

Linear scalability of density functional theory calculations without imposing electron localization

Marcel D. Fabian, Ben Shpiro, and Roi Baer*

Fritz Haber Research Center for Molecular Dynamics and the Institute of Chemistry, The Hebrew University of Jerusalem, Israel

E-mail: roi.baer@huji.ac.il

Abstract

Linear scaling density functional theory approaches to electronic structure are often based on the tendency of electrons to localize even in large atomic and molecular systems. However, in many cases of actual interest, for example in semiconductor nanocrystals, system sizes can reach very large extension before significant electron localization sets in and the scaling of the numerical methods may deviate strongly from linear. Here, we address this class of systems, by developing a massively parallel density functional theory (DFT) approach which doesn't rely on electron localization and is formally quadratic scaling, yet enables highly efficient linear wall-time complexity in the weak scalability regime. The approach extends from the stochastic DFT method described in Fabian et. al. WIREs: Comp. Mol. Science, e1412 2019 but is fully deterministic. It uses standard quantum chemical atom-centered Gaussian basis sets for representing the electronic wave functions combined with Cartesian real space grids for some of the operators and for enabling a fast solver for the Poisson equation. Our main conclusion is, that when a processor-abundant high performance computing (HPC) infrastructure is available, this type of approach has the potential to allow the study of large systems in regimes where quantum confinement or electron delocalization prevents linear-scaling.

1 Introduction

In the past few decades, the supercomputers' massive number-crunching power, measured in floating-point operations per second (FLOPS), has grown a million-fold¹ and is currently pushing towards the exaflop (10^{18} FLOPS) realm. Combining this new technology with electronic structure calculations can revolutionize computational materials science and biochemistry, provided we complement it with algorithms that can efficiently exploit its massively parallel-based infrastructure.

One of the key questions then becomes how to quantify the efficiency of a certain algorithm on a massively parallel machine. A crucial measure in this regard is the *speedup*, which we define as the ratio

$$\mathbb{S}(W, M) \equiv \frac{T_1(W)}{T_M(W)} \quad (1.1)$$

between the wall-times, $T_1(W)$ for executing a given computational work W using a single processor and $T_M(W)$ for its execution using M processors working in parallel. In operational regimes where the speedup is nearly proportional to M , i.e. $\mathbb{S} = \mathbb{E} \times M$ there is a clear advantage in using a parallel multiprocessor approach where \mathbb{E} is the efficiency, with $\mathbb{E} = 1$ being ideal.

The efficient use of parallel computing was discussed by Amdahl in his seminal paper²,

where he identified in W an inherently serial (subscript s) and parallelizable (subscript p) part, $W = W_s + W_p$. He assumed that the execution wall-time is independent of M for completing W_s and decreases linearly with M for W_p . Amdahl defined the *serial fraction* as $s_A = \frac{T_1(W_s)}{T_1(W)}$, measured on a single processor machine for a given job independent of M . With this definition, the speedup can be expressed as: $\mathbb{S}_A(W, M) = (s_A + \frac{1-s_A}{M})^{-1}$ (*Amdahl’s law*, also called *strong scalability*) and saturates once M exceeds the value of $1/s_A$.

Gustafson pointed out^{3,4} that in real-world usage s_A isn’t independent of M , due to the fact, that one does not generally take a fixed-sized problem but rather scales the workload W with the available computing power. He then defined the serial fraction $s_G = \frac{T_M(W_s)}{T_M(W)}$ as measured on the M -processor system and showed that the speedup can be expressed as $\mathbb{S}(M) = s_G + M(1 - s_G)$ (*Gustafson’s law* also called *weak scalability*), enabling linear speed up which does not inherently saturate as M increases.

These considerations can be applied to electronic structure calculations of extended systems in DFT codes that lower the cubic scaling by taking advantage of electron localization⁵⁻²⁵. For linear-scaling schemes, the Amdahl serial fraction $s_A = \frac{T_1(W_s)}{T_1(W)}$ is expected to be system-size independent (since both timings in the numerator and the denominator scale linearly with system size) while for codes of higher algorithmic complexity, s_A decreases as system size increases²⁶. In a weak scalability analysis of the linear scaling codes Gustafson’s serial fraction $s_G = \frac{T_M(W_s)}{T_M(W)}$ is also expected to be system-size independent (since both timings in the numerator and the denominator scale linearly with system size) and therefore take to form: $s_G = (1 + \frac{M_0}{M})^{-1}$, where M_0 is a constant (depending on the hardware, algorithm). For large M , the speedup saturates to $\mathbb{S}_G(M) \rightarrow 1 + M_0$, but if M_0 is very large there is a sizable regime where $M \ll M_0$ and the s_G is essentially zero so *an ideal linear speedup* emerges, as reported, for example, for the CONQUEST code^{6,27}, even up to $M = 200,000$ cores on the Fujitsu-made

K-computer. It is clear from the previous studies mentioned above that it is important to determine the strong and weak scalability properties of codes that can use massively parallel machines, because they are sensitive to many details concerning hardware, systems size, algorithmic scaling etc.

In this paper we develop an efficiently parallelizable, (semi)local DFT approach which offers quadratic scaling with system size and does not involve approximations derived from assuming electron localization. It uses atom-centered Gaussian basis-sets combined with calculations performed on Cartesian grids for providing the electrostatic and exchange-correlation energies. We describe the theory and implementation in section 2, where we also provide an illustration of the non-localized nature of electrons in the large benchmarking systems we use (see Fig. 2.1). Next, we present the algorithmic complexity and the parallel strong/weak scalability properties of our approach in section 3, and finally, we summarize and discuss the conclusions in section 4.

2 Method

In our method, we work with standard quantum chemistry basis-sets, composed of atom-centered local functions $\phi_\alpha(\mathbf{r})$, $\alpha = 1, \dots K$. For calculating the necessary integrals, solving the Poisson equations, and generating the exchange-correlation potentials, we use a 3D Cartesian real-space grid of equidistant points spanning a simulation box, containing the system’s atoms and electronic density. For this purpose, we developed an efficient method for evaluating the basis functions on a relevant set of grid points, outlined in section A of the supplementary material. Our method of combining basis functions and real-space grids is similar in spirit to those existing in literature, such as SIESTA⁹ and CP2K/Quickstep⁷, but differs in important details. Unlike SIESTA, we use standard non-orthogonal Gaussian basis sets and unlike Quickstep we represent the basis functions on the grid where all integrals are performed as summations. The first type of inte-

gral that we have to evaluate on the grid then, is the overlap matrix:

$$S_{\alpha\beta} = h^3 \sum_{\mathbf{g}} \phi_{\alpha}(\mathbf{r}_{\mathbf{g}}) \phi_{\beta}(\mathbf{r}_{\mathbf{g}}), \quad (2.1)$$

where $\mathbf{r}_{\mathbf{g}}$ are the grid points and h is the grid-spacing. Next, the kinetic energy integrals are evaluated as

$$T_{\alpha\beta} = \frac{1}{2} h^3 \sum_{\mathbf{g}} \nabla \phi_{\alpha}(\mathbf{r}_{\mathbf{g}}) \cdot \nabla \phi_{\beta}(\mathbf{r}_{\mathbf{g}}), \quad (2.2)$$

where the derivatives of the basis functions are calculated analytically and then placed on the grid (see section A.2.3 of the supplementary material for details). To avoid an excessive number of grid points, the equally-spaced grid is complemented with norm-conserving pseudopotentials²⁸, representing the effects of the tightly bound core electrons (which are not treated explicitly) and taken into account in the KS Hamiltonian, represented by the Fock matrix

$$F^{KS} = T + V^{NL} + V^{KS}, \quad (2.3)$$

where

$$V_{\alpha\beta}^{NL} = h^3 \sum_{\mathbf{g}} \phi_{\alpha}(\mathbf{r}_{\mathbf{g}}) \sum_{C \in \text{nuclei}} \hat{v}_{nl}^C \phi_{\beta}(\mathbf{r}_{\mathbf{g}}) \quad (2.4)$$

are the integrals for the non-local pseudopotential and

$$V_{\alpha\beta}^{KS} = h^3 \sum_{\mathbf{g}} \phi_{\alpha}(\mathbf{r}_{\mathbf{g}}) v_{KS}(\mathbf{r}_{\mathbf{g}}) \phi_{\beta}(\mathbf{r}_{\mathbf{g}}) \quad (2.5)$$

are the KS potential integrals, where:

$$v_{KS}(\mathbf{r}_{\mathbf{g}}) = \sum_{C \in \text{nuclei}} v_{loc}^C(\mathbf{r}_{\mathbf{g}} - \mathbf{R}_C) + v_H[n](\mathbf{r}_{\mathbf{g}}) + v_{xc}[n](\mathbf{r}_{\mathbf{g}}). \quad (2.6)$$

In Eq. (2.6), $v_H[n](\mathbf{r}_{\mathbf{g}})$ is the Hartree potential on the grid which is evaluated directly from the grid representation of the electron density $n(\mathbf{r}_{\mathbf{g}})$ by a reciprocal space-based method for treating long range interactions²⁹. The exchange-correlation potential $v_{xc}[n](\mathbf{r}_{\mathbf{g}})$ (within the local density approximation (LDA)) is also determined on the grid directly from the

electron density. From the grid representation of the pseudopotentials¹ we obtain the potential $v_{loc}^C(\mathbf{r}_{\mathbf{g}} - \mathbf{R}_C)$ appearing in Eq. (2.6) for nucleus C at position \mathbf{R}_C and, by grid integration, the matrix V^{NL} appearing in Eq. (2.3). All integral calculations are performed in parallel for different basis function pairs; for more details see the supplementary material C.

The electron density on the grid is formally defined as

$$n(\mathbf{r}_{\mathbf{g}}) = 2 \sum_{\alpha, \beta}^K P_{\alpha\beta} \phi_{\alpha}(\mathbf{r}_{\mathbf{g}}) \phi_{\beta}(\mathbf{r}_{\mathbf{g}}), \quad (2.7)$$

where P is the density matrix (DM) and the factor of two comes from integration over spin degrees of freedom. The DM must obey an electron conserving criterion, namely that the integral over all grid points evaluates to the total number of electrons in the system: $h^3 \sum_{\mathbf{g}} n(\mathbf{r}_{\mathbf{g}}) = N_e$. Indeed, performing this integral and using Eq. (2.1) and (2.7) we find

$$N_e = 2 \text{Tr}[PS]. \quad (2.8)$$

This relation is part of a more general requirement, that the Kohn-Sham eigenstates are populated according to the Fermi-Dirac function $f_{FD}(\varepsilon) = \frac{1}{1+e^{\beta(\varepsilon-\mu)}}$ where ε is the corresponding energy eigenvalue. For the DM, this condition can be satisfied by defining²¹:

$$P = f_{FD}(S^{-1}F^{KS})S^{-1}. \quad (2.9)$$

For finite-temperature DFT, β is the inverse temperature and μ is the chemical potential. For ground-state calculations β obeys $\beta(\varepsilon_L - \varepsilon_H) \gg 1$, where ε_L (ε_H) is the Kohn-Sham eigenvalue of the lowest unoccupied (highest occupied) molecular orbital. The chemical potential in the Fermi-Dirac function is adjusted to reproduce the systems' number of electrons N_e through Eq. (2.8).

