uparrow,\downarrow} \left(\frac{\Omega}{8\pi^3} \right)^2 \int_{BZ} d\mathbf{k} \int_{BZ} d\mathbf{l} \left[\sum_{n=1}^{N_{occ}^{\sigma}} \sum_{m=1}^{N_{occ}^{\sigma}} \sum_{\mathbf{G}} \frac{4\pi}{|\mathbf{G}-\mathbf{k}+\mathbf{l}|^2} \rho_{ml;n\mathbf{k}}^{\sigma}(-\mathbf{G}) \rho_{n\mathbf{k};ml}^{\sigma}(\mathbf{G}) \right] \quad (9)$$

where

$$\rho_{n\mathbf{k};ml}^{\sigma}(\mathbf{G}) = \sum_{\mathbf{G}'} \psi_{\sigma,n\mathbf{k}}^*(\mathbf{G}') \psi_{\sigma,n\mathbf{k}}(\mathbf{G}' + \mathbf{G}) \quad (10)$$

As pointed out by Gygi and Balderechi²⁵⁻²⁸ and others^{8,29-32}, this expression must be evaluated with some care especially for small Brillouin zone samplings and small unit cell size, because of the singularity at $\mathbf{G} - \mathbf{k} - \mathbf{l} = 0$.

In this work, we chose instead to implement exact exchange for periodic crystals using a formalism based on localized Wannier orbitals^{33,34}. The set of Wannier orbitals, $\{w_{\sigma,n}(\mathbf{r}-\mathbf{R})\}$, at different Bravais lattice vectors, $\mathbf{R}$, are related to the complete set of Bloch functions, $\{\psi_{\sigma,n\mathbf{k}}\}$, by the following unitary transformation over $\mathbf{k}$.

$$w_{\sigma,n}(\mathbf{r}-\mathbf{R}) = \frac{\Omega}{8\pi^3} \int_{BZ} e^{-i\mathbf{k}\cdot\mathbf{R}} \psi_{\sigma,n\mathbf{k}}(\mathbf{r}) d\mathbf{k} \quad (11)$$

where $8\pi^3/\Omega$ is the volume of the first Brillouin zone²⁴. The inverse Wannier transform can be readily written as a unitary transform over $\mathbf{R}$ and is

$$\psi_{\sigma,n\mathbf{k}}(\mathbf{r}) = \sum_{\mathbf{R}} e^{i\mathbf{k}\cdot\mathbf{R}} w_{\sigma,n}(\mathbf{r}-\mathbf{R}) \quad (12)$$

These orbitals form a complete and orthonormal set. In addition, since these orbitals are generated by a unitary transformation, the exact exchange energy, as well as other quantities such as the total energy and density, from this complete set is equal to that from the original Bloch orbitals. Similarly, any set of orbitals related to the Bloch orbitals by a unitary transformation over the occupied band orbital index n can be used in Eq. 11.

In general, a complete set of Bloch orbitals (complete sampling of the Brillouin zone) is needed to evaluate the complete set of Wannier orbitals (Eq. 11). However, for non-metallic systems with sufficiently large unit cells, it turns out that one can obtain at least one of the complete sets of Wannier orbitals contained in the manifold of the sets of Wannier orbitals from having only the Γ -point Bloch functions, $\psi_{\sigma,n\mathbf{k}=0}$ of all the occupied bands. The strategy for doing this is quite simple. The trick is to apply a Marzari-Vanderbilt localization unitary transformation (which is the counterpart of the Foster-Boys transformation for finite systems)³³⁻³⁵ over n to $\psi_{\sigma,n\mathbf{k}=0}$ to produce a new set of Γ -point Bloch functions to $\tilde{\psi}_{\sigma,n\mathbf{k}=0}$. These new orbitals, which are maximally localized, are extremely localized within each cell for non-metallic systems with a sufficiently large unit cells³⁴. If this is the case, then $\tilde{\psi}_{\sigma,n\mathbf{k}=0}$ can be represented as a sum of piece-wise localized functions, $\tilde{w}_{\sigma,n}(\mathbf{r}-\mathbf{R})$, on the Bravais lattice by

$$\tilde{\psi}_{\sigma,n\mathbf{k}=0}(\mathbf{r}) = \sum_{\mathbf{R}} \tilde{w}_{\sigma,n}(\mathbf{r}-\mathbf{R}) \quad (13)$$

Since Eq. 13 is just Eq. 12 with $\mathbf{k}=0$, $\tilde{w}_{\sigma,n}(\mathbf{r}-\mathbf{R})$ form a complete set of Wannier orbitals contained in the manifold of the sets of Wannier functions. With this new set of localized orbitals the exchange term per unit cell can be written as

$$E_{x-exact}[\{\tilde{w}\}] = -\frac{1}{2} \sum_{\sigma=\uparrow,\downarrow} \sum_{n=1}^{N_{occ}^\sigma} \sum_{m=1}^{N_{occ}^\sigma} \iint_{\Omega} \frac{\tilde{w}_{\sigma,n}^*(\mathbf{r}) \tilde{w}_{\sigma,m}(\mathbf{r}) \tilde{w}_{\sigma,m}^*(\mathbf{r}') \tilde{w}_{\sigma,n}(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|} d\mathbf{r} d\mathbf{r}' \quad (14)$$

To evaluate this integral in a plane-wave basis set requires some care, since representing overlap densities $(\tilde{w}_{\sigma,n}^*(\mathbf{r}) \tilde{w}_{\sigma,m}(\mathbf{r}))$ in Fourier space results in the inclusion of redundant periodic images (see Fig. 1). The interaction between these images, due to the long-range nature of the Coulomb potential, would result in an incorrect integration of the exchange term. Therefore, in order to use the plane-wave method, one has to limit integration in the exchange term to a single cell so that interactions between images are eliminated^{36,37}. To do this, we replace the standard Coulomb kernel in Eq. (14) by the following cutoff Coulomb kernel

$$f_{cutoff}(r) = \frac{1 - \left(1 - e^{-(r/R_c)^{N+2}}\right)^N}{r} \quad (15)$$

Where $r = |\mathbf{r}-\mathbf{r}'|$, and N and R_c are adjustable parameters. The kernel defined by Eq. 15 is shown in Fig. 2. This kernel rapidly decays to zero at distances larger than R_c . Hence Eq. 14 can be transformed to

$$E_{x-exact} = -\frac{1}{2} \sum_{\sigma=\uparrow,\downarrow} \sum_{n=1}^{N_{occ}^\sigma} \sum_{m=1}^{N_{occ}^\sigma} \iint \tilde{\psi}_{\sigma,n}^*(\mathbf{r}) \tilde{\psi}_{\sigma,m}(\mathbf{r}) f_{cutoff}(|\mathbf{r}-\mathbf{r}'|) \tilde{\psi}_{\sigma,m}^*(\mathbf{r}') \tilde{\psi}_{\sigma,n}(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \quad (16)$$

That is replacing $\tilde{w}_{\sigma,n}(\mathbf{r})$ with $\tilde{\psi}_{\sigma,n}(\mathbf{r})$, combined with using Eq. 15, in Eq. 16 will give the same energy, since the cutoff coulomb interaction is nearly $1/r$ with itself, and zero with its periodic images. The parameter R_c must be chosen carefully. It has to exceed the size of each Wannier orbital to include all of the orbital in the integration, while concurrently having $2R_c$ be smaller than the shortest linear dimension of the unit cell to exclude periodic interactions. Finally we note that

when one uses the cutoff Coulomb kernel, localized orbitals are not needed to calculate the exchange term since Eq. 16 can be unitary transformed, resulting in

$$\begin{aligned}
E_{x-exact}[\{\psi\}] &= -\frac{1}{2} \sum_{\sigma=\uparrow,\downarrow} \sum_{n=1}^{N_{occ}^{\sigma}} \sum_{m=1}^{N_{occ}^{\sigma}} \iint_{\Omega} \psi_{\sigma,n}^*(\mathbf{r}) \psi_{\sigma,m}(\mathbf{r}) f_{cutoff}(|\mathbf{r}-\mathbf{r}'|) \psi_{\sigma,m}^*(\mathbf{r}') \psi_{\sigma,n}(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \\
&= -\frac{1}{2} \sum_{\sigma=\uparrow,\downarrow} \sum_{n=1}^{N_{occ}^{\sigma}} \sum_{m=1}^{N_{occ}^{\sigma}} \left\{ \frac{\Omega}{2} \sum_{\mathbf{G}} |\rho_{\sigma,nm}(\mathbf{G})|^2 f_{cutoff}(|\mathbf{G}|) \right\}
\end{aligned} \tag{17}$$

While using the localized functions here is not required in this formulation, one should still evaluate the set of maximally localized Wannier functions in order to estimate their extent and consequently the minimal size of the unit cell. We note however, that the decay of the Wannier functions is a rather conservative estimate and a more reasonable measure of this decay is the extent of the localized overlap densities, which will decay at least twice as fast as the Wannier orbitals. The justification for this assertion can be obtained by recasting the above formulation (i.e. Eqs. 11-17) in terms of delocalized and localized overlap densities instead of molecular and Wannier orbitals.

