
ONETEP Tutorials

Release 0.0.1

ONETEP Collective

Aug 22, 2023

CONTENTS:

1	Introduction	1
2	Tutorial 1: Setting up Simple ONETEP Calculations	3
3	Tutorial 4: Geometry optimization	11
4	Tutorial 5: Analysis and visualization	19
5	Tutorial 6: Time-Dependent DFT	29
6	Tutorial 8: Implicit solvation, visualisation and properties: Protein-ligand free energy of binding for the T4 lysozyme	35
7	Tutorial 9: DFT+U on strongly correlated magnetic materials: A case study on antiferromagnetic Hematite	47
8	Tutorial 10: Simulation cell relaxation	57
9	Tutorial 11: Electrified electrode-electrolyte interfaces under potentiostatic control	61
10	Tutorial 12: Quantum embedding with (time-dependent) embedded mean-field theory: hydrogenation and excitations of pentacene	71
11	Tutorial 13: ASE ONETEP interface	79
12	PDF version of all tutorials	85
13	Indices and tables	87
	Bibliography	89
	Index	93

CHAPTER
ONE

INTRODUCTION

The current set of tutorials for ONETEP can be found here..

TUTORIAL 1: SETTING UP SIMPLE ONETEP CALCULATIONS

Nicholas D.M. Hine and Chris-Kriton Skylaris

2.1 Input files

Setting up a ONETEP job involves creating a main input file with the suffix **.dat** which contains all the required information to describe both the system and the parameters of the job. This requires the user to provide input in the form of keywords and blocks. Keywords are written in the form

keyword: value [unit]

For example, to specify that the task we wish the code to perform is a Single-Point energy calculation, we would add:

task : SinglePoint to our input file (note that capitalisation is irrelevant).

If we wish to specify a cutoff energy of 500 eV for our standard grid, we would add:

cutoff_energy : 500.0 eV

The value in eV's will be converted internally to atomic units (Eh in this case). If a keyword is not specified in the input file, it is given a default value which is intended to work across a broad range of systems. A full list of keywords and blocks, giving their meaning, syntax and default values, can be found on the ONETEP keyword database: <https://onetep.org/resources/keyword-database/>

Blocks are used to define the values of input parameters which need to contain multiple records, such as the definition of the unit cell. They take the form:

%block blockname

a1 a2 a3

b1 b2 b3

...

`%endblock blockname`

Most blocks tend not to have a meaningful default value, and must be specified if the related functionality is to be used. Comments can be added to input files using the `# ' or !'` characters. Anything after these characters on a given line will be ignored.

2.1.1 Setting up the Input File

We will start by running a simple job on a silane molecule SiH_4 . Create a working directory in which to run ONETEP

```
> mkdir silane
```

```
> cd silane
```

Create a new input file called `silane.dat` in your favourite text editor e.g.

```
> vi silane.dat &
```

You might like to put a comment at the top explaining what this input file is for e.g.

```
$ # $ Simple ONETEP input file for a silane molecule
```

The first thing is to specify the simulation cell.

The simplest choice is a cubic box with sides of about 40.0 bohr. Enter the 3-component cell vectors, one per line, between the `%block lattice_cart` and `%endblock lattice_cart` keywords.

Second, the atomic species need to be specified, in this case silicon and hydrogen.

This information needs to be provided between `%block species` and `%endblock species` keywords.

In this block, we need to specify five pieces of information per species, separated by spaces: (i) your symbol for the atomic species (this can be the same as the element symbol)

(ii) the element symbol itself

(iii) the atomic number Z

(iv) the number of NGWFs per atom

(v) the NGWF radius.

The number of NGWFs required can usually be judged from the symmetry of the atomic orbitals involved: In this case four for silicon and one for hydrogen will be adequate (can you think why?).

For this molecule, 6.0 bohr should be a reasonable starting point for the NGWF radii.

Each atomic species in our calculation needs a pseudopotential file. The pseudopotential files are specified between `%block species_pot` and `%endblock species_pot` keywords. You can use the `hydrogen.recpot` and `silicon.recpot` files from the ONETEP pseudopotentials *pseudo* directory. Copy them to your working directory now (or make a symbolic link).

Next, we need to specify the atomic positions, between `%block positions_abs` and `%endblock positions_abs` keywords.

There is one line per atom. Remember to use your symbol for the atomic species as defined in the species block. The coordinates are assumed to be given in Bohr unless specified otherwise. While it is not requirement in ONETEP that all the atoms should lie within the simulation cell, it is best (for visualisation purposes) to start by placing the silicon atom at the centre of the cell.

The Si-H bond length is about 2.76 bohr and silane is a tetrahedral molecule. The simplest way to work out the coordinates is to note that tetrahedral bonds can be chosen to lie along unit vectors

$(a, b, 0)$, $(-a, b, 0)$, $(0, -b, a)$ and $(0, -b, -a)$ where $a = p \sqrt{2/3}$ and $b = p \sqrt{1/3}$.

For example, the vector for the first Si-H bond is $(2.2535, 1.5935, 0.0000)$ Bohr. Add these offsets to the position of your silicon atom to create the SiH_4 molecule.

The last essential parameter to specify is the kinetic energy cutoff parameter for the psinc basis set. A reasonable value to start with is 300 eV. Use the `cutoff_energy` keyword and remember to specify the energy unit as well as the value.

2.1.2 Running the Job

Examine the output: if you have followed these instructions it should converge very quickly (8 iterations) to a total energy of around -6.1897Eh.

2.1.3 Convergence Convergence Convergence

Just as in any form of traditional DFT, we must ensure that our calculation results are converged with respect to the size of the basis. In ONETEP, convergence with basis size is controlled by a small number of parameters, with respect to which the total energy is variational. In this context, that means the total energy at a given value of the parameter will be an upper bound to the true, converged total energy, and increasing the parameter will monotonically decrease the total energy, which asymptotically tends to its converged value.

2.2 Cutoff Energy

The first parameter will be familiar to anyone who has carried out plane-wave DFT calculations: the cutoff energy. This specifies the kinetic energy of the maximum G-vector of the reciprocal-space grid, and therefore the spacing of the real-space grid. With a 40 bohr cell and a 300eV cutoff, ONETEP will have chosen a $48 \times 48 \times 48$ grid, hence a grid spacing of 0.833 Bohr. This may be too coarse: move your old output file to a new name (eg `SiH4.out_Ec300`) and try changing the cutoff energy to 500eV, then re-run the job script. You may wish to add `output_detail: VERBOSE` to your input file, to see exactly what grids are being used at each cutoff.

Comparing the two outputs, you should see that the total energy has decreased by around 0.03Eh

(nearly 1eV, or 0.2eV/atom). This suggests 300eV was too low initially. Try increasing the cutoff in steps of 100eV (You may wish to automate this, by having a loop in your job script in which the input file is updated and the job run for each update, if you are sufficiently familiar with bash scripting)

Plot the total energy (ET) as a function of cutoff energy. You should see a monotonic decrease in ET as a function of Ecut: try to evaluate at what value you think the total energy is converged to about 0.1eV/atom of its asymptotic limit. Note that the calculation time increases rapidly with cutoff energy, because the number of grid points in each FFTbox is growing rapidly with cutoff energy, and thus each FFT takes longer, so do not try going beyond around 1200eV.

In few cases in reality do we require strict convergence of the total energy. It is more usual that we require convergence of some measurable quantity such as a binding energy, which is based on energy differences.

In that case, we do not require the total energy to be converged, only the difference between total energies of very similar systems. This may converge much faster than the total energy itself, presuming the same species are present in both systems. Always consider what it is that you need converging before you start running enormous calculations!

2.2.1 NGWF radius

Next, we will investigate convergence with respect to the NGWF radius. Pick a value of cutoff energy for which you can perform reasonably fast calculations (say, 500.0 eV) and try increasing the NGWF radius from 6.0 to 10.0 in 1.0 Bohr steps. Plot the total energy against NGWF radius. Again, you should see a monotonic decrease. Note that above 6.0 bohr the FFT box is as large as the simulation cell, in a larger cell this would keep growing, and the calculation time would increase rapidly. Also, you should notice that the number of NGWF Conjugate Gradients iterations grows with the size of the localisation region, this is natural since with larger spheres there are more NGWF coefficients to simultaneously optimise.

You may also wish to try converging with respect to the number of NGWFs per atom (eg try 9 NGWFs on the Silicon). In some systems, notably crystalline solids, this can be crucial to achieving good convergence of the NGWFs themselves.

2.2.2 Kernel Cutoff

This SiH₄ system is too small to investigate convergence with respect to the cutoff of the density kernel. In larger systems truncation of the density kernel can be a good way to speed up the calculation. Indeed, asymptotically it is only by truncating the kernel that true 'linear-scaling' behaviour of the computational effort will be observed.

The kernel cutoff is controlled by the `kernel_cutoff` keyword. This defaults to 1000 Bohr (i.e. effectively infinite). Density kernel truncation should be used with a degree of caution: generally speaking, one would want to be able to run a full calculation for a fairly large system first, with an infinite cutoff, to establish a known baseline. Then, try decreasing the kernel cutoff from that point

and see what the effect is on the total energy, on the level of NGWF convergence (as measured by the NGWF RMS gradient), and on the computation time. If significant time savings can be achieved without trading in too much accuracy, it may be worthwhile to bring down the cutoff for all similar calculations. Proceed with care, though as calculations with a truncated kernel tend to converge in a less stable manner.

Crystalline Silicon

You may wish to try out also a calculation on a periodic solid.

As it is fairly well-behaved but illustrates some interesting concepts, let's try crystalline silicon, in the diamond (f.c.c.) structure. We will build a $2 \times 2 \times 2$ version of the 8-atom simple-cubic unit cell.

Copy your SiH_4 input to a new file (eg Si64.dat) in a new directory (eg SILICON) and remove the references to hydrogen from the species and species_pot blocks. Copy silicon.recpot to this directory as well. In the new input file, set the NGWF radius to 7.0 bohr, the number of NGWFs per atom to 4, and the cutoff energy to 600eV. Edit the cell side length so that it is $2 \times$ the lattice parameter of crystalline silicon in the LDA (around 10.1667 bohr). For reasons that will become clear if you read the last section of this tutorial, on Common Problems, also set `ngwf_cg_max_step`: 8.0 to prevent the CG line step being capped unnecessarily and `maxit_ngwf_cg`: 30 to terminate the NGWF CG after 30 iterations (in case it's not converging). Typing out the positions would be rather time-consuming and error-prone with 64 atoms in the cell, so use your favourite scripting/programming language (bash, awk, python, perl, etc would all be suitable or even C or FORTRAN) to write a list of the positions. You will need to repeat the basis (atoms at (0, 0, 0) and $(1/4, 1/4, 1/4)a$) at each of the positions of the f.c.c. lattice: (0, 0, 0), $(1/2, 1/2, 0)$, $(0, 1/2, 1/2)$, and $(1/2, 0, 1/2)$. Copy the result into your positions_abs block. An example input file for this job can be found on the tutorial web page.

The calculation should take around a few minutes. Feel free to stop it as soon as you see what is happening, since you will find that the calculation fails to converge: the RMS gradient remains stuck above the threshold for convergence. Likewise, the total energy will not converge to a fixed value. Make a copy of your output and modify the NGWF radius in the input file to 8.0 bohr and the number of NGWFs per Si atom to 9. This introduces NGWFs with d-like symmetry rather than just s and p, allowing much more variational freedom. You should now find the calculation converges nicely, but will take rather longer to run.

Now try activating `write_forces`: T to calculate the forces on each atom. All the forces should be small: in principle they are constrained by the symmetry of the crystal to be exactly zero.

However, you will see that they are not exactly zero because the symmetry of the system is broken by the psinc grid, which is not necessarily commensurate with the unit cell. However, in this small cell, it will not be possible to fix this as the number of points across the FFT box must be odd, and in a small cell the simulation cell and the FFT box coincide, so the number of points across the simulation cell must also be odd.

Adjust your script to write a $5 \times 5 \times 5$ supercell of the crystal (1000 atoms). Reduce the kernel cutoff to 25 bohr with `kernel_cutoff`: 25.0 and set the code to perform 1 NGWF iteration only `maxit_ngwf_cg`: 1 (otherwise the calculation would take longer to run, you can try this if you have

time). To restore the symmetry, adjust the `psinc_spacing` value to be a divisor of the supercell length such that an exact number of grid points spans each unit cell of the crystal (pick a value which gives an effective cutoff energy close to 600.0 eV so as not to increase the run time too much) and set off the 1000 atom job. This should not take too long on 32 cores.

Beyond around 500 atoms, the calculation should be into the so-called 'linear-scaling' regime, so the 8000 atom calculation should only take a little over 8 times the 1000 atom calculation. This is rather better than the nearly 512 times longer it would take with traditional cubic-scaling DFT!

2.2.3 Diagnosing Common Failures

With badly-chosen input settings, even fairly standard calculations in ONETEP will not converge, or may even converge to the wrong result. Fortunately, many of these problems are easy to fix with a bit of experience. In general, it is advisable to run with full output verbosity (`output_detail: VERBOSE`) the first few times you run a new kind of system, and to be on the lookout for any warnings or garbage numbers in the output (eg `****s` in place of what should be real numbers). Remember that for the energy to be accurate, we must have simultaneous convergence of both the density kernel and the NGWFs. If either of these are not converging well by the end of the calculation, there may be a problem. In this section, we will briefly examine some reasons behind common types of convergence failure, and what to do to eliminate those failures and perform accurate simulations.

- Problem: Repeated 'safe' steps (of 0.150 or 0.100) during NGWF Conjugate Gradients optimisation, leading to poor or no convergence. This often means that the step length cap for NGWF CG is too short. **Solution:** increase `ngwf_cg_max_step`, eg to 8.0.
- Problem: Repeated 'safe' steps (of 0.150 or 0.100) during LNV Conjugate Gradients optimisation, leading to poor or no convergence. This often means that the step length cap for LNV CG is too short. **Solution:** increase `lnv_cg_max_step`, eg to 8.0.
- Problem: Occupancies 'break' during LNV optimisation of kernel. Examine the output with `output_detail: VERBOSE` and look at the occupancy error and occupancy bounds during the "Penalty functional idempotency correction" section of each LNV step. Check for occupancies outside the stable range (approx -0.3:1.3) or RMS occupancy errors not decreasing (particularly if no kernel truncation is applied). **Solution:** Activate LNV line step checking with `lnv_check_trial_steps: T`. This checks that the kernel is still stable after the proposed line step is taken.
- Problem: Occupancies are 'broken' from start of calculation. Symptoms as above. Palser Manolopoulos may be unstable due to degeneracy or near-degeneracy at the Fermi level. Check the output of Palser Manolopoulos for warnings. **Solution:** If there is an initial degeneracy at the Fermi level, an $O(N^3)$ diagonalisation may be required to get a good starting kernel. Set `maxit_palser_mano: -1`.
- Problem: RMS Commutator (HKS-SKH) of kernel and Hamiltonian stagnates (stops going down with each iteration) during LNV optimisation. This is a sign that the current set of NGWFs is not able to represent a density matrix that both reproduces the electron density that generated the Hamiltonian while simultaneously describing the occupied eigenstates of that Hamiltonian. If this problem does not start to go away after a few steps of NGWF

optimisation, a better or larger initial set of NGWFs may be required. **Solutions:** Increase number of NGWFs per atom, increase radius of NGWFs, improve initial guess for NGWFs.

