

Original citation:

Zuehlsdorff, T. J., Hine, Nicholas, Payne, M. C. and Haynes, P. D.. (2015) Linear-scaling time-dependent density-functional theory beyond the Tamm-Dancoff approximation : obtaining efficiency and accuracy with in situ optimised local orbitals. The Journal of Chemical Physics, 143 (20). 204107.

Permanent WRAP url:

<http://wrap.warwick.ac.uk/75015>

Copyright and reuse:

The Warwick Research Archive Portal (WRAP) makes this work by researchers of the University of Warwick available open access under the following conditions. Copyright © and all moral rights to the version of the paper presented here belong to the individual author(s) and/or other copyright owners. To the extent reasonable and practicable the material made available in WRAP has been checked for eligibility before being made available.

Copies of full items can be used for personal research or study, educational, or not-for-profit purposes without prior permission or charge. Provided that the authors, title and full bibliographic details are credited, a hyperlink and/or URL is given for the original metadata page and the content is not changed in any way.

Publisher statement:

© (2015) AIP Publishing. This article may be downloaded for personal use only. Any other use requires prior permission of the author and AIP Publishing.

The following article appeared in (citation above) and may be found at

<http://dx.doi.org/10.1063/1.4936280>

A note on versions:

The version presented here may differ from the published version or, version of record, if you wish to cite this item you are advised to consult the publisher's version. Please see the 'permanent WRAP url' above for details on accessing the published version and note that access may require a subscription.

For more information, please contact the WRAP Team at: wrap@warwick.ac.uk

warwick**publications**wrap
highlight your research

<http://go.warwick.ac.uk/lib-publications>

**Linear-scaling time-dependent density-functional theory beyond
the Tamm-Dancoff approximation: obtaining efficiency and
accuracy with *in situ* optimised local orbitals**

T. J. Zuehlsdorff,^{1,*} N. D. M. Hine,² M. C. Payne,¹ and P. D. Haynes^{3,4,5}

¹*Cavendish Laboratory, J. J. Thomson Avenue, Cambridge CB3 0HE, UK*

²*Department of Physics, University of Warwick, Coventry, CV4 7AL*

³*Department of Materials, Imperial College London,
Exhibition Road, London SW7 2AZ, UK*

⁴*Department of Physics, Imperial College London,
Exhibition Road, London SW7 2AZ, UK*

⁵*Thomas Young Centre for Theory and Simulation of Materials,
Imperial College London, Exhibition Road, London SW7 2AZ, UK*

(Dated: October 6, 2015)

Abstract

We present a solution of the full time-dependent density-functional theory (TDDFT) eigenvalue equation in the linear response formalism exhibiting a linear-scaling computational complexity with system size, without relying on the simplifying Tamm-Dancoff approximation (TDA). The implementation relies on representing the occupied and unoccupied subspace with two different sets of *in situ* optimised localised functions, yielding a very compact and efficient representation of the transition density matrix of the excitation with the accuracy associated with a systematic basis set. The TDDFT eigenvalue equation is solved using a preconditioned conjugate gradient algorithm that is very memory-efficient. The algorithm is validated on a small test molecule and a good agreement with results obtained from standard quantum chemistry packages is found, with the preconditioner yielding a significant improvement in convergence rates. The method developed in this work is then used to reproduce experimental results of the absorption spectrum of bacteriochlorophyll in an organic solvent, where it is demonstrated that the TDA fails to reproduce the main features of the low energy spectrum, while the full TDDFT equation yields results in good qualitative agreement with experimental data. Furthermore, the need for explicitly including parts of the solvent into the TDDFT calculations is highlighted, making the treatment of large system sizes necessary that are well within reach of the capabilities of the algorithm introduced here. Finally, the linear-scaling properties of the algorithm are demonstrated by computing the lowest excitation energy of bacteriochlorophyll in solution. The largest systems considered in this work are of the same order of magnitude as a variety of widely studied pigment-protein complexes, opening up the possibility of studying their properties without having to resort to any semiclassical approximations to parts of the protein environment.

*Electronic address: tjz21@cam.ac.uk

I. INTRODUCTION

The study of optical properties of large complex systems is of increasing interest in computational biology, with most efforts being focused on understanding large pigment-protein complexes (PPCs)[1–5]. These systems turn up in a variety of different roles in nature, from biosensors to light-harvesting and linker complexes in photosynthetic bacteria and *ab-initio* computational studies can play a key role in gaining a deeper insight into the mechanisms governing them. However, PPCs are generally characterised by the fact that the protein environment plays an important role in influencing the absorption properties of the pigment, creating the need for large-scale quantum mechanical calculations that are computationally challenging[2–4]. In general, time-dependent density-functional theory (TDDFT)[6], the time-dependent extension to ground-state density-functional theory (DFT)[7, 8], is considered the method of choice when treating this class of systems, mainly due to the good balance between computational cost and achievable accuracy for most common choices of exchange-correlation functionals.

In recent years there have been a number of developments in computational algorithms[9–11] that have helped to make medium-sized systems routinely accessible to TDDFT. However, most common approaches to solving the low energy spectrum of a system using TDDFT show a computational complexity of at least $\mathcal{O}(N^3)$ with system size, imposing an upper limit on the system sizes that can be realistically studied and effectively ruling out a full treatment of the PPCs mentioned above. To treat these large biological systems explicitly in TDDFT, it is necessary to make use of computational approaches that scale linearly with system size.

TDDFT is generally considered in two different flavours. The time domain approach, where the Kohn-Sham equations are propagated explicitly in time[12], and the linear response approach[13], where the excitation energies of the system can be recast as the solutions to an effective eigenvalue equation[14, 15]. The time-domain approach can yield the entire spectrum of the system via a Fourier transform to the frequency domain, is non-perturbative and can thus be applied to problems beyond the linear response regime. However, it comes with the disadvantage that the Kohn-Sham equations have to be propagated sufficiently long to obtain narrow line-widths and dark states cannot be resolved. The direct solution of the TDDFT eigenvalue problem on the other hand can be used to obtain

dark states and triplet transitions that are of interest in some photochemical processes. For this reason, we focus on the linear response flavour of TDDFT for the purpose of this work. In the time-domain TDDFT approach, an $\mathcal{O}(N)$ computational effort with system size can be achieved by extending linear-scaling techniques known from ground-state DFT to the time-dependent Kohn-Sham equations[16–18]. In the linear response approach, algorithms capable of solving for the lowest eigenvalues of the TDDFT eigenvalue equation are also known[19–21], opening up the possibility of a direct computation of excited states in large pigment-protein complexes without relying on additional semi-classical approximations.

The linear response TDDFT equation is a non-Hermitian eigenvalue problem, causing it to be difficult to solve using standard off-the-shelf eigenvalue solvers. A simplifying approximation, known as the Tamm-Dancoff approximation (TDA)[22, 23], recasts the problem into an Hermitian one but its effect on excitation energies and oscillator strengths is not straightforwardly understood. In this work, we introduce a linear-scaling implementation of full TDDFT in the framework of the `ONETEP` code[24], without relying on the TDA, as was required in a previous approach[20]. We test the performance of the algorithm on a number of large scale systems and specifically investigate the quality of the full TDDFT eigenstates to those obtained within the TDA. The largest systems considered explicitly in this work are of similar size as a number of widely studied PPCs (see for example [2, 3]), thus highlighting the capabilities of the algorithm developed here to enable a fully *ab-initio* treatment of this class of systems.

This work is organised as follows: Section II focuses on providing a short overview of the theoretical background necessary for the main results presented in this work, with sections II A and II B introducing the linear response formalism to TDDFT, both in the form of a matrix eigenvalue equation and an effective variational principle. Section II C then provides a short summary of the `ONETEP` code in which the algorithm presented is implemented, as well as an overview over a solution to the Tamm-Dancoff eigenvalue problem[20]. In section III, the linear-scaling solution to the full TDDFT eigenvalue equation is outlined, with a special focus being placed on the appropriate choice of preconditioner (III B) for the conjugate gradient algorithm. The power of the methodology developed here is demonstrated on a small test system by comparison to accurate benchmark results (IV A), before moving on to a realistic system of bacteriochlorophyll *a* in an organic solvent. It is shown that the algorithm developed here scales fully linearly with system size and allows the treatment of

systems inaccessible by conventional approaches.

II. THEORETICAL BACKGROUND

In this section we briefly introduce the theoretical background of linear response TDDFT. We consider a Kohn-Sham system with ground-state density $\rho^{\{0\}}$ and occupied and unoccupied Kohn-Sham states $\{\psi_{v\sigma}^{\text{KS}}\}$ and $\{\psi_{c\sigma}^{\text{KS}}\}$ respectively, where σ denotes a spin index. We limit the discussion to isolated systems, such that the Kohn-Sham eigenstates can be chosen to be real. Furthermore, only semi-local exchange-correlation functionals in the adiabatic approximation will be considered, thus ignoring any long-range and memory effects. While memory effects are routinely ignored in standard TDDFT implementations, long-range interactions can be included in form of hybrid functionals and are known to yield a better description for excitations in infinite systems and charge-transfer states, where semi-local exchange correlation functionals are known to fail[15, 25]. However, for the purpose of this work the main focus is on excitations that retain a localised character and that are thus well described by semi-local functionals.

A. Linear response TDDFT

In linear response TDDFT, the individual excitation energies of the system can be obtained by solving a non-Hermitian eigenvalue equation of the form[14, 15]

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ -\mathbf{B} & -\mathbf{A} \end{pmatrix} \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix} = \omega \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix} \quad (1)$$

Here, \mathbf{A} and \mathbf{B} denote block matrices that can be conveniently expanded in a basis of unoccupied and occupied eigenstates of the ground state Kohn-Sham system:

$$A_{cv\sigma,c'v'\sigma'} = \delta_{\sigma\sigma'}\delta_{cc'}\delta_{vv'}(\epsilon_{c'\sigma'}^{\text{KS}} - \epsilon_{v'\sigma'}^{\text{KS}}) + K_{cv\sigma,c'v'\sigma'} \quad (2)$$

$$B_{cv\sigma,c'v'\sigma'} = K_{cv\sigma,c'v'\sigma'}. \quad (3)$$

where $\{\epsilon_{c\sigma}^{\text{KS}}\}$ and $\{\epsilon_{v\sigma}^{\text{KS}}\}$ denote the eigenvalues associated with the unoccupied and occupied Kohn-Sham states respectively. The eigenvectors are made up of two different components \mathbf{X} and \mathbf{Y} that can be thought of as excitation and de-excitation contributions to the corresponding eigenstates. As can be seen from Eq. 1, the effective non-Hermitian TDDFT

eigenvalue matrix can be characterised by a diagonal part consisting of Kohn-Sham transitions between occupied and unoccupied states and off-diagonal coupling terms described through the coupling matrix \mathbf{K} . The exact form of \mathbf{K} depends upon the exchange-correlation functional used. Here, we will limit our attention to (semi)-local functionals in the adiabatic approximation, in which case the matrix elements of \mathbf{K} can be expressed as

$$K_{cv\sigma,c'v'\sigma'} = \int d^3r d^3r' \psi_{c\sigma}^{\text{KS}}(\mathbf{r}) \psi_{v\sigma}^{\text{KS}}(\mathbf{r}) \times \left[\frac{1}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta^2 E_{\text{xc}}}{\delta\rho_{\sigma}(\mathbf{r})\delta\rho_{\sigma'}(\mathbf{r}')}\bigg|_{\rho\{0\}} \right] \psi_{c'\sigma'}^{\text{KS}}(\mathbf{r}') \psi_{v'\sigma'}^{\text{KS}}(\mathbf{r}') \quad (4)$$

where E_{xc} is the exchange-correlation energy. In the remaining part of this work, all spin indices will be dropped for convenience.

From the structure of Eq. 1 it can be seen that the TDDFT eigenvalue equation has solutions in the positive and negative frequency domain. These positive and negative frequencies can be interpreted as excitation and de-excitation energies [15]. Note that due to the structure of the equation, the positive and negative eigenvalue solutions are coupled via the block matrix \mathbf{B} . Assuming that the coupling of excitations to the de-excitation part of the full TDDFT spectrum is small, one can set the coupling matrices \mathbf{B} to zero, which causes a complete decoupling of the excitation and de-excitation part of the spectrum. The positive excitation energies can then be solved for via the Hermitian eigenvalue equation

$$\mathbf{A}\mathbf{X} = \omega\mathbf{X}. \quad (5)$$

Solving Eq. 5 instead of Eq. 1 is referred to as the Tamm-Dancoff approximation (TDA)[22, 23]. While the TDA is often reported to yield reliable excitation energies in many situations[23], the structure of the equation violates both time-reversal symmetry and important sum rules related to the oscillator strengths[26] of the excitations. Furthermore, there are known cases where the TDA yields to significant errors in excitation energies[27], making a treatment of the full eigenvalue problem desirable. However, as will be discussed in more detail in the next section, the main disadvantage of a full treatment of the TDDFT eigenvalue equation originates from the fact that Eq. 1 constitutes a non-Hermitian eigenvalue problem, meaning that a variety of standard numerical methods for computing the eigenvalues of large matrices cannot be straightforwardly applied to it.