To demonstrate the convergence of our proposed cutoff Coulomb procedure we have tested it on a theoretical periodic model of linear alkanes ($-\text{CH}_2-\text{CH}_2-$). This system was chosen, because it has previously been shown that the extent of the localized orbitals of a full Hartree-Fock solution is anywhere from 7 $-\text{CH}_2-$ units for a threshold of 10^{-3} to 29 $-\text{CH}_2-$ units for a threshold of 10^{-1038} . However, we expect that the exact exchange energy, Eq. 17, will be nearly converged by 7 $-\text{CH}_2-$ units, since by squaring the localized orbitals a threshold of only 10^{-6} will be seen for the localized overlap densities. The hybrid DFT calculations (PBE0, $E_{\text{cut}}=70$ Ry) were performed in simple cubic unit cells containing 2, 4, 6, 8 and 10 $-\text{CH}_2-$ units. The respective unit cells had dimensions of $2.5\text{\AA} \times 12\text{\AA} \times 12\text{\AA}$, $5.0\text{\AA} \times 12\text{\AA} \times 12\text{\AA}$, $7.5\text{\AA} \times 12\text{\AA} \times 12\text{\AA}$, $10\text{\AA} \times 12\text{\AA} \times 12\text{\AA}$, and $12.5\text{\AA} \times 12\text{\AA} \times 12\text{\AA}$. Within these cells, the carbon-carbon and

carbon-hydrogen bond lengths were set to 1.496Å and 1.095Å, and the C-C-C and H-C-H bond angles were set to 113.34° and 105.48°. For each of the 5 unit cells, calculations were performed using a series of 5 cutoff Coulomb kernels with increasing values of R_c ($N=8$ and $R_c=2, 4, 6, 8$, and 10 bohrs). In Fig. 3, the exact exchange energy per electron contribution of the hybrid exchange term and the value of the HOMO-LUMO gap for each of the 5 unit cells a function of R_c is shown. As expected, that both the exchange energy and band gap were found to converge very rapidly with respect to system size (8 $-CH_2-$ units) and cutoff Coulomb kernel size ($R_c=8$ bohrs).

We also note that for very large systems (>200 atoms) the overall cost of the computation can be reduced by using maximally localized Wannier functions since each individual maximally localized Wannier function only overlaps with a (in many cases small part) part of other Wannier functions in the same cell. Although this advantage is not discussed in this work, we explored it in our work and found that the amount of computation could be reduced significantly.

In Appendix B we present the formulae for the Fourier representation of Eq. 15, and in Appendix C, the formulae for analytical derivatives with respect to the lattice parameters.

III. Parallel Algorithm for Hybrid-DFT Calculations in a Plane-Wave Basis

During the course of total energy minimization or molecular dynamics simulation the electron gradient $\delta E_{Hybrid-DFT} / \delta \psi_{\sigma,n}^*$, (Eq. 3) needs to be calculated as efficiency as possible. For a pseudopotential plane-wave calculation the main parameters that determine the cost of a calculation are N_g , N_{occ} , N_a , and N_{proj} , where N_g is the size of the three-dimensional FFT grid, N_{occ} is the number of occupied orbitals, N_a is the number of atoms, and N_{proj} is the number of

projectors per atom. The evaluation of the gradient, when exact exchange is included, involves four major computational pieces with varying degrees of parallelism:

1. The Hartree potential V_H and the local exchange and correlation potentials $V_x + V_c$.
 - The main computational kernel in these computations is the calculation of N_{occ} three-dimensional FFTs
2. The exact exchange operator, $\sum_m K_{nm} \psi_m$
 - The major computational kernel in this computation involves the calculation of $(N_{occ}+1) \times N_{occ}$ three-dimensional FFTs.
3. The non-local pseudopotential, $\hat{V}_{NL}$
 - The major computational kernel in this computation can be expressed by the following matrix multiplications: $G = P^t * C$, and $C_2 = P * G$, where P is an $N_g \times (N_{proj} * N_a)$ matrix, C and C_2 are $N_g \times N_{occ}$ matrices, and G is an $(N_{proj} * N_a) \times N_{occ}$ matrix. We note that for most pseudopotential plane-wave calculations $N_{proj} * N_a \approx N_{occ}$.
4. Enforcing orthogonality
 - The major computational kernels in this computation are following matrix multiplications: $S = C^t * C$ and $C_2 = C * S$, where C and C_2 are $N_g \times N_{occ}$ matrices, and S is an $N_{occ} \times N_{occ}$ matrix.

Algorithm 1: Serial algorithm for calculating exact exchange in a plane-wave basis

Input: ψ - $N_g \times N_{occ}$ array

Output: $K\psi$ - $N_g \times N_{occ}$ array

for $m=1, N_{occ}$

 for $n=1, m$

 for $\mathbf{r}=1, N_g$ $\rho_{mn}(\mathbf{r}) = \psi_m(\mathbf{r})^* \psi_n(\mathbf{r})$

```

     $\rho_{mn}(\mathbf{G}) \leftarrow \mathbf{FFT\_rc}(\rho_{mn}(\mathbf{r}))$ 
    for  $\mathbf{G}=1, N_g$   $V_{mn}(\mathbf{G}) = f_{\text{screened}}(\mathbf{G}) * \rho_{mn}(\mathbf{G})$ 
     $V_{mn}(\mathbf{r}) \leftarrow \mathbf{FFT\_cr}(V_{mn}(\mathbf{G}))$ 
    for  $\mathbf{r}=1, N_g$   $K\psi_m(\mathbf{r}) = V_{mn}(\mathbf{r}) * \psi_n(\mathbf{r})$ 
    if  $m < n$  for  $\mathbf{r}=1, N_g$   $K\psi_n(\mathbf{r}) = V_{mn}(\mathbf{r}) * \psi_m(\mathbf{r})$ 
  end for
end for

```

The computation of the exact exchange operator, $\sum_m K_{nm} \psi_m$, with a cost of $O(N_{occ}^2 * N_g * \log(N_g))$, is by far the most demanding term in a pseudopotential plane-wave hybrid-DFT calculation. A simple serial algorithm for calculating exact exchange operator is given in Algorithm 1. In this algorithm the routines **FFT_rc**, and **FFT_cr** are defined to be the forward and reverse three-dimensional Fast Fourier Transforms (FFTs). The cost of this algorithm is quite high, since it requires the computation of $N_{occ} * (N_{occ} + 1)$ three-dimensional FFTs. For a modest size problem containing $N_{occ}=200$ filled orbitals, the above algorithm will need to compute $\sim 40,000$ three-dimensional FFTs.

There are several possible ways to parallelize this algorithm. A straightforward way is to do a spatial decomposition of the orbitals, ψ_n , as shown in Fig. 4a. This approach is versatile, and straightforward to implement. However it requires using parallel three-dimensional FFTs, whose parallel efficiencies are effectively limited to $\sim N_g^{1/3}$ processors (or processor groups). Another way to parallelize this algorithm is to distribute the orbitals across the processors as shown in Fig 4b. An advantage of this type of distribution is that parallel FFTs are not required. Instead the parallelization is done by moving the orbitals to set up the $O(N_{occ}^2)$ wave-function pairs. In particular, a multicast of $O(N_{occ})$ orbitals across the processors (or processor row) followed by a multi-reduction which reverses the pattern. We note that with this type of algorithm one could readily fill a very large parallel machine by assigning each a few FFTs to each processor. However, in order to get this type of algorithm to work it will be vital to mask latency, since the

interconnects between the processors will be flooded with $O(N_{occ})$ streams, each on long messages comprising N_g floating point words of data.

Algorithm 2: Parallel algorithm for calculating exact exchange in a plane-wave basis using a two-dimensional processor grid

Input: ψ - $(N_g/Np_i) \times (N_{occ}/Np_j)$ array

Output: $K\psi$ - $(N_g/Np_i) \times (N_{occ}/Np_j)$ array

Work Space: Ψ - $(N_g/Np_i) \times N_{occ}$ array

$K\Psi$ - $(N_g/Np_i) \times N_{occ}$ array

Np = total number of processors, where $Np = Np_i * Np_j$

Np_i = size of column processor group

Np_j = size of row processor group

$taskid_i$ = rank along the column processor group

$taskid_j$ = rank along the row processor group

$\Psi \leftarrow 0$

$K\Psi \leftarrow 0$

Gather ψ onto Ψ

Row_Reduction(Ψ)

counter = 0

for $m=1, N_{occ}$

for $n=1, m$

if counter == $taskid_j$ then

for $r=1, (N_g/Np_i)$ $\rho_{mn}(\mathbf{r}) = \Psi_m(\mathbf{r})^* \Psi_n(\mathbf{r})$

$\rho_{mn}(\mathbf{G}) \leftarrow \text{Column_FFT_rc}(\rho_{mn}(\mathbf{r}))$

for $\mathbf{G}=1, (N_g/Np_i)$ $V_{mn}(\mathbf{G}) = f_{\text{screened}}(\mathbf{G}) * \rho_{mn}(\mathbf{G})$

