

Linear scaling density matrix real time TDDFT: Propagator unitarity and matrix truncation

Conn O'Rourke and David R. Bowler

Citation: *The Journal of Chemical Physics* **143**, 102801 (2015); doi: 10.1063/1.4919128

View online: <http://dx.doi.org/10.1063/1.4919128>

View Table of Contents: <http://scitation.aip.org/content/aip/journal/jcp/143/10?ver=pdfcov>

Published by the [AIP Publishing](#)

Articles you may be interested in

[Cholesky-decomposed density MP2 with density fitting: Accurate MP2 and double-hybrid DFT energies for large systems](#)

J. Chem. Phys. **140**, 224112 (2014); 10.1063/1.4881144

[Daubechies wavelets for linear scaling density functional theory](#)

J. Chem. Phys. **140**, 204110 (2014); 10.1063/1.4871876

[Energy conserving, linear scaling Born-Oppenheimer molecular dynamics](#)

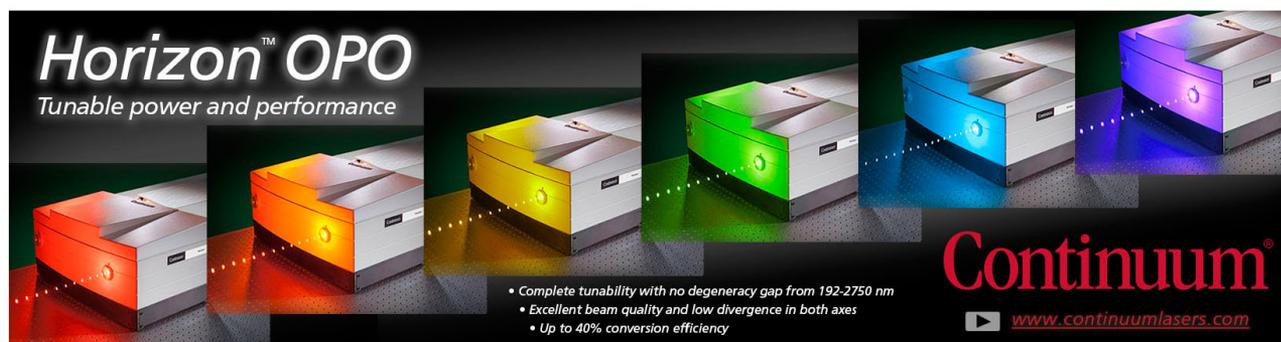
J. Chem. Phys. **137**, 134105 (2012); 10.1063/1.4755991

[A density matrix-based method for the linear-scaling calculation of dynamic second- and third-order properties at the Hartree-Fock and Kohn-Sham density functional theory levels](#)

J. Chem. Phys. **127**, 204103 (2007); 10.1063/1.2794033

[Hartree-Fock and Kohn-Sham atomic-orbital based time-dependent response theory](#)

J. Chem. Phys. **113**, 8908 (2000); 10.1063/1.1318745



Horizon™ OPO
Tunable power and performance

• Complete tunability with no degeneracy gap from 192-2750 nm
• Excellent beam quality and low divergence in both axes
• Up to 40% conversion efficiency

Continuum®
www.continuumlasers.com

Linear scaling density matrix real time TDDFT: Propagator unitarity and matrix truncation

Conn O'Rourke^{a),b)} and David R. Bowler^{b)}

London Centre for Nanotechnology, University College London, 17-19 Gordon St., London WC1H 0AH, United Kingdom

(Received 7 November 2014; accepted 19 January 2015; published online 29 April 2015)

Real time, density matrix based, time dependent density functional theory (TDDFT) proceeds through the propagation of the density matrix, as opposed to the Kohn-Sham orbitals. It is possible to reduce the computational workload by imposing spatial cutoff radii on sparse matrices, and the propagation of the density matrix in this manner provides direct access to the optical response of very large systems, which would be otherwise impractical to obtain using the standard formulations of TDDFT. Following a brief summary of our implementation, along with several benchmark tests illustrating the validity of the method, we present an exploration of the factors affecting the accuracy of the approach. In particular, we investigate the effect of basis set size and matrix truncation, the key approximation used in achieving linear scaling, on the propagator unitarity and optical spectra. Finally, we illustrate that, with an appropriate density matrix truncation range applied, the computational load scales linearly with the system size and discuss the limitations of the approach. © 2015 AIP Publishing LLC. [<http://dx.doi.org/10.1063/1.4919128>]

I. INTRODUCTION

Linear scaling or $O(N)$ density functional theory (DFT), in which the computational workload scales linearly with the number of atoms in the system N , is now well established.¹ In the standard approach to DFT, diagonalisation of an eigenvalue equation, or alternatively the orthogonalisation of the Kohn-Sham states during minimisation of the energy, results in a severe computational bottleneck that limits the size of systems which can be studied. Working with the density matrix, upon which a truncation radius is applied, allows the computational workload to be made to scale linearly with N . Circumventing the size limitations of the standard approach in this manner allows vastly larger systems to be studied; for example, calculations have now been performed on millions of atoms,^{2,3} in comparison to the upper limit of around a thousand for the standard approach.

While density functional theory is a ubiquitous tool in the arsenal of the electronic structure theorist, it is limited to the study of *ground-state* properties. Extending DFT to the time domain results in its excited state counterpart, time dependent density functional theory (TDDFT). Linear response TDDFT (LR-TDDFT), as developed by Casida,⁴ again suffers from a computational bottleneck which forces it to scale poorly with system size. LR-TDDFT requires the solution of an eigenvalue equation for a matrix written in the space of electron-hole pairs, which ostensibly scales as poorly as $O(N^6)$. In practice, this scaling can be reduced, through efficient implementation and methods employing the Liouville-Lanczos approach, to be as low as $O(N^3)$.^{5,6} For small systems, LR-TDDFT is

computationally feasible and has been widely used, while for larger systems, the scaling renders it unsuitable. It is also worth noting that linear scaling density matrix based LR-TDDFT, avoiding the propagation of the density matrix, has also been recently demonstrated.⁷

An alternative approach to LR-TDDFT is the real time propagation of the time-dependent Kohn-Sham equations, pioneered by Yabana and Bertsch.⁸ Real time TDDFT (RT-TDDFT) proceeds by the construction of an effective Hamiltonian, followed by the direct propagation of the Kohn-Sham orbitals using this Hamiltonian. Assuming both the number of occupied states (N_{KS}) and the number of mesh points (N_M) scale linearly with system size, RT-TDDFT will scale with the number of atoms, N , as $N_{KS}N_M \sim N^2$. A significant prefactor in the form of the number of time steps and the computational effort for construction of the Hamiltonian exist, making this method unsuitable for systems of small size. However, the $O(N^2)$ scaling has made it the natural choice for tackling systems of large size, and a complementary partner to Casida's approach.

In a similar manner to $O(N)$ DFT, it is possible to improve upon the scaling of RT-TDDFT by propagating the density matrix, as opposed to propagating the Kohn-Sham orbitals directly. By applying a spatial truncation radius upon the density matrix, the computational workload can be reduced, opening up the possibility of studying excited states in large systems that cannot feasibly be examined with other methods. Although not widely employed, this approach has been demonstrated to scale linearly with system size and has been used to study several large systems, fullerene, sodium clusters, polyacetylene oligomers, carbon nanotubes, and silicon clusters to name a few.⁹⁻¹⁵

Several factors must be taken into consideration when employing this method; for example, the accuracy of results

^{a)}Electronic mail: ucapcor@ucl.ac.uk

^{b)}Also affiliated with UCL Satellite, MANA International Centre for Materials Nanoarchitectonics (MANA), National Institute for Materials Science (NIMS), 1-1 Namiki, Tsukuba, Ibaraki 305-0044, Japan.

produced will depend strongly on the range of truncation of the density matrix. Also when working in a non-orthogonal basis, as is the case in the CONQUEST code, the overlap matrix will be well-ranged. However, the inverse overlap, which features in the density matrix propagators, will not necessarily be. In order to ensure the unitarity of the propagation, the propagators must be carefully tested for matrix truncation errors, and little discussion on the effect of matrix truncation upon propagator unitarity have been presented elsewhere.

