## Faster Exact Exchange in Periodic Systems using Single-precision Arithmetic

John Vinson<sup>a)</sup>

Material Measurement Laboratory, National Institute of Standards and Technology, 100 Bureau Drive, Gaithersburg, MD 20899

(Dated: 4 November 2020)

Density-functional theory simplifies many-electron calculations by approximating the exchange and correlation interactions with a one-electron operator that is a functional of the density. Hybrid functionals incorporate some amount of exact exchange, improving agreement with measured electronic and structural properties. However, calculations with hybrid functionals require substantial computational resources, limiting their use. By calculating the exchange interaction of periodic systems with single-precision arithmetic, the computation time is cut nearly in half with a negligible loss in accuracy. This improvement makes exact exchange calculations quicker and more feasible, especially for high-throughput calculations. Example hybrid density-functional theory calculations of band energies, forces, and x-ray absorption spectra show that this single-precision implementation maintains accuracy with significantly reduced runtime and memory requirements.

<sup>&</sup>lt;sup>a)</sup>Electronic mail: john.vinson@nist.gov

### I. INTRODUCTION

Within density-functional theory (DFT), both the exchange and the correlation interactions are simplified and included as potentials that typically depend only on the local electron density (and its derivatives). While DFT has proven to be widely successful, standard local or semi-local functionals have a number of shortcomings that limit accuracy and general applicability. One particular problem is that within DFT, each electron interacts with the Coulomb potential of the total electron density. This means that each electron is repelled by its own charge density, *e.g.*, self-interaction error. This is mitigated somewhat by including some fraction of the exact exchange interaction (replacing the functional exchange), and a variety of hybrid or screened-hybrid functionals incorporating a mix of functional and exact exchange have been proposed.<sup>1</sup> Despite the success of hybrid functionals, their use has been hindered by the significantly higher computational cost of calculating the exact exchange operator compared to calculations using local or semi-local density functionals.

In this paper I show that the use of single-precision arithmetic substantially reduces the cost of hybrid density functional calculations while maintaining accuracy. First, the exchange operator is reviewed. Then the Adaptively Compressed Exchange operator approximation is introduced,<sup>2</sup> and the implementation of single-precision exact exchange is outlined. In Sec. III example calculations are shown that compare both electronic and structural properties between the original, double-precision and new, single-precision implementations of the exact exchange: total electronic energies, band energies, lattice constants, bulk moduli, forces, and x-ray absorption spectra. The examined systems include both metals and insulators with unit cells ranging from 1 to 192 atoms. Finally, a summary and outlook are presented in Sec. IV.

### II. REVIEW

The exchange operator  $V_X$  is straight-forward,

$$V_X(\mathbf{r}, \mathbf{r}', [\{\psi_i\}]) = -\sum_{i=1}^{N_{\text{occ}}} \frac{\psi_i^*(\mathbf{r})\psi_i(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}$$
(1)

where  $V_X$  is an operator in real-space coordinates **r** and **r**', and is a functional of the  $N_{\text{occ}}$ occupied electron orbitals  $\psi_i$ , requiring self-consistency. (In this case self-consistency requires that the eigenfunctions of a Hamiltonian including  $V_X$  are the same as those that were used to construct  $V_X$ .) The evaluation of the exact exchange operator, especially in periodic systems, is time consuming. This is due in part to the large size of the basis. The number of plane waves is much larger than the number of occupied electron orbitals,  $N_G \gg N_{\text{occ}}$ . Furthermore, the exchange operator must be calculated repeatedly until self-consistency is reached.

A number of methods to reduce the computational expense of evaluating exact exchange in periodic systems have been proposed based on localizing the electron orbitals. Localization techniques reduce the spatial extent of individual orbitals in Eq.  $1,^{3-6}$  and are aided by the exponential localization of occupied orbitals in insulating systems.<sup>7,8</sup> In large systems these techniques can lead to linear scaling in evaluating the exchange integral, but the localization procedure itself can be costly and its applicability is limited to systems with a band gap.

In non-periodic or molecular systems, exact exchange is much less costly, and Hartree-Fock or post-Hartree-Fock calculations are common-place. In simulations of molecular systems, localized basis sets such as Gaussian-type orbitals are typically used. Not only do these basis functions have finite extent, but the total number of basis functions  $N_B$  is much smaller,  $N_G \gg N_B > N_{occ}$ . Substantial work has been done to reduce the computational cost of four-center integrals like the exchange. The use of reduced-precision has been investigated in quantum chemistry calculations including Møller-Plesset perturbation theory<sup>9</sup> and coupled-cluster calculations.<sup>10</sup> In both cases it was found to be sufficiently accurate. Finitebasis set codes have also shown the utility of reducing the precision for storing pre-computed four-center integrals.<sup>11,12</sup> Within the context of periodic systems, single-precision arithmetic can be used to calculate the electron self-energy within the ABINIT code.<sup>13</sup> However, for periodic, plane-wave approaches reduced precision is rare.

For either finite or periodic systems, with or without a band gap, the Adaptively Compressed Exchange (ACE) operator method reduces the computational cost of exact exchange without compromising accuracy.<sup>2</sup> The single-precision implementation of exact exchange is built on top of the ACE method. A brief overview is presented here.