$V_{mn}(\mathbf{r}) \leftarrow \text{Column_FFT_cr}(V_{mn}(\mathbf{G}))$

for $r=1, (N_g/Np_i)$ $K\Psi_m(\mathbf{r}) -= V_{mn}(\mathbf{r}) * \Psi_n(\mathbf{r})$

if $m > n$ for $r=1, (N_g/Np_i)$ $K\Psi_n(\mathbf{r}) -= V_{mn}(\mathbf{r}) * \Psi_m(\mathbf{r})$

end if

counter = mod(counter + 1, Np_j)

end for

end for

Row_Reduce($K\Psi$)

Scatter $K\Psi$ back to $K\psi$

In our current implementation we chose to parallelize the exchange algorithm by distributing both the spatial and orbital dimension over a two-dimensional processor grid as shown in Fig 4c. This type of distribution is slightly more complicated in that it requires both performing parallel three-dimension FFTs along the processor grid columns, as well as broadcasting the orbitals along the processor grid rows. A parallel algorithm for calculating the exchange operator with this distribution is given in Algorithm 2. In this algorithm, the routines **Column_FFT_rc**, and **Column_FFT_cr** are the forward and reverse parallel 3-dimensional Fast Fourier Transforms (FFTs) along the processor columns, and **Row_Reduction** is a parallel reduction along the processor rows. This algorithm is very simple to implement, and it is perfectly load balanced since each CPU only computes $(N_e*(N_e+1)/N_p)$ three-dimension FFTs. However, this simple algorithm can be improved. One problem is that it uses a lot of workspace. Another is that each CPU in the **Row_Reduction** subroutine receives and sends a total of $(Npj-1)*(N_g/Npi)*(N_e/Npj) \sim N_g N_e / Npi$ amount of data, which is approximately twice as much as is necessary to compute all pairs of wave functions.

A slightly more ambitious parallel algorithm for calculating the exchange operator, which halves the workspace and communication overhead, while maintaining a good load balance, is given in Algorithm 3. In this algorithm the communication between CPUs is performed in the **Brdcst_Step** and **Reduce_Step** subroutines. These subroutines basically implement an asynchronous radix-2 butterfly diagram except that instead of transferring a $2^{**}(Levels-1)*(N_g/Npi)$ chunk of data at the last step, i.e. $l=(Levels-1)$, in the **for** loop, they transfer only a $(N_g/Npi)(N_e/Npj)(1+(Npj-2^{**}Floor(Log_2(Npj)))) \sim N_g N_e / N_p$ chunk data, instead of a $(Npj/2)N_g N_e / N_p$ chunk of data. The following broadcast and reduction routines implement this special radix-2 butterfly algorithm.

We note that compared to a full broadcast or reduction, these subroutines receive and send only a total of $\sim(Npj/2)*(N_g/Npi)*(N_e/Npj)=N_gN_e/(2Npi)$ chunks of data. The partial exact exchange computations are written using the **Index_Compute** (shown in the Appendix) and **HFX_Compute** subroutines. The **Index_Compute** subroutine is used to set up the index ranges used to compute the exact exchange terms. The **Exchnng_Compute** subroutine is used to compute the three-dimensional FFTs and point-wise products. This subroutine is computationally very expensive since it computes $(j2-j1+1)*(i2-i1+1)$ (or $(j2-j1+1)(i2-i1+2)/2$, if $i1=j1$ and $i2=j2$) three-dimensional FFTs. Thus it is important to make the various calls to this subroutine as load balanced as possible. It is trivial to make the first $\log_2(Npj)$ calls to **HFX_Compute** subroutine compute exactly the same number three dimensional FFTs on each CPU. However, the indexing for computing the remaining pairs of wave functions after the last special butterfly step requires some special logic. The **Index_Compute** subroutine has three different options. The first option is used for determining the indexes for the first 0 to $(Levels-1)$ steps. The second and third options are used for determining the index ranges to the **HFX_Compute** subroutine after the special last butterfly broadcast step.

Algorithm 3: Incomplete butterfly parallel algorithm for calculating exact exchange, in a plane-wave basis using a two-dimensional processor grid. The parameters N_e , N_p , N_{pi} , N_{pj} , $taskid_i$ and $taskid_j$, as well as the input and output are the same as in Algorithm 2 (subroutines **Brdcst_Step, **Index_Compute**, and **Exchng_Compute** are given in Appendix D)**

Work Space: $\Psi - (N_g/N_{pi}) \times (wsiz)$ array
 $K\Psi - (N_g/N_{pi}) \times (wsiz)$ array
 $wsiz = (N_e/N_{pj}) * (N_{pj}/2 + 1 + (N_{pj}-2 * \text{Floor}(\text{Log}_2(N_{pj}))))$

```

Ψ(:,1:(Ne/Npj))= Ψ̃; KΨ=
Levels = Floor(Log2(Npj))
for l=0,(Levels-1)
  Brdcst_Step( l, taskid_j, Npj, Ψ, Ng, Ne, requests, reqcnt)
  Index_Compute( l, l, taskid_j, Npj, i1, i2, j1, j2, im, it)
  Exchng_Compute( i1, i2, j1, j2, im, it, Ψ, KΨ)
  Row_WaitAll(requests, reqcnt)
end for
Index_Compute(2, l, taskid_j, Npj, i1, i2, j1, j2, im, it)
if (j2 >= j1) then
  Exchng_Compute(i1,i2,j1,j2,im,it,Ψ,KΨ)
end if
Index_Compute(3,l,taskid_j,Npj,i1,i2,j1,j2,im,it)
Exchng_Compute(i1,i2,j1,j2,im,it,Ψ,KΨ)
for l=(Levels-1),0
  Reduce_Step(l,taskid_j,Npj,KΨ,Ng,Ne,requests,reqcnt)
end for
KΨ += KΨ(:,1:(Ne/Npj))

```

We have implemented both Algorithms 2 and 3 in the NWPW module of the NWChem 5.1 program package³⁹. The overall and major constituent timings for hybrid-DFT calculations of an 80 atom supercell of hematite (Fe₂O₃) with an FFT grid of $N_g=72^3$ (wavefunction cutoff energy = 100 Ry and density cutoff energy = 200 Ry) and orbital occupations of $N_{occ}^\uparrow = 272$ and $N_{occ}^\downarrow = 272$ are shown in Table I at various combinations of processor sizes (N_p) and processor grids (N_{p_j}) for Algorithm 2. The best overall timings with Algorithm 2 at each processor size (N_p) are shown in Fig. 5 as well as in bold face in Table I. These calculations were performed

on a quad-core Cray-XT4 system (NERSC Franklin) composed of a 2.3 GHz single socket quad-core AMD Opteron processors (Budapest). The NWChem program was compiled using Portland Group FORTRAN 90 compiler, version 7.2.4, and linked with the Cray MPICH2 library, version 3.02, for message passing.

The overall performance of our hybrid-DFT calculations using Algorithm 2 is found to fairly reasonable for this size problem. The best time per step takes 1,190 seconds (20 minutes) for 16 processors down to 21.2 seconds with an 88% parallel efficiency on 1024 processors and 12.3 seconds with a 75% parallel efficiency on 2048 processors. As shown in Table I, these timings are found to be very sensitive to the layout of the two-dimensional processor grid. For 16, 32, 64, 128, 256, 512, 1024, and 2048 processors, the best timings are found with 16x1, 32x1, 32x2, 32x4, 32x8, 32x16, 32x32, and 32x64 processor grids, respectively. For illustrative purposes, we also report in Table I and 5.4 the largest contributions to the total electronic update step, including the exact exchange operator, the non-local pseudopotential operator, and orthogonality. As expected the exact exchange computation dominates the overall timing at all the processor sizes used in our study and it is scaling to large processor sizes. In addition, these results show that the non-local pseudopotential operator computations are also scaling very well. Unfortunately, the orthogonality computations are not scaling very well at large processor sizes, with no speedups found beyond 100 processors. Since this computation takes less than a second per step it is not dominating the overall timings at the processor sizes used. However, by 2048 processors it is taking 10% of the overall computation.

We also checked the performance of Algorithm 3 for the hybrid-DFT calculations of an 80, as well as a 160, atom supercell of hematite (Fe_2O_3) (see Fig. 6). We found that the parallel efficiencies beyond a 1000 CPUs were appreciably better than for Algorithm 2. With this

algorithm the overall best timing per step found for the 80 atom supercell was 3.6 seconds on 9792 CPUs, and for the 160 atom supercell of hematite was 17.7 seconds on 23,936 CPUs. The timings results for the 160 atom supercell are somewhat uneven, since limited numbers of processor grids were tried at each processor size. However, even with this limited amount of sampling, these calculations were found to have speedups to at least 25,000 CPUs. We expect that further improvements will be obtained by trying more processor geometry layouts.

IV. Hybrid-DFT Calculations

All DFT and hybrid-DFT calculations in this study were performed using the pseudopotential plane-wave program (NWPW module) contained in the NWChem computational chemistry package³⁹. Both the gradient corrected PBE96⁴⁰ and hybrid PBE0⁵ exchange-correlation potentials were used. The valence electron interactions with the atomic core are approximated with a generalized norm-conserving Hamann⁴¹ or Troullier-Martins⁴² pseudopotentials modified to a separable form suggested by Kleinman and Bylander⁴³. The hybrid-DFT calculations were performed using the screened Coulomb kernel with parameters $R_c=8$ bohrs and $N=8$.