In this paper, we briefly summarize our implementation of RT-TDDFT in the CONQUEST code, for completeness, and confirm its reliability. We then present several tests probing the limitations of the method and factors affecting accuracy. In particular, we examine the effect of matrix truncation, the key approximation used in achieving linear scaling, on the unitarity of the propagators used and optical spectra generated.

II. COMPUTATIONAL APPROACH

Linear scaling approaches for excited state properties have existed for well over a decade (for a review see Ref. 16), with the first approach utilising the locality inherent in the density matrix and being carried out at the semi-empirical level.⁹ Subsequent efforts again all tend to employ the nearsightedness of the density matrix, with the first full linear scaling TDDFT being done by Yam *et al.*^{17,18} Our approach follows that of Yam *et al.* closely, with a few differences; most notably, we choose not to perform the orthogonalisation procedure via the Cholesky decomposition and rather work in our non-orthogonal basis. As mentioned, linear scaling approaches to calculating the excited state properties in the frequency domain have also been presented, by Yokojima *et al.*,^{9,11} and more recently by Zuehlsdorff *et al.* in the ONETEP code.⁷ It is also worth noting that an approach for calculating the unoccupied Kohn-Sham states, via a basis optimisation approach which is also linear scaling, has also been implemented in the ONETEP code.¹⁹

In a similar vein to the standard approaches to TDDFT in the time and frequency domain, the reformulations using the density matrix can be viewed as complementary to one another. The frequency domain approach is suitable for calculating the lowest optical excitations in the system, but if the density matrix response involves higher excitations it will not be suitable. While the real-time density matrix approach employed here and by Yam *et al.* calculates the full optical spectrum, it has a significant prefactor in the form of the number of time steps needed for the numerical integration.

In this section, we briefly give an overview of the approach in our non-orthogonal basis set, and in Sec. III we illustrate the effect of the basis set on the results and the reliability of the method with several tests on small molecules.

A. Density matrix RT-TDDFT

Rather than working with the conventional single particle Kohn-Sham orbitals, CONQUEST works with the density matrix written in a separable form in terms of a localised

basis of support functions $\phi_{i\alpha}$,

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{i\alpha, j\beta} \phi_{i\alpha}(\mathbf{r}) \mathbf{K}_{i\alpha, j\beta} \phi_{j\beta}(\mathbf{r}') \quad (1)$$

where $\phi_{i\alpha}$ is the α th support function centred on atom i . Support functions are a non-orthogonal basis set of localised orbitals and have an overlap matrix given by

$$\mathbf{S}_{\alpha, \beta} = \int \phi_{i\alpha}(\mathbf{r}) \phi_{j\beta}(\mathbf{r}) \mathbf{d}\mathbf{r}. \quad (2)$$

Linear scaling behaviour can be obtained through applying a spatial cutoff on the density matrix. Beyond this cutoff radius, the matrix elements are set to zero which, along with the spatial limitation of the support functions, ensures that the number of non-zero density matrix elements increases linearly with system size (for a fuller overview of the CONQUEST code, see Ref. 20).

RT-TDDFT is now well established,⁸ and implementations of density matrix RT-TDDFT have been reported elsewhere.^{16,18} Rather than employing an orthogonalisation procedure via a Cholesky or Löwdin decomposition, which will increase the range of the sparse matrices and is done elsewhere, we work in our non-orthogonal basis. Expanding the time-dependent Kohn-Sham equations in this basis of non-orthogonal support functions, in the instance where the support functions are stationary with time, gives

$$i \frac{\partial}{\partial t} \mathbf{c}(t) = \mathbf{S}^{-1} \mathbf{H} \mathbf{c}(t) \quad (3)$$

and

$$i \frac{\partial}{\partial t} \mathbf{c}^\dagger(t) = -\mathbf{c}^\dagger(t) \mathbf{H} \mathbf{S}^{-1} \quad (4)$$

which describe the time dependence of the coefficients of our basis set expansion, $\mathbf{c}(t)$. This allows us to write the quantum Liouville equation of motion for our auxiliary density matrix \mathbf{K} in the non-orthogonal support function basis,

$$i \dot{\mathbf{K}} = \mathbf{S}^{-1} \mathbf{H} \mathbf{K} - \mathbf{K} \mathbf{H} \mathbf{S}^{-1}. \quad (5)$$

The formal solution to this equation can be expressed as

$$\mathbf{K}(t) = \mathbf{U}(t, t_0) \mathbf{K}(t_0) \mathbf{U}^\dagger(t_0, t) \quad (6)$$

where $\mathbf{U}(t, t_0)$ is a propagator satisfying both

$$\mathbf{c}(t) = \mathbf{U}(t, t_0) \mathbf{c}(t_0) \quad (7)$$

$$i \frac{\partial}{\partial t} \mathbf{U}(t, t_0) = \mathbf{S}^{-1} \mathbf{H} \mathbf{U}(t, t_0). \quad (8)$$

Expressing the propagator in integral form, we have

$$\mathbf{U}(t, t_0) = \mathcal{T} \exp \left\{ -i \int_{t_0}^t \mathbf{d}\tau \mathbf{S}^{-1} \mathbf{H}(\tau) \right\} \quad (9)$$

where \mathcal{T} is the time ordering operator. Evolution of the system for a total time, $T = n\Delta t$, may be carried out piecewise in smaller intervals, allowing us to express the total evolution operator as the product of small time operators,

$$\mathbf{U}(t, t_0) \approx \prod_{n=0}^{N-1} \mathbf{U}((n+1)\Delta t, n\Delta t) \quad (10)$$

where

$$\mathbf{U}(t + \Delta t, t) = \exp[-i\mathbf{S}^{-1} \mathbf{H}(t) \Delta t]. \quad (11)$$

Evolution of the time dependent system is then reduced to the problem of approximating the propagator $\mathbf{U}(t + \Delta t, t)$. Two approximations exist in the definition of $\mathbf{U}(t + \Delta t, t)$, first that of approximating the matrix exponential $\exp(\mathbf{A})$ and second the exact form of the matrix for which we wish to calculate the exponential. There are several methods for calculating the exponential of a matrix,²¹ here we use the simplest approximation, a Taylor expansion,

$$\exp(\mathbf{A}\Delta t) = \mathbf{I} + \sum_{n=1}^{\infty} \frac{(\mathbf{A}\Delta t)^n}{n!}. \quad (12)$$

Similarly, there are many different approaches for deciding which matrix exponential to use as a propagator. Three approximations have been implemented: the so called *exponential-midpoint* propagator (EM), the *enforced time-reversal symmetry* (ETRS) propagator, and the fourth order Magnus (M4) propagators, all of which are taken from the work of Marques *et al.*²² on RT-TDDFT propagators and are briefly described in our non-orthogonal basis for completeness.