- **Problem:** RMS NGWF gradient stagnates (stops going down) during NGWF CG optimisation, while energy is still going down slowly. This often suggests that the NGWFs may have expanded away from their centres to have significant value near the edge of their localisation region, and thus cannot optimise successfully. **Solution:** Increase NGWF radius. Sometimes increasing the kinetic energy cutoff helps as well. For smaller systems and initial tests, a useful check on the accuracy of the final result is to perform a full $O(N^3)$ diagonalisation at the end of the calculation, if it is computationally feasible to do so. To activate this, turn on a properties calculation with `do_properties: T`, and then ask for an eigenvalue calculation of the first 100 eigenvalues either side of the Fermi energy, for the kernel and Hamiltonian matrices, by setting `num_eigenvalues: 100`. If all is well, then the occupation eigenvalues should all be close to 0.00000 or 1.00000 (empty or full) and the Hamiltonian eigenvalues should all be within a sensible range.

One final note if you're not getting the result you expect – check the units on your atomic positions! ONETEP expects positions in Bohr if the units are not specified, so if your positions are in Angstroms, you will need to add 'ang' as the first line of the `positions_abs` block.

This completes tutorial 1.

Files for this tutorial:

- `SiH4.dat`
- `Si8000.dat`
- `Si64.dat`
- `Si1000.dat`

TUTORIAL 4: GEOMETRY OPTIMIZATION

Author

Simon M.-M. Dubois

Date

June 2021 (revised Aug 2023)

3.1 Introduction

This tutorial aims at showing how to run a simple geometry optimization with ONETEP.

Geometry optimization is one of the primary tasks in quantum simulation. The essence of the calculation is for the constituting atoms to be moved to the positions where the total energy is minimal. In general, this can be tackled efficiently if the forces on the atoms can be computed. Over the past twenty years, various schemes have been derived to solve this problem in the framework of ab initio calculations. These range from simple approaches based on molecular dynamics, such as the steepest descent and damped dynamics methods, to the more sophisticated conjugated gradient and Quasi-Newton methods.

The geometry optimization scheme implemented in ONETEP relies on the isolation of the atomic and electronic subsystems (i.e. the Born-Oppenheimer approximation). For a given configuration of the ionic positions, the electronic degrees of freedom are completely relaxed so that the electronic subsystem stays on the Born-Oppenheimer surface. All the possible configurations of the ionic positions therefore define a multi-dimensional potential energy surface for which we want to find the global minimum. The atomic forces are calculated by application of the Hellmann-Feynman theorem and the ionic positions are moved around by means of the Broyden-Fletcher-Goldfarb-Shanno (BFGS) method in order to find the minimum of the potential energy. At this point, one has to keep in mind that several local minima may be present in the configuration space and the algorithm can get trapped in one of those. Therefore, despite the sophistication of the minimization method, the location of a global minimum still requires the intuition of a good starting point.

The calculation flow of a geometry optimization in ONETEP is a three step process:

- Given an ionic configuration, the electronic degrees of freedom are relaxed (cfr. self-consistent optimization of the density kernel and NGWFs).
- The total energy and atomic forces are computed and compared with those of previous ionic configurations. The threshold chosen as stopping criterion for the geometry optimization is

tested.

- The atomic position are updated by means of the BFGS algorithm.

3.1.1 The Ethene Molecule

As a first example, we will deal with the geometry optimization run a geometry optimization of the ethene molecule. To do this, we need to set:

```
task : GeometryOptimization
```

For the SCF part, try running a calculation with an energy cutoff of about 650 eV, NGWF radii of about 6.0 Bohr and a cubic simulation cell of side-length 40 Bohr:

```
cutoff_energy      : 650.0 eV
output_detail : VERBOSE

%block species
H  H  1  1  6.75
C  C  6  4  6.75
%endblock species

%block lattice_cart
Ang
20.0 0.0 0.0
0.0 20.0 0.0
0.0 0.0 20.0
%endblock lattice_cart

%block positions_abs
Ang
C  5.0000000000000000  5.9228319999999997  6.9051750000000007
C  5.0000000000000000  5.9228319999999997  5.5702150000000001
H  5.0000000000000000  6.8456639999999993  7.4753900000000009
H  5.0000000000000000  5.0000000000000000  7.4753900000000009
H  5.0000000000000000  6.8456639999999993  5.0000000000000000
H  5.0000000000000000  5.0000000000000000  5.0000000000000000
%endblock positions_abs
```

You will also need the pseudopotential block:

```
%block species_pot
H  hydrogen.recpot
C  carbon.recpot
%endblock species_pot
```

A full example input/output file can be found `ethene.tar.gz`.

Now you are ready to run ONTEP:

```
mpirun -n 2 onetep ethene.dat | tee ethene.out
```

The calculation should take ~15 min on 2 MPI processes (on a Intel Xeon Silver 4114 cpu). In the meantime you may want to repeat the procedure with varying parameters in order to converge the calculation with respect to the cutoff energy, the NGWF radii, as well as the size of the simulation cell. Besides, if you aim to compute a properties (e.g. the C-C bond length) with a given computational accuracy (e.g. 0.005 Ang), you should also check that the `geom_max_iter` and `ngwf_treshold_orig` parameters do not prevent to reach the desired accuracy.

The output of ONETEP consists principally of two files: `ethene.out` (the main output file) and `ethene.geom`. This latter contains one block of information for each iteration of the geometry optimization. Each block looks like:

```

          4
      -1.37265351E+001  -1.37265351E+001
-> <-- E
      3.77945227E+001  0.00000000E+000  0.
->00000000E+000 <-- h
      0.00000000E+000  3.77945227E+001  0.
->00000000E+000 <-- h
      0.00000000E+000  0.00000000E+000  3.
->77945227E+001 <-- h
C      1      9.44793440E+000  1.11928879E+001  1.
->30335653E+001 <-- R
C      1      9.44790699E+000  1.11928871E+001  1.
->05404458E+001 <-- R
H      1      9.44899454E+000  1.29429511E+001  1.
->41255628E+001 <-- R
H      1      9.44900396E+000  9.44173342E+000  1.
->41253623E+001 <-- R
H      1      9.44897328E+000  1.29435030E+001  9.
->45003514E+000 <-- R
H      1      9.44897086E+000  9.44121996E+000  9.
->45024026E+000 <-- R
C      1      -1.75749466E-005  -2.60718796E-004  -3.74725290E-
->004 <-- F
C      1      -1.20705498E-005  -2.14434259E-004  2.65511453E-
->004 <-- F
H      1      1.13344672E-005  2.37875189E-004  1.82501034E-
->005 <-- F
H      1      9.56687732E-006  1.39393204E-005  -1.03600237E-
->004 <-- F
H      1      4.64762662E-006  1.58612511E-004  4.33446848E-
->005 <-- F
H      1      4.09652530E-006  6.47260346E-005  1.51219285E-
->004 <-- F

```

where all values are in Hartree atomic units and

- The first line is the iteration number.

- The second line is the total energy.
- The next three lines are the lattice vectors expressed in Cartesian coordinates.
- The next N lines (where N is the number of atoms) give the atomic coordinates.
- The following N lines give the atomic forces.

The main informations regarding the geometry optimization are gathered in the `ethene.geom` file, however you may want to visualize the results in a glimpse. You can use the perl script `geom2xyz` to generate a `.xyz` file containing the atomic coordinates at each iteration of the geometry optimization:

```
chmod 700 geom2xyz
geom2xyz ethene.geom
```

This should produce a file `ethene.xyz` that you can visualize with your favourite package (e.g. XCrysDEN). Though the film of the relaxation provides you with crucial information such as the appearance of dissociation, symmetry breaking, etc. It is a good practice to keep track of the energy and forces at each iteration in order to assess the relaxation process. The `-- E` tag which labels the total energies in the `ethene.geom` file may be used for that purpose. Create a new file `ethene_energy.dat` and plot the evolution of the total energy using:

```
$ grep ' E' ethene.geom | awk '{print $1}' > ethene_energy.dat
$ gnuplot plot with lines 'ethene_energy.dat'
```

You should notice that the total energy of the system decreases monotonically. Similarly, you can keep track of the maximum rms force on the ions at each iteration by running:

```
$ grep "<-- BFGS" ethene.out | grep "|F|max"
```

This should produce you something like:

```
| |F|max | 2.038842E-002 | 2.000000E-003 | Eh/Bohr | No  ↵
↵| <-- BFGS
| |F|max | 3.567221E-003 | 2.000000E-003 | Eh/Bohr | No  ↵
↵| <-- BFGS
| |F|max | 5.188186E-003 | 2.000000E-003 | Eh/Bohr | No  ↵
↵| <-- BFGS
| |F|max | 1.375629E-003 | 2.000000E-003 | Eh/Bohr | Yes↵
↵| <-- BFGS
| |F|max | 4.568394E-004 | 2.000000E-003 | Eh/Bohr | Yes↵
↵| <-- BFGS
```

The second column is the calculated value of the maximum rms force on the atoms, the third column is the force threshold that the code is trying to achieve, the fourth column provides the units, and the fifth column informs you as to whether convergence of the force has been achieved or not. You may visualize this information using `gnuplot`:

```
$ grep "<-- BFGS" ethene.out | grep "|F|max" | awk '{print $4}' >
↵ethene_force.dat
$ gnuplot plot with lines 'ethene_force.dat'
```

You are now familiar with the geometry optimization scheme in ONETEP. You might examine in more details the input variables that allow to control the process. The keywords associated with the geometry optimization all start with the `geom_` prefix. Their description is found on the [ONETEP Documentation](#). In particular, take a few minutes to have a look at the variables:

```
geom_max_iter
geom_convergence_win
geom_disp_tol
geom_energy_tol
geom_force_tol
```

Though their default values may appear to be convenient in most circumstances, these latter are the very basic input variables to master before launching a geometry optimization. Here, it is important to note that the three tolerance criteria (`geom_disp_tol`, `geom_energy_tol`, and `geom_force_tol`) are not exclusive. The three criteria have to be satisfied in order for the optimization to stop. You might have noticed that during the relaxation of ethene molecule, the default threshold imposed on the atomic forces (`geom_force_tol : 0.02 Ha/Bohr`) has been reached before the one associated with the convergence of the energy (`geom_energy_tol : 10e-06 Ha/Atom`).

Like all the quasi-Newton schemes, the BFGS algorithm accumulates information about the Hessian matrix. As the the number of iteration increases, BFGS improves its knowledge of the the potential energy surface around the minimum and the matrix used to build the quadratic model of the potential energy surface converges towards the true Hessian matrix corresponding to the local minimum. However, the Hessian matrix is poorly approximated during the first few relaxation steps. It is therefore important to properly initialize the BFGS scheme. This may be conveniently done by means of a unique parameter `geom_frequency_est`. For the best efficiency, its value should corresponds to a rough estimation of the average of the optical phonon frequencies at the center of the Brillouin zone.

In the case of the ethene molecule, the average of the experimentally reported vibration frequencies is 0.0081 Hartree. This value is very close to the default setting of `geom_frequency_est` and we do not expect any speed-up of the relaxation process by adjusting it.

In various circumstances, it may appears convenient to impose some constraints to the atomic positions during the geometry optimization. Note that in the case of molecular systems it is often a good idea to keep an atom or an axis fixed during the optimization process in order to avoid losing computational time due to the rotations and/or translations of the system. Therefore it is worth having a quick look at the meaning of the variables:

```
species
species_constraints
```

Finally, it is worth noting that, when running a geometry optimization, ONETEP produces a `.continuation` file. This latter contains all the information regarding the optimization process and can be very helpful to restart an optimization from a previous run. In such a case, the only thing you will need is to turn on the flag `geom_continuation`. In the same line of thought, a appropriate use of the keywords that control the reading/writing actions of the code, may help you to save some precious computational time:

```
write_converged_dkngwfs
read_denskern
read_tightbox_ngwfs
```

For example, to use `write_converged_dkngwfs` : T is a good practice when running a molecular optimization as it avoid you to lose time in writing the density kernels and NGWFs on the disk.

3.1.2 The Sucrose Molecule

At this point, you should be familiar with most of the keywords needed to run a proper geometry optimization. Therefore, we suggest you to leave the ethene molecule and to try to optimize a larger organic molecule. You can find an example input file for the sucrose molecule [here](#). You should edit and read it carefully. You see that the `write_converged_dkngwfs` flag has been activated. In addition, the values of `ngwf_cg_max_step` and `lnv_cg_max_step` have been increased in order to allow unconstrained line search during the conjugate gradient optimization of the density kernel and NGWFs respectively.

The calculation should take a bit more than an hour (with 16 MPI processes on a Intel Xeon Silver 4114 CPU). Keeping trace of the atomic forces, you should notice a rapid decrease of the maximum rms force on the ions during the first relaxation steps. However, the hydrogen atoms tend to wiggle quite a lot and it takes a some time for the positions to settle down according to relaxation criteria.

```
| |F|max | 1.212754E-002 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 1.407689E-002 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 1.253042E-002 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 4.859480E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 1.052111E-002 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 5.953036E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 6.344620E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 6.587572E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 5.241521E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 5.455889E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
| |F|max | 3.623343E-003 | 2.000000E-003 | Ha/Bohr | No | <-- BFGS
```

3.1.3 Periodic Crystals

Here above, the geometry optimization scheme has been illustrated by means of two molecular systems. Obviously, the same scheme holds for periodic crystals. As an example, we will investigate the adsorption of ammonia on a (10,8) carbon nanotube.