B. Iterative solutions to the TDDFT equation

The dimensions of the TDDFT eigenvalue equation grow as $\mathcal{O}(N^2)$ with system size, making a direct diagonalisation of the matrices in Eq. 1 or Eq. 5 undesirable for larger systems. Furthermore, one is often interested in a relatively small number of low energy excited states in the visible or ultraviolet energy range of a system, such that the computation of high energy excited states is unnecessary. The best approach to tackle the TDDFT eigenvalue problem for real systems of interest is thus to use iterative methods in order to calculate the lowest few excited states.

Within the TDA, such an iterative scheme is straightforwardly defined, since the Hermitian properties of the block matrix \mathbf{A} allow for the definition of the lowest excitation of the system via a variational principle:

$$\omega_{\min}^{\text{TDA}} = \min_{\mathbf{X}} \Omega_{\text{TDA}}(\mathbf{X}) = \min_{\mathbf{X}} \frac{\mathbf{X}^\dagger \mathbf{A} \mathbf{X}}{\mathbf{X}^\dagger \mathbf{X}}. \quad (6)$$

From this definition, the gradient of the Rayleigh-Ritz functional $\Omega_{\text{TDA}}(\mathbf{X})$ with respect to \mathbf{X} can then be straightforwardly computed

$$\frac{\partial \Omega_{\text{TDA}}(\mathbf{X})}{\partial \mathbf{X}} = \mathbf{g}_{\text{TDA}} = \frac{2}{\mathbf{X}^\dagger \mathbf{X}} \left[\mathbf{A} \mathbf{X} - \frac{\mathbf{X}^\dagger \mathbf{A} \mathbf{X}}{\mathbf{X}^\dagger \mathbf{X}} \mathbf{X} \right]. \quad (7)$$

This gradient can be used as a steepest-descent search direction in a conjugate gradient algorithm to optimise a random starting vector $\mathbf{X}_{\text{guess}}$ until the lowest eigenstate $\omega_{\min}^{\text{TDA}}$ has been obtained. Note that in order to compute the gradient in Eq. 7, it is only required to evaluate the matrix-vector product $\mathbf{A} \mathbf{X}$. This means that \mathbf{A} does not have to be explicitly calculated or stored during the calculation, which is prohibitive for large systems.

The definition of the lowest eigenvalue of the Tamm-Dancoff TDDFT equation via the variational principle relies on the fact that it constitutes an Hermitian eigenvalue problem and is thus not generally possible for the full TDDFT equation. It was however pointed out by Thouless[28] that since the blocks \mathbf{A} and \mathbf{B} of the full TDDFT equation are Hermitian, a variational principle can be formulated for its lowest *positive* eigenstate via:

$$\begin{aligned} \omega_{\min} &= \min_{(\mathbf{X}, \mathbf{Y})} \Omega_{\text{Thou}}(\mathbf{X}, \mathbf{Y}) \\ &= \min_{(\mathbf{X}, \mathbf{Y})} \frac{\begin{pmatrix} \mathbf{X}^\dagger & \mathbf{Y}^\dagger \end{pmatrix} \begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{A} \end{pmatrix} \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix}}{|\mathbf{X}^\dagger \mathbf{X} - \mathbf{Y}^\dagger \mathbf{Y}|}. \end{aligned} \quad (8)$$

The numerator of the above expression is guaranteed to be positive semi-definite for (semi)-local exchange-functionals, while the denominator is forced to be positive by taking the absolute value.

While it is possible to use the above variational principle to directly obtain the lowest excited state of the full TDDFT equation, the Thouless functional is not the most ideal formulation of the problem in the context of semi-local exchange-correlation functionals (see Sec. III A). A more computationally efficient reformulation was introduced by Tsiper [29] by noticing the equivalence of the full TDDFT eigenvalue problem to that of a set of classical harmonic oscillators. As a first step, one introduces the effective vectors $\mathbf{p} = \mathbf{X} - \mathbf{Y}$ and $\mathbf{q} = \mathbf{X} + \mathbf{Y}$. It can then be easily shown that the Thouless functional can be rewritten as:

$$\begin{aligned} \omega_{\min} &= \min_{(\mathbf{p}, \mathbf{q})} \Omega_{\text{Tsiper}}(\mathbf{p}, \mathbf{q}) \\ &= \min_{(\mathbf{p}, \mathbf{q})} \frac{\begin{pmatrix} \mathbf{p}^\dagger & \mathbf{q}^\dagger \end{pmatrix} \begin{pmatrix} \mathbf{A} - \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \mathbf{A} + \mathbf{B} \end{pmatrix} \begin{pmatrix} \mathbf{p} \\ \mathbf{q} \end{pmatrix}}{|\mathbf{p}^\dagger \mathbf{q} + \mathbf{q}^\dagger \mathbf{p}|}. \end{aligned} \quad (9)$$

Here, the variational principle is again only defined for the lowest *positive* excitation of the system. Just like in the Thouless functional, the numerator of the functional is guaranteed to be positive-semidefinite. The denominator has to be forced to stay positive by taking the absolute value, since $\mathbf{p}^\dagger \mathbf{q} + \mathbf{q}^\dagger \mathbf{p}$ is not guaranteed to be positive-semidefinite.

While the variational principles in Eqns. 6 and 9 are written in terms of the lowest excitation of the system only, the concept can be straightforwardly extended to higher excitations. In order to converge the second-lowest excitation ω_2 of the system, the same effective functional of Eqn. 9 can be used in full TDDFT, with the additional constraint that \mathbf{p}_2 and \mathbf{q}_2 , the trial vectors associated with ω_2 obey an effective orthogonality constraint of the form

$$\left| \mathbf{p}_1^\dagger \mathbf{q}_2 + \mathbf{q}_1^\dagger \mathbf{p}_2 \right| = 0 \quad (10)$$

where \mathbf{p}_1 and \mathbf{q}_1 are taken to denote the vectors associated with ω_{\min} , the lowest excitation of the system. The principle can be extended to an arbitrary number of excited states to be converged, where the vectors \mathbf{p}_i and \mathbf{q}_i associated with the i^{th} excitation are constrained to be orthogonal to the vectors of all other excitations via Eqn. 10.

C. Tamm-Dancoff TDDFT in the ONETEP code

The linear-scaling solution to the full TDDFT equation developed in this work (see Sec. III) is implemented in the ONETEP code[24]. As with other linear-scaling DFT approaches, any reference to individual Kohn-Sham states $\{\psi_i^{\text{KS}}\}$ is given up in favour of a collective representation in the form of the single particle density matrix $\rho^{\{v\}}(\mathbf{r}, \mathbf{r}')$ such that

$$\begin{aligned}\rho^{\{v\}}(\mathbf{r}, \mathbf{r}') &= \sum_v^{N_{\text{occ}}} \psi_v^{\text{KS}}(\mathbf{r})\psi_v^{\text{KS}}(\mathbf{r}') \\ &= \sum_{\alpha\beta} \phi_\alpha(\mathbf{r})P^{\{v\}\alpha\beta}\phi_\beta(\mathbf{r}')\end{aligned}\quad (11)$$

where $\{\phi_\alpha(\mathbf{r})\}$ denotes a set of *in situ* optimised[24] localised atom-centered support functions referred to as non-orthogonal generalised Wannier functions (NGWFS)[33] and $\mathbf{P}^{\{v\}}$ is the single particle density matrix in the representation of those NGWFs. Linear scaling of computational cost with system size is then obtained by exploiting the fact that the ground-state density matrix decays exponentially for any system with a band gap, causing $\mathbf{P}^{\{v\}}$ to be sparse for sufficiently large system size[34, 35].

The *in situ* optimisation of the support functions $\{\phi_\alpha\}$ means that only a minimal number of functions is required to accurately span the occupied subspace, but the unoccupied subspace is generally badly represented[30]. This issue is overcome by optimising a second set of NGWFs $\{\chi_\alpha\}$ for a low energy subset of the unoccupied subspace that is represented by the effective density matrix $\mathbf{P}^{\{c\}}$ [30]. It has been demonstrated[20] that the compact sets of support functions $\{\phi_\alpha\}$ and $\{\chi_\beta\}$ provide a very good representation for low energy excited states as calculated in the TDA. Defining an effective response density $\rho^{\{1\}}(\mathbf{r})$ associated with a TDA eigenvector \mathbf{X} such that

$$\begin{aligned}\rho^{\{1\}}(\mathbf{r}) &= \sum_{c,v} \psi_c^{\text{KS}}(\mathbf{r})\mathbf{X}_{cv}\psi_v^{\text{KS}}(\mathbf{r}) \\ &= \sum_{\alpha\beta} \chi_\alpha(\mathbf{r})P^{\{1\}\alpha\beta}\phi_\beta(\mathbf{r})\end{aligned}\quad (12)$$

it becomes clear that the effective response density matrix $\mathbf{P}^{\{1\}}$ is the representation of \mathbf{X} in mixed unoccupied-occupied NGWF space[20]. The matrix-vector product $\mathbf{f} = \mathbf{A}\mathbf{X}$ can then be directly constructed in NGWF space as

$$\begin{aligned}\mathbf{f}_{\text{TDA}}^{\chi\phi} &= \mathbf{P}^{\{c\}}\mathbf{H}\mathbf{X}\mathbf{P}^{\{1\}} - \mathbf{P}^{\{1\}}\mathbf{H}\phi\mathbf{P}^{\{v\}} \\ &\quad + \mathbf{P}^{\{c\}}\left(\mathbf{V}_{\text{SCF}}^{\{1\}\chi\phi}\left[\mathbf{P}^{\{1\}}\right]\right)\mathbf{P}^{\{v\}}.\end{aligned}\quad (13)$$

Here, \mathbf{H}^χ and \mathbf{H}^ϕ denote the ground state Kohn-Sham Hamiltonian in the $\{\chi_\alpha\}$ and $\{\phi_\beta\}$ representation respectively. $\mathbf{V}_{\text{SCF}}^{\{1\}\chi\phi}[\mathbf{P}^{\{1\}}]$ is the self-consistent field response of the system due to a perturbation $\rho^{\{1\}}(\mathbf{r})$ in the ground state density[20] and is the result of \mathbf{X} acting on the coupling matrix \mathbf{K} in mixed unoccupied-occupied NGWF space. Note that $\mathbf{f}_{\text{TDA}}^{\chi\phi}$ represents a contravariant tensor quantity and has to be multiplied by \mathbf{S}^χ and \mathbf{S}^ϕ from the left and right respectively to obtain a covariant quantity (see [20] for further details). The lowest excitation energy of the system can then be written as

$$\omega_{\min}^{\text{TDA}} = \min_{\mathbf{P}^{\{1\}}} \Omega_{\text{TDA}}[\mathbf{P}^{\{1\}}] = \frac{\text{Tr}[\mathbf{P}^{\{1\}\dagger} \mathbf{S}^\chi \mathbf{f}^{\chi\phi} \mathbf{S}^\phi]}{\text{Tr}[\mathbf{P}^{\{1\}\dagger} \mathbf{S}^\chi \mathbf{P}^{\{1\}} \mathbf{S}^\phi]} \quad (14)$$

where \mathbf{S}^χ and \mathbf{S}^ϕ denote the overlap matrices of the $\{\chi_\alpha\}$ and $\{\phi_\beta\}$ NGWF representation respectively. Higher excited states can be obtained from the same variational principle by enforcing an orthogonality constraint between all excited states. If all involved density matrices $\mathbf{P}^{\{1\}}$, $\mathbf{P}^{\{v\}}$ and $\mathbf{P}^{\{c\}}$ can be treated as sparse for sufficiently large system size[52], evaluating Eqs. 13 and 14 scales as $\mathcal{O}(N)$ with system size and $\omega_{\min}^{\text{TDA}}$ can be computed in linear-scaling effort using standard iterative approaches.

The above formulation yields accurate excitation energies if \mathbf{X} is well-represented by the low energy subset of unoccupied states for which $\mathbf{P}^{\{c\}}$ is optimised. However, in many scenarios it is desirable to include higher energy unoccupied states in an approximate manner to achieve convergence[20, 31]. One straightforward way of doing so is to introduce the joint unoccupied-occupied representation $\{\varphi_\alpha\} = \{\phi_\beta\} \oplus \{\chi_\gamma\}$ and to redefine $\mathbf{P}^{\{c\}}$ via a projector onto the entire unoccupied subspace representable by $\{\varphi_\alpha\}$ [31]:

$$\mathbf{P}^{\{c\}} = (\mathbf{S}^\varphi)^{-1} - (\mathbf{S}^\varphi)^{-1} \mathbf{S}^{\varphi\phi} \mathbf{P}^{\{v\}} \mathbf{S}^{\phi\varphi} (\mathbf{S}^\varphi)^{-1}. \quad (15)$$

Here, the elements of $\mathbf{S}^{\varphi\phi}$ are given by $S_{\alpha\beta}^{\varphi\phi} = \langle \varphi_\alpha | \phi_\beta \rangle$. While using the joint representation for the unoccupied subspace instead of $\{\chi_\alpha\}$ roughly doubles the computational cost compared to using $\{\chi_\alpha\}$, it yields consistently good TDDFT excitation energies[31, 32] and will be used throughout in Sec. IV. However, for the purpose of clarity in outlining the linear-scaling TDDFT algorithm in Sec. III, we shall use $\{\chi_\alpha\}$ to label the unoccupied space, noting that, if desired, the representation can be replaced by $\{\varphi_\beta\}$ and the projector of Eq. 15.