For our initial application, we examine the reaction and activation energies for several reactions involving silanes using DFT and hybrid-DFT.

This reaction set is non-trivial as it involves breaking of Si-H, Si-Si and H-H bonds, along with significant electronic and structural rearrangement. Reactions containing such dramatic changes in bonding are rarely able to give reliable results in current implementations of DFT. Not surprisingly, prior DFT calculations reported by Nachtigall *et al.* showed a wide variance in predicted reaction and activation energies¹². Moreover, their results showed that exchange-correlation functionals usually employed in plane-wave DFT calculations are not reliable. Gas-phase reaction energies were determined for all the reactions (13)-(16) and activation energies were calculated for reactions (13) and (14). The electronic wavefunctions were expanded in a plane-wave basis with a cut-off energy of $E_{\text{cut}}=70\text{Ry}$. The supercell used in these calculations was a face-centered-cubic lattice with a lattice constant of $a=20.1\text{\AA}$. In Fig. 7 the structural parameters for H_2 , SiH_2 , SiH_4 , and Si_2H_6 , as well as for the transition state species $\text{SiH}_4\text{-TS}$, $\text{Si}_2\text{H}_6\text{-TS1}$, and $\text{Si}_2\text{H}_6\text{-TS2}$ calculated at the DFT and hybrid-DFT levels are reported, and in Table II the reaction and activation energies for the silanes reactions (13)-(16) involving silanes calculated at the levels of DFT and hybrid-DFT theory are reported. Comparisons are made to B3LYP, MP2, and E-QCISD(T) calculations reported by Nachtigall *et al.*¹². For the reaction energies very little difference was found between the PBE96 and PBE0 calculations with the average absolute differences from E-QCISD(T) were found to be 4.2 kcal/mol and 3.6 kcal/mol for PBE96, and PBE0 methods respectively. For the activation energies average the PBE0 calculations increase the height of the reaction barriers and the absolute differences from E-QCISD(T) were found to be 11.0 kcal/mol and 5.3 kcal/mol for PBE96 and PBE0 methods respectively.

For our next application, we calculate band gaps in insulators using DFT and hybrid-DFT. It is well known that DFT calculations using local and semi-local exchange-correlation

functionals underestimate band gaps for semiconductors and insulators by as much as 60%. The main source of error has been ascribed to the lack of a discontinuity at the Fermi level in exchange-correlation potentials of only the density. This type of exchange-correlation functional contain at least some form of the discontinuity since the effective potential is qualitatively different when going from occupied states to unoccupied states. In Table III, the band gaps of Al_2O_3 , TiO_2 , and SiO_2 at the conventional PBE96 and hybrid PBE0 DFT levels are reported. As one can see the band gaps at the hybrid-DFT level are considerably better than at the conventional DFT level. The band gaps were estimated by taking the difference between the HOMO and LUMO single particle energies at the Γ -point using large supercells. With this approach only the minimal gaps at the Γ -point or at any point folded into the Γ -point were obtained. However, since large unit cells were used, many of the high-symmetry points in the Brillouin zone were covered. In all these calculations a cutoff energy of 100 Ry was used, and the sizes of the supercells were 80, 72, and 72 atoms respectively for the three crystals. The default pseudopotentials of NWChem were used, except for the Ti. It was found that to obtain reliable structures for TiO_2 the Ti pseudopotential needed to include 3s and 3p states in the valence.

Finally we illustrate our hybrid-DFT formalism by applying it to charge trapped states in annite, an Fe(II) containing trioctahedral mica. There is evidence that oxidation of annite results in mobile charge states inside the single-particle band gap, associated with localized Fe^{3+} ion states. Rosso and Ilton have shown using small cluster models within Hartree-Fock approximation that removing an electron from the Fe^{2+} octahedral layer results in a localized Fe^{3+} state which is fairly mobile. A major failing of commonly used DFT functionals, such as PBE96, is their inability to model these types of charge trapped states. This by-now well known

failure is a byproduct of the self-interaction contained in many functionals, with self-interaction creating an artificial Coulomb barrier to charge localization. In most cases, this barrier overly favors delocalized electronic states, even for systems where significant charge localization is expected. In Fig. 8, we present results for the oxidation of annite at the conventional PBE96 and hybrid PBE0 DFT levels are reported. In these calculations a cutoff energy of 100 Ry was used, and the supercell contained 88 atoms ($\text{K}_4\text{Fe}_{12}\text{Si}_{12}\text{Al}_4\text{O}_{48}\text{H}_8$). Default pseudopotentials of NWChem were used. These calculations show that when an electron is removed from the conventional PBE96 calculation, the spin down density is completely delocalized over the octahedral layer. That is the oxidized octahedral layer in this calculation consists of $12\text{Fe}^{+2.0833}$ rather than $11\text{Fe}^{2+} + 1\text{Fe}^{3+}$. On the other hand, a significant degree of charge localization was seen in the hybrid PBE0 calculation.

V. Conclusions

We have developed a parallel algorithm for implementing exact exchange into Γ -point pseudopotential plane-wave density functional theory, and we have used this new development to perform hybrid-DFT calculation on several systems for which conventional DFT methods do not work well including reaction barriers in silane systems, band gaps in oxides and the electronic structure of a polaron in hematite. The technique developed can readily be employed in plane-wave DFT programs. Our implementation uses a formalism based on maximally localized orbitals, which allows the integration over the infinite periodic space to be limited to a single unit cell. This approach makes atomic forces trivial to implement. Lattice stresses are

also straightforward to implement, making it applicable to both confined and extended systems, as well as Car-Parrinello molecular dynamics simulations⁴⁴.

We have presented two different parallel algorithms (Algorithm 2 and Algorithm 3) for computing exact exchange that are based on using a two dimensional processor grid. A unique aspect of Algorithm 3 is the development of an incomplete butterfly that halves the amount of data communicated, as well as making judicious use of data replication, in the exact exchange compared to a standard Broadcast algorithm (Algorithm 2). For less than 1,000 CPUs both these algorithms were found to have similar performances, however, beyond 1,000 CPUs Algorithm 3 performed appreciably better. For CPU sizes beyond a 1,000 CPUs the overall performance of our hybrid-DFT calculations using Algorithm 3 was found to be fairly reasonable. For an 80 atom Fe_2O_3 hybrid- DFT calculation (544 valence electrons), the fastest time per step found was 3.6 seconds on 9792 CPUs, and for the 160 atom supercell (1088 valence electrons) of hematite was 17.7 seconds on 23,936 CPUs.

We note that the parallel algorithms developed in this work can also be readily applied to other standard plane-wave formulations of exact exchange^{29,30,32}, screened exchange⁴⁵, or screened exchange hybrids such as HSE^{8,46,47}. Finally, while incorporating exact exchange in a pseudopotential plane-wave calculation is still indeed computationally challenging, we have shown that even on modest size parallel machines, pseudopotential plane-wave hybrid-DFT simulations containing 300 orbitals at a cutoff energy for the density of 200 Ry can be computed in roughly 10 seconds per step. With the advent of new parallel machines along with additional algorithmic refinements and code optimizations, we expect these timings can be further improve in the very near future.

Acknowledgements

This research was supported by the ASCR Multiscale Mathematics program and the BES Geosciences program of the U.S. Department of Energy, Office of Science ~ DE-AC06-76RLO 1830. The Pacific Northwest National Laboratory is operated by Battelle Memorial Institute. We wish to thank the Scientific Computing Staff, Office of Energy Research, and the U. S. Department of Energy for a grant of computer time at the National Energy Research Scientific Computing Center (Berkeley, CA). Some of the calculations were performed on the Chinook and Spokane computing systems at the Molecular Science Computing Facility in the William R. Wiley Environmental Molecular Sciences Laboratory (EMSL) at PNNL. EMSL operations are supported by the DOE's Office of Biological and Environmental Research. SBB and JHW also acknowledge support by the ASCR Multiscale Mathematics program of the U.S. Department of Energy, Office of Science ~ DE-FG02-05ER25707. KT acknowledges support by the Nanoscale Science, Engineering, and Technology program and the Environmental Management Sciences program of the U.S. Department of Energy, Office of Science ~ DE-AC06-76RLO 1830, and from LSI Inc. through the Semiconductor Research Corporation.