¹Here we use the Kleinman-Bylander (KB) form³⁰, which produces two types of operators, a non-local potential operator $\hat{v}_{nl}^C \equiv v_{nl}^C(\mathbf{r} - \mathbf{R}_C, \mathbf{r}' - \mathbf{R}_C)$ which is defined in a small sphere around each atomic core (\mathbf{R}_C is the location of atom C) and a scalar potential $v_{loc}^C(\mathbf{r} - \mathbf{R}_C)$ containing the long-range electron-shielded nucleus Coulomb attraction.

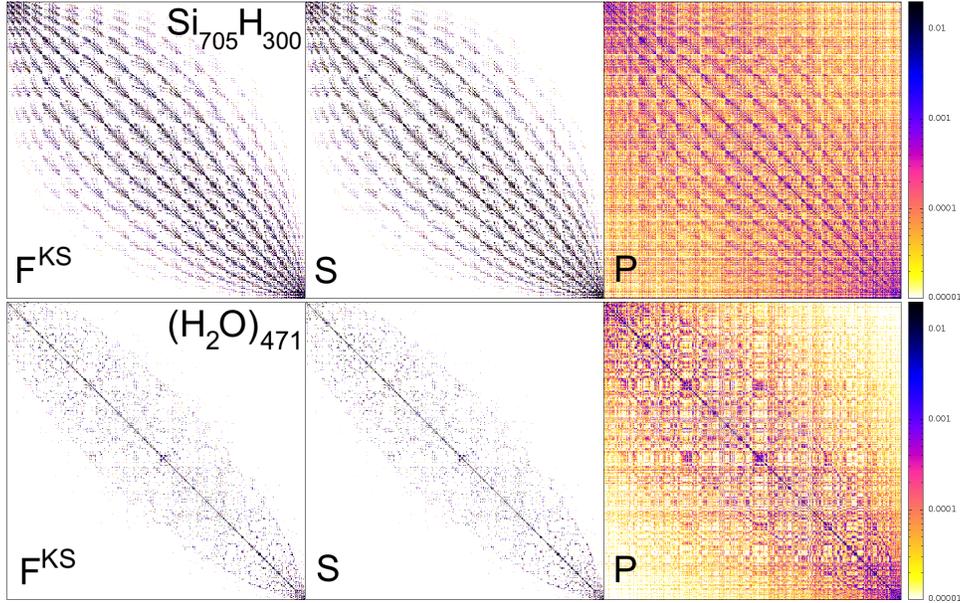


Figure 2.1: The Fock (F^{KS}), overlap (S) and DM (P) matrices of $\text{Si}_{705}\text{H}_{300}$ and $(\text{H}_2\text{O})_{471}$ cluster (calculated within the LDA) are shown using a color-coded plot. The basis set in both systems is similar in size, with $K \approx 6000$. For clearer inspection of the sparsity pattern the rows and columns are permuted so as to achieve a minimum bandwidth around the diagonal. We applied the `MinimumBandWidthOrdering` command of Mathematica^{®31} to the atomic proximity matrix $D_{AB} = \Theta(R_0 - R_{AB})$ (where R_{AB} is the distance between any pair of atoms A, B and $R_0 = 10a_0$ is the proximity distance), giving a permutation which is then used to order the atom-centered basis functions.

The use of atom-centered local basis functions allows for sparsity in the basic matrices F^{KS} and S , as illustrated in Fig. 2.1 for two systems of similar size but different chemical nature, a 2.5nm (diameter) semiconductor nanocrystal $\text{Si}_{705}\text{H}_{300}$ and a 3nm water cluster $(\text{H}_2\text{O})_{471}$. For the matrix representation in Fig. 2.1, we have ordered the atoms (and the basis functions associated with them) in a way that takes into account their spatial proximity (near atoms tend to have similar indices). Therefore, it is clear by mere inspection that F^{KS} and S have a relatively small spatial range and are therefore quiet sparse. Our approach makes an effort to exploit this property by using sparse matrix algebra. Despite the spatial locality of F^{KS} and S , P in these large systems is highly non-local, expressing the physical fact, that the electronic coherence in these systems is long ranged. For the silicon system, this fits our intuition, namely that silicon is by nature a semiconductor, with properties which are close to those of metals. Although water is a large

band-gap system, it is known that under LDA it exhibits very small HOMO-LUMO gaps³²⁻³⁵ (see also Fig. 2.2).

The various expectation values of relevant observables (i.e., operators in the grid representation) can be expressed as trace operations:

$$\langle \hat{O} \rangle = 2\text{Tr} [PO] \quad (2.10)$$

where

$$O_{\alpha\beta} = h^3 \sum_{\mathbf{g}} \phi_{\alpha}(\mathbf{r}_{\mathbf{g}}) \hat{O} \phi_{\beta}(\mathbf{r}_{\mathbf{g}}) \quad (2.11)$$

is the matrix representation of the one body operator \hat{O} in the atomic basis. In order to expedite the calculation we need to parallelize the computational work, and this can be done by representing the trace operations as a sum over unit column vectors u_{α} (with coordinates $(u_{\alpha})_{\beta} = \delta_{\alpha\beta}$, i.e., zeros in all positions except

at α), computed column by column:

$$\langle \hat{O} \rangle = 2 \sum_{\alpha}^K u_{\alpha}^T O P u_{\alpha}. \quad (2.12)$$

For achieving this, we treat the DM as an operator, i.e. we devise a linear-scaling method for applying it to the column vector u_{α} , based on Eq. (2.9): $P u_{\alpha} = f_{FD}(S^{-1} F^{KS}) S^{-1} u_{\alpha}$. The operation $S^{-1} u_{\alpha}$ is performed by the linear-scaling preconditioned conjugate-gradient approach involving repeated application of the sparse overlap matrix S on column vectors². The operation of $f_{FD}(S^{-1} F^{KS})$ on the column vector $S^{-1} u_{\alpha}$ employs a Chebyshev expansion of the function $f_{FD}(\varepsilon)$, which results in repeated applications of the operator $S^{-1} F^{KS}$ to column vectors. Details are described in the supplementary material B. The entire procedure can be readily distributed over several processors in parallel, each commissioned with a distinct set of u_{α} column vectors. This calculation method has the additional benefit that it avoids storage of the non-sparse DM. We discuss the algorithmic complexity of the approach, as well as its weak and strong scalability in section 3.

Equations (2.1)-(2.9) and the techniques of their application discussed above form a series of nonlinear equations that must be solved together, to give the self-consistent-field (SCF) solution. The procedure is iterative and uses the direct inversion of the iterative subspace (DIIS) convergence acceleration method³⁸. Once converged various expectation values such as charges and multipoles, density of states and polarizability can be calculated, as well as forces on the nuclei³⁹, which can be used for structure optimization.

In order to check and validate the implementation of the algorithm outlined above, we show in figure 2.2 the density of states (DOS) of a cluster of 100 water molecules, obtained with our program, and with the commercially avail-

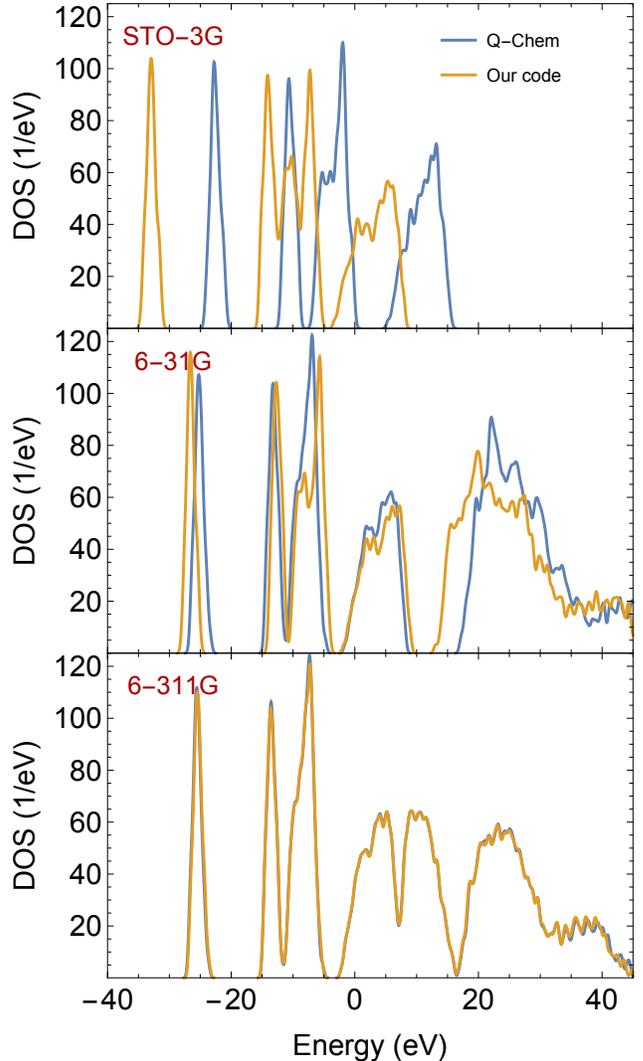


Figure 2.2: The DOS as a function of energy for a water cluster $(\text{H}_2\text{O})_{100}$ calculated using the all-electron Q-CHEM³⁷ and our code (with pseudo-potentials). Comparison is made for three standard Gaussian basis-sets as indicated in the panels within the LDA and a grid point spacing for our code of $\Delta x = 1/3 a_0$. Both calculations plot the DOS using $k_B T = 0.01 E_h$.

²We use the incomplete Cholesky preconditioning³⁶ for the conjugated gradient approach implemented in the HSL-MI28 and MI21 codes, respectively, where HSL is a collection of FORTRAN codes for large scale scientific computation (<http://www.hsl.rl.ac.uk/>).

able quantum chemistry program Q-Chem³⁷, using three different basis-sets, ranging from single to triple zeta quality (STO-3G, 6-31G and 6-311G). Overall, there is a clear trend visible that the two different DOSs become almost identical as the basis set quality grows. For the smallest STO-3G basis-set the results differ the most, due to a shift in energy between the two DOSs. The DOSs computed in the bigger 6-31G basis set are much closer and for the 6-311G basis set both our and Q-Chem’s DOSs are in close to perfect agreement. It is important to note, that whereas Q-Chem is an all-electron program, meaning that all electrons are described explicitly, our implementation treats core electrons with pseudopotentials. The discrepancies for the smaller basis-sets between the two codes can very likely be attributed to this fact. We stress that our results change very little for the different basis-sets in the gap region (-4 eV) which is usually of the highest interest in the energy spectrum.

3 Scaling properties of the method

In this section we study the method’s algorithmic complexity and analyze the speedup achievable by parallelization in terms of strong and weak scalability.