In the ACE method, the exchange operator  $V_X[\{\psi_i\}]$  is still applied to each orbital to create a set

$$W_j(\mathbf{r}) = [V_X(\mathbf{r}, \mathbf{r}', [\{\psi_i\}])\psi_j(\mathbf{r}')](\mathbf{r})$$
(2)

As noted in Ref. 2, the construction of W requires  $N_{\text{occ}}N_e$  solutions to a Poisson-like problem,

where  $N_e \ge N_{\text{occ}}$  is the number of orbitals being solved for. Using fast Fourier transforms to solve the Poisson problem gives a computational cost that scales as  $N_{\text{occ}}N_eN_G \log N_G$ .

From W, overlaps are taken,

$$M_{jk} = \int d\mathbf{r} \,\psi_j^*(\mathbf{r}) W_k(\mathbf{r}) \tag{3}$$

where M has the dimension  $N_e$ , and this scales as  $N_e^2 N_G$ . After a Cholesky factorization  $-M = LL^{\dagger}$ , the ACE approximation to the exchange operator is given by

$$\xi_i(\mathbf{r}) = W_j(\mathbf{r})(L^{\dagger})_{ji} \tag{4}$$

$$V_X^{\text{ACE}}(\mathbf{r}, \mathbf{r}') = -\sum_{i=1}^{N_e} \xi_i(\mathbf{r}) \xi_i(\mathbf{r}')$$
(5)

Two nested self-consistent loop are used to converge the electron orbitals. In the inner loop  $V_X^{ACE}$  is held constant and only updated once the inner loop has reach self-consistency. This greatly reduces the number of times the exchange operator needs to be constructed, but, despite the improvement of the ACE method, including exact exchange still adds significant cost to the calculation. The evaluation of Eq. 2 scales roughly cubically with system size. Additionally, significant memory is required to store both W and  $\xi$ .

Here single-precision calculations of the exact exchange for plane wave basis set calculations are introduced, building on top of the ACE method. Both memory usage and calculation time are reduced by using single-precision arithmetic to construct and store Win Eq. 2 and to store  $\xi$  in Eqs. 4 and 5. In Eq. 3 the double-precision electron orbitals are used, and the matrices M and L are built in double-precision. Mixed-precision matrix operations have not been used. Instead, matrices are temporarily converted between precision, *e.g.*, W to double in Eq. 3 and  $L^{\dagger}$  to single in Eq. 4. The conversions were not found to be a significant burden, and mixed-precision linear algebra libraries are not widely available. This single-precision approximation reduces the runtime and maximum memory load by almost a factor of 2, with little to no loss of accuracy compared against the double-precision version.

I modified version 6.5 of the QUANTUM ESPRESSO code.<sup>14,15</sup> Taking advantage of the iterative nature of the ACE method, an option was also implemented to fall back to the double precision routines after the estimated relative error in the exact exchange energy is below some threshold. This allows the speed savings to be realized for the first several iterations, while still producing results at full accuracy. This functionality might be necessary

for high-precision calculations of large unit cells, *e.g.*, for forces. However, the error from using single-precision is small, and double-precision calculations of the exact exchange are largely unnecessary. In what follows I will refer to results and timings of the unmodified, fully double-precision code (DP), the modified code that uses single-precision for the exact exchange (SP), and calculations that start with SP and fall back to SP (SDP). To show the success of this single-precision implementation I compare SP and DP calculations for both insulating and metallic systems and both large and small cell sizes.

## III. RESULTS

Example calculations are shown for a variety of systems and covering several different calculation objectives. In each case, the goal is to reproduce the results of a full double-precision calculation. First, the accuracy of SP calculations of the electron energies is assessed by comparing total and individual band energies, structural parameters of ZnS polymorphs, and relative energies of Mn doping sites in lithium tintanate. Next, the accuracy of calculated forces is examined using 64-molecule water cells, and the electron orbitals themselves are investigated through x-ray absorption calculations of  $VO_2$ . Finally, the limitations of single precision are explored using Si supercells.

For all calculations additional convergence parameters are listed in the appendix. Pseudopotentials were taken from the PseudoDojo collection<sup>16,17</sup> with non-linear core corrections removed and generated using the ONCVPSP code.<sup>18,19</sup> To converge the calculations planewave cut-offs and k-point sampling were chosen to ensure a target total energy convergence of  $1.4 \times 10^{-4}$  Ry per atom (2 meV per atom), and each self-consistent calculation was set to a tolerance of  $5 \times 10^{-8}$  Ry per atom.