Appendix A: Fourier Representation of the Free-Space Coulomb Kernel

Free-space boundary conditions can be performed, provided the density has decayed to zero at the edge of the supercell, by restricting the integration to just one isolated supercell, Ω ,

$$\begin{aligned} E_{Coulomb} &= \frac{1}{2} \int_{\Omega} \int_{\Omega} \rho_1(\mathbf{r}) g(|\mathbf{r} - \mathbf{r}'|) \rho_2(\mathbf{r}') d\mathbf{r} d\mathbf{r}' \\ V_H(\mathbf{r}) &= \int_{\Omega} g(|\mathbf{r} - \mathbf{r}'|) \rho_2(\mathbf{r}') d\mathbf{r}' \end{aligned} \quad (\text{A.1})$$

This essentially defines a modified coulomb interaction

$$g(|\mathbf{r} - \mathbf{r}'|) = \begin{cases} \frac{1}{|\mathbf{r} - \mathbf{r}'|} & \text{for } \mathbf{r}, \mathbf{r}' \in \Omega \\ 0 & \text{otherwise} \end{cases} \quad (\text{A.2})$$

Hockney and Eastwood showed that an interaction of the form of Eq. A2 could still be used in conjunction with the Fast-Fourier Transform convolution theorem^{48,49}. In their algorithm, the interaction between neighboring supercells is removed by padding the density with an external region of zero density, or in the specific case of a density defined in cubic supercell of length L , the density is extended to a cubic supercell of length $2L$, where the original density is defined as before on the $[0, L]^3$ domain and the remainder of the $[0, 2L]^3$ domain is set to zero. The grid is 8 times larger than the conventional grid. The Coulomb potential is calculated by convoluting the density with the Green's function kernel on the extended grid. The density on the extended grid is defined by expanding the conventional grid to the extended grid and putting zeros where the conventional grid is not defined. After the aperiodic convolution, the free-space potential is obtained by restricting the extended grid to the conventional grid.

The convolution theorem suggests that defining $g(\mathbf{r})$ in reciprocal space will lead to a much higher accuracy. A straightforward definition in reciprocal space is

$$\begin{aligned} g(\mathbf{r}) &= \sum_{\mathbf{G}} g_{\text{uniform}}(\mathbf{G}) e^{i\mathbf{G} \cdot \mathbf{r}} \\ g_{\text{uniform}}(\mathbf{G}) &= \frac{1}{h^3} \int_{\Omega'} \frac{e^{-i\left(\frac{\mathbf{G} \cdot \mathbf{r}}{2}\right)}}{|\mathbf{r}|} d\mathbf{r} \end{aligned} \quad (\text{A.3})$$

where Ω' is the volume of the extended unit cell and h^3 is the volume of the unit cell divided by the number of conventional FFT grid points. The reciprocal space definition gains accuracy because the singularity at $\bar{\mathbf{r}} = \bar{\mathbf{r}}'$ in Eq. A.1 is analytically integrated out. Even when Eq. A.3 is used to define the kernel, a slight inexactness in the calculated electron-electron Coulomb energy will always present due to the discontinuity introduced in the definition of the extended density where the extended density is forced to be zero in the extended region outside of Ω . However, this discontinuity is small, since the densities we are interested in decay to zero within Ω , thus making the finite Fourier expansion of the extended densities extremely close to zero in extended region outside of Ω .

Equation A.3 could be calculated numerically; however, alternative definitions can be used with little loss of numerical accuracy^{50,51}. In an earlier work^{50,51}, we suggested that the cutoff Coulomb kernel could be defined as

$$\begin{aligned}
g(\vec{r}) &= g_{short-range}(\vec{r}) + g_{long-range}(\vec{r}) \\
g_{short-range}(\vec{r}) &= \sum \bar{g}_{short-range}(\vec{G}) e^{i\vec{G} \cdot \vec{r}} \\
\bar{g}_{short-range}(\vec{G}) &= \begin{cases} \frac{4\pi}{h^3 |\vec{G}|^2} \left(1 - e^{-\left(\frac{|\vec{G}|^2}{4\epsilon^2}\right)} \right) & \text{for } |\vec{G}| \neq 0 \\ \frac{\pi}{h^3 \epsilon^2} & \text{for } |\vec{G}| = 0 \end{cases} \\
\bar{g}_{long-range}(\vec{r}) &= \begin{cases} \frac{erf(\epsilon |\vec{r}|)}{|\vec{r}|} & \text{for } |\vec{r}| \neq 0 \\ \frac{2\epsilon}{\sqrt{\pi}} & \text{for } |\vec{r}| = 0 \end{cases}
\end{aligned} \tag{A.4}$$

Appendix B: Fourier Representation of the Cutoff Coulomb Kernel

Directly using Eq. 15 in a plane-wave program can lead to inaccurate results because of the singularity contained in it at $\mathbf{r} = 0$. However, very high accuracy can be obtained by using the convolution theorem to define the cutoff Coulomb kernel $f(\mathbf{r})$. A straightforward definition in reciprocal space is

$$\begin{aligned}
f_{cutoff}(\mathbf{r}) &= \sum_{\mathbf{G}} f_{uniform}(\mathbf{G}) e^{i\mathbf{G} \cdot \mathbf{r}} \\
f_{uniform}(\mathbf{G}) &= \frac{1}{h^3} \int_{\Omega} e^{-i(\frac{\mathbf{G} \cdot \mathbf{r}}{2})} \frac{1 - \left(1 - \exp \left\{ - \left(|\mathbf{r}| / R_c \right)^{N+2} \right\} \right)^N}{|\mathbf{r}|} d\mathbf{r}
\end{aligned} \tag{B.1}$$

where Ω is the volume of the unit cell and h^3 is Ω divided by the number of conventional FFT grid points. The reciprocal space definition gains accuracy because the singularity at $\mathbf{r} = \mathbf{r}'$ in Eq. 15 is analytically integrated out. Eq. B.1 could be calculated numerically; however we have found that the following alternative definition can be used with little loss of numerical accuracy.

$$\begin{aligned}
f_{cutoff}(\mathbf{r}) &= \left(f_{short-range}(\mathbf{r}) + f_{long-range}(\mathbf{r}) \right) \left(1 - \left(1 - \exp \left\{ -(|\mathbf{r}| / R_c)^{N+2} \right\} \right)^N \right) \\
f_{short-range}(\mathbf{r}) &= \sum_{\mathbf{G}} f_{short-range}(\mathbf{G}) \exp(i\mathbf{G} \cdot \mathbf{r}) \\
f_{short-range}(\mathbf{G}) &= \begin{cases} \frac{1}{\Omega} \frac{4\pi}{|\mathbf{G}|^2} \left(1 - e^{-\left(\frac{|\mathbf{G}|^2}{4\varepsilon^2} \right)} \right) & \text{for } |\mathbf{G}| \neq 0 \\ \frac{1}{\Omega} \frac{\pi}{\varepsilon^2} & \text{for } |\mathbf{G}| = 0 \end{cases} \\
f_{long-range}(\mathbf{r}) &= \begin{cases} \frac{\text{erf}(\varepsilon|\mathbf{r}|)}{|\mathbf{r}|} & \text{for } |\mathbf{r}| \neq 0 \\ \frac{2\varepsilon}{\sqrt{\pi}} & \text{for } |\mathbf{r}| = 0 \end{cases}
\end{aligned} \tag{B.2}$$

Appendix C: Analytical Derivatives With Respect To The Lattice Parameters Of The Screened Coulomb Kernel

The screened Coulomb energy may be written in plane-wave basis by

$$E_{cutoff_Coulomb} = \xi \sum_{\sigma,i} \left\{ \frac{\Omega}{2} \sum_{\mathbf{G}} |\rho_{\sigma,i}(\mathbf{G})|^2 \tilde{f}_{cutoff}(\mathbf{G}) \right\} \tag{C.1}$$

where Ω is the unit cell volume, N_1 , N_2 , and N_3 are the number of FFT gridpoints in the first, second, and third dimensions,

$$\rho_{\sigma,i}(\mathbf{G}) = \frac{1}{N_1 N_2 N_3} \sum_{\mathbf{r}} \rho_{\sigma,i}(\mathbf{r}) \exp(-i\mathbf{G} \cdot \mathbf{r}) \tag{C.2}$$

is the inverse Fourier transform of the density.

$$\tilde{f}_{cutoff}(\mathbf{G}) = \sum_{\mathbf{r}} \tilde{f}_{cutoff}(\mathbf{r}) \exp(-i\mathbf{G} \cdot \mathbf{r}) \tag{C.3}$$

is the inverse Fourier transform of the modified screened Coulomb kernel. Which we define here by

$$\tilde{f}_{cutoff}(\mathbf{r}) = \tilde{f}_{long-short}(\mathbf{r}) h_{cut}(|\mathbf{r}|) \quad (\text{C.4})$$

with

$$h_{cut}(|\mathbf{r}|) = \left(1 - \left(1 - \exp\left\{ - \left(|\mathbf{r}| / R \right)^{p+2} \right\} \right)^p \right) \quad (\text{C.5})$$

$$\tilde{f}_{long-short}(\mathbf{r}) = \tilde{f}_{long-range}(|\mathbf{r}|) + \frac{1}{N_1 N_2 N_3} \sum_{\mathbf{G}} \tilde{f}_{short-range}(|\mathbf{G}|) \exp(i\mathbf{G} \cdot \mathbf{r}) \quad (\text{A.6})$$

$$\tilde{f}_{short-range}(|\mathbf{G}|) = \begin{cases} \frac{4\pi}{|\mathbf{G}|^2} \left(1 - \exp\left(- \frac{|\mathbf{G}|^2}{4\varepsilon^2} \right) \right) & \text{for } |\mathbf{G}| \neq 0 \\ \frac{\pi}{\varepsilon^2} & \text{for } |\mathbf{G}| = 0 \end{cases} \quad (\text{C.7})$$

$$\tilde{f}_{long-range}(|\mathbf{r}|) = \begin{cases} \frac{\Omega}{N_1 N_2 N_3} \frac{\text{erf}(\varepsilon|\mathbf{r}|)}{|\mathbf{r}|} & \text{for } |\mathbf{r}| \neq 0 \\ \frac{\Omega}{N_1 N_2 N_3} \frac{2\varepsilon}{\sqrt{\pi}} & \text{for } |\mathbf{r}| = 0 \end{cases} \quad (\text{C.8})$$