3.1 Algorithmic complexity

To understand the algorithmic complexity of our method, we have to examine how each part of our code scales as we increase the system size K . Here we are especially interested in the asymptotic behavior, meaning that the program part with the largest scaling will determine the overall algorithmic complexity. Our entire SCF cycle, that is described in detail in the supplementary material A.3, includes different integral calculations, solving the Poisson equation and calculating the density. The integral calculation is expected to scale linearly with system size K , i.e. $O(K)$, because the relevant matrices (F^{KS} , S) are expected to become sparse (see also Fig. (2.1)). The Poisson

equation is solved by a fast Fourier transform (FFT) which scales as $O(N_g \log N_g)$, where N_g are the grid points, expected to scale linearly with system size. This leaves only the density calculation which is done according to equation 2.12. The application of the DM P to a column vector u_α , expressed through a Chebyshev series, involves repeated applications of the operator $S^{-1}F^{KS}$ to the column vector $v = S^{-1}u_\alpha$ (see supplementary material B for details). The length of the Chebyshev expansion, N_C , is independent of the system size K and so the algorithmic complexity of the Pu_α operation is identical to that of one $S^{-1}F^{KS}v$ operation, namely linear with K . There are a total of K different Pu_α operations (see Eq. (2.12)), so that the overall algorithmic complexity of the method is asymptotically quadratic, i.e. $O(K^2)$. As the system size grows our algorithm could be modified to take advantage of the emerging sparsity of the DM, allowing for a K -independent complexity of each Pu_α operation. In such situations one can expect an overall linear-scaling numerical complexity, i.e. $O(K)$. However, in the present paper, we focus on the broad class of systems which are very large but for which the DM has not yet localized. Hence we are in the formally quadratic complexity regime.

To show that quadratic complexity is indeed what we achieve with this method, we plot, in figure 3.1 the wall-time per SCF cycle vs. system size for water clusters (taken from <http://www.ergoscf.org/xyz/h2o.php>) and hydrogen-terminated silicon nanocrystals (we use a series of nanocrystals, starting from $\text{Si}_{35}\text{H}_{36}$ reaching $\text{Si}_{2785}\text{H}_{780}$, for details, see supplementary material D), using STO-3G and the larger 6-31G basis-sets. As seen in the figure, all cases show overall quadratic algorithmic complexity. It is noteworthy to state that the small and intermediate sized systems in the figure exhibit a varying algorithmic complexity with system size associated with the interplay between linear complexity processes having a large prefactor and cubic stages due to the non-sparse nature of the Hamiltonian and overlap matrices.

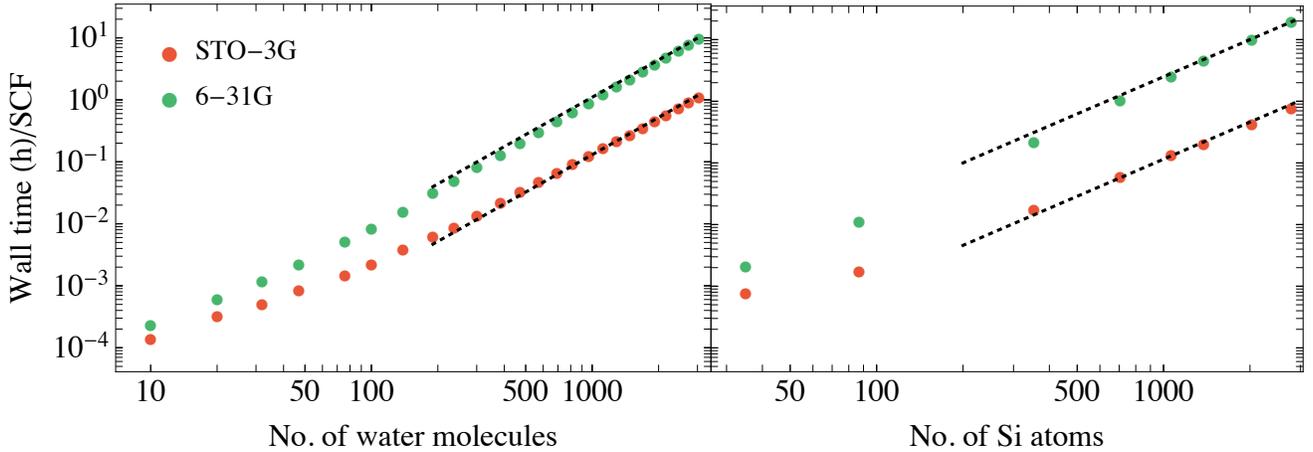


Figure 3.1: The wall-time as a function of system size, for the water clusters and the silicon nanocrystals, calculated using two basis sets within the LDA. The calculations were performed on eight Intel Xeon Gold 6132 CPU @ 2.60GHz 755GB RAM (connected through Infiniband), using 112 cores for all systems. The dotted lines in the figures are guides to the eye with designated quadratic scaling. Fitting the function $t = Ax^n$ to the data at the larger time range, one obtains the exponent $n = 1.9$ (2.1) for both the water and the silicon systems in the STO-3G (6-31G) basis.

3.2 Strong scalability

In figure 3.2 we study the strong scalability properties of our code, i.e. the scalability achievable when increasing the number of processors for a *given* task. We show in the figure the speedup and efficiency for a single SCF iteration of the $\text{Si}_{1379}\text{H}_{476}$ nanocrystal. Our definition for the speedup in Eq. 1.1 requires the knowledge of the elapsed wall time it takes a single processor (more accurately 1 core) to finish this nanocrystal calculation. Due to (human) time constraints we had to extrapolate this timing from a calculation on 36 cores on one single compute node by $T_1 = 36T_{36}$. The results can be analyzed in terms of the Amdahl law finding that the the serial fraction is $s_A = 9 \times 10^{-5}$ showing a high degree of parallelization. Accordingly, the parallelization efficiency drops very slowly as the number of processors increases, with 96% efficiency even at $M = 500$ (see the inset in the top panel). We emphasize that this is achieved with a 10Gb ethernet network communication. Potentially, the decay of efficiency may be slowed down by employing a faster communication solution. According to the Amdahl law, efficiency will drop to 0.5 when $M \approx \frac{1}{s_A} = 10^4$. In the supplementary material E we show results for

a smaller system, $\text{Si}_{705}\text{H}_{300}$ where the Amdahl serial fraction is larger, $s_A = 2 \times 10^{-4}$, a system-size dependency due to the quadratic complexity of our method (see our discussion in Section 1).

3.3 Weak scalability

In this section we focus on the weak scalability properties of our method, namely how the wall time changes with system size K when the number of processors afforded to the calculation M grows in fixed proportion $r = K/M$. In the left panel of Fig. 3.3 we present the wall-time T as a function of system size K for six series of runs we made with different fixed ratios ranging from $r = 4$ up to $r = 120$ (in the actual calculation, r is the number of vectors u_α assigned to each processor (see Eq. (2.12))). The markers of each series fall on asymptotically straight lines in the log-log plot which appear parallel to the dark-dashed line indicating a constant slope of 1. This confirms the claim of achieving linear-scaling wall-time in this regime of operation, where r is held constant. We would also like to examine the *speedup* in order to determine the degree of efficiency of our calculation on the parallel machine. For calculating the speedup under our definition in Eq. (1.1) we need to be

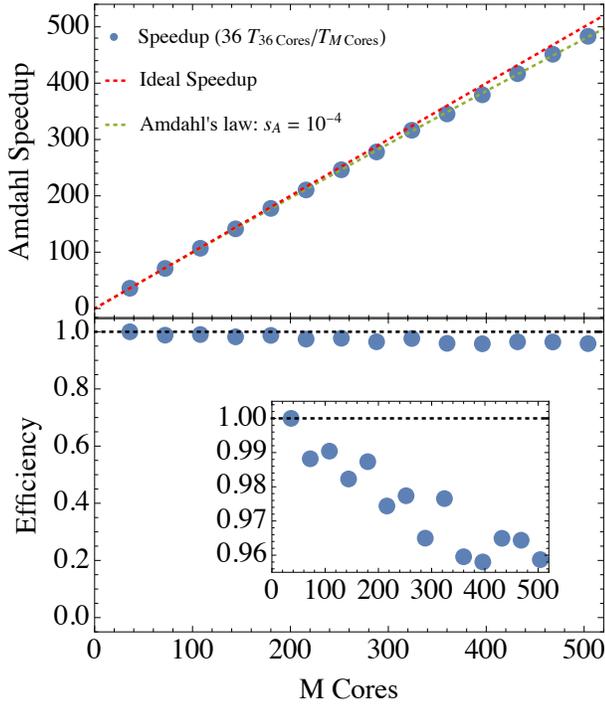


Figure 3.2: Strong scalability speedup analysis (upper panel) and efficiency $\mathbb{E}(M) = \mathbb{S}(M)/M$ (lower panel) for $\text{Si}_{1379}\text{H}_{476}$. The reference time for 1 processor for the speedup is extrapolated from $T_1 = 36T_{36}$. The inset in the lower panel enables a higher resolution of the efficiency regime close to unity. The calculation used the 6-31G basis-set (11984 basis functions) within the LDA and were performed on several 2.60GHz Intel Xeon Gold 6240 with 256 GB using 10Gb Ethernet networking communications.

able to estimate the wall time $T_1(W)$, which for the large systems is not easily accessible due to (human) time constraints. Therefore, we developed the following model for the wall time, with which we will estimate the $M = 1$ wall times:

$$T(K, M) = \frac{\tau_2 K^3 / (K_0 + K)}{M} + \tau_1 K \log M. \quad (3.1)$$

The first term on the right is the dominant parallelizable part of the calculation run on M processors (electron density calculation, see the supplementary material C for more information). For $K \gg K_0$ it exhibits quadratic scaling while for $K \ll K_0$ the scaling is cubic due to insufficient sparsity of the Hamiltonian and overlap matrices for small K . The second term in Eq. (3.1) reflects the timing of the serial part of the calculation, dominated by the communication time needed for specific MPI functions (reduce and broadcast) and scales linearly with K and logarithmically with M .

Using the analytical model, the speedup can now be obtained by plugging Eq. (3.1) into Eq. (1.1), resulting in the following closed form expression,

$$\mathbb{S}(r, M) = \frac{M}{\left(\frac{K_0}{rM} + 1\right) \frac{\tau_1}{\tau_2 r} \log M + 1}. \quad (3.2)$$

From this equation, it can be seen, that for asymptotically large values of M , the speedup approaches the limit $\mathbb{S}(r, M) \rightarrow \frac{M}{\frac{\tau_1}{\tau_2 r} \log M + 1}$ and as long as r is not too small,

$$r > \frac{\tau_1}{\tau_2} \log M, \quad (3.3)$$

the speedup is close to ideal $\mathbb{S}(r, M) \rightarrow M$.

