Massively Parallel Fragment-Based Quantum Chemistry for Large Molecular Systems: The SERESTIPY Software

Patrick Eschenbach[†], Niklas Niemeyer[†] and Johannes Neugebauer[*]

Theoretische Organische Chemie, Organisch-Chemisches Institut

and Center for Multiscale Theory and Computation,

Westfälische Wilhelms-Universität Münster,

Corrensstraße 36, 48149 Münster, Germany

[†]Both authors contributed equally.

Date: September 9, 2022

[*]email: j.neugebauer@uni-muenster.de

Abstract

We describe the SERESTIPY software, which is an add-on to the quantum-chemistry program SERENITY. SERESTIPY is a representational-state transfer-oriented application programming interface written in the Python programming language enabling parallel subsystem density-functional theory calculations. We introduce approximate strategies in the context of frozen-density embedding time-dependent density-functional theory to make parallel large-scale excited-state calculations feasible. Their accuracy is carefully benchmarked with calculations for large assemblies of porphine molecules. We apply this framework to a theoretical model nanotube consisting of rings of porphine monomers, with 12,160 atoms (or 264,960 basis functions) in total. We obtain its electronic structure and absorption spectrum in less than a day of computation time.

1 Introduction

Computational studies on large molecular systems of chemical interest demand high computational effort employing *ab-initio* descriptions in quantum chemistry [1–4] (QC). Traditionally, high-performance computing (HPC) infrastructure is utilized to perform calculations requiring a large amount of computational resources. This type of calculations, however, is often limited by the capacity and availability of the chosen HPC center. As the internet has seen a drastic development in the last years [5–21] these limitations can be mitigated by deploying QC software into the available cloud infrastructure [22–24]. Cloud-based QC opens the possibility to dynamically scale the needed resources depending on the system of interest. With these seemingly infinite resources of the cloud at hand, it is necessary to integrate traditional QC software into that framework. Therefore, it is mandatory to employ application programming interfaces (APIs) that provide a uniform gateway between a possible user and the QC software running as backend. Recently, the open chemistry project [25, 26] has been launched. It provides an end-to-end platform for chemical-science research that integrates data from computational and experimental approaches through a modern web-based interface. In this project, the focus was to facilitate the interplay between experimental and computational researchers employing a shared platform that is easy to use for both parties. On a technical level, this platform uses a RESTful [27] API design which easily integrates into cloud-based infrastructure, because already large parts of the internet rely on these technologies. The QC backend, however, is the open-source NWChem [28] program.

In this work we will present the Serestipy software, which is inspired by design guidelines employed in the open-chemistry project. In contrast to the original idea of creating a uniform platform for experimental and computational researchers, we

will use such a framework to drastically improve the efficiency of QC studies on large molecular systems using the SERENITY [29, 30] program as a backend. To this end, we employ fragmentation approaches such as Frozen-Density Embedding [31] (FDE) that follow a divide-and-conquer strategy. The concept of "horizontal scaling" integrates perfectly with this fragmentation approach as each fragment is treated as a separate entity. The possibility to dynamically scale the needed resources in a cloud-based infrastructure enables massively-parallel calculations. Breaking down the full computational effort into smaller fragment-based calculation pieces drastically reduces the overall computational time. Therefore, in this work we will exemplarily tackle at an acceptable time scale massively-parallel computations of excited states of nano-scale systems in the form of model porphine nano-tubes that are usually completely out of reach for first-principles electronic-structure calculations. In particular, we will employ subsystem time-dependent density-functional theory [32–35] (subsystem TDDFT) with parallel FDE calculations as a reference [36, 37] to efficiently obtain excited states. Besides the presented example, this massively-parallel framework could be employed to efficiently perform tasks such as embarrassingly-parallel potential-energy surface constructions [38–43], parallel mode- and intensity-tracking or general semi-numerical frequency calculations [44–47], or many-body expansion calculations [48–52].

This article is structured as follows. First, we outline FDE theory in the context of ground and excited states. We propose an approximate inter-subsystem coupling strategy for excited states based on the simplified TDDFT [53, 54]. Next, technical tests are performed using SERESTIPY for clusters of 30 water molecules to compare the computational timings and the convergence behavior of the parallel and serial approaches. Next, the porphine-nanotube model system used in this study is introduced. We discuss the absorption spectrum of porphine and assess the accuracy

of approximate transition-charge couplings. We demonstrate the efficacy of the implementation and the parallel FDE approach for FaT and subsystem TDDFT calculations for the nanotube system. Finally, we conclude from our results.

2 Theory

2.1 Frozen-Density Embedding

The starting point for the derivation of FDE [31] is subsystem density-functional theory [55, 56] (subsystem DFT) as a fragmentation ansatz within density-functional theory [57, 58]. The main idea of subsystem DFT is to partition the total electron density ρ_{tot} of some system into subsystem densities ρ_I

$$\rho_{\mathrm{tot}}(\mathbf{r}) = \sum_I \rho_I(\mathbf{r}), \tag{1}$$

where I labels the subsystems and $\mathbf{r}$ collects the spatial coordinates. On the grounds of DFT being an exact theory, this partitioning is intrinsically not an approximation. As a result, the total energy of the system can be written formally as a sum over intra- and inter-subsystem contributions

$$E[\{\rho_I\}] = \sum_I E^{\mathrm{intra}}[\rho_I] + \sum_{I<J} E^{\mathrm{inter}}[\rho_I, \rho_J], \tag{2}$$

assuming that the intra- and inter-subsystem energies are strictly additive.

FDE can be considered a specific realization of subsystem DFT that can actually be put to practical use. The essential prerequisites for an FDE-based description of a supersystem partitioned into subsystems are: (i) for each subsystem there exists a separate non-interacting Kohn–Sham-like reference system [59] (along with

the corresponding auxiliary Slater determinant and orbitals), (ii) in the limit of a vanishing environment, FDE reduces to regular KS theory, (iii) the density of each subsystem "experiences" an effective KS-like potential, and (iv) the number of electrons of each subsystem is constrained to be constant.

Consideration (i) directly implies that each ρ_I can be written as a sum over simple one-particle densities

$$\rho_I(\mathbf{r}) = \sum_i \left|\phi_i^I(\mathbf{r})\right|^2, \tag{3}$$

where i labels the occupied orbitals ϕ of subsystem I. The total electronic supersystem FDE energy can be derived as

$$\begin{aligned} E[\{\rho_I\}, \rho^{\mathrm{tot}}] = &\sum_I T_{\mathrm{s}}[\phi_i^I] + \sum_I E_{\mathrm{xc}}[\rho_I] + V_{\mathrm{nuc}}[\rho_{\mathrm{tot}}] + J[\rho_{\mathrm{tot}}] \\ &+ T_{\mathrm{s}}^{\mathrm{nadd}}[\{\rho_I\}, \rho^{\mathrm{tot}}] + E_{\mathrm{xc}}^{\mathrm{nadd}}[\{\rho_I\}, \rho^{\mathrm{tot}}], \end{aligned} \tag{4}$$

where the kinetic energy of subsystem I is written in terms of its KS orbitals

$$T_{\mathrm{s}}[\rho_I] = \min_{\{\phi_i^I\} \to \rho_I} \left(-\sum_i \left\langle \phi_i^I \right| \frac{\boldsymbol{\nabla}^2}{2} \left| \phi_i^I \right\rangle \right), \tag{5}$$

and the non-additive energy contributions are introduced to make the expression formally equal to the supersystem energy

$$F^{\mathrm{nadd}}[\{\rho_I\}, \rho^{\mathrm{tot}}] = F[\rho^{\mathrm{tot}}] - \sum_I F[\rho_I] \tag{6}$$

for $F = T_{\mathrm{s}}$, E_{xc}. To obtain the density of subsystem I that minimizes its energy in the presence of all other subsystems, a self-consistent solution of some set of KS-like one-particle Schrödinger equations that accounts for both the KS potential

of subsystem I as well as all other subsystems is required. These local potential contributions enter the one-particle equations via a Fock-like operator for each subsystem I. The local potential whose orbitals fulfill these equations can be shown to be the derivative of the FDE energy with respect to the density of subsystem I. The corresponding Fock-type operator is given as,

$$\hat{F}^I = \left(-\frac{\boldsymbol{\nabla}^2}{2} + v^I_{\text{eff}}[\rho_I](\mathbf{r}) + v^I_{\text{emb}}[\rho_I, \rho^{\text{tot}}](\mathbf{r}) \right), \tag{7}$$

which, accounting for (iv), enters the Kohn–Sham equations with constrained electron density (KSCED) for each subsystem I as

$$\hat{F}^I \phi^I_i(\mathbf{r}) = \epsilon^I_i \phi^I_i(\mathbf{r}). \tag{8}$$

It is understood that v^I_{eff} is the effective KS potential evaluated for the density of subsystem I (and the corresponding nuclear contribution) and v^I_{emb} is the embedding potential which arises from the frozen environment densities (and the corresponding nuclei) and acts on subsystem I

$$v^I_{\text{emb}}[\rho_I, \rho^{\text{tot}}](\mathbf{r}_1) = \sum_{J \neq I} \left(v^{\text{nuc}}_J(\mathbf{r}_1) + \int \frac{\rho_J(\mathbf{r}_2)}{r_{12}} \, \mathrm{d}\mathbf{r}_2 \right) + \frac{\delta (T^{\text{nadd}} + E^{\text{nadd}}_{\text{xc}})[\rho_I, \rho^{\text{tot}}]}{\delta \rho_I(\mathbf{r}_1)}. \tag{9}$$

Here v^{nuc}_J is the potential arising from the nuclei of subsystem J and the density derivative of the non-additive contributions is given as

$$\frac{\delta F^{\text{nadd}}[\rho_I, \rho^{\text{tot}}]}{\delta \rho_I(\mathbf{r})} = \left. \frac{\delta F[\rho]}{\delta \rho(\mathbf{r})} \right|_{\rho = \rho_{\text{tot}}} - \left. \frac{\delta F[\rho]}{\delta \rho(\mathbf{r})} \right|_{\rho = \rho_I}. \tag{10}$$

Both KS-DFT and subsystem DFT are formally exact in the limit that exact functionals are used, provided that the corresponding representability conditions of

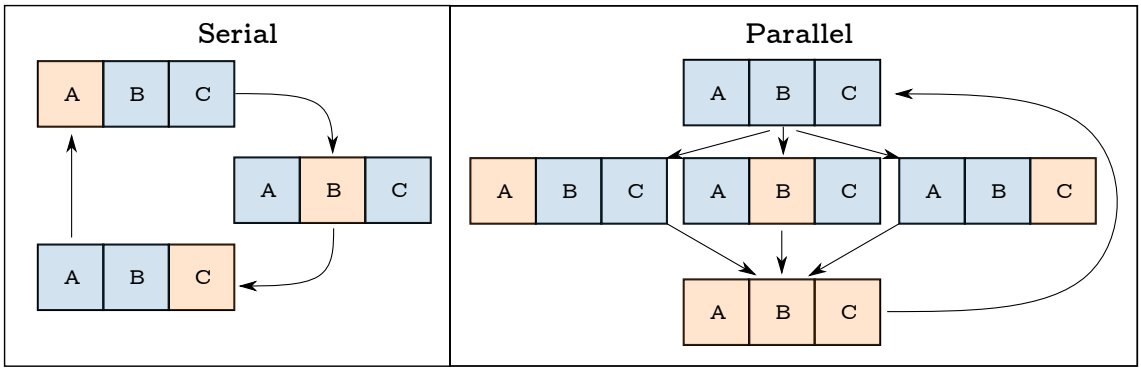

Figure 1: Comparison of serial and parallel FaT. Each box represents a subsystem, and orange and blue denote relaxed and frozen subsystems, respectively.

the densities are met. However, FDE introduces additional approximations through the usage of approximate non-additive functionals (this also holds for subsystem DFT) and not fully relaxing the densities of the environment subsystems. More specifically, if the same exchange–correlation (XC) functional approximation is used for evaluating the intra-subsystem and non-additive XC energy, an approximation is needed for the non-additive kinetic energy (NAKE) contributions. This is due to the fact that the intra-subsystem kinetic energy is evaluated with the subsystem KS orbitals while for the non-additive kinetic contributions one must resort to approximate density functionals because the supersystem orbitals are unknown. As a result, the main area of application of subsystem DFT using approximate NAKE functionals are weakly-interacting subsystems. For more strongly-interacting subsystem connected via for example covalent bonds, other ways of describing the subsystem interaction such as potential-reconstruction [60–68] or projection techniques [69–72] are usually more suitable.

In practice, FDE calculations are performed by solving Eq. (8) for each subsystem in the presence of all other, frozen subsystems in an alternating fashion until some criterion for the self consistency of *e.g.* the subsystem densities is met. This procedure is commonly called *Freeze-and-Thaw* [73] (FaT) and will be referred to as such in

the following. There are two modes in which a FaT procedure can be performed, which are schematically depicted in Fig. 1. In a cycle of a serial FaT, the densities of each subsystem are relaxed subsequently and are made available for the updates of the other subsystems immediately, on the one hand. In a cycle of a parallel FaT, on the other hand, the densities of each subsystem are only updated after all subsystem densities have been relaxed. In the parallel case, each individual FDE calculation can be carried out on independent worker compute nodes (WCN), which implies that the parallel mode has some conceptual advantages regarding cross-node parallelization.