The carbon nanotube considered here contains 488 carbon atoms in its unit-cell and its chiral periodicity is of 62.87 Bohr.

Following the prescriptions stated above, you should be able to write an input file for the nanotube (an example is given in [here](#)). Note that for large systems, the spatial expansion of the density

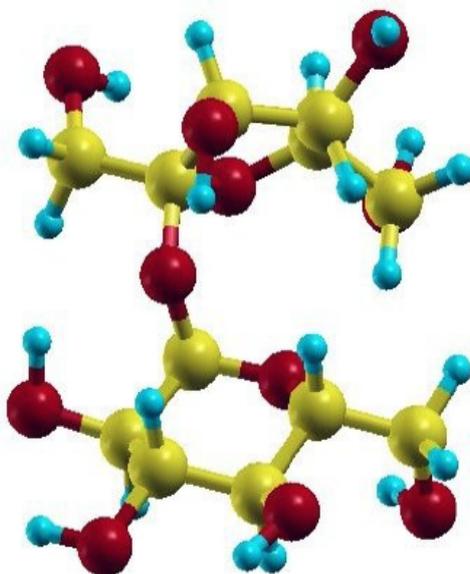


Fig. 3.1: Ball stick representation of the sucrose molecule.

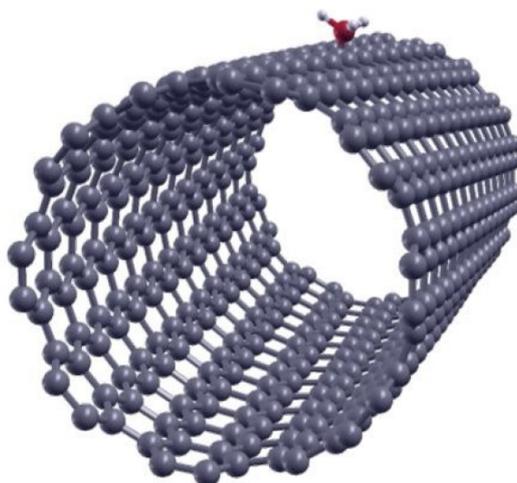


Fig. 3.2: Ball-stick representation of the ammonia adsorbed on a CNT (10,8)

kernel has to be truncated in order to achieve the linear scaling. This can be done with the `kernel_cutoff` variable. Obviously, stringent truncation of the density kernel is expected to affect the accuracy of the calculation. Therefore, the cutoff length has to be carefully adjusted. You should already notice a significant decrease of the forces after the first few iterations.

TUTORIAL 5: ANALYSIS AND VISUALIZATION

Author

Jacek Dziejcz, Chris-Kriton Skylaris

Date

July 2008 (revised April 2010, July 2023)

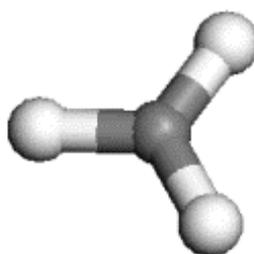
4.1 Introduction

This tutorial demonstrates how to:

- Use ONETEP to calculate various electronic properties,
- Instruct ONETEP to generate files needed for later visualization of orbitals, electronic densities and potentials,
- Visualize these properties using VMD¹,
- Set up and run a calculation on a nanostructure using a cut-off for the density kernel.

4.2 Density, spin density, Kohn-Sham orbitals and the electrostatic potential for CH₃

In this part we will perform a calculation on the CH₃ radical:



¹ VMD (Visual Molecular Dynamics) is a free of charge visualization package available from <https://www.ks.uiuc.edu/Research/vmd>

As this molecule contains an odd number of electrons we need to perform a spin-polarised (unrestricted) calculation. In ONETEP this is achieved by optimising a different density kernel for the “up” and the “down” spin:

$$\rho(\mathbf{r}, \mathbf{r}') = \sum_{\alpha\beta} \phi_{\alpha}(\mathbf{r}) K^{\alpha\beta(\uparrow)} \phi_{\beta}^*(\mathbf{r}') + \sum_{\alpha\beta} \phi_{\alpha}(\mathbf{r}) K^{\alpha\beta(\downarrow)} \phi_{\beta}^*(\mathbf{r}')$$

The ONETEP input is in the file `methyl.dat` and the coordinates (in angstrom) are in the file `methyl.pdb`. The ONETEP input file contains the coordinates as well (in atomic units), but not in a form directly readable by visualization packages. The `.pdb` file can be directly visualized in VMD.

The `methyl.dat` file specifies a single point energy calculation (`TASK SINGLEPOINT`) with a psinc kinetic energy cutoff of 800 eV (`CUTOFF_ENERGY 800.0 eV`), the Perdew-Zunger variant of the LSDA exchange-correlation functional (`XC_FUNCTIONAL CAPZ`) and the spin-polarised option (`SPINPOLARIZED TRUE`). Also notice the input flag `DO_PROPERTIES TRUE`, which proceeds with the calculation of various electronic properties at the end of the single point energy calculation.

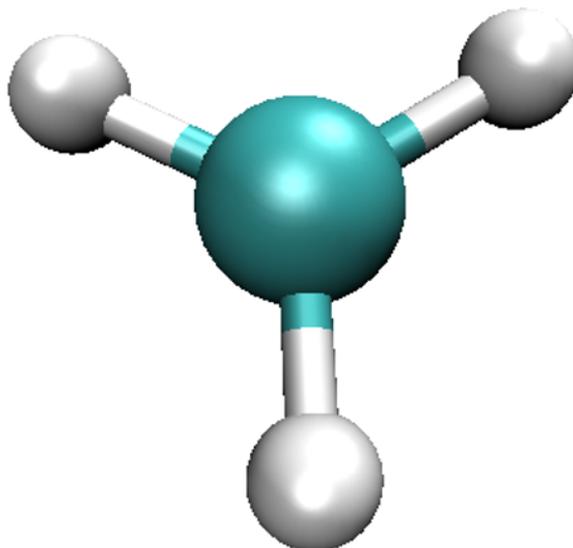
Run the input, redirecting the output to a file such as `methyl.out`. We also provide a reference `methyl.out` file. The calculation should take a minute or two to run. Once it completes, you will notice that a number of `.cube` files have been created, including the file `methyl_spindensity.cube`. Let us examine this first. ONETEP can output volumetric data (such as spin densities, charge densities, potentials, etc.) in Gaussian `.cube` format (`CUBE_FORMAT TRUE`), Materials Studio `.grd` format (`GRD_FORMAT TRUE`) and OpenDX `.dx` format (`DX_FORMAT TRUE`). The `.cube` format has the advantage of having the ionic positions output in addition to the volumetric data. In this tutorial we will use the `.cube` format which can be viewed with a number of free molecular visualisation programs. The instructions that follow are assuming that the VMD program can be used to visualize the files but in principle you can use any other software that can display `.cube` files (such as VESTA, Molekel, gOpenMol, XCrySDens, etc).

Start VMD by typing `vmd` in the terminal, use `File/New molecule/Browse` to find `methyl_spindensity.cube`, then click on `Load` to load the molecule. You should be able to see a crude, line-based representation of the molecule in a separate window. You can now get rid of the `Molecule` file browser window. Choosing `Graphics/Representations...` opens another window which lets you control the look of your molecule. In this window, change the `Drawing Method` from `Lines` to `CPK`, which will render your molecule in a ball-and-stick fashion, with the customary colouring². Increase both `Sphere Resolution` and `Bond Resolution` (30 is a good value) to get rid of the jagged edges. You may wish to adjust `Sphere Scale` and `Bond Radius` to your liking as well.

Try dragging with your mouse over the window that shows the molecule to rotate it. Try scrolling the mouse wheel to get closer or further away from the molecule. You may press the `=` key at any

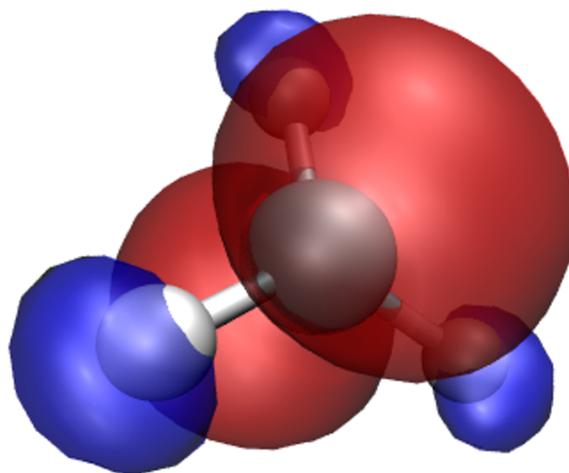
² The colouring is described here: https://en.wikipedia.org/wiki/CPK_coloring.

time to reset the view. Pressing the T key will get you to `Translate Mode`, where dragging with the mouse translates the molecule, instead of rotating it. To go back to `Rotate Mode`, press R. If your mouse lacks the scroll wheel, pressing S to go to `Scale Mode` might be of use. You should be able to obtain a representation similar to the one shown here.



So far we've only looked at the nuclei in the system. Let's try some electronic properties, starting from the spin density which we have already loaded, but not visualized yet. A neat thing about VMD is that you can use several representations at once. Thus, we can overlay the spin density isosurfaces on top of the CPK representation of the ions. In the `Graphics/Representations...` window click on `Create Rep`. This will clone the CPK representation, leaving you with two identical representations. Now change one of them to `Isosurface`. Not much will appear initially, because the default way of showing the isosurface is by using *points*. This is computationally cheap, but visually so as well. You can change this under `Draw`, by choosing `Solid Surface`. *Before you do it*, however, make sure to move the `Isovalue` slider to something different than the default 0.0 (or type a value in the box). This is because there is a huge number of points in our system (some 400000) where the spin density is exactly or almost exactly zero (everywhere outside our molecule). Trying to draw a surface through these points usually confuses VMD to the point of crashing or at least stuttering. For this reason it is best to pick any value other than the default of 0.0 to start from, before choosing `Solid Surface`.

Experiment with the settings (`Coloring Method`, `Material`, `Isovalue`) to get a feel for how they work. It makes sense to set `Coloring Method` to `ColorID` here, as this lets us to manually pick a colour for the isosurface (from the drop-down box near `ColorID`). After some adjustments you should obtain an isosurface similar to the one shown here. Do not worry if you cannot get the transparency right – it's – only possible when you render “production quality” images, think of what you see as a draft.

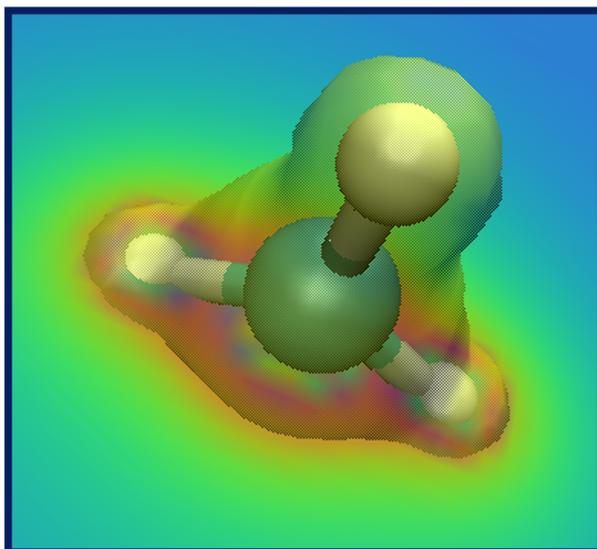


What we have obtained is the textbook picture of the spin density of a methyl radical. It has positive as well as negative regions which is a consequence of the fact that the spatial parts of the Kohn-Sham orbitals for each spin are allowed to be different, even for doubly occupied states.

The properties calculation also produces *Kohn-Sham orbitals*. Their energies for each spin are printed in the output file (try to find them, they are towards the very end, copy them into the table below) and `.cube` files for the squares of some of the orbitals are also produced. HOMO orbitals are written, separately for each spin, to `methyl_HOMO_DN.cube` and `methyl_HOMO_UP.cube`, and their LUMO counterparts to `methyl_LUMO_DN.cube` and `methyl_LUMO_UP.cube`. Similarly named files contain the orbitals just below the HOMO and just above the LUMO (not provided here, but generated during the calculation).

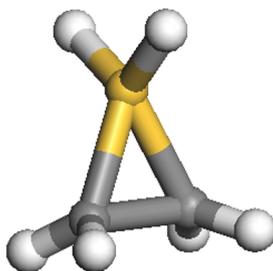
Quantity	Value
Energy of HOMO for spin 1 (UP)	
Energy of LUMO for spin 1 (UP)	
Energy of HOMO for spin 2 (DOWN)	
Energy of LUMO for spin 2 (DOWN)	
Energy of the 1 st orbital above LUMO for spin 1 (UP)	
Energy of the 1 st orbital below HOMO for spin 1 (UP)	

Finally, let's try visualizing the local potential (sum of the ionic, Hartree (Coulomb) and XC potentials), which is written out to `methyl_potential.cube`. Isosurface plots of potentials can be obtained similarly to the isosurface plots of densities. Let's also try to do a contour plot. This can be accomplished by choosing `VolumeSlice` for `Drawing Method`. Try playing with `Slice Axis` and `Slice Offset` to get the hang of it. Admittedly, the quality of the contour plot is not too good, even if you set `Render Quality` to `High`. It is improved, however, when you create a production image. Try obtaining a composite CPK + isodensity + contour plot similar to the one shown here.