We now fit our model to the calculation's timing results from the six constant- r series shown in the left panel of Fig. 3.3 (a total of 32 data points). This leads to a best-fit set of parameters (in hours): $\tau_1 \rightarrow 5.16 \times 10^{-6}$ h and $\tau_2 \rightarrow 5.63 \times 10^{-7}$ h and $K_0 \rightarrow 2292.6$ for our

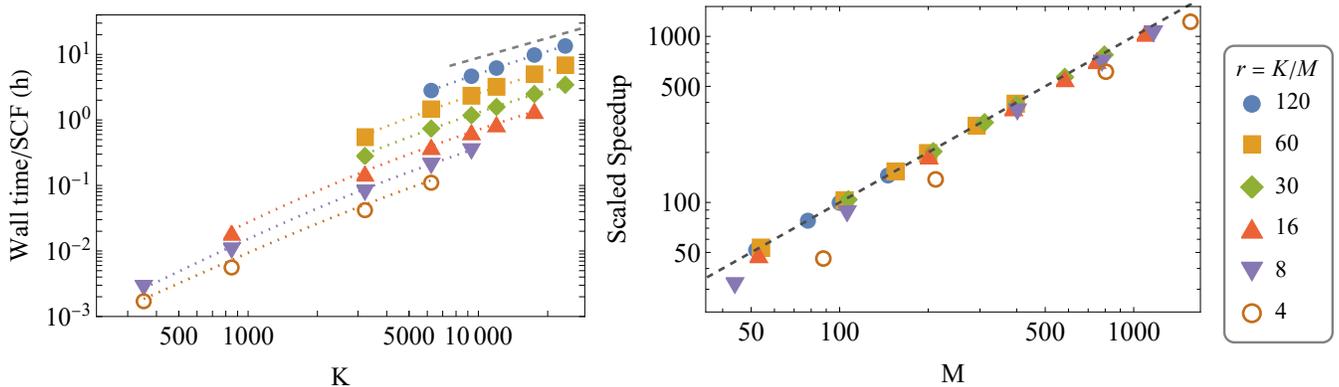


Figure 3.3: Weak scalability speedup analysis. Left: The wall-time for a single SCF cycle vs. the number of basis set functions K in calculations given for several fixed values of r (K/M where M is the number of processors) on a series of eight hydrogen-terminated silicon nanocrystals (detailed in the supplementary material D) using the 6-31G basis-set within the LDA. The black-dashed line is a guide to the eye showing linear scaling wall-time. The calculations were run on several Intel Xeon Gold 6240 CPUs @ 2.60GHz 256GB RAM connected through a 10Gb ethernet networking communications. We had access to at most 1584 cores, therefore for $r = 4$ we could not treat systems greater than $K = 6420$ and similar though less stringent limitations appeared for $r = 8$ and 16. The colored dotted lines are the best-fit of the data to our model in Eq. 3.4 Right: The scaled speedup as a function of the number of processors M , calculated for the six values of r from Eq. 3.2 using the best-fit parameters of our model. The black-dashed line indicates the “perfect” speedup $S = M$.

model and the resulting fit functions,

$$\frac{T_{\text{fit}}(K; r)}{10^{-6}\text{h}} = K \left(\frac{0.563r}{\frac{2292.6}{K} + 1} + 5.16 \log(K/r) \right) \quad (3.4)$$

are plotted in the left panel of the figure as dotted colored lines, one for each values of r . It can be seen that these fit functions indeed reproduce the actual data (given as points) quite closely.

Having the best-fit parameters, let us now discuss the actual estimated values for the (scaled) speedup in the Gustafson sense. These estimates, based on Eq. (3.2) are plotted in the right panel of Fig. 3.3. We see that for $r > 16$ the the speedup is not too far from ideal, in accordance with the analysis presented above, however as indicated in Eq. (3.3) the speedup is smaller when r decreases as is clearly visible for $r = 8$ and small M and for $r = 4$ for all values of M . However, even for these small r cases, the speedup is maintained as M increases and the calculation is still quite efficient.

4 Summary and Conclusions

In this paper we presented a parallelizable electronic structure approach to finite temperature density functional theory under (semi)local functionals, using atom-centered Gaussian basis sets which offers linear wall-time complexity as a function of system size in the weak scalability regime. The inherent time complexity of the method is quadratic $O(K^2)$, as discussed in section 3.1 and it does not involve truncation of density matrix elements, characteristics of linear-scaling approaches. Our trace-based calculation combined with Chebyshev expansions allows for efficient parallelization in the strong scalability sense, as shown in subsection 3.2. Due to the quadratic complexity, we found that the value of the Amdahl parameter was system-size dependent, with $s_A = 2 \times 10^{-4}$ for the $\text{Si}_{705}\text{H}_{300}$ system and $s_A = 9 \times 10^{-5}$ for $\text{Si}_{1379}\text{H}_{476}$. The overall weak scalability performance shows that linear scaling wall time is achievable, as demonstrated in section 3.3 and

is highly efficient when the number of orbitals per processors r is not smaller than ~ 10 and beyond that efficiency drops by a factor of ~ 1.5 . Our main conclusion is, that this type of approach has the potential to be a useful and efficient tool for studying large systems in regimes where quantum confinement or electron delocalization prevents traditional linear-scaling to set in. In the future, we aim to develop force evaluation for this approach, enabling the determination of the system's structure and other properties. Furthermore, for even larger systems, where electrons localize, we plan to enable linear scaling either through stochastic orbital methods²¹ or by exploiting directly the DM's finite range.

Acknowledgments

RB gratefully thanks the Israel Science Foundation grant 800/19.

References

- (1) Heldens, S.; Hijma, P.; Werkhoven, B. V.; Maassen, J.; Belloum, A. S. Z.; Van Nieuwpoort, R. V. The Landscape of Exascale Research: A Data-Driven Literature Analysis. *ACM Comput. Surv.* **2020**, *53*, 1–43.
- (2) Amdahl, G. M. Validity of the single processor approach to achieving large scale computing capabilities. Proceedings of the April 18-20, 1967, spring joint computer conference on - AFIPS '67 (Spring). Atlantic City, New Jersey, 1967; p 483.
- (3) Gustafson, J. L. Reevaluating Amdahl's law. *Commun. ACM* **1988**, *31*, 532–533.
- (4) Gustafson, J. L.; Montry, G. R.; Benner, R. E. Development of parallel methods for a 1024-processor hypercube. *SIAM journal on Scientific and Statistical Computing* **1988**, *9*, 609–638, Publisher: SIAM.
- (5) Ratcliff, L. E.; Dawson, W.; Fiscaro, G.; Caliste, D.; Mohr, S.; Videau, B.; Cristiglio, V.; Stella, M.; D'Alessandro, M.; Goedecker, S.; Nakajima, T.; Deutsch, T.; Genovese, L. Flexibilities of wavelets as a computational basis set for large-scale electronic structure calculations. *The Journal* **2020**, *29*.
- (6) Nakata, A.; Baker, J. S.; Mujahed, S. Y.; Poulton, J. T. L.; Arapan, S.; Lin, J.; Raza, Z.; Yadav, S.; Truflandier, L.; Miyazaki, T.; Bowler, D. R. Large scale and linear scaling DFT with the CONQUEST code. *J. Chem. Phys.* **2020**, *152*, 164112.
- (7) Kühne, T. D.; Iannuzzi, M.; Del Ben, M.; Rybkin, V. V.; Seewald, P.; Stein, F.; Laino, T.; Khaliullin, R. Z.; Schütt, O.; Schiffmann, F.; Golze, D.; Wilhelm, J.; Chulkov, S.; Bani-Hashemian, M. H.; Weber, V.; Borštnik, U.; Taillefumier, M.; Jakobovits, A. S.; Lazzaro, A.; Pabst, H.; Müller, T.; Schade, R.; Guidon, M.; Andermatt, S.; Holmberg, N.; Schenter, G. K.; Hehn, A.; Bussy, A.; Belleflamme, F.; Tabacchi, G.; Glöß, A.; Lass, M.; Bethune, I.; Mundy, C. J.; Plessl, C.; Watkins, M.; VandeVondele, J.; Krack, M.; Hutter, J. CP2K: An electronic structure and molecular dynamics software package - Quickstep: Efficient and accurate electronic structure calculations. *J. Chem. Phys.* **2020**, *152*, 194103.
- (8) Olsen, J. M. H.; Reine, S.; Vahtras, O.; Kjellgren, E.; Reinholdt, P.; Dundas, K. O. H.; Li, X.; Cukras, J.; Ringholm, M.; Hedegård, E. D.; Remigio, R. D.; List, N. H.; Faber, R.; Tenorio, B. N. C.; Pedersen, T. B.; Rinkevicius, Z.; Sauer, S. P. A.; Mikkelsen, K. V.; Kongsted, J.; Coriani, S.; Ruud, K.; Helgaker, T.; Jensen, H. J. A.; Norman, P. Dalton Project: A Python platform for molecular- and electronic-structure simulations of complex systems. *The Journal* **2020**, *19*.

- (9) García, A.; Papior, N.; Akhtar, A.; Artacho, E.; Blum, V.; Bosoni, E.; Brandimarte, P.; Brandbyge, M.; Cerdá, J. I.; Corsetti, F.; Cuadrado, R.; Dikan, V.; Ferrer, J.; Gale, J.; García-Fernández, P.; García-Suárez, V. M.; García, S.; Huhs, G.; Illera, S.; Korytár, R.; Koval, P.; Lebedeva, I.; Lin, L.; López-Tarifa, P.; Mayo, S. G.; Mohr, S.; Ordejón, P.; Postnikov, A.; Pouillon, Y.; Pruneda, M.; Robles, R.; Sánchez-Portal, D.; Soler, J. M.; Ullah, R.; Yu, V. W.-z.; Junquera, J. SIESTA: Recent developments and applications. *J. Chem. Phys.* **2020**, *152*, 204108.
- (10) Prentice, J. C. A.; Aarons, J.; Womack, J. C.; Allen, A. E. A.; Andrinopoulos, L.; Anton, L.; Bell, R. A.; Bhandari, A.; Bramley, G. A.; Charlton, R. J.; Clements, R. J.; Cole, D. J.; Constantinescu, G.; Corsetti, F.; Dubois, S. M.-M.; Duff, K. K. B.; Escartín, J. M.; Greco, A.; Hill, Q.; Lee, L. P.; Linscott, E.; O'Regan, D. D.; Phipps, M. J. S.; Ratcliff, L. E.; Serrano, Á. R.; Tait, E. W.; Teobaldi, G.; Vitale, V.; Yeung, N.; Zuehlsdorff, T. J.; Dziedzic, J.; Haynes, P. D.; Hine, N. D. M.; Mostofi, A. A.; Payne, M. C.; Skylaris, C.-K. The ONETEP linear-scaling density functional theory program. *J. Chem. Phys.* **2020**, *152*, 174111.
- (11) Rudberg, E.; Rubensson, E. H.; Sałek, P.; Kruchinina, A. Ergo: An open-source program for linear-scaling electronic structure calculations. *SoftwareX* **2018**, *7*, 107–111.
- (12) Goedecker, S. Linear scaling electronic structure methods. *Rev. Mod. Phys.* **1999**, *71*, 1085–1123.
- (13) Yang, W.; Lee, T.-S. A density-matrix divide-and-conquer approach for electronic structure calculations of large molecules. *The Journal of Chemical Physics* **1995**, *103*, 5674–5678.
- (14) Galli, G. Large-scale electronic structure calculations using linear scaling methods. *Physica Status Solidi B-Basic Research* **2000**, *217*, 231–249.
- (15) Scuseria, G. E. Linear scaling density functional calculations with Gaussian orbitals. *The Journal of Physical Chemistry A* **1999**, *103*, 4782–4790, Publisher: ACS Publications.
- (16) Baer, R.; Head-Gordon, M. Sparsity of the Density Matrix in Kohn-Sham Density Functional Theory and an Assessment of Linear System-Size Scaling Methods. *Phys. Rev. Lett.* **1997**, *79*, 3962–3965.
- (17) Baer, R.; Head-Gordon, M. Energy renormalization-group method for electronic structure of large systems. *Physical Review B-Condensed Matter* **1998**, *58*, 15296–15299.
- (18) Baer, R.; Neuhauser, D.; Rabani, E. Self-Averaging Stochastic Kohn-Sham Density-Functional Theory. *Phys. Rev. Lett.* **2013**, *111*, 106402.
- (19) Osei-Kuffuor, D.; Fattebert, J.-L. Accurate and Scalable $O(N)$ Algorithm for First-Principles Molecular-Dynamics Computations on Large Parallel Computers. *Phys. Rev. Lett.* **2014**, *112*, 046401.
- (20) Cytter, Y.; Rabani, E.; Neuhauser, D.; Baer, R. Stochastic Density Functional Theory at Finite Temperatures. *Phys. Rev. B* **2018**, *97*, 115207.
- (21) Fabian, M. D.; Shpiro, B.; Rabani, E.; Neuhauser, D.; Baer, R. Stochastic density functional theory. *Wiley Interdisciplinary Reviews: Computational Molecular Science* **2019**, *10.1002/wcms.1412*, e1412.
- (22) Li, W.; Chen, M.; Rabani, E.; Baer, R.; Neuhauser, D. Stochastic embedding DFT: Theory and application to p-nitroaniline in water. *J. Chem. Phys.* **2019**, *151*, 174115.
- (23) Chen, M.; Baer, R.; Neuhauser, D.; Rabani, E. Energy window stochastic density