2.2 Frozen-Density Embedding for Excited States

Linear-response subsystem time-dependent density-functional theory (subsystem TDDFT or FDE-TDDFT) is a variant of subsystem DFT for the description of excited states [32–35]. The linear-response of the density $\delta\rho_{\mathrm{tot}}$ to a frequency(ω)-dependent perturbation is partitioned into individual subsystem contributions

$$\delta\rho_{\mathrm{tot}}(\mathbf{r},\omega) = \sum_I \delta\rho_I(\mathbf{r},\omega), \tag{11}$$

where $\delta\rho_I(\mathbf{r})$ is the density response of subsystem I which is expanded in elements of that subsystem's first-order density matrix δP:

$$\delta\rho_I(\mathbf{r},\omega) = \sum_{(pq)_I} \delta P^I_{pq}(\omega)\phi^I_p(\mathbf{r})\phi^I_q(\mathbf{r}). \tag{12}$$

Using a time-dependent perturbation treatment, one finds that (i) only the occupied–virtual and virtual–occupied blocks of δP contribute to the linear response of density and (ii) they can be written in terms of matrix elements of the change in the

perturbing potential δv^I_{pq}

$$\delta P^I_{pq}(\omega) = \frac{1}{\omega - (\epsilon^I_p - \epsilon^I_q)} \delta v^I_{pq}(\omega). \tag{13}$$

Within the electric-dipole and adiabatic approximations, the perturbation of δv^I_{pq} consists of two parts, *i.e.* one that is a result of the (spatially-uniform) external perturbation, such as an electric-dipole perturbation, and one part that is due to an induced change in the potential caused by the linear response of the subsystem I itself and *all* other subsystems. Considering the matrix elements of these changes in potentials, one finds

$$\delta v_{(ia)_I}(\omega) = \delta v^{\text{ext}}_{(ia)_I}(\omega) + \delta v^{\text{ind}}_{(ia)_I}(\omega), \tag{14}$$

where $\delta v^{\text{ext}}_{ia}$ describes the external perturbation in the orbital-transition space of subsystem I and $\delta v^{\text{ind}}_{ia}$ describes the change of the potential due to a change in the subsystem densities. This is assumed to be instantaneous in the following, *i.e.* retardation effects will be neglected. The change in the induced potential can further be derived in terms of the first-order density matrix as

$$\delta v^{\text{ind}}_{(ia)_I}(\omega) = \sum_J \sum_{(jb)_J} K_{(ia)_I,(jb)_J} \left(\delta P^J_{jb}(\omega) + \delta P^J_{bj}(\omega) \right), \tag{15}$$

where the subsystem TDDFT coupling matrix K is introduced. Following the generalization of subsystem TDDFT with exact exchange (see Refs. [74] and [75]) and employing the (11|22) notation for the two-electron integrals, its elements read

$$\begin{aligned} K_{(ia)_I,(jb)_J} = (\phi^I_i \phi^I_a | \phi^J_j \phi^J_b) + \delta_{IJ} \left[(\phi^I_i \phi^I_a | f^I_{\text{xc}} | \phi^I_j \phi^I_b) - c_{\text{HF}} (\phi^I_i \phi^I_j | \phi^I_a \phi^I_b) \right] \\ + (\phi^I_i \phi^I_a | f^{\text{nadd}}_{\text{xc}} | \phi^J_j \phi^J_b) + (\phi^I_i \phi^I_a | f^{\text{nadd}}_{\text{kin}} | \phi^J_j \phi^J_b). \end{aligned} \tag{16}$$

The first part describes the classical Coulomb interaction between the responses of subsystems I and J. The second part includes the non-classical intra-subsystem XC interaction via the intra-subsystem XC kernel f_{xc}^{I} as well as Hartree–Fock exchange with an amount of c_{HF} as used in the respective approximation for the XC functional. Within the adiabatic approximation, f_{xc}^{I} reduces to

$$f_{\mathrm{xc}}^{I}(\mathbf{r}_1, \mathbf{r}_2) = \left. \frac{\delta^2 E_{\mathrm{xc}}[\rho]}{\delta\rho(\mathbf{r}_1)\delta\rho(\mathbf{r}_2)} \right|_{\rho=\rho_I} . \tag{17}$$

The non-additive XC and kinetic kernels stem from the non-additive XC and kinetic energy potentials and read

$$f^{\mathrm{nadd}}(\mathbf{r}_1, \mathbf{r}_2) = \left. \frac{\delta^2 F[\rho]}{\delta\rho(\mathbf{r}_1)\delta\rho(\mathbf{r}_2)} \right|_{\rho=\rho_{\mathrm{tot}}} - \delta_{IJ} \left. \frac{\delta^2 F[\rho]}{\delta\rho(\mathbf{r}_1)\delta\rho(\mathbf{r}_2)} \right|_{\rho=\rho_I} . \tag{18}$$

with $F = E_{\mathrm{xc}}, T_{\mathrm{s}}$. Inserting Eq. (15) into Eq. (13) and subsequently rearranging the equations for the external potential, one finds [74, 76]

$$\left[\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{pmatrix} - \omega \begin{pmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & -\mathbf{1} \end{pmatrix} \right] \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix} = - \begin{pmatrix} \mathbf{Q} \\ \mathbf{R} \end{pmatrix}, \tag{19}$$

where the following definitions apply

$$A_{(ia)_I,(jb)_J} = \delta_{IJ}\delta_{ab}\delta_{ij}(\epsilon_a^I - \epsilon_i^I) + K_{(ia)_I,(jb)_J}, \tag{20}$$

$$B_{(ia)_I,(jb)_J} = K_{(ia)_I,(bj)_J}, \tag{21}$$

for the matrix elements and

$$X_{(ia)_I} = \delta P_{ia}^I, \qquad Q_{(ia)_I} = \delta v_{(ia)_I}^{\mathrm{ext}}, \tag{22}$$

$$Y_{(ia)_I} = \delta P_{ai}^I, \qquad R_{(ia)_I} = \delta v_{(ai)_I}^{\mathrm{ext}}, \tag{23}$$

for the vector elements. For $\omega \to \omega^{0n}$, where ω^{0n} is the excitation energy to the n-th excited state, the linear response of the density diverges irrespective of the perturbation strength. As a result, excitation energies and corresponding transition densities can be obtained by setting $Q, R = 0$ and therefore as a solution of the generalized non-Hermitian eigenvalue problem [74–76]

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{pmatrix} \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix} = \omega \begin{pmatrix} \mathbf{1} & \mathbf{0} \\ \mathbf{0} & -\mathbf{1} \end{pmatrix} \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix}. \tag{24}$$

Scrutiny of Eqs. (20) and (21) makes the inherent subsystem block structure of the subsystem TDDFT response equations apparent [34]. Within the uncoupled approximation of subsystem TDDFT (FDEu), the contributions of the change in the induced potential where $I \neq J$ are neglected [see Eq. (15)]. This amounts to setting the inter-subsystem blocks of the K-matrix to zero and thus implies that the block of each subsystem of the fully coupled response matrix is decoupled from the response of all other subsystems. As a result, the transition densities of the uncoupled excitations are only expanded in the orbital-transition space of the respective subsystem and are therefore localized on subsystem only (assuming that monomer basis sets are employed). The FDEu approximation can be expected to give accurate results only for transitions that are mainly localized on only one subsystem [34].

This work is restricted to the Tamm–Dancoff approximation [77] of subsystem TDDFT [78] for computational convenience, *i.e.* we neglect the B-matrix in Eq. (24). Subsystem TDA-TDDFT (and subsystem TDDFT) calculations are generally performed in a two-step procedure: In the first step, some of the lowest-lying excitations of each subsystem are determined within the respective uncoupled approximation (FDEu step). In the second step, the corresponding uncoupled transition density

matrices are used to setup a coupled FDE (FDEc step) Hamilton-like matrix H [34]

$$H_{(m)_I(n)_J} = \sum_{(ia)_I} \sum_{(jb)_J} X^m_{(ia)_I} A_{(ia)_I,(jb)_J} X^n_{(jb)_J}, \tag{25}$$

which is defined in the basis of the uncoupled excitations m and n of subsystems I and J as determined in the FDEu step. Delocalized excitation energies with corresponding contributions from the local excitations can be extracted from eigenvalues and eigenvectors, respectively, of the H-matrix by solving the reduced eigenvalue equation

$$\mathbf{Hx} = \omega^{\mathrm{FDEc}}\mathbf{x}. \tag{26}$$

The FDEc procedure is somewhat ambiguous in the sense that coupled excitations energies and transition moments may highly depend on the number and choice of uncoupled excitations used to setup the reduced H-matrix. We note, however, that in the limit of *all* uncoupled excitations of *all* subsystems, Eq. (25) can be interpreted as an orthogonal transformation of the A-matrix, *i.e.* the eigenvalues of the H-matrix will be identical to the eigenvalues of the coupled A-matrix.

2.3 Approximate Inter-Subsystem Couplings

The TDA renders both the response matrix A and the FDEc Hamilton matrix H as well as their eigenvalue problems symmetric. As a result, the uncoupled eigenvectors of a particular subsystem trivially diagonalize the intra-subsystem block of the FDEc matrix belonging to that subsystem. It follows that the diagonal blocks of the coupling Hamilton matrix consist of the uncoupled excitation energies ω^I_m

$$H_{(m)_I(n)_J} = \delta_{IJ}\delta_{mn}\omega^I_m + (1-\delta_{IJ})V_{(m)_I(n)_J}, \tag{27}$$

which are coupled by virtue of FDEc coupling elements V on the off-diagonal blocks

$$V^{\mathrm{FDEc}}_{(m)_I(n)_J} = \sum_{(ia)_I} \sum_{(jb)_J} X^m_{(ia)_I} A_{(ia)_I,(jb)_J} X^n_{(jb)_J}. \tag{28}$$

In practice, the FDEc step can thus be simplified straightforwardly by (i) only calculating the coupling elements V explicitly and (ii) doing that only for $I < J$ to exploit the symmetry of the matrix. Nonetheless, this can be one of the most computationally demanding steps in subsystem TDDFT since the number of couplings to be calculated, in principle, grows quadratically with both the number of subsystems and uncoupled transitions determined beforehand (in the case that *all* excitations are coupled). In view of the fact that the excitonic coupling between the transitions of two subsystems may decay rapidly as a function of the subsystem separation, it is plausible that certain approximations can be used to simplify the calculation of these couplings for spatially distant subsystems.

2.3.1 Simplified Subsystem TDDFT

Our approximate coupling strategy can be considered a subsystem extension to the simplified TDA/TDDFT approaches originally published by Grimme an co-workers [53, 54]. In simplified TDDFT, the four-center integrals entering the response matrices are approximated with a simple monopole-type approximation. We propose the same approximation for entries of the subsystem TDDFT response matrices with reference orbitals from subsystem DFT. The elements of the A and B matrix would then read

$$\tilde{A}_{(ia)_I,(jb)_J} = \delta_{IJ}\delta_{ab}\delta_{ij}(\epsilon^I_a - \epsilon^I_i) + (\phi^I_i \widetilde{\phi^I_a|\phi^J_j} \phi^J_b) - (\phi^I_i \widetilde{\phi^J_j|\phi^J_a} \phi^I_b), \tag{29}$$

$$\tilde{B}_{(ia)_I,(jb)_J} = (\phi^I_i \widetilde{\phi^I_a|\phi^J_b} \phi^J_j) - (\phi^I_i \widetilde{\phi^J_b|\phi^J_a} \phi^I_j), \tag{30}$$

with the definitions for the approximate integrals

$$(\phi_i^I \widetilde{\phi_a^I|\phi_j^J} \phi_b^J) = \sum_{A\in I, B\in J} q^A_{(i)_I(a)_I} \Gamma^K_{AB} q^B_{(j)_J(b)_J}, \tag{31}$$

$$(\phi_i^I \widetilde{\phi_j^J|\phi_a^I} \phi_b^J) = \sum_{A\in I, B\in J} q^A_{(i)_I(j)_J} \Gamma^J_{AB} q^B_{(a)_I(b)_J}, \tag{32}$$

$$(\phi_i^I \widetilde{\phi_b^J|\phi_a^I} \phi_j^J) = c_{\mathrm{HF}} \sum_{A\in I, B\in J} q^A_{(i)_I(b)_J} \Gamma^K_{AB} q^B_{(a)_I(j)_J}, \tag{33}$$

where A and B denote atoms on either subsystem I or J and, again, c_{HF} the amount of Hartree–Fock exchange in the corresponding XC functional. We obtain the Γ-matrices as proposed in Refs. [53, 54],

$$\Gamma^K_{AB} = \left(\frac{1}{\left| \vec{R}_{(A)_I} - \vec{R}_{(B)_J} \right|^{\alpha} + \eta^{-\alpha}} \right)^{1/\alpha}, \tag{34}$$

$$\Gamma^J_{AB} = \left(\frac{1}{\left| \vec{R}_{(A)_I} - \vec{R}_{(B)_J} \right|^{\beta} + (c_{\mathrm{HF}} \cdot \eta)^{-\beta}} \right)^{1/\beta}, \tag{35}$$

where $\vec{R}_{(A)_I}$ and $\vec{R}_{(B)_J}$ are the position vectors of atoms A and B, respectively, and η is their average chemical hardness. The atomic partial charges q are obtained based on a Löwdin population analysis of the subsystem orbitals with orthogonalized molecular-orbital coefficients $\tilde{C}$ from

$$q^A_{(p)_I(q)_J} = \sum_{\mu\in A} \tilde{C}^I_{\mu p} \tilde{C}^J_{\mu q}, \tag{36}$$

and

$$\tilde{C}^I_{\mu p} = \sum_{\lambda} \sqrt{\mathbf{S}^I}_{\mu\lambda} C^I_{\lambda p}, \tag{37}$$

where $\mathbf{S}^I$ and $\mathbf{C}^I$ are the atomic-orbital overlap and coefficient matrices of subsystem I, respectively. In a subsystem basis (as it is employed in this work), basis functions are not shared between subsystem. As a result, the approximate exchange-like integrals $(\phi_i^I \widetilde{\phi_j^J | \phi_a^I} \phi_b^J)$ and $(\phi_i^I \widetilde{\phi_b^J | \phi_a^I} \phi_j^J)$ vanish (see Eq. 36) and the inter-subsystem coupling thus reduces to Coulomb-like integrals only

$$V^{\mathrm{sTDA}}_{(m)_I(n)_J} = \sum_{(ia)_I} \sum_{(jb)_J} X^m_{(ia)_I} \tilde{A}_{(ia)_I,(jb)_J} X^n_{(jb)_J} \tag{38}$$

$$= \sum_{(ia)_I} \sum_{(jb)_J} X^m_{(ia)_I} (\phi_i^I \widetilde{\phi_a^I | \phi_j^J} \phi_b^J) X^n_{(jb)_J}. \tag{39}$$

The formulation of inter-subsystem couplings given in Eq. (39) is particularly useful if approximate integrals are needed explicitly. In the context of subsystem TDDFT, this is for example the case for the calculation of linear-response properties which involves solving a linear equation system in the fully coupled orbital-transition space [76]. The present work, however, is concerned with couplings of specific transitions directly and it is thus advantageous to use a different formalism (see below).