Note that we have modified the definition of screened Coulomb kernel in Eq. 15 by setting

$$\tilde{f}_{cutoff}(\mathbf{r}) = \frac{\Omega}{N_1 N_2 N_3} f_{cutoff}(\mathbf{r}) \quad (\text{C.9})$$

The derivatives of with respect to the components of the cell matrix $\mathbf{h}_{uv}$ can be calculated for a plane-wave basis by

$$\begin{aligned}
\frac{\partial E_{cutoff-Coulomb}}{\partial \mathbf{h}_{uv}} &= \xi \sum_{\sigma,i} \left\{ \frac{\partial \Omega}{\partial \mathbf{h}_{uv}} \frac{1}{2} \sum_{\mathbf{G}} |\rho_{\sigma,i}(\mathbf{G})|^2 \tilde{f}_{cutoff}(\mathbf{G}) + \frac{\Omega}{2} \sum_{\mathbf{G}} \frac{\partial |\rho_{\sigma,i}(\mathbf{G})|^2}{\partial \mathbf{h}_{uv}} \tilde{f}_{cutoff}(\mathbf{G}) \right. \\
&\quad \left. + \frac{\Omega}{2} \sum_{\mathbf{G}} |\rho_{\sigma,i}(\mathbf{G})|^2 \frac{\partial \tilde{f}_{cutoff}(\mathbf{G})}{\partial \mathbf{h}_{uv}} \right\} \quad (C.10) \\
&= \xi \sum_{\sigma,i} \left\{ -\frac{\Omega}{2} \sum_{\mathbf{G}} |\rho_{\sigma,i}(\mathbf{G})|^2 \tilde{f}_{cutoff}(\mathbf{G}) (\mathbf{h}_{uv}^t)^{-1} + \frac{\Omega}{2} \sum_{\mathbf{G}} |\rho_{\sigma,i}(\mathbf{G})|^2 \left(\frac{\partial \tilde{f}_{cutoff}(\mathbf{G})}{\partial \mathbf{h}_{uv}} \right) \right\}
\end{aligned}$$

The derivative of the inverse Fourier transformed screened Coulomb kernel can be computed using the following formulae

$$\frac{\partial \tilde{f}_{cutoff}(\mathbf{G})}{\partial \mathbf{h}_{uv}} = \sum_{\mathbf{r}} \left(\frac{\partial \tilde{f}_{cutoff}(\mathbf{r})}{\partial \mathbf{h}_{uv}} \right) \exp(-i\mathbf{G} \cdot \mathbf{r}) \quad (C.11)$$

$$\frac{\partial \tilde{f}_{cutoff}(\mathbf{r})}{\partial \mathbf{h}_{uv}} = \frac{\partial \tilde{f}_{long-short}(\mathbf{r})}{\partial \mathbf{h}_{uv}} h_{cut}(|\mathbf{r}|) + \tilde{f}_{long-short}(\mathbf{r}) \frac{\partial h_{cut}(|\mathbf{r}|)}{\partial |\mathbf{r}|} \frac{\partial |\mathbf{r}|}{\partial \mathbf{h}_{uv}} \quad (C.12)$$

$$\begin{aligned}
\frac{\partial \tilde{f}_{long-short}(\mathbf{r})}{\partial \mathbf{h}_{uv}} &= \tilde{f}_{long-range}(\mathbf{r}) (\mathbf{h}_{uv}^t)^{-1} + \frac{\partial \tilde{f}_{long-range}(|\mathbf{r}|)}{\partial |\mathbf{r}|} \frac{\partial |\mathbf{r}|}{\partial \mathbf{h}_{uv}} \\
&\quad + \frac{1}{N_1 N_2 N_3} \sum_{\mathbf{G}} \left(\frac{\partial \tilde{f}_{short-range}(|\mathbf{G}|)}{\partial |\mathbf{G}|^2} \frac{\partial |\mathbf{G}|^2}{\partial \mathbf{h}_{uv}} \right) \exp(i\mathbf{G} \cdot \mathbf{r}) \quad (C.13)
\end{aligned}$$

$$\frac{\partial \tilde{f}_{long-range}(|\mathbf{r}|)}{\partial |\mathbf{r}|} = \begin{cases} \frac{\Omega}{N_1 N_2 N_3} \left(\frac{2\varepsilon}{\sqrt{\pi}} \frac{\exp(-\varepsilon^2 |\mathbf{r}|^2)}{|\mathbf{r}|} - \frac{\text{erf}(\varepsilon |\mathbf{r}|)}{|\mathbf{r}|^2} \right) & \text{for } |\mathbf{r}| \neq 0 \\ 0 & \text{for } |\mathbf{r}| = 0 \end{cases} \quad (C.14)$$

$$\frac{\partial \tilde{f}_{short-range}(|\mathbf{G}|)}{\partial |\mathbf{G}|^2} = \begin{cases} -\frac{4\pi}{|\mathbf{G}|^4} \left(1 - \exp\left(-\frac{|\mathbf{G}|^2}{4\varepsilon^2}\right) \right) + \frac{\pi}{\varepsilon^2} \frac{\exp\left(-\frac{|\mathbf{G}|^2}{4\varepsilon^2}\right)}{|\mathbf{G}|^2} & \text{for } |\mathbf{G}| \neq 0 \\ 0 & \text{for } |\mathbf{G}| = 0 \end{cases} \quad (C.15)$$

$$\frac{\partial h_{cut}(|\mathbf{r}|)}{\partial |\mathbf{r}|} = \frac{-\exp\left\{-\left(|\mathbf{r}|/R\right)^{p+2}\right\}\left(1-\exp\left\{-\left(|\mathbf{r}|/R\right)^{p+2}\right\}\right)^{p-1}p(p+2)\left(\frac{|\mathbf{r}|}{R}\right)^{p+1}}{R} \quad (\text{C.16})$$

$$\frac{\partial |\mathbf{r}|}{\partial \mathbf{h}_{uv}} = \sum_s \frac{r_u r_s}{|\mathbf{r}|} \left(\mathbf{h}_{sv}^t\right)^{-1} \quad (\text{C.17})$$

$$\frac{\partial |\mathbf{G}|^2}{\partial \mathbf{h}_{uv}} = -2G_u \sum_s G_s \left(\mathbf{h}_{sv}^t\right)^{-1} \quad (\text{C.18})$$

Appendix D: Pseudocode for Brdcst_Step, Exchng_Compute, and Index_Compute subroutines

Subroutine Brdcst_Step(l, taskid_j, Npj, Ψ , N_g , N_e , requests, reqcnt)

```

Levels = Floor(Log2(Npj)); shft = 2**l;
pfrom = (taskid_j+shft)%Npj; pto=(taskid_j+Npj-shft)%Npj
if l==(Levels-1) then
    sz = (Npj - 2**Level)/2 + 1
else
    sz = 2**l
end if
rindx=( $N_e/N_{pj}$ )*shft; rsize=( $N_g/N_{pi}$ )*( $N_e/N_{pj}$ )*sz; ssize=rsize;
reqcnt = 0;
reqcnt+=1; Row_Isend( pto, ssize,  $\Psi(:,1)$ , requests(reqcnt))
reqcnt+=1; Row_Irecv( pfrom, rsize,  $\Psi(:,rindx)$ , requests(reqcnt))

```

Subroutine Reduce_Step(l, taskid_j, Npj, $K\Psi$, N_g , N_e , tmp)

```

Levels = Floor(Log2(Npj)); shft = 2**l;
pfrom = (taskid_j+shft)%Npj; pto=(taskid_j+Npj-shft)%Npj
if l==(Levels-1) then
    sz = (Npj - 2**Level)/2 + 1
else
    sz = 2**l
end if
sindx=( $N_e/N_{pj}$ )*shft; ssize=( $N_g/N_{pi}$ )*( $N_e/N_{pj}$ )*sz; rsize=ssize;
reqcnt = 0;
reqcnt += 1; Row_Isend( pto, ssize,  $K\Psi(:,sindx)$ , requests(reqcnt))
reqcnt += 1; Row_Irecv( pfrom, rsize, tmp, requests(reqcnt))
Row_WaitAll( requests, reqcnt)
 $K\Psi(:,1:rsize) += tmp(:,1:rsize)$ 