- functional theory. *J. Chem. Phys.* **2019**, *151*, 114116.
- (24) Chen, M.; Baer, R.; Neuhauser, D.; Rabani, E. Overlapped embedded fragment stochastic density functional theory for covalently-bonded materials. *J. Chem. Phys.* **2019**, *150*, 034106.
- (25) Chen, M.; Baer, R.; Neuhauser, D.; Rabani, E. Stochastic density functional theory: Real- and energy-space fragmentation for noise reduction. *J. Chem. Phys.* **2021**, *154*, 204108.
- (26) Corsetti, F. Performance Analysis of Electronic Structure Codes on HPC Systems: A Case Study of SIESTA. *PLoS ONE* **2014**, *9*, e95390.
- (27) Arita, M.; Arapan, S.; Bowler, D. R.; Miyazaki, T. Large-scale DFT simulations with a linear-scaling DFT code CONQUEST on K-computer. *Journal of Advanced Simulation in Science and Engineering* **2014**, *1*, 87–97.
- (28) Troullier, N.; Martins, J. L. Efficient Pseudopotentials for Plane-Wave Calculations. *Phys. Rev. B* **1991**, *43*, 1993–2006.
- (29) Martyna, G. J.; Tuckerman, M. E. A reciprocal space based method for treating long range interactions in ab initio and force-field-based calculations in clusters. *J. Chem. Phys.* **1999**, *110*, 2810–2821.
- (30) Kleinman, L.; Bylander, D. M. Efficacious Form for Model Pseudopotentials. *Phys. Rev. Lett.* **1982**, *48*, 1425–1428.
- (31) Inc, W. R. Mathematica, Version 12.1.
- (32) Herbert, J. M.; Head-Gordon, M. Calculation of Electron Detachment Energies for Water Cluster Anions: An Appraisal of Electronic Structure Methods, with Application to $(\text{H}_2\text{O})_{20}^-$ and $(\text{H}_2\text{O})_{24}^-$. *J. Phys. Chem. A* **2005**, *109*, 5217–5229.
- (33) Rudberg, E. Difficulties in applying pure Kohn–Sham density functional theory electronic structure methods to protein molecules. *J. Phys.: Condens. Matter* **2012**, *24*, 072202.
- (34) Lever, G.; Cole, D. J.; Hine, N. D. M.; Haynes, P. D.; Payne, M. C. Electrostatic considerations affecting the calculated HOMO–LUMO gap in protein molecules. *J. Phys.: Condens. Matter* **2013**, *25*, 152101.
- (35) Sosa Vazquez, X. A.; Isborn, C. M. Size-dependent error of the density functional theory ionization potential in vacuum and solution. *The Journal of Chemical Physics* **2015**, *143*, 244105.
- (36) Scott, J.; Tuma, M. HSL_MI28: An Efficient and Robust Limited-Memory Incomplete Cholesky Factorization Code. *ACM Trans. Math. Softw.* **2014**, *40*, 1–19.
- (37) Shao, Y.; Gan, Z.; Epifanovsky, E.; Gilbert, A. T. B.; Wormit, M.; Kussmann, J.; Lange, A. W.; Behn, A.; Deng, J.; Feng, X.; Ghosh, D.; Goldey, M.; Horn, P. R.; Jacobson, L. D.; Kaliman, I.; Khaliullin, R. Z.; Kus, T.; Landau, A.; Liu, J.; Proynov, E. I.; Rhee, Y. M.; Richard, R. M.; Rohrdanz, M. A.; Steele, R. P.; Sundstrom, E. J.; III, H. L. W.; Zimmerman, P. M.; Zuev, D.; Albrecht, B.; Alguire, E.; Austin, B.; Beran, G. J. O.; Bernard, Y. A.; Berquist, E.; Brandhorst, K.; Bravaya, K. B.; Brown, S. T.; Casanova, D.; Chang, C.-M.; Chen, Y.; Chien, S. H.; Closser, K. D.; Crittenden, D. L.; Diedenhofen, M.; Jr, R. A. D.; Do, H.; Dutoi, A. D.; Edgar, R. G.; Fatehi, S.; Fusti-Molnar, L.; Ghysels, A.; Golubeva-Zadorozhnaya, A.; Gomes, J.; Hanson-Heine, M. W. D.; Harbach, P. H. P.; Hauser, A. W.; Hohenstein, E. G.; Holden, Z. C.; Jagau, T.-C.; Ji, H.; Kaduk, B.; Khistyayev, K.; Kim, J.; Kim, J.; King, R. A.; Klunzinger, P.; Kosenkov, D.; Kowalczyk, T.; Krauter, C. M.; Lao, K. U.; Laurent, A. D.; Lawler, K. V.; Levchenko, S. V.; Lin, C. Y.; Liu, F.;

Livshits, E.; Lochan, R. C.; Luenser, A.; Manohar, P.; Manzer, S. F.; Mao, S.-P.; Mardirossian, N.; Marenich, A. V.; Maurer, S. A.; Mayhall, N. J.; Neuscamman, E.; Oana, C. M.; Olivares-Amaya, R.; O'Neill, D. P.; Parkhill, J. A.; Perrine, T. M.; Peverati, R.; Prociuk, A.; Rehn, D. R.; Rosta, E.; Russ, N. J.; Sharada, S. M.; Sharma, S.; Small, D. W.; Sodt, A.; Stein, T.; Stock, D.; Su, Y.-C.; Thom, A. J. W.; Tsuchimochi, T.; Vanovschi, V.; Vogt, L.; Vydrov, O.; Wang, T.; Watson, M. A.; Wenzel, J.; White, A.; Williams, C. F.; Yang, J.; Yeganeh, S.; Yost, S. R.; You, Z.-Q.; Zhang, I. Y.; Zhang, X.; Zhao, Y.; Brooks, B. R.; Chan, G. K. L.; Chipman, D. M.; Cramer, C. J.; III, W. A. G.; Gordon, M. S.; Hehre, W. J.; Klamt, A.; III, H. F. S.; Schmidt, M. W.; Sherrill, C. D.; Truhlar, D. G.; Warshel, A.; Xu, X.; Aspuru-Guzik, A.; Baer, R.; Bell, A. T.; Besley, N. A.; Chai, J.-D.; Dreuw, A.; Dunietz, B. D.; Furlani, T. R.; Gwaltney, S. R.; Hsu, C.-P.; Jung, Y.; Kong, J.; Lambrecht, D. S.; Liang, W.; Ochsenfeld, C.; Rassolov, V. A.; Slipchenko, L. V.; Subotnik, J. E.; Voorhis, T. V.; Herbert, J. M.; Krylov, A. I.; Gill, P. M. W.; Head-Gordon, M. Advances in molecular quantum chemistry contained in the Q-Chem 4 program package. *Mol. Phys.* **2015**, *113*, 184–215.

- (38) Pulay, P. Improved SCF convergence acceleration. *J. Comput. Chem.* **1982**, *3*, 556–560.
- (39) Shpiro, B.; Fabian, M. D.; Rabani, E.; Baer, R. Forces from stochastic density functional theory under nonorthogonal atom-centered basis sets. 2021; _eprint: 2108.06770.

Supplementary material: “Linear scalability of density functional theory calculations without imposing electron localization”

A Basis-set / grid representation

A.1 Basis-set \rightarrow grid transformation

Each basis function $\phi_\alpha(\mathbf{r})$ is centered on an atom, located at point \mathbf{R}_α and associated with l_α , and m_α , the angular momentum and magnetic quantum number, respectively. The basis functions employed are non-orthogonal gaussians that are contracted to Gaussian type orbitals (GTOs):

$$\phi_\alpha(\mathbf{r}) = \sum_{p=1}^{N_{cn}} d_\alpha^p \Phi(\mathbf{r} - \mathbf{R}_\alpha; \zeta_\alpha^p, l_\alpha, m_\alpha), \quad (1)$$

where ζ_α^p are the Gaussian exponents, d_α^p are the contraction coefficients N_{cn} is the contraction length, and

$$\Phi(\mathbf{r}; \zeta, l, m) = \sum_{i,j,k} w_{lm}^{i,j,k} \varphi^{ijk}(\mathbf{r}; \zeta) \quad (2)$$

are the primitive spherical GTOs which are constructed from the following Cartesian GTOs

$$\varphi^{ijk}(\mathbf{r}; \zeta) = x^i e^{-\zeta x^2} \times y^j e^{-\zeta y^2} \times z^k e^{-\zeta z^2}, \quad (3)$$

where $\mathbf{r} = (x, y, z)$, $i + j + k = l$ and the weights w^{ijk} are given in Table 1.

In order to evaluate the GTO primitive at a grid-point $\mathbf{r}_g = (x_g, y_g, z_g)$, for each Cartesian

dimension, e.g. x , we store in memory an array containing, the values of $(x_g - R_{\alpha x})^i e^{-\zeta_{\alpha p} (x_g - R_{\alpha x})^2}$ (and its derivatives) in the interval $[-L_{\alpha x}, L_{\alpha x}]$ around $R_{\alpha x}$. The memory cost for the three arrays is small since L_{α} is small due to the compact nature of the Gaussian function. With these considerations, the evaluation of $\varphi^{ijk}(\mathbf{r}_g, \zeta)$ involves just two multiplications, at least an order of magnitude faster than a direct evaluation of the GTO primitive.