The exponential midpoint propagator approximates the $\mathbf{U}(t + \Delta t, t)$ by the exponential taken at $\tau = t + \Delta t/2$,

$$\mathbf{U}_{EM}(t + \Delta t, t) = \exp \left\{ -i\mathbf{S}^{-1}\mathbf{H} \left(t + \frac{\Delta t}{2} \right) \right\}. \quad (13)$$

Implicitly enforcing time-reversibility, such that propagating forward from t and backwards from $t + \Delta t$ by $\Delta t/2$ produces the same result, provides the so called *enforced time-reversal symmetry* method,

$$\begin{aligned} \mathbf{U}_{ETRS}(t + \Delta t, t) = & \exp \left\{ -i\frac{\Delta t}{2}\mathbf{S}^{-1}\mathbf{H}(t + \Delta t) \right\} \\ & \times \exp \left\{ -i\frac{\Delta t}{2}\mathbf{S}^{-1}\mathbf{H}(t) \right\}. \end{aligned} \quad (14)$$

Using the Magnus operator, the exponential solution to Schrödinger equation for a time-dependent Hamiltonian may be written as²³

$$\mathbf{U}_{M4}(t + \Delta t, t) = \exp \{ \mathbf{M}_{G4} \}, \quad (15)$$

where \mathbf{M}_{G4} is an infinite series of integrals providing an exact solution. Truncating this expansion to fourth order and approximating the integrals using Gauss-Legendre points as in Ref. 22 give in our non-orthogonal basis,

$$\begin{aligned} \mathbf{M}_{G4} = & -i\frac{\Delta t}{2} [\mathbf{S}^{-1}\mathbf{H}(t_1) + \mathbf{S}^{-1}\mathbf{H}(t_2)] \\ & - \frac{\sqrt{3}\Delta t^2}{12} [\mathbf{S}^{-1}\mathbf{H}(t_2), \mathbf{S}^{-1}\mathbf{H}(t_1)], \end{aligned} \quad (16)$$

where $t_{2,1} = t + [1/2 \pm \sqrt{3}/6]\Delta t$.

It is important to note the presence of the inverse overlap matrix \mathbf{S}^{-1} in these propagators, and again consider that while the overlap matrix will be well-ranged and suitable for truncation, the inverse overlap is not necessarily so. We therefore need to carefully test the sparsity of the product $\mathbf{S}^{-1}\mathbf{H}$ and its effect on the unitarity of our propagators.

B. Linear response

The idea behind extracting optical transitions from the linear response of a system to an external electric field is well known.^{8,15} Propagating in real time provides direct access to the time-dependent charge density and therefore the electronic response to external fields. Applying a time dependent external electric field polarised along axis j ,

$$\delta v_{ext}(\mathbf{r}, t) = -\mathbf{E}_j(t) \cdot \mathbf{r}$$

allows us to examine the time-dependent response of the system. Application of this electric field will produce an induced time-dependent dipole moment,

$$\mathbf{P}(t) = \mathbf{P}(0) - \int d\mathbf{r} n(\mathbf{r}, t) \mathbf{r}. \quad (17)$$

As an example of the calculated response of a system to an applied electric field, Figure 1 illustrates the induced dipole response of a benzene molecule on application of a field with a Gaussian time profile, centred at $t = 0$.

Access to the time-dependent dipole moment allows us to calculate the time dependent polarisability,

$$\alpha_{ij}(\omega) = \frac{\int dt e^{i\omega t} P_i(t)}{\int dt e^{i\omega t} E_j(t)}.$$

The imaginary part of the polarisability is directly proportional to the absorption cross section, $\sigma(\omega)$, and the experimentally observed strength function, $S(\omega)$,

$$S(\omega) = \frac{2\omega}{\pi} \text{Im} \left(\frac{1}{3} \text{Tr}(\alpha_{\mu j}(\omega)) \right). \quad (18)$$

As noted by Tsolakidis *et al.*, the approach satisfies the f-sum rule and the integration of the strength function over energy gives the number of electrons, which may be used as a measure of the completeness of the basis set.¹⁵

Density matrix RT-TDDFT therefore has the potential to be an extremely useful tool for theoretically predicting the electronic absorption spectra of large system.

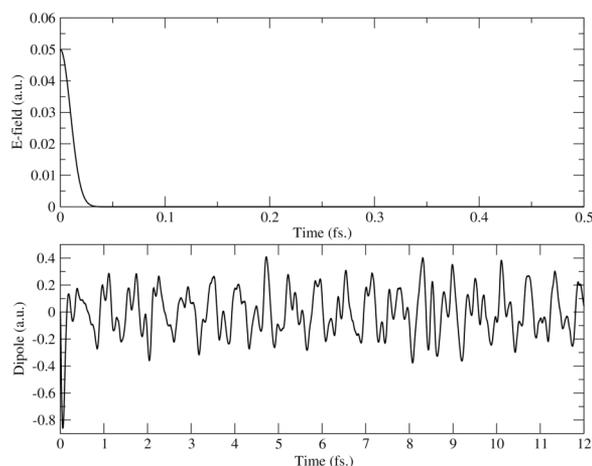


FIG. 1. Applied electric field and induced dipole moment for a benzene molecule ($\Delta t = 0.03$ a.u. ≈ 0.00073 fs).

III. SMALL MOLECULES

In order to verify that our implementation is correct, we have performed tests on several systems for which the electronic transitions have been studied experimentally and theoretically elsewhere, allowing us to make direct comparisons. For this purpose, we have chosen four small molecules (carbon monoxide, methane, ethylene, and benzene) and used our implementation to calculate the optical absorption spectra within the TDLDA approximation.

Meaningful comparison of our results with experiment requires the identification the electronic transitions to which the peaks in our calculated absorption spectra correspond. As we have mentioned in Casida's approach, information about electronic transitions is inherently produced, while in RT-TDDFT it is not.

It is often possible to identify the corresponding transition by examining the polarisation and energy of peaks and comparing to that of optically allowed transitions experimentally. Where possible, in order to more confidently assign peaks of our calculated absorption spectra to particular electronic transitions, we have followed the procedure in Ref. 26 whereby a sinusoidal electric field tuned to a particular excitation mode is applied. A resulting electronic resonance is set up, allowing us to examine the difference between ground state charge density and excited state charge density and thereby infer the electronic transition.

A. Basis sets

Our support functions are expanded in a basis of numerical orbitals, in this case pseudo-atomic orbitals, which are generated following the approach of the Siesta code.³² These PAOs are eigenfunctions of the atomic pseudopotentials with a confinement energy shift used to determine a radial cutoff for the orbitals, beyond which they are zero. This confinement energy provides a single parameter to define the cutoff radii for different orbitals and is the energy each orbital obtains on being confined by an infinite potential to a particular radius. It is clear that a minimal basis with which ground state properties are accurately reproduced will generally not be satisfactory for calculating excited state properties, and therefore we illustrate the basis set dependence of two selected transitions for the C_2H_4 molecule.

Multiple orbitals per angular momentum channel can be used (multiple- ζ), with the shape of multiple orbitals

determined by a split norm procedure.³² This procedure uses a parameter to define the norm of a numerical orbital outside some radius where they match the tail of the first zeta PAO, and within this radius, they vary smoothly to the origin. Subtracting this numerical orbital from the original PAO gives the multiple-zeta orbital. Of course, it is possible to define these radii by hand and fine tune the basis set. In addition to multiple zeta, polarisation orbitals can be included within the basis set and are obtained by solving the same pseudo-atomic problem but with an applied electric field.

We use the notations SZ, 2Z, 3Z, and 4Z to describe single zeta, double zeta, triple zeta, and so on. Similarly, we describe the number of polarisation orbitals included in the basis by SZP, SZ2P, and SZ3P (one, two, and three polarised orbitals, respectively).

To first gauge the effect of varying our basis set on the results, we have performed calculations on the ethylene molecule with varying numbers of PAOs and two different confinement energies. The basis sets have been generated with a confinement energy of 1 meV and 5 meV, resulting in confinement radii of 4.93 and 4.24 Å for the carbon atoms, respectively, and radii of 4.77 and 4.21 Å for the hydrogen atoms, respectively. The total run time was 14.51 fs (600 a.u.) with a time step of ~ 0.0242 fs (0.1 a.u.). The results can be seen in Table I.

Calculated energies for the $\pi \rightarrow 3s$ transition show a wide variation with basis set choice, while the $\pi \rightarrow \pi^*$ valence transition varies less. This is in line with expectation, given the more diffuse nature of the Rydberg transition, we would expect its description to require larger basis. The effect of systematically increasing the number of basis functions is to improve our results with respect to that of the reference values. Similarly, increasing the cutoff radii, by reducing the confinement energy, tends to improve the quality of the result. This is to be expected, as increasing the size of our basis set, while systematically increasing the range, will maximise the variational degrees of freedom available to describe our time dependent density matrix.

