# Improving the performance of hybrid functional-based molecular dynamics simulation through screening of Hartree-Fock exchange forces

Guido Falk von Rudorff,[†] Rasmus Jakobsen,[†] Kevin M Rosso,[‡] and Jochen Blumberger[*,†]

[†]*Department of Physics and Astronomy, University College London, London WC1E 6BT, UK*

[‡]*Pacific Northwest National Laboratory, Richland, Washington, United States.*

E-mail: j.blumberger@ucl.ac.uk

**Abstract**

Density functional theory-based molecular dynamics calculations of condensed phase systems often benefit from the use of hybrid functionals. However, their use is computationally very demanding and severely limits the system size and time scale that can be simulated. Several methods have been introduced to accelerate hybrid functional molecular dynamics including Schwarz screening and the auxiliary density matrix method (ADMM). Here we present a simple screening scheme that can be applied in addition to these methods. It works by examining Hartree-Fock exchange (HFX) integrals and subsequently excluding those that contribute very little to any nuclear force component. The resultant force error is corrected by a history-dependent extrapolation scheme. We find that for systems where the calculation of HFX forces is a major

bottleneck, a large fraction of the integrals can be neglected without introducing significant errors in the nuclear forces. For instance, for a sample of bulk CoO, 92% of the HFX integrals that have passed Schwarz screening within the ADMM approach can be neglected leading to a performance gain of a factor of 3 at a negligible error in nuclear forces ($\leq 5 \times 10^{-4}$ H bohr$^{-1}$). We also show that total energy conservation and solvation structures are not adversely affected by the screening method.

# 1 Introduction

Density functional theory (DFT) has become a standard tool for investigating the structure and energetics of a wide range of materials. Depending on the physical property in question, different levels of approximations to the exchange-correlation functional are required. While functionals based on the Generalized Gradient Approximation (GGA) are computationally inexpensive and typically give good results for lattice constants[1] and absorption energies,[2] they often fail to predict band gaps,[3] defect levels and electronic coupling matrix elements.[4] Those failures have been attributed to the electron self interaction error (SIE) of GGA functionals[5] and several correction schemes have been devised to mitigate this problem including e.g. Perdew-Zunger self-interaction correction,[6,7] DFT+U,[8] constrained DFT (CDFT),[9] and enforcement of the integer-derivative discontinuity within a generalized Koopman's theorem formulation[10] or by optimal tuning of range-separated hybrid functionals.[11]

Most (but not all) of the DFT methods that successfully address the problem of SIE rely on the calculation of Hartree-Fock-Exchange (HFX).[12] Yet, this comes at a significantly increased computational effort compared to GGA-DFT due to the quartic scaling of the Hartree-Fock (HF) method with respect to the number of basis functions. This is particularly problematic for large condensed phase systems. Several schemes are already in use to mitigate the high computational cost. Firstly, integral screening based on the Schwarz inequality[13] typically improves the scaling from exponent 4 to 2.2-2.3.[14] For non-conducting materials, the density matrix falls off exponentially[15] with distance, which means that any two or

four-center integrals are dominated by local interactions. This improves the performance of hybrid functionals both in short-range and long-range regimes[16] which can be also applied to conducting materials[17] and–together with Schwarz screening–offer linear scaling.[18] While the range separation introduces additional empirical parameters, their precise values are often not critical to the results[19] and may be obtained from first principles in some formulations.[20]

Besides the density functional, another factor contributing to the balance between quality and efficiency is the basis set. In many cases, converging HFX calculations requires prohibitively large basis sets.[21] Here approaches like the Auxiliary Density Matrix Method (ADMM)[22] have proven very successful. In this approach a small basis set is used for the calculation of HFX integrals and the error made is corrected by calculating the difference between small and large basis set at the density functional GGA level of theory. Both the Schwarz screening and ADMM can be combined, as is done e.g. in the CP2K simulation package, resulting in speed-ups of more than an order of magnitude compared to a straightforward HFX calculation.[22]

While Schwarz screening plus ADMM render geometry optimisations and molecular dynamics (MD) simulation of large extended systems feasible at the hybrid DFT level, the overall cost can still be significantly higher than for GGA DFT. In particular for strongly correlated materials,[23] semiconductors[24] and transition metal oxides,[25] where employing HFX drastically improves the results, the calculation of nuclear forces via the molecular orbital derivatives[22] become the bottleneck for the overall calculation. In order to improve the computational efficiency for the force calculation, we suggest here an interaction screening method of HFX force components that is applied in addition to and after Schwarz screening and ADMM. Considering the HFX energy expression

$$E_X^{\mathrm{HF}} = -\frac{1}{2}\sum_{\lambda\sigma\mu\nu}\left[\sum_i C_{\mu i}C_{\sigma i}\right]\left[\sum_j C_{\nu j}C_{\lambda j}\right]\int d\mathbf{r_1}\int d\mathbf{r_2}\frac{\phi_\mu(\mathbf{r_1})\phi_\nu(\mathbf{r_1})\phi_\lambda(\mathbf{r_2})\phi_\sigma(\mathbf{r_2})}{|\mathbf{r_2}-\mathbf{r_1}|}, \quad (1)$$

where $\lambda, \sigma, \mu, \nu$ run over all orbital combinations that have passed Schwarz screening, each

integral is screened against its contributions to the HFX force components. Hence the criterion depends on both the density matrix coefficients $C_{\mu i}$ as well as the overlap of all four orbitals. The error made due to neglect of HFX force components that fall below the screening threshold is estimated and corrected by a history dependent interpolation scheme. We show that the method gives additional efficiency savings of up to a factor of three at negligible errors in the nuclear forces.

# 2 Methods

## 2.1 Profiling of hybrid functional-based MD

We performed short MD simulations of condensed phase systems reflecting our current research interest: liquid water, hematite ($\alpha$-$Fe_2O_3$) and a water/hematite interface to identify the factors limiting computational performance (see section 2.5 for simulation details). We used the profiling components included in CP2K for the Quickstep module. Figure 1 shows the results for sample systems of various size. It is evident that for extended systems the relative costs for HFX-related calculations increases due to the unfavourable scaling behaviour. Moreover, the force evaluation is by far the most expensive part of the HFX calculations for hematite and for the interface with liquid water. Based on this, we aimed to understand how physical properties of the system like local electron density and the distribution of the atoms contribute to the cost of the MD runs and how it can be reduced.

## 2.2 HFX force component analysis

The significant proportion of the time spent on evaluating the HFX forces makes it desirable to identify small contributions that can be safely neglected. However, neglecting any small contributions will lead to large errors if there are many small collinear contributions that sum up. Therefore, we calculated the individual force contributions resulting from each single electron repulsion integral (ERI) in Eqn. 1 for the hematite/liquid water interface.