III. FULL TDDFT IN ONETEP

In this section, we will outline a conjugate gradient algorithm to compute the lowest N_ω excited states of the full TDDFT equation. While other iterative eigensolvers like the Lanczos and Davidson algorithms and multishift methods have been applied to this problem, both in the framework of standard cubic scaling[9, 11] and $\mathcal{O}(N)$ [19] approaches, the conjugate gradient method is chosen here for both its good performance in the TDA[20] and its low memory requirements suitable for large scale applications. A special focus will be put on an effective preconditioning scheme, as well as the linear-scaling properties of the algorithm. As in the previous section, the discussion is limited to semi-local exchange-correlation functionals only.

A. The Tsiper functional in mixed $\{\chi_\alpha\}$ - $\{\phi_\beta\}$ NGWF space

The key to obtaining a linear-scaling implementation of the full TDDFT equation is to rewrite the Tsiper functional in the same compact, localised NGWF representation that has been used to rewrite the Tamm-Dancoff functional in Eq. 14. For this purpose, we define $\underline{\mathbf{f}}_{\text{Tsip}}$ as

$$\underline{\mathbf{f}}_{\text{Tsip}} = \begin{pmatrix} \mathbf{f}_p \\ \mathbf{f}_q \end{pmatrix} = \begin{pmatrix} \mathbf{A} - \mathbf{B} & \mathbf{0} \\ \mathbf{0} & \mathbf{A} + \mathbf{B} \end{pmatrix} \begin{pmatrix} \mathbf{p} \\ \mathbf{q} \end{pmatrix} = \begin{pmatrix} \mathbf{A}\mathbf{p} - \mathbf{B}\mathbf{p} \\ \mathbf{A}\mathbf{q} + \mathbf{B}\mathbf{q} \end{pmatrix} \quad (16)$$

Following analogous steps to the derivation of the linear-scaling solution to the TDA eigenvalue equation[20], we define the effective response density matrices $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ and response densities $\rho^{\{p\}}(\mathbf{r})$ and $\rho^{\{q\}}(\mathbf{r})$ such that

$$\begin{aligned} \rho^{\{p\}}(\mathbf{r}) &= \sum_{c,v} \psi_c^{\text{KS}}(\mathbf{r}) p_{cv} \psi_v^{\text{KS}}(\mathbf{r}) \\ &= \sum_{\alpha\beta} \chi_\alpha(\mathbf{r}) P^{\{p\}\alpha\beta} \phi_\beta(\mathbf{r}) \end{aligned} \quad (17)$$

with $\rho^{\{q\}}(\mathbf{r})$ following an analogous definition. Thus $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ are the matrices \mathbf{p} and \mathbf{q} in $\{\chi_\alpha\}$ - $\{\phi_\beta\}$ NGWF representation. Just like $\mathbf{P}^{\{1\}}$ [20], $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ have to follow an effective invariance constraint in order to be valid response density matrices, which originates from the orthogonality between the unoccupied and occupied Kohn-Sham spaces in which \mathbf{p} and \mathbf{q} are represented. For $\mathbf{P}^{\{p\}}$, the invariance constraint can be written as

$$\mathbf{P}^{\{p\}'} = \mathbf{P}^{\{c\}} \mathbf{S} \mathbf{X} \mathbf{P}^{\{p\}} \mathbf{S}^\phi \mathbf{P}^{\{v\}} = \mathbf{P}^{\{p\}}, \quad (18)$$

with an identical statement for $\mathbf{P}^{\{q\}}$.

From Sec. IIC, the action of \mathbf{A} acting on some vector \mathbf{X} written in $\{\chi_\alpha\}$ - $\{\phi_\beta\}$ NGWF space is already known. Using Eq. 13, it is straightforward to rewrite $\underline{\mathbf{f}}_{\text{Tsiper}}$ in $\{\chi_\alpha\}$ and $\{\phi_\beta\}$ representation such that

$$\begin{aligned} \underline{\mathbf{f}}_{\text{Tsiper}}^{\chi\phi} &= \begin{pmatrix} \mathbf{f}_{\{p\}}^{\chi\phi} \\ \mathbf{f}_{\{q\}}^{\chi\phi} \end{pmatrix} = \begin{pmatrix} \mathbf{P}^{\{c\}} \mathbf{H} \chi \mathbf{P}^{\{p\}} - \mathbf{P}^{\{p\}} \mathbf{H} \phi \mathbf{P}^{\{v\}} \\ \mathbf{P}^{\{c\}} \mathbf{H} \chi \mathbf{P}^{\{q\}} - \mathbf{P}^{\{q\}} \mathbf{H} \phi \mathbf{P}^{\{v\}} \end{pmatrix} \\ &+ \begin{pmatrix} \mathbf{0} \\ 2\mathbf{P}^{\{c\}} \left(\mathbf{V}_{\text{SCF}}^{\{1\}\chi\phi} \left[\mathbf{P}^{\{q\}} \right] \right) \mathbf{P}^{\{v\}} \end{pmatrix}. \end{aligned} \quad (19)$$

The advantage of the reformulation of the Tsiper functional in terms of \mathbf{p} and \mathbf{q} now becomes apparent. In order to evaluate Eq. 19 for any semi-local exchange-correlation functional, it is sufficient to evaluate $\mathbf{V}_{\text{SCF}}^{\{1\}\chi\phi}$ only once. Since calculating $\mathbf{V}_{\text{SCF}}^{\{1\}\chi\phi}$ is generally the most expensive part of applying the TDDFT operator, it follows that computing Eq. 19 is not significantly more expensive than evaluating Eq. 13, suggesting that a full solution of to the TDDFT equations can be of similar computational complexity to the solution to the TDA[36].

Using Eq. 19, the lowest excitation energy of the system as specified by the Tsiper functional (9) can then be rewritten in $\{\chi_\alpha\}$ and $\{\phi_\beta\}$ representation

$$\begin{aligned} \omega_{\min} &= \min_{\{\mathbf{P}^{\{p\}}, \mathbf{P}^{\{q\}}\}} \left\{ \frac{\text{Tr} \left[\mathbf{P}^{\{p\}\dagger} \mathbf{S} \chi \mathbf{f}_{\{p\}}^{\chi\phi} \mathbf{S}^\phi \right]}{2 \left| \text{Tr} \left[\mathbf{P}^{\{p\}\dagger} \mathbf{S} \chi \mathbf{P}^{\{q\}} \mathbf{S}^\phi \right] \right|} \right. \\ &\quad \left. + \frac{\text{Tr} \left[\mathbf{P}^{\{q\}\dagger} \mathbf{S} \chi \mathbf{f}_{\{q\}}^{\chi\phi} \mathbf{S}^\phi \right]}{2 \left| \text{Tr} \left[\mathbf{P}^{\{p\}\dagger} \mathbf{S} \chi \mathbf{P}^{\{q\}} \mathbf{S}^\phi \right] \right|} \right\} \end{aligned} \quad (20)$$

where the minimisation is carried out under the normalisation constraint

$$\text{Tr} \left[\mathbf{P}^{\{p\}\dagger} \mathbf{S} \chi \mathbf{P}^{\{q\}} \mathbf{S}^\phi \right] = 1. \quad (21)$$

Specifying the normalisation constraint allows us to drop the absolute value from the denominator of the Tsiper functional when computing the gradient of Eq. 20 with respect to changes in $\left(\mathbf{P}^{\{p\}} \ \mathbf{P}^{\{q\}} \right)$. Using Eq. 19 the contravariant gradient of the Tsiper functional

can be, in close analogy to the gradient in the TDA (7), written as

$$\begin{aligned} \underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} &= \begin{pmatrix} \mathbf{g}_{\{p\}}^{\chi\phi} \\ \mathbf{g}_{\{q\}}^{\chi\phi} \end{pmatrix} = \begin{pmatrix} \mathbf{f}_{\{p\}}^{\chi\phi} \\ \mathbf{f}_{\{q\}}^{\chi\phi} \end{pmatrix} - \left\{ \text{Tr} \left[\mathbf{P}^{\{p\}\dagger} \mathbf{S}^{\chi} \mathbf{f}_{\{p\}}^{\chi\phi} \mathbf{S}^{\phi} \right] \right. \\ &\quad \left. + \text{Tr} \left[\mathbf{P}^{\{q\}\dagger} \mathbf{S}^{\chi} \mathbf{f}_{\{q\}}^{\chi\phi} \mathbf{S}^{\phi} \right] \right\} \begin{pmatrix} \mathbf{P}^{\{q\}} \\ \mathbf{P}^{\{p\}} \end{pmatrix}. \end{aligned} \quad (22)$$

This result exploits the fact that $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ follow the normalisation constraint. The above gradient can be used as a steepest-descent search direction in a conjugate gradient algorithm[37]. Note however, that since the TDDFT operator in the TDA can be computed in linear-scaling effort[20], the evaluation of both $\underline{\mathbf{f}}_{\text{Tsip}}^{\chi\phi}$ and $\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi}$ also scales fully linearly with system size, as long as all involved density matrices $\mathbf{P}^{\{c\}}$, $\mathbf{P}^{\{v\}}$, $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ can be treated as sparse for sufficiently large system sizes. Furthermore, it is worth pointing out that for any matrix pair $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ obeying the invariance constraint of Eq. 18, the gradient $\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi}$ follows the invariance constraint by construction. This condition is vital as it means that any pair of trial response density matrices updated with a search direction derived from $\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi}$ will also obey the appropriate invariance constraint by construction.

It should be noted, that once a truncation of the density matrices $\mathbf{P}^{\{q\}}$ and $\mathbf{P}^{\{p\}}$ is introduced, the invariance constraint of Eqn. 18 can only hold approximately. While it is possible to iteratively apply the projection at the end of each conjugate gradient step until some measure of the violation of the invariance constraint is kept below a certain threshold[20, 32], in practice this can lead to convergence problems as it destroys the variational nature of the algorithm presented here. This problem is overcome by introducing a set of auxiliary density matrices $\mathbf{L}^{\{q\}}$ and $\mathbf{L}^{\{p\}}$ [32] in the spirit of ground state linear-scaling approaches[38]. The auxiliary matrices can be arbitrarily truncated and are used to define the real density matrices $\mathbf{P}^{\{q\}}$ and $\mathbf{P}^{\{p\}}$ used in the algorithm via

$$\mathbf{P}^{\{p\}} = \mathbf{P}^{\{c\}} \mathbf{S}^{\chi} \mathbf{L}^{\{p\}} \mathbf{S}^{\phi} \mathbf{P}^{\{v\}} \quad (23)$$

at every step of the calculation. While this scheme comes at a computational cost as $\mathbf{P}^{\{p\}}$ is less sparse than $\mathbf{L}^{\{p\}}$, it guarantees that every point in the algorithm, $\mathbf{P}^{\{q\}}$ and $\mathbf{P}^{\{p\}}$ fulfil their respective invariance constraints to the degree that $\mathbf{P}^{\{v\}}$ and $\mathbf{P}^{\{c\}}$ fulfil their respective idempotency constraint[32]. Since linear-scaling DFT calculations employ a number of techniques that ensure that for sensible truncation schemes, $\mathbf{P}^{\{v\}}$ and $\mathbf{P}^{\{c\}}$ retain near-

idempotency[38], the above scheme yields a robust convergence of TDDFT calculations when truncations are applied to the respective response density matrices.

B. Preconditioning

The TDDFT eigenvalue problem generally has a large condition number associated with it, causing iterative eigensolvers to show a relatively slow convergence. This is most easily appreciated by considering that the elements of the coupling matrix \mathbf{K} are generally small compared to the diagonal elements of Kohn-Sham eigenvalue differences and the condition number of the full TDDFT matrix is reasonably well approximated by the condition number of $\begin{pmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D} \end{pmatrix}$, where the elements of the block matrix \mathbf{D} are given by

$$D_{cv,c'v'} = \delta_{cc'}\delta_{vv'}(\epsilon_{c'}^{\text{KS}} - \epsilon_{v'}^{\text{KS}}). \quad (24)$$

Clearly, \mathbf{D} has a condition number that is much larger than 1, resulting in relatively slow convergence of iterative eigensolvers.

For these reasons, it has long been appreciated that $\begin{pmatrix} \mathbf{D} & \mathbf{0} \\ \mathbf{0} & \mathbf{D} \end{pmatrix}$ should form an efficient preconditioner for the full TDDFT eigenvalue problem. However, applying the preconditioner requires the computation of \mathbf{D}^{-1} , which can only easily be constructed in a Kohn-Sham eigenstate representation where \mathbf{D} is diagonal. In linear-scaling TDDFT, a diagonal representation of \mathbf{D} is not available and the explicit construction of \mathbf{D}^{-1} via matrix inversion is undesirable.