However, our values are still far from those computed elsewhere, and we find generally that for small molecules, it is essential to use a large basis with multiple extended polarisation orbitals in order to produce results in line with other works. In addition, we find that fine tuning the radial cutoffs by hand, as opposed to using the confinement energy and split norm procedure, can allow us to improve the quality of our results for small molecules.

TABLE I. Basis set dependence of calculated TDLDA transition energies (eV) for first valence ($\pi \rightarrow \pi^*$) and Rydberg ($\pi \rightarrow 3s$) excitations for the C_2H_4 molecule.

Transition	Basis set								Reference 24	Experiment ²⁵
	2ZP	2Z2P	3ZP	3Z2P	4ZP	4Z2P	5ZP	5Z2P	3ZP	
1 meV										
$\pi \rightarrow \pi^*$	7.84	7.62	7.73	7.62	7.67	7.62	7.67	7.62	7.45	8.0
$\pi \rightarrow 3s$	8.43	7.95	7.78	7.67	7.46	7.40	7.46	7.29	6.69	7.11
5 meV										
$\pi \rightarrow \pi^*$	7.82	7.73	7.75	7.69	7.70	7.68	7.67	7.67	7.45	8.0
$\pi \rightarrow 3s$	10.64	8.03	7.88	7.76	7.57	7.51	7.46	7.45	6.69	7.11

TABLE II. Comparison of calculated TDDA transition energies for small molecules with other values and experiment. Conquest results obtained with 5Z4P basis sets, with the exception of benzene (2Z2P).

Molecule	Transition	RT-TDDFT (eV)	Reference	Experiment
CO	$\sigma \rightarrow \pi^*$	8.17	8.20 ²⁷	8.51 ²⁸
CH ₄	T ₂ → 3s	9.22	9.27 ²⁴	9.70 ²⁹
C ₂ H ₄	$\pi \rightarrow \pi^*$	7.48	7.45 ²⁴	8.00 ²⁵
C ₆ H ₆	$\pi \rightarrow \pi^*$	6.87	~6.90 ³⁰	6.90 ³¹

B. Small molecule results

Exhibited in Table II are the calculated transitions for our four test molecules. In the case of the smallest molecules (carbon monoxide, ethylene, and methane), a hand tuned 5Z4P basis set is employed, while for benzene, the result is obtained using a 2Z2P basis with a 5 meV confinement energy (all the calculations satisfy the f-sum rule to >94%). Also presented in Figure 2 are the optical absorption spectra for the benzene and carbon-monoxide molecules, along with the experimental data.

We can see a strong agreement between our results and that of other studies, giving us confidence in our implementation. Very good agreement is exhibited between the calculated benzene absorption spectra and the experimental values using

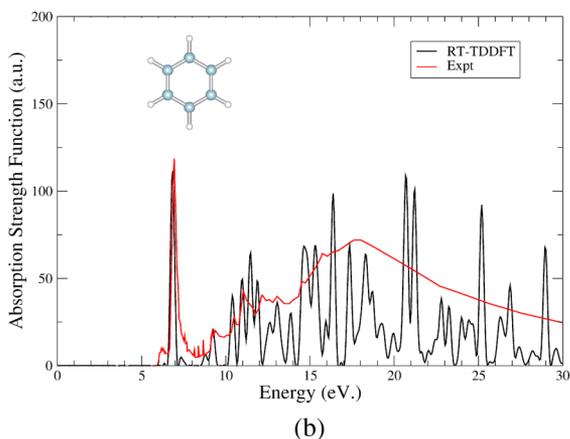
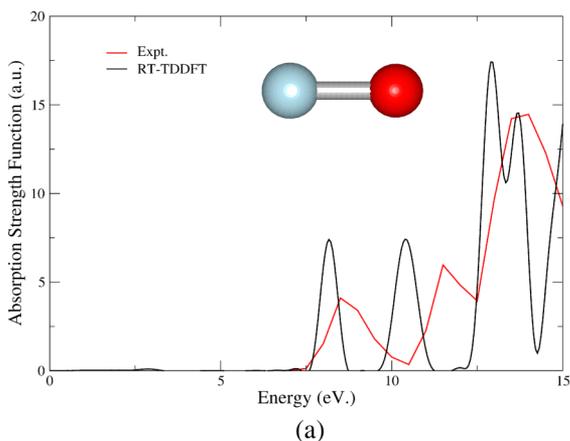


FIG. 2. (a): Absorption strength function for carbon monoxide from RT-TDDFT and experiment. Experimental data taken from Ref. 28. (b) Absorption strength function for Benzene from RT-TDDFT. Experimental data taken from Ref. 31.

a reasonably modest 2Z2P basis set. This highlights the point that for larger molecules, we have generally found that the need for large hand tuned basis sets, as is necessary for the smaller molecules, is reduced. Typically, results in agreement with those in the literature and experiment are found using smaller basis sets, a point that is important to bear in mind, given the context of linear scaling methods.

IV. PROPAGATOR UNITARITY

Having demonstrated the correctness of our implementation and explored the influence of basis sets, we now turn to our main concern, the effects of localisation in linear scaling methods on the accuracy of results.

We wish the total charge in our system to remain stable, and in order for this to be the case, the propagators must be unitary with respect to the non-orthogonal basis set,

$$U^\dagger U - I = 0, \quad (19)$$

where U is our propagator matrix and I is the identity matrix.

From our approximation for the matrix exponential, Eq. (12), it can be shown that, if it were exact, our propagators would indeed exhibit this property. However, as it is impossible for us to store an infinite sum on our computer, we must truncate our Taylor expansion at some point. Doing so will introduce errors, with two factors affecting the scale of the break from unitarity; the time step and the number of terms in our summation. While we can extend our expansion arbitrarily and reduce the time step arbitrarily, we wish to avoid excess computational expense by keeping the expansion as small as possible and the time step as large as possible within some acceptable margin of accuracy. We can directly examine the unitarity of our propagators through Eq. (19).

A. Time-step dependence

As a test, we have examined the extent of the break from unitarity for a range of time-steps and number of terms in the matrix exponential expansion. We have used a small molecule for the purpose, benzene, with a small applied electric field perturbation with a Gaussian profile centered on $t = 0$.

Exhibited in Figure 3, we can see the dependence on simulation time step of the propagator unitarity, with the obvious trend being that as the time step is reduced, the propagator approaches unitarity. We can see that even for time steps up to ~ 0.15 a.u., the propagator maintains its unitarity to a high degree (similar results were obtained for each of the propagators). The corresponding effect on the charge conservation can be seen in Figure 4 and, as expected, we see that as the time step increases, the conservation of charge deteriorates with the propagation eventually becoming unstable for large timesteps. While the maximum permissible timestep will depend on the system under study, we found that generally a timestep of 0.06 a.u. or below provided satisfactory charge conservation.