2.3.2 Transition Charges

For the second approach for coupling excitations on different subsystems, we use the concept of *atomic transition charges*. The subsystem TDDFT transition-charge approach can be summarized as follows: The uncoupled transition density matrices are first transformed to the (monomer) atomic-orbital basis they are defined in

$$X^{m,\mathrm{AO}}_{(\mu\nu)_I} = \sum_{(ai)_I} C^I_{\mu i} X^m_{(ia)_I} C^I_{\nu a} \tag{40}$$

and then a Löwdin population analysis [79] is performed for the transition density to obtain atomic transition charges $Q^m_{(A)_I}$ for each transition m and atom A on

subsystem I

$$Q^m_{(A)_I} = \sum_{\mu \in (A)_I} \left(\sqrt{\mathbf{S}^I} \mathbf{X}^{m,\mathrm{AO}}_I \sqrt{\mathbf{S}^I} \right)_{\mu\mu} . \tag{41}$$

The quantum-mechanical subsystem TDDFT description of the inter-subsystem couplings is then effectively replaced by a mere Coulomb-like coupling based on the obtained atomic charges via

$$V^{\mathrm{TC}}_{(m)_I (n)_J} = \sum_{(A)_I (B)_J} \frac{Q^m_{(A)_I} Q^n_{(B)_J}}{\left| \vec{R}_{(A)_I} - \vec{R}_{(B)_J} \right|}, \tag{42}$$

where, again, $\vec{R}_{(A)_I}$ and $\vec{R}_{(B)_J}$ are the position vectors of atom A and B, respectively. We note the close relationship of these two approaches as both of them rely on partial charges obtained from a Löwdin orthogonalization. In particular, for the specific choice of $\eta = 0$ entering the Γ^K-matrix, transition-charge and Coulomb couplings from simplified subsystem TDDFT context become equivalent (see Sec. S1 of the Supporting Information).

Closely-related methods are the transition density cube method [80], the transition-density fragment interaction method [81–83] as well as the transition charge from electrostatic potential (TrESP) method [84] and extensions thereof [85]. In fact, the approach used here can be regarded as a simplified TrESP method, since we do not obtain transition charges from a least-squares fit involving the electrostatic potential but from a simple population analysis. In this work, we also compare the approximate couplings introduced above to the well-known dipole–dipole approximation as the leading term of the multipole expansion of the electronic coupling for dipole-allowed

transitions [86, 87]. In the subsystem TDDFT context, this coupling reads

$$V^{\mathrm{DD}}_{(m)_I(n)_J} = \frac{\left(\vec{\mu}_{(m)_I} \cdot \vec{\mu}_{(n)_J}\right)}{|\vec{R}_{IJ}|^3} - 3\frac{\left(\vec{\mu}_{(m)_I} \cdot \vec{R}_{IJ}\right)\left(\vec{\mu}_{(n)_J} \cdot \vec{R}_{IJ}\right)}{|\vec{R}_{IJ}|^5}, \tag{43}$$

where $\vec{\mu}_{(m)_I}$ is the transition dipole moment of the m-th excitation of I and $\vec{R}_{IJ}$ is the difference vector between the geometric centers of the two monomers I and J.

3 Computational Details

All calculations were performed with locally-modified versions of the Serenity [29, 30] and Serestipy programs. The approximate couplings strategies have been implemented in the former of the two and will be part of one of its upcoming releases. In all FaT and subsystem TDDFT calculations, PBE0 [88, 89], PW91 [90, 91], and PW91k [92] were used as approximations for the intra-subsystem XC, non-additive XC, and non-additive kinetic energy functionals, respectively (abb. PBE0/PW91/PW91k). Monomer def2-TZVP [93] basis sets were used throughout. Coulomb contributions of both the Fock matrix and response matrix-vector products were approximated with the RI approximation and the def2-universal-JFIT [94] and RI-FIT [95] (RI-C) basis sets corresponding to def2-TZVP, respectively. Grid accuracy settings of 4 as implemented in Serenity were used and potentials and kernels for the FDE and FDEu calculations, respectively, were evaluated employing grids located only on the considered subsystems. FDEc calculations employ grids on both of the coupled monomers. All FDEu eigenvectors were converged to residual norms of $< 10^{-5}$. In all FDEu calculations, 24 of the lowest-lying excitations were determined and the TDA [77, 78] is applied throughout.

4 Results and Discussion

4.1 Water Clusters

As already stated in Sec. 2.1, FaT calculations can be carried out in a serial or a parallel mode. In both modes we must consider two possible types of parallelization schemes. In the first case the QC calculation is carried out on one single WCN and parallelized employing OpenMP [96] or similar algorithms. This case, however limits the calculation to the resources available on the chosen WCN. Those calculations can only be scaled "vertically". Vertical scaling implies upgrading the available WCN by adding CPU cores or more main memory, which might be very costly. The second parallelization scheme makes use of the possibility to split the overall calculation into individual steps that can be carried out simultaneously on different WCNs. Each individual calculation might still be parallelized using algorithms such as OpenMP [96], as in the first mentioned case. These cases, therefore, can be scaled "horizontally" according to the system size. In this regard, horizontal scaling means

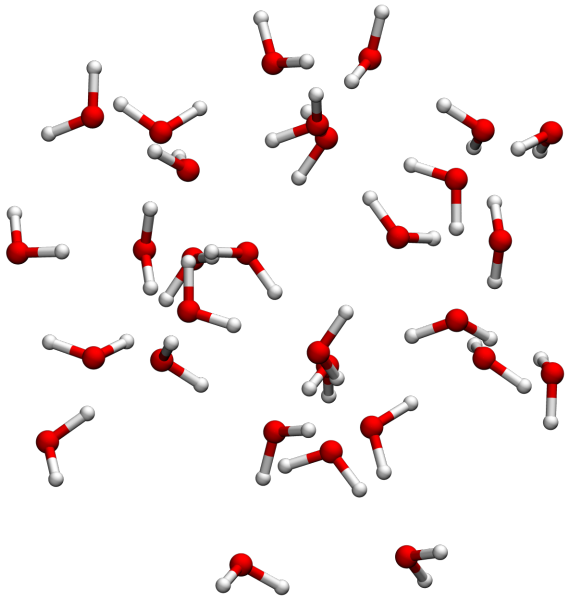

Figure 2: Water cluster consisting of 30 water molecules.

the possibility to easily add or remove WCNs that run the calculations. Serial FaT calculations might be carried out efficiently using the first parallelization scheme. The second scheme, however, is not available for serial FaT calculations, because of the fact that each FDE calculation must be carried out subsequently. Parallel FaT calculations are the most efficient when employing the second parallelization scheme, because each FDE calculation might be carried out simultaneously on different WCNs. In the following, we will refer to the first scheme as *intra-fragment* parallelization, while the second scheme is referred to as *inter-fragment* parallelization, in the context of subsystem DFT. Apart from that there exist intra-subsystem parallelization schemes, where the complete workload of a calculation is distributed over several WCNs employing principles such as the message passing interface [97] (MPI) protocol, but these will not be the focus of this work. The inter-fragment parallelization scheme is beneficial in a cloud-computing environment, because a seemingly infinite and easily scalable number of (small) WCNs is available. Our Serestipy framework can distribute the parallel FaT calculations efficiently by making use of the Serenity program as a backend employing an inter-fragment parallelization. The Serenity program only supports OpenMP parallelization, while an MPI scheme is not available. For implementation details of the Serestipy API, the reader is referred to Sec. S2 in the Supporting Information. To assess the wall-clock timings and convergence behavior of parallel and serial FaT in the following sections, calculations were carried out with Serestipy for a model water cluster consisting of 30 molecules (see Fig. 2).

4.1.1 Computational Timings

All presented calculations were performed on machines with a dual AMD EPYC 7552 48-core processor and 528 GB of main memory. Because the individual FDE calculations in parallel FaT runs can be performed on independent WCNs, each

FDE calculation can be performed with individual hardware settings. For serial FaT calculations, however, it is not possible to make use of cross-node parallelization in the case of the SERENITY program and the hardware requirements are, thus, limited by the individual WCN. All parallel FaT calculations use 2 CPU cores and 8 GB of main memory for each individual FDE calculation. Thus, the total number of CPU cores throughout the calculation is $2N$ where N is the number of water molecules included in the model. To provide a suitable comparison between the capabilities of parallel and serial FaT calculations, we chose two different strategies for the serial runs: (i) provide the same total amount ($2N$) of CPU cores for all the serial runs and (ii) keep the number of CPU cores fixed at a value of 2 for all the serial runs. Parallel FaT calculations are especially beneficial in cases where the available computational hardware does not allow for efficient serial runs and vertical scaling is not possible. This case is mimicked by the second strategy. Wall-clock timings following the first strategy are presented in Fig. 3, while wall-clock timings following

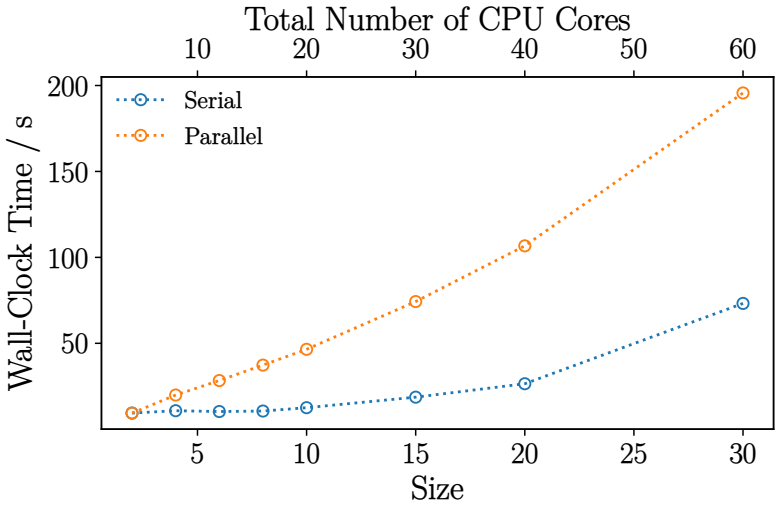

Figure 3: Wall-clock times for serial and parallel FaT runs for an increasing number of included water molecules. For the parallel runs 2 CPU cores have been used for the individual FDE calculations, while for the serial run $2 \cdot N$ CPU cores have been used for each individual FDE step, where N is the number of water molecules. This procedure ensured a comparable total number of CPU cores for parallel and serial runs. A total number of three cycles has been applied in both cases.

the second strategy are presented in Fig. 4. As can be seen from Fig. 3, where both types of FaT calculations use the same total amount of CPU cores, the serial FaT is clearly faster than the parallel FaT for all system sizes. This is because of the smaller amount of CPU cores for each FDE calculation carried out in parallel compared to the serial case. By contrast to the serial case, the communication overhead for the parallel runs slows down the whole procedure. Note that the serial FaT calculations following the first strategy are limited by the resources of the available WCNs: When increasing the system size for systems beyond 48 fragments, the wall-clock timings will drastically increase compared to smaller systems. The specific number of 48 fragments is a result of the available WCNs we used throughout this work, but can of course differ depending on the hardware. In case of the parallel FaT calculations, this hardware limitation is de facto not present (in our assumption of an effectively unlimited number of [small] nodes), as the user can choose to run additional parallel FDE steps on additional WCNs. This case is beneficial for the parallel FaT as can be seen from Fig. 4. In the second strategy, we simulated the hardware limitation,

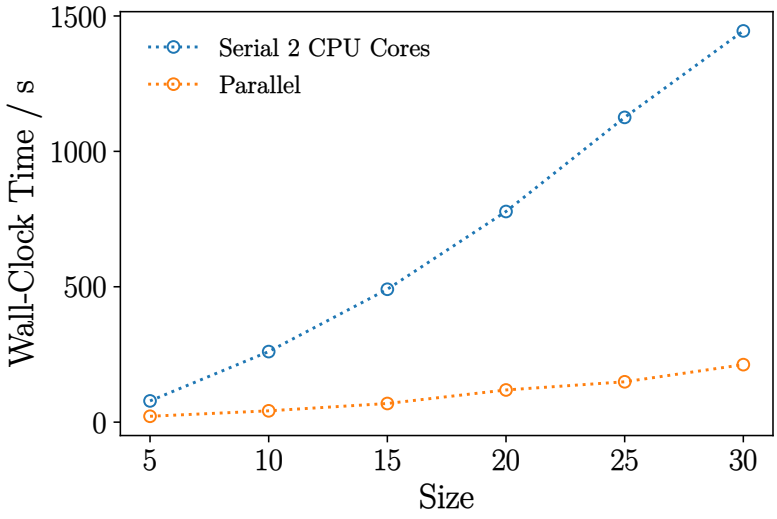

Figure 4: Wall-clock times for serial and parallel FaT runs for an increasing number of included water molecules. For the parallel runs 2 CPU cores have been used for the individual FDE calculations, while for the serial run 2 CPU cores have been used for each individual FDE step. A total number of three cycles has been applied in both cases.

by fixing the number of CPU cores to exactly 2 for every serial FaT calculation. As can be seen, the wall-clock timings of the serial FaT calculations increase more drastically for larger systems compared to those in Fig. 3. By contrast, the parallel FaT calculations are much faster than the corresponding serial ones.