#### A. Energies

As an initial assessment, the effect of single-precision exchange on the precision of electron energies in three crystalline materials are examined: Si, rocksalt ZnS, and Cu. For silicon, the lattice constants are taken from experiment,<sup>20</sup> while for both ZnS and Cu the structures were relaxed. For all three materials the HSE exchange-correlation potential was used.<sup>21,22</sup> The HSE functional is a screened hybrid, where the exact exchange in Eq. 1 is reduced by a

	Prec.	Time	$\mathrm{E}_{\mathrm{tot}}$	RMS	$\Delta_{max}$	Mem.
		(s)	(Ry)	(eV)	(eV)	(GB)
Si	DP	808	-16.8099	-	-	24.8
	$\operatorname{SP}$	292	-16.8099	$2.1 \times 10^{-4}$	$1.7 \times 10^{-3}$	13.3
ZnS	DP	5424	-684.937	-	-	45.3
(rs)	$\operatorname{SP}$	2724	-684.937	$2.5 \times 10^{-6}$	$7.4 \times 10^{-6}$	23.6
Cu	DP	5640	-376.314	-	-	43.5
	$\operatorname{SP}$	2612	-376.314	$6.1 \times 10^{-7}$	$2.5 \times 10^{-6}$	23.4

TABLE I. The runtime, total energies, RMS errors, maximum difference, and approximate peak memory usage for calculating the exact exchange using double and single-precision. The RMS error of using single precision is calculated using 8 total bands for Si (4 valence and 4 conduction), 21 for ZnS (8 semi-core, 9 valence, and 4 conduction), and 14 for copper (including 4 semi-core). The maximum absolute difference  $\Delta_{max}$  is determined over the same range of bands.

range-dependent function. While screened hybrids tend to converge with fewer k-points than hybrid functionals, the computational cost is otherwise the same. The runtime, total energy, and root-mean-square (RMS) band energies between the two methods of running the exact exchange for each crystalline material are shown in Table I. The deviations between the double-precision and single-precision runs are negligible, but the time and memory savings of single precision are substantial.

Next, I show the effect of the single-precision exact exchange on the equation of state for Cu and several phases of ZnS. For each, the energy as a function of volume was fit to the 3rd order Birch–Murnaghan equation using 9 volumes from approximately 94 % to 108 % of the equilibrium volume.<sup>23</sup> These results are summarized in Table II. While phase diagrams are a useful output of high-throughput studies, their utility is hampered by shortcomings of standard density functionals. Calculated lattice parameters are often accurate to a few percentage points, but relative stability between phases can be mis-ordered. These shortcomings can be partially mitigated using hybrid density functionals. The SP calculations for both metallic copper and the three insulating phases of ZnS are able to reproduce the structural parameters to better than 0.05 %, exceeding the predictive accuracy of the DFT calculations.

	Prec.	a	c	B <sub>0</sub>	B' <sub>0</sub>	$\Delta E_{tot}$
		(a.u.)	(a.u.)	(GPa.)		$(\mathrm{meV})$
Cu	DP	3.423	-	132.8	5.244	-
	$\operatorname{SP}$	3.423	_	132.8	5.243	_
ZnS	DP	10.281	-	72.29	4.374	0.0
(zb)	$\operatorname{SP}$	10.281	-	72.29	4.373	0.1
ZnS	DP	7.244	11.867	74.88	4.261	160.9
(w)	$\operatorname{SP}$	7.244	11.867	74.88	4.288	160.9
ZnS	DP	9.614	-	92.96	4.408	699.1
(rs)	$\operatorname{SP}$	9.614	-	92.98	4.408	699.2

TABLE II. Calculated lattice vectors (a, c), bulk modulus  $(B_0)$ , and the first derivative of the bulk modulus with respect to pressure  $(B'_0)$  for copper and ZnS [zincblende (zb), wurtzite (w), and rocksalt (rs)]. For the polymorphs of ZnS, the energy per formula unit  $(\Delta E_{tot})$  is listed relative to the DP calculation of the zincblende phase.

Third, the SP implementation was tested for moderately-sized, 42-atom, metallic systems by investigating the relative energies of dopant sites in AB<sub>2</sub>O<sub>4</sub>-type spinel lithium titanate. In the case of insulating systems, localization techniques such as Wannier projection yield substantial improvements to the computational time needed for exact exchange calculations.<sup>3,24</sup> For metallic systems, localization is less effective, but, as shown previously for copper, there is little change in accuracy when using single precision on gapless systems. Spinel lithium titanate has been suggested as a possible lithium battery anode material, though its poor electrical conductivity has lead to attempts to add dopants.<sup>25,26</sup> Previous work has shown that Mn dopants likely reside on Li sites through a combination of firstprinciples modeling and x-ray absorption measurements.<sup>27</sup> A small portion of the analysis of Ref. 27 is reproduced here. The relative energies of substituting Mn on various Li sites in  $Li_7Ti_{10}MnO_{24}$  is calculated using both HSE and the PBE-sol functionals,<sup>28</sup> the latter including simplified Hubbard-U corrections applied to the Ti and Mn d-orbitals.<sup>29</sup> As shown in Table III the relative binding of the single-precision and double-precision HSE are almost the same. The PBE-sol calculations show that Mn doping the  $Li-8a_2$  site is about 0.1 eV per formula unit more stable than doping the  $Li-8a_1$  site. In the HSE calculations, these

Site	PBE-sol	HSE-DP	HSE-SP	$\Delta_{\mathrm{HSE}}$
Li-8a <sub>3</sub>	-	-	-	-0.001
$Li-8a_2$	0.104	0.176	0.175	0.000
$Li-8a_1$	0.219	0.173	0.173	0.000
$Li-16d_2$	0.313	0.276	0.275	0.000
$Li-16d_1$	0.509	0.376	0.375	0.000

TABLE III. Calculated energy penalties in eV for substituting Mn onto various Li sites with respect to the energy of the Li-8a<sub>3</sub> site which is found to be the most favorable. Site labels are taken from Ref. 27. The absolute energy difference between the DP and SP HSE runs,  $\Delta_{\text{HSE}}$ , is found to be nearly 0 eV for each structure.

two sites are degenerate. However, Mn residing on the  $\text{Li-8a}_3$  site is still found to be the most stable with both functionals.