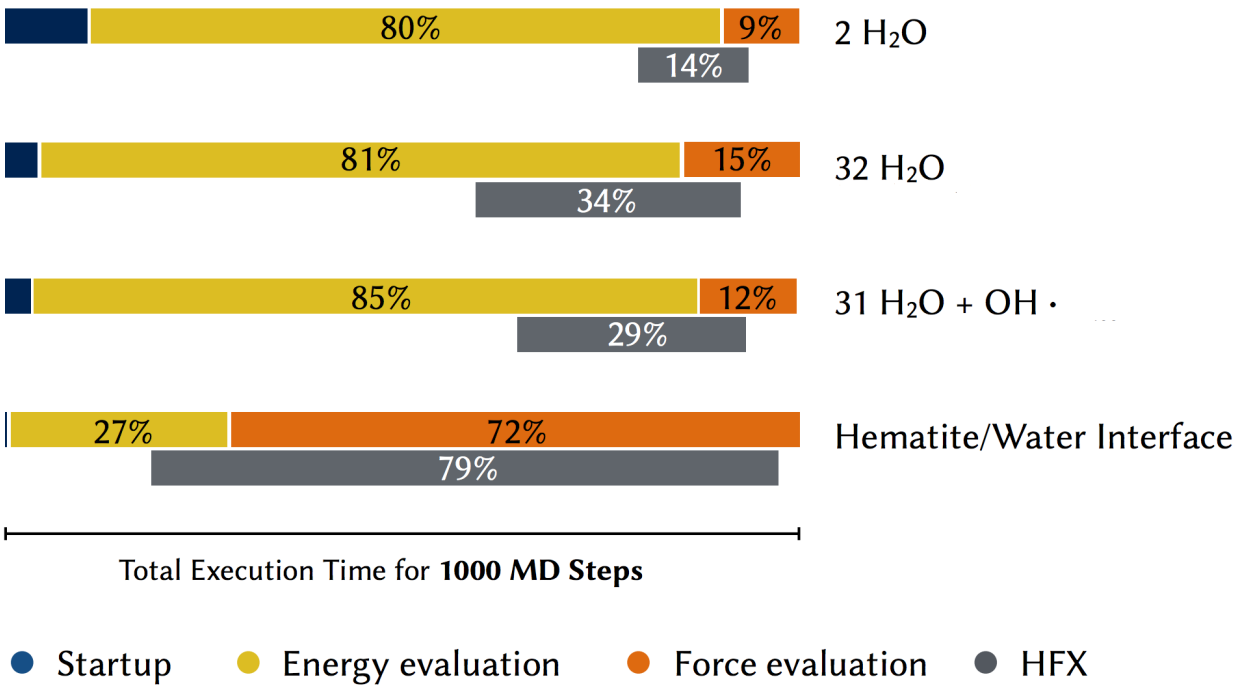


Figure 1: Computational cost for hybrid density functional MD runs. The top bar shows the fraction of the total run time required for startup, energy and force calculation, for each system simulated. The bottom bar shows the fraction of total run time required by HFX-related calculations. Text below system labels specifies length of trajectory generated, number of 24-core nodes, average number of SCF steps and total walltime.

With about 2000 electrons, this produces a tremendous amount of data, which is why we analysed only force components larger than $10^{-10}$ H/bohr for a single MD snapshot, resulting in about $6 \times 10^9$ individual force contributions in total or approx. $1.4 \times 10^7$ per atom which have been stored in a database for analysis. It is vital to note that our setup already uses Schwarz screening and ADMM, and so all interactions presented here would not be screened by these methods.

In our analysis, the HFX force contribution due to each HFX integral is summed up in groups defined by the order of magnitude of the absolute force value. This means that the many small force contributions are aggregated and can be compared to the few large force contributions. From the results in Figure 2, it becomes clear that HFX forces with a magnitude between $10^{-2}$ and $10^{-6}$ H/bohr give comparable contributions when summed up. However, the contribution of forces smaller than $10^{-6}$ H/bohr steadily decreases and amounts to only a few percent to the total HFX force vector. This implies that it is likely that only a small error is made when those small contributions are neglected.

## 2.3 Screening of HFX force components

Making use of the results of the force component analysis requires a way to identify the four-center integrals that have shown to contribute only little without actually calculating them in every time step during the MD. The proximity of configurations for consecutive MD steps allows for transferring the results as long as the conformation has not changed too much in the meantime. Figure 1 in the SI illustrates the similarity of the forces from two consecutive time steps. It is clearly visible that the forces do not change significantly both in the case of light and heavier elements. This is a necessary prerequisite for the method outlined in the following.

Based on our analysis we devise a screening algorithm that is schematically shown in Figure 3. All HFX force contributions are evaluated only at every $n$th MD time step, which we term "update step". At each update step, an exclusion list is constructed listing