In order to obtain a linear-scaling preconditioner, we consider the Tsiper functional of Eq. 9 when preconditioned from the left:

$$\Omega_{\text{Tsip}}(\mathbf{p}, \mathbf{q}) = \frac{\begin{pmatrix} \mathbf{p}^\dagger & \mathbf{q}^\dagger \end{pmatrix} \begin{pmatrix} \mathbf{D}^{-1}(\mathbf{A} - \mathbf{B}) & \mathbf{0} \\ \mathbf{0} & \mathbf{D}^{-1}(\mathbf{A} + \mathbf{B}) \end{pmatrix} \begin{pmatrix} \mathbf{p} \\ \mathbf{q} \end{pmatrix}}{\left| \begin{pmatrix} \mathbf{p}^\dagger & \mathbf{q}^\dagger \end{pmatrix} \begin{pmatrix} \mathbf{0} & \mathbf{D}^{-1} \\ \mathbf{D}^{-1} & \mathbf{0} \end{pmatrix} \begin{pmatrix} \mathbf{p} \\ \mathbf{q} \end{pmatrix} \right|} \quad (25)$$

While the action of \mathbf{D}^{-1} on a matrix cannot be straightforwardly constructed in $\{\chi_\alpha\}$ -

$\{\phi_\beta\}$ NGWF space, the action of \mathbf{D} is trivially known. Denoting $\underline{\mathbf{G}}_{\text{Tsip}}^{\chi\phi} = \begin{pmatrix} \mathbf{G}_{\{p\}}^{\chi\phi} \\ \mathbf{G}_{\{q\}}^{\chi\phi} \end{pmatrix}$ as the preconditioned version of the gradient $\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi}$ of Eq. 22, it can be seen that applying the preconditioner \mathbf{D}^{-1} to the Tsiper functional is equivalent to solving the linear system

$$\begin{pmatrix} \mathbf{P}^{\{c\}} \mathbf{H} \chi \mathbf{G}_{\{p\}}^{\chi\phi} - \mathbf{G}_{\{p\}}^{\chi\phi} \mathbf{H} \phi \mathbf{P}^{\{v\}} \\ \mathbf{P}^{\{c\}} \mathbf{H} \chi \mathbf{G}_{\{q\}}^{\chi\phi} - \mathbf{G}_{\{q\}}^{\chi\phi} \mathbf{H} \phi \mathbf{P}^{\{v\}} \end{pmatrix} = \underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} \quad (26)$$

for $\underline{\mathbf{G}}_{\text{Tsip}}^{\chi\phi}$. This linear system can be solved iteratively to a chosen degree of numerical accuracy using a standard conjugate gradient algorithm (see algorithm 2 detailed in [37]). While applying the preconditioner in NGWF space is therefore not as straightforward as in Kohn-Sham space, it scales fully linearly with system size and only requires a number of comparatively inexpensive matrix-matrix multiplications. If the computational overhead from solving the linear system in every step of the conjugate gradient calculation is significantly less than the time saved by constructing $\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi}$ fewer times due to a faster convergence rate, the preconditioning proposed here becomes highly efficient. This point will be addressed in more detail in Sec. IV A.

C. Optimising multiple excited states

In most situations, we are not interested in converging only the lowest excited state of the full TDDFT equation, but rather the subspace spanning the lowest N_ω excitations $\{\omega_i\}$. Following the considerations in Sec. II B regarding the convergence of higher excited states, we introduce a set of N_ω TDDFT trial vectors $\{\underline{\mathbf{P}}_i^{\text{Tsip}}; i = 1, \dots, N_\omega\}$, where

$$\underline{\mathbf{P}}_i^{\text{Tsip}} = \begin{pmatrix} \mathbf{P}_i^{\{p\}} \\ \mathbf{P}_i^{\{q\}} \end{pmatrix}. \quad (27)$$

In order to span the appropriate subspace, $\{\underline{\mathbf{P}}_i^{\text{Tsip}}\}$ is required to follow an orthonormality condition that is written as

$$\frac{1}{2} \left(\text{Tr} \left[\mathbf{P}_i^{\{p\}\dagger} \mathbf{S} \chi \mathbf{P}_j^{\{q\}} \mathbf{S} \phi \right] + \text{Tr} \left[\mathbf{P}_j^{\{p\}\dagger} \mathbf{S} \chi \mathbf{P}_i^{\{q\}} \mathbf{S} \phi \right] \right) = \delta_{ij}. \quad (28)$$

In the algorithm presented here, in close analogy to the way the Tamm-Dancoff eigenvalue problem is treated[20], the orthonormality condition is enforced using a Gram-Schmidt procedure. However, some additional care has to be taken, since the quantity

$\text{Tr} \left[\mathbf{P}_i^{\{p\}\dagger} \mathbf{S}^\chi \mathbf{P}_i^{\{q\}} \mathbf{S}^\phi \right]$ is not required to be positive-semidefinite. Following the convention established in Ref. [19], prior to orthogonalising the set $\{\underline{\mathbf{P}}_i^{\text{Tsip}}\}$, $\text{Tr} \left[\mathbf{P}_i^{\{p\}\dagger} \mathbf{S}^\chi \mathbf{P}_i^{\{q\}} \mathbf{S}^\phi \right]$ is computed for all i and, if found negative, the vector $\underline{\mathbf{P}}_i^{\text{Tsip}}$ is transformed according to

$$\begin{pmatrix} \mathbf{P}_i^{\{p\}} \\ \mathbf{P}_i^{\{q\}} \end{pmatrix} \rightarrow \begin{pmatrix} \mathbf{P}_i^{\{p\}} \\ -\mathbf{P}_i^{\{q\}} \end{pmatrix}. \quad (29)$$

Once an orthonormal set $\{\underline{\mathbf{P}}_i^{\text{Tsip}}; i = 1, \dots, N_\omega\}$ has been obtained, it can be used to construct $\left\{ \left(\underline{\mathbf{f}}_{\text{Tsip}}^{\chi\phi} \right)_i; i = 1, \dots, N_\omega \right\}$ from Eq. 19. It is then possible to write down an effective functional for the sum of the lowest N_ω excitations of the system, such that

$$\min_{\{\underline{\mathbf{P}}_i^{\text{Tsip}}\}} \Omega_{\text{Tsip}}^{N_\omega} \left(\{\underline{\mathbf{P}}_i^{\text{Tsip}}\} \right) = \sum_i^{N_\omega} \omega_i \quad (30)$$

where

$$\begin{aligned} \Omega_{\text{Tsip}}^{N_\omega} \left(\{\underline{\mathbf{P}}_i^{\text{Tsip}}\} \right) = \sum_i^{N_\omega} \left\{ \frac{\text{Tr} \left[\mathbf{P}_i^{\{p\}\dagger} \mathbf{S}^\chi \left(\underline{\mathbf{f}}_{\{p\}}^{\chi\phi} \right)_i \mathbf{S}^\phi \right]}{2 \left| \text{Tr} \left[\mathbf{P}_i^{\{p\}\dagger} \mathbf{S}^\chi \mathbf{P}_i^{\{q\}} \mathbf{S}^\phi \right] \right|} \right. \\ \left. + \frac{\text{Tr} \left[\mathbf{P}_i^{\{q\}\dagger} \mathbf{S}^\chi \left(\underline{\mathbf{f}}_{\{q\}}^{\chi\phi} \right)_i \mathbf{S}^\phi \right]}{2 \left| \text{Tr} \left[\mathbf{P}_i^{\{p\}\dagger} \mathbf{S}^\chi \mathbf{P}_i^{\{q\}} \mathbf{S}^\phi \right] \right|} \right\}. \quad (31) \end{aligned}$$

and the minimisation in Eqn. 30 is to be carried out under the effective orthonormality constraint placed on $\{\underline{\mathbf{P}}_i^{\text{Tsip}}\}$.

Differentiating the above expression with respect to $\underline{\mathbf{P}}_i^{\text{Tsip}}$, it is possible to construct a contravariant gradient $\left(\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} \right)_i$ that is orthogonal to all current subspace vectors $\{\underline{\mathbf{P}}_i^{\text{Tsip}}\}$. The gradient can be written as

$$\begin{aligned} \left(\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} \right)_i = \left(\underline{\mathbf{f}}_{\text{Tsip}}^{\chi\phi} \right)_i - \sum_j^{N_\omega} \left\{ \frac{\text{Tr} \left[\mathbf{P}_j^{\{p\}\dagger} \mathbf{S}^\chi \left(\underline{\mathbf{f}}_{\{p\}}^{\chi\phi} \right)_i \mathbf{S}^\phi \right]}{2 \text{Tr} \left[\mathbf{P}_j^{\{p\}\dagger} \mathbf{S}^\chi \mathbf{P}_j^{\{q\}} \mathbf{S}^\phi \right]} \right. \\ \left. + \frac{\text{Tr} \left[\mathbf{P}_j^{\{q\}\dagger} \mathbf{S}^\chi \left(\underline{\mathbf{f}}_{\{q\}}^{\chi\phi} \right)_i \mathbf{S}^\phi \right]}{2 \text{Tr} \left[\mathbf{P}_j^{\{p\}\dagger} \mathbf{S}^\chi \mathbf{P}_j^{\{q\}} \mathbf{S}^\phi \right]} \right\} \begin{pmatrix} \mathbf{P}_j^{\{q\}} \\ \mathbf{P}_j^{\{p\}} \end{pmatrix}. \quad (32) \end{aligned}$$

The fact that $\left(\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} \right)_i$ is orthogonal to all current subspace vectors $\{\underline{\mathbf{P}}_i^{\text{Tsip}}\}$ is crucial for the correct performance of the algorithm. It ensures that when using $\left(\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} \right)_i$ as a steepest descent direction to update $\underline{\mathbf{P}}_i^{\text{Tsip}}$ such that

$$\underline{\mathbf{P}}_i^{\text{Tsip}} \rightarrow \underline{\mathbf{P}}_i^{\text{Tsip}} + \lambda \left(\underline{\mathbf{g}}_{\text{Tsip}}^{\chi\phi} \right)_i \quad (33)$$

for some given line step λ , there is no violation of the orthonormality constraint of Eq. 28 to first order in λ .

The above outline contains all of the basic ingredients to construct a preconditioned conjugate gradient algorithm capable of solving for the lowest N_ω TDDFT excitations of the system in linear-scaling effort. The exact algorithm used in the implementation discussed in this work is adapted from [39] for the purpose of solving the full TDDFT eigenvalue problem (See algorithm 1 in [37]). Here we limit ourselves to the comment that while minimising the Tsiper functional of Eq. 31 yields a set of vectors spanning the same subspace as the TDDFT eigenvectors corresponding to the lowest N_ω eigenvalues, these eigenvalues can only be obtained through a subspace diagonalisation that scales as $\mathcal{O}(N_\omega^3)$. However, since N_ω is generally taken to be small, this single cubic scaling step is not considered to be a bottleneck in any practical calculation. The subspace diagonalisation can be carried out by diagonalising the $N_\omega \times N_\omega$ dimensional matrix \mathbf{Q} with matrix elements

$$Q_{ij} = \frac{\text{Tr} \left[\mathbf{P}_j^{\{p\}\dagger} \mathbf{S}^\chi \left(\mathbf{f}_{\{p\}}^{\chi\phi} \right)_i \mathbf{S}^\phi \right] + \text{Tr} \left[\mathbf{P}_j^{\{q\}\dagger} \mathbf{S}^\chi \left(\mathbf{f}_{\{q\}}^{\chi\phi} \right)_i \mathbf{S}^\phi \right]}{2}. \quad (34)$$

While the subspace diagonalisation only has to be carried out at the end of the calculation, the Gram-Schmidt orthonormalisation has an $\mathcal{O}(N_\omega^2)$ scaling associated with it and thus, like the TDA implementation[20], the algorithm only shows a linear-scaling with system size for a constant number of excitations. However, we point out that for many systems, especially the large pigment-protein complexes mentioned as the main focus of this work, low energy absorption properties are dominated by a small number of excitations of interest that stays constant with system size and for this class of systems, the method presented here allows for truly linear-scaling calculations.

IV. RESULTS

In this section we will present some of the strengths of the algorithm developed in Sec. III and specifically address the question of whether full TDDFT produces significantly more accurate results compared to the TDA for a certain class of systems. We will first focus on *trans*-azobenzene, a molecule small enough to be easily treatable with standard cubic-scaling implementations of TDDFT. After comparing the linear-scaling TDDFT approach presented here to these benchmark calculations we will discuss the importance of preconditioning the

full TDDFT eigenvalue equation in order to speed up convergence. Section IV B then aims at reproducing experimental results of bacteriochlorophyll *a* in an organic solvent, addressing again the question of whether full TDDFT provides a significant advantage over the TDA in this system, as well as the influence of an explicit treatment of solvent molecules on the spectrum. Finally, it is demonstrated that the algorithm is capable of obtaining low energy excitations in linear-scaling effort by computing the excitations in system sizes of up to ≈ 7000 atoms.

All calculations performed in this section are done using the PBE functional[40]. Norm-conserving pseudopotentials[41], as well as the projection operator and the joint NGWF set of Eq. 15 are used throughout for the ONETEP TDDFT calculations.

A. *Trans*-azobenzene

As a first test system, we choose *trans*-azobenzene ($C_{12}H_{10}N_2$) as its moderate size allows for detailed benchmark comparisons to conventional TDDFT implementations showing an $O(N^3)$ scaling. Furthermore, the optical spectrum of this system has already been studied to a high degree of accuracy using the *GW* approximation and the Bethe-Salpeter equation (BSE), where it was found that the TDA generated considerable errors compared to a full solution to the Bethe-Salpeter equation[27]. While no straightforward comparison can be drawn between the *GW*+BSE and TDDFT, the similar structure of the equations leads us to expect that TDDFT in the TDA and full TDDFT will also yield significantly different results for this system, making it an ideal test case.