### B. Forces

Switching from crystalline systems to liquid water, the effect of reduced precision on the calculation of atomic forces will now be evaluated. First-principles simulations of liquid water and aqueous systems are hampered by the requirements of large cells to incorporate density fluctuations and the need to include nuclear quantum effects to properly capture the hydrogen motion. Recently, this has been addressed by training machine-learning models to simulate the atomic forces with DFT-level accuracy.<sup>30</sup> These models can be trained on moderately sized simulations, and then the trained model can be run with significantly lower computation cost, allowing longer simulation times and larger cells.

Snapshots were generated by the Deep Potential Molecular Dynamics model.<sup>30–32</sup> This model was trained using the PBE0 functional<sup>33</sup> and Tkatchenko-Scheffler approximation to the van der Waals interactions.<sup>34</sup> The calculations here use both. The snapshots were generated using path-integral molecular dynamics with 8 beads and NPT conditions: 64 molecules, 300 K, and 100 kPa. (1 bar).<sup>35</sup> For a DP (SP) run the total memory required per snapshot is 33 GB (18 GB), and the runtime is approximately 169 min. (81 min.). The forces on each atom are compared between the double-precision and single-precision exact exchange using 4 snapshots spaced more than 1.2 ps apart across all 8 beads (for a total of 6144 forces).



FIG. 1. (a) The difference in the magnitudes of the forces between the DP and SP calculations binned into histograms 0.3 meV/a.u. wide. (b) The angle between the DP and SP force vectors binned into histograms 0.6 mrads wide. Note the logscale for both.

The difference between the two methods is negligible with a RMS of only 0.142 meV/a.u. and a maximum discrepancy of 2.56 meV/a.u., while the average magnitude of the forces was 800 times larger at 0.12 eV/a.u. For comparison, the differences in the forces with and without exact exchange (PBE0 vs. PBE) are substantial with an RMS of over 0.25 eV/a.u. Fig. 1(a) shows the distribution of the deviations in the magnitude of the forces between the double-precision and single-precision calculations: for each atom i,  $\Delta F_i = \|F_i^{\text{DP}}\| - \|F_i^{\text{SP}}\|$ . In Fig. 1(b) the angle between the force vectors is shown:  $\theta_i = \cos^{-1}[F_i^{\text{DP}} \cdot F_i^{\text{SP}}/(\|F_i^{\text{DP}}\| \|F_i^{\text{SP}}\|)]$ . The average deviation in the angle is less than 19  $\mu$ rads and the RMS is 193  $\mu$ rads.

### C. Electron orbitals

So far the SP results have been shown to be accurate for the electron energies (both band energies and the total energy) as well as density response (equation of state and forces). As a final test I consider the effects of the precision of the exact exchange on x-ray absorption calculations. In x-ray absorption, the electron orbitals are probed directly in the form of the transition matrix elements where the electron-photon operator promotes an electron from a core level into the conduction band. Using the OCEAN code,<sup>36,37</sup> the x-ray absorption spectra of the  $M_1$  insulating phase of VO<sub>2</sub> are calculated, looking at the O K and V L<sub>23</sub> edges. The lattice constants and atomic positions were taken from experiment.<sup>38</sup> The OCEAN



FIG. 2. Calculated V L and O K edge x-ray absorption of insulating VO<sub>2</sub> compared with experiment.<sup>39</sup> Separate calculations were carried out using PBE (blue, dashed) or HSE (red, dotted) functionals for the basis states for the Bethe-Salpeter equation. Along the bottom, the solid cyan line shows the difference between the DP and SP HSE calculated spectra, multiplied by a factor of 1000.

calculation uses DFT electron orbitals as the basis for solving the Bethe-Salpeter equation. The measured data are compared to calculations using the PBE functional with both SP and DP HSE calculations in Fig. 2. All three calculations use the same screened corehole potential, taken from the PBE calculation. The difference between the DP and SP calculations is negligible. In the plot the difference curve is multiplied by a factor of 1000 to make any differences visible.