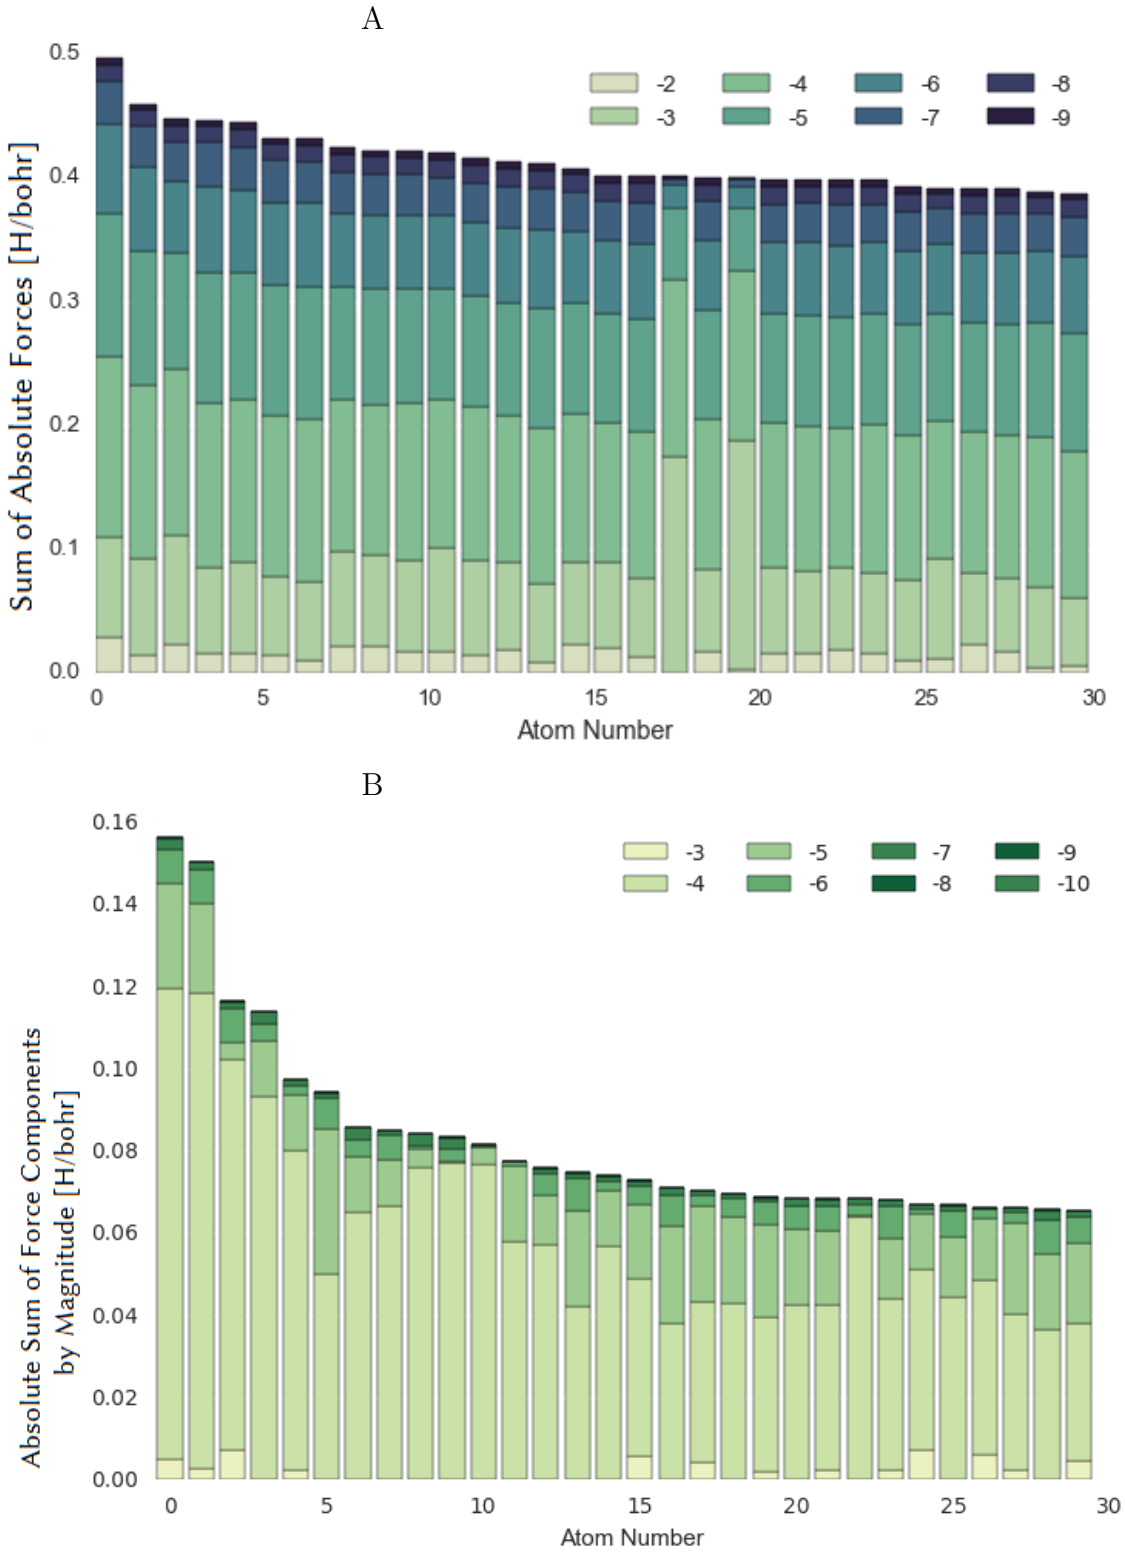


Figure 2: A: Sum of absolute HFX force contributions grouped by order of magnitude of each HFX integral. The 30 atoms from the hematite/liquid water interface with the largest cumulative absolute HFX force are shown. Dark bar segments denote the cumulative effect of small contributions while light segments show the effect of few but larger contributions. All force components have been projected onto the x axis for this graph. B: Absolute sum of HFX force contributions, so force contributions of opposite sign but same magnitude cancel each other out.

all combinations of atoms $(ijkl)$ where all $x$, $y$ and $z$-HFX force components of all four atoms are below a threshold $f_T$. For the timesteps between two update steps, only those combinations $(i'j'k'l')$ are evaluated that are not present in the exclusion list. The latter is updated at every update step. Hence, for a fraction $(n-1)/n$ of time steps only the subset of HFX force contributions is calculated that is above the threshold $f_T$. As we will see below, this leads to a significant gain in efficiency.

The error made by neglecting the HFX force components below the threshold $f_T$ is corrected using a history-dependent extrapolation scheme. We use the fact that along the MD trajectory the HFX forces and their derivatives are smooth and that at each update step the error made for each force component is exactly known. Assuming the system is at update step $q$, the error of each force component at the next update step $q'$ is estimated by linear extrapolation of the errors at $q$ and the $k$ updates preceding $q$. Then the force error at the $n-1$ time steps between $q$ and $q'$ is estimated by a cubic spline interpolation using the estimated errors at $q'$ and the exact errors at $q$ and the $k$ update steps preceding $q$.

## 2.4 Choice of screening parameters

The screening method has three adjustable parameters: the force threshold $f_T$, the repeat rate of the updates $(n)$, and the length of the force history $(k)$. The force threshold $f_T$ is crucial to the overall performance. It should be chosen large enough so as to screen a significant fraction of the four-centre integrals at the time steps between two updates. On the other hand, it should be small enough such that the screening of the force calculation does not produce force errors too large to be mitigated by the force correction scheme. A HFX force component analysis described in section 2.2 may help determine a sensible value for $f_T$ for a given system.

