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

Dominik B. Jochym

STFC Rutherford Appleton Laboratory

July 8, 2010

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. Including the use of hybrid functionals promises to address some of the limitations that have previously hindered the application of TDDFT to extended systems. [3, 1, 7]

2 Testbed implementation of TDDFT

Milestone 1 - *Implement a straightforward scheme based on the existing DFPT module code in CASTEP to compute the electronic response to an external electric field of a set frequency. This will provide a reference calculation against which the later, more sophisticated calculations may be benchmarked. Test results on small molecular systems will be compared against previous calculations published in scientific literature. This work comprises stages 1 and 2 of the work plan.*

The above refers to equation 17 of Hutter's paper [6], which describes the first-order response[1] of electrons to an electric field at a particular frequency. The density functional perturbation theory implementation in CASTEP[9] already had much of the functionality required for this. Hutter's equation 17 is:

$$\left(H^{(0)} - \epsilon_i\right)\left|\Phi_i^{(\pm)}\right\rangle + P_c V^{(1)}(\pm\omega)\left|\Phi_i^{(0)}\right\rangle = \mp\omega\left|\Phi_i^{(\pm)}\right\rangle, \tag{1}$$

where $H^{(0)}$ is the ground state (Kohn-Sham) Hamiltonian, ϵ_i are the ground state Kohn-Sham eigenvalues for band index i, $|\Phi_i^{(\pm)}\rangle$ are the response wavefunctions to a perturbation of frequency ω, P_c is the projector on the subspace

[1] Throughout this document when referring to the first-order (or linear) response, we will simply use "response", for brevity.

of unperturbed unoccupied states, $|\Phi_i^{(0)}\rangle$ are the ground state Kohn Sham orbitals and $V^{(1)}(\pm\omega)$ is the response potential, containing contributions from the Hartree, exchange-correlation, and electric field perterbation terms. In the above we have neglected reference to electronic spin, for brevity. The reponse wavefunctions throughout this report are taken to be orthogonal to the occupied ground state, so there is an implied projector.

In the Tamm-Dancoff approximation[4], occupied–virtual contributions to equation (1) are disregarded but the virtual–ocupied ones are kept, under the assumption that the contribution from the former is small. This amounts to setting $|\Phi_i^{(+)}\rangle = 0$, so Hutter's equation 17 becomes

$$\left(H^{(0)} - \epsilon_i\right)\left|\Phi_i^{(-)}\right\rangle + P_c V^{(1)}(-\omega)\left|\Phi_i^{(0)}\right\rangle = \omega\left|\Phi_i^{(-)}\right\rangle, \tag{2}$$

which is the equation implemented for electric field response in CASTEP (with the right hand side normally zero). A chosen frequency, ω, can be set through the `excited_state_scissors` parameter.

In the `secondd` module, we have two different methods available for calculating the response. Namely, a variational solver, and a Green's function solver. The variational solver is more stable than the Green's function solver when approaching the first excitation energy. However, the variational solver cannot converge for values of ω above the first excitation energy, whether the polarisability is negative or otherwise.

We chose an isolated methane molecule for our test system. To get the energy of the excited states, we calculate the polarisability at a number of ω values, and extrapolate for the divergence. Extensive tests were performed to investigate convergence of the first excitation energy with respect to plane wave cutoff energy and size of supercell. To get the excitation energy to two decimal places, a cutoff energy of 750 eV[2] and a cell 21 Å^3 was required.

We found the first excitation energy to be 9.10 eV, using the LDA. This compares well with a literature value of 9.053 eV [8], where they used GAUSSIAN 98. As we are using the Tamm-Dancoff approximation, this value for the first excitation energy can be compared directly to that obtained by a direct calculation of the poles, covered in the next section.

3 Initial demonstration implementation of Hutter Solver

Milestone 2 - *This will comprise much of the framework of a releasable implementation of TDDFT using Hutter's published method. Issues of data distribution and parallel efficiency will not be addressed at this stage and an 'off-the-shelf' eigensolver will be used. Initial testing and debugging. Code will be benchmarked against a set of existing calculations on molecular systems chosen from scientific literature. This work comprises stages 3-5 of the work plan.*

Hutter reformulated the equations of time-dependent Hartree-Fock (TDHF) theory (applied to TDDFT with pure density functionals) such that they could be efficiently implemented in a plane-wave basis set. The TDHF equations are a

[2]The C_00 and H_00 norm-conserving pseudopotentials were used.

non-Hermitian eigenvalue equation, which gives the excitation energies directly

$$\begin{pmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^* & \mathbf{A}^* \end{pmatrix} \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix} = \omega \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix} \begin{pmatrix} \mathbf{X} \\ \mathbf{Y} \end{pmatrix}. \tag{3}$$

In the Tamm-Dancoff approximation, the $\mathbf{B}$ matrices are set to zero, which leads to a Hermitian eigenvalue problem $\mathbf{AX} = \omega\mathbf{X}$. It is this Hermitian eigenvalue problem that will be implemented for the "Hutter solver".