Within DFT, local or semi-local exchange-correlation potentials do not correctly describe the insulating phases of VO<sub>2</sub>. Several attempts to rectify this have been attempted, such as including exact exchange<sup>40,41</sup> or self-consistent self-energy calculations.<sup>42</sup> As criteria, structure and the band gap only reflect the occupied states and lowest-lying conduction band. Electronic excitations such as optical or x-ray spectra provide the additional comparison of how hybrid functionals modify the unoccupied states of material. These calculations come with an increase in computational cost since they require also determining the unoccupied orbitals. Optical excitations in VO<sub>2</sub> were calculated using self-consistent self-energy showing a dramatic change in the absorption spectra.<sup>43</sup> However, x-ray absorption calculations of the V L<sub>23</sub> edge using a similar method found, despite the opening of a gap in the ground state, little effect on the calculated spectra.<sup>44</sup> Here clear differences are evident between the BSE x-ray absorption using PBE versus HSE. In VO<sub>2</sub>, the unoccupied Vd orbitals are split by the symmetry depending on their overlap with the oxygen orbitals.<sup>45</sup> The HSE decreases the occupancy of the lower-energy orbitals, resulting in a stronger transition near 515 eV at the V L<sub>3</sub> (522 eV at the L<sub>2</sub>). There is also an increase in the crystal field splitting from the PBE to the HSE calculation, apparent in the shift of the peaks near 517.5 eV and 524.5 eV to higher energy. The features from 527 eV to 533 eV also reflect the structure of the Vd orbitals, but these are transitions from the O 1s to p-like orbitals that have hybridized with the Vd states. The HSE calculation gives reduced intensity in this region, indicating less hybridization between the Op and Vd orbitals. In general, the HSE calculation shows worse agreement with the measured x-ray absorption. It was suggested in Ref. 41 that the vanadium oxides require a reduced percentage of exact exchange, but a full investigation of the exchange dependence of the x-ray spectra is beyond the scope of this work.

#### D. Limitations with system size

So far, the SP implementation has produced results in agreement with calculations using the unmodified DP implementation. However, if the precision target for the total energy is set small enough, the errors from rounding to single precision will prevent convergence. In such cases, falling back to double precision is required. To demonstrate this, silicon cells with 2–64 atoms were calculated using SP, DP, and SDP for the exact exchange, using total energy convergence targets ranging from  $10^{-4}$  Ry to  $10^{-10}$  Ry. In contrast, the previously discussed calculations used a per atom convergence target. In table IV, the number of iterations required to converge the exchange operator is shown for several supercells of silicon. Here an iteration is every time a new exchange operator is constructed (Eqs. 2–5), following the convergence of the electron orbitals and densities using the previous exchange operator. As expected, when the target precision is very small the SP calculation does not converge or converges to the wrong value. The SPD runs take at most one additional iteration to converge as compared to the DP ones, but the majority of the iterations are done at single precision resulting in time savings over a DP calculation.

From Table IV, the SP calculations are able to correctly converge to a target accuracy down to about 1 part in  $10^6$  of the exact exchange energy. For large system sizes, and

	Prec.	$E_{EXX}$	Iterations for target precision			
		(Ry)	$10^{-4}$	$10^{-6}$	$10^{-8}$	$10^{-10}$
	SP		3	4	$8^{\ddagger}$	*
$Si_2$	DP	$-8.4 \times 10^{-1}$	3	4	6	8
	SDP		3	4/-	5/2	5/3
	SP		3	5	$6^{\ddagger}$	*
$\mathrm{Si}_8$	DP	$-3.4  imes 10^0$	3	5	6	8
	SDP		3/-	5/-	5/2	5/3
	SP		4	5	$7^{\ddagger}$	*
$\mathrm{Si}_{16}$	DP	$-6.7 \times 10^0$	4	5	7	9
	SDP		4/-	5/-	5/2	5/4
	SP		4	$6^{\ddagger}$	*	*
$\mathrm{Si}_{64}$	DP	$-2.7  imes 10^1$	4	6	7	9
	SDP		4/-	5/2	5/2	5/4

TABLE IV. The number of iterations needed to converge the total energy of various silicon super cells to a given absolute precision. For reference the exact exchange energy is listed as well. An asterisk denotes a calculation that will *not converge* due to the limited precision of the SP. The SDP calculations switch from SP to DP once the relative accuracy of the exact exchange has dropped below  $10^{-7}$ . For SDP calculations the number of iterations shown are split into SP/DP. The SP calculations marked with a ‡ did not converge to the DP total energy within the target precision.

correspondingly large energies, this will limit the achievable level of absolute convergence. The total energy is extensive. When considering super cells or polymorphs, like the ZnS phases shown earlier, equivalent levels of precision per formula unit are achievable using the SP exact exchange irrespective of system size. The same does not necessarily hold when considering a small subset of a larger system, such as adsorbates on a surface, or impurities like dopants. In these cases, strict convergence and SDP may be required to calculate forces or total energies. As an example, in phosphorous-doped silicon the energy difference between holding the atoms at their bulk silicon positions and allowing them to relax in response to the dopant has been calculated to be approximately 30 meV.<sup>46</sup> Extrapolating the results in Table IV, an SP calculation of Si<sub>383</sub>P would entail an error approaching 1 meV.

### IV. SUMMARY

I have shown that the computational cost of hybrid DFT calculations, both total run time and required memory, can be nearly cut in half through the use of single-precision arithmetic. Any errors from this approximation are below the expected accuracy of the calculations themselves. This improvement does not require insulating systems nor localization methods, and is easily adaptable to any plane-wave DFT method. While the current implementation requires norm-conserving pseudopotentials and collinear spin, removing these limitations should be straightforward. For small to moderate unit cells, the single-precision method introduced here is a substantial improvement for hybrid DFT calculations.