The form of our propagators requires the extrapolation of the Hamiltonian matrix to some unknown point beyond the current time t , H_+ . As suggested by Marques *et al.*²² in order

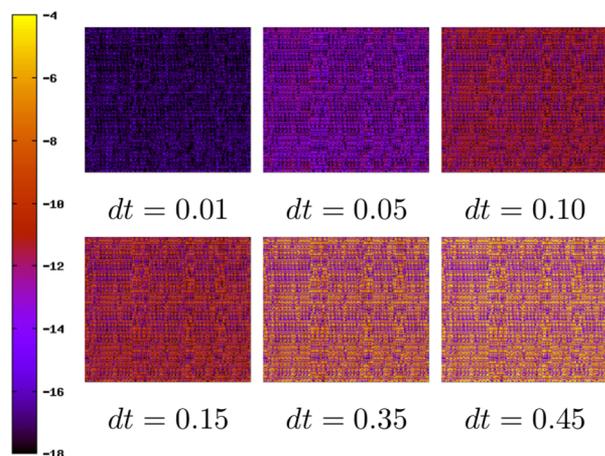


FIG. 3. Plot of the absolute values of matrix $U^\dagger U - I$ (on a base 10 log scale), illustrating the propagator unitarity for the exponential midpoint propagator, for varying time step sizes (in a.u.). The system studied is a single benzene molecule, and the matrix is shown at the end of a 10 a.u. run.

to minimise errors, it is possible to carry this procedure out self-consistently. In our case meaning that we propagate $K(t)$ to $K(t + \Delta t)$ based on an extrapolated Hamiltonian. We then construct a new Hamiltonian matrix $H(t + \Delta t)$ using $K(t + \Delta t)$. H_+ can then be interpolated from Hamiltonian matrices for times up to and including $(t + \Delta t)$, and the whole procedure is iterated until some self-consistency criteria is obtained. Generally speaking, this procedure is performed three times in the early stages of a run, following a perturbation, and reduces to two as the run progresses. The effect of not performing this self-consistency procedure on the charge conservation can be seen in Figure 4. While the self-consistency cycle is found to improve the charge conservation, in reality for small time steps, the difference in charge conservation and calculated properties is not found to be significant enough to warrant the extra computational load of constructing the Hamiltonian matrix several times per time-step. As a compromise, we enforce the self-consistency only for a small number of steps (~ 50 – 100) at the beginning of a run, typically, when our external electric field is applied for the study of the linear response and the external perturbation is largest.

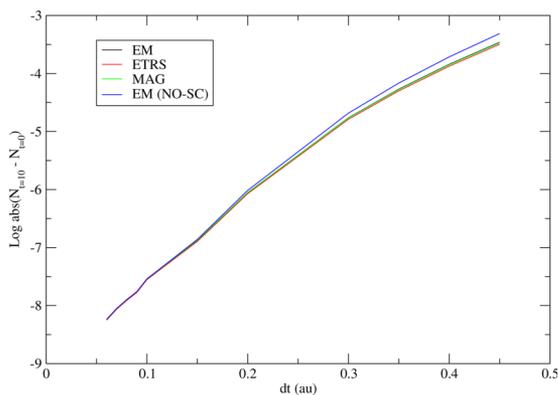


FIG. 4. Variation in total charge (on a base 10 log scale) with time step size, following a 10 a.u. run for benzene using all three propagators. Also included is charge variation for the EM propagator without the self-consistent propagator step (see text for details).

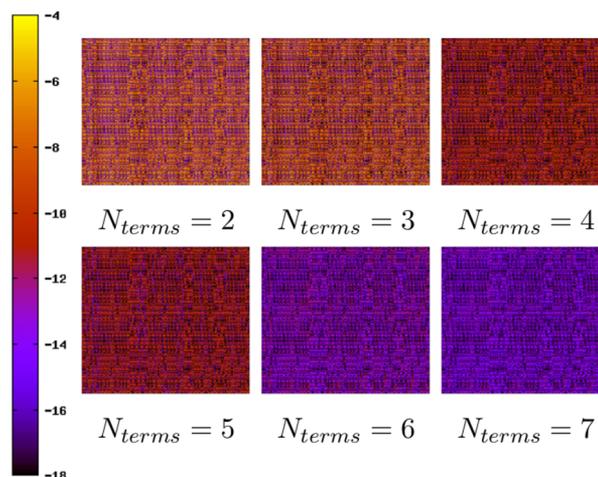


FIG. 5. Plot of the absolute values of the matrix $U^\dagger U - I$ (on a base 10 log scale), illustrating the propagator unitarity for the exponential midpoint propagator, for differing number of terms in the Taylor expansion for our propagator. The system studied is a single benzene molecule, and the matrix is shown at the end of a 10 a.u. run ($dt = 0.04$ a.u.).

A significant point to note is that little difference is exhibited between the calculated results using each of the three propagators in terms of charge conservation, and in general, we have found this to be the case. It is reported that for systems with strongly time-dependent Hamiltonians, the fourth order Magnus propagator, U_{M4} , is advantageous,²² but for our present work, this is not the case and we have opted for the simplest exponential midpoint propagator throughout.

B. Matrix exponential truncation

The effect of truncating the Taylor expansion used to evaluate the matrix exponential on the unitarity of the propagator can be seen in Figure 5. We see that reducing the number of terms reduces the unitarity of the propagator, as expected. Looking at Figure 6, the convergence of the charge conservation with the number of terms in the exponential expansion can be seen. We find that we reach good convergence with six terms included in the expansion, and we opt for this level of accuracy throughout the remainder of the paper.

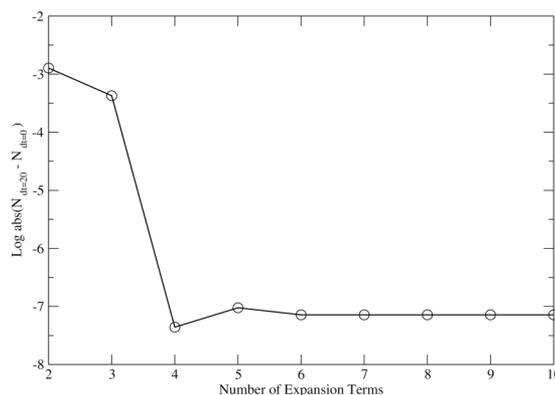


FIG. 6. Absolute variation in total charge (on a base 10 log scale) with the number of terms in our matrix exponential expansion, following a 20 a.u. run for benzene using the EM propagator with a time step of 0.04 a.u.

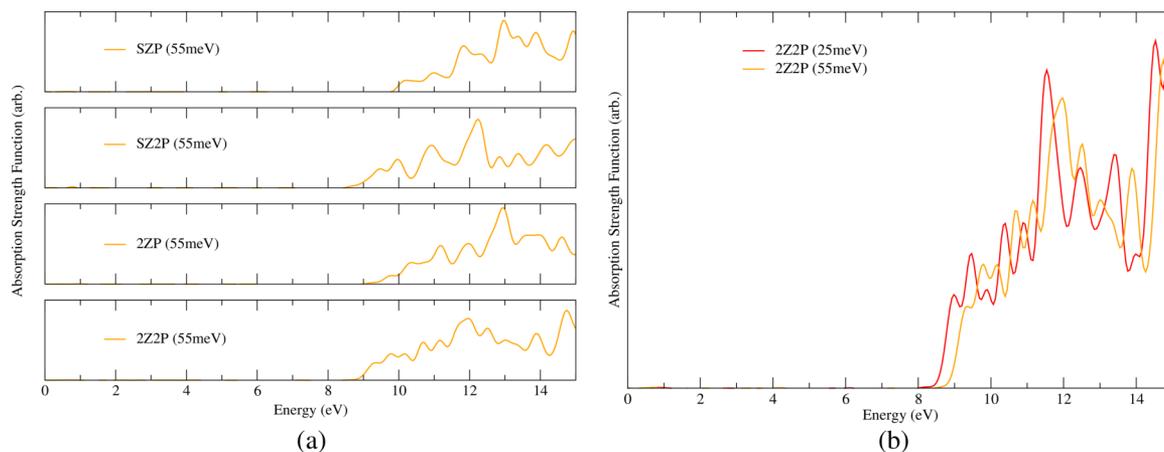


FIG. 7. Basis set variation of the calculated alkane optical absorption spectra. (a) Effect of increasing the number of PAOs in the basis set and (b) the effect of extending the radii of the basis functions are shown.

V. ALKANE MOLECULES: TESTING MATRIX TRUNCATION EFFECTS

In this section, we perform calculations on long chain alkane molecules. Our aim is to examine the effect of matrix truncation on the propagation of the density matrix and propagator unitarity, along with the computational scaling with system size.