First the ionic positions of *trans*-azobenzene are optimised[42] in ONETEP, after which the lowest eight excitations are computed both in the TDA and in full TDDFT. The TDA results are obtained using the algorithm introduced in Ref. [20], while the full TDDFT results are computed using the algorithm introduced in this work. Calculations are performed using a kinetic energy cutoff of 1000 eV and a box size of $56.69 \times 56.69 \times 56.69 \text{ \AA}^3$. In order to avoid any interaction between periodic images, the TDDFT calculations are carried out in open boundary conditions[43]. A minimal set of one NGWF per H and four NGWFs for each C and N atom is used both for $\{\chi_\alpha\}$ and $\{\phi_\beta\}$. A localisation radius of $10 a_0$ is applied to $\{\phi_\beta\}$ representing the occupied subspace, while the localisation is relaxed to $13 a_0$ for $\{\chi_\alpha\}$ in order to better represent the more delocalised unoccupied states. In some molecules, low energy

	ONETEP TDA	NWChem TDA	ONETEP RPA	NWChem RPA
1	2.233	2.192	2.184	2.149
2	3.520(0.047)	3.524(0.091)	3.516(0.286)	3.518(0.186)
3	3.536(0.001)	3.546(0.001)	3.489(0.002)	3.499(0.003)
4	3.720(1.094)	3.681(1.025)	3.408(0.499)	3.379(0.751)
5	3.822	3.866	3.821	3.865
6	3.875(0.003)	3.923(0.001)	3.874	3.922
7	4.234(0.001)	4.268(0.001)	4.229(0.001)	4.262(0.001)
8	4.315	4.305	4.241	4.230

TABLE I: Lowest eight excited states of azobenzene, as computed with **ONETEP** and **NWChem**. Energies are given in eV, oscillator strengths are shown in brackets. States without specified oscillator strengths are dark. Where necessary, the states have been reordered according to their character, such that the order is the same as for the **ONETEP** TDA results. For the **NWChem** calculations, an aug-cc-pVTZ Gaussian basis set is used. Here, TDA denotes the Tamm-Dancoff approximation, while RPA is used to denote a solution to the full TDDFT equation.

excitations can be found that are of very delocalised Rydberg-state character and provide a challenge for localised basis set representations. However, tests performed using TDA have shown that these states can be systematically converged by increasing localisation radius of $\{\chi_\alpha\}$, and converged values are generally found to be in good agreement with results obtained from real-space methods[20]. For the purpose of this work, we have performed a number of convergence tests, increasing the localisation radius to $18 a_0$ and found the lowest eight excitations of alizarin to be well converged at the localisation radius of $13 a_0$.

The TDDFT results from **ONETEP** are compared to results obtained using **NWChem**[44]. The **NWChem** calculations are performed using the same ionic positions as in **ONETEP** and an aug-cc-pVTZ Gaussian basis set containing diffuse functions that are designed to yield a good description of weakly bound unoccupied states[45]. In order to make the all-electron calculations more comparable to the pseudopotential calculations in **ONETEP**, the Kohn-Sham states corresponding to the core electrons are excluded from the occupied subspace when calculating the TDDFT excitation energies.

The results for the lowest eight excitations for the TDA and full TDDFT as calculated

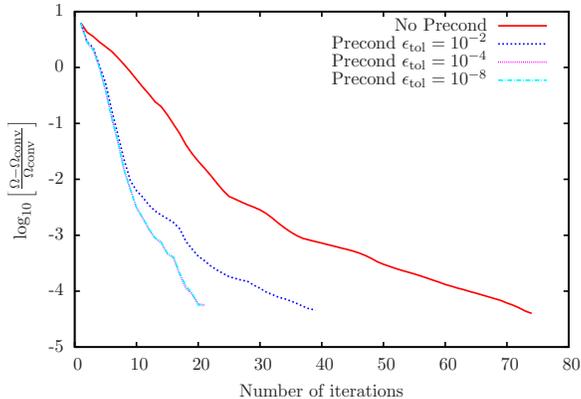


FIG. 1: Convergence of the 2 lowest states of Azobenzene for different degrees of preconditioning applied. “Precond ϵ_{tol} ” describes the tolerance to which the linear system (26) is solved in order to apply the preconditioning.

using **ONETEP** in comparison to the NWChem results can be found in Table I. As can be seen, there is a generally good agreement between the **ONETEP** and the NWChem results, both in terms of the excitation energy and the oscillator strength. The largest discrepancy is found in the the sixth excited state, with the **ONETEP** results being 48 meV lower in energy for both the TDA and full TDDFT. Most other results show a significantly smaller discrepancy.

Generally, the full TDDFT results compare very similarly to the NWChem benchmark as the TDA results, showing that the algorithm indeed performs correctly. Remaining discrepancies are most likely due to the all-electron treatment in NWChem compared to the pseudopotential treatment in **ONETEP**, as well as basis set differences (see Ref. [20] for a detailed comparison between NWChem and **ONETEP** regarding TDA results). It is thus evident that the minimal NGWF representation is capable of obtaining excitation energies of a comparable quality to those from considerably larger Gaussian basis set representations, highlighting the advantages of using an *in situ* optimised representation when performing TDDFT calculations.

Regarding the impact of full TDDFT on the low energy excited states of *trans*-azobenzene, it becomes apparent that when solving the full TDDFT equations, the dominant low energy excited state, which is mainly made up of a HOMO-1 \rightarrow LUMO transition, decreases in energy by more than 0.31 eV as compared with the TDA results. The peak with the second highest oscillator strength, corresponding to a HOMO-2 \rightarrow LUMO transition,

stays almost constant in energy. In the full TDDFT, a significant part of spectral weight is shifted from the dominant peak to the second peak. This result is in good qualitative agreement with the $GW+BSE$ results, where it was found that the TDA blue-shifts the dominant peak by 0.2 eV and overestimates its oscillator strength[27]. It is thus clear that solving the full TDDFT equation instead of the TDA can lead to significant changes in the computed optical spectrum, with bright peaks being shifted by tenths of electronvolts and a significant redistribution of spectral weight occurring.

The small size of *trans*-azobenzene makes it an ideal system to verify the effectiveness of the preconditioning introduced in this work. For this purpose, we compute the two lowest excited states of the system for a number of different convergence tolerances ϵ_{tol} of the iterative preconditioner[37], as well as the case when no preconditioning is applied to the conjugate gradient algorithm. The convergence rate of the two lowest excited states with respect to different levels of preconditioning applied can be found in Fig. 1. Note that iteratively applying the preconditioner to a tolerance of 10^{-4} cuts the number of iterations needed to reach convergence by almost a factor of four compared to the case where no preconditioning is applied. This highlights the ill-conditioning of the TDDFT eigenvalue equation mentioned in section III B and shows that preconditioning is vital in achieving good convergence rates. However, note that even when the iterative preconditioning tolerance is set as high as 10^{-2} , corresponding to only applying the preconditioner approximately at each iteration, the number of iterations needed to reach convergence is decreased by a factor of two. This finding is vital as solving the linear system of Eq. 26 iteratively in each conjugate gradient step has a computational overhead associated with it, which can be minimised if the solution is only obtained approximately rather than to a high degree of numerical accuracy. It should be pointed out however, that the computational overhead associated with the preconditioning is negligible for the system at hand, making up 0.17% and 0.53% of the total calculation time for $\epsilon_{\text{tol}} = 10^{-2}$ and $\epsilon_{\text{tol}} = 10^{-4}$, which can again be attributed to the compact size of the NGWF representation. Furthermore, it is found that the total calculation time is reduced by a factor of 2.86 when comparing the unconditioned system to a preconditioned system with $\epsilon_{\text{tol}} = 10^{-4}$, thus showing that the preconditioner introduced here indeed leads to significant reductions in computational effort.

We note that the convergence of the preconditioned system with the high convergence tolerance of $\epsilon_{\text{tol}} = 10^{-2}$ closely follows the fast convergence of the tight tolerance results

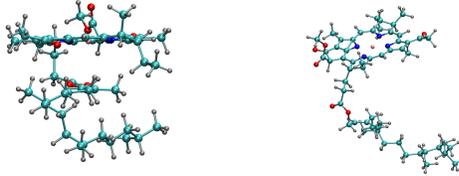


FIG. 2: Structure of bacteriochlorophyll *a*, as obtained from a geometry optimisation in vacuum (left) and a single snapshot from an MD simulation in toluene (right). This figure was created using VMD[50].

of $\epsilon_{\text{tol}} = 10^{-8}$ for the first ten iterations. This suggests that the convergence tolerance of the preconditioner can be chosen adaptively. Starting off with a high tolerance for the first few iterations and tightening it closer to convergence has the potential to provide the ideal balance between reducing the computational overhead of the preconditioning and increasing the convergence rate. In conclusion it is demonstrated that the preconditioned conjugate gradient algorithm introduced in this work yields a very good agreement with existing TDDFT implementations and shows excellent convergence rates.

B. Bacteriochlorophyll

We now shift the focus to bacteriochlorophyll *a* ($\text{MgN}_4\text{O}_6\text{C}_{55}\text{H}_{74}$), a chromophore that is of great interest in computational biology due to its role in light-harvesting complexes[1–3]. Bacteriochlorophyll (BChl) is a medium-sized system still within the range of conventional cubic scaling TDDFT implementations. Therefore, the focus in this section is not to demonstrate the linear-scaling capabilities of the algorithm developed in this work but rather to address the question of whether full TDDFT yields a better description of the low energy absorption spectrum than the TDA when compared to experimental results.

All TDDFT calculations performed in this section are carried out using a minimal set of NGWFs for all atomic species involved and a localisation radius of $10 a_0$ and $13 a_0$ for the NGWFs representing the occupied and unoccupied space respectively. A kinetic energy cutoff of 1020 eV is used in all calculations.

We aim to compare to the experimental results for the low energy absorption spectrum of BChl *a* in a toluene organic solvent[46, 47]. As a first step, we take the atomic positions of BChl *a* site 1 in the Fenna-Matthews-Olson complex (see Ref. [3] for an explanation of

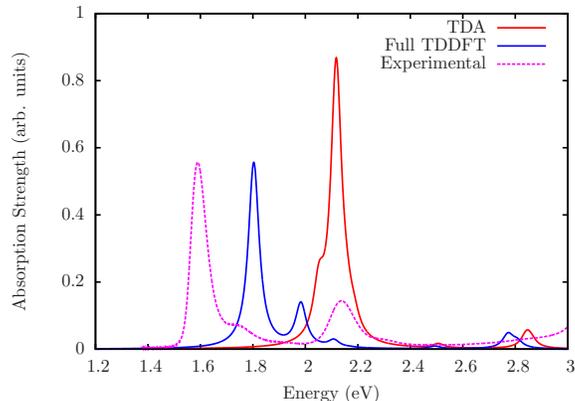


FIG. 3: Absorption spectrum of bacteriochlorophyll in implicit solvent as calculated with both the TDA and full TDDFT in comparison to experimental data of bacteriochlorophyll in toluene. A Lorentzian broadening of 0.025 eV is used for the TDDFT results and the experimental data is scaled such that the peak height of the Q_y transition agrees with that of the Q_y peak obtained from full TDDFT.

how the input structure was obtained from X-ray diffraction results), optimise the atomic positions in vacuum and then calculate the TDA and full TDDFT spectra of the system within an implicit solvent model[48], where the static dielectric function ϵ is chosen to be 2.38 in order to match the dielectric function of toluene at room temperature. The final structure of BChl *a* in vacuum can be found in Fig. 2. As can be seen, the porphine ring is entirely flat in this configuration, while the alkane tail folds underneath the ring structure.

The results of the TDDFT calculation as compared with the experimental results[53] can be found in Fig. 3. As can be seen the experimental spectrum shows three main features: A main absorption peak at around 1.6 eV, a shoulder between 1.65 and 1.75 eV and a second peak at around 2.1 eV. The first and second peaks are commonly referred to as the Q_y and Q_x transitions respectively and can be characterised as HOMO→LUMO and HOMO-1→LUMO transitions in a single-particle picture. As can be seen from Fig. 3, the TDA results generate only a single main peak at 2.12 eV that is of Q_y character. This main peak shows a shoulder at 2.05 eV that is mainly of HOMO-1→LUMO and HOMO-2→LUMO character. The full TDDFT spectrum on the other hand produces a Q_y peak at 1.80 eV that is the lowest excitation of the system, as well as a second peak of Q_x character at 1.98 eV, but fails to reproduce a shoulder to the Q_y peak. It also shows a third peak with small oscillator strength at 2.11 eV that is of HOMO-2→LUMO character.

It can therefore be concluded that the TDA fails in correctly reproducing the absorption spectrum of BChl *a* in toluene. Not only is the Q_y transition overestimated by 0.47 eV compared to the experimental results, it does not correspond to the lowest excitation of the system and there is no clean Q_x transition. The full TDDFT results show a considerable improvement. While the Q_y transition is still overestimated by 0.2 eV, it now corresponds to the lowest excitation of the system and there is a considerable splitting between the Q_y and Q_x transitions. However, the full TDDFT results underestimate the energy of the Q_x transition by around 0.14 eV compared to experimental results and fail to exhibit a shoulder to the Q_y transition.