4.1.2 Convergence Behavior

In addition to the wall-clock timings, we assessed the convergence behavior of parallel and serial FaT calculations with respect to the number of needed FaT cycles. Therefore, we chose a convergence threshold of 10^{-5} for the mean over all subsystems of the RMSD of the density-matrix change between two FaT cycles. As can be seen from Fig. 5, the parallel FaT requires the same amount or fewer FaT cycles compared to the serial case for all system sizes. This may seem surprising at first glance, but is connected to the way how the electron densities of the subsystems are optimized and updated during the FaT procedure (see Sec. 2.1). The sequential update scheme makes a serial FaT procedure dependent on the order of the applied

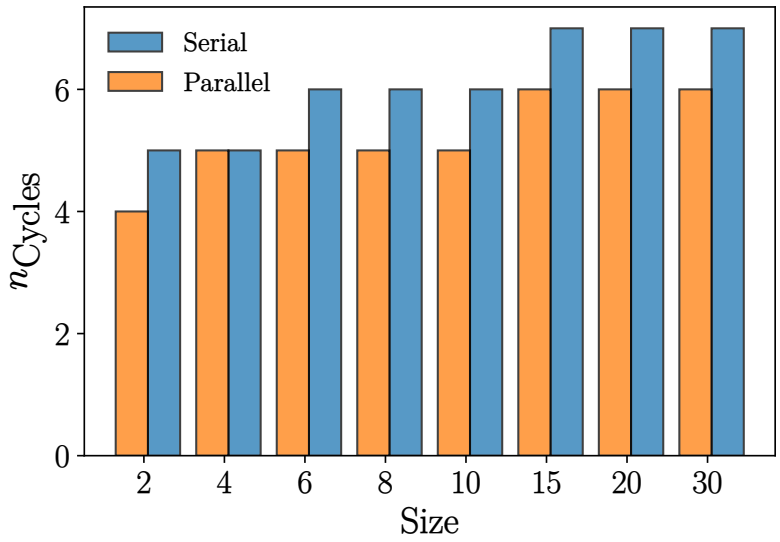

Figure 5: Number of needed parallel and serial FaT cycles until convergence threshold is reached for an increasing number of water molecules in the calculation.

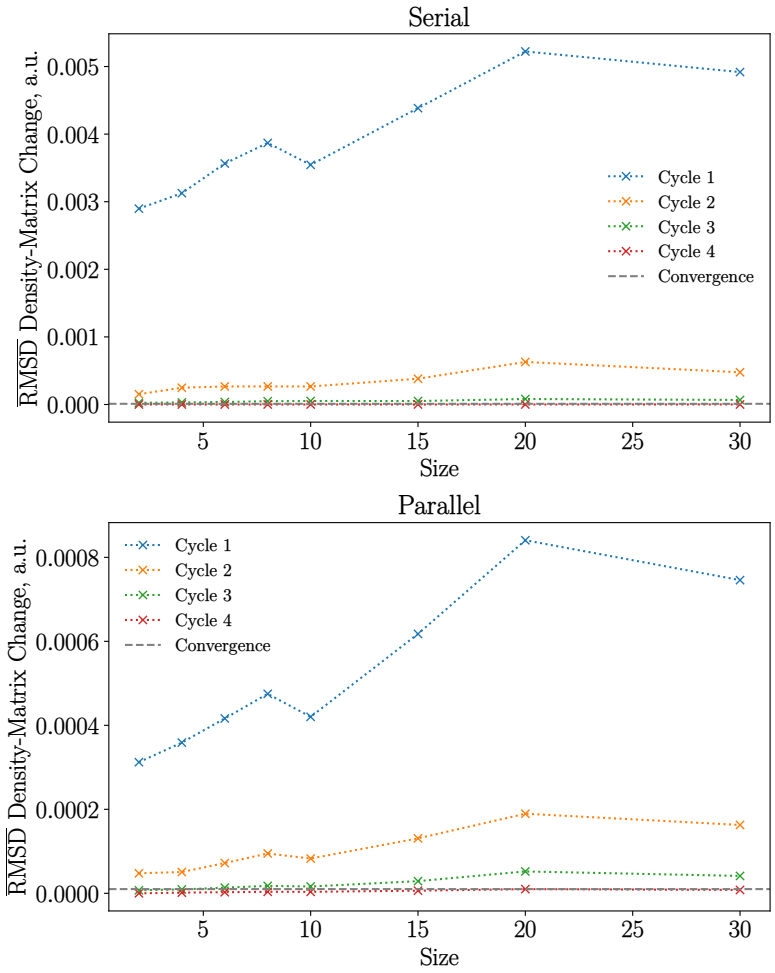

Figure 6: Mean over all subsystems of the RMSD of the density-matrix change between two FaT cycles calculated for the first four serial (top) and parallel (bottom) FaT cycles including an increasing number of water molecules.

FDE steps in case it is not fully converged. In the parallel FaT, each subsystem electron density is updated simultaneously during each FaT cycles, which makes the procedure independent of the order of FDE calculations. This difference in the density update scheme leads to the generally better convergence in parallel calculations. As can be seen from Fig. 6 (top) the electron densities change more strongly between the first three FaT cycles in the serial case than for the parallel case (bottom). The change in the electron density is one order of magnitude smaller in the parallel case than for the serial case. This trend does not change for increasing system sizes. In fact, the parallel density update between two parallel FaT cycles

leads to smaller density changes than in case of the serial procedure. Additionally, it is important to note that both methods converge to the same subsystem electron densities.

As can be seen from the presented results, serial FaT calculations are computationally more efficient provided the system under study is small- or medium-sized, in a sense that it consists only of a few fragments. On the downside, however, serial FaT calculations are dependent on the order when not fully converged. This is not beneficial, because the serial FaT calculations converge more slowly compared to the parallel counter part. We note that in calculations employing the three-partitioning version of FDE, the convergence behavior was reported to be better for the serial version [98, 99]. However, these calculations may not directly be comparable, since these calculations employ additional capping fragments and cap molecules, which may change the convergence behavior. In case of FDE, the parallel FaT is computationally most efficient for a large amount of similarly-sized fragments or in case the available hardware is not suitable for efficient serial calculations. It can be easily scaled to, in principle, infinitely large calculations (assuming a de factor unlimited number of WCNs). Additionally, individual hardware requirements can be specified during the calculation for differently sized fragments.

4.2 Porphine Nanotube

The porphine molecule is a basic building block in many cofactors/chromophores in nature. For example, chlorophyll molecules in photosynthetic reaction centers [100] or heme groups in hemoglobin [101], myoglobin [102] or cytochromes [103] all rely on porphine as structural motif because of its beneficial properties in terms of charge and excitation-energy transfer [104–109]. Apart from nature, this structural motif is also very useful in many compounds investigated in materials science [110–112]. This

is because of its π-stacking capabilities often leading to large-scale liquid crystals that have been reported several times in the literature [113–116]. As such, the porphine molecule is well-studied [104–109] both experimentally and theoretically in the literature and thus constitutes a good starting point for our purposes.

In the following, we first give an overview over the general excited-state properties of the porphine molecule. Subsequently, we describe the model setup of our artificial example of a porphine nanotube and give details about the massively-parallel calculations. Next, an assessment of the transition-charge approximation used in this work to compute inter-subsystem couplings is given. Finally, we present results for a porphine nanotube model system introduced below.

4.2.1 Absorption Spectrum

In this work, we are concerned with calculating excited-state properties of porphine nanotubes exploiting the high level of redundancy within the supermolecular system with the parallel approach. Therefore, we discuss briefly the absorption spectrum of the monomeric porphine molecule as a first step to determine an XC functional to use for the nanotube system.

In Fig. 7, we compare the experimental absorption spectrum of porphine with three computed ones, each employing a different XC functional [PBE0, LC-BLYP [117], and CAM-B3LYP [118]/def2-TZVP]. The experimental absorption spectrum of porphine in the UV/vis region consists of two bands with little absorption in the visible wavelength window. These are commonly called Q_x and Q_y bands [107]. Additionally, one finds a strong band in the near-UV region (called Soret or B band) along with two shoulders in the tail of higher energy (called N and L bands). Due to its dominant character in the absorption spectrum, in this work we will center our discussions around the Soret band. The involved transition wavelengths are

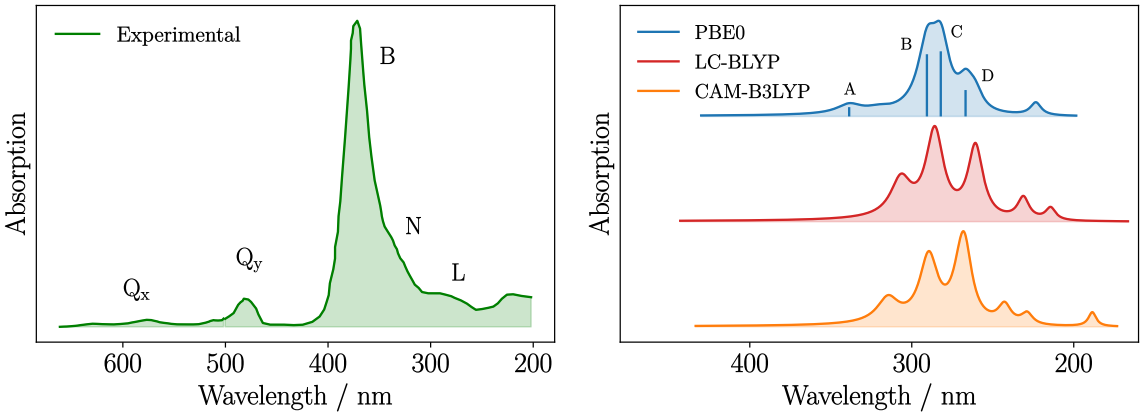

Figure 7: Comparison between the experimental absorption spectrum of free base porphine (left, digitized based on Ref. [105]) and three computed ones in dipole-length representation (right) [(PBE0, LC-BLYP, CAM-B3LYP)/def2-TZVP]. For the computed spectra, a def2-TZVP basis was used and 24 of the lowest-lying excitations were determined. The oscillator strengths were broadened using Lorentzian functions with a full width at half maximum of 0.2 eV.

underestimated by about 1.1 eV (corresponding to 100 nanometers in this wavelength regime) by all employed XC functionals. The general shape of the experimental absorption spectrum is fairly well reproduced with the PBE0 functional, whereas the deviation of the LC-BLYP and CAM-B3LYP spectra is somewhat larger. In particular, the agreement of the Soret band in relation to the bands in the lower-energy region is much better in the PBE0 case. Note, however, that for PBE0 the Soret band convolutes two close-lying transitions with a similar oscillator strength. Due to the better agreement of the PBE0 spectrum compared to the other functionals tested, in this work we will continue to employ the PBE0 functional as an XC functional and corresponding potentials and kernels.