Besides force threshold, the repeat rate sensitively affects the speed-up of the calculation. As before, it should be chosen large enough to reduce the computational effort but sufficiently small so as to keep the force error correction accurate. Assuming that for a given $f_T$ the

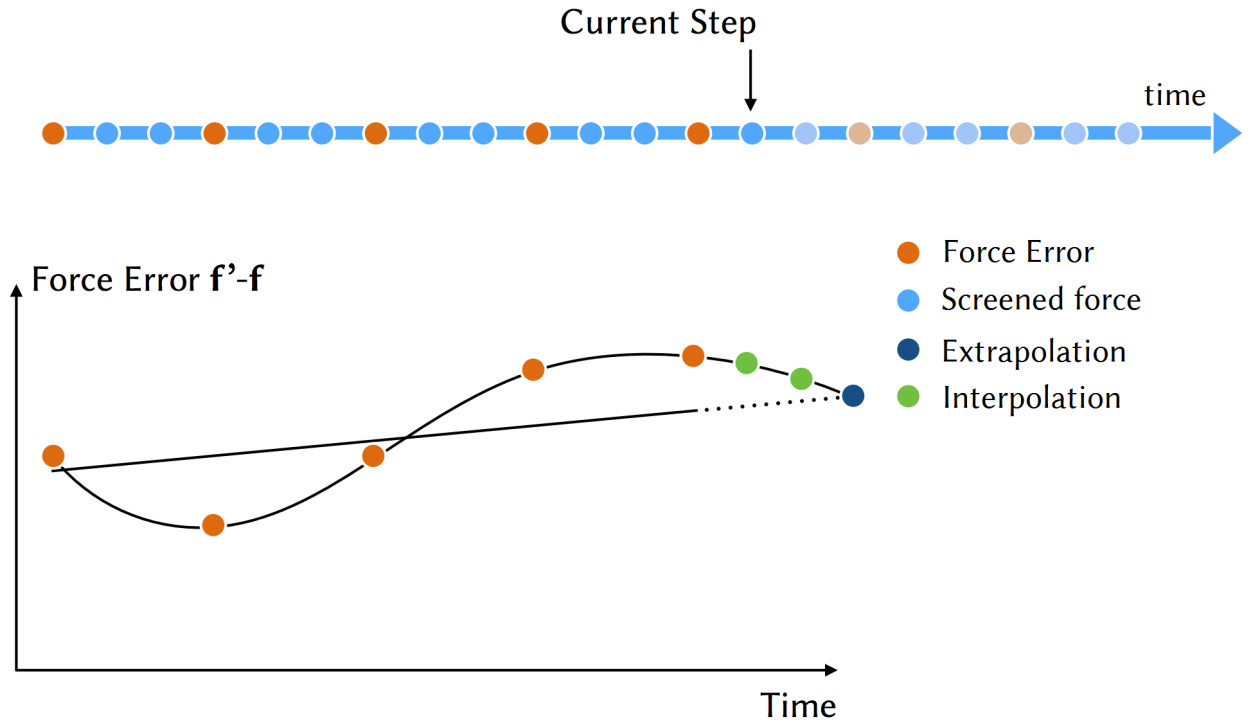


Figure 3: Top panel: Timeline of MD steps. Full HFX calculations ("updates") are carried out at steps shown in orange and HFX integral screening is applied at steps shown in blue. Bottom panel: Schematic of the error correction. The history of differences between full HFX force, $\mathbf{f}'$, and screened HFX force, $\mathbf{f}$, is shown in orange and used to extrapolate the difference for the next full force calculation (dark blue) followed by interpolation to obtain an error estimate for the next MD step (green).

full HFX force calculation for an update takes $m$ times the calculation of a screened force calculation ($m > 1$), the speed-up $S$ achieved by force screening is given by

$$S = \frac{nm}{m + n - 1}. \tag{2}$$

This gives $S \approx m/2$ for $n = m$ and $S \approx m$ for $n = m^2$. Hence, the maximum theoretical speed-up for given $m$, $S = m$, is almost reached for $n = m^2$. Therefore, one would chose $n$ to be between 2 and $m^2$ depending on the performance of the force error corrector and the accuracy requirements for the force calculation.

The third parameter, the history of $k$ previous updates to be included in the correction scheme, is related to the frequency of the fastest motion of the system, e.g. the OH stretch vibration for aqueous systems. This is because the HFX forces are susceptible to the orbital overlap which exponentially decreases with distance. The largest HFX forces result from bonded atoms, such that bond vibrations are dominating the time-dependence of HFX force components. If the history is too long, the linear force extrapolation will be stationary, thereby not leveraging the full potential of the correction, since the current curvature of the error is averaged out. Choosing a very short history makes the system overly sensitive to sudden steep force alterations. In general, the optimum $k$ value will be system dependent. For aqueous solutions we found $k = 5$ to be a good choice (see also simulation details). The impact of $k$ on the computational cost is negligible.

## 2.5 Simulation Details

For benchmarking and validation purposes, we used the following systems: cobalt(II) oxide (CoO),[26] hematite ($\alpha$-$Fe_2O_3$),[27] water dimer, bulk water, an OH radical in bulk water,[7] and a pre-equilibrated hematite/liquid water interface.[28] The transition-metal oxides (TMO) are good test systems because of their complex electronic structure and their strongly correlated properties[29] on one side and the practical interest in their applications, e.g. water

splitting,[30,31] on the other side. The hematite/liquid water interface system combines two disparate regions into one system where hybrid functionals are necessary for a good description of the electronic structure.[32]

In our calculations, we used the Quickstep module[33] from the CP2K software package with the Auxiliary Density Matrix Method.[22] The HSE06 functional[34] was employed for all systems except for hematite and the hematite/water interface for which a modified HSE06 functional was used where the fraction of HFX exchange is reduced to 12%. This functional, denoted HSE06(12%) in the following, was shown to improve the description[35] of the electronic structure compared to the original HSE06 functional. Core-electrons were replaced by Goedecker-Teter-Hutter (GTH) pseudopotentials.[36] For Fe and Co a semi-core GTH was used treating 3s, 3p, 3d and 4s electrons explicitly, and for O and H the valence electrons were treated explicitly. Becke-Johnson damped[37] D3 dispersion correction[38] was applied for each system.

Determining the improvement in computational efficiency for this method requires comparing the runtime of a regular MD trajectory to the one with the screening scheme devised being active. Since the startup of the software package can take significant time as well (see Figure 1), we performed three runs in total for each system. Two of which do not make use of the new screening method, but have a different number of MD steps. In the third run screening is active and this run can have an arbitrary number of steps. From the walltime of these three runs, we determined the startup time, the mean duration of a regular/update MD step and the mean duration of the screened force step.

All calculations were performed with the development version of cp2k on an Microsoft Azure G5 instance with 32 cores and 448 GB memory.

# 3 Results

The performance of the HFX force screening method is summarised in Table 1. For each system the force threshold $f_T$ was chosen to be $10^{-5}$ H/bohr, a full force calculation was done every five MD times steps (equal to one update every 2.5 fs) and a history of five force vectors was chosen for the force error correction scheme. These settings resulted in a maximum error in any force components of $\leq 5 \times 10^{-4}$ Hartree bohr$^{-1}$. We find that a significant fraction of the ERIs can be excluded leading to efficiency gains for a range of model systems, in particular for extended systems with a high electron density. For instance, for a $2 \times 2 \times 2$ super cell of CoO only 8% of the ERIs need to be included resulting in a speed-up of about a factor of three. Small systems with low average electron density like a water box benefit only little by this approach. This is expected because the contribution of the HFX calculation to the overall computational cost is comparably small (see Figure 1). Yet, for systems containing TMOs, where inclusion of HFX is often required but computationally expensive, the data in Table 1 suggest a major improvement in MD performance.