Table 1: Cartesian angular momenta and weight derived from the real solid harmonics¹ for $l \leq 2$. Note that only cases where $i + j + k = l$ give non-zero values for d and w .

l	s	p			d							
m	0	1	0	-1	2		1	0			-1	-2
i	0	1	0	0	2	0	1	2	0	0	0	1
j	0	0	0	1	0	2	1	0	2	0	1	0
k	0	0	1	0	0	0	0	0	0	2	1	1
w^{ijk}	1	1	1	1	$\sqrt{3}/2$	$-\sqrt{3}/2$	$\sqrt{3}$	$-1/2$	$-1/2$	1	$\sqrt{3}$	$\sqrt{3}$

A.1.1 The GTO window

The length L_{α} of the window encompassing the basis function $\phi_{\alpha}(\mathbf{r})$ is determined such that the most protruding primitive GTO has $1 - \eta$ of its “charge” included, where η is a small cutoff parameter. This can be expressed as the solution of the charge equation

$$c\left(L_{\alpha}\sqrt{\zeta_{\alpha}}, l_{\alpha}\right) = 1 - \eta, \quad (4)$$

where $c(d, l) \equiv \frac{\int_0^d x^l e^{-x^2} dx}{\int_0^{\infty} x^l e^{-x^2} dx}$ and $\zeta_{\alpha} = \min_p \zeta_{\alpha}^p$ identify the exponent of the most protruding primitive GTO. An important issue is the choice of η , small values give larger windows, improving accuracy at the expense of a higher computational cost. In Fig. 1 the wall-time and accuracy is plotted vs the η parameter and the window length L for a deterministic calculation of $\text{Si}_{353}\text{H}_{196}$. As η drops from 10^{-3} to 10^{-6} , L grows from 5 to 7Å (the volume and number of grid-points grow by a factor of 2.7), the single SCF iteration wall-time increases

by a factor of 3, while the relative energy error drops by nearly 3 orders of magnitude.

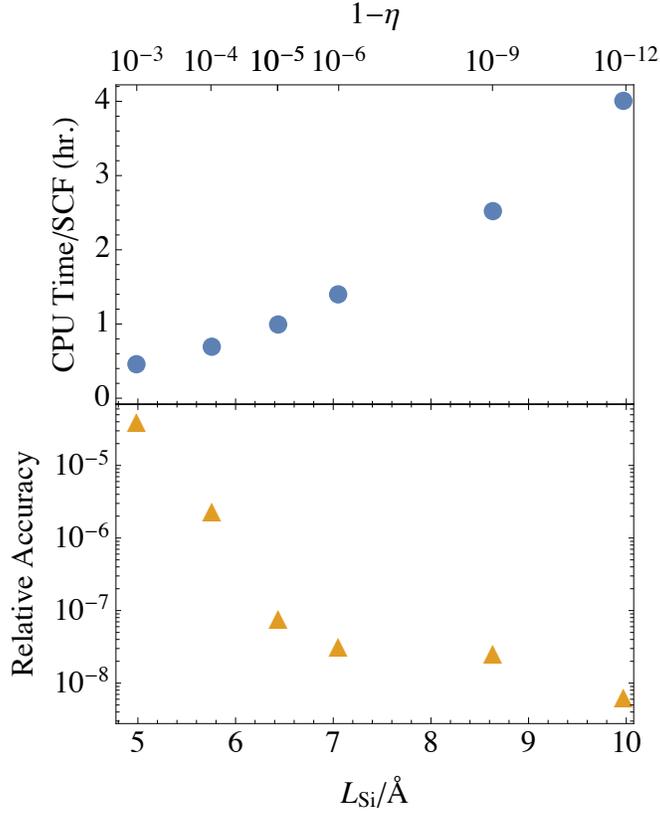


Figure 1: The wall-time of a deterministic LDA calculation (cubic-scaling) of $\text{Si}_{353}\text{H}_{196}$ and the accuracy in the relative energy vs. η and the silicon atom window length L_{Si} . The grid spacing was $\Delta h = 0.5a_0$ and the basis set was 6-31G.

Furthermore, not all grid points of the entire grid will be relevant to evaluate the numerical integrals for each pair of basis function. We therefore introduce for each contracted spherical GTO $\phi_\alpha(\mathbf{r}_g)$ a small grid-axes aligned cubic window W_α of lengths $\mathbf{L}_\alpha = (L_{\alpha x}, L_{\alpha y}, L_{\alpha z})$, in which the basis function is defined and nonzero. In figure 2 two such grid windows are displayed for the 2 dimensional case for illustrative purpose (our program uses of course 3D windows): two atomic centers are present (R_α, R_β), each with their own window (W_α, W_β). For each of the windows only a subset of points is relevant and even more so for the overlap of both windows, $W_{\alpha\beta}$, where only four grid points are relevant in this example.

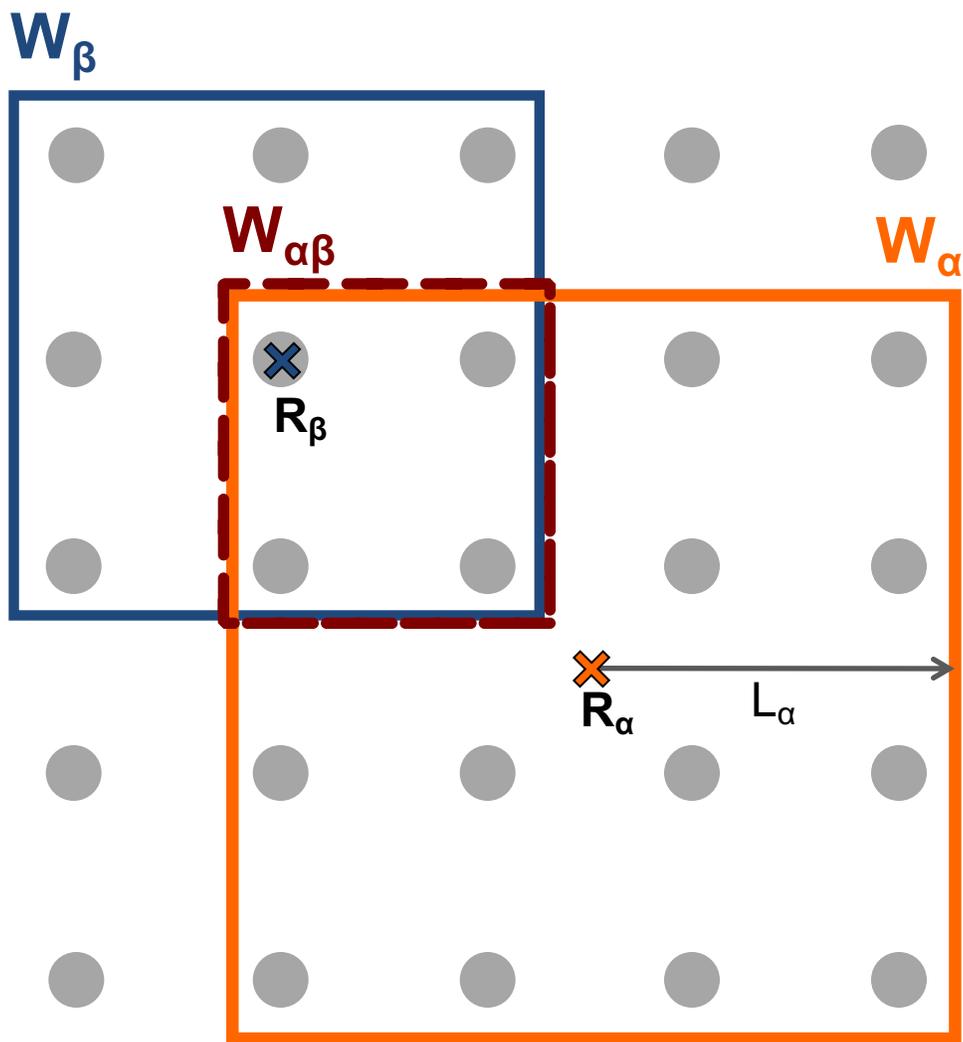


Figure 2: Schematic of grid windows in 2D. Shown are the grid points displayed as grey disks and two atomic centers, R_α and R_β , with their respective GTO windows, W_α and W_β of length L_α and L_β (not shown). The window $W_{\alpha\beta}$ is the region of overlap between W_α and W_β .

A.1.2 The density $n(\mathbf{r}_g)$

One important step in the DFT algorithm is to evaluate the electronic density on the grid, based on the density matrix. In order to calculate the electronic density, we define

$$n_\alpha(\mathbf{r}_g) = 2\phi_\alpha(\mathbf{r}_g) \sum_\beta \phi_\beta(\mathbf{r}_g) P_{\beta\alpha} \quad (5)$$

as the density contributed by the α 's column of the density matrix P , which can be calculated independently on parallel processors. The calculation is done by the following schematic algorithm:

1. For all $\mathbf{r}_g \in W_\alpha$, set $n_\alpha(\mathbf{r}_g) = 0$.
2. For each β such that $W_{\alpha\beta} = W_\alpha \cap W_\beta$ (see Fig. 2) and $\mathbf{r}_g \in W_{\alpha\beta}$, we update the density
$$n_\alpha(\mathbf{r}_g) \leftarrow n_\alpha(\mathbf{r}_g) + \phi_\alpha(\mathbf{r}_g) \phi_\beta(\mathbf{r}_g) P_{\beta\alpha}.$$

The total density $n(\mathbf{r}) = \sum_\alpha n_\alpha(\mathbf{r})$ is then obtained as a reduce operation for all processors. For more details on the parallelization see section C.

A.2 Grid \rightarrow basis-set transformation

Here we will explain how to construct the AO matrices from the corresponding grid-operators, the KS potential $v_{KS}(\mathbf{r}) \rightarrow V_{\alpha\beta}^{KS}$, the non-local pseudopotential $\hat{v}_{nl}(\mathbf{r}, \mathbf{r}') \rightarrow V_{\alpha\beta}^{NL}$ and the kinetic energy $-\frac{\hbar^2}{2m_e} \nabla^2 \rightarrow T_{\alpha\beta}$.

A.2.1 KS potential V^{KS} and overlap S

In order to generate a matrix element $V_{\alpha\beta}^{KS}$ in the basis set representation from a grid-based function $v_{KS}(\mathbf{r})$, we only have to sum over the grid points, where both windows W_α and W_β overlap ($W_{\alpha\beta} = W_\alpha \cap W_\beta$, see also figure 2) :

$$V_{\alpha\beta}^{KS} = \sum_{\mathbf{r}_g \in W_{\alpha\beta}} \phi_\alpha(\mathbf{r}_g) v_{KS}(\mathbf{r}_g) \phi_\beta(\mathbf{r}_g) h^3, \quad (6)$$

with h as the grid spacing. A list is kept in advance for tracking the pairs of orbitals ϕ_α and ϕ_β that actually have a non empty $W_{\alpha\beta}$, expediting the evaluation. The overlap matrix can be calculated this way taking $v_s(\mathbf{r}_g) = 1$ at each gridpoint.