In the context of a plane-wave code, the vectors $\mathbf{X}$ and $\mathbf{Y}$ are the plane-wave coefficients of electronic wavefunctions, i.e. $|\Phi_i^{(\pm)}\rangle$ from above. In CASTEP, these are four-dimensional arrays, the dimensions being G-vectors (plane-waves), k-points, bands and spin. For a TDDFT calculation, one typically has $\sim 10^5 - 10^6$ G-vectors, $\sim 100 - 1000$ bands, a single k-point and spin of 1 or 2 dpending on whether the system is being treated as spin-degenerate. Of course, storing and diagonalising a matrix of this size is innapropriate, especially as one is only interested in the lowest few excited states. Instead we turn to iterative methods, where only the result of the operator on a vector is required.

Hutter defines this operator in equation 35 of his paper, where he chooses to separate it into two components, such that $\mathbf{A} = \mathcal{A} + \mathcal{B}$. The first contribution is from Kohn-Sham orbital energy differences

$$\mathcal{A} \left|\Phi_i^{(-)}\right\rangle = (H^{(0)} - \epsilon_i) \left|\Phi_i^{(-)}\right\rangle. \tag{4}$$

The action of this operator was already implemented for DFPT in CASTEP. The second contribution is

$$\mathcal{B} \left|\Phi_i^{(-)}\right\rangle = P_c \delta V_{\mathrm{SCF}}[n^{(-)}] \left|\Phi_i^{(0)}\right\rangle, \tag{5}$$

where $\delta V_{\mathrm{SCF}}[n^{(-)}]$ is the self-consistent reponse potential for a change in electron density

$$n^{(-)}(\mathbf{r}) = \sum_{i=1}^{N} \Phi_i^{(0)*}(\mathbf{r}) \Phi_i^{(-)}(\mathbf{r}), \tag{6}$$

and

$$\delta V_{\mathrm{SCF}}[n^{(-)}] = \int d\mathbf{r}' \left\{ \frac{1}{|\mathbf{r} - \mathbf{r}'|} + \frac{\delta^2 E_{\mathrm{XC}}}{\delta n(\mathbf{r}) \delta n(\mathbf{r}')} \right\} n(\mathbf{r}'), \tag{7}$$

which is the same as used in DFPT, but for a factor of 2 in the response density.

Given the effect of the $\mathbf{A}$ operator on a given wavefunction, an 'off-the-shelf' iterative eigensolver with a reverse communication interface can be used. By using a library routine, we can thoroughly test our implementation of the operator. We used two different solvers, namely ARPACK[3] and EA19, the latter being from the HSL (formerly the Harwell Subroutine Library)[4]. ARPACK is written in Fortran77 and implements the Arnoldi process. HSL-EA19 is written in F95 + TR 15581 and implements a Jacobi-conjugate preconditioned gradients scheme.

[3] http://www.caam.rice.edu/software/ARPACK/
[4] http://www.hsl.rl.ac.uk

3.1 Results

At the end of section 2, we calculated the first excitation energy of methane to be 9.10 eV by scanning through frequency. Calculating the lowest eigenvalue directly on the same system, we get a value of 9.08 eV. Table 2 compares the eigenvalues for a selection of molecules calculated using CASTEP and CPMD[5], where Hutter's method was originally implemented.