The accuracy of the forces after screening is shown in Figure 4. The deviation with the exact HFX forces is within the target error of $5 \times 10^{-4}$ H/bohr over a range of five orders of magnitude. This means that while the calculations become up to three times faster, the resulting error in forces is very small and comparable to typical residual errors in geometry optimizations.

The force error correction scheme proved vital for the success of our screening method. Figure 5 shows the mean absolute force error for different force thresholds $f_T$ with and without the correction scheme. For a more aggressive (larger) threshold value, the mean force error increases in both cases as expected. Importantly, when the correction scheme is applied the errors are reduced by one order of magnitude and the sensitivity of the error on the threshold value becomes weaker (smaller slope in Figure 5). This allows for a higher force thresholds to be used leading to exclusion of more ERIs in the steps between updates, and an improved computational efficiency.

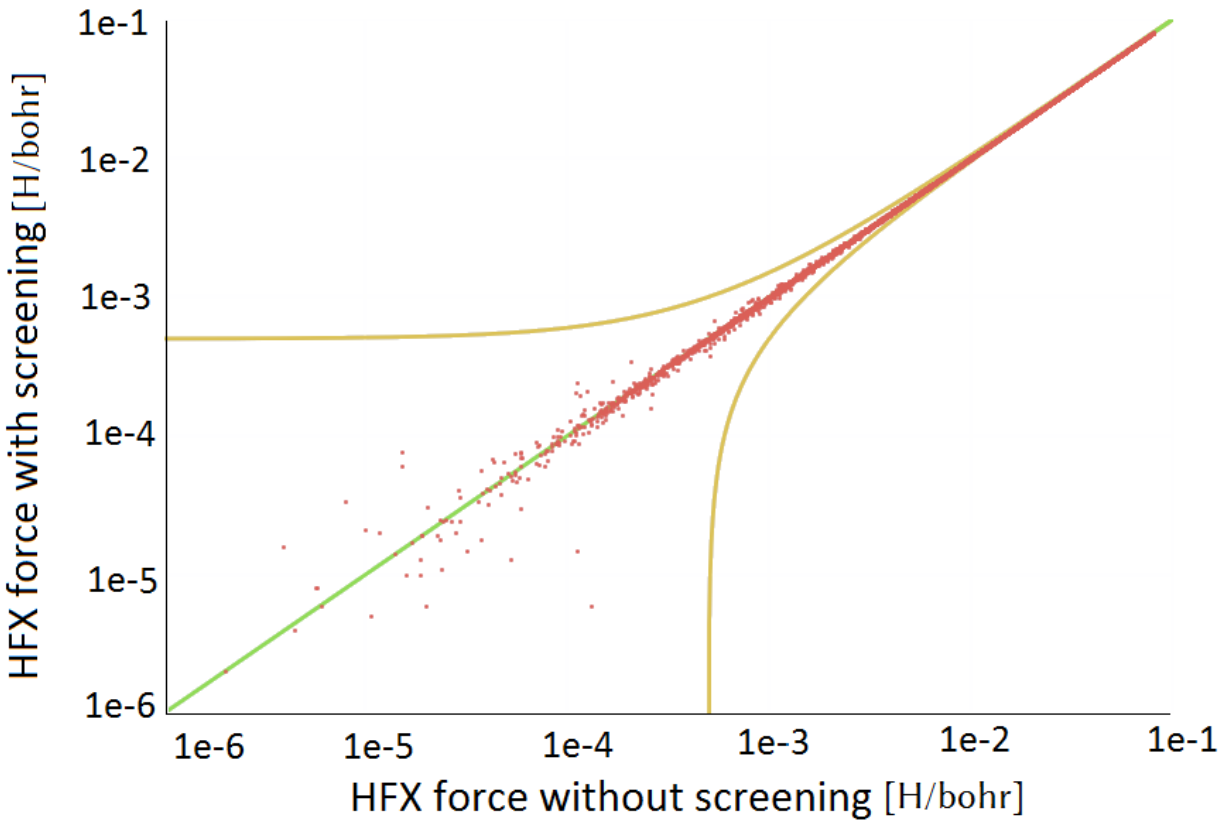


Figure 4: Magnitude of HFX force components after integral screening vs (exact) HFX force without screening (red dots). All points fall within the envelope indicating the maximum permissible force error ($\leq 5 \cdot 10^{-4}$ H/bohr, line in orange). Ideal agreement is indicated by a green line. The data were obtained for 1000 MD steps of a box of 32 water molecules.

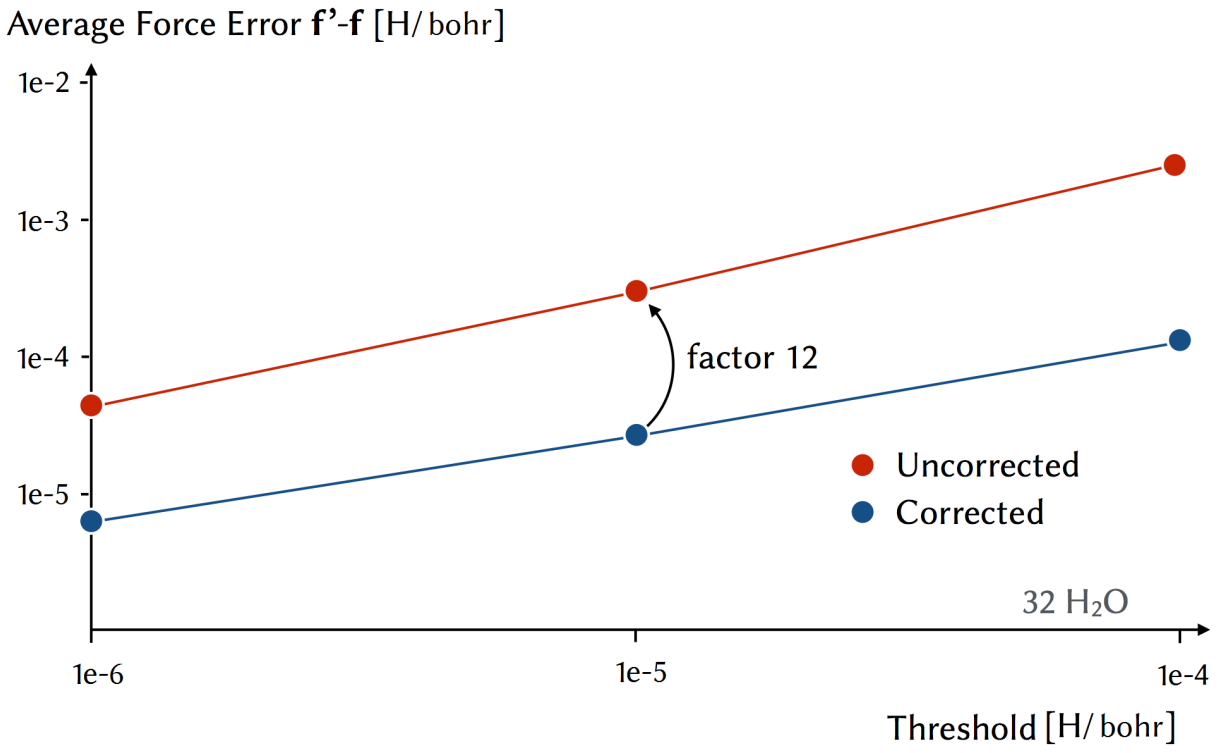


Figure 5: Mean absolute force error with and without the correction scheme. Forces taken from 1 ps of trajectory of a 32 water molecule box. Distributions of the absolute force errors with the correction scheme are shown in the SI, Figure 2.