A.2.2 Non-local pseudopotential V^{NL}

While the grid enables a very efficient evaluation of the KS-potential, it is not ideal for describing core electrons. Therefore these electrons are taken into account through norm-conserving pseudopotentials. For atom C at the location \mathbf{R}_C , there is, similar to the window for a basis function, a window W^C , i.e. a set of grid points for which the nonlocal part of the pseudopotential, $v_{nl}^C(\mathbf{r}_g - \mathbf{R}_C, \mathbf{r}'_g - \mathbf{R}_C)$ operates. The matrix V^{NL} representing the non-local part of the pseudopotential in the AO basis has non zero elements $V_{\alpha\beta}^{NL}$ only for those orbital pairs for which there exists an atom C for which $W_\alpha^C \equiv W^C \cap W_\alpha \neq \emptyset$ and $W_\beta^C \neq \emptyset$ (where \emptyset is the empty set of grid points). In practice, the calculation of V^{NL} is done by locating, for each β all atoms C for which $W_\beta^C \neq \emptyset$ and then calculating the localized function:

$$\psi_{C\beta}(\mathbf{r}_g) = \sum_{\mathbf{r}'_g \in W_\beta^C} v_{nl}^C(\mathbf{r}_g - \mathbf{R}_C, \mathbf{r}'_g - \mathbf{R}_C) \phi_\beta(\mathbf{r}'_g) h^3, \quad \mathbf{r}_g \in W^C. \quad (7)$$

Then one searches for all orbitals ϕ_α for which $W_\alpha^C \neq \emptyset$ and calculates

$$V_{\alpha\beta}^C = \sum_{\mathbf{r}_g \in W_\alpha^C} \phi_\alpha(\mathbf{r}_g) \psi_{C\beta}(\mathbf{r}_g) h^3 \quad (8)$$

finally, the matrix element is obtained as a sum over all atoms C for which $W_\beta^C \neq \emptyset$:

$$V_{\alpha\beta}^{NL} = \sum_{C: W_\beta^C \neq \emptyset} V_{\alpha\beta}^C. \quad (9)$$

The pseudo-potential matrix V^{NL} has to be calculated only once in the beginning of the calculation.

A.2.3 Kinetic energy T

The final type of integral considered here is the kinetic energy integral.

$$T_{\alpha\beta} = \frac{1}{2} \sum_{\mathbf{r}_g \in W_{\alpha\beta}} [\nabla\phi_\alpha(\mathbf{r}_g)] \cdot [\nabla\phi_\beta(\mathbf{r}_g)] h^3. \quad (10)$$

Here we need to take the gradient of the basis function ϕ_α and ϕ_β and multiply them. Taking the gradient of both basis functions rather than the Laplace operator for only one ensures that T will also numerically be positive. For the gradient we need the partial derivatives of the basis functions calculated at each grid point \mathbf{r}_g . Currently derivatives up to second order are implemented. They are taken with respect to equation 3.

A.3 Putting everything together: The SCF cycle

In order to solve the non-linear KS equation, one has to do so iteratively, until the SCF solution is found. This process is depicted in figure 3, with a focus on our unique basis-set/grid representation. All integrals are calculated from their respective grid operators, as discussed in section A.2.1-A.2.3. This grid→basis-set transformation is depicted in figure 3 by purple arrows and the corresponding equation number. Except for V^{KS} , all other matrices are only calculated once in the beginning of the SCF cycle. For the start of the SCF cycle the KS-Fock matrix, F^{KS} , is calculated by summing $F^{KS} = T + V^{NL} + V^{KS}$, where the first V^{KS} matrix is determined by an initial guess of the density $n(\mathbf{r})$ (for us usually the sum of atomic densities). The next step in the SCF cycle is to update the density $n(\mathbf{r})$ on the grid. This step is the only basis-set→grid transformation and is shown in figure 3 by a red arrow. As discussed in section A.1.2, the density matrix P is put column-wise on the grid. Each column of P is calculated according to $P = f_{FD}(S^{-1}F^{KS})S^{-1}$, where the application of the inverse overlap matrix S^{-1} is handled by a preconditioned conjugate gradient (PCG) approach¹,

¹We used the incomplete Cholesky preconditioning² for the conjugated gradient approach implemented in the HSL-MI28 and MI21 codes, respectively, where HSL is a collection of FORTRAN codes for large scale scientific computation (<http://www.hsl.rl.ac.uk/>).

thereby only requiring the repeated application of S (which is advantageous, since S is much sparser than S^{-1}) and the matrix exponential ($\exp(S^{-1}F^{KS})$) is evaluated by a Chebyshev expansion (see section B). Once the density is calculated, the Hartree potential $v_H(\mathbf{r})$ is obtained by solving Poisson’s equation through a fast Fourier transform (FFT)³ and the exchange-correlation potential $v_{XC}(\mathbf{r})$ by employing the relevant functional form (in our case the local-density approximation (LDA)). The KS potential on the grid $v_{KS}(\mathbf{r})$ is then calculated according to $v_{KS}(\mathbf{r}_g) = \sum_{C \in \text{nuclei}} v_{loc}^C(\mathbf{r}_g - \mathbf{R}_C) + v_H[n](\mathbf{r}_g) + v_{xc}[n](\mathbf{r}_g)$, where v_{loc} is only calculated once in the beginning of the SCF cycle. The KS potential matrix V^{KS} is updated according to eq. 6 before the next SCF cycle starts with a new KS Fock matrix.

B Chebyshev expansion

The inverse temperature parameter β in the Fermi-Dirac function $f_{FD}(\varepsilon) = \frac{1}{1+e^{\beta(\varepsilon-\mu)}}$ needs to be chosen high enough so that $\beta(\varepsilon_L - \varepsilon_H) \gg 1$ where ε_L (ε_H) is the Kohn-Sham eigenvalue of the lowest unoccupied (highest occupied) molecular orbital. The chemical potential parameter μ must be adjusted to reproduce the number of electrons in the system. For the actual application of the Fermi-Dirac function on a vector u_α we need to evaluate a matrix exponential of a matrix product ($\exp(S^{-1}F^{KS})$) and this is achieved by using the Chebyshev polynomial expansion of the Fermi-Dirac function

$$f_{FD}(\varepsilon) = \sum_{l=0}^{N_C-1} a_l(T, \mu) T_l\left(\frac{H - \bar{E}}{\Delta E}\right) \quad (11)$$

where $T_l(x)$ are the Chebyshev polynomials obeying the recursion relation $T_{l+1}(x) = 2xT_l(x) - T_{l-1}(x)$, $N_C \approx 2\Delta E\beta$ is the expansion length, and finally: $\Delta E = \frac{E_{max} - E_{min}}{2}$, $\bar{E} = \frac{E_{max} + E_{min}}{2}$, where E_{max} (E_{min}) is the largest (smallest) eigenvalue of $H = S^{-1}F^{KS}$. The Chebyshev expansion coefficients $a_l(T, \mu)$ can be computed using fast Fourier transform methods⁴. The expansion, when applied to the Fermi-Dirac function of the Hamiltonian H allows us to approximate the α ’s column of the DM $\rho_\alpha \equiv P u_\alpha$ by the following procedure: First, we set

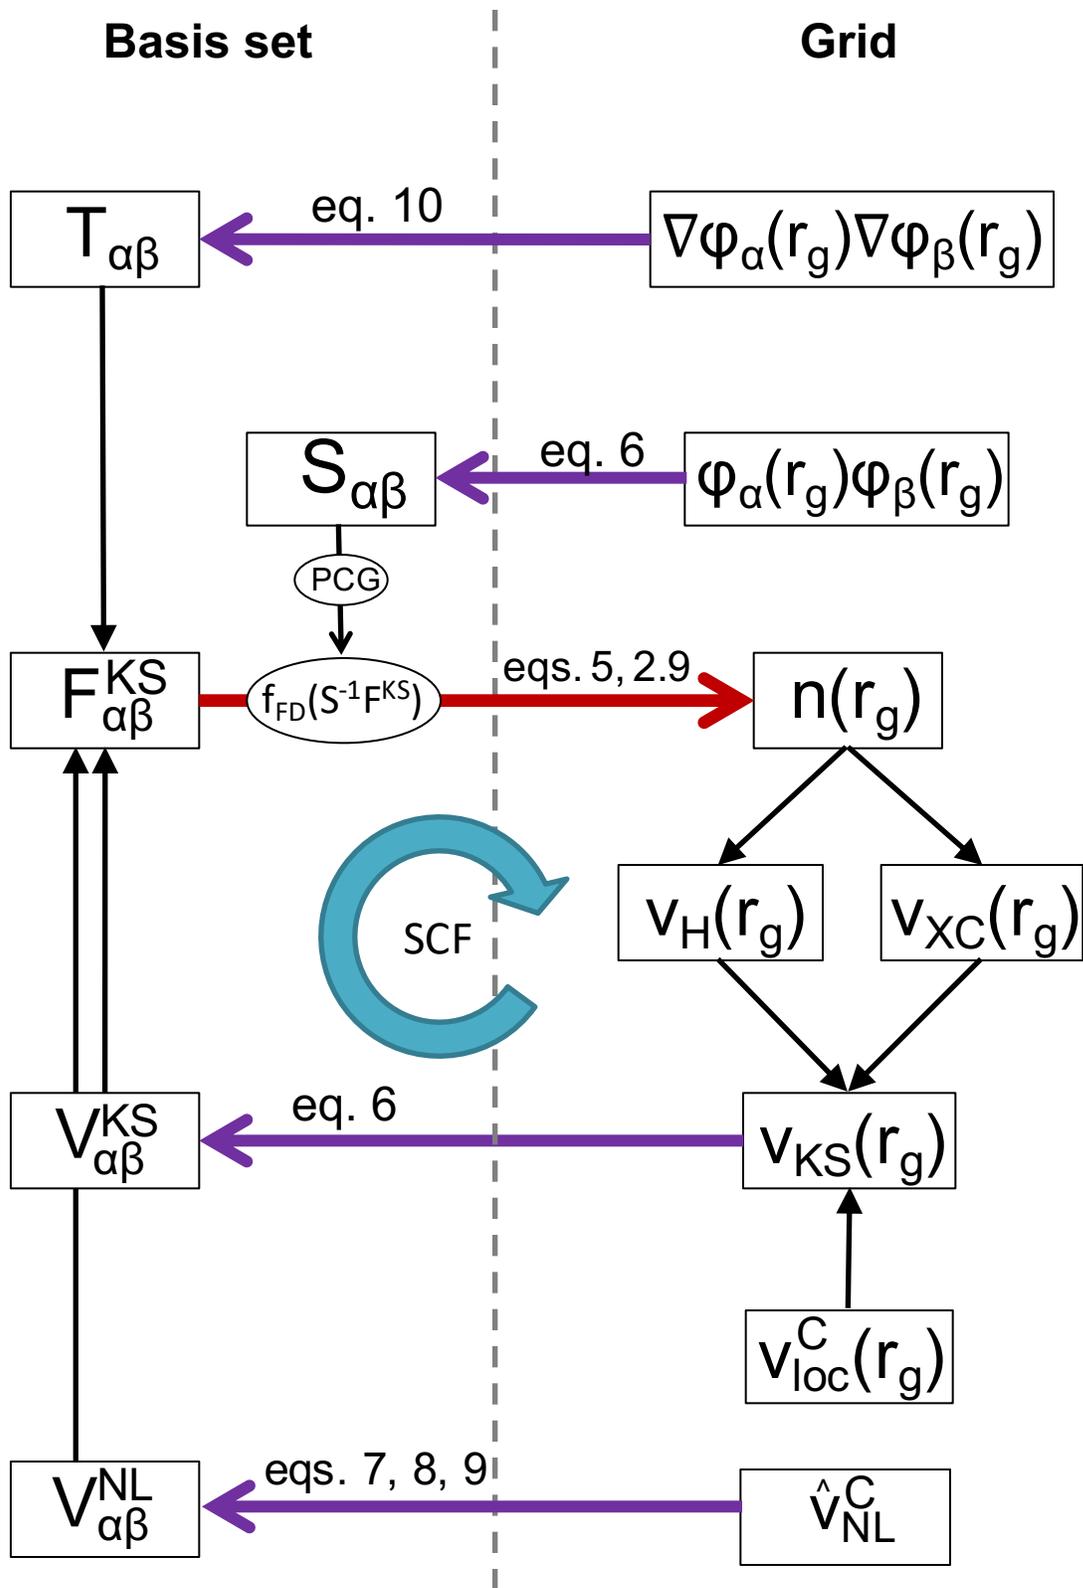


Figure 3: Flowchart for the KS-DFT calculation indicating the stages of the processing and the transformation of data back and forth between the grid and basis set representation during the SCF procedure