```

Subroutine Exchng_Compute(i1, i2, j1, j2, im, it, Ψ , $K\Psi$)

```

special = (i1==j1) and (i2==j2); counter=0
if (special) then
    for m=i1,i2
        for n=j1,m
            if (counter%im)==it then
                 $\rho(:) \leftarrow \text{Column\_FFT\_rc}(\Psi(:,m) \cdot \Psi(:,n))$ 
                 $V(:) \leftarrow \text{Column\_FFT\_cr}(f_{\text{screened}}(:) \cdot \rho(:))$ 
                 $K\Psi(:,m) = V(:) \cdot \Psi(:,n)$ ; if  $m < n$   $K\Psi(:,n) = V(:) \cdot \Psi(:,m)$ 
            end if
            counter+=1
        end for
    end for

```



```

    end if
    counter += 1
end for
end for
else
    for m=i1,i2
    for n=j1,j2
        if (counter%im)==it then
             $\rho(\cdot) \leftarrow \text{Column\_FFT\_rc}(\Psi(\cdot, m) \cdot \Psi(\cdot, n))$ 
             $V(\cdot) \leftarrow \text{Column\_FFT\_cr}(f_{\text{screened}}(\cdot) \cdot \rho(\cdot))$ 
             $K\Psi(\cdot, m) = V(\cdot) \cdot \Psi(\cdot, n); K\Psi(\cdot, n) = V(\cdot) \cdot \Psi(\cdot, m)$ 
        end if
        counter += 1
    end for
end for
end if

Subroutine Index_Compute( opt, l, taskid_j, Npj, i1, i2, j1, j2, im, it)
Levels = Floor(Log2(Npj)); im=1; it=0
if (opt==1) then
    if (l==0) then
        i1=1; i2=(Ne/Npj); j1=i1; j2=i2
    else
        shft=2**(l-1); sz=shft;
        if (l==Levels) sz=(Npj-2**Levels)/2+1
        i1=1; i2=(Ne/Npj); j1=shft*(Ne/Npj); j2=j1+sz*(Ne/Npj)
        if (l==Levels) and (Npj%2)==0 then
            im = 2; it = Floor(taskid_j/(Npj/2))
            if (it==1) then
                swap(i1,j1); swap(i2,j2)
            end if
        end if
    end if
else if (opt==2) then
    shft = 2**(Levels-1)
    sz = (Npj-2**Levels)/2
    i1=1; i2=(Ne/Npj); j1=shft*(Ne/Npj); j2=j1+sz*(Ne/Npj)-1
else if (opt==3) then
    shft = 2**(Levels-1)
    sz = (Npj-2**Levels)/2+1

```

```

i1=1; i2=(Ne/Npj); j2=(shft+sz)*(Ne/Npj)-1
j1=jend-(Ne/Npj)+1
if (Npj%2)==0 then
    im = 2; it = Floor(taskid_j/(Npj/2))
    if (it==1) then
        swap(i1,j1); swap(i2,j2)
    end if
end if
end if

```

Table I: Timing in seconds for 80 atom Fe₂O₃ hybrid-DFT energy calculations at various combinations of processor sizes (Np) and processor grids (Np_j, Np_i=Np/Np_j). Bold face timings are the overall fastest simulations at each processor size. Timings determined from calculations on the Franklin Cray-XT4 computer system at NERSC.

Np	Np _j =1	Np _j =2	Np _j =4	Np _j =8	Np _j =16	Np _j =32	Np _j =64
Total Time							
16	1.19E+03						
32	5.29E+02	6.21E+02					
64	3.37E+02	2.68E+02					
128	4.88E+02	1.61E+02	1.32E+02	1.58E+02			
256		2.21E+02	8.59E+01	7.11E+01	8.21E+01	8.43E+01	9.29E+01
512			1.21E+02	4.33E+01	3.72E+01	4.48E+01	4.67E+01
1024			2.33E+02	7.28E+01	3.08E+01	2.12E+01	2.66E+01
2048					4.69E+01	1.32E+01	1.23E+01
Exact Exchange							
16	1.16E+03						
32	5.10E+02	6.09E+02					
64	3.28E+02	2.62E+02					
128	4.82E+02	1.57E+02	1.28E+02	1.53E+02			
256		2.17E+02	8.35E+01	6.85E+01	7.83E+01	7.86E+01	8.40E+01
512			1.19E+02	4.17E+01	3.54E+01	4.21E+01	4.07E+01
1024			2.31E+02	7.14E+01	2.93E+01	1.93E+01	2.34E+01
2048					4.57E+01	1.22E+01	1.04E+01
Non-Local Pseudopotential							
16	2.31E+01						
32	1.15E+01	7.07E+00					
64	6.41E+00	3.52E+00					
128	2.44E+00	2.05E+00	1.78E+00	1.76E+00			
256		1.37E+00	1.08E+00	8.37E-01	8.34E-01	8.87E-01	1.16E+00
512			5.25E-01	3.70E-01	3.21E-01	3.23E-01	4.97E-01
1024			3.72E-01	2.16E-01	1.48E-01	1.32E-01	1.82E-01
2048					1.10E-01	8.04E-02	6.57E-02
Orthogonality							
16	3.38E+00						
32	1.68E+00	1.59E+00					
64	7.84E-01	8.64E-01					
128	3.02E-01	5.30E-01	6.41E-01	1.15E+00			
256		3.36E-01	3.57E-01	8.39E-01	1.46E+00	3.11E+00	5.54E+00
512			2.60E-01	4.45E-01	7.37E-01	1.49E+00	3.90E+00
1024			2.61E-01	3.37E-01	5.63E-01	9.96E-01	2.19E+00
2048					3.56E-01	4.50E-01	1.25E+00

Table II Reaction and activation energies (kcal/mol) of silanes. Zero-point corrections not included.

Method	$\text{SiH}_4 \rightarrow \text{SiH}_2 + \text{H}_2$		$\text{Si}_2\text{H}_4 \rightarrow 2\text{SiH}_2$	$\text{Si}_2\text{H}_6 \rightarrow \text{SiH}_4 + \text{SiH}_2$	$\text{Si}_2\text{H}_6 \rightarrow \text{Si}_2\text{H}_4 + \text{H}_2$			$\Delta\Delta E_{\text{rxn}}$	$\Delta\Delta E_{\text{ts}}$
	E_{rxn}	E_{ts}	E_{rxn}	E_{rxn}	E_{rxn}	E_{ts1}	E_{ts2}		
PBE96^a	59.4	51.4	70.4	58.5	47.6	72.2	46.6	4.2	11.0
PBE0^a	62.1	56.7	69.5	59.9	52.4	79.3	51.3	3.6	5.3
B3LYP^{b,c}	59.8	58.0	61.0	53.0	51.8	84.6	54.2	1.0	2.1
MP2^{b,c}	63.5	61.8	67.2	57.8	54.1	88.6	55.4	3.3	0.9
E-QCISD(T)^{b,c}	60.3	60.0	62.8	54.5	52.0	87.8	55.4	-----	-----
^a Geometries were optimized									
^b B3LYP/6-311++G(2df,2pd), MP2/6-311++G(2df,2pd), and extrapolated QCISD(T) calculations from Nachtigall ¹² .									
^c Geometries not optimized, taken from Gordon et al. ⁵²									

Table III: Band gaps (eV) of selected systems.

Crystal	Exp.	PBE96	PBE0
Al₂O₃	9.0	5.9	8.4
TiO₂	3.0	1.8	3.1
SiO₂	8.9	6.1	8.7

Figure Captions

Figure 1. Localized orbital of a 72 atom unit cell of a SiO₂ crystal, along with its artificial periodic images.

Figure 2. Plot of the cutoff coulomb kernel, Eq. 14 ($R=6.0$ bohrs, $N=8$).

Figure 3. Hybrid DFT results for periodic supercells containing 2, 4, 6, 8 and 10 -CH₂- units using a series of 5 cutoff Coulomb kernels with increasing values of R_c ($N=8$ and $R_c=2, 4, 6, 8$, and 10 bohrs). The exact exchange energy per electron contribution of the hybrid exchange term (left) and the HOMO-LUMO gap (right) for each of the 5 unit cells a function of R_c is shown.

Figure 4. Three parallel distributions for the occupied molecular orbitals: (a) parallel decomposition in which only the spatial dimension N_g is distributed over a one-dimensional processor grid, (b) parallel decomposition in which the orbital dimension N_{occ} is distributed over a one-dimensional processor grid, (c) parallel decomposition in which both the spatial and orbital dimensions are distributed over a two-dimensional processor grid.

Figure 5. The overall fastest and the corresponding constituent timings taken from Table I for an 80 atom Fe₂O₃ hybrid-DFT energy calculations. Calculations were performed on the Franklin Cray-XT4 computer system at NERSC.

Figure 6. The overall fastest timings taken for an 80 (Algorithms 2 and 3) and 160 atom (Algorithm 3) Fe₂O₃ hybrid DFT energy calculations. Calculations were performed on the Franklin Cray-XT4 computer system at NERSC.

Figure 7. Geometries of silanes calculated at the PBE96, PBE0 levels. Bond lengths are in Å and angles are in degrees. The values in normal font are the PBE96 results, italicized are for PBE0, and the values in parenthesis are from Nachtigall¹².