Molecule	State	CASTEP	CPMD
N_2	1	9.282	9.283
	2	9.282	9.283
	3	9.692	9.692
	4	10.259	10.251
	5	10.270	10.269
	6	10.270	10.269
	7	11.495	11.488
	8	11.495	11.488
H_2	1	9.988	9.997
	2	10.840	10.831
	3	11.024	11.014
	4	11.024	11.014
	5	11.403	11.392

Table 1: Eigenvalue comparison

4 Implementation of parallel distribution solver for true HPC use on HECToR

Milestone 3 - *Port of code to HECToR. More extensive testing, debugging and benchmarking against previous calculations. Demo calculations on larger systems than feasible in stage 2. This work comprises stages 6 and 7 of the work plan.*

While using stock serial eigensolvers was invaluable while developing the code for the operator, a parallel solver is required. HSL-EA19 currently has no parallel version available. ARPACK does have a parallel version, but it is not one that can be used with the existing parallel distribution schemes in CASTEP. We chose to implement two algorithms, a state-by-state conjugate gradients solver and a block Davidson solver, with help from Phil Hasnip (University of York and member of the CASTEP Development Group).

Hutter's formulation assumes that the special k-point Γ is used.[6] This means that the real-space representation of the wavefunctions can always be real (as opposed to complex). CASTEP already has Γ-point optimisations and we can take advantage of these also. Using a single k-point restricts our choice of

[5] http://www.cpmd.org

[6] It is reasonably straightforward to extend Hutter's derivation to a single arbitrary k-point. The complex representation of the response density in equation 6 (above) allows for this extension. The equations for multiple k-points are not trivial and would require further research, and is therefore beyond the scope of this dCSE.

existing parallel stategies to G-vector and bands. Our TDDFT code is currently based on the v4.4 release, which does not contain the bands-parallel code.

4.1 Comparison of Solvers

At this point, it would be beneficial to compare the serial performance of the four eigensolvers we have available. The table below gives the calculation time and number of times the TDDFT operator is applied to converge the lowest 20 eigenvalues of our H_2 test system, containing ~10,000 plane-waves. This trend

Solver	Time (s)	# of operations
ARPACK	171	756
EA19	195	848
Conjugate Gradients	291	1290
Block Davidson	114	520

in the performance is as expected from Tretiak's paper [10]. The state-by-state CG solver is not without its benefits, however, as it has very modest memory usage.

4.2 Parallel benchmark

For a parallel benchmark, we chose a larger molecular system in the form of Buckminster-Fullerene, C_{60}. We ran the calculations on HECToR (phase 2a) for 5 iterations of the Davidson solver, and used 16 cores as a baseline. See figure 1 below. We have used the prototype shared memory extension coded by Chris Armstrong (NAG) under a core CSE call. With 4-way SMP, a parallel efficiency of approximately 80% could be achieved with 256 processing elements. Initial tests with HECToR phase 2b gave us comparable calculation times only when using 4 cores per node (one per die), again with 4-way SMP.

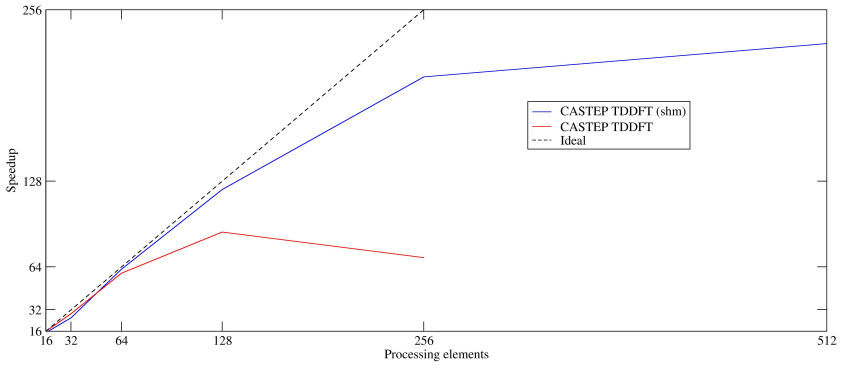

Figure 1: Parallel speedup for the C_{60} benchmark on HECToR phase 2a.

We plan to merge our TDDFT code with the main CASTEP branch in the near future. Minor modifications will be required for bands-parallel. Improve-

ments to the parallel scaling can be expected to be inline with those for the ground state calculation.

4.3 Demonstration calculation of an optical spectrum

Below is the optical spectrum produced from the first 100 eigenstates of C_{60}.