$U_0 \leftarrow S^{-1}u_\alpha$, $U_1 \leftarrow H_N U_0$ (where $H_N = \frac{H-\bar{E}}{\Delta E}$ is the “normalized” Hamiltonian) and $l = 2$, then we compute $\rho_\alpha \leftarrow a_0(T, \mu)U_0 + a_1(T, \mu)U_1$, and loop until $l = N_C$ the following steps:

$$U_2 \leftarrow 2H_N U_1 - U_0, \quad (12)$$

$$\rho_\alpha \leftarrow \rho_\alpha + a_l(T, \mu)U_2, \quad (13)$$

$$U_0 \leftarrow U_1, \quad U_1 \leftarrow U_2 \quad (14)$$

$$l \leftarrow l + 1 \quad (15)$$

C Parallelization

In this section we will discuss the main advantage of this algorithm, which is the straightforward path to parallelization. To understand the parallelization strategy a flowchart of the program is depicted in figure 4. All parallelization shown here is done through the *message passing interface* (MPI).

There are a total number of M processors that can be used and between which the workload in the SCF cycle is split. At first all processors are producing the same initial guess density and an initial KS potential, $n_0(\mathbf{r}_g)$ and $v_{KS_0}(\mathbf{r}_g)$ respectively. This $v_{KS_0}(\mathbf{r}_g)$ is then transferred from the grid to the matrix representation with equation 6. This step is parallelized in such a way that each processor gets a batch of matrix elements $V_{\alpha\beta}^{KS}$ that it needs to calculate and in the end these separate matrix elements are combined and distributed to every processor. The KS-Fock matrix is updated next for every processor, where $F^{KS} = T + V^{NL} + V^{KS}$. T and V^{NL} are calculated once before the SCF cycle according to equation 9 and 10 and are parallelized in the same way as described for V^{KS} . Each processor is independently calculating an electron density on the grid according to equation 5. For this, each processor gets a different set of unit vectors u_α , where the total number of vectors is just the basis set size K . Each processor therefore gets to calculate a fraction $r = K/M$

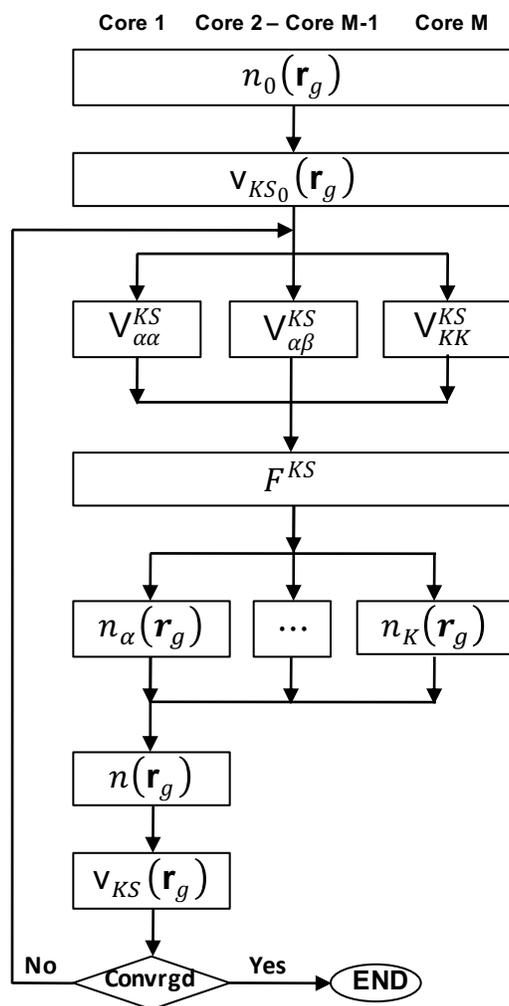


Figure 4: Parallelization strategy for the SCF cycle employed in the program

of the densities. The maximal parallelization that can be achieved in this step is therefore for $M = K$ and $r = 1$. All the densities from the different processors are summed on the first (“master”) processor. Only the master process calculates the new potential $v_{KS}(\mathbf{r}_g)$ and checks if the calculation reached the defined convergence criteria ($\Delta E = E^{SCF} - E^{SCF-1}$). If this is the case the calculation finishes, otherwise $v_{KS}(\mathbf{r}_g)$ is distributed to all processors and the SCF cycle is repeated from the calculation of the KS potential matrix. For more details on the SCF cycle see also subsection A.3.

D Systems

Table 2: Silicon clusters employed in the calculations

System	K (STO-3G)	K (6-31G)	$\#e^-$	$\#$ atoms
Si ₃₅ H ₃₆	176	352	176	71
Si ₈₇ H ₇₆	424	848	424	163
Si ₃₅₃ H ₁₉₆	1608	3216	1608	549
Si ₇₀₅ H ₃₀₀	3120	6240	3120	1005
Si ₁₀₆₃ H ₄₁₂	4664	9328	4664	1475
Si ₁₃₇₉ H ₄₇₆	5992	11984	5992	1855
Si ₂₀₃₁ H ₆₂₈	8752	17504	8752	2659
Si ₂₇₈₅ H ₇₈₀	11920	23840	11920	3565

Table 3: Water clusters employed in the calculations (taken from <http://www.ergoscf.org/xyz/h2o.php>)

System	K (STO-3G)	K (6-31G)	$\#e^-$	$\#$ atoms
(H ₂ O) ₁₀	60	120	80	30
(H ₂ O) ₂₀	120	240	160	60
(H ₂ O) ₃₂	192	384	256	96
(H ₂ O) ₄₇	282	564	376	141
(H ₂ O) ₇₆	456	912	608	228
(H ₂ O) ₁₀₀	600	1200	800	300
(H ₂ O) ₁₃₉	834	1668	1112	417
(H ₂ O) ₁₉₀	1140	2280	1520	570
(H ₂ O) ₂₃₇	1422	2844	1896	711
(H ₂ O) ₃₀₁	1806	3612	2408	903
(H ₂ O) ₃₈₄	2304	4608	3072	1152
(H ₂ O) ₄₇₁	2826	5652	3768	1413
(H ₂ O) ₅₇₃	3438	6876	4584	1719
(H ₂ O) ₆₉₂	4152	8304	5536	2076
(H ₂ O) ₈₁₆	4896	9792	6528	2448
(H ₂ O) ₉₆₄	5784	11568	7712	2892
(H ₂ O) ₁₁₁₂₀	6720	13440	8960	3360
(H ₂ O) ₁₂₉₃	7758	15516	10344	3879
(H ₂ O) ₁₄₈₁	8886	17772	11848	4443
(H ₂ O) ₁₆₉₈	10188	20376	13584	5094
(H ₂ O) ₁₉₂₄	11544	23088	15392	5772
(H ₂ O) ₂₁₆₅	12990	25980	17320	6492
(H ₂ O) ₂₄₆₉	14814	29628	19752	7407
(H ₂ O) ₂₇₃₇	16422	32844	21896	8211
(H ₂ O) ₃₀₅₀	18300	36600	24400	9150

E Strong scalability $\text{Si}_{705}\text{H}_{300}$

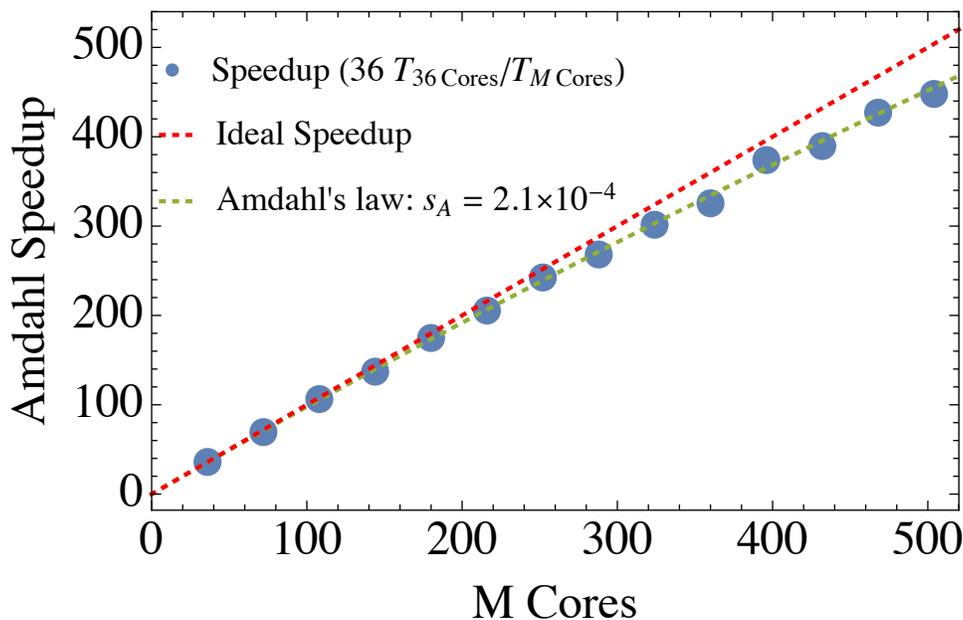


Figure 5: Strong scalability speedup analysis for $\text{Si}_{705}\text{H}_{300}$. The reference time for 1 processor for the speedup is extrapolated from $T_1 = 36T_{36}$. The calculations used the 6-31G basis-set (6240 basis functions) within the LDA and were performed on several 2.60GHz Intel Xeon Gold 6240 with 256 GB using 10Gb Ethernet networking communications.

References

- (1) Helgaker, T.; Jorgensen, P.; Olsen, J. *Molecular Electronic-Structure Theory*; Wiley, 2014.
- (2) Scott, J.; Tuma, M. HSL_MI28: An Efficient and Robust Limited-Memory Incomplete Cholesky Factorization Code. *ACM Trans. Math. Softw.* **2014**, *40*, 1–19.
- (3) Martyna, G. J.; Tuckerman, M. E. A reciprocal space based method for treating long range interactions in ab initio and force-field-based calculations in clusters. *J. Chem. Phys.* **1999**, *110*, 2810–2821.
- (4) Baer, R.; Head-Gordon, M. Chebyshev expansion methods for electronic structure calcu-

lations on large molecular systems. *The Journal of Chemical Physics* **1997**, *107*, 10003–10013.