Table 1: Performance of the screening method for various model systems. The number of electron repulsion integrals (ERI) in full HFX force calculations after Schwarz screening (and employing ADMM) is given as well as the percentage of ERI included in calculations with screening. A speed-up of 2 would mean the calculation with integral screening is twice as fast than the same calculation without screening. The remaining error of any force component is $\leq 5 \times 10^{-4}$ after screening and force error correction. All calculations are done with $n, k = 5, f_T = 10^{-5}$ H/bohr. Performance measurements consist of a MD long enough to complete an filter update cycle, i.e. 50 steps.

| System | Number of ERI | Included ERI | Speedup |
|---|---|---|---|
| CoO bulk $1 \times 1 \times 1$ (92 electrons) | $4.2 \cdot 10^8$ | 55% | 1.4 |
| CoO bulk $2 \times 1 \times 1$ (184 electrons) | $3.1 \cdot 10^9$ | 30% | 2.0 |
| CoO bulk $2 \times 2 \times 1$ (368 electrons) | $1.5 \cdot 10^{10}$ | 16% | 2.7 |
| CoO bulk $2 \times 2 \times 2$ (736 electrons) | $5.2 \cdot 10^{10}$ | 10% | 3.0 |
| CoO bulk $4 \times 2 \times 2$ (1.472 electrons) | $1.2 \cdot 10^{11}$ | 8% | 2.9 |
| Water dimer (16 electrons) | $8.9 \cdot 10^6$ | 89% | 1.1 |
| 32 water box (256 electrons) | $1.5 \cdot 10^9$ | 32% | 1.3 |
| OH radical in 31 water box (255 electrons) | $2.1 \cdot 10^9$ | 50% | 1.3 |
| Hematite bulk $1 \times 1 \times 1$ (300 electrons) | $1.6 \cdot 10^{10}$ | 54% | 2.4 |
| Hematite bulk $2 \times 1 \times 1$ (600 electrons) | $4.7 \cdot 10^{10}$ | 28% | – |
| Hematite bulk $2 \times 2 \times 1$ (1.200 electrons) | $1.1 \cdot 10^{11}$ | 21% | – |

Among the three parameters that define our screening method, the threshold $f_T$ affects the speed-up most. While the optimal choice is likely to be system dependent, here we used the same value for all systems for better comparison. The choice of the update frequency parameter $n$ depends on the force correction scheme used: the better the error prediction, the fewer updates are needed. Ultimately, the minimal update frequency is limited by the force-force correlation time. Yet, fewer updates than every $m^2$ MD steps do not lead to further efficiency gains according to Eq. 2. For the systems studied, we found that one update every 5 MD steps gives a good compromise between computational improvements and accuracy. This results in speed-ups $S$ between 1.1× and 3×, which is close to the maximum for given $f_T$, $S{=}m$. The precise number of history steps to be included in the force error prediction is not very important with regard to both computational cost and accuracy. Hence, the error correction is close to optimal in our calculations and further speed-ups can only be gained by further increasing $f_T$ at the cost of a larger error in the nuclear forces.

We have also investigated the effect of force screening on the energy conservation during

a Born-Oppenheimer MD run. Employing a history of previous densities as initial guess to accelerate SCF convergence[39] generally breaks the time-reversibility[40] which leads to a small energy drift. The total energy in NVE is a good indicator of the energy drift. Figure 6 shows this quantity in a water dimer MD. We observe a small energy drift of $4.7 \cdot 10^{-7}$ H/atom/ps as compared to $3.2 \cdot 10^{-7}$ H/atom/ps without our method.

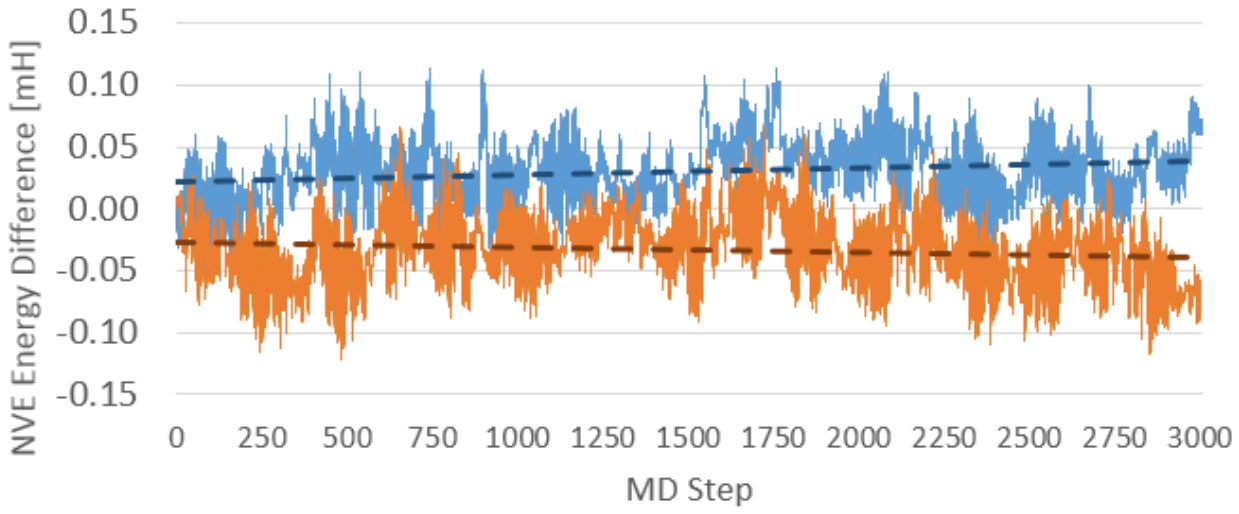


Figure 6: Energy conservation of an MD run for a water dimer in the NVE ensemble (HFX forces threshold $10^{-5}$ H/bohr, time step 0.5 fs) with (top) and without (bottom) HFX integral screening. Linear regression shows a total energy drift of $4.7 \cdot 10^{-7}$ H/atom/ps with screening compared to $3.2 \cdot 10^{-7}$ H/atom/ps without screening.