Figure 8. Top illustrates the structure of annite and the location of its ferromagnetic octahedral layers. Bottom shows the electron densities in the octahedral layer from DFT and Hybrid-DFT simulations of oxidized annite.

Figure 1

Figure 2

Figure 3

Figure 4

Figure 5

Figure 6

Figure 7

Figure 8

References

1. Parr, R. G.; Yang, W. Density-functional theory of atoms and molecules; Oxford University Press ; Clarendon Press: New York; Oxford [England], 1989.
2. Dreizler, R. M.; Gross, E. K. U. Density functional theory : an approach to the quantum many-body problem; Springer-Verlag: Berlin; New York, 1990.
3. Becke, A. D. J Chem Phys 1988, 88, 2547.
4. Lee, C.; Yang, W.; Parr, R. G. Phys Rev B 1988, 37, 785.
5. Adamo, C.; Barone, V. J Chem Phys 1997, 110(13), 6158-6170.
6. Lundberg, M.; Siegbahn, P. E. M. Journal of Chemical Physics 2005, 122(22), -.
7. Perdew, J. P.; M., E. In Electronic Density Functional Theory: Recent Progress and New Directions; Dobson, J. F.; Vignale, G.; M.P., D., Eds.; Plenum: New York, 1997.
8. Marsman, M.; Paier, J.; Stroppa, A.; Kresse, G. Journal of Physics-Condensed Matter 2008, 20(6), -.
9. Stadele, M.; Majewski, J. A.; Vogl, P.; Gorling, A. Physical Review Letters 1997, 79(11), 2089-2092.
10. Stadele, M.; Moukara, M.; Majewski, J. A.; Vogl, P.; Gorling, A. Physical Review B 1999, 59(15), 10031-10043.
11. Cinquini, F.; Giordano, L.; Pacchioni, G.; Ferrari, A. M.; Pisani, C.; Roetti, C. Physical Review B 2006, 74(16), 165403.
12. Nachtigall, P.; Jordan, K. D.; Smith, A.; Jonsson, H. Journal of Chemical Physics 1996, 104(1), 148-158.
13. van Wullen, C. Physical Chemistry Chemical Physics 2000, 2(10), 2137-2144.
14. Wiggs, J.; Jonsson, H. Computer Physics Communications 1995, 87(3), 319-340.
15. Wiggs, J.; Jonsson, H. Computer Physics Communications 1994, 81(1-2), 1-18.
16. Nelson, J. S., S. J. Plimpton, and M. P. Sears. Phys Rev B 1993, 47(4), 1765.
17. Kendall, R. A.; Apra, E.; Bernholdt, D. E.; Bylaska, E. J.; Dupuis, M.; Fann, G. I.; Harrison, R. J.; Ju, J. L.; Nichols, J. A.; Nieplocha, J.; Straatsma, T. P.; Windus, T. L.; Wong, A. T. Computer Physics Communications 2000, 128(1-2), 260-283.
18. Bylaska, E. J.; Valiev, M.; Kawai, R.; Weare, J. H. Comput Phys Comm 2002, 143, 11-28.
19. Windus, T. L.; Bylaska, E. J.; Dupuis, M.; Hirata, S.; Pollack, L.; Smith, D. M.; Straatsma, T. P.; Apra, E. Computational Science - Iccs 2003, Pt Iv, Proceedings 2003, 2660, 168-177.
20. Gygi, F.; Draeger, E. W.; Schulz, M.; de Supinski, B. R.; Gunnels, J. A.; Austel, V.; Sexton, J. C.; Franchetti, F.; Kral, S.; Ueberhuber, C. W.; Lorenz, J. SC '06: Proceedings of the 2006 ACM/IEEE conference on Supercomputing, 2006.
21. Bylaska, E. J.; Weare, J. H. Abstracts of Papers of the American Chemical Society 2005, 230, U1321-U1321.
22. Kohanoff, J. Electronic structure calculations for solids and molecules : theory and computational methods; Cambridge University Press: Cambridge, 2006.
23. Szabo, A.; Ostlund, N. S. Modern quantum chemistry : introduction to advanced electronic structure theory; Dover Publications: Mineola, N.Y., 1996.
24. Ashcroft, N. W.; Mermin, N. D. Solid state physics; Holt, Rinehart and Winston: New York, 1976.
25. Gygi, F.; Baldereschi, A. Physical Review Letters 1989, 62(18), 2160-2163.
26. Gygi, F.; Baldereschi, A. Helvetica Physica Acta 1986, 59(6-7), 972-972.
27. Gygi, F.; Baldereschi, A. Physical Review B 1986, 34(6), 4405-4408.
28. Gygi, F.; Baldereschi, A. Helvetica Physica Acta 1985, 58(5), 928-928.
29. Chawla, S.; Voth, G. A. Journal of Chemical Physics 1998, 108(12), 4697-4700.
30. Sorouri, A.; Foulkes, W. M. C.; Hine, N. D. M. Journal of Chemical Physics 2006, 124(6), -.
31. Gorling, A. Physical Review B 1996, 53(11), 7024-7029.
32. Carrier, P.; Rohra, S.; Gorling, A. Physical Review B 2007, 75(20), -.
33. Marzari, N.; Vanderbilt, D. Phys Rev B 1997, 56(20), 12847-12865.
34. Silvestrelli, P. L. Phys Rev B 1999, 59(15), 9703-9706.

35. Foster, J. M.; Boys, S. F. *Reviews of Modern Physics* 1960, 32(2), 300-302.
36. Bylaska, E. J.; Tsemekhman, K.; Gao, F. *Physica Scripta* 2006, T124, 86-90.
37. Du, J. C.; Corrales, L. R.; Tsemekhman, K.; Bylaska, E. J. *Nuclear Instruments & Methods in Physics Research Section B-Beam Interactions with Materials and Atoms* 2007, 255(1), 188-194.
38. Maslen, P. E.; Ochsenfeld, C.; White, C. A.; Lee, M. S.; Head-Gordon, M. *Journal of Physical Chemistry A* 1998, 102(12), 2215-2222.
39. Bylaska, E. J.; de Jong, W. A.; Kowalski, K.; Straatsma, T. P.; Valiev, M.; Wang, D.; Aprà, E.; Windus, T. L.; Hirata, S.; Hackler, M. T.; Zhao, Y.; Fan, P.-D.; Harrison, R. J.; Dupuis, M.; Smith, D. M. A.; Nieplocha, J.; Tipparaju, V.; Krishnan, M.; Auer, A. A.; Nooijen, M.; Brown, E.; Cisneros, G.; Fann, G. I.; Früchtl, H.; Garza, J.; Hirao, K.; Kendall, R.; Nichols, J. A.; Tsemekhman, K.; Wolinski, K.; Anchell, J.; Bernholdt, D.; Borowski, P.; Clark, T.; Clerc, D.; Dachsel, H.; Deegan, M.; Dyall, K.; Elwood, D.; Glendening, E.; Gutowski, M.; Hess, A.; Jaffe, J.; Johnson, B.; Ju, J.; Kobayashi, R.; Kutteh, R.; Lin, Z.; Littlefield, R.; Long, X.; Meng, B.; Nakajima, T.; Niu, S.; Pollack, L.; Rosing, M.; Sandrone, G.; Stave, M.; Taylor, H.; Thomas, G.; van Lenthe, J.; Wong, A.; Zhang, Z.; Pacific Northwest National Laboratory, Richland, Washington 99352-0999, USA., 2007.
40. Perdew, J. P.; Burke, K.; Ernzerhof, M. *Physical Review Letters* 1996, 77(18), 3865-3868.
41. Hamann, D. R. *Phys Rev B* 1989, 40(5), 2980-2987.
42. Troullier, N.; Martins, J. L. *Physical Review B* 1991, 43(3), 1993-2006.
43. Kleinman, L.; Bylander, D. M. *Physical Review Letters* 1982, 48(20), 1425-1428.
44. Car, R.; Parrinello, M. *Phys Rev Lett* 1985, 55(22), 2471.
45. Shimazaki, T.; Asai, Y. *Journal of Chemical Physics* 2009, 130(16), -.
46. Heyd, J.; Scuseria, G. E.; Ernzerhof, M. *Journal of Chemical Physics* 2003, 118(18), 8207-8215.
47. Heyd, J.; Scuseria, G. E. *Journal of Chemical Physics* 2004, 120(16), 7274-7280.
48. Hockney, R. W. *Methods Comput Phys* 1970, 9, 135.
49. Hockney, R.; Eastwood, J. W. *Computer simulations using particles*; McGraw-Hill: New York, 1981.
50. Lubin, M. I.; Bylaska, E. J.; Weare, J. H. *Chem Phys Lett* 2000, 322, 447.
51. Bylaska, E. J.; Taylor, P. R.; Kawai, R.; Weare, J. H. *J Phys Chem* 1996, 100(17), 6966.
52. Gordon, M. S.; Truong, T. N.; Bonderson, E. K. *Journal of the American Chemical Society* 1986, 108(7), 1421-1427.