4.3 Visualizing NGWFs and NNHOs for C_2SiH_6

In this example we will perform two sets of calculations on the C_2SiH_6 molecule:

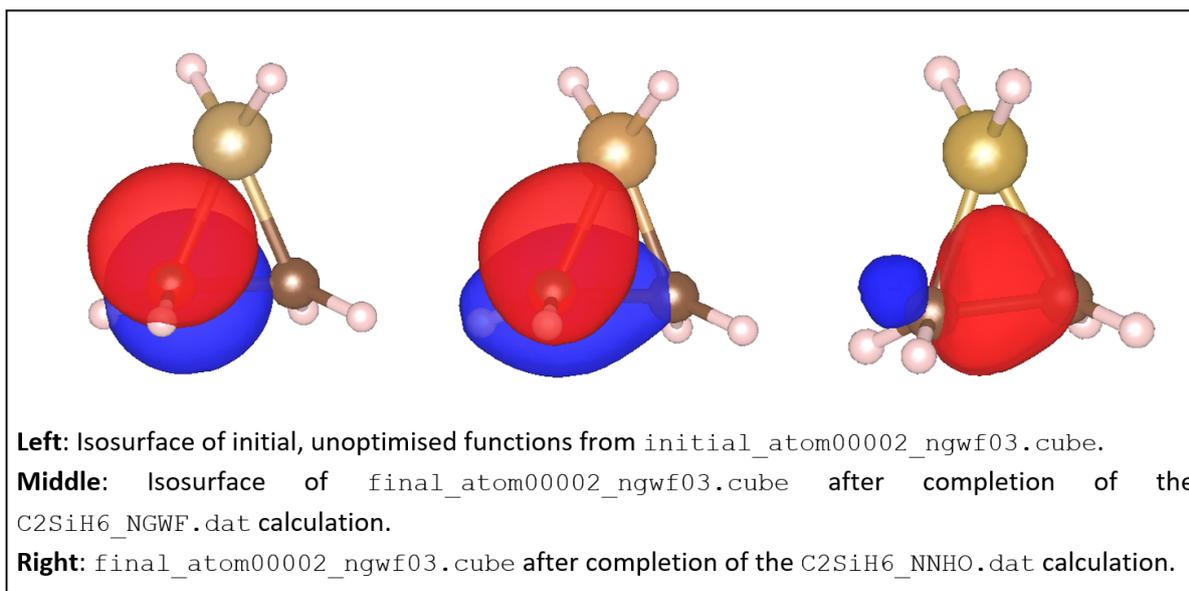


The first calculation will use the input file `C2SiH6_NGWF.dat`, which has similar parameters (and, thus, keywords) to the previous example but also contains the `WRITE_NGWF_PLOT TRUE` keyword that allows output of selected NGWFs in the scalarfield formats we discussed earlier (`.cube` by default). The NGWFs that will be outputted are selected by the `species_ngwf_plot` block in which the *species* of atoms whose NGWFs are to be outputted are listed. In this example we output NGWFs of the Si atom and of the first H and C atoms (as written in the input coordinates). The second input file is `C2SiH6_NNHO.dat`, which contains the additional keyword `NNHO TRUE` which instructs ONETEP to perform a same-centre rotation of the NGWFs to transform them to non-orthogonal natural hybrid orbitals (NNHOs). These contain the same information as the NGWFs but are more “natural” as they conform with chemical concepts, such as being directed towards chemical bonds, and physical concepts, as in several of their properties

they resemble proper Wannier functions. The mixing of NGWFs to NNHOs is done according to the procedure by Foster and Weinhold (J. P. Foster and F. Weinhold, *J. Am. Chem. Soc.* **102**, 7211 (1980)). For this calculation we will use the PBE GGA exchange-correlation functional (XC_FUNCTIONAL PBE).

Run the calculation to completion with the two inputs (in separate directories), it should take no more than five minutes for each of them. Reference outputs are provided here: `C2SiH6_NGWF.out` and here: `C2SiH6_NNHO.out`.

Examine some of the NGWF and NNHO output files. As an example, below we show plots of the third function (NGWF or NNHO) of atom 2 (one of the carbons). Try to obtain similar plots.



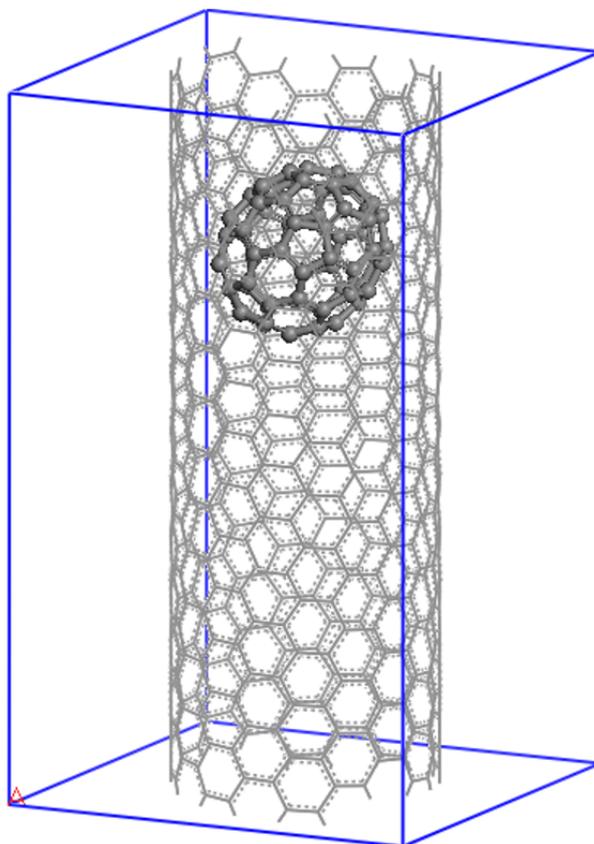
You can observe that initially the function is a p-atomic orbital (as it is initialised by ONETEP). After the calculation the NGWF is rather distorted but still contains quite a lot of p character. The NNHO however is a mixture of all the 4 NGWFs of the carbon atom and is optimally pointed along the C-C bond. You can quantify these observations by comparing the two output files, `C2SiH6_NGWF.out` and `C2SiH6_NNHO.out`, which contain an NGWF s/p/d/f Character Analysis section towards the bottom of the file (thanks to the `NGWF_ANALYSIS TRUE` keyword in the input). You will see how much the NGWFs differ from the NNHOs. Of course all the other quantities (energies, Kohn-Sham orbitals, orbital energies, etc.) are independent of whether you use NGWFs or NNHOs. Check this by completing the table below.

Quantity	Value
Total energy of the system	
Energy of HOMO	
Energy of LUMO for spin 2 (DOWN)	

Finally, examine the atomic population in the output files (we have asked for it using the keyword `POPN_CALCULATE TRUE` in the input) and confirm that the charges on each atom are consistent with their relative electronegativities.

4.4 A calculation on a nanostructure

Let us now see how to set up and visualize a calculation on a nanostructure whose size is in the region where conventional cubic scaling codes become very inefficient, while linear-scaling codes like ONETEP are still at the beginning of their capabilities. We will perform a calculation on the following “nano-peapod” structure, which consists of a C_{70} fullerene inside a single repeat-unit of a (10,8) carbon nanotube.



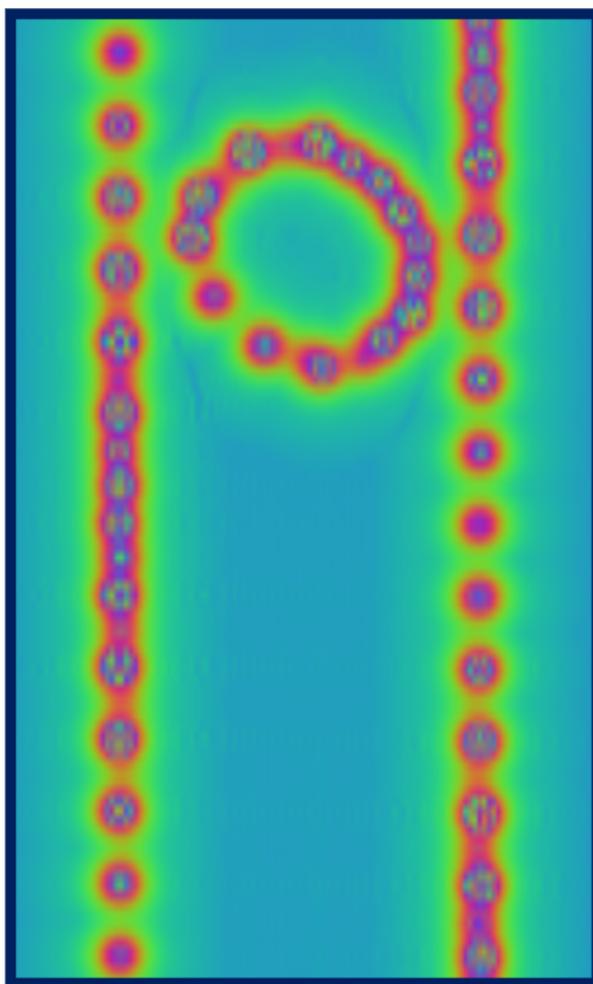
The (10,8) is a chiral nanotube with 488 atoms in each repeat-unit, so the peapod input consists of 558 atoms, with no symmetry, in a unit cell of 20.0 x 20.0 x 33.27 (angstrom), which is equivalent to 37.795 x 37.795 x 62.874 (bohr). The ONETEP input is in the file `C70_in_10-8.dat`.

We impose a density kernel cut-off of 30.0 bohr (`KERNEL_CUTOFF 30.0 bohr`) in order to achieve linear-scaling behaviour.

This calculation is best run on a parallel computer, but you can run it on a desktop machine where it should complete in about two-three hours. It took just under 8 minutes when run on 5 nodes (360 CPU cores) in 2023. If you do not want to wait or do not have the sufficient resources, here's the reference output: `C70_in_10-8.out`.

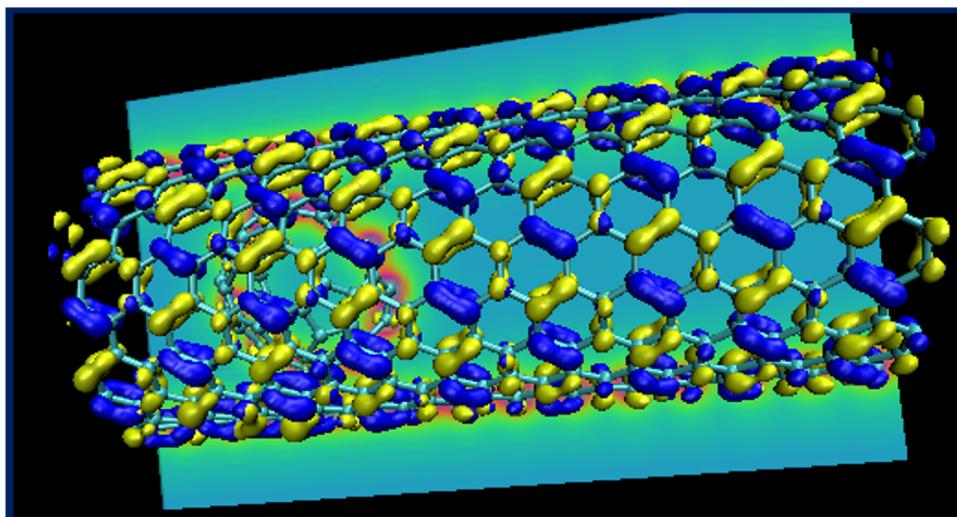
Let us start by examining this file. At the beginning of the calculation the *filling* (the opposite of sparsity) of various matrices is reported. You will notice that the density kernel is not 100% full as a consequence of the cut-off that is imposed in the input. Information about the psinc grid sizes is also provided, including the actual plane-wave cut-off to which they correspond and the size of the FFT box. The calculation converges in 7 NGWF iterations, which is the point where the NGWF gradient threshold set in the input (`NGWF_THRESHOLD_ORIG 0.00003`) has been satisfied. Normally you'd likely use a tighter threshold for extra accuracy (the default is 2E-6).

As before, a range of properties are calculated (`DO_PROPERTIES T`). As an example, you can examine the total potential (the sum of ionic, Hartree and exchange-correlation potentials) which is outputted to the file `C70_in_10-8_PROP_potential.cube`. We do not provide this file here due to size considerations. A contour plot on a plane containing the nanotube axis of the potential will look similar to what you see below, which is compatible with the chiral nature of the nanotube and reveals also the asymmetric way in which the oblong C_{70} is located inside it.



Red regions correspond to large and positive values of the potential (standard electrostatic conventions) and reveal the location of nuclei, whose distance from the plane varies along the axis of the tube, as a result of the chirality. You can go on and explore other properties of the nano-peapod from the `C70_in_10-8.out` file and the other output files that were produced by the properties calculation.

If you are in an ambitious mood, try creating a fancy plot showing the structure of the nano-peapod system with its HOMO and LUMO orbitals and a contour plot of the potential, similar to the one below.



This completes tutorial 5.

TUTORIAL 6: TIME-DEPENDENT DFT

Author

Tim Zuehlsdorf

Date

January 2016 (revised August 2023)

5.1 TDDFT in ONETEP

This tutorial aims at showing how to run a simple linear-response TDDFT calculation in ONETEP. Linear-response TDDFT is the main method of choice when computing optical excitations for large system sizes. In this formalism, excitation energies are directly obtained as the solutions to an effective eigenvalue equation. The approach implemented in ONETEP can be made to scale fully linearly with systems size, allowing for the computation of excitations in systems containing thousands of atoms.

In linear-response TDDFT, an excitation is expressed through a transition vector that can be written in the basis of all possible Kohn-Sham transitions from occupied to unoccupied states. Thus unlike in standard ground state DFT, where a good representation of the occupied Kohn-Sham states is sufficient to obtain accurate results, in a TDDFT calculation an equally good representation of a subset of the unoccupied Kohn-Sham space is vital. Since the ONETEP support functions are optimised *in situ* to ideally represent the occupied manifold, they generally provide a relatively poor description of the unoccupied manifold, making them insufficient for describing excitations in a linear-response framework. In practice, this problem is removed by optimising a second set of support functions to ideally span a low energy subset of the conduction space in a post-processing step after a `Singlepoint` calculation. Together, the two sets of support functions form an ideal representation of any given low energy excitation, thus enabling efficient and accurate TDDFT calculations in very large systems.

5.2 Setting up a TDDFT calculation in ONETEP

A TDDFT calculation in ONETEP thus proceeds in three separate steps:

- A standard ground state calculation using `Task: Singlepoint`. The code will write out `.dkn` and `.tightbox_ngwfs` files containing the valence density kernel and the NGWF support functions.
- A conduction optimisation using `Task: Cond`. The code will read in the converged `.dkn` and `.tightbox_ngwfs` files from the ground-state optimisation and perform a postprocessing `conduction optimisation` of some low energy part of the conduction space manifold. It will write out `.dkn_cond` and `.tightbox_ngwfs_cond` files containing an effective density kernel for the low energy conduction space, as well as the NGWF support functions for the conduction space.
- A TDDFT calculation using `Task: Lr_tddft`. The code will read in the density kernels and support functions for both the ground state and the conduction space calculations and then calculate the lowest n excitations of the system, where n is determined by the keyword `lr_tddft_num_states`.

In practice, it is possible to string these tasks together using `Task: Singlepoint Cond Lr_tddft` in an input file that lists appropriate keywords for all three calculation types. The tasks will be performed one after another, without the user having to restart the calculation in between.

