

Hybrid Time-Dependent Density Functional Theory in CASTEP: a dCSE Project

Part 2: Atomic Forces in Excited Electronic States

Dominik B. Jochym

STFC Rutherford Appleton Laboratory

February 13, 2012

Abstract

Time-dependent density functional theory (TDDFT) is a formally exact method for calculating electronic excitations. In particular, linear-response TDDFT can be used to determine atomic forces in these excited states so that the structure and molecular dynamics can be probed. This report discusses the implementation of the above methods in CASTEP [2], one of the most popular electronic structure codes in use on HECToR. Making this functionality available opens up avenues for new science in a number of areas, including, but not limited to, photovoltaics and laser chemistry.

Contents

1	Introduction	2
2	Hybrid TDDFT	2
3	HPC I/O Bottlenecks	2
4	Excited State Forces	3
4.1	Geometry Optimisation of an Excited State	5
4.2	Molecular Dynamics	7
5	Closing remarks	7
6	Feature list of current code	8

1 Introduction

Time-dependent density functional theory (TDDFT) has become a well-established technique for modelling excited state properties in molecular systems, and has been implemented in several quantum-chemistry codes. An implementation of TDDFT in CASTEP [2] will give the UK electronic structure community an opportunity to address cutting-edge scientific problems in areas such as inorganic and organic photovoltaic materials, catalytic reactions at surfaces, light-emitting polymer materials for optical displays, and femtosecond laser chemistry. CASTEP in particular has an excellent record of use in high performance computing and is one of the most used codes on HECToR. The availability of TDDFT in CASTEP provides a platform for complementing a wide range of other property calculations at the same level of theory, within the same code. The implementation of hybrid functionals promises to address some of the limitations that have previously hindered the application of TDDFT to extended systems. [3, 1, 5]

This report is the second part of a 2-year dCSE project and focusses on an implementation for computing atomic forces in electronic excited states that are determined through linear response TDDFT in the Tamm-Dancoff approximation. For discussion on the iterative eigensolvers used in determining these excited states, please refer to the previous report available at <http://www.hector.ac.uk/cse/distributedcse/reports/castep02/>, and references therein.

2 Hybrid TDDFT

Milestone 1 - *Implement XC response kernel for hybrid functionals.*

The work for this milestone has already been discussed in the previous dCSE report (<http://www.hector.ac.uk/cse/distributedcse/reports/castep02/>). In summary, the Hartree-Fock exchange contribution to the excitation energy, and its operation on a vector have been implemented, allowing available hybrid exchange correlation functionals to be used as an adiabatic exchange correlation kernel for TDDFT.

3 HPC I/O Bottlenecks

As discussed in the previous dCSE report on TDDFT in CASTEP, the eigensolvers have been tested on HECToR. However, problems were encountered when attempting to produce photoabsorption spectra of even moderately sized molecular systems (~ 50 atoms). Specifically, restarts of a TDDFT calculation with a large number of states (> 100) were unfeasibly long, reading in the order of 100GB data. Each TDDFT state requires an entire wavefunction object and all lower states are required to enforce the orthogonality of the next states. The cause of the bottleneck was in the code used to read in CASTEP wavefunction files from disk. This problem was addressed as part of another dCSE project “Boosting the scaling performance of CASTEP: enabling next generation HPC for next generation science”. In summary, a restart of a particular TDDFT calculation on HECToR that previously took ~ 2.5 hours to read in 100 files of 800MB size was reduced to approximately 15 minutes, a factor of 10 improvement.