Part of the success here is due to the relatively small contribution of the exact exchange to the total energy of the electrons. For the systems investigated here, the ratio of the Fock energy to total energy ranges from 2% in Cu to 6% in liquid water. On the other hand the differences in Fock energies between SP and DP calculations are below 1 part in  $10^{-6}$  (compared to the relative precision limit of single-precision numbers of approximately  $10^{-7}$ ). Taken together, the relative error in the total energy of using the SP exact exchange is around  $10^{-7}$  or better. This type of reduced precision approach could also be used for evaluating other parts of the DFT Hamiltonian, such as the higher-order semi-local terms for meta-GGA functionals.

The fundamental system-size scaling of hybrid DFT is unchanged by the reduced-precision implementation. For large systems, some other method or approximation, such as the localization techniques mentioned previously, is necessary to make the calculation feasible. The use of mixed or reduced precision may also be applicable to these methods and should be investigated.

### Appendix A: Convergence parameters

The water calculations were carried out using a single k-point and the  $\Gamma$ -point specific routines within QUANTUM ESPRESSO. For the water and ZnS systems only the occupied bands were included in the calculation, while for Si, rocksalt ZnS, and Cu the number of bands is specified in Table I. In LTO, 200 bands were used for 300 valence electrons. The ions in the LTO calculation were relaxed to reduce the forces below  $10^{-3}$  Ry/a.u. For VO<sub>2</sub>

	E <sub>cut</sub> (Ry.)	k-points
$Si_2, Si_8$	50	$8 \times 8 \times 8$
$\mathrm{Si}_{16}$	50	$6 \times 6 \times 6$
$\mathrm{Si}_{64}$	50	$3 \times 3 \times 3$
Cu	110	$8 \times 8 \times 8$
ZnS $(rs)$	100	$6 \times 6 \times 6$
ZnS $(zb)$	100	$4 \times 4 \times 4$
ZnS $(w)$	100	$6 \times 6 \times 3$
LTO	100	$2 \times 2 \times 1$
$\rm H_2O$	90	$1 \times 1 \times 1$
$VO_2$	110	$6 \times 6 \times 6$

TABLE V. The plane-wave energy cutoff and k-point sampling for each of the systems studied in this paper.

120 bands were calculated (100 valence electrons), but the number of ACE projectors was set to only 80. For Cu, LTO, and VO<sub>2</sub>, Fermi-Dirac smearing was used with a broadening parameter of 0.02 Ry.

### Appendix B: Data availability

The data that support the findings of this study are available from the corresponding author upon reasonable request.

# REFERENCES

- <sup>1</sup>S. Kümmel and L. Kronik, "Orbital-dependent density functionals: Theory and applications," Rev. Mod. Phys. **80**, 3–60 (2008).
- <sup>2</sup>L. Lin, "Adaptively compressed exchange operator," Journal of Chemical Theory and Computation **12**, 2242–2249 (2016).
- <sup>3</sup>X. Wu, A. Selloni, and R. Car, "Order-*n* implementation of exact exchange in extended insulating systems," Phys. Rev. B **79**, 085102 (2009).

- <sup>4</sup>F. Gygi and I. Duchemin, "Efficient computation of hartree–fock exchange using recursive subspace bisection," Journal of Chemical Theory and Computation **9**, 582–587 (2013).
- <sup>5</sup>A. Damle, L. Lin, and L. Ying, "Compressed representation of kohn–sham orbitals via selected columns of the density matrix," Journal of Chemical Theory and Computation **11**, 1463–1469 (2015).
- <sup>6</sup>I. Carnimeo, S. Baroni, and P. Giannozzi, "Fast hybrid density-functional computations using plane-wave basis sets," Electronic Structure **1**, 015009 (2019).
- <sup>7</sup>N. Marzari, A. A. Mostofi, J. R. Yates, I. Souza, and D. Vanderbilt, "Maximally localized wannier functions: Theory and applications," Rev. Mod. Phys. **84**, 1419–1475 (2012).
- <sup>8</sup>C. Brouder, G. Panati, M. Calandra, C. Mourougane, and N. Marzari, "Exponential localization of wannier functions in insulators," Phys. Rev. Lett. **98**, 046402 (2007).
- <sup>9</sup>V. P. Vysotskiy and L. S. Cederbaum, "Accurate quantum chemistry in single precision arithmetic: Correlation energy," Journal of Chemical Theory and Computation 7, 320–326 (2011).
- <sup>10</sup>P. Pokhilko, E. Epifanovsky, and A. I. Krylov, "Double precision is not needed for manybody calculations: Emergent conventional wisdom," Journal of Chemical Theory and Computation 14, 4088–4096 (2018).
- <sup>11</sup>M. Guidon, F. Schiffmann, J. Hutter, and J. VandeVondele, "Ab initio molecular dynamics using hybrid density functionals," The Journal of Chemical Physics **128**, 214104 (2008).
- <sup>12</sup>H. Shang, Z. Li, and J. Yang, "Implementation of exact exchange with numerical atomic orbitals," The Journal of Physical Chemistry A **114**, 1039–1043 (2010), pMID: 20070129.
- <sup>13</sup>X. Gonze, B. Amadon, G. Antonius, F. Arnardi, L. Baguet, J.-M. Beuken, J. Bieder, F. Bottin, J. Bouchet, E. Bousquet, N. Brouwer, F. Bruneval, G. Brunin, T. Cavignac, J.-B. Charraud, W. Chen, M. Côté, S. Cottenier, J. Denier, G. Geneste, P. Ghosez, M. Giantomassi, Y. Gillet, O. Gingras, D. R. Hamann, G. Hautier, X. He, N. Helbig, N. Holzwarth, Y. Jia, F. Jollet, W. Lafargue-Dit-Hauret, K. Lejaeghere, M. A. L. Marques, A. Martin, C. Martins, H. P. C. Miranda, F. Naccarato, K. Persson, G. Petretto, V. Planes, Y. Pouillon, S. Prokhorenko, F. Ricci, G.-M. Rignanese, A. H. Romero, M. M. Schmitt, M. Torrent, M. J. van Setten, B. V. Troeye, M. J. Verstraete, G. Zérah, and J. W. Zwanziger, "The abinit project: Impact, environment and recent developments," Comput. Phys. Commun. **248**, 107042 (2020).