The origin of some of the failures of the TDA can be traced by breaking down the excitations into individual Kohn-Sham transitions. In full TDDFT, the Q_y transition is almost exclusively (to 95%, as compared to 60% in the TDA) a transition between the HOMO and the LUMO. In Bacteriochlorophyll, this transition has a strong dipole moment associated with it, which in turn causes $\mathbf{V}_{\text{SCF}}^{\{1\}}$ to be large and the TDDFT energy to show a large increase compared to the HOMO-LUMO energy difference. In full TDDFT, this large dipole is screened by the de-excitation vector \mathbf{Y} which is almost entirely made up of the same HOMO-LUMO transition, thus significantly lowering both the excitation energy and the oscillator strength. In the TDA, $\mathbf{Y} = \mathbf{0}$ and instead the large dipole moment of the HOMO-LUMO transition is screened by mixing in smaller fractions of higher energy transitions, including fractions of Q_x transition. This causes the Q_y transition to have a significantly higher energy and larger oscillator strength in the TDA and also contributes to the absence of a clean Q_x transition.

Since optimising the atomic positions of BChl *a* in vacuum might have lead to a structure that is unrealistic for the system solvated in toluene, it is not clear how much of the failure of full TDDFT to reproduce the experimental results is due to the choice of exchange-correlation functional used in this work. In order to obtain a more realistic structure of BChl *a* in toluene, we make use of the classical molecular dynamics package AMBER[49]. We solvate BChl *a* in 704 molecules of toluene (corresponding to 10,700 atoms for the total system) and equilibrate the system to 300 K and a pressure of 1 atm, followed by a 300 ps simulation in the NVE ensemble. From this MD run we take 8 snapshots 10 ps apart that are then used as input atomic positions in the TDDFT calculations. In order to include solvent effects at a quantum mechanical level, we explicitly include all toluene molecules

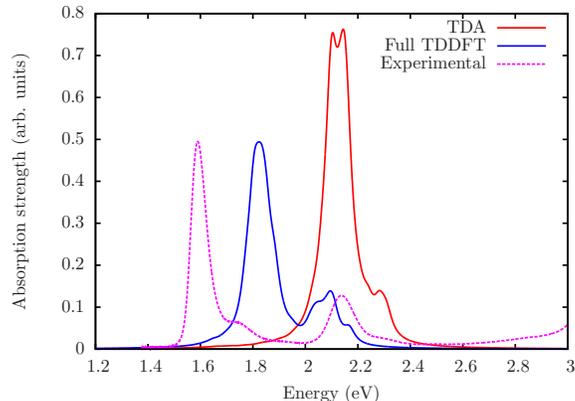


FIG. 4: Absorption spectrum of bacteriochlorophyll *a* in 15 Å of toluene as calculated on 8 MD snapshots with both the TDA and full TDDFT in comparison to experimental data of bacteriochlorophyll *a* in toluene. A Lorentzian broadening of 0.025 eV is used for the TDDFT results and the experimental data is scaled such that the peak height of the Q_y transition agrees with that of the Q_y peak obtained from full TDDFT.

within a 15 Å radius from the Mg atom in the calculation, while representing the rest of the solvent by an implicit solvation model. This process yields a system size of 700-800 atoms depending on the snapshot, which is closer to the limit of sizes that can be treated by conventional $\mathcal{O}(N^3)$ methods. The structure of the BChl *a* molecule as obtained from a single MD snapshot can be found in Fig. 2. As can be seen, the alkane tail extends away from the porphine ring in this configuration, while the ring itself is no longer perfectly flat.

The averaged full TDDFT and TDA spectra of the solvated MD snapshots as compared with the same experimental dataset are shown in Fig. 4. Note that the full TDDFT results now show the Q_y transition at 1.83 eV and the Q_x transition at 2.10 eV. While the Q_y transition is still overestimated by 0.2 eV, the shape of the feature is considerably improved, as the Q_x transition is now in very good agreement with experimental results, both in its positioning and in its intensity. It is worth pointing out that a number of snapshots also show a shoulder to the main Q_y transition that is of HOMO-2→LUMO character. However, in the averaged spectrum this feature is not as pronounced as in the experimental spectrum, which can be due to the fact that the splitting between the Q_x and Q_y transitions is too small compared to experiment. The TDA results on the other hand still fail to reproduce main spectral features. The Q_y transition is overestimated by 0.5 eV, although the shoulder

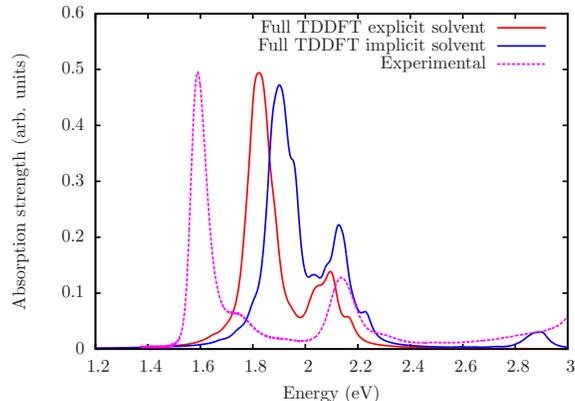


FIG. 5: Comparison between the full TDDFT spectrum of bacteriochlorophyll in 15 Å of toluene, averaged over 8 snapshots, and the averaged spectrum from 8 implicit solvent calculations using the same atomic positions for the bacteriochlorophylls. A Lorentzian broadening of 0.025 eV is used for the TDDFT results. Both spectra are plotted against experimental data that is scaled such that the peak height of the Qy transition agrees with that of the Qy peak obtained from full TDDFT in explicit toluene.

present in Fig. 3 has disappeared. The spectrum shows a new peak at approximately 2.3 eV that does however have a different character to the Qx transition in full TDDFT. A clearly identifiable Qx transition is still absent from the TDA results. It can therefore be summarised that full TDDFT yields a much improved representation of the experimental results at the PBE level as long as a realistic structure for the solute and the solvent environment is obtained. The main failure of full TDDFT for this system is the overestimation of the Qy transition by 0.2 eV which can most likely be ascribed to errors in the exchange-correlation functional used.

While the 700-800 atom systems obtained from classical MD yield a relatively good spectral shape for full TDDFT at the PBE level, they are considerably larger than the 140 atoms of the solute alone. It is therefore worth investigating how much of the improvement of the spectrum from Fig. 3 to Fig. 4 is due to the different ionic positions of the solute and how much is due to an explicit quantum mechanical treatment of solvent molecules in the calculation. For this purpose we take the atomic positions of BChl *a* from the 8 MD snapshots and compute the absorption spectrum in implicit solvent, without including any explicit representation of toluene molecules. The result can be found in Fig. 5. As can be seen,

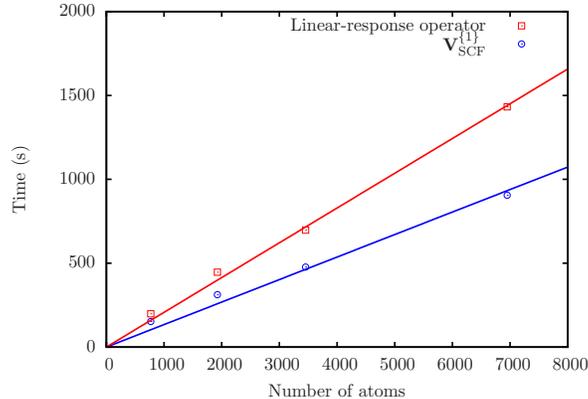


FIG. 6: Time taken for applying the TDDFT operator and calculating $\mathbf{V}_{SCF}^{\{1\}}$ for different system sizes of bacteriochlorophyll *a* in toluene. The lines shown are linear fits. Calculations were performed on eight SandyBridge nodes containing 16 cores each.

the positioning of the Qx transition in implicit solvent is in very good agreement with the experimental data. However, its oscillator strength is significantly overestimated. Furthermore, the explicit solvent environment causes the Qy transition to red-shift by about 0.1 eV. It can be concluded that while most of the improvements in spectral features compared to Fig. 3 are due to the more realistic atomic positions of the Bacteriochlorophyll in toluene, the explicit inclusion of the toluene environment at the TDDFT level yields a spectrum that is in closest agreement with experimental results.

In conclusion it can be summarised that full TDDFT at the PBE level reproduces experimental results to a satisfactory degree, while the TDA completely fails in this system. The best agreement between experiment and calculation is obtained when taking the atomic positions of an MD snapshot of BChl in toluene and including the local solvent environment explicitly in the TDDFT calculation. While an explicit treatment of the solvent molecules requires large scale TDDFT calculations, it is demonstrated that the method presented in this work is well suited for tackling these systems, opening up the possibility of more detailed studies of solvent effects on chromophores.

C. Linear-scaling capabilities

We now focus on demonstrating the linear-scaling capabilities of the TDDFT method presented in this work. For this purpose, we take one of the classical MD snapshots studied in

the previous section and compute the Qy transition of BChl *a*, including progressively larger regions of the solvent environment. Calculations are performed using the same parameters as in the previous section, apart from using a smaller radius of 10 a_0 for the NGWF set $\{\chi_\alpha\}$. In order to reach linear-scaling computational effort with system size, it is necessary to truncate all involved density matrices $\mathbf{P}^{\{v\}}$, $\mathbf{P}^{\{c\}}$, $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ to make matrix-matrix operations scale as $\mathcal{O}(N)$. Here, we choose a tight truncation radius of 20 a_0 , causing the unoccupied and ground-state density matrices to have the same sparsity pattern as the NGWF representation overlap matrices \mathbf{S}^x and \mathbf{S}^ϕ . However, by performing a full TDDFT calculation on the 770-atom system of the previous section, we have confirmed that the error introduced to the energy of the Qy transition for such a cutoff is only 0.03 eV.

While the Qy transition of interest in this system retains a relatively localised character and the error introduced through truncating $\mathbf{P}^{\{p\}}$ and $\mathbf{P}^{\{q\}}$ can be expected to be relatively small, the total error of 0.03 eV is a combination of errors introduced by the response density matrices and by the truncation of the ground state density matrix. The calculations show, that for relatively localised excitations, fully linear-scaling calculations on realistic systems are possible with only introducing minor errors. For very delocalised excitations, the truncation of the response density matrix becomes more difficult and the effects of applying a truncation on long-range charge-transfer excitations has been discussed elsewhere in more detail[20]. Here we note, that for a large class of systems like pigment-protein complexes, excitations of interest are expected to retain a relatively localised character and fully linear-scaling calculations are indeed possible.

We choose to perform a linear-scaling test of the full TDDFT method on four different system sizes of Bchl in Toluene, each specified by the radius, as measured from the Mg atom, up to which solvent molecules are included in the calculation. The four different radii chosen are 15, 20, 24 and 30 Å, corresponding to system sizes of 770, 1925, 3455 and 6950 atoms. The timings, both for evaluating the self-consistent field response $\mathbf{V}_{\text{SCF}}^{\{1\}}$ and applying the full TDDFT operator (19) are shown in Fig. 6. In general, a clear linear trend can be observed for both the calculation of $\mathbf{V}_{\text{SCF}}^{\{1\}}$ and the TDDFT operator, with slight discrepancies for the two smaller systems. These discrepancies are likely originating from the fact that matrices in the two smallest systems are still relatively dense and a full transition to linear scaling effort only occurs at larger system sizes. However, it becomes clear that the algorithm is capable of solving the full TDDFT eigenvalue equation for systems of thousands of atoms in linear-

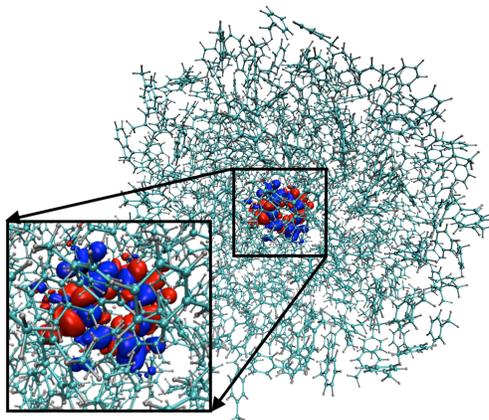


FIG. 7: The electron-hole difference density of the Q_y transition of bacteriochlorophyll surrounded by a sphere of 30 Å radius of toluene solvent molecules. The calculation includes 454 toluene molecules, corresponding to a total system size of 6950 atoms. The figure was created using VMD[50].

scaling effort. Furthermore, recent work implementing hybrid OpenMP-MPI approaches to parallelism mean that these calculations scale efficiently to many thousands of CPU cores[51].

Figure 7 shows a plot of the electron-hole difference density for the fully converged Q_y transition of the 6950 atom system. It should be noted that these system sizes are inaccessible by standard $\mathcal{O}(N^3)$ approaches, both at the DFT and the TDDFT level. The calculation presented here is to be seen for demonstration purposes only, given that the Q_y transition retains a relatively localised character and a fully quantum mechanical treatment of such a large region of the solvent environment can be considered unnecessary. However, it should be pointed out that the system treated here is of similar size as a variety of pigment-protein complexes, most notably the Fenna-Matthews-Olson complex, where seven BChl molecules are embedded in a complex protein environment and site energy variations due to environmental screening effects are both subtle and important[1–3]. The method presented in this work has the potential of treating these systems fully quantum mechanically, without relying on semi-classical approximations.

V. CONCLUSION

In conclusion, we have outlined a preconditioned conjugate gradient algorithm capable of solving for the lowest eigenvalues of the full TDDFT equation with linear-scaling ef-

fort. We have demonstrated the efficiency of the compact NGWF representation and shown that it yields results comparable to those obtained with well-converged Gaussian basis sets. Furthermore, the vital importance of preconditioning the iterative solution to the TDDFT equation has been demonstrated, yielding a four-fold speedup of the convergence rate in the case of *trans*-azobenzene.