4 Excited State Forces

Milestone 2 - *Implement calculation of forces.*

Computation of the atomic forces is performed by computing the derivative of the total energy with respect to each atomic coordinate. Following Hutter's paper [4], computation of the forces needs to take in to account that a change in the Kohn-Sham orbitals also affects the response orbitals, complicating the derivative. The Lagrangian method used leads to three contributions to the forces: from the ground state orbitals, from the linear response TDDFT orbitals and from the so-called Z vector. The Z vector is "the matrix of Lagrange multipliers associated with the stationarity of the Kohn-Sham orbitals". [4] This vector is akin to a response wavefunction and requires the solution of a Sternheimer-like equation. The ground state force contribution is, of course, already implemented in CASTEP, so we will omit this term in the following discussion. Therefore, the contribution to the force on atom I from TDDFT has two terms

$$F_{\text{TDDFT}}^I = F_x^I + F_Z^I, \quad (1)$$

where F_x^I is the contribution from the TDDFT response orbital (referred to as x in Hutter's paper and as $\Phi^{\{-\}}$ in the previous dCSE report), and F_Z^I is the contribution from the Z vector. In Dirac notation (Hutter's formulation was written in a matrix element notation) the force terms are

$$F_x^I = \sum_i^{\text{occ}} \left\langle \Phi_i^{\{-\}} \left| \frac{\partial H^{\{0\}}}{\partial R_I} \right| \Phi_i^{\{-\}} \right\rangle - \sum_i^{\text{occ}} \left\langle \Phi_i^{\{0\}} \left| \frac{\partial H^{\{0\}}}{\partial R_I} \right| \sum_j^{\text{occ}} \left\langle \Phi_i^{\{-\}} \left| \Phi_j^{\{-\}} \right\rangle \Phi_j^{\{0\}} \right\rangle, \quad (2)$$

and

$$F_Z^I = \sum_i^{\text{occ}} \left\langle Z_i \left| \frac{\partial H^{\{0\}}}{\partial R_I} \right| \Phi_i^{\{-\}} \right\rangle + \left\langle \Phi_i^{\{-\}} \left| \frac{\partial H^{\{0\}}}{\partial R_I} \right| Z_i \right\rangle. \quad (3)$$

In the above equations, $H^{\{0\}}$ is the ground state Kohn-Sham Hamiltonian and R_I is the atomic coordinate of atom I . Note that the last term of equation 2 contains a rotation in bands of the ground state wavefunction by the matrix with elements $\left\langle \Phi_i^{\{-\}} \left| \Phi_j^{\{-\}} \right\rangle$. The operator $\partial H^{\{0\}}/\partial R_I$ was already implemented in CASTEP, although care was needed to correctly implement the non-local pseudopotential contribution.

All of the objects above are available after a TDDFT calculation, except for the Z vector. This is determined according to the Sternheimer-like Handy-Schaefer Z vector equation

$$(H^{\{0\}} - \varepsilon_i) |Z_i^*\rangle + P_c \delta V_{\text{SCF}}[n^{\{z\}}] \left| \Phi_i^{\{0\}} \right\rangle = |u_i\rangle, \quad (4)$$

where ε_i are the Kohn-Sham eigenvalues, the projector on unoccupied unperturbed states, P_c , and Hartree-exchange-correlation response potential, δV_{SCF} , are defined in the previous dCSE report on this project (<http://www.hector.ac.uk/cse/distributedcse/reports/castep02/>). The new electron density-like object $n^{\{x\}}$ is

$$n^{\{z\}}(\mathbf{r}) = \sum_i^{\text{occ}} Z_i(\mathbf{r}) \Phi_i^{\{0\}}(\mathbf{r}), \quad (5)$$

which can be computed using the existing first order response density routine for TDDFT. The wavefunction-like object u is

$$|u_i\rangle = P_c \left\{ \delta V_{\text{SCF}}[n^{\{x\}}] \left| \phi_i^{\{0\}} \right\rangle + \delta V_{\text{SCF}}[n^{\{-\}}] \left| \Phi_i^{\{-\}} \right\rangle + \sum_j^{\text{occ}} \left| \Phi_j^{\{-\}} \right\rangle \left\langle \Phi_j^{\{0\}} \left| \delta V_{\text{SCF}}[n^{\{-\}}] \left| \Phi_i^{\{0\}} \right\rangle + V_{\text{XC}}^{\{2\}}[n^{\{-\}}] \left| \phi_i^{\{0\}} \right\rangle \right\rangle \right\}. \quad (6)$$