As a first step, we calculate the absorption spectra for the C₁₁H₂₄ molecule for several different basis sets using the generalised gradient PBE functional³³ (all further calculations in this section are performed with this functional), and the results can be seen in Figure 7. Experimentally as the length of the alkane carbon chain increases, the absorption onset is found to reduce, and the reported adsorption onset for C₁₀H₂₂ is ~ 175 nm³⁴ (~ 7.1 eV). We see that as the number of PAOs in the basis set is increased the calculated absorption onset approaches this value. Particularly noticeable is the change of the absorption energy caused by the addition of polarisation orbitals. Similarly, a significant shift is induced by extending the range of the PAOs (a variation from 55 meV to 25 meV in the confinement energy extends the radii of the carbon and hydrogen basis sets by ~ 0.35 Å and 0.33 Å, respectively). This is understandable, given that the first transitions in the alkane molecules are reported as being Rydberg in character,³⁴ we would expect the addition of more diffuse PAOs to improve the description of these excitations. Given the well documented difficulties of TDDFT to accurately describe Rydberg transitions,³⁵ and given that this is not our aim in any case, we proceed to carry out our tests with the SZP and SZ2P basis sets generated using a confinement energy of 55 meV (radial cutoff for the PAOs is 3.31 Å and 3.12 Å for carbon and hydrogen, respectively).

Yam *et al.* have previously studied the long chain alkanes within the linear scaling excited state regime,¹⁸ calculating the absorption onset at around 8 eV for C₄₀H₈₂ with the LDA functional. However, little discussion of the effects of matrix truncation on propagator unitarity has been presented elsewhere.

A. Propagator truncation

The use of a basis of non-orthogonal atomic orbitals requires the inverse overlap matrix for our propagation (indeed,

this matrix is required for ground state calculations in any case), as seen in Eq. (11). In order to compute the inverse overlap matrix, Conquest uses Hotelling's method,³⁶ however, for poorly conditioned overlap matrices, computing the inverse overlap matrix can prove difficult. In our current implementation of TDDFT, the atoms remain stationary and so, therefore, does the overlap matrix. Therefore, we have also included the possibility of computing the inverse overlap with the SCALAPACK routines. Although the scaling will not be linear, computing the inverse overlap in this way makes only a relatively small contribution to our total TDDFT runtime, as we only calculate the inverse overlap once at $t = 0$.

While it is apparent that the overlap matrix will be sparse, allowing it to be truncated, the inverse of a sparse matrix will not in general be sparse itself. We have therefore tested the effect of truncating both the S^{-1} matrix and the $S^{-1}H$ matrix on the propagation. Figure 8 shows the average absolute error in the matrix elements of S^{-1} and the $S^{-1}H$ matrices caused by truncation (the error in S^{-1} elements given is the average of the elements of the $S^{-1}S-I$ matrix and the error in the $S^{-1}H$ is calculated with the values from an untruncated S^{-1} matrix).

As the range of the matrices increases, the error caused by the truncation converges towards zero, as we expect. The S^{-1} matrix converges less quickly than the $S^{-1}H$ matrix,

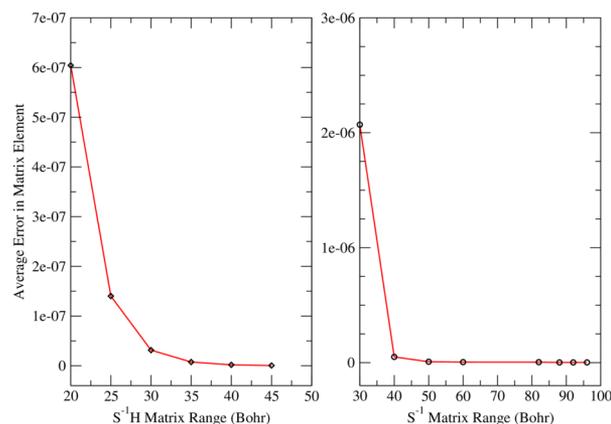


FIG. 8. Average absolute error in the $S^{-1}H$ (left) and S^{-1} (right) matrix elements with matrix range for the C₄₇H₉₆ molecule. SZP basis set is used, generated with a 55 meV confinement potential.

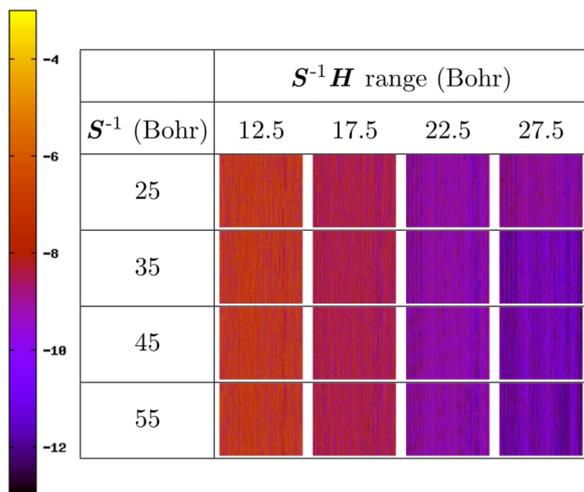


FIG. 9. Plot of the absolute values of the matrix $U^\dagger U - I$ (on a base 10 log scale), illustrating the propagator unitarity for differing truncation ranges of the S^{-1} and $S^{-1}H$ matrices for the $C_{47}H_{96}$ molecule.

indicating that it is more dense than the $S^{-1}H$ matrix. The effect truncation of these matrices has on the unitarity of the propagators can be seen in Figure 9. We see that the unitarity converges as the $S^{-1}H$ range increases, and the propagators are converged with a range of around ~ 22.5 - 27.5 Bohr. This indicates that the $S^{-1}H$ matrix is indeed sparse, while the S^{-1} matrix is less so, and we can safely truncate it. It is important to not that we do not explicitly use the S^{-1} in our propagators, only the $S^{-1}H$ matrix. Although it makes sense to truncate the S^{-1} matrix, given that we are truncating $S^{-1}H$ and that the Hamiltonian matrix is sparse. We can see this by noting that the unitarity of the propagator in Figure 9 is also well converged for each of the truncation ranges imposed on the inverse overlap.

As additional atoms are added the Hamiltonian matrix, overlap matrix, and the inverse overlap will vary. Increasing the system size may therefore affect the ranges of these matrices. While we only use the $S^{-1}H$ matrix in our calculation, comparison of the density of both matrices has been included. We have tested this effect by fixing the S^{-1} and $S^{-1}H$ ranges at 30 and 35 Bohr, respectively, and examined the error in the truncated $S^{-1}H$ matrix with system size with the results shown in Figure 10. We see that the error changes slightly on increasing system size but converges as the size

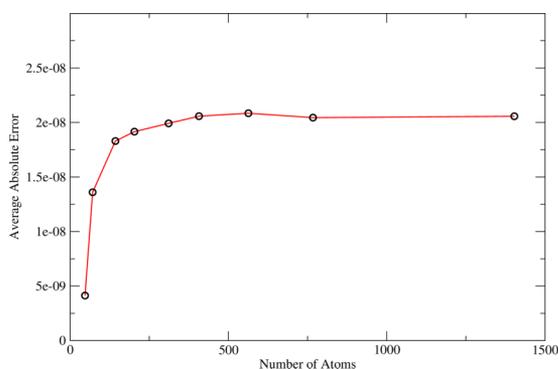


FIG. 10. Average absolute error in the $S^{-1}H$ matrix elements with system size.

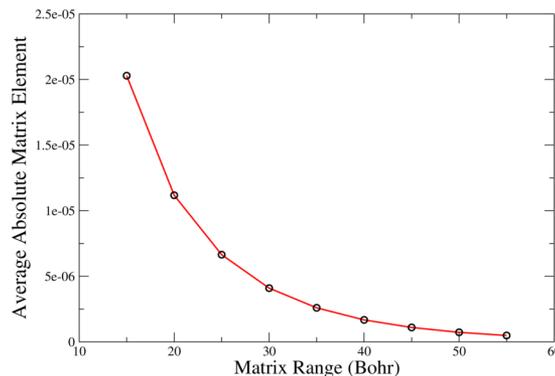


FIG. 11. Average value of the $U^\dagger U - I$ matrix with $S^{-1}H$ matrix range for the $C_{103}H_{208}$ molecule calculated with a SZ2P basis set.

increases. Consequently, the propagator unitarity was found to exhibit the same trend. This illustrates that the $S^{-1}H$ is well-ranged, irrespective of system size, allowing us to impose a cutoff radii on both of these matrices. In effect, this ensures that as the system size increases, the computational load can be made to scale linearly.