We have furthermore shown that the TDA fails to reproduce experimental results of BChl *a* in toluene solution at the PBE level, while full TDDFT yields a significant improvement of all spectral features. It was also shown that the best results compared to experimental data are achieved when treating a certain part of the solvent explicitly at TDDFT level, making it necessary to perform TDDFT calculations of several hundreds of atoms. While these system sizes can reach the limits of standard $\mathcal{O}(N^3)$ methods, they are straightforwardly treated with the method introduced here, opening up the possibility of more detailed studies on the effects of pigment-solvent interactions on excitation energies of chromophores.

Finally, we have shown that the algorithm scales fully linearly with system size as long as all involved density matrices are truncated. The largest full TDDFT calculation performed in this work treats a system of 6950 atoms, far larger than systems that can be realistically addressed with cubic-scaling approaches. These large-scale systems are of the same order of magnitude as a large variety of pigment-protein complexes that are studied in the field of computational biology, opening up the possibility of computing their excitation spectra without the need to rely on any semi-classical approximations.

Acknowledgments

TJZ acknowledges the support of EPSRC Grant EP/J017639/1 and the ARCHER eCSE programme. MCP and PDH acknowledge the support of EPSRC grant EP/J015059/1. The underlying data of this publication can be accessed via the following persistent URI: <https://www.repository.cam.ac.uk/handle/1810/251293>

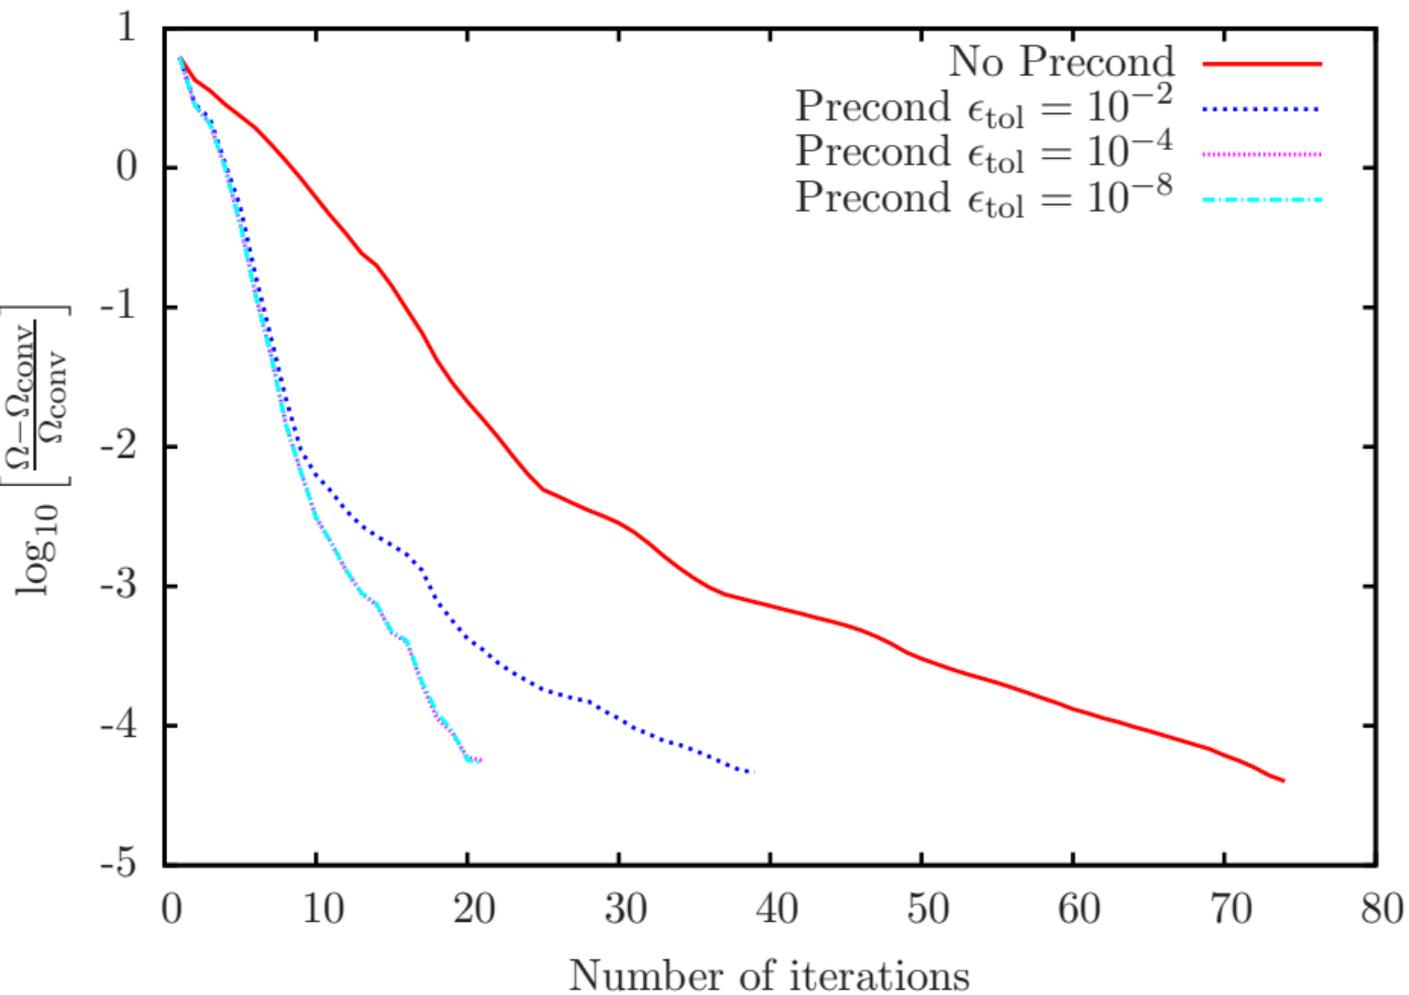
-
- [1] S. Shim, P. Rebentrost, S. Valleau, and A. Aspuru-Guzik, *Biophys. J.* **102**, 649-660 (2012)
 - [2] C. König, and J. Neugebauer, *J. Chem. Theory Comput.*, **9** 1808-1820 (2013)

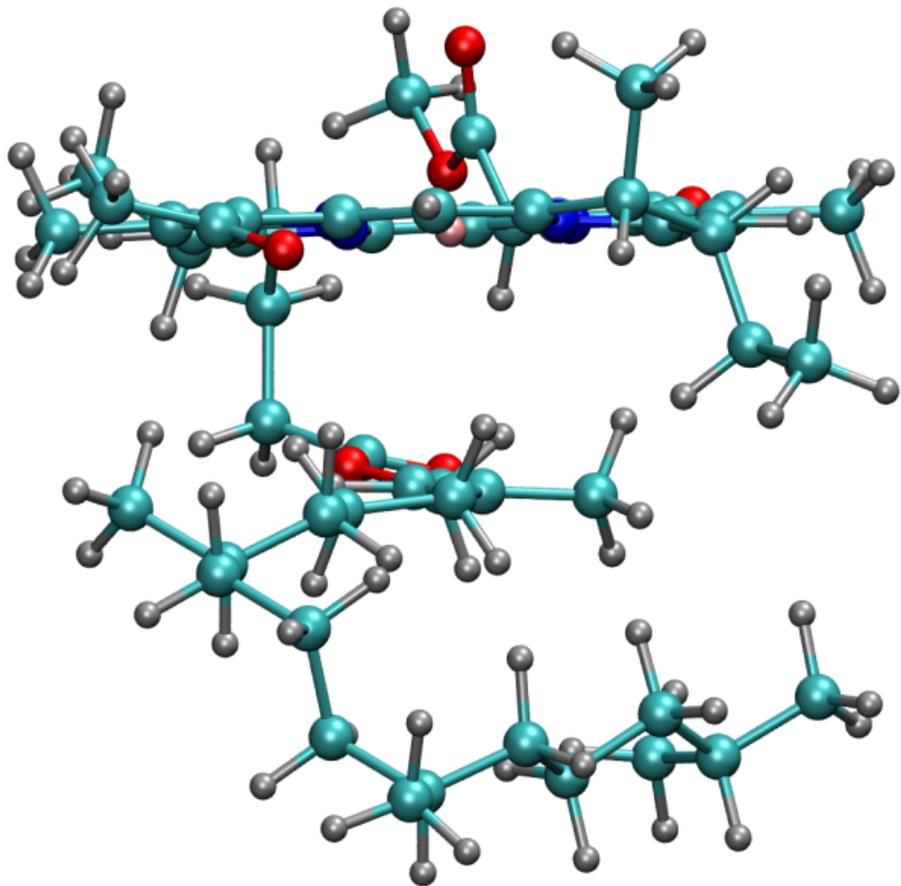
- [3] D. J. Cole, A. W. Chin, N. D. M. Hine, P. D. Haynes, and M. C. Payne, *J. Phys. Chem. Lett.* **4**, 4206 (2013)
- [4] C. M. Isborn, A. W. Götz, M. A. Clark, R. C. Walker, T. J. Martínez, *J. Chem. Theory Comput.* **8**, 5092-5106 (2012)
- [5] C. König, and J. Neugebauer, *Chem. Phys. Chem.* **13**, 386-425 (2012)
- [6] E. Runge, and E. K. U. Gross, *Phys. Rev. Lett.* **52**, 997 (1984)
- [7] P. Hohenberg, and W. Kohn, *Phys. Rev.* **136**, B864-B871 (1964)
- [8] W. Kohn, and L. J. Sham, *Phys. Rev.* **140**, 1133 (1965)
- [9] D. Rocca, R. Gebauer, Y. Saad, and S. Baroni, *J. Chem. Phys.* **128**, 154105 (2008)
- [10] H. Hübener, and F. Giustino, *Phys. Rev. B* **89**, 085129 (2014)
- [11] H. Hübener, and F. Giustino, *J. Chem. Phys.* **141**, 044117 (2014)
- [12] K. Yabana, and G. F. Bertsch, *Phys. Rev. B* **54**, 4484 (1996)
- [13] M. Petersilka, U. J. Grossmann, and E. K. U. Gross, *Phys. Rev. Lett.* **76**, 1212 (1996)
- [14] M. E. Casida, *Recent Advances in Density Functional Methods*, In: Chong, D.P. (ed.) vol. 1. World Scientific, Singapore (1995)
- [15] G. Onida, L. Reining, and A. Rubio, *Rev. Mod. Phys.* **74**, 601 (2002)
- [16] C. Yam, S. Yokojima, and G. Chen, *Phys. Rev. B* **68**, 153105 (2003)
- [17] G. Cui, W. Fang, and W. Yang, *Phys. Chem. Chem. Phys.* **12**, 416-421 (2010)
- [18] C. O'Rourke, and D. R. Bowler, *J. Chem. Phys.* **143**, 102801 (2015)
- [19] S. Tretiak, C. M. Isborn, A. M. N. Niklasson, and M. Challacombe, *J. Chem. Phys.* **130**, 054111 (2009)
- [20] T. J. Zuehlsdorff, N. D. M. Hine, J. S. Spencer, N. M. Harrison, D. J. Riley, and P. D. Haynes, *J. Chem. Phys.* **139**, 064104 (2013)
- [21] M. Challacombe, *Computation* **2**, 111 (2014)
- [22] A. L. Fetter, and J. D. Walecka, *Quantum Theory of Many-Particle Systems* (McGraw-Hill, New York, 1971)
- [23] S. Hirata, and M. Head-Gordon, *Chem. Phys. Lett.* **314**, 291-299 (1999)
- [24] P. D. Haynes, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, *J. Chem. Phys.* **122**, 084119 (2005)
- [25] A. Dreuw, J. L. Weisman, and M. Head-Gordon, *J. Chem. Phys.* **119**, 2943 (2003)
- [26] J. Berkowitz, *Photoabsorption, Photoionization, and Photoelectron Spectroscopy* (Academic