Two new terms, $n^{\{x\}}$ and $V_{\text{XC}}^{\{2\}}$, are introduced. The density object is¹

$$n^{\{x\}}(\mathbf{r}) = \sum_i^{\text{occ}} \Phi_i^{\{-\}*}(\mathbf{r}) \Phi_i^{\{-\}}(\mathbf{r}) - \sum_i^{\text{occ}} \Phi_i^{\{0\}*}(\mathbf{r}) \sum_j^{\text{occ}} \left\langle \Phi_i^{\{-\}} \left| \Phi_j^{\{-\}} \right\rangle \Phi_j^{\{0\}}(\mathbf{r}), \quad (7)$$

where, as in equation 2, there is a rotation in bands of the ground state wavefunction in the last term. The second order exchange correlation potential is obtained from the third (functional) derivative of the exchange correlation energy with respect to the electron density

$$V_{\text{XC}}^{\{2\}}[n^{\{-\}}](\mathbf{r}) = \iint d\mathbf{r}d\mathbf{r}' \frac{\delta E_{\text{XC}}}{\delta n(\mathbf{r})\delta n(\mathbf{r}')\delta n(\mathbf{r}'')} \Big|_{n^{\{0\}}} n^{\{-\}}(\mathbf{r}')n^{\{-\}}(\mathbf{r}''). \quad (8)$$

We evaluate this using a two-point finite difference of the first order exchange correlation response potential, $\delta V_{\text{SCF}}[n^{\{-\}}]$.

The Z vector equation (4) can be solved using the existing CASTEP density functional perturbation theory (DFPT) E-field solvers [7], replacing the $id\Phi_i^{\{0\}}/dk_\alpha$ term with u_i . We employed the variational E-field solver in our prototype code, extending the functionality to take a general ‘‘perturbation’’, rather than just the electric field. In order to keep the results consistent between two runs, it was found that the convergence tolerance for the E-field minimiser had to be tightened considerably for this use, from 10^{-5} to 10^{-8}\AA^3 . This is required because the initial trial wavefunctions are randomised. Such an extreme tolerance may be problematic because as a general rule of thumb, for the result to be accurate from the E-field solver, the ground state wavefunction needs to be converged to the square of the E-field tolerance. The full ramifications of this have yet to be explored.

The computed forces were tested for correctness by performing a numerical derivative of the excitation energy with respect to an atomic displacement, i.e. computing the excitation energy at two small atomic displacements and performing a two-point derivative. The resulting forces matched to within the error expected from the numerical derivative. This test was successfully performed for a set of small molecules. The hydrogen molecule was a particularly useful test case as the pseudopotential used was purely local. This allowed for separate testing of local and non-local pseudopotential contributions to the $\partial H^{\{0\}}/\partial R_I$ operator.

¹This is a corrected form of Hutter’s equation 44, where the minus sign preceding the second term was omitted.

4.1 Geometry Optimisation of an Excited State

Incorporating the TDDFT forces into the existing geometry optimiser required some careful consideration in terms of module dependencies. Specifically, the `geometry` module requires a total energy for the electronic state and the atomic forces for that state. Previously this was done only for the ground state, with the energy and forces being provided by the `electronic` and `firstd` modules, respectively. Through discussion with the CASTEP Development Group it was decided to have the `tddft` module called from `electronic` and `firstd` when the total energy or forces of an excited state is requested. However, the TDDFT code makes use of some of the higher level DFPT code in the `secondd` module, essentially using it in a way not envisioned in its initial code design, leading to some circular compilation dependencies. The `secondd` module (and one of its dependants, `magres_utils`) makes direct reference to routines in the `electronic` module. It was decided that the routines in question were more utilities in nature and would be moved from the `electronic` module to the low-level `algor` module, hence removing any circular dependencies.