4.2.2 Model and Calculation Setup

Since the goal of this study is to demonstrate the accuracy, efficiency, and scalability of subsystem (TD-)DFT calculations, we employ a technical model system consisting of porhpine molecules in a easily scalable way, and in which monomers interacting with various strengths are present. For these purposes, porphine molecules are first

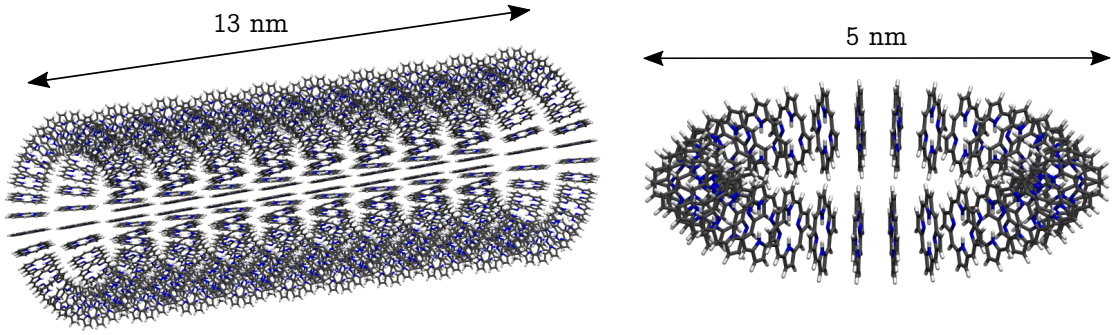

Figure 8: Atomic representation of the porphine rings as employed in this study. One ring consists of 32 monomers with identical distance and angle between them. The distance between the center of mass of one porphine and the center of the ring is 20 Å and the separation between the center of mass of porphines between two layers is 14 Å.

arranged in a circular fashion and the resulting porphine rings are stacked on top of each other to create nano-scale porphine tubes (see Fig. 8). From a parallelism point of view, this is especially convenient since it drastically simplifies the distribution of jobs to different nodes because all subsystems are of the same size and thus the resulting computational demand will be highly similar. We note, however, that we do not exploit the symmetry of the system explicitly in any way. Systems with subsystems of varying sizes (*e.g.* large proteins) are, in principle, also accessible in such calculations but would require to determine the computational resources for a particular job based on a metric concerning the subsystem size (like the number of basis functions).

We first optimized one porphine molecule using PBE0 as an approximation for the XC functional and a def2-TZVP basis set in combination with the D3 [119]-BJ [120] correction using TURBOMOLE 7.5 [121]. One ring of porphines is then defined by two parameters: The first one is the distance between the center of mass of each individual porphine molecule and the center of mass of the entire ring, which was chosen to be 20 Å. The second parameter is the number of porphine molecules in one ring (defining the angle between the porphine molecules) which was chosen to be 32.

The resulting distance between the π systems of two neighboring porphine molecules is then about 5–6 Å. To setup the nanotube, we have stacked the porphine rings on top of each other so that the distance between the center of mass of molecules of two different rings is 14 Å. We would like to stress that this model setup is artificial and entirely technically motivated, since the interaction strengths are in a typical range also observed in for example photosynthetic antenna systems (couplings of the order of 10^2 to 10^3 cm^{-1} for the nearest neighbors.)

4.2.3 Accuracy of the Transition-Charge Approximation

To assess the accuracy of the transition-charge (TC) approximation instead of subsystem TDDFT to compute inter-subsystem couplings, we consider one layer of porphine rings as depicted in Fig. 8 and compare the difference between the couplings computed following Eqs. (28) and (42) as a function of porphine molecules between the coupled monomers. For these calculations, the electronic structures of the monomers were converged in five FaT cycles and then 24 of the lowest-lying excitations were determined for each monomer. For the FDEu and FDEc calculations, all but the active monomers are used as environment systems. The data for the transitions indicated by bars in the PBE0 spectrum in Fig. 7 is displayed in Fig. 9.

We find a relatively strong coupling between close monomers for all considered transition pairs which decays rapidly with the inter-subsystem separation for both the FDEc, TC and simplified TDA couplings. For all transition pairs, a non-zero coupling can be identified up to the fith or sixth neighbor. We note in passing that the equivalence of the TC couplings and simplified subsystem TDA couplings with $\eta = 0$ is verified numerically by all excitations. Comparing the TC and simplified TDA couplings with the FDEc couplings for direct neighbors, we find that the former two are underestimated throughout. These deviations are caused by short-range

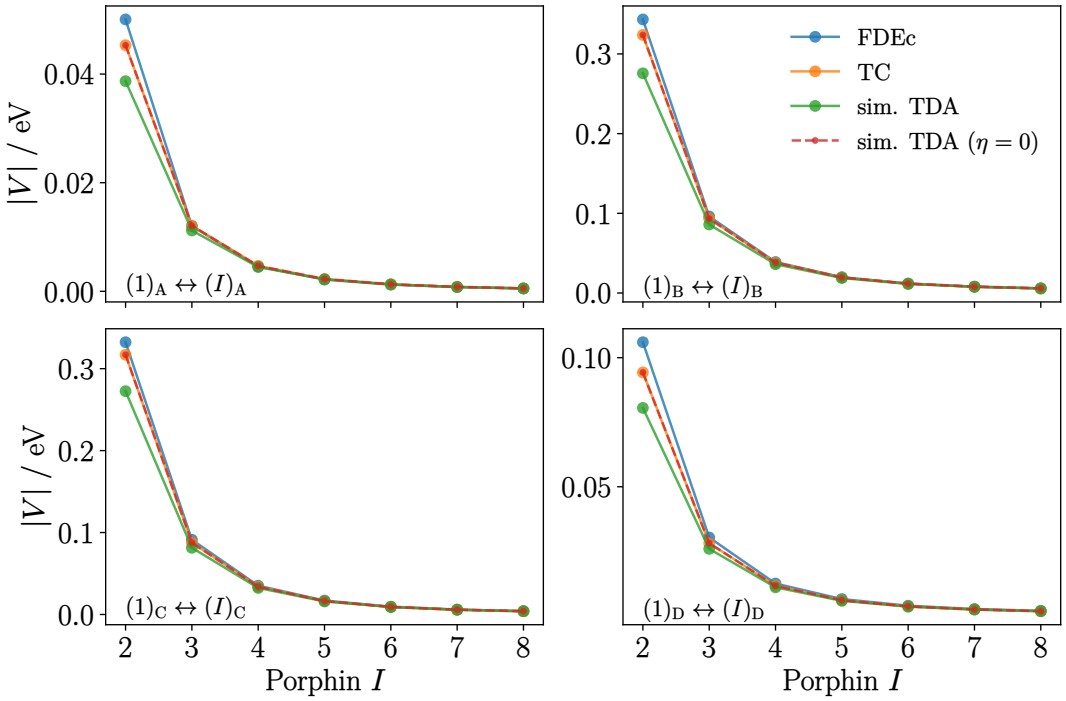

Figure 9: Comparison of couplings of some excitations with the same excitation on a neighboring porphine molecule between FDEc [PBE0/PW91/PW91k/def2-TZVP], transition charges computed based on the uncoupled transition density matrices and couplings based on the simplified TDA approach. For numerical validation of the derivation in Sec. S1 of the Supporting Information, the simplified TDA couplings with $\eta = 0$ are also shown in addition. The considered excitations are indicated by bars and labeled in the PBE0 spectrum in Fig. 7.

non-classical interactions. While these are, in principle, described by the non-additive kernels in the FDEc case, they are neglected by the transition-charge and simplified TDA coupling by construction. From the second neighbor on, FDEc couplings are reproduced by the TC couplings with very high accuracy for all of the considered transitions. As Coulomb interactions decay more slowly with the inter-subsystem distance, neglecting all other interactions for intermediate distances is a well-justified approximation. We find no advantage of sTDA couplings explicitly including the chemical hardnesses via a non-zero η-parameter for any of the considered excitations as they, in turn, further underestimate the TC couplings throughout. A comparison of the couplings shown in Fig. 9 to transition-dipole–transition-dipole couplings [see Eq. (43)] is given in Sec. S2 of the Supporting Information.

For our approximate coupling strategy for the nanotube consisting of ten rings, we conclude the following: (i) due to the strong deviations of FDEc and TC couplings between direct neighbors, we employ subsystem TDDFT without further approximations in the couplings according to Eq. (28) for these pairs (that is, direct neighbors of the same ring and direct neighbors between two rings), (ii) for all but direct neighbors in the same and adjacent rings, we employ approximate TC couplings as given in Eq. (42), due to their minor computational cost while being sufficiently accurate. For neighbors that are not part of the adjacent rings, we neglect the coupling altogether. We do not calculate simplified TDA or transition-dipole couplings for this example.

4.2.4 Workflow for Parallel FDE-TDDFT

The workflow for a parallel FDE-TDDFT procedure is presented in Fig. 10. As we are working on a traditional HPC cluster, first a meta-process is needed that spawns an API instance on each of the N WCNs, where N is the number of subsystems in the calculation. For bottom-up FDE calculations we first need initial, isolated subsystem densities for each subsystem. Those can be calculated completely independently, so that the isolated calculations are sent to individual WCNs (see Fig. 10 top). Once all calculations are finished, the results of the isolated runs are stored in a database directory. Subsequent calculation steps can be started on the basis of the stored results. After the isolated runs have finished, the parallel FaT cycles start. For this purpose, FDE calculations are sent to the individual WCNs. The individual FDE calculations are based on the densities stored in the database from the isolated runs. Note that on each WCN for the FDE calculation, a different subsystem is considered to be active while a specific set of other subsystems are considered set to be the frozen environment subsystems. The choice for these subsystems is depicted in Fig. 11 for all types of tasks employed in this article. We note that, while we

have hand-picked the environment systems in this study, this is a task which can be trivially automated via, for example, distance criteria. After all FDE calculations on all WCNs have finished, only the electron densities of each active system is stored in the database. These active system electron densities serve as starting point for the next parallel FaT cycle. We perform three such cycles as is generally enough to obtain converged properties (see Fig. 6). After the parallel FaT has converged for each system, individual FDEu-TDDFT calculations are sent to the WCNs. As in the parallel FaT runs, a different subsystem is considered to be the active one on each WCN (see Fig. 10 bottom). In each FDEu-TDDFT calculation, uncoupled excitation energies, transition moments and transition charges [see Eq. (41)] are determined for each excitation and are stored on disk. The position vector of the atoms is stored along with the transition charges. Next, FDEc tasks are sent to the WCNs to obtain the couplings of the excited states for all requested pairs of subsystems. The couplings are also stored on disk. Finally, after all FDEu and FDEc runs have finished, the Hamilton-like matrix H from Eq. (25) is assembled in a separate Python program by initializing its diagonal with the uncoupled excitation energies of each subsystem. The inter-subsytem blocks are then either read from disk

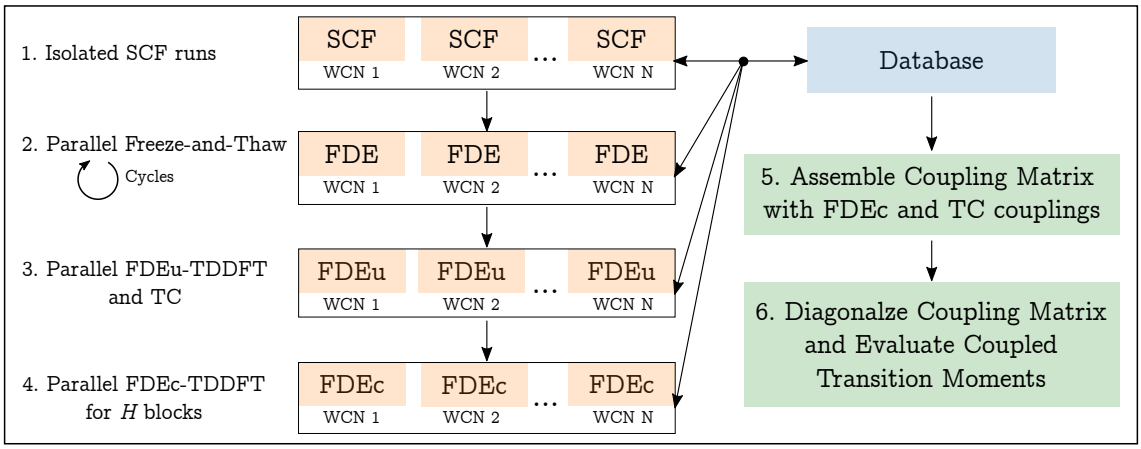

Figure 10: Worklow of parallel FDE-TDDFT calculation. For a more detailed description see the main text.

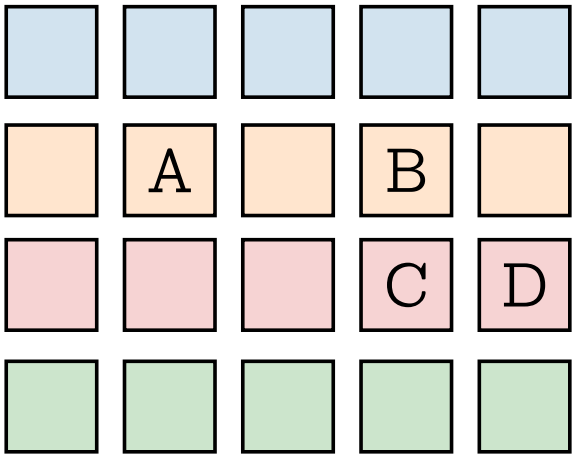

Figure 11: Schematic depiction of the choice of the environment systems in the calculations for the porphine nanotubes where each row represents one ring consisting of 32 porphine molecules. In FDE and FDEu calculations there is only a single active system. All subsystems of the same ring and both adjacent rings are chosen as environment systems, *i.e.* for **A** being the active system, all blue, yellow and red systems would be environment systems. In FDEc calculations there are two active systems. All subsystems of the same ring of both subsystems and in turn all adjacent rings are chosen as environment systems (having an effect on the couplings via the non-additive kernels). That is, for **B** and **C** being active systems, all depicted subsystems are chosen as environment systems and for **C** and **D** being the active systems, all yellow, red and green subsystems would be chosen as environment systems.

for FDEc couplings, calculated in parallel according to Eq. (42) for TC couplings, or neglected depending on the requested coupling pattern. The H-matrix is then diagonalized and coupled transition moments are evaluated from its eigenvectors and uncoupled transition moments.