5.3 An NN-substituted nanoribbon

We turn to a simple example of an excited state calculation in a periodic system, namely that of an infinitely long graphene nanoribbon with two nitrogen substitutions. We will aim to compute excited states that are associated with the substitution and that are thus relatively localised in character. In a first step, we aim to build an input file for a pristine graphene nanoribbon in periodic boundary conditions. A primitive unit cell for the nanoribbon in question can be found below:

```
%block positions_abs
C 10.90 10.0 0.000
C 13.58 10.0 0.000
H 7.50 10.0 2.327
C 9.56 10.0 2.327
C 14.93 10.0 2.327
H 16.99 10.0 2.327
%endblock positions_abs
```

To generate an appropriate input file, copy the above primitive unit cell block into a new input file called `ribbon.dat`. ONETEP does not make use of any k -point sampling and all calculations are effectively done at the Γ -point. Therefore, in order to accurately simulate periodic structures it is necessary to simulate supercells rather than primitive cells as would be done in standard plane wave DFT codes. An appropriate supercell from the above primitive cell can be created by repeating the primitive cell 7 times along the z-axis, translating the z-coordinates of the atoms by the lattice vector of the primitive unit cell. The resulting supercell should contain 48 atoms in total.

Now create an appropriate `%block lattice_cart` for the system. The length of the supercell is specified by the length of your primitive cell and the number of repeats of that cell. However, the size of the unit cell in the x and y direction is free to specify by the user. Since the ONETEP method uses periodic boundary conditions in all three directions, it is necessary to include enough vacuum in the x and y direction of the simulation cell to prevent it from interacting with its periodic neighbours. In this case, about $20 a_0$ of vacuum between any atoms of different periodic images is a reasonable length.

Using the lessons learnt from previous tutorials add a `%block species` and `%block species_pot` block to your input file. Choose an appropriate NGWF radius (try $8 a_0$) and kinetic energy cutoff (around 600 eV) for your calculation. When choosing the NGWF radius, note that ONETEP will not allow you to pick an NGWF diameter that is larger than the dimensions of your simulation cell, as this would cause an NGWF to interact with its periodic image. Finally, perform a singlepoint calculation of the system. You can add a keyword `do_properties: T` to the input file. This will trigger the code to perform a properties calculation, write out a file `.val_bands` containing all Kohn-Sham energies and write cube files of selected Kohn-Sham states around the band gap.

5.3.1 Conduction optimisation

Once the ground state calculation is finished, it is necessary to perform a conduction optimisation. First, set `Task: Cond` and add a new block `%block species_cond` to the input file. This block is made up in an identical way as `%block species` and specifies the number of support functions and their radius in the conduction optimisation. In many practical examples, it is often advisable to choose a larger NGWF radius for the conduction optimisation than for the ground state calculation, especially if one is interested in lightly bound states, as unoccupied Kohn-Sham states tend to be more diffuse than occupied ones (try $10 a_0$ for example).

A second keyword required for the conduction optimisation is the number of conduction states that should be explicitly optimised (`cond_num_states`). It is normally advisable to optimise only well-bound states as unbound states are difficult to describe with localised orbitals. In order to do so, have a look at the file `.val_bands` and count the number of Kohn-Sham states with negative energy that are unoccupied. Note that `cond_num_states` expects the number of electrons to be optimised as an input, thus in order to optimise the five lowest unoccupied Kohn-Sham states in a spin-degenerate system, one should choose `cond_num_states: 10`.

Once you have run the calculation, have a look at the output. You will find that ONETEP has generated a number of files such as cube files of unoccupied Kohn-Sham states expressed in terms of the optimised conduction NGWFs.

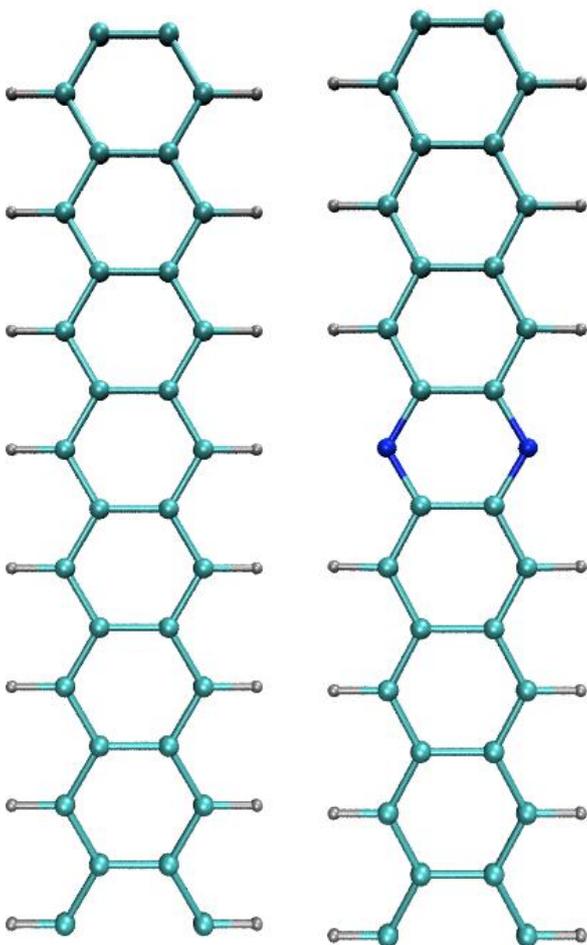
5.3.2 TDDFT calculation

We can now perform a full ONETEP TDDFT calculation of the system at hand. To do so, set `Task : Lr_tddft` in the input file. Furthermore, add the keywords `lr_tddft_num_states : 6`, `lr_tddft_write_densities : T` and `lr_tddft_analysis : T`. The code will compute the lowest 8 singlet excitations of the system and generate cube files for the electron, hole and transition density for each excitation that can be visualised. Furthermore `lr_tddft_analysis : T` triggers a breakdown of the converged TDDFT eigenvectors into Kohn-Sham transitions, allowing you to study which are the dominant transitions for each excitations.

Once you have performed the TDDFT calculation, look at the output file. You will see that the excitation energies and oscillator strengths for each of the excitations are printed out, as well as a detailed breakdown of excitation energies into Kohn-Sham transitions. Have a look at some of the cube files produced. Where are the excitations located in the system?

5.3.3 Nitrogen substitution

We can now move on from the case of the pristine nanoribbon to one with two nitrogen substitutions. For this purpose, copy the input file `ribbon.dat` to a new file `ribbon_NN.dat`. In that file, remove two C-H from the `%block positions_abs` that are opposite to each other in the ribbon, and replace them by two N at the same positions where the C were located.



Left: Pristine nanoribbon. Right: Nanoribbon with two carbons and two hydrogens substituted for two nitrogens

Note that in order to run the calculation, you will have to add the nitrogen species to the `%block species_pot`, `%block species` and `%block species_cond` blocks. Change the task to `Task: Singlepoint Cond Lr_tddft`. The code will run a ground state and conduction optimisation, followed by a TDDFT calculation for the full system. Have a look at the output. How do the excited states change due to the nitrogen substitutions? Where is each excited state located within the system?

5.4 Additional tasks

Substituting nitrogen atoms in the same place as carbon atoms does not yield a relaxed ground state structure, as the N-C bond is not of the same length as the C-C bond. Thus in order to obtain more realistic results for the substituted system, perform a geometry optimisation (see tutorial 4), followed by a ground state, conduction and TDDFT calculation of the full system. How do the results change? Furthermore, create a system where the nitrogen atoms are not substituted at exactly opposite positions in the structure, but an asymmetry along the z-axis is introduced. How does the character of the low energy excitations change?

5.5 Input files

All the files needed for the simulation can be downloaded from

- `ribbon_pristine.dat`
- `ribbon_pristine_NN.dat`
- `ribbon_pristine.out`
- `ribbon_pristine_NN.out`
- `carbon.recpot`
- `hydrogen.recpot`
- `nitrogen.recpot`

TUTORIAL 8: IMPLICIT SOLVATION, VISUALISATION AND PROPERTIES: PROTEIN-LIGAND FREE ENERGY OF BINDING FOR THE T4 LYSOZYME

Author

Lennart Gundelach, Jacek Dziedzic

Date

June 2021 (revised June 2023)

6.1 Introduction

6.1.1 Protein-Ligand Free Energies of Binding

The binding free energy is a measure of the affinity of the process by which two molecules form a complex by non-covalent association. An example of this, of central importance in biology, is the binding of a ligand to a protein. Many methods to computationally approximate the binding free energies of protein-ligand interactions have been proposed with the ultimate goal of computationally predicting small molecule drug candidates which bind strongly to the protein of interest.

6.1.2 Quantum Mechanics in Binding Free Energies

A key limitation common to most computational methods of estimating binding free energies is the assumption of the validity of classical mechanics. The atoms and electrons that constitute biological molecules, like proteins, are, however, governed by the laws of quantum mechanics. Charge transfer, polarization and non-local interactions are not captured by traditional classical mechanical force-fields. Thus, a true description of protein-ligand binding requires a quantum mechanical (QM) treatment of the problem. In theory, a full, *ab-initio* QM approach would be system-independent, parameter-free and would describe the full spectrum of physical phenomena at work.

Unfortunately, high-level QM methods like coupled-cluster (CC) are prohibitively expensive and often have cubic or worse scaling with system size. Thus, even the ligands alone are often too large for routine calculations with these methods.

6.1.3 Linear Scaling Density Functional Theory

Due to the cubic scaling of conventional density functional theory, full-protein calculations on many thousands of atoms are not feasible. To study larger systems, linear-scaling versions of DFT have been developed [Bowler2012]. The ONETEP code [Prentice2020_T8] is one such linear-scaling DFT implementation, exploiting hybrid MPI-OMP parallelism [Wilkinson2014] for efficient and scalable calculations. The unique characteristic of ONETEP is that even though it is linear-scaling, it is able to retain large basis set accuracy as in conventional cubic-scaling DFT calculations. The implicit solvation model is a minimal-parameter Poisson-Boltzmann (PB) based model which is implemented self-consistently as part of the DFT calculation [Dziedzic2011] [Womack2018] and uses the smeared-ion formalism and electron-density iso-surfaces to construct solute cavities.

6.1.4 T4 Lysozyme

The protein under investigation in this tutorial is a double mutant of the T4 lysozyme (L99A/M102Q). This protein has been artificially mutated to form a buried polar binding site and has served as a model or benchmark system for various protein-ligand binding free energy studies [Mobley2017]. Although this protein is not directly pharmaceutically relevant, it is a useful model system due to its relatively small size (2500 atoms), structural rigidity and well-defined, buried binding site, which can accommodate a wide variety of ligands. Fig. 6.1 shows the ligand catechol inside the buried binding site of the T4 lysozyme L99A/M102Q mutant. PDB files of the complex, host and ligand are provided as part of this tutorial for you to visualize the system. The picture shown uses the NewCartoon representation for the protein with coloring based on secondary structure and CPK (ball-and-stick) for the ligand with element based coloring.

6.1.5 QM-PBSA Binding Free Energies

In this tutorial we will calculate the binding free energy of catechol to the T4 lysozyme L99A/M102Q mutant. We will employ a simplified QM-PBSA approach [Fox2014] [Gundelach2021] on a single snapshot of the protein-ligand complex.

The QM-PBSA approach is a quantum-mechanical adaptation of traditional MM-PBSA, which is an end-point, implicit solvent, binding free energy method. In this approach, the binding free energy is given by

$$\Delta G_{\text{bind}} = G_{\text{complex}} - G_{\text{host}} - G_{\text{ligand}}, \quad (6.1)$$

where G_{complex} , G_{host} , and G_{ligand} is the free energy of, respectively, the complex, host and ligand in an implicit solvent. Each of these can be decomposed into three terms,

$$G = E + \Delta G_{\text{solv}} - TS, \quad (6.2)$$

where E is the total gas-phase energy, ΔG_{solv} is the free energy of solvation and $-TS$ is an entropy correction term. In this tutorial, the entropy term will be ignored, as it is usually calculated in other programs using normal mode analysis. The linear-scaling DFT code ONETEP will be used to calculate the gas-phase and solvation free energy of the complex, host and ligand at a fully quantum-mechanical level.



Fig. 6.1: Catechol bound in the buried binding site of the T4 lysozyme L99A/M102Q double mutant. Visualization in VMD.

6.2 Setting up the calculations

We will set up three separate calculations, one each for the protein-ligand complex, the protein (host) and catechol (ligand). The structure of the complex was taken from a molecular dynamics simulation of the complex used in two QM-PBSA studies on this system [Fox2014] [Gundelach2021]. The structure of the unbound ligand and host were obtained from the complex by deletion of the respective molecules. Apart from the atomic coordinates, we must specify the details of the ONETEP single-point calculations, provide pseudopotentials for the atoms present in the system and adapt job submission scripts to run the calculations on the supercomputer of choice.

6.2.1 The input files

The ONETEP input file, referred to as the `.dat` file, contains two main elements: 1) the coordinates and atom types of the system (i.e the structural information) and 2) the details of the calculation. Due to the large system size, we have split these two components across separate files: the `.dat` file, which contains the structural information, and a `.header` file which contains instructions for ONETEP. This header file is included in the `.dat` file via the command `includefile`. All information could also be contained in a single `.dat` file; however, the use of a separate header file can make it easier to set up hundreds or even thousands of calculations which differ only in the coordinates and not the calculation settings.

`.dat` file

The two blocks included in the `.dat` file are `lattice_cart` and `positions_abs`, which specify the simulation cell and absolute positional coordinates of each atom within the simulation cell, respectively. The `includefile` command on the first line specifies the header file to include for the calculation.

`.header` file

This `.header` file contains all further details of the ONETEP calculation. The `species` block specifies the name, element, atomic number, number of NGWFs and the NGWF radius for each atom type in the system. The `species_pot` lists the names of the pseudopotential files for each atom type. The rest of the file consists of ONETEP keywords which control the details of the calculation. The provided header files are fully commented, and details on each keyword are given in the ONETEP keyword directory (<http://onetepkeywords.icedb.info/onetepdoc>). We will be performing single-point energy calculations using the PBE exchange-correlation functional, the D2 dispersion correction and ONETEP's minimal parameter implicit solvent model. The calculation will output verbose detail and an `.xyz` file for easy visualization. The total system charge is +9 for the complex and host and 0 of the ligand. The implicit solvent is set to use the default parameters for water.