Similarly, increasing the number of basis functions will directly affect the overlap matrix and consequently the inverse overlap and the propagator. In order to gauge the extent of this effect, we have examined the $C_{103}H_{208}$ molecule with a larger basis set (SZ2P as opposed to SZP). Exhibited in Figure 11 is the absolute value of the $U^\dagger U - I$ matrix with $S^{-1}H$ matrix truncation range. Despite the larger number of basis set functions, we see that the $S^{-1}H$ matrix is still well-ranged, although the range is wider when compared to the SZP results of Figure 9, and again a truncation will lead to a computational load that scales linearly with system size.

A further point to note is that it is possible to avoid the use of the inverse overlap matrix in the TDDFT propagation altogether. Yam *et al.* have employed a Cholesky orthogonalisation scheme to bypass the need for the inverse overlap.¹⁸ However, using this scheme requires the inverse of the Cholesky decomposition, and it is not apparent that it will be more sparse than the inverse overlap. It is possible that this scheme might improve the calculation of the propagator, as the orthogonalised Hamiltonian may be more localised than our $S^{-1}H$ matrix. Calculating the Cholesky decomposition can be made to scale linearly, and implementation of this alternative method has already begun in order to contrast the two approaches. However, the parallelisation of Cholesky inversion is difficult given the Conquest matrix storage, and inversion of the overlap matrix remains important.

B. Density matrix truncation, scaling and limits

Finally, we examine the effect of truncating the density matrix and have performed calculations generating spectra for the $C_{47}H_{96}$ molecule at varying truncation radii, R_{Cut} , of the density matrix. Typically for ground state calculations, a suitable typical density matrix truncation range is around 16-20 Bohr. The results can be seen in Figure 12, and generally, we find that as the density matrix cutoff increases the spectra tend to converge, as expected, with higher lying states requiring a larger cutoff to converge. We can see from the comparison of

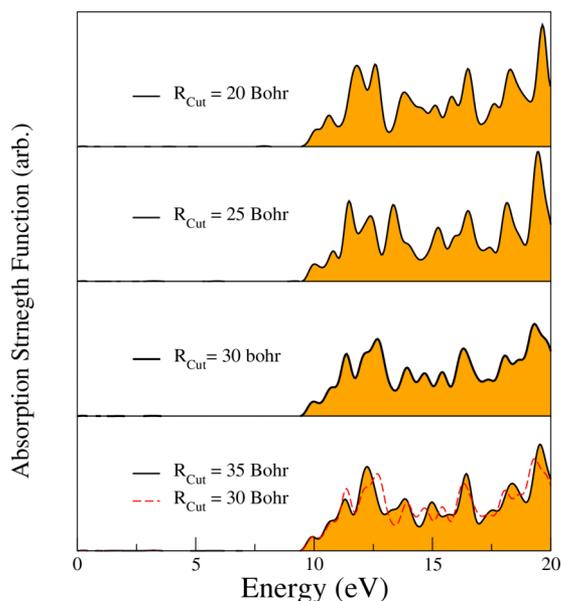


FIG. 12. K matrix truncation radii dependence: spectra generated for the $C_{47}H_{96}$ molecule at varying density matrix cutoff radii (total run time of 400 a.u. at a time step of 0.05 a.u.).

$R_{\text{Cut}} = 30$ and $R_{\text{Cut}} = 35$ that there is good agreement for the initial transitions, as well as the general shape of the spectra.

Applying this $R_{\text{Cut}} = 35$ Bohr cutoff (along with a cutoff of 35 Bohr on the $S^{-1}H$ matrix) we can examine the computational scaling with system size, with the results exhibited in Figure 13. Clear linear scaling of the computational workload up to well over 1000 atoms is exhibited, illustrating the potential power of the method.

Finally, a few comments on the limits of the approach must be made. TDDFT for long-range charge transfer is well known to be poorly described by local and semi-local functionals.³⁷ While we have employed LDA and GGA functionals here, linear scaling exact exchange has also been recently implemented in the Conquest code, allowing the use of non-local functionals with this approach in the future.

While the near-sightedness principle dictates that the ground-state density matrix is exponentially localised for well gapped systems, there is no formal justification for the localisation of the response density matrix. As noted in

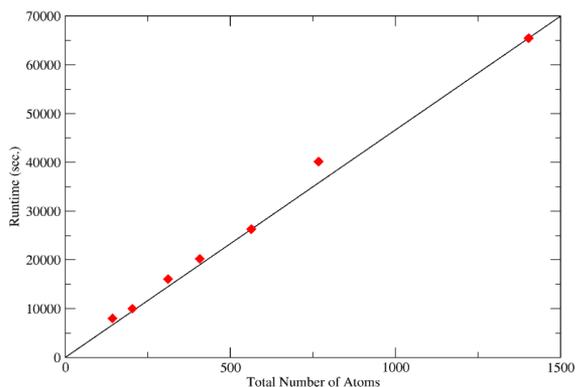


FIG. 13. Computational TDDFT run time versus system size for long chain alkane molecules. The system was run with a timestep of 0.05 a.u. for a total time of 10 a.u. A matrix truncation range, $R_{\text{Cut}} = 35$ a.u., has been applied.

Ref. 7, for systems with well localised excitations, it would be expected that the response density matrix could be truncated safely and linear scaling achieved, while for systems with delocalised excitations, this will not be the case.

VI. CONCLUSIONS

We have outlined our implementation of real-time time dependent density functional theory in the Conquest $O(N)$ code. We have demonstrated the soundness of the implementation through benchmark tests for small molecules and also discussed the effect of basis set and system sizes on the results.

$O(N)$ approaches utilise the density matrix, as opposed to working directly with Kohn-Sham orbitals, providing a route to the linear scaling computational time with system size by its truncation. We have discussed the range of our propagator matrices for an alkane chain test system, and the implications of this matrix truncation on the unitarity of the propagation. Similarly, we have examined the effect of truncating the density matrix on the calculated optical absorption spectra, showing that the range required is much more extended than that required for converged ground state properties. Nevertheless, we have shown that accurate linear scaling TDDFT calculations are practical. While the impact of localisation cutoff in the charge density matrix on these TDDFT calculations is a topic warranting further study, we have shown that in truncating these matrices at a suitable point, we obtain a computational load that increases linearly with system size. This offers a complementary approach to the usual Casida linear response approach: linear response TDDFT is well suited to relatively small systems, while linear scaling RT-TDDFT offers a viable method for studying excitations in large systems. We have shown linear scaling beyond 1000 atoms, and 10 000+ atoms are perfectly practical with the excellent parallel scaling available in Conquest.

ACKNOWLEDGMENTS

C.O'R. is supported by the MANA-WPI project and D.R.B. was funded by the Royal Society. We thank Umberto Terranova for useful discussions. This work made use of the facilities of ARCHER, the UK's national high-performance computing service, which is provided by UoE HPCx Ltd at the University of Edinburgh, Cray, Inc., and NAG Ltd, and funded by the Office of Science and Technology through EPSRC's High End Computing Programme. Calculations were performed at HECToR through the UKCP Consortium. The authors acknowledge the use of the UCL Legion High Performance Computing Facility, and associated support services, in the completion of this work.

¹D. R. Bowler and T. Miyazaki, "O (n) methods in electronic structure calculations," *Rep. Prog. Phys.* **75**, 036503 (2012).

²D. R. Bowler and T. Miyazaki, "Calculations for millions of atoms with density functional theory: Linear scaling shows its potential," *J. Phys.: Condens. Matter* **22**, 074207 (2010).