4.2.5 Ten-Layer Porphine Nanotube

For the calculation on the ten-layer model, we used a heterogenous mixture of machines equipped with 528 GB of main memory and (i) a dual AMD EPYC 7552 48-core processor (10 machines) or (ii) a Intel(R) Xeon(R) E5-2696 v4 44-core processor (5 machines). The Python program for the H-matrix assembly and diagonalization was run on one of the nodes with a dual AMD EPYC 7552 48-core processor. As can be seen from Fig. 10, we employed 320 Serestipy instances, where

each instance used 3 CPU cores and 20 GB of main memory. Therefore, in total we utilized 960 CPU cores (distributed over 15 machines) for the overall calculation.

As can be seen from Tab. 1, the overall computational time needed for the calculation (including ground-state and excited-state calculations) of the ten-layer model is about 19 hours. The parallel FaT cycles (7.86 hours) and parallel FDEu calculations (9.11 hours) take ca. 41% and 48% of the overall computation time, respectively. The parallel isolated SCF and FDEc calculations only take a minor portion of the overall calculation time. Likewise, the calculation of the TC couplings as well as the H-matrix assembly and diagonalization took a negligible part of the computation time. We note that, in principle, the overall computation time will not increase for larger system sizes provided that additional WCNs are available, as we do not include all subsystems in the environment of each individual FDE calculation (see Fig. 11). In case enough computational resources are available, corresponding to enough WCNs to perform the parallel calculations simultaneously, the overall calculation can be scaled horizontally according to the system size. In case not enough hardware is available to perform all parallel calculations in an acceptable time frame it is still possible

Table 1: Total wall-clock timings (h) for the different steps in the ten-layer porphine nanotube calculation. For more detail about each step see Fig. 10. For the FDEc step, $24 \cdot 24 \cdot 608 = 350,208$ overall excitonic couplings have been calculated. The H-matrix assembly includes the calculation of TC charge couplings of 8,816 subsystem pairs totaling $24 \cdot 24 \cdot 8,816 = 5,078,016$ excitonic couplings.

Calculation	Time, h
Isolated SCFs	1.16
FaT Cycle 1	2.54
FaT Cycle 2	2.65
FaT Cycle 3	2.67
FDEu step	9.11
FDEc step	0.91
H-matrix assembly/diagonalization	0.14
Total	19.18

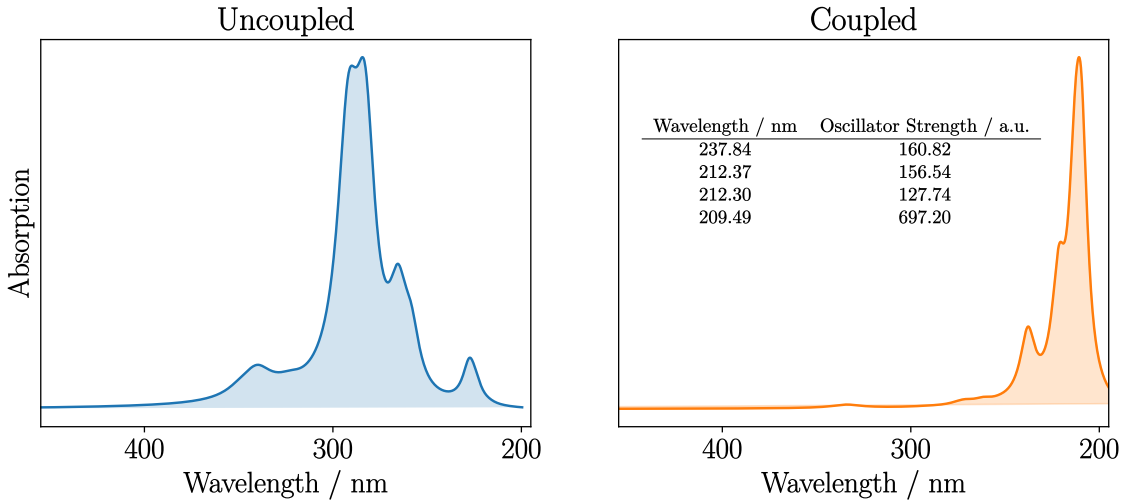

Figure 12: Comparison of the uncoupled absorption spectra of all 320 monomers of the ten-layer porphin nanotube and the coupled one.

to perform batch-wise parallel calculations. This will of course increase the overall computational time, but is still computationally more efficient than serial calculations that can only be scaled vertically. As has been mentioned before, each WCN used only 3 CPU cores and 20 GB of main memory, thus, the parallel calculations can be carried out on rather inexpensive commodity hardware and would still achieve an acceptable time scale.

Next, we consider the coupled absorption spectrum of the porphin nanotubes. In Fig. 12, we compare the sum of uncoupled absorption spectra of all monomers to the coupled absorption spectra. We find a strong blue shift in excitation energies of about 1.6 eV (corresponding to 80 nanometers in this wavelength regime), which holds true for the main band in particular. The blue shift is most likely due to the face-to-face arrangement of identical monomers, which leads to a shift to shorter wavelengths of the allowed one of the coupled excitations. Inspecting the oscillator strengths of the uncoupled spectra, we find that the largest individual oscillator strength belongs to the corresponding excitations labeled as B and C in Fig. 7. Going over to the coupled spectra, however, these couple into very few transitions with very

large oscillator strengths. Those excitations with oscillator strengths larger than 100 a.u. are listed in the table of Fig. 12. This is an important sanity check because it provides evidence that the invididual excitations on all monomers were, in fact, coupled to yield excitations with such high intensity.

5 Conclusion

In this article, we have presented the Serestipy software as an add-on for the quantum-chemistry program Serenity. Serestipy enables parallel large-scale subsystem density-functional theory calculations by providing the interface between Serenity and modern HPC architecture. Serestipy is open-source and available on GitHub [122].

We have introduced two closely-related approximate coupling strategies in the context of subsystem TDDFT. The starting point is a subsystem DFT variant of the simplified TDDFT approach originally devised by Grimme and co-workers [53, 54]. For the second approach, we have reformulated the couplings obtained from simplified subsystem TDDFT in terms of atomic transition charges to fit more seamlessly into the massively-parallel framework. The equivalence of two strategies was shown both analytically in Sec. S1 of the appendix and numerically in Section 4.2.3. The transition-charge couplings were in good agreement with those obtained using subsystem TDDFT for intermediate inter-subsystem distances for which Coulomb contributions dominate the subsystem interaction. For close distances, however, larger deviations occurred most likely due to more pronounced non-classical interactions such as correlation which cannot be described using these approximate couplings.

Testing the Serestipy framework, we have contrasted the serial and parallel FaT approaches for clusters of 30 water molecules. Here, we demonstrated the convergence

advantage that the parallel approach can have. While producing a larger computational overhead on a single compute node, the parallel approach is computationally more efficient in cases where more than one compute node can be used. This opens the possibility to overcome hardware limitations that can hinder efficient serial calculations. In addition, the parallelism provided by the SERESTIPY software allows for more efficient horizontal scaling of QC applications, while running serial calculations can only be enhanced by vertical scaling of the available HPC infrastructure.

As an application of the SERESTIPY framework, we have considered a model system of porphine rings which are arranged in ten layers. This system consists of 12,160 atoms corresponding to 264,960 basis functions for the production-level triple-zeta basis we have employed. The description of the subsystem interaction in subsystem DFT results in a formal linear scaling (for a more detailed discussion of the scaling behavior, see, for instance, Ref. [29]). This allowed us to only include a limited number of environment systems in the FDE and FDE-TDDFT calculations. While these were hand-picked in this study, this task might be readily automatized based on, for example, simple distance criteria. We proposed and used transition charges to approximate subsystem TDDFT couplings for fragments of intermediate distance. We have broken down the overall computation time of less than one day into the parallel FaT ground-state part as well as the excited state parts including FDEu- and FDEc-TDDFT. We have furthermore obtained an absorption spectrum by coupling 24 uncoupled excitations on each subsystem totaling 7,680 overall excitations. This work paves the way to large-scale applications of FDE-TDDFT for excited states of systems of challenging size. One can think of, for example, excitation-energy transfer in chlorosome light-harvesting antennae from green sulfur bacteria [123].

We are convinced that the full potential of subsystem DFT can only be reached by

employing the parallel ansatz. In this respect, we think that SERENITY in combination with SERESTIPY provides a highly valuable tool for vastly extending the applicability of subsystem DFT. Additionally, in combination with subsystem TDDFT gradients it opens the possibility for efficient applications such as investigating excited-state dynamics of very large (bio)molecular systems in the future.

Supporting Information

Equivalence of atomic transition charge couplings and simplified subsystem TDDFT; comparison of couplings presented in the main text to transition-dipole couplings; technical details of the SERESTIPY program.

Acknowledgments

P.E. and J.N. gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through IRTG 2678 Münster–Nagoya (GRK 2678 – 437785492). N.N. and J.N. gratefully acknowledge funding by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) through SFB 1459 (Project A03, Project-ID 433682494).

Data Availability

The data supporting the findings of this study are available either within the Supplementary Material or upon reasonable request from the authors.

References

[1] Krishnan Raghavachari, Arjun Saha. Accurate composite and fragment-based quantum chemical models for large molecules. *Chem. Rev.*, **115**(12) (2015) 5643–5677.

[2] Richard Catlow, Philip Burke, Julia Goodfellow, Dominic Tildesley, Maggie Wilson, Keiji Morokuma. New challenges in quantum chemistry: quests for accurate calculations for large molecular systems. *Philos. Trans. Royal Soc. A*, **360**(1795) (2002) 1149–1164.

[3] Richard A. Friesner, Barry D. Dunietz. Large-Scale ab Initio Quantum Chemical Calculations on Biological Systems. *Acc. Chem. Res.*, **34**(5) (2001) 351–358.

[4] Nityananda Sahu, Shridhar R. Gadre. Molecular Tailoring Approach: A Route for ab Initio Treatment of Large Clusters. *Acc. Chem. Res.*, **47**(9) (2014) 2739–2747.

[5] D. Reed, D. Gannon, J. Dongarra. Reinventing High Performance Computing: Challenges and Opportunities. *arXiv:2203.02544*, 2022.

[6] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, et. al. TensorFlow : Large-scale machine learning on heterogeneous systems. https://www.tensorflow.org/, March 2022.

[7] Edo Liberty, Zohar Karnin, Bing Xiang, Laurence Rouesnel, Baris Coskun, Ramesh Nallapati, Julio Delgado, Amir Sadoughi, Yury Astashonok, Piali Das, Can Balioglu, Saswata Chakravarty, Madhav Jha, Philip Gautier, David Arpin, Tim Januschowski, Valentin Flunkert, Yuyang Wang, Jan Gasthaus, Lorenzo Stella, Syama Rangapuram, David Salinas, Sebastian Schelter, Alex Smola. Elastic machine learning algorithms in amazon sagemaker. In *Proceedings of*

the 2020 ACM SIGMOD International Conference on Management of Data, p. 731–737. Association for Computing Machinery, 2020.

[8] Norman P. Jouppi, Cliff Young, Nishant Patil, David Patterson, Gaurav Agrawal, et. al. In-Datacenter Performance Analysis of a Tensor Processing Unit. *SIGARCH Comput. Archit. News*, **45**(2) (2017) 1–12.

[9] Open compute project. https://www.opencompute.org/, March 2022.

[10] Ramesh Govindan, Ina Minei, Mahesh Kallahalla, Bikash Koley, Amin Vahdat. Evolve or Die: High-Availability Design Principles Drawn from Googles Network Infrastructure. In *Proceedings of the 2016 ACM SIGCOMM Conference*, p. 58–72. Association for Computing Machinery, 2016.

[11] Phillipa Gill, Navendu Jain, Nachiappan Nagappan. Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications. *SIGCOMM Comput. Commun. Rev.*, **41**(4) (2011) 350–361.

[12] Bianca Schroeder, Eduardo Pinheiro, Wolf-Dietrich Weber. DRAM Errors in the Wild: A Large-Scale Field Study. *Commun. ACM*, **54**(2) (2011) 100–107.

[13] Peter H. Hochschild, Paul Turner, Jeffrey C. Mogul, Rama Govindaraju, Parthasarathy Ranganathan, David E. Culler, Amin Vahdat. Cores That Don't Count. In *Proceedings of the Workshop on Hot Topics in Operating Systems*, p. 9–16, New York, NY, USA, 2021. Association for Computing Machinery.

[14] Daniel Firestone, Andrew Putnam, Sambhrama Mundkur, Derek Chiou, Alireza Dabagh, Mike Andrewartha, Hari Angepat, Vivek Bhanu, Adrian Caulfield, Eric Chung, Harish Kumar Chandrappa, Somesh Chaturmohta, Matt Humphrey, Jack Lavier, Norman Lam, Fengfen Liu, Kalin Ovtcharov, Jitu Padhye, Gautham Popuri, Shachar Raindel, Tejas Sapre, Mark Shaw, Gabriel Silva, Madhan

Sivakumar, Nisheeth Srivastava, Anshuman Verma, Qasim Zuhair, Deepak Bansal, Doug Burger, Kushagra Vaid, David A. Maltz, Albert Greenberg. Azure Accelerated Networking: SmartNICs in the Public Cloud. In *15th USENIX Symposium on Networked Systems Design and Implementation (NSDI 18)*. USENIX Association, 2018.