6.2.2 Submission Scripts

Due to the large system size of over 2500 atoms, these single-point calculations can only be run on a supercomputer. Thus, a submission script appropriate for the HPC environment you are working on will be necessary. The standard distribution of ONETEP provides sample submission scripts for a variety of HPC systems. These can be found in your ONETEP directory under `hpc_resources`.

We recommend to run the complex and host calculations on multiple compute nodes, making full use of the hybrid MPI-OMP capabilities of ONETEP. On the national supercomputer ARCHER2, the use of 4 compute nodes (128 cores each) with 32 MPI processes and 16 OMP threads per process results in a wall-time of about 8 hours. Due to the much smaller size of the ligand, the calculation on the ligand in solvent should be limited to a single node, with at most 10 MPI processes.

6.3 Evaluating the Outputs

Upon successful completion of the calculations, we will examine the three `.out` files created. Each of these files contains the full details and log of the calculation, as well as the final results and some timing information. While much information about the system can be gained from the output files, we will focus first only on the final results necessary to estimate the binding free energy of the ligand, catechol, to the protein.

kcal/mol	Complex	Host	Ligand	Complex–Host–Ligand
E	-7372184.3	-7328209.2	-43940.1	-35.0
ΔG_{solv}	-2615.0	-2613.3	-9.7	8.0
G	-7374799.3	-7330822.5	-43949.7	-27.1

Fig. 6.2: Calculating the binding free energy of catechol to the protein.

As outlined in equations (6.1) and (6.2) we need to calculate the total free energy of the complex, host and ligand before subtracting the total energy of the host and ligand from that of the complex. As stated before, we will be ignoring any entropy contributions in this tutorial. The total energy is then the sum of the total gas phase energy and the solvation free energy. These energies are summarized in an easy to read section at the very end of the output files, just before the timing information. To find it, search the output file for `Total energy in solvent`. This section breaks down the different energy contributions and states the total energies in vacuum (gas phase) and in solvent as well as the solvation free energy. Fig. 6.2 summarizes the energy values obtained. To estimate the binding free energy we simply apply equation (6.1) to yield:

$$\begin{aligned}\Delta G_{\text{bind}} &= G_{\text{complex}} - G_{\text{host}} - G_{\text{ligand}} = \\ &= -7374799.3 - (-7330822.5) - (-43949.7) = -27.1 \text{ kcal/mol}\end{aligned}$$

Thus, the estimated binding energy of catechol to the T4 lysozyme is -27.1 kcal/mol. However, there are a number of severe limitations of this estimate: 1) the entropy correction term $-TS$ has been neglected; 2) only a single snapshot was evaluated; 3) the implicit solvent model incorrectly

interprets the buried cavity in the T4 lysozyme, and 4) the QM-PBSA method is designed to calculate relative binding free energies between similar sets of ligands. For an in depth look at the full application of the QM-PBSA binding free energy method to 7 ligands binding to the T4 lysozyme and a discussion of the errors, convergence and limitations of the method, please consult our recent publication [Gundelach2021].

6.3.1 Cavity Correction

The minimal-parameter PBSA solvent-model implemented in ONETEP incorrectly handles the buried cavity in the T4 lysozyme (L99A/M102Q). This is a known issue for solvent models based on the solvent accessible surface area, and has been described in detail in 2010 by Genheden *et al.* [Genheden2010], and in 2014 by Fox *et al.* [Fox2014].

In the un-complexed protein calculation, i.e the host, the surface area of the interior of the buried binding site is counted towards the solvent accessible surface area (SASA) used to calculate the non-polar solvation term. Thus, the non-polar term of just the protein is larger than that of the complex indicating the formation of a larger cavity in the solvent. Conceptually, the SASA model creates an additional, fictitious, cavity in the solvent with the SASA of the buried binding site. Because the non-polar term of both the protein and complex are known, a post-hoc cavity-correction may be applied to remove the additional (spurious) contribution of the buried cavity to the non-polar solvation energy. A full derivation is provided in [Fox2014].

$$E_{\text{cav-cor}} = 7.116(E_{\text{non-polar}}^{\text{host}} - E_{\text{non-polar}}^{\text{complex}}) = 7.116(289.5 - 286.2) = 23.5 \text{ kcal/mol.}$$

Applying the cavity correction term calculated above to the binding free energy, we obtain a cavity-corrected binding free energy of $-27.1 + 23.5 = -3.6$ kcal/mol. For comparison, the experimental binding energy of catechol to the T4 lysozyme is -4.4 kcal/mol. It should however be noted, that the close correspondence of this single snapshot QM-PBSA binding free energy to the absolute experimental energy is likely a lucky coincidence, as the QM-PBSA method is mainly applicable to relative binding free energies and the entropy correction term has not yet been included.

6.4 Properties

We will now show how a number of useful properties of the system can be studied through a *properties calculation*. In the interest of saving computational time, and for clarity of presentation, we will use the ligand system as an example.

Add the following keywords to the `.header` file of the ligand calculation:

```
do_properties T
dx_format T
cube_format F
```

and run it again.

The first of these keywords instructs ONETEP to perform a *properties calculation* towards the end of the run. This will calculate, among others, Mulliken charges on the atoms, bond lengths, the HOMO-LUMO gap, the density of states (DOS) and some grid-based quantities, such as the HOMO and LUMO canonical molecular orbitals, electronic charge density and potential. The grid-based quantities (often called *scalarfields*) can be output in three different formats: `.cube`, `.dx`, and `.grid`. By default `.cube` files are written, and not the other two formats. In this example we switch off `.cube` output and turn on `.dx` output. This is effected by the last two keywords.

Once your calculation finishes, you will see that quite a number of `.dx` files have been produced:

- `..._HOMO.dx` – density of the canonical HOMO orbital.
- `..._LUMO.dx` – density of the canonical LUMO orbital.
- `..._HOMO-n.dx` – density of the n -th canonical orbital below HOMO.
- `..._LUMO+n.dx` – density of the n -th canonical orbital above LUMO.
- `...density.dx` – the electronic density of the entire system.
- `...potential.dx` – the total potential (ionic + Hartree + XC) in the system.
- `...electrostatic_potential.dx` – the electrostatic potential (ionic + Hartree) in the system.

These files correspond to the calculation in solvent. There will be a second set of `.dx` files with `vacuum` in their names – these correspond to the calculation in vacuum. This lets you study and visualize in-vacuum and in-solvent properties separately and to perform comparisons between the two. Here, you can expect the scalarfields to be rather similar between in-vacuum and in-solvent because the ligand is charge-neutral and polarizes the solvent only very slightly.

There is a separate tutorial (Tutorial 5) devoted to visualization. You can use the skills taught there to create fancy visualizations of the properties of your choice. Here we will only show how to produce a neat visualization of the electronic density coloured by the electrostatic potential using VMD.

Load the electronic density and the electrostatic potential into one molecule, and the atomic coordinates into a separate molecule. This will make it easier to treat the scalarfields and the atomic coordinates separately. To achieve this, issue:

```
vmd ligand_2001_density.dx ligand_2001_electrostatic_potential.dx -
  ↪m ligand_2001.xyz
```

Once VMD loads the files, go to Graphics/Representations. Ensure Selected Molecule (at the top of the window) is the `.xyz` file (atomic coordinates). Under Drawing Method Choose CPK – this will create a ball-and-stick drawing of the ligand. Switch Selected Molecule to the `.density.dx` file to operate on the electronic density scalarfield. Under Drawing Method choose Isosurface if it is not chosen already. Choose an Iso-value of 0.1 to pick a reasonable density isovalue to plot. Under Coloring Method choose Volume (you might need to scroll to the very bottom to get there). In the tiny drop-down window to the right of Coloring Method switch from scalarfield 0 (the density itself) to scalarfield 1 (the potential) – this will colour the density *with* the potential. For Material (further to the right) choose Glass2 – this will choose a somewhat translucent material that will let us see

both the ball-and-stick model and the electronic density. Under `Draw` in the bottom-right of the window, choose `Solid Surface` instead of `Points`. Finally, let's change the range of the potential to the kinds of values that occur at the distance from the molecule at which our electronic density isosurface lies. These have been determined by trial and error. There are four tabs just above `Coloring Method`. Somewhat counterintuitively, switch to `Trajectory`, where, under `Color Scale Data Range` you can enter the minimum and maximum values for the potential (in eV). Enter `-1` in the left field and `1.5` in the right field and click `Set`. This should give a nice representation, which you can then rotate and translate to your liking using the mouse in the `OpenGL Display` window. Once you are satisfied, you can render the final image by going to `File/Render`. In the top drop-down menu choose `Tachyon` and click on `Start Rendering`. After a short while you will get a `.tga` ("TARGA format") file in the directory you are working in. It will look more or less like the graphics in Fig. 2. Most graphics manipulation programs and graphics viewers read `.tga` files. If you have `ImageMagick` installed, you can use it to convert the image to a more common format, like `.png`:

```
convert vmdscene.dat.tga vmdscene.dat.png
```

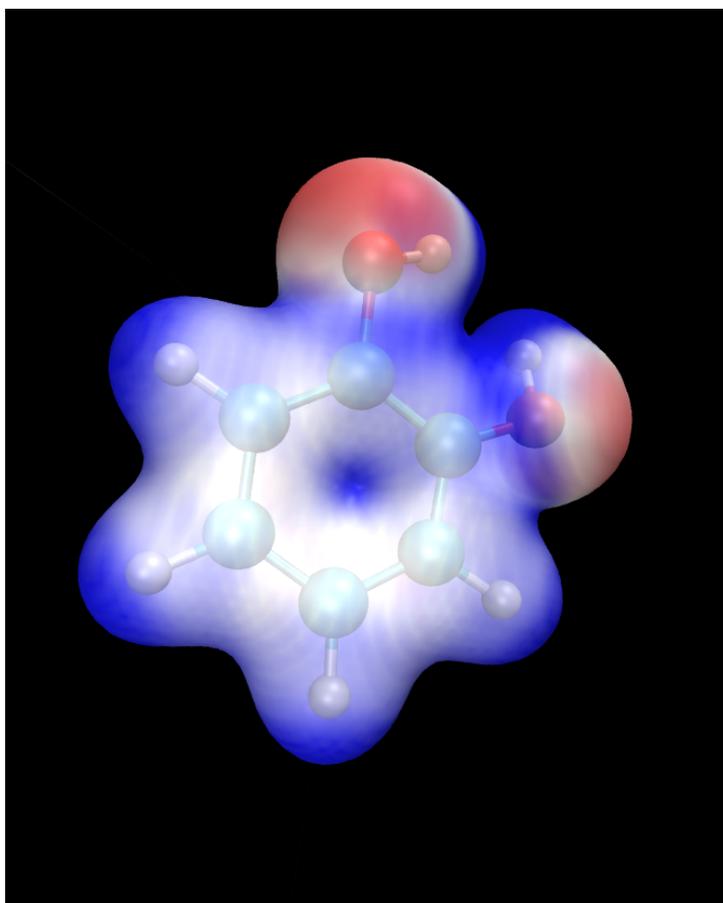


Fig. 6.3: Visualization of the ligand in VMD. A ball-and-stick model of the molecule is shown, together with an isosurface of the electronic density, coloured by the electrostatic potential.

6.4.1 Atomic charges

Mulliken population analysis

By default, during a properties calculation, ONETEP performs Mulliken population analysis, calculating partial atomic charges. The charges are written to the output file, in a table that looks like this:

Mulliken Atomic Populations			
Species	Ion	Total	Charge (e)
O	1	6.750	-0.750
H	2	0.448	0.552
C	3	3.817	0.183
...			

The partial charges (in the electrons-are-negative sign convention) are output in the last column.

Mulliken population analysis has a number of drawbacks, chief among which is that it depends on the basis set used and there is no well-defined complete basis set limit. Below we discuss two alternative schemes that can be used in ONETEP: Natural Population Analysis (NPA) and Density-Derived Electrostatic and Chemical (DDEC) analysis.

Natural Population Analysis

In Natural Population Analysis the set of non-orthogonal, optimized NGWFs is transformed into a set of orthogonal atom-centered Natural Atomic Orbitals (NAOs). This approach lets empty, highly-diffuse orbitals distort to achieve orthogonality with their more highly-preserved occupied counterparts, ensuring the final NAO population is stable with respect to basis set size. More details, and references to papers on the method, can be found in the documentation for this functionality – chapter “Population Analysis” in the main ONETEP documentation.

To perform Natural Population Analysis *in lieu* of Mulliken population analysis, add the following keyword to your previous ligand calculation:

```
write_nbo T
```

and run it again. Keep the three keywords you added last time.

Once your calculation completes you will find the results of NPA in your output file. They will look like this:

Natural Population	
Summary	

(continues on next page)

(continued from previous page)

Atom		Population (e)	Charge (e)
O	1	6.7313861	-0.7313861
H	2	0.4487370	0.5512630
C	3	3.7852506	0.2147494
...			

Density-Derived Electrostatic and Chemical (DDEC) analysis

ONETEP uses the DDEC3 method [Manz2012] to effect atoms-in-molecule electron density partitioning, producing partial charges, as well as higher multipoles (if desired), which are both chemically meaningful and give a faithful reproduction of the electrostatic potential of the QM system. More details, and references to papers on the method, can be found in the documentation at www.onetep.org/pmwiki/uploads/Main/Documentation/ddec.pdf.

To perform DDEC analysis *in lieu* of Mulliken population analysis, add the following keyword to your previous ligand calculation:

```
ddec_calculate T
```

You will also need to add a block `ddec_rcomp` that will specify where the reference ion densities can be found. You will need *two* reference density files for every atomic species in your system – one for the core and one for the total density, except for H and He which only require the total density file. The reference density files for a number of often-found elements can be found in the `c2_refdens` directory of your ONETEP installation. Fortunately all the files necessary for our ligand calculation (so, reference densities for C, H and O) are already there. Add the following block to your ligand input file:

```
%block ddec_rcomp
H ALL "H_c2.refconf"
O ALL "O_c2.refconf"
O CORE "O_c2.coreconf"
C ALL "C_c2.refconf"
C CORE "C_c2.coreconf"
%endblock ddec_rcomp
```

and copy the five files listed in the block from the `c2_refdens` directory to where your calculation resides. The documentation explains where you can find reference density files for other elements, should you ever need them.

Once you re-run your ligand calculation, you will find the results of DDEC analysis towards the end of your output file. They will look like this:

DDEC Charges (X=0.21)		
Atom	Population (e)	Charge (e)

(continues on next page)

(continued from previous page)

O	1	8.5534066	-0.5534066
H	2	0.5775414	0.4224586
C	3	5.8305022	0.1694978
...			