A test calculation on the formaldehyde molecule was performed so that comparison with the results in Hutter’s paper could be made. Table 1 gives the comparative results. An exact match is not expected because Hutter’s work included the Martyna-Tuckerman method to decouple periodic images of the electrostatic potential [6], which is not implemented in CASTEP. Different pseudopotentials were also used. Those details aside, the geometries generated by CASTEP compare very well with the previously published results. It should be noted that in the case of the first excitation, the non-planar configuration is only obtained when the planar symmetry of the ground state geometry is broken by adding some random noise to the positions of the initial configuration. This is because the geometry optimisation algorithm can stop at stationary points in the energy landscape, not necessarily a global minimum. One expects the excited state energy landscape to be complicated even for small molecules therefore some care will have to be taken when running excited state geometry optimisations so as not to converge on metastable configurations. Developing more sophisticated geometry optimisation methods are outside the scope of this dCSE.

Table 1: Geometry optimised structures of formaldehyde (CH_2O), bond lengths in Å and angles in degrees, where Φ is the out-of-plane angle. The ground state calculations used the PBE exchange correlation functional and TDDFT calculations used its adiabatic counterpart for the exchange correlation kernel.

Electron configuration	C–O	C–H	\angle HCH	Φ
Ground state (Hutter)	1.211	1.118	116.1	0
Ground state (This work)	1.208	1.120	116.1	0
1st singlet excitation (Hutter)	1.308	1.103	116.8	30.0
1st singlet excitation (This work)	1.305	1.103	117.5	28.9
2nd singlet excitation (Hutter)	1.204	1.115	119.0	0
2nd singlet excitation (This work)	1.203	1.114	120.4	0

In terms of performance, carrying out an excited state optimisation is more expensive than that for the ground state for a number of reasons. The obvious

one is the added expense of solving the TDDFT eigenvalue problem at each trial geometry. Another is that the energy tolerance used to determine the ground state wavefunction needs to be tighter when used in conjunction with a linear response method. This is in order to guarantee the accuracy of the response wavefunctions and therefore increases the number of iterations to obtain the ground state. The computation of the TDDFT forces is also not near negligible compared to solving the eigenvalue problem, as is the case in the ground state. To obtain the Z vector, an extra calculation equivalent to a one-shot, single-direction E-field DFPT run is required, comparable in computation time to the ground state calculation in our tests. However the single longest step is still obtaining the TDDFT eigenvalues/vectors. This step could be reduced by using the TDDFT eigenvectors from the previous geometry iteration as the starting point for the current one. Of course, care must be taken to preserve normalisation and orthogonality to both the (new) ground state and between the eigenvectors themselves. An implementation of this reduced the number of iterations used by the TDDFT solver by a factor of at least two in our tests. Some stability can be gained by adding a small random component (0.001) to the previous eigenvector, thus ensuring the initial trial vector spans the required subspace. In our tests this did not increase the number of TDDFT iterations. Using the earlier formaldehyde case as an example, the time for the structure optimisation of the 2nd excited state in 7 BFGS iterations was reduced by 40% when compared to randomised initial trial eigenvectors.




Figure 1: Left: The ground state optimised structure of formaldehyde viewed from the top and side. Right: As for the left, but for the first singlet excitation. Carbon, oxygen and hydrogen atoms are coloured grey, red and white, respectively. The total energy for the first singlet excitation is 0.3 eV higher when the atoms are in the ground state configuration than when in the relaxed configuration.

4.2 Molecular Dynamics

Milestone 3 - *Implement Born-Oppenheimer molecular dynamics.*

With the code abstraction already employed in the prototype TDDFT geometry optimisation, Born-Oppenheimer molecular dynamics of an excited state is already in place, but has yet to be tested. A potential issue beyond the scope of this dCSE is that of excited state crossings that are almost certain to occur during the course of a calculation.

5 Closing remarks

The importance of working with a well structured and documented code cannot be emphasised enough. Throughout this project it was found that similarities with existing methods meant that code originally written for a specific purpose in CASTEP could be used more generally with little (or occasionally no) further code modification. This is thanks to the high coding standards used in the CASTEP source. In the cases where changes were required for existing code, the work was expedited greatly by the documentation in the subroutine headers and clear comments throughout.