[15] P.A. Morreale, J.M. Anderson. *Software Defined Networking: Design and Deployment (1st ed.)*. CRC Press, 2014.

[16] Expanse system architecture. https://www.sdsc.edu/services/hpc/expanse/, March 2022.

[17] D. Merkel. Docker: Lightweight Linux Containers for Consistent Development and Deployment. *Linux J.*, **2**(239) (2014.

[18] Production-grade container orchestration. https://kubernetes.io/, March 2022.

[19] Tiago Rosado, Jorge Bernardino. An Overview of Openstack Architecture. In *Proceedings of the 18th International Database Engineering & Applications Symposium*, p. 366–367. Association for Computing Machinery, 2014.

[20] Daniel Nurmi, Rich Wolski, Chris Grzegorczyk, Graziano Obertelli, Sunil Soman, Lamia Youseff, Dmitrii Zagorodnov. The Eucalyptus Open-Source Cloud-Computing System. In *2009 9th IEEE/ACM International Symposium on Cluster Computing and the Grid*, p. 124–131, 2009.

[21] D. Sanderson. *Programming Google App Engine with Python: Build and Run Scalable Python Apps on Google's Infrastructure*. O'Reilly Press, 2015.

[22] Gartner & Statista. Public cloud application services/software as a service (saas) end-user spending worldwide from 2015 to 2022 [graph], 2022.

[23] Gartner & Statista. Public cloud services end-user spending worldwide from 2017 to 2023 [graph], 2022.

[24] Gartner & Statista. Market growth forecast for public cloud services worldwide from 2011 to 2022 [graph], 2022.

[25] M. D. Hanwell, C. J. Harris, W. A. de Jong. Open chemistry: RESTful web APIs, JSON, NWChem and the modern web application. *J. Cheminform.*, **55** (2017) 13321.

[26] M. D. Hanwell, C. J. Harris, A. Genova, M. Haghighatlari, M. El Khatib, P. Avery, J. Hachmann, W. A. de Jong. Open Chemistry, JupyterLab, REST, and quantum chemistry. *Int. J. Quantum Chem.*, **121**(1) (2021) e26472.

[27] R. T. Fielding. *REST: Architectural Styles and the Design of Network-based Software Architectures*. Doctoral dissertation, University of California, Irvine, 2000.

[28] E. Apra, E. J. Bylaska, W. A. de Jong, N. Govind, K. Kowalski, T. P. Straatsma, M. Valiev, H. J. J. van Dam, Y. Alexeev, J. Anchell, V. Anisimov, F. W. Aquino, R. Atta-Fynn, J. Autschbach, N. P. Bauman, J. C. Becca, D. E. Bernholdt, K. Bhaskaran-Nair, S. Bogatko, P. Borowski, J. Boschen, J. Brabec, A. Bruner, E. Cauët, Y. Chen, G. N. Chuev, C. J. Cramer, J. Daily, M. J. O. Deegan, T. H. Dunning, M. Dupuis, K. G. Dyall, G. I. Fann, S. A. Fischer, A. Fonari, H. Früchtl, L. Gagliardi, J. Garza, N. Gawande, S. Ghosh, K. Glaesemann, A. W. Götz, J. Hammond, V. Helms, E. D. Hermes, K. Hirao, S. Hirata, M. Jacquelin, L. Jensen, B. G. Johnson, H. Jonsson, R. A. Kendall, M. Klemm, R. Kobayashi, V. Konkov, S. Krishnamoorthy, M. Krishnan, Z. Lin, R. D. Lins, R. J. Littlefield, A. J. Logsdail, K. Lopata, W. Ma, A. V. Marenich, J. Martin del Campo, D. Mejia-Rodriguez, J. E. Moore, J. M. Mullin, T. Nakajima, D. R. Nascimento,

J. A. Nichols, P. J. Nichols, J. Nieplocha, A. Otero-de-la Roza, B. Palmer, A. Panyala, T. Pirojsirikul, B. Peng, R. Peverati, J. Pittner, L. Pollack, R. M. Richard, P. Sadayappan, G. C. Schatz, W. A. Shelton, D. W. Silverstein, D. M. A. Smith, T. A. Soares, D. Song, M. Swart, H. L. Taylor, G. S. Thomas, V. Tipparaju, D. G. Truhlar, K. Tsemekhman, T. Van Voorhis, A. Vazquez-Mayagoitia, P. Verma, O. Villa, A. Vishnu, K. D. Vogiatzis, D. Wang, J. H. Weare, M. J. Williamson, T. L. Windus, K. Woliński, A. T. Wong, Q. Wu, C. Yang, Q. Yu, M. Zacharias, Z. Zhang, Y. Zhao, R. J. Harrison. NWChem: Past, present, and future. *J. Chem. Phys.*, **152**(18) (2020) 184102.

[29] Jan P. Unsleber, Thomas Dresselhaus, Kevin Klahr, David Schnieders, Michael Böckers, Dennis Barton, Johannes Neugebauer. Serenity: A subsystem quantum chemistry program. *J. Comput. Chem.*, **39**(13) (2018) 788–798.

[30] Dennis Barton, Moritz Bensberg, Michael Böckers, Thomas Dresselhaus, Patrick Eschenbach, Lars Hellmann, Kevin Klahr, Anja Massolle, Niklas Niemeyer, Anton Rikus, David Schnieders, Johannes Tölle, Jan Patrick Unsleber, Johannes Neugebauer. qcserenity/serenity: Release 1.4.0, 2021.

[31] Tomasz Adam Wesołowski, Arieh Warshel. Frozen density functional approach for ab initio calculations of solvated molecules. *J. Phys. Chem.*, **97**(30) (1993) 8050–8053.

[32] Mark E. Casida, Tomasz A. Wesołowski. Generalization of the kohn–sham equations with constrained electron density formalism and its time-dependent response theory formulation. *Int. J. Quantum Chem.*, **96**(6) (2004) 577–588.

[33] Tomasz A Wesołowski. Hydrogen-bonding-induced shifts of the excitation energies in nucleic acid bases: an interplay between electrostatic and electron density overlap effects. *J. Am. Chem. Soc.*, **126**(37) (2004) 11444–11445.

[34] Johannes Neugebauer. Couplings between electronic transitions in a subsystem formulation of time-dependent density functional theory. *J. Chem. Phys.*, **126**(13) (2007) 134116.

[35] Johannes Tölle, Johannes Neugebauer. The seamless connection of local and collective excited states in subsystem time-dependent density functional theory. *J. Phys. Chem. Lett.*, **13** (2022) 1003–1018.

[36] M. Iannuzzi, B. Kirchner, J. Hutter. Density functional embedding for molecular systems. *Chem. Phys. Lett*, **421**(1) (2006) 16–20.

[37] K. Kiewisch, C. R. Jacob, L. Visscher. Quantum-Chemical Electron Densities of Proteins and of Selected Protein Sites from Subsystem Density Functional Theory. *J. Chem. Theory Comput.*, **9**(5) (2013) 2425–2440.

[38] T. Hrenar, H.-J. Werner, G. Rauhut. Accurate calculation of anharmonic vibrational frequencies of medium sized molecules using local coupled cluster methods. *J. Chem. Phys.*, **126**(13) (2007) 134108.

[39] G. Rauhut, B. Hartke. Modeling of high-order many-mode terms in the expansion of multidimensional potential energy surfaces: Application to vibrational spectra. *J. Chem. Phys.*, **131**(1) (2009) 014108.

[40] M. Sparta, M. B. Hansen, E. Matito, D. Toffoli, O. Christiansen. Using Electronic Energy Derivative Information in Automated Potential Energy Surface Construction for Vibrational Calculations. *J. Chem. Theory Comput.*, **6**(10) (2010) 3162–3175.

[41] C. König, M. B. Hansen, I. H. Godtliebsen, O. Christiansen. FALCON: A method for flexible adaptation of local coordinates of nuclei. *J. Chem. Phys.*, **144**(7) (2016) 074108.

[42] G. Schmitz, D. G. Artiukhin, O. Christiansen. Approximate High Mode Coupling Potentials using Gaussian Process Regression and Adaptive Density Guided Sampling. *J. Chem. Phys.*, **150**(13) (2019) 131102.

[43] D. G. Artiukhin, E. L. Klinting, C. König, O. Christiansen. Adaptive Density-Guided Approach to Double Incremental Potential Energy Surface Construction. *J. Chem. Phys.*, **152**(19) (2020) 194105.

[44] Markus Reiher, Johannes Neugebauer. A mode-selective quantum chemical method for tracking molecular vibrations applied to functionalized carbon nanotubes. *J. Chem. Phys.*, **118**(4) (2003) 1634–1641.

[45] Thomas Weymuth, Moritz P Haag, Karin Kiewisch, Sandra Luber, Stephan Schenk, Christoph R Jacob, Carmen Herrmann, Johannes Neugebauer, Markus Reiher. Movipac: Vibrational spectroscopy with a robust meta-program for massively parallel standard and inverse calculations. *J. Comput. Chem.*, **33**(27) (2012) 2186–2198.

[46] C. Herrmann, J. Neugebauer, M. Reiher. Finding a needle in a haystack: direct determination of vibrational signatures in complex systems. *New J. Chem.*, **31** (2007) 818–831.

[47] S. Luber, M. Iannuzzi, J. Hutter. Raman spectra from ab initio molecular dynamics and its application to liquid S-methyloxirane. *J. Chem. Phys.*, **141**(9) (2014) 094503.

[48] Ryan M Richard, John M Herbert. A generalized many-body expansion and a unified view of fragment-based methods in electronic structure theory. *J. Chem. Phys.*, **137**(6) (2012) 064113.

[49] Ryan M Richard, John M Herbert. Many-body expansion with overlapping

fragments: Analysis of two approaches. *J. Chem. Theory Comput.*, **9**(3) (2013) 1408–1416.

[50] Daniel Schmitt-Monreal, Christoph R Jacob. Frozen-density embedding-based many-body expansions. *Int. J. Quantum Chem.*, **120**(21) (2020) e26228.

[51] Daniel Schmitt-Monreal, Christoph R Jacob. Density-based many-body expansion as an efficient and accurate quantum-chemical fragmentation method: Application to water clusters. *J. Chem. Theory Comput.*, **17**(7) (2021) 4144–4156.

[52] Linus Scholz, Johannes Neugebauer. Protein response effects on cofactor excitation energies from first principles: Augmenting subsystem time-dependent density-functional theory with many-body expansion techniques. *J. Chem. Theory Comput.*, **17**(10) (2021) 6105–6121.

[53] Stefan Grimme. A simplified tamm-dancoff density functional approach for the electronic excitation spectra of very large molecules. *J. Chem. Phys.*, **138**(24) (2013) 244104.

[54] Christoph Bannwarth, Stefan Grimme. A simplified time-dependent density functional theory approach for electronic ultraviolet and circular dichroism spectra of very large molecules. *Comput. Theor. Chem.*, **1040** (2014) 45–53.

[55] Christoph R. Jacob, Johannes Neugebauer. Subsystem density-functional theory. *Wiley Interdiscip. Rev.: Comput. Mol. Sci.*, **4**(4) (2014) 325–362.

[56] Tomasz A Wesołowski, Sapana Shedge, Xiuwen Zhou. Frozen-density embedding strategy for multilevel simulations of electronic structure. *Chem. Rev.*, **115**(12) (2015) 5891–5928.

[57] P. Hohenberg, W. Kohn. Inhomogeneous electron gas. *Phys. Rev.*, **136** (1964) B864–B871.

[58] R.G. Parr, Y. Weitao. *Density-Functional Theory of Atoms and Molecules*. International Series of Monographs on Chemistry. Oxford University Press, 1994.

[59] Walter Kohn, Lu Jeu Sham. Self-consistent equations including exchange and correlation effects. *Phys. Rev.*, **140**(4A) (1965) A1133.

[60] R. van Leeuwen, E. J. Baerends. Exchange-correlation potential with correct asymptotic behavior. *Phys. Rev. A*, **49** (1994) 2421–2431.

[61] Q. Wu, W. Yang. A direct optimization method for calculating density functionals and exchange–correlation potentials from electron densities. *J. Chem. Phys.*, **118**(6) (2003) 2498–2509.

[62] O. Roncero, M. P. de Lara-Castells, P. Villarreal, F. Flores, J. Ortega, M. Paniagua, A. Aguado. An inversion technique for the calculation of embedding potentials. *J. Chem. Phys.*, **129** (2008) 184104.

[63] Samuel Fux, Christoph R Jacob, Johannes Neugebauer, Lucas Visscher, Markus Reiher. Accurate frozen-density embedding potentials as a first step towards a subsystem description of covalent bonds. *J. Chem. Phys.*, **132**(16) (2010) 164101.

[64] J. D. Goodpaster and N. Ananth and F. R. Manby and T. F. Miller III. Exact nonadditive kinetic potentials for embedded density functional theory. *J. Chem. Phys.*, **133** (2010) 084103.