Comparison of Mulliken, NPA and DDEC charges

The three approaches for calculating partial charges are compared in Table [tab:charges]. Mulliken charges are, in general, the most pronounced out of the three, while DDEC partial charges are overall smaller in absolute value. The predictions of NPA are rather close to Mulliken analysis, while DDEC differs more from the first two.

Atom number	Species	Mulliken charge	NPA charge	DDEC charge
1	O	-0.750	-0.731	-0.553
2	H	0.552	0.551	0.422
3	C	0.183	0.215	0.169
4	C	-0.319	-0.301	-0.229
5	H	0.311	0.251	0.160
6	C	-0.320	-0.261	-0.158
7	H	0.295	0.237	0.130
8	C	-0.313	-0.252	-0.124
9	H	0.298	0.241	0.131
10	C	-0.309	-0.300	-0.243
11	H	0.296	0.240	0.146
12	C	0.230	0.246	0.216
13	O	-0.711	-0.685	-0.510
14	H	0.557	0.549	0.444

Fig. 6.4: Comparison of three approaches for calculating partial charges for the ligand.

But... tables are boring. How can we visualize the charges using VMD? This is not as straightforward as we would like. The structure (atomic coordinates) is contained in the `.xyz` file, but the charges are not. Some programs can visualize a quantity added in an extra column in the `.xyz` file (which would become something like an `.xyzq` file), but not VMD, at least not easily.

Fortunately VMD can read a different format named `.vtf`, which contains both the atomic coordinates and some scalar quantity, like charge. It is easy to convert an `.xyz` file and a list of charges to a `.vtf` file. We provide a simple `bash` script with this tutorial that does exactly that. It scans a ONETEP `.out` file for charge information (be it Mulliken, NPA or DDEC charges) and extracts the values of the charges on all atoms. It then looks for a corresponding `.xyz` file and, if found, it produces a `.vtf` file ready for visualizing with VMD.

To use it, download the provided script called `out2charge`, put it in your `$PATH`, and run it on your output:

```
out2charge ligand_2001.out
```

If everything goes well, you should see the following output:

```
Charges were output to ligand_2001.charge.  
The files ligand_2001.xyz and ligand_2001.charge will be used  
to construct ligand_2001.vtf.  
Load ligand_2001.vtf into VMD and select 'Coloring method -> charge  
↔'.
```

Indeed, a new file `ligand_2001.charge` will be produced, containing the charges extracted from the `.out` file. These charges, together with the information in the `.xyz` file will be used to construct a `.vtf` file readable by VMD. Load this file into VMD:

```
vmd ligand_2001.vtf
```

and go to Graphics/Representation. For Drawing Method choose CPK and for Coloring Method choose Charge. You will get a nice ball-and-stick model of your ligand, with the atoms coloured according to charge. In Fig. 3 we show a comparison of the plots for the three ways of partitioning charge that we described earlier.

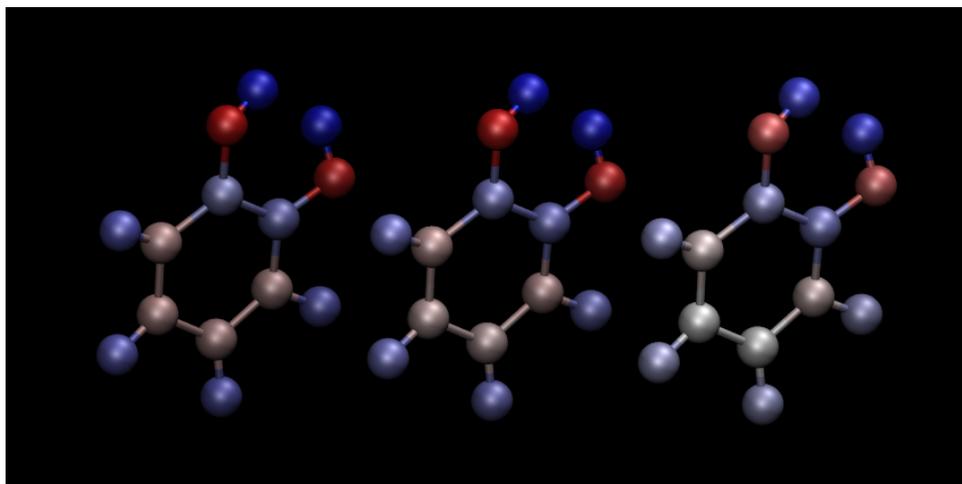


Fig. 6.5: Comparison of atomic charges on the ligand: Mulliken (left), NPA (middle) and DDEC (right). Warm colours correspond to negative charges. Visualization in VMD.

This completes tutorial 8.

Files for this tutorial:

- `out2charge`
- `T8_files.zip`

References:

TUTORIAL 9: DFT+U ON STRONGLY CORRELATED MAGNETIC MATERIALS: A CASE STUDY ON ANTIFERROMAGNETIC HEMATITE

Author

Davide Sarpa

Date

July/Aug 2023

7.1 Introduction

The goal of the tutorial is to provide a working example on how it is possible to model strongly correlated magnetic materials applying DFT+U methods [Anisimov1991] [Anisimov1997] [Dudarev1998]. We will be working with hematite as an example of such materials but the method can be applied to a plethora of materials.

7.1.1 Hematite

Hematite is a blood-red iron oxide with formula $\alpha - \text{Fe}_2\text{O}_3$ with a melting point of 1350 °C. It belongs to the hexagonal crystal family, in particular, it is a Ditrigonal scalenohedral with a $R\bar{3}C$ space group (167), sharing the same structure as corundum. The lattice parameters a, b, c are $a = b = 5.0356 \text{ \AA}$ and $c = 13.7489 \text{ \AA}$ with 6 formula units per cell with a band gap of 1.9 – 2.2 eV. Its structure is an hpc anion stacking of O^{2-} along the [001] direction with Fe^{3+} occupying 2/3 of the interstitial octahedral positions [Cornell2003]. Below the Néel temperature ($T_N = 963\text{K}$), Fe_2O_3 is antiferromagnetic with weak ferromagnetism. The high-spin $d^5\text{Fe}^{3+}$ cations within one bilayer in the (001) planes are ferromagnetically coupled to each other while antiferromagnetically coupled to the adjacent Fe bilayers [Parkinson2016]. The magnetic moment is determined to be $4.6 \mu_B$ per atom. The upper edge of the valence band is dominated by oxygen p states, while the bottom of the conduction band is dominated by Fe d states. Hematite is generally considered to be a charge transfer rather than a Mott-Hubbard insulator [Naveas2023] [Huang2016] [Si2020].

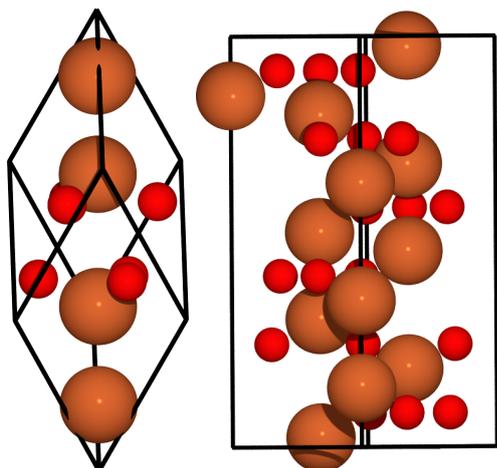


Fig. 1: Primitive rhombohedral cell (left), Conventional hexagonal cell (right). Iron atoms in brown, oxygen in red.

7.1.2 Magnetism

The weak ferromagnetism is due to spin-canting which is a relativistic effect. Luckily for us it is possible to obtain Hematite in an antiferromagnetic state with magnetic moments close to the experimental values by properly setting up the simulation without the need to include relativistic effects explicitly.

If we consider the primitive hematite cell along the Z axis there are 3 possible different antiferromagnetic states

- $+-+$ (up-down-up-down)
- $++-$ (up-up-down-down)
- $+--$ (up-down-down-up)

Out of all these states the last one ($+--$) is the ground state, we would like to be able to force our system to end up in the ground state.

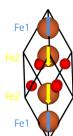


Fig. 7.1: Fig.2: Ground state antiferromagnetic state for hematite in primitive cell, Fe1 labels iron with up spin configuration, Fe2 for down spin configuration.

Magnetic systems are challenging to model due to the existence of very many different local minima which, in most cases, are very close in energy. Forcing the system into a specific state may not be easy, but there are methods that can help us achieve what we want.

7.1.3 DFT+U

The problem of DFT to describe correlated systems can be attributed to the tendency of xc functionals to over-delocalize valence electrons and to over-stabilize metallic ground states, this prevents materials like Hematite to be described by DFT. LDA and GGA both predict hematite to be a metallic system and they also underestimate local magnetic moments. The reason behind this delocalization is rooted to the inability of the approximated xc to completely cancel out the electronic self-interaction contained in the Hartree term. The main advantage of the DFT+U method is that it is within the realm of DFT, thus does not require significant effort to be implemented in the existing DFT codes and its computational cost is only slightly higher than that of normal DFT computations.

DFT+U treats the d and f orbitals as localised while keep the rest at the delocalised LDA/GGA level. In order to do that it projects the bands into localised orbitals and compute a new potential. The DFT+U method can be used to penalise the non-integer occupancy of these orbitals, tending to fill states with occupancy greater than 0.5 and empty states with occupancy less than 0.5.

$$\hat{V}_{DFT+U}^{(\sigma)} = \sum_I U^{(I)} |\varphi_m^{(I)}\rangle \left(\frac{1}{2} \delta_{mm'} - n_{mm'}^{(I)(\sigma)} \right) \langle \varphi_{m'}^{(I)} |$$

The U and J values are screened Coloumb and Exchange terms which are system and implementation dependent. In general, you are not able to plug and play an U or J value from the literature. What is usually done it's empirically testing different values (run multiple calculations with different combination of U and J) or most software (including ONETEP) have a linear response theory implementation to calculate the parameters from first principles [O-Regan2010] [O-Regan2012] [Cococcioni2005].

7.2 Setting up the calculations

Now we start with the actual tutorial, we will setup a bulk hematite simulation where we apply a DFT+U correction on the d orbitals of the iron atom. We also have to label atoms based on their spin which allow us to assign different parameters for spin down and spin up Fe atoms. This will help us achieve the AFM state we are looking for. We will see what we should look out in an output and what are the interesting properties for this material.

You will see that the cell and atoms we are using are neither from a primitive or a conventional cell, It is a orthogonal 3x2x1 supercell and this because orthogonal cell are easier to deal with and due to having 11 bohr radius NGWFs we need to have bigger cells.

7.2.1 Input files

ONETEP requires different files to work properly.

1. A .dat file which contains all the information about your sytem (positions and cell) and the simulations parameters
2. Pseudopotentials files, we will be using on the fly generated by CASTEP, but you could use your favourites.

All the files needed for the simulation can be downloaded from

- Fe_NCP19_PBE_OTF.usp
- O_NCP19_PBE_OTF.usp
- hematite.out
- hematite.dat

Dat

The first two blocks are the cell and atoms positions. You might see that iron atoms are labelled Fe1 or Fe2 depending if they will be treated as spin up atoms or spin down atoms.

The third block is

```
%BLOCK SPECIES
  Fe1 Fe 26 13 11.000000
  Fe2 Fe 26 13 11.000000
  O O 8 4 11.000000
%ENDBLOCK SPECIES
```

This block tells us which atom type we should assign to Fe1, Fe2, and O, their atomic number and how many NGWFS we should use for each atom type as well as their radius. For strongly correlated systems NGWFs radius of 11.00 bohr is suggested. The next block is

```
%BLOCK SPECIES_ATOMIC_SET
  Fe1 "SOLVE conf=3s2 3p6 3d5 4s0 4p0 SPIN=+5 CHARGE=+3"
  Fe2 "SOLVE conf=3s2 3p6 3d5 4s0 4p0 SPIN=-5 CHARGE=+3"
  O "SOLVE INIT SPIN=0 CHARGE=-1"
%ENDBLOCK SPECIES_ATOMIC_SET
```

This block setup the initial configurations for the atoms. Fe1 and Fe2 atoms will have a charge of +3 and a spin of + or - 5. We have also setup the electronic configurations. The atomic solver generates the first guess for the density kernel for the first scf iteration, it does it by calculating the atomic density.

The next block is the Hubbard block where we setup the DFT+U parameters

```
%block hubbard
  Fe1 2 6.0 0.0 -10.0 0.00 5.0
  Fe2 2 6.0 0.0 -10.0 0.00 -5.0
%endblock hubbard
```

In this block we setup on the d ($l=2$) orbitals a U value of 6 the rest are the default parameters except the last column (5,-5). The last column is the spin-splitting functionality that help certain system to break the spin symmetry in AFM configurations.

The rest of the blocks are to tell ONETEP for which atom types it has to compute the ldos and dos. Related to this functionality we also have to parameters called dos_smear and pdos_max_l. The First relates to the gaussian smearing applied to the dos and the max l quantum number to compute the pdos for.

The rest of the parameters are self-explanatory with some exception such as `maxit_palsen_mano` and `maxit_hotelling` which are related to the diagonalisation library.

If require more explanation you can refer to the ONETEP keyword database.

Pseudopotentials

It is important to know is that the number of NGWFs depends on your Pseudopotentials. If you use your own be sure to modify them.

7.3 Evaluating the outputs

ONETEP will generate many files based on we configured the calculations, but for this tutorial we will be focusing on few ones.