Finally, we evaluate the affect of HFX force screening on the solvation structure of the OH radical. We have chosen this system because GGA functionals such as BLYP fail to describe the correct solvation structure of this species, predicting a OH-water hemibond not seen in experiment.[7] Hybrid functionals such as HSE06 do not exhibit this artifact. Using the latter functional, we performed a molecular dynamics run with and without the force screening method. In both cases, the simulation started from a solvation structure pre-equilibrated with the BLYP functional and exhibiting the hemi-bond. After 2 ps, the peak in the radial distribution function corresponding to the hemibond vanishes – see Figure 7. This means that also for this system where the HFX forces are of crucial importance, the dynamic behaviour is not altered by the screening.

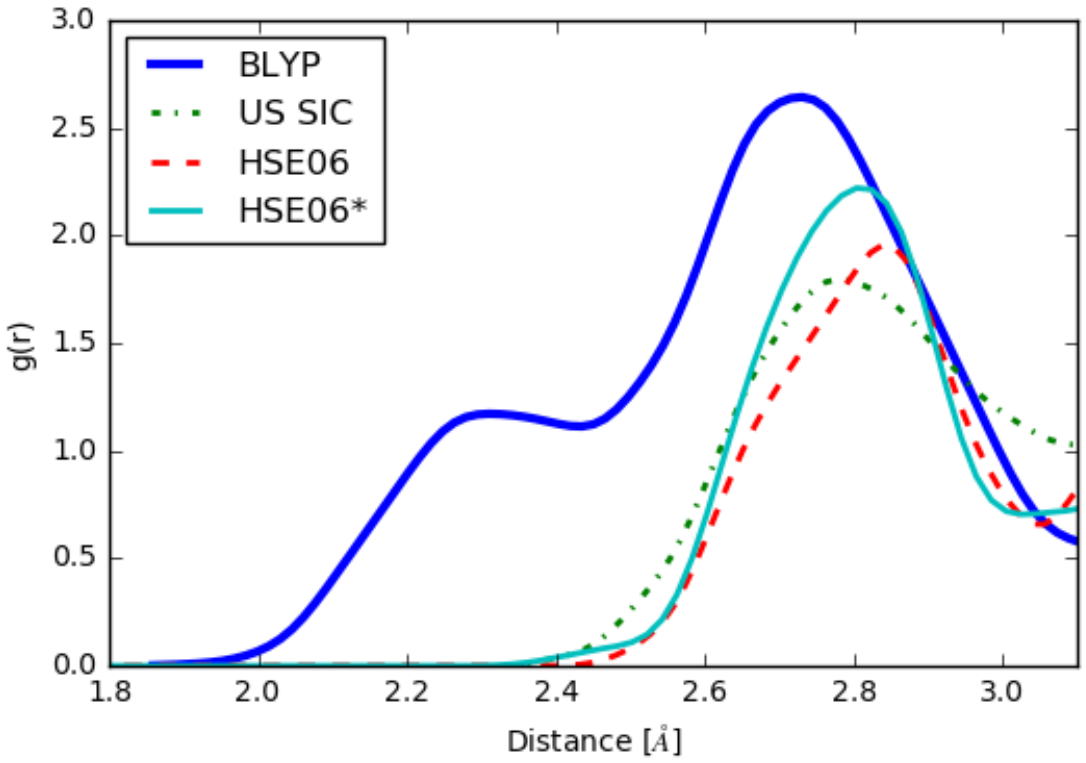


Figure 7: Radial distribution functions between the oxygen atoms of the OH radical and solvent water molecules. The peak at 2.25 Å corresponding to the hemibond (BLYP data, blue line) is no longer present in HSE06 after 2 ps of equilibration with screening (turquoise line) and without screening (red). Data for BLYP and US SIC (green) from Ref.[7] Deviations between HSE06 with and without screening are more likely related to statistical errors than integral screening.

# 4 Conclusion

In this work we have devised a simple method for further performance improvement of hybrid functional-based molecular dynamics simulation. In addition to Schwarz screening, each integral is screened against its contribution to the nuclear forces and is neglected if the contribution falls below a given threshold. We have shown that in systems where force evaluation is the major bottleneck, a large fraction of integrals (90%) can be excluded resulting in significant speed-ups (up to a factor of 3 in systems investigated). Due to the nearsightedness of exchange interactions, we expect the efficiency gain to further increase with increasing system size. While our current implementation in CP2K has a memory scaling of $\mathcal{O}(N^4)$, this can be overcome by using a sparse matrix for the filter list instead. Work in this direction is planned in our group.

An essential ingredient of our method is the correction scheme we have devised for correction of the error due to neglect of integrals. The error is reduced by an order of magnitude at no computational overhead resulting in essentially negligible residual force errors. Since comparing forces against reference forces grasps the static picture only, we also checked the effect of force screening on thermal fluctuations for properties like energy conservation and the solvation structure of OH radical in water. We found that the force screening method has little to no effect on the results. For the energy conservation, we observe a larger variance but no systematic energy drift.

The screening method presented offers a systematic speed-up for hybrid functional MD in systems which are currently inaccessible to this level of theory due to the high computational costs associated with it. While all systems benefit to a varying extent, this method aims at systems where force evaluation is the computational bottleneck, like the transition metal containing materials investigated in this work. For molecular dynamics runs where energy rather than force evaluation is dominating the computational cost, applying the same integral exclusion list to the energy calculation could offer similar performance gains.

# 5 Acknowledgements

G. v. R. gratefully acknowledges a PhD studentship co-sponsored by University College London and Pacific Northwest National Laboratory (PNNL) through its BES Geosciences program supported by the U.S. Department of Energys Office of Science, Office of Basic Energy Sciences, Chemical Sciences, Geosciences and Biosciences Division. MD simulations without screening were carried out on ARCHER, the UK national HPC facility (Edinburgh), to which access was granted via the ARCHER Leadership pilot call and the Materials Chemistry Consortium (EPSRC grant EP/L000202). Data analysis and calculations with screening were carried out with computing resources provided through a Microsoft Azure Sponsorship as well as AWS Cloud Credits for Research and supported by software made available by Tableau Inc.