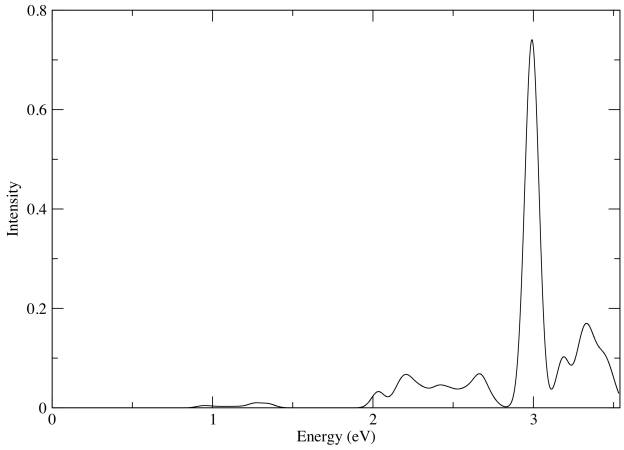

Figure 2: Optical spectrum of C_{60} generated from the first 100 excited states.

5 Hybrid TDDFT

When a hybrid functional is used in TDDFT, the contribution to the eigenvalue (excitation energy) is

$$-c_{\mathrm{HF}} \sum_{i}^{\mathrm{occ}} \sum_{j}^{\mathrm{occ}} \int \int d\mathbf{r} d\mathbf{r}' \frac{\phi_i^{(-)*}(\mathbf{r}) \phi_j^{(-)}(\mathbf{r}) \phi_j^{(0)*}(\mathbf{r}') \phi_i^{(0)}(\mathbf{r}')}{|\mathbf{r} - \mathbf{r}'|}, \tag{8}$$

where c_{HF} is the fraction of Hartree-Fock exchange prescribed by the approximation at hand. This expression is very similar to those already implemented in the CASTEP nlxc module. Some minor modifications were made to include the required response wavefunctions while keeping the number of FFTs at a minimum, with a modest storage overhead. Even so, calculations including Hartree-Fock exchange are often an order of magnitude more expensive than a ground state calculation. In our case, the expense comes from the double loop over bands. Such calculations will benefit greatly from parallel distribution over bands. Note that the implementation of the Hartree-Fock response term should also be applicable to DFPT, opening up the use of hybrid functionals in those calculations.

The response kernel for the hybrid PBE0 approximation has been implemented and tested. In doing so we came across a numerical instability that

seems to occur only for systems with so-called 'spurious' excited states [5]. The problem manifests in the production of negative eigenvalues.[7] The instability appears to arise from a subtle interplay or the Hartree-Fock response and the PBE exchange-correlation response when the response density is very small, but non-zero. The problem is circumvented by increasing the granularity of CASTEP's fine grid (to `fine_grid_scale=4.0`), which improves the quality of the numerical derivatives in the GGA response term.

While preparing for the implementation of the response terms for the B3LYP hybrid, it was discovered that the existing code used an incorrect calling sequence for the Becke 88 and LYP correlation routines. These routines are key components for the B3LYP implementation. The code was fixed and will be released in CASTEP version 5.5. The implementation of the B3LYP response terms is in progress.

Using purely Hartree-Fock exchange for TDDFT in the Tamm-Dancoff approximation is equivalent to CIS (configuration interaction with single substitutions). This function is also implemented.

5.1 Results

Below is a comparison of eigenvalues using the PBE0 functional.

Molecule	State	CASTEP	CPMD
N_2	1	9.342	9.343
	2	9.483	9.460
	3	9.483	9.460
	4	9.897	9.897
	5	9.897	9.897

Table 2: Eigenvalue comparison

6 Feature list of current code

- Calculation of singlet states in Tamm-Dancoff approximation
- 'Pure' and hybrid-DFT adiabatic XC kernel
- 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

[7] While this should never happen in the case of pure functionals, negative eigenvalues are known to occur in TD Hartree-Fock theory. It is unclear whether they can occur in hybrid TDDFT.

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] A.L. Fetter and J.D. Walecka, *Quantum Theory of Many-Particle Systems*, McGraw-Hill, New York, 1971.

[5] A. Heßelmann and A. Görling, *Blindness of the Exact Density Response Function to Certain Types of Electronic Excitations: Implications for Time-Dependent Density-Functional Theory*, Physical review letters **102** (2009), no. 23, 233003.

[6] 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.

[7] 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.

[8] A.R. Porter, O.K. Al-Mushadani, M.D. Towler, and R.J. Needs, *Electronic excited-state wave functions for quantum Monte Carlo: Application to silane and methane*, The Journal of Chemical Physics **114** (2001), 7795.

[9] 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.

[10] S. Tretiak, C.M. Isborn, A.M.N. Niklasson, and M. Challacombe, *Representation independent algorithms for molecular response calculations in time-dependent self-consistent field theories*, The Journal of chemical physics **130** (2009), 054111.