Also invaluable during code development was the availability of a CVS mirror of the main CASTEP source. This allowed work to be done on a branch of the code base while being able to pull in changes from major version releases. Subsequent migration of the CASTEP repository to the Mercurial versioning system makes such a process even easier for the coder.

The modular nature of the code introduced during this dCSE means that improvements to the geometry optimisation methods from the recent LBFGS dCSE project can immediately be taken advantage of in TDDFT calculations. It is hoped that in adhering to CASTEP coding standards that at least some of the code developed during this dCSE will be used for unforeseen functionality in the future. As an extension to this project, it is already planned to integrate the TDDFT forces work with CASTEP's existing finite-displacement phonon code, allowing investigation of vibrational properties of systems in an electronic excited state.

A demonstration project based on polyfluorene is under way. This conjugated polymer is of interest as a material for organic light emitting diodes that can be colour tuned. Results will be reported in due course.

The code developed during this dCSE project will be merged with the main CASTEP source for a future major version release. Access to a beta of the TDDFT code is available to licensed CASTEP users upon request, contact Keith Refson at keith.refson@stfc.ac.uk.

6 Feature list of current code

- Calculation of singlet and triplet states in Tamm-Dancoff approximation for both open and closed-shell systems
- ‘Pure’ and hybrid-DFT adiabatic XC kernels
- Solvers: Conjugate gradient and block Davidson (both with preconditioning)
- G-vector parallel
- Optimisations for Γ -point
- Calculation checkpointing and restart
- Oscillator strengths (for computing spectra)
- Characterisation of eigenvectors by decomposing into KS orbitals
- Calculation of atomic forces using the Handy-Schaefer Z-vector method
- Structure optimisation of a chosen excited state
- Born-Oppenheimer molecular dynamics of a chosen excited state (untested)

This project was funded under the HECToR Distributed Computational Science and Engineering (CSE) Service operated by NAG Ltd. HECToR - A Research Councils UK High End Computing Service - is the UK’s national supercomputing service, managed by EPSRC on behalf of the participating Research Councils. Its mission is to support capability science and engineering in UK academia. The HECToR supercomputers are managed by UoE HPCx Ltd and the CSE Support Service is provided by NAG Ltd. <http://www.hector.ac.uk>

References

- [1] L. Bernasconi, M. Sprik, and J. Hutter, *Hartree-Fock exchange in time dependent density functional theory: application to charge transfer excitations in solvated molecular systems*, Chemical Physics Letters **394** (2004), no. 1-3, 141–146.
- [2] S.J. Clark, MD Segall, CJ Pickard, PJ Hasnip, MJ Probert, K. Refson, and MC Payne, *First principles methods using CASTEP*, Z. Kristallogr **220** (2005), no. 5-6, 567–570.
- [3] A. Dreuw, J.L. Weisman, and M. Head-Gordon, *Long-range charge-transfer excited states in time-dependent density functional theory require non-local exchange*, The Journal of Chemical Physics **119** (2003), 2943.
- [4] J. Hutter, *Excited state nuclear forces from the Tamm–Dancoff approximation to time-dependent density functional theory within the plane wave basis set framework*, The Journal of Chemical Physics **118** (2003), 3928.
- [5] A.F. Izmaylov and G.E. Scuseria, *Why are time-dependent density functional theory excitations in solids equal to band structure energy gaps for semilocal functionals, and how does nonlocal Hartree–Fock-type exchange introduce excitonic effects?*, The Journal of Chemical Physics **129** (2008), 034101.
- [6] G.J. Martyna and M.E. Tuckerman, *A reciprocal space based method for treating long range interactions in ab initio and force-field-based calculations in clusters*, The Journal of Chemical Physics **110** (1999), 2810.
- [7] K. Refson, P.R. Tulip, and S.J. Clark, *Variational density-functional perturbation theory for dielectrics and lattice dynamics*, Physical Review B **73** (2006), no. 15, 155114.