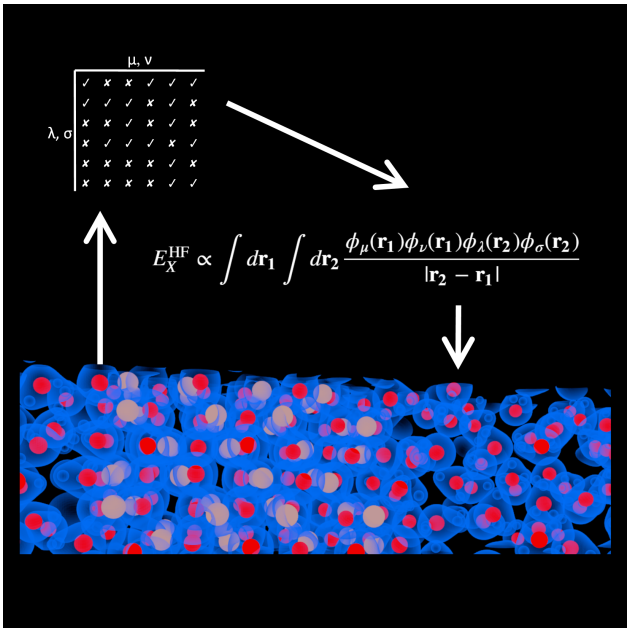


Figure 8: TOC graphic

# References

(1) Goerigk, L.; Grimme, S. *J Chem Theory Comput* **2010**, *6*, 107–126.

(2) Jacquemin, D.; Wathelet, V.; Perpete, E. A.; Adamo, C. *J Chem Theory Comput* **2009**, *5*, 2420–2435.

(3) Janesko, B. G.; Henderson, T. M.; Scuseria, G. E. *Phys Chem Chem Phys* **2009**, *11*, 443–454.

(4) Blumberger, J. *Phys. Chem. Chem. Phys* **2008**, *10*, 5651.

(5) Cohen, A. J.; Mori-Sanchez, P.; Yang, W. *Science* **2008**, *321*, 792–794.

(6) Perdew, J. P.; Zunger, A. *Phys Rev B* **1981**, *23*, 5048–5079.

(7) VandeVondele, J.; Sprik, M. *Phys Chem Chem Phys* **2005**, *7*, 1363.

(8) Dudarev, S. L.; Savrasov, S. Y.; Humphreys, C. J.; Sutton, A. P. *Phys Rev B* **1998**, *57*, 1505–1509.

(9) Wu, Q.; Van Voorhis, T. *J Chem Phys* **2006**, *125*, 164105.

(10) Dabo, I.; Ferretti, A.; Poilvert, N.; Li, Y.; Marzari, N.; Cococcioni, M. *Phys Rev B* **2010**, *82*, 115121.

(11) Kronik, L.; Stein, T.; Refaely-Abramson, S.; Baer, R. *J Chem Theory Comput* **2012**, *8*, 1515–1531.

(12) Tsuneda, T.; Hirao, K. *J Chem Phys* **2014**, *140*, 18A513.

(13) Hser, M.; Ahlrichs, R. *J Comput Chem* **1989**, *10*, 104–111.

(14) Strout, D. L.; Scuseria, G. E. *J Chem Phys* **1995**, *102*, 8448.

(15) Kohn, W. *Int J Quantum Chem* **1995**, *56*, 229–232.

(16) Yanai, T.; Tew, D. P.; Handy, N. C. *Chem Phys Lett* **2004**, *393*, 51–57.

(17) Heyd, J.; Scuseria, G. E.; Ernzerhof, M. *J Chem Phys* **2003**, *118*, 8207.

(18) Izmaylov, A. F.; Scuseria, G. E.; Frisch, M. J. *J Chem Phys* **2006**, *125*, 104103.

(19) Vydrov, O. A.; Heyd, J.; Krukau, A. V.; Scuseria, G. E. *J Chem Phys* **2006**, *125*, 074106.

(20) Kronik, L.; Stein, T.; Refaely-Abramson, S.; Baer, R. *J Chem Theory Comput* **2012**, *8*, 1515–1531.

(21) Heyd, J.; Scuseria, G. E. *J Chem Phys* **2004**, *121*, 1187.

(22) Guidon, M.; Hutter, J.; VandeVondele, J. *J Chem Theory Comput* **2010**, *6*, 2348–2364.

(23) Rivero, P.; Moreira, I. d. P. R.; Scuseria, G. E.; Illas, F. *Phys Rev B* **2009**, *79*, 245129.

(24) Heyd, J.; Peralta, J. E.; Scuseria, G. E.; Martin, R. L. *J Chem Phys* **2005**, *123*, 174101.

(25) Furche, F.; Perdew, J. P. *J Chem Phys* **2006**, *124*, 044103.

(26) Sasaki, S.; Fujino, K.; Takuchi, Y. *Proc. Jpn. Acad. B Phys. Biol. Sci.* **1979**, *55*, 43–48.

(27) Blake, R. L.; Hessevick, R. E.; Zoltai, T.; Finger, L. W. *Am Mineral* **1966**, *51*, 123–129.

(28) von Rudorff, G. F.; Jakobsen, R.; Rosso, K. M.; Blumberger, J. *J. Phys. Chem. Lett.* **2016**, 1155–1160.

(29) Liao, P.; Carter, E. A. *Phys Chem Chem Phys* **2011**, *13*, 15189.

(30) Zhang, X.; Bieberle-Htter, A. *ChemSusChem* **2016**, *9*, 1223–1242.

(31) Liao, L.; Zhang, Q.; Su, Z.; Zhao, Z.; Wang, Y.; Li, Y.; Lu, X.; Wei, D.; Feng, G.; Yu, Q.; Cai, X.; Zhao, J.; Ren, Z.; Fang, H.; Robles-Hernandez, F.; Baldelli, S.; Bao, J. *Nat Nanotechnol* **2013**, *9*, 69–73.

(32) von Rudorff, G. F.; Jakobsen, R.; Rosso, K. M.; Blumberger, J. *J Phys Condens Matter* **2016**, *28*, 394001.

(33) VandeVondele, J.; Krack, M.; Mohamed, F.; Parrinello, M.; Chassaing, T.; Hutter, J. *Comp Phys Comm* **2005**, *167*, 103–128.

(34) Krukau, A. V.; Vydrov, O. A.; Izmaylov, A. F.; Scuseria, G. E. *J Chem Phys* **2006**, *125*, 224106.

(35) Pozun, Z. D.; Henkelman, G. *J Chem Phys* **2011**, *134*, 224706.

(36) Goedecker, S.; Teter, M.; Hutter, J. *Phys Rev B* **1996**, *54*, 1703–1710.

(37) Johnson, E. R.; Becke, A. D. *J Chem Phys* **2006**, *124*, 174104.

(38) Grimme, S.; Antony, J.; Ehrlich, S.; Krieg, H. *J Chem Phys* **2010**, *132*, 154104.

(39) Kolafa, J. *J Comput Chem* **2004**, *25*, 335–342.

(40) Niklasson, A. M. N.; Tymczak, C. J.; Challacombe, M. *Phys Rev Lett* **2006**, *97*, 123001.