³J. VandeVondele, U. Bortnik, and J. Hutter, "Linear scaling self-consistent field calculations with millions of atoms in the condensed phase," *J. Chem. Theor. Comput.* **8**, 3565–3573 (2012).

- ⁴M. E. Casida, *Recent Developments and Applications of Modern Density Functional Theory* (Elsevier, Amsterdam, 1996).
- ⁵L. Kronik, A. Makmal, M. L. Tiago, M. M. G. Alemany, M. Jain, X. Huang, Y. Saad, and J. R. Chelikowsky, "PARSEC—The pseudopotential algorithm for real-space electronic structure calculations: Recent advances and novel applications to nano-structures," *Phys. Status Solidi B* **243**, 1063–1079 (2006).
- ⁶D. Rocca, R. Gebauer, Y. Saad, and S. Baroni, "Turbo charging time-dependent density-functional theory with Lanczos chains," *J. Chem. Phys.* **128**, 154105 (2008).
- ⁷T. J. Zuehlsdorff, N. D. M. Hine, J. S. Spencer, N. M. Harrison, D. J. Riley, and P. D. Haynes, "Linear-scaling time-dependent density-functional theory in the linear response formalism," *J. Chem. Phys.* **139**, 064104 (2013).
- ⁸K. Yabana and G. F. Bertsch, "Time-dependent local-density approximation in real time," *Phys. Rev. B* **54**, 4484–4487 (1996).
- ⁹S. Yokojima and G. Chen, "Time domain localized-density-matrix method," *Chem. Phys. Lett.* **292**, 379–383 (1998).
- ¹⁰S. Yokojima and G. Chen, "Linear scaling calculation of excited-state properties of polyacetylene," *Phys. Rev. B* **59**, 7259–7262 (1999).
- ¹¹S. Yokojima and G. Chen, "Linear-scaling localized-density-matrix method for the ground and excited states of one-dimensional molecular systems," *Chem. Phys. Lett.* **300**, 540–544 (1999).
- ¹²S. Yokojima, D. Zhou, and G. Chen, "Linear-scaling computation of ground state with time-domain localized-density-matrix method," *Chem. Phys. Lett.* **302**, 495–498 (1999).
- ¹³W. Liang, S. Yokojima, and G. Chen, "Generalized linear-scaling localized-density-matrix method," *J. Chem. Phys.* **110**, 1844–1855 (1999).
- ¹⁴W. Liang, S. Yokojima, D. Zhou, and G. Chen, "Localized-density-matrix method and its application to carbon nanotubes," *J. Phys. Chem. A* **104**, 2445–2453 (2000).
- ¹⁵A. Tsolakidis, D. Sánchez-Portal, and R. M. Martin, "Calculation of the optical response of atomic clusters using time-dependent density functional theory and local orbitals," *Phys. Rev. B* **66**, 235416 (2002).
- ¹⁶C. Yam, Q. Zhang, F. Wang, and G. Chen, "Linear-scaling quantum mechanical methods for excited states," *Chem. Soc. Rev.* **41**, 3821–3838 (2012).
- ¹⁷C. Y. Yam, S. Yokojima, and G. Chen, "Localized-density-matrix implementation of time-dependent density-functional theory," *J. Chem. Phys.* **119**, 8794–8803 (2003).
- ¹⁸C. Yam, S. Yokojima, and G. Chen, "Linear-scaling time-dependent density-functional theory," *Phys. Rev. B* **68**, 153105 (2003).
- ¹⁹L. E. Ratcliff, N. D. M. Hine, and P. D. Haynes, "Calculating optical absorption spectra for large systems using linear-scaling density functional theory," *Phys. Rev. B* **84**, 165131 (2011).
- ²⁰C. Goringe, E. Hernández, M. Gillan, and I. Bush, "Linear-scaling dft-pseudopotential calculations on parallel computers," *Comput. Phys. Commun.* **102**, 1–16 (1997).
- ²¹G. Moore, "Orthogonal polynomial expansions for the matrix exponential," *Linear Algebra Appl.* **435**, 537–559 (2011), Special Issue: Dedication to Pete Stewart on the occasion of his 70th birthday.
- ²²A. Castro, M. A. L. Marques, and A. Rubio, "Propagators for the time-dependent Kohn–Sham equations," *J. Chem. Phys.* **121**, 3425–3433 (2004).
- ²³W. Magnus, "On the exponential solution of differential equations for a linear operator," *Commun. Pure Appl. Math.* **7**, 649–673 (1954).
- ²⁴N. N. Matsuzawa, A. Ishitani, D. A. Dixon, and T. Uda, "Time-dependent density functional theory calculations of photoabsorption spectra in the vacuum ultraviolet region," *J. Phys. Chem. A* **105**, 4953–4962 (2001).
- ²⁵V. J. Hammond and W. C. Price, "Oscillator strengths of the vacuum ultraviolet absorption bands of benzene and ethylene," *Trans. Faraday Soc.* **51**, 605–610 (1955).
- ²⁶S. K. Min, Y. Cho, and K. S. Kim, "Efficient electron dynamics with the planewave-based real-time time-dependent density functional theory: Absorption spectra, vibronic electronic spectra, and coupled electron-nucleus dynamics," *J. Chem. Phys.* **135**, 244112 (2011).
- ²⁷C. Hu, O. Sugino, and Y. Miyamoto, "Modified linear response for time-dependent density-functional theory: Application to Rydberg and charge-transfer excitations," *Phys. Rev. A* **74**, 032508 (2006).
- ²⁸W. Chan, G. Cooper, and C. Brion, "Absolute optical oscillator strengths for discrete and continuum photoabsorption of carbon monoxide (7–200 eV) and transition moments for the $x^1\sigma^+ \rightarrow a^1\pi$ system," *Chem. Phys.* **170**, 123–138 (1993).
- ²⁹M. G. Curtis and I. C. Walker, "Low-energy electron-impact excitation of methane, silane, tetrafluoromethane and tetrafluorosilane," *J. Chem. Soc., Faraday Trans. 2* **85**, 659–670 (1989).
- ³⁰K. Yabana and G. F. Bertsch, "Time-dependent local-density approximation in real time: Application to conjugated molecules," *Int. J. Quantum Chem.* **75**, 55–66 (1999).
- ³¹E. Koch and A. Otto, "Optical absorption of benzene vapour for photon energies from 6 eV to 35 eV," *Chem. Phys. Lett.* **12**, 476–480 (1972).
- ³²J. M. Soler, E. Artacho, J. D. Gale, A. García, J. Junquera, P. Ordejón, and D. Sánchez-Portal, "The siesta method for ab initio order-n materials simulation," *J. Phys.: Condens. Matter* **14**, 2745 (2002).
- ³³J. P. Perdew, K. Burke, and M. Ernzerhof, "Generalized gradient approximation made simple," *Phys. Rev. Lett.* **77**, 3865–3868 (1996).
- ³⁴E. A. Costner, B. K. Long, C. Navar, S. Jockusch, X. Lei, P. Zimmerman, A. Campion, N. J. Turro, and C. G. Willson, "Fundamental optical properties of linear and cyclic alkanes: Vuv absorbance and index of refraction," *J. Phys. Chem. A* **113**, 9337–9347 (2009).
- ³⁵M. E. Casida, C. Jamorski, K. C. Casida, and D. R. Salahub, "Molecular excitation energies to high-lying bound states from time-dependent density-functional response theory: Characterization and correction of the time-dependent local density approximation ionization threshold," *J. Chem. Phys.* **108**, 4439–4449 (1998).
- ³⁶A. H. R. Palser and D. E. Manolopoulos, "Canonical purification of the density matrix in electronic-structure theory," *Phys. Rev. B* **58**, 12704–12711 (1998).
- ³⁷A. Dreuw, J. L. Weisman, and M. Head-Gordon, "Long-range charge-transfer excited states in time-dependent density functional theory require non-local exchange," *J. Chem. Phys.* **119**, 2943–2946 (2003).