- <sup>14</sup>P. Giannozzi, O. Andreussi, T. Brumme, O. Bunau, M. B. Nardelli, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, M. Cococcioni, N. Colonna, I. Carnimeo, A. D. Corso, S. de Gironcoli, P. Delugas, R. A. DiStasio, A. Ferretti, A. Floris, G. Fratesi, G. Fugallo, R. Gebauer, U. Gerstmann, F. Giustino, T. Gorni, J. Jia, M. Kawamura, H.-Y. Ko, A. Kokalj, E. Küçükbenli, M. Lazzeri, M. Marsili, N. Marzari, F. Mauri, N. L. Nguyen, H.-V. Nguyen, A. O. de-la Roza, L. Paulatto, S. Poncé, D. Rocca, R. Sabatini, B. Santra, M. Schlipf, A. P. Seitsonen, A. Smogunov, I. Timrov, T. Thonhauser, P. Umari, N. Vast, X. Wu, and S. Baroni, "Advanced capabilities for materials modelling with quantum ESPRESSO," Journal of Physics: Condensed Matter **29**, 465901 (2017).
- <sup>15</sup>P. Giannozzi, S. Baroni, N. Bonini, M. Calandra, R. Car, C. Cavazzoni, D. Ceresoli, G. L. Chiarotti, M. Cococcioni, I. Dabo, *et al.*, "Quantum espresso: a modular and open-source software project for quantum simulations of materials," J. Phys. Condens. Matter **21**, 395502 (2009).
- <sup>16</sup>M. van Setten, M. Giantomassi, E. Bousquet, M. Verstraete, D. Hamann, X. Gonze, and G.-M. Rignanese, "The pseudodojo: Training and grading a 85 element optimized norm-conserving pseudopotential table," Computer Physics Communications **226**, 39 54 (2018).
- <sup>17</sup>http://www.pseudo-dojo.org PBE-sol Scalar-relativstic v. 0.4.
- <sup>18</sup>D. R. Hamann, "Optimized norm-conserving vanderbilt pseudopotentials," Phys. Rev. B 88, 085117 (2013).
- <sup>19</sup>The open-source code ONCVPSP is avaiable at http://www.mat-simresearch.com v. 3.3.1.
- <sup>20</sup>R. W. G. Wyckoff, *Crystal Structures*, 2nd ed. (Interscience Publishers, New York, 1963).
- <sup>21</sup>J. Heyd, G. E. Scuseria, and M. Ernzerhof, "Hybrid functionals based on a screened coulomb potential," The Journal of Chemical Physics **118**, 8207–8215 (2003).
- <sup>22</sup>J. Heyd, G. E. Scuseria, and M. Ernzerhof, "Erratum: "hybrid functionals based on a screened coulomb potential" [j. chem. phys. 118, 8207 (2003)]," The Journal of Chemical Physics **124**, 219906 (2006).
- <sup>23</sup>F. Birch, "Finite elastic strain of cubic crystals," Phys. Rev. **71**, 809–824 (1947).
- <sup>24</sup>H.-Y. Ko, J. Jia, B. Santra, X. Wu, R. Car, and R. A. DiStasio Jr., "Enabling large-scale condensed-phase hybrid density functional theory based ab initio molecular dynamics. 1. theory, algorithm, and performance," Journal of Chemical Theory and Computation 0,

null (0).