[65] C. Huang, M. Pavone, E. A. Carter. Quantum mechanical embedding theory based on a unique embedding potential. *J. Chem. Phys.*, **134** (2011) 154110.

[66] Xing Zhang, Emily A. Carter. Kohn-sham potentials from electron densities using a matrix representation within finite atomic orbital basis sets. *J. Chem. Phys.*, **148**(3) (2018) 034105.

[67] Y. Shi, A. Wasserman. Inverse Kohn–Sham Density Functional Theory: Progress and Challenges. *J. Phys. Chem. Lett.*, **12**(22) (2021) 5308–5318.

[68] David Schnieders, Johannes Neugebauer. Accurate embedding through potential reconstruction: A comparison of different strategies. *J. Chem. Phys.*, **149**(5) (2018) 054103.

[69] Frederick R. Manby, Martina Stella, Jason D. Goodpaster, Thomas F. Miller. A simple, exact density-functional-theory embedding scheme. *J. Chem. Theory Comput.*, **8**(8) (2012) 2564–2568.

[70] Bence Hégely, Péter R. Nagy, György G. Ferenczy, Mihály Kállay. Exact density functional and wave function embedding schemes based on orbital localization. *J. Chem. Phys.*, **145**(6) (2016) 064107.

[71] Patrick K. Tamukong, Yuriy G. Khait, Mark R. Hoffmann. Accurate dissociation of chemical bonds using dft-in-dft embedding theory with external orbital orthogonality. *J. Phys. Chem. A*, **121**(1) (2017) 256–264.

[72] Moritz Bensberg, Johannes Neugebauer. Automatic basis-set adaptation in projection-based embedding. *J. Chem. Phys.*, **150**(18) (2019) 184104.

[73] Tomasz Adam Wesołowski, Jacques Weber. Kohn-sham equations with constrained electron density: an iterative evaluation of the ground-state electron density of interacting molecules. *Chem. Phys. Lett.*, **248**(1) (1996) 71–76.

[74] Michael Böckers, Johannes Neugebauer. Excitation energies of embedded open-shell systems: Unrestricted frozen-density-embedding time-dependent

density-functional theory. *J. Chem. Phys.*, **149**(7) (2018) 074102.

[75] Johannes Tölle, Michael Böckers, Niklas Niemeyer, Johannes Neugebauer. Inter-subsystem charge-transfer excitations in exact subsystem time-dependent density-functional theory. *J. Chem. Phys.*, **151**(17) (2019) 174109.

[76] Niklas Niemeyer, Johannes Tölle, Johannes Neugebauer. Approximate versus exact embedding for chiroptical properties: Reconsidering failures in potential and response. *J. Chem. Theory Comput.*, **16**(5) (2020) 3104–3120.

[77] So Hirata, Martin Head-Gordon. Time-dependent density functional theory within the tamm–dancoff approximation. *Chem. Phys. Lett.*, **314**(3-4) (1999) 291–299.

[78] Carolin König, Nicolas Schlüter, Johannes Neugebauer. Direct determination of exciton couplings from subsystem time-dependent density-functional theory within the tamm–dancoff approximation. *J. Chem. Phys.*, **138**(3) (2013) 034104.

[79] Per-Olov Löwdin. On the non-orthogonality problem connected with the use of atomic wave functions in the theory of molecules and crystals. *J. Chem. Phys.*, **18**(3) (1950) 365–375.

[80] Brent P Krueger, Gregory D Scholes, Graham R Fleming. Calculation of couplings and energy-transfer pathways between the pigments of lh2 by the ab initio transition density cube method. *J. Phys. Chem. B*, **102**(27) (1998) 5378–5386.

[81] Kazuhiro J. Fujimoto, Shigehiko Hayashi. Electronic coulombic coupling of excitation-energy transfer in xanthorhodopsin. *J. Am. Chem. Soc.*, **131**(40) (2009) 14152–14153.

[82] Kazuhiro J. Fujimoto. Transition-density-fragment interaction approach for exciton-coupled circular dichroism spectra. *J. Chem. Phys.*, **133**(12) (2010) 124101.

[83] Kazuhiro J. Fujimoto. Transition-density-fragment interaction combined with transfer integral approach for excitation-energy transfer via charge-transfer states. *J. Chem. Phys.*, **137**(3) (2012) 034101.

[84] ME Madjet, A Abdurahman, Th Renger. Intermolecular coulomb couplings from ab initio electrostatic potentials: application to optical transitions of strongly coupled pigments in photosynthetic antennae and reaction centers. *J. Phys. Chem. B*, **110**(34) (2006) 17268–17281.

[85] Kazuhiro J. Fujimoto. Electronic coupling calculations with transition charges, dipoles, and quadrupoles derived from electrostatic potential fitting. *J. Chem. Phys.*, **141**(21) (2014) 214105.

[86] T. Förster, O. Sinanoglu. *Modern Quantum Chemistry*, Volume 3. Academic Press, New York, 1965.

[87] William W Parson. *Modern Optical Spectroscopy*, Volume 2. Springer, 2007.

[88] John P. Perdew, Kieron Burke, Matthias Ernzerhof. Generalized gradient approximation made simple. *Phys. Rev. Lett.*, **77** (1996) 3865–3868.

[89] John P. Perdew, Matthias Ernzerhof, Kieron Burke. Rationale for mixing exact exchange with density functional approximations. *J. Chem. Phys.*, **105**(22) (1996) 9982–9985.

[90] J. P. Perdew, Y. Wang. *Electronic Structure of Solids'91*. Academie, Berlin, 1991.

[91] J. P. Perdew, J. A. Chevary, S. H. Vosko, K. A. Jackson, M. R. Pederson, D. J. Singh, C. Fiolhais. Atoms, molecules, solids, and surfaces: Applications of the generalized gradient approximation for exchange and correlation. *Phys. Rev. B*, **46** (1992) 6671.

[92] A. Lembarki, H. Chermette. Obtaining a gradient-corrected kinetic-energy functional from the perdew-wang exchange functional. *Phys. Rev. A*, **50** (1994) 5328–5331.

[93] Florian Weigend, Reinhart Ahlrichs. Balanced basis sets of split valence, triple zeta valence and quadruple zeta valence quality for h to rn: Design and assessment of accuracy. *Phys. Chem. Chem. Phys.*, **7** (2005) 3297–3305.

[94] Florian Weigend. Accurate coulomb-fitting basis sets for h to rn. *Phys. Chem. Chem. Phys.*, **8**(9) (2006) 1057–1065.

[95] Florian Weigend, Marco Häser, Holger Patzelt, Reinhart Ahlrichs. Ri-mp2: optimized auxiliary basis sets and demonstration of efficiency. *Chem. Phys. Lett.*, **294**(1-3) (1998) 143–152.

[96] R. Chandra, L. Dagum, D. Kohr, R. Menon, D. Maydan, J. McDonald. *Parallel programming in OpenMP*. Morgan Kaufmann, 2001.

[97] Message P Forum. MPI: A Message-Passing Interface Standard. Technical report, 1994.

[98] A. Goez, J. Neugebauer. A Local Variant of the Conductor-Like Screening Model for Fragment-Based Electronic-Structure Methods. *J. Chem. Theory Comput.*, **11**(11) (2015) 5277–5290.

[99] A. Goez, J. Neugebauer. Benchmarking Electron Densities and Electrostatic Potentials of Proteins from the Three-Partition Frozen Density Embedding

Method. *J. Chem. Theory Comput.*, **12**(10) (2016) 4843–4855.

[100] R. E. Blankenship. *Molecular Mechanisms of Photosynthesis*. Blackwell Science, Oxford, 2002.

[101] S.-Y. Park, T. Yokoyama, N. Shibayama, Y. Shiro, J. R. H. Tame. 1.25 Å Resolution Crystal Structures of Human Haemoglobin in the Oxy, Deoxy and Carbonmonoxy Forms. *J. Mol. Biol.*, **360**(3) (2006) 690–701.

[102] S. R. Hubbard, W. A. Hendrickson, D. G. Lambright, S. G. Boxer. X-ray Crystal Structure of a Recombinant Human Myoglobin Mutant at 2.8 Å Resolution. *J. Mol. Biol.*, **213**(2) (1990) 215–218.

[103] I. Ishigami, N. A. Zatsepin, M. Hikita, C. E. Conrad, G. Nelson, J. D. Coe, S. Basu, T. D. Grant, M. H. Seaberg, R. G. Sierra, M. S. Hunter, P. Fromme, R. Fromme, S. Yeh, D. L. Rousseau. Crystal Structure of CO-bound Cytochrome *c* Oxidase Determined by Serial Femtosecond X-ray Crystallography at Room Temperature. *Proc. Natl. Acad. Sci.*, **114**(30) (2017) 8011–8016.

[104] Martin Gouterman. Study of the effects of substitution on the absorption spectra of porphin. *J. Chem. Phys.*, **30**(5) (1959) 1139–1161.

[105] L. Edwards, DH Dolphin, Martin Gouterman, Alan D Adler. Porphyrins xvii. vapor absorption spectra and redox reactions: Tetraphenylporphins and porphin. *J. Mol. Spectrosc*, **38**(1) (1971) 16–32.

[106] Luis Serrano-Andrés, Manuela Merchán, Mercedes Rubio, Björn O Roos. Interpretation of the electronic absorption spectrum of free base porphin by using multiconfigurational second-order perturbation theory. *Chem. Phys. Lett.*, **295**(3) (1998) 195–203.

[107] S. J. A. Van Gisbergen, A Rosa, Giampaolo Ricciardi, E. J. Baerends. Time-dependent density functional calculations on the electronic absorption spectrum of free base porphin. *J. Chem. Phys.*, **111**(6) (1999) 2499–2506.

[108] Dage Sundholm. Interpretation of the electronic absorption spectrum of free-base porphin using time-dependent density-functional theory. *Phys. Chem. Chem. Phys.*, **2**(10) (2000) 2275–2281.

[109] F. Hernández-Fernández, M. Pavanello, L. Visscher. Effect of Metallation, Substituents and Inter/Intra-Molecular Polarization on Electronic Couplings for Hole Transport in Stacked Porphyrin Dyads. *Phys. Chem. Chem. Phys.*, **18** (2016) 21122–21132.

[110] M. O. Senge, M. Fazekas, E. G. A. Notaras, W. J. Blau, M. Zawadzka, O. B. Locos, E. M. Ni Mhuircheartaigh. Nonlinear Optical Properties of Porphyrins. *Adv. Mater.*, **19**(19) (2007) 2737–2774.

[111] M. K. Panda, K. Ladomenou, A. G. Coutsolelos. Porphyrins in Bio-Inspired Transformations: Light-Harvesting to Solar Cell. *Coord. Chem. Rev.*, **256**(21) (2012) 2601–2627.

[112] I. Beletskaya, V. S. Tyurin, A. Y. Tsivadze, R. Guilard, C. Stern. Supramolecular Chemistry of Metalloporphyrins. *Chem. Rev.*, **109**(5) (2009) 1659–1713.

[113] B. A. Gregg, M. A. Fox, A. J. Bard. Porphyrin Octaesters: New Discotic Liquid Crystals. *J. Chem. Soc., Chem. Commun.*, 1987) 1134–1135.

[114] S. Kugimiya, M. Takemura. Novel Liquid Crystals Consisting of Tetraphenylporphyrin Derivatives. *Tetrahedron Lett.*, **31**(22) (1990) 3157–3160.

[115] C. M. Drain, A. Varotto, I. Radivojevic. Self-Organized Porphyrinic Materials. *Chem. Rev.*, **109**(5) (2009) 1630–1658.

[116] J. Miao, L. Zhu. Columnar Liquid Crystalline Assembly of Doubly Discotic Supermolecules Based on Tetra-Triphenylene-Substituted Phthalocyanine. *Soft Matter*, **6** (2010) 2072–2079.

[117] Hisayoshi Iikura, Takao Tsuneda, Takeshi Yanai, Kimihiko Hirao. A long-range correction scheme for generalized-gradient-approximation exchange functionals. *J. Chem. Phys.*, **115**(8) (2001) 3540–3544.

[118] Takeshi Yanai, David P Tew, Nicholas C Handy. A new hybrid exchange–correlation functional using the coulomb-attenuating method (cam-b3lyp). *Chem. Phys. Lett.*, **393**(1-3) (2004) 51–57.

[119] Stefan Grimme, Jens Antony, Stephan Ehrlich, Helge Krieg. A consistent and accurate ab initio parametrization of density functional dispersion correction (dft-d) for the 94 elements h-pu. *J. Chem. Phys.*, **132**(15) (2010) 154104.

[120] Stefan Grimme, Stephan Ehrlich, Lars Goerigk. Effect of the damping function in dispersion corrected density functional theory. *J. Comput. Chem.*, **32**(7) (2011) 1456–1465.

[121] TURBOMOLE V7.5 2020, a development of University of Karlsruhe and Forschungszentrum Karlsruhe GmbH, 1989-2020, TURBOMOLE GmbH, since 2007; available from `http://www.turbomole.com`.

[122] Serestipy, GitHub Repository, `https://github.com/qcserestipy/serestipy`, 2022.

[123] Takatoshi Fujita, Joonsuk Huh, Semion K Saikin, Jennifer C Brookes, Alán Aspuru-Guzik. Theoretical characterization of excitation energy transfer in chlorosome light-harvesting antennae from green sulfur bacteria. *Photosynth. Res.*, **120**(3) (2014) 273–289.