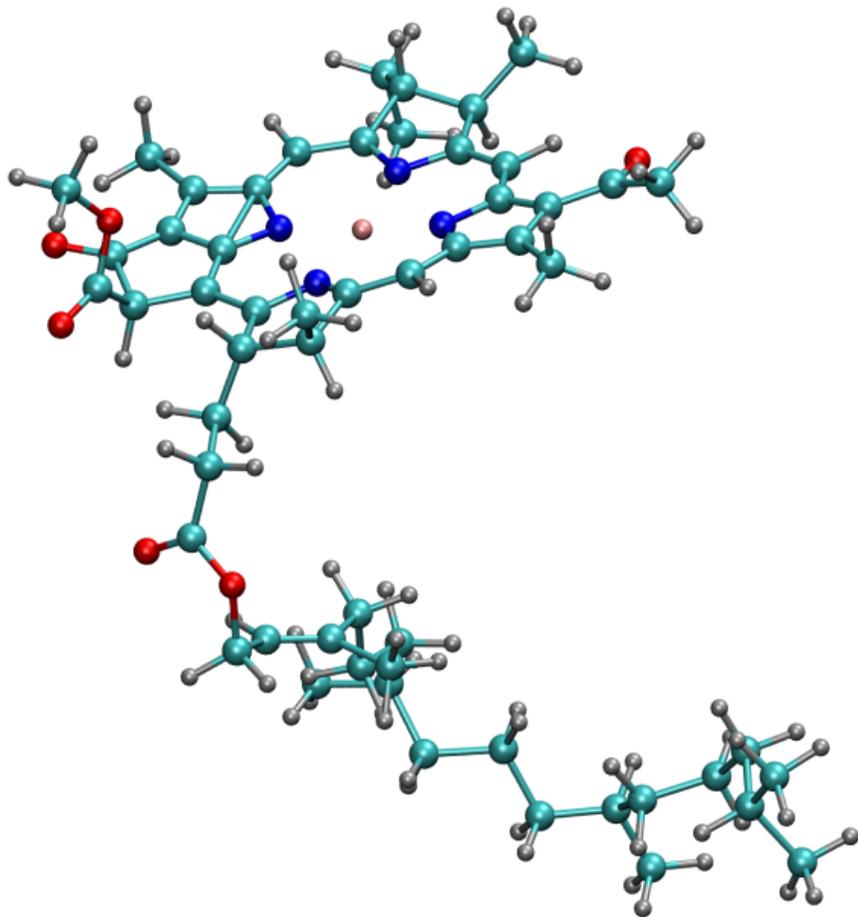
- Press, New York, 1979)
- [27] M. Grüning, A. Marini, and X. Gonze, *Nano Lett.* **9**(8), 2820-2824(2009)
 - [28] D. J. Thouless, *Nucl. Phys.* **22**, 78 (1961)
 - [29] E. V. Tsiper, *J. Phys. B* **34**, L401 (2001)
 - [30] L. E. Ratcliff, N. D. M. Hine, and P. D. Haynes, *Phys. Rev. B* **84**, 165131 (2011)
 - [31] J.-H. Li, T. J. Zuehlsdorff, M. C. Payne, and N. D. M. Hine, *Phys. Chem. Chem. Phys.* **17**, 12065-12079 (2015)
 - [32] T. J. Zuehlsdorff, *Computing the Optical Properties of Large Systems* (Springer International Publishing, Switzerland, 2015)
 - [33] A. A. Mostofi, C.-K. Skylaris, P. D. Haynes, and M. C. Payne, *Comput. Phys. Commun.* **147**, 788 (2002)
 - [34] G. Nenciu, *Commun. Math. Phys.* **91**, 81 (1983)
 - [35] S. Ismail-Beigi, and T. A. Arias, *Phys. Rev. Lett.* **82**, 2127 (1999)
 - [36] Stratmann, R. E., Scuseria, G., and Frisch, M. J., *J. Chem. Phys.* **109**, 8218 (1998)
 - [37] See supplemental material at [URL will be inserted by AIP] for an outline of the conjugate gradient algorithm used in this work.
 - [38] X.-P. Li, R. W. Nunes, and D. Vanderbilt, *Phys. Rev. B* **47**, 10891 (1993)
 - [39] C. K. Gan, P. D. Haynes, and M. C. Payne, *Comput. Phys. Commun.* **134**, 33-40 (2001)
 - [40] J. P. Perdew, K. Burke, and M. Ernzerhof, *Phys. Rev. Lett.* **77**, 3865 (1996)
 - [41] N. Troullier, and J. L. Martins, *Phys. Rev. B* **43**, 1993 (1991)
 - [42] N. D. M. Hine, M. Robinson, P. D. Haynes, C.-K. Skylaris, M. C. Payne, and A. A. Mostofi, *Phys. Rev. B* **83**, 195102 (2011)
 - [43] N. D. M. Hine, J. Dziedzic, P. D. Haynes, and C.-K. Skylaris, *J. Chem. Phys.* **135**, 204103 (2011)
 - [44] M. Valiev, E. J. Bylaska, N. Govind, K. Kowalski, T. P. Straatsma *et al.*, *Comput. Phys. Commun.* **181**, 1477 (2010)
 - [45] T. H. Dunning Jr., *J. Chem. Phys.* **90**, 1007 (1989)
 - [46] J. S. Connolly, E. B. Samuel, and A. F. Janzen, *Photochem. Photobiol.* **36**, 565-574(1982)
 - [47] H. Du, R.-C. A. Fuh, J. Li, L. A. Corkan, and J. S. Lindsey, *Photochem. Photobiol.* **68**, 141-142
 - [48] J. Dziedzic, H. H. Helal, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, *Europhys. Lett.* **95**,

43001 (2011)

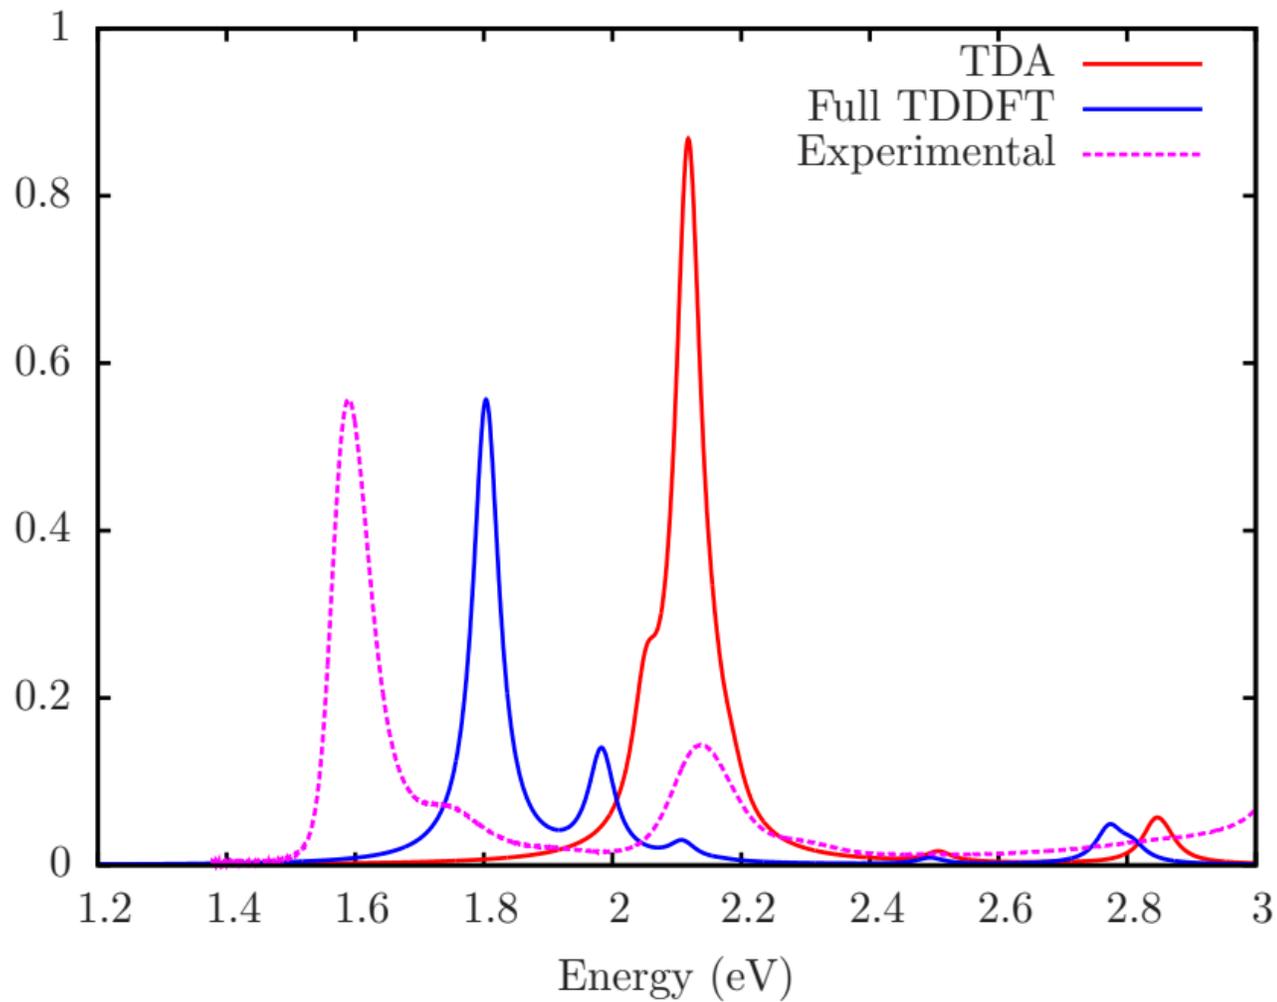
- [49] D. A. Case, J. T. Berryman, R. M. Betz, D. S. Cerutti *et al.* (2015), AMBER 2015, University of California, San Francisco.
- [50] W. Humphrey, A. Dalke, and D. Schulten, *J. Molec. Graphics* **14**, 3338 (1996)
- [51] K. A. Wilkinson, N. D. M. Hine, and C.-K. Skylaris, *J. Chem. Theory Comput.* **10**, 4782 (2014)
- [52] See the discussion in [20] regarding the question when a truncation of $\mathbf{P}^{\{1\}}$ is strictly justified
- [53] The raw data of the experimental results is taken from <http://omlc.org/spectra/PhotochemCAD/html/135.html>



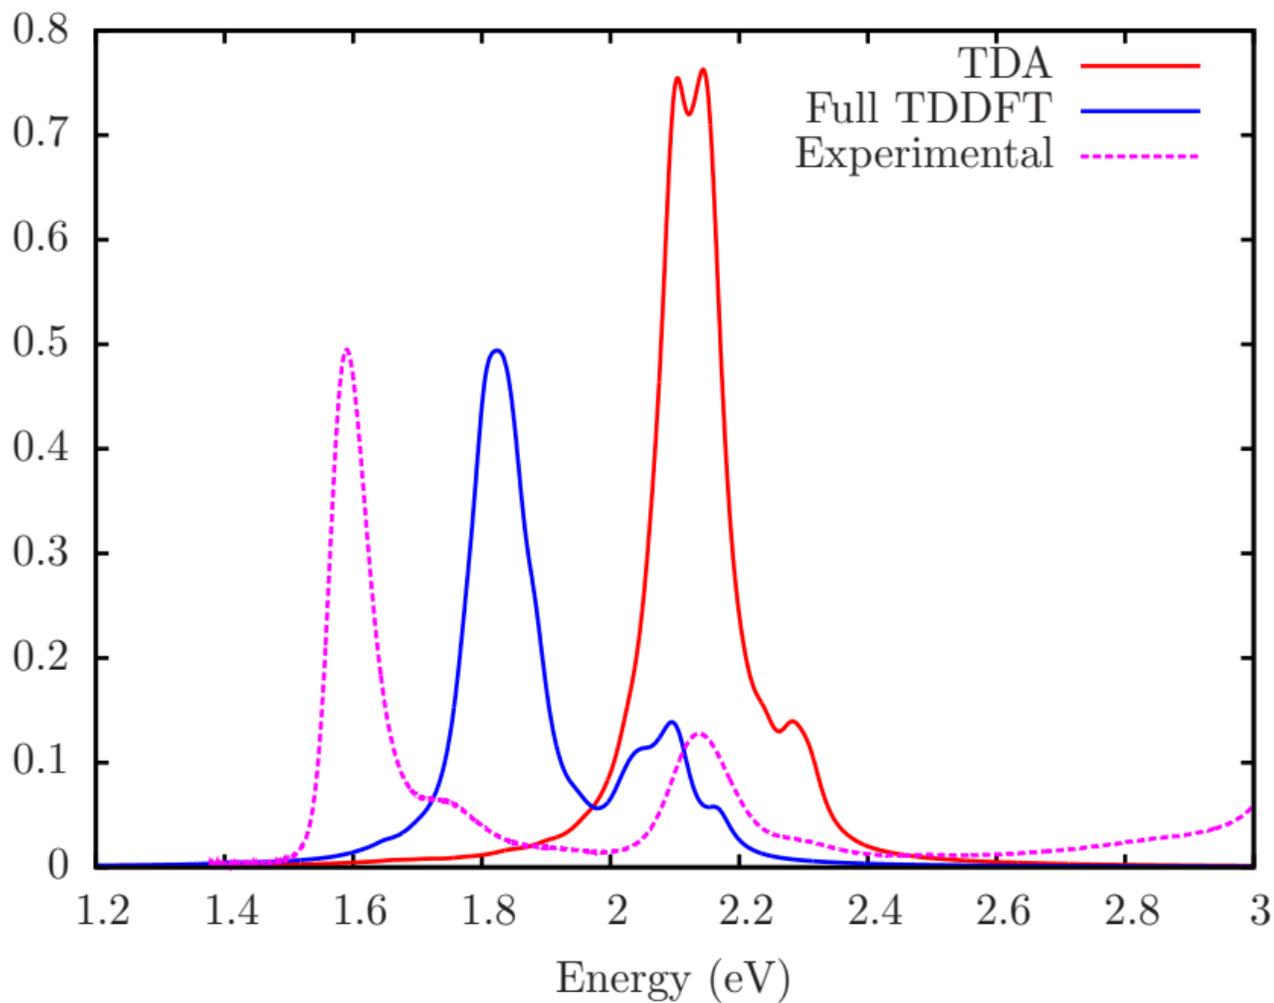




Absorption Strength (arb. units)



Absorption strength (arb. units)



Absorption strength (arb. units)