- <sup>25</sup>D. Capsoni, M. Bini, V. Massarotti, P. Mustarelli, G. Chiodelli, C. B. Azzoni, M. C. Mozzati, L. Linati, and S. Ferrari, "Cations distribution and valence states in mn-substituted li4ti5o12 structure," Chemistry of Materials **20**, 4291–4298 (2008).
- <sup>26</sup>H. Kaftelen, M. Tuncer, S. Tu, S. Repp, H. Göçmez, R. Thomann, S. Weber, and E. Erdem, "Mn-substituted spinel li4ti5o12 materials studied by multifrequency epr spectroscopy," J. Mater. Chem. A 1, 9973–9982 (2013).
- <sup>27</sup>H. Singh, M. Topsakal, K. Attenkofer, T. Wolf, M. Leskes, Y. Duan, F. Wang, J. Vinson, D. Lu, and A. I. Frenkel, "Identification of dopant site and its effect on electrochemical activity in mn-doped lithium titanate," Phys. Rev. Materials 2, 125403 (2018).
- <sup>28</sup>J. P. Perdew, A. Ruzsinszky, G. I. Csonka, O. A. Vydrov, G. E. Scuseria, L. A. Constantin, X. Zhou, and K. Burke, "Restoring the density-gradient expansion for exchange in solids and surfaces," Phys. Rev. Lett. **100**, 136406 (2008).
- <sup>29</sup>M. Cococcioni and S. de Gironcoli, "Linear response approach to the calculation of the effective interaction parameters in the LDA+U method," Phys. Rev. B **71**, 035105 (2005).
- <sup>30</sup>L. Zhang, J. Han, H. Wang, R. Car, and W. E, "Deep potential molecular dynamics: A scalable model with the accuracy of quantum mechanics," Phys. Rev. Lett. **120**, 143001 (2018).
- <sup>31</sup> "Deep potential: A general representation of a many-body potential energy surface," Communications in Computational Physics 23, 629–639 (2018).
- <sup>32</sup>H.-Y. Ko, L. Zhang, B. Santra, H. Wang, W. E, R. A. D. Jr, and R. Car, "Isotope effects in liquid water via deep potential molecular dynamics," Molecular Physics **117**, 3269–3281 (2019).
- <sup>33</sup>J. P. Perdew, M. Ernzerhof, and K. Burke, "Rationale for mixing exact exchange with density functional approximations," The Journal of Chemical Physics **105**, 9982–9985 (1996).
- <sup>34</sup>A. Tkatchenko and M. Scheffler, "Accurate molecular van der waals interactions from ground-state electron density and free-atom reference data," Phys. Rev. Lett. **102**, 073005 (2009).
- <sup>35</sup>M. Calegari and R. Car, Private communication.
- <sup>36</sup>J. Vinson, J. J. Rehr, J. J. Kas, and E. L. Shirley, "Bethe-salpeter equation calculations of core excitation spectra," Phys. Rev. B 83, 115106 (2011).

- <sup>37</sup>K. Gilmore, J. Vinson, E. Shirley, D. Prendergast, C. Pemmaraju, J. Kas, F. Vila, and J. Rehr, "Efficient implementation of core-excitation bethe-salpeter equation calculations," Comput. Phys. Comm. **197**, 109 – 117 (2015).
- <sup>38</sup>G. Andresson, "Studies on vanadium oxides ii. the crystal structure of vanadium dioxide," Acta Chemica Scandinavica **10**, 623 (1956).
- <sup>39</sup>M. Kumar, J. P. Singh, K. H. Chae, J. Park, and H. H. Lee, "Annealing effect on phase transition and thermochromic properties of vo2 thin films," Superlattices and Microstructures **137**, 106335 (2020).
- <sup>40</sup>V. Eyert, "vo<sub>2</sub>: A novel view from band theory," Phys. Rev. Lett. **107**, 016401 (2011).
- <sup>41</sup>S. Xu, X. Shen, K. A. Hallman, R. F. Haglund, and S. T. Pantelides, "Unified bandtheoretic description of structural, electronic, and magnetic properties of vanadium dioxide phases," Phys. Rev. B 95, 125105 (2017).
- <sup>42</sup>M. Gatti, F. Bruneval, V. Olevano, and L. Reining, "Understanding correlations in vanadium dioxide from first principles," Phys. Rev. Lett. **99**, 266402 (2007).
- <sup>43</sup>M. Gatti, F. Sottile, and L. Reining, "Electron-hole interactions in correlated electron materials: Optical properties of vanadium dioxide from first principles," Phys. Rev. B 91, 195137 (2015).
- <sup>44</sup>J. Vinson and J. J. Rehr, "Ab initio bethe-salpeter calculations of the x-ray absorption spectra of transition metals at the *l*-shell edges," Phys. Rev. B 86, 195135 (2012).
- <sup>45</sup>M. Abbate, F. M. F. de Groot, J. C. Fuggle, Y. J. Ma, C. T. Chen, F. Sette, A. Fujimori, Y. Ueda, and K. Kosuge, "Soft-x-ray-absorption studies of the electronic-structure changes through the vo<sub>2</sub> phase transition," Phys. Rev. B **43**, 7263–7266 (1991).
- <sup>46</sup>L. Greenman, H. D. Whitley, and K. B. Whaley, "Large-scale atomistic density functional theory calculations of phosphorus-doped silicon quantum bits," Phys. Rev. B 88, 165102 (2013).