- `.out`: the main output file
- `DOS.txt`: Density of states file
- `LDOS.txt`: Local density of states file
- `PDOS.txt`: Partial density of states file
- `spindensity.cube`: Cube file necessary to visualise the spin density

7.3.1 What to look for in the main output file

First thing is to check is the whether the atoms are in the configuration you wanted to (in our case a Fe^{3+} with spin UP or DOWN). This can be seen by looking for this block for each atom (shown here the down Fe atom)

```
Config String: 3s2 3p6 3d5 4s0 4p0 SPIN=-5 CHARGE=+3
Orbitals (num,occ): 5 2.00 6.00 5.00 0.00 0.00
Orbitals (num,l): 5 0 1 2 0 1
```

As you can see it was properly set as we have the occupancies as we wanted them to be

The second step is, as explained in the DFT+U part, the occupancies for the majority spin (either UP or DOWN for different Fe atoms) has to be > 0.5 while < 0.5 for the minority spin. This is very important to allow DFT+U to do its job and it can be checked in the following table by looking at the diagonal elements.

```
#####
DFT+U information on atom 1 of Hubbard species Fe1
#####
Occupancy matrix of Hubbard site 1 and spin 1 is
m_1 = -2 -1 0 1 2
0.98677750 0.00064944 0.00008254 -0.00024296 -0.00055787
0.00064545 0.99097329 -0.00010565 0.00019818 0.00042874
0.00006971 -0.00009587 0.97935271 0.00006742 0.00560449
```

```

-0.00024459 0.00020075 0.00009975 0.99062614 -0.00144761
-0.00055154 0.00042561 0.00602074 -0.00142407 0.96101391
#####
Occupancy matrix of Hubbard site 1 and spin 2 is
m_1 = -2 -1 0 1 2
0.17854088 -0.10543254 -0.00014018 -0.01900359 -0.00209684
-0.10478392 0.27925352 -0.00100259 0.00081665 -0.01300318
-0.00018286 -0.00127346 0.07647309 0.00442738 0.00630459
-0.01913097 0.00082721 0.00279433 0.25746916 0.08014266
-0.00208260 -0.01299767 0.00576137 0.07901457 0.10302424
#####
Total occupancy of Hubbard site 1 is 5.80350444 e
Local magnetic moment of Hubbard site 1 is 4.01398264 mu_B
DFT+U energy of Hubbard site 1 is 0.08324891 Ha
#####

```

Another important thing to check are the bands occupancies. Hematite is a semiconductor with a 2 eV band gap, we would then expect to have fully occupied bands and unoccupied virtual bands. If we were to treat it as a metal we could expect fractional occupancies occurring, but that would be physically wrong for our system.

If you look at the band occupancies for both spin up and down, you can see that we indeed obtain fully occupied bands and unoccupied bands. This reassure us that the structure we obtained is chemically and physically sensible. The small number you see in the “unoccupied” bands is due to numerical errors.

This is for spin 1 (up)

```

Orbital | Energy (Eh) | Occupancy

1 -3.063849888 1.0000000
.....
1791 0.357172741 0.9999996
1792 0.357225976 0.9999996
1793 0.357724185 0.9999996
1794 0.357733628 0.9999996
1795 0.358380935 0.9999997
1796 0.358393032 0.9999997
1797 0.360316821 0.9999996
1798 0.360327604 0.9999996
1799 0.366378308 0.9999993
1800 0.366387961 0.9999993
..... - gap at 0K - .....
1801 0.438702297 0.0000005
1802 0.438704363 0.0000005

```

```
1803 0.438802525 0.0000005
1804 0.438803496 0.0000005
```

This is for spin 2 (down)

```
Orbital | Energy (Eh) | Occupancy
1 -3.062779703 1.0000000
.....
1791 0.357161543 0.9999996
1792 0.357225094 0.9999996
1793 0.357691641 0.9999996
1794 0.357691940 0.9999996
1795 0.358354526 0.9999997
1796 0.358355045 0.9999997
1797 0.360301824 0.9999996
1798 0.360302485 0.9999996
1799 0.366369520 0.9999993
1800 0.366387326 0.9999993
..... - gap at 0K - .....
1801 0.438975662 0.0000005
1802 0.438982779 0.0000005
1803 0.439082674 0.0000005
1804 0.439094865 0.0000005
```

And as last we should also check that we obtain a band gap and its value is close to experiment. This can be seen from the output by looking for these lines.

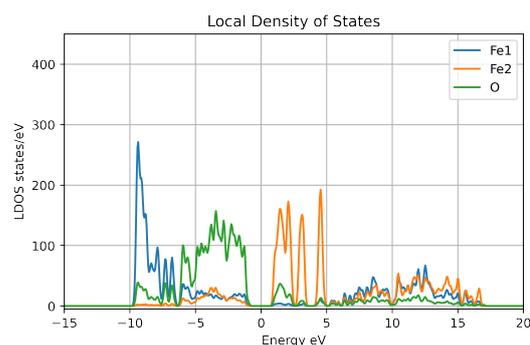
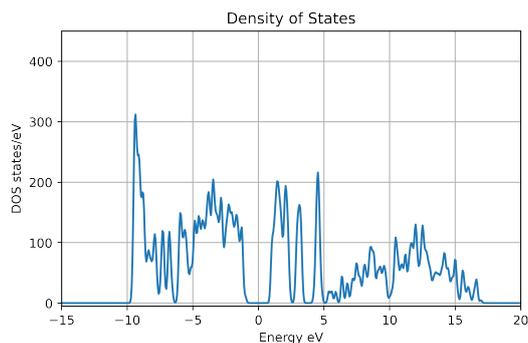
Why do we get two band gaps? Because we are studying a magnetic system, we get a band gap for each spin channel and for an AFM material the bandgap should be the same (numerical errors aside).

```
HOMO-LUMO gap: 0.072314337 Eh
HOMO-LUMO gap: 0.072588336 Eh
```

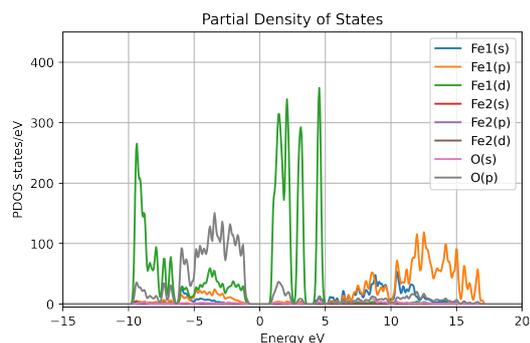
7.3.2 DOS and PDOS

Next step is to plot the density of states. It will tell us the distribution of electrons and states in our system we would expect to have gap around 0 of about 2 eV as this is the band gap of the material.

We indeed obtain a gap between the states but it does not tell us much more. To obtain more information we will be plotting the local density of states (LDOS)



and the partial density of states (PDOS).



From the local density of states we can immediately notice that the lowest lying bands in the plot are mostly made of Fe1 bands but, this is very important, the top of the valence band is made mostly by O bands. The bottom of the conduction band is made of Fe2 orbitals. This allow us to classify hematite as a charge transfer insulator between the oxygen and the iron atoms. What if we would like to know which atomic orbitals contribute the most in this charge transferring? We need to plot the PDOS.

It will project the bands into the atomic components, in this way, as you can see in the graph the top of the valence band is dominated by O 2p states while the bottom of the conduction band by Fe2 3d states.

7.3.3 Mulliken population analysis

The Mulliken population analysis is a very good tool to understand if our system is behaving correctly. In an AFM material the total spin should be 0 and the local spin should be the same for the same atoms. In this case we have two different types The spin up and down Fe atoms. The absolute value of the local spin should be the same with just different signs.

The material is also charge neutral and we would expect that the similar atoms should carry similar charges.

Species	Ion	Total	Charge (e)	Spin (hbar)
0	1	6.906	-0.906	0.01
0	2	6.906	-0.906	0.01
0	3	6.906	-0.906	-0.01
0	4	6.905	-0.905	-0.01
0	5	6.906	-0.906	-0.00
0	6	6.907	-0.907	-0.00
Fe	7	14.641	1.359	2.20
Fe	8	14.641	1.359	2.20
Fe	9	14.641	1.359	-2.20
Fe	10	14.641	1.359	-2.20

As you can see from this snapshot we do indeed obtain the same charge and same spin for all similar atoms as we would expect.

7.3.4 Spin Density

Now it is time to visualise the spin density which is the total electron density of electrons of one spin minus the total electron density of the electrons of the other spin. We would like to visualise it to know if we obtained the afm state we wanted, the up-down-down-up configuration.

You can directly open and visualise The .cube generated at the end of the calculations with VESTA, VMD or lots of other softwares.

You can see from the picture that we did get the AFM states with +--+ configuration as we wanted.

7.3.5 What to do next

The tutorial is now complete, but you could still move forward. What can you do next?

ONETEP outputs many. more information than the one showed you here.

You can plot

- The electrostatic potential
- The orbitals
- The electron density

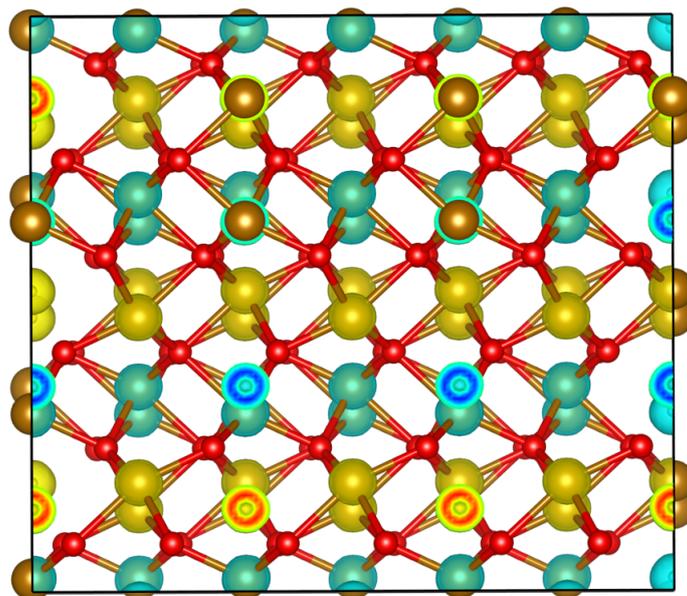


Fig. 7.2: Fig.2: Hematite spin density, blu spheres refers to atom with up spin and yellow to down spin

- The band plot

You can then relax the structure and recompute the properties to see what changed and how.

We have chosen to use $U=6$, but you could try different U values and see how it affects the system.

7.3.6 References

TUTORIAL 10: SIMULATION CELL RELAXATION

Author

Chris-Kriton Skylaris

Date

August 2023

8.1 Introduction

This tutorial demonstrates how to:

- Use ONETEP to relax the simulation cell of a crystalline material

8.2 Cell relaxation of bulk crystalline silica

This calculation will relax the lattice of a silica (SiO_2) simulation cell, which is depicted below:

The simulation cell of silica used in this tutorial. The silicon atoms are beige and the oxygen atoms are red.

8.2.1 Input file keywords

The input file, which is provided, is called “silica96.dat” and contains 96 atoms in total.

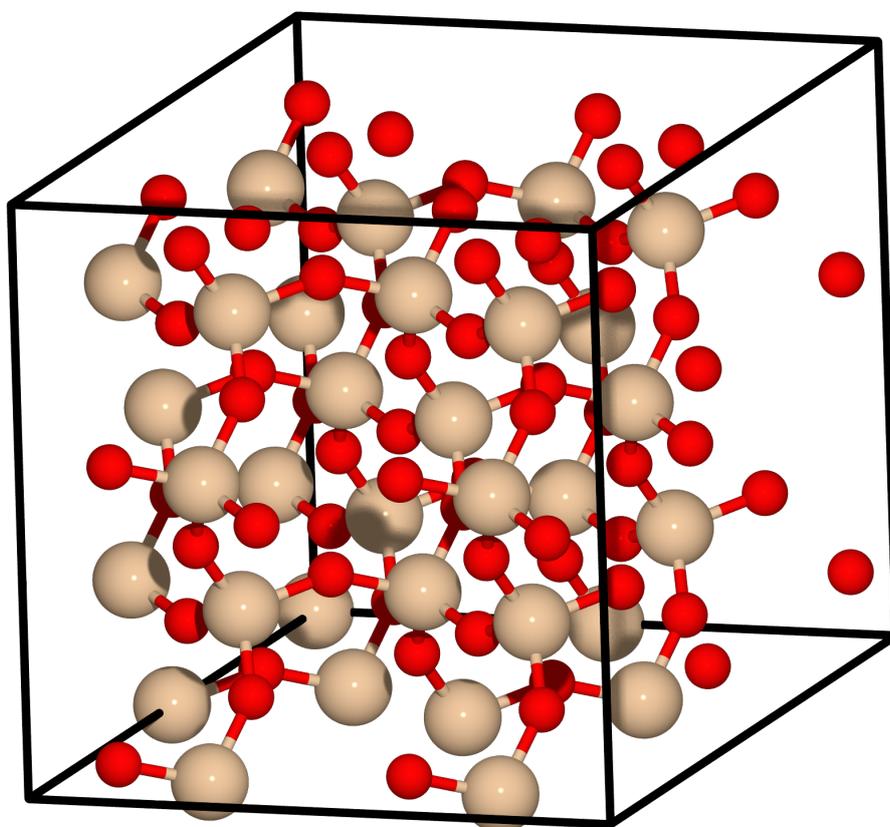
To perform cell relaxation the

task STRESS

keyword is required.

You will notice in the input file also the following keywords related to the cell relaxation:

stress_tensor T



stress_elasticity F
stress_relax T
stress_assumed_symmetry tetra1
stress_relax_atoms T

Where *stress_tensor T* instructs the code to compute the stress tensor, while the calculation of elastic constants is turned off with *stress_elasticity F*. In this calculation the simulation cell will be relaxed (in order to determine the optimal lattice vectors) and this is denoted by *stress_relax T*. In this calculation, in addition to the simulation cell we want to relax also the coordinates of the atoms and for this we use the keyword *stress_relax_atoms T*. It is worth noting here that the atomic coordinates would also be relaxed if the *stress_relax_atoms* was set to *F* (False), but in this case they would only be “stretched” to be commensurate with the change in the lattice vectors, in other words they would retain the same fractional coordinates. On the other hand if the *stress_relax_atoms* is activated the coordinates of the atoms are fully relaxed and are not constrained to remain equal to the fractional coordinates they had at the start of the calculation.

Finally the *stress_assumed_symmetry tetra1* instructs the code to assume a particular symmetry for the simulation cell and maintain this symmetry during the cell relaxation. If the symmetry of the cell is known and is supported by the code (see the user manual) it is important to activate it with this keyword as it will significantly reduce the number of single point energy calculations that will be performed.

8.2.2 Running the calculation

Now run the calculation and examine the output. Let's examine the output step by step noting the various stages of the calculation.

At the very beginning some information about the initialisation of the calculation is produced such as:

- *PSINC grid sizes*: information about the grids used for the psinc basis functions
- *Atom SCF Calculation for...*: here the code initialises the NGWFs with atomic orbitals created specifically for the valence electrons of the chosen pseudopotentials and confined within the NGWF spherical regions.
- *STRESS: undistorted cell* : this is the beginning of the very first energy calculation from which calculations with applied stains will be subtracted to compute the stress tensor.
- *Atomic positions optimised prior to stress calculation*: a geometry relaxation is performed first since we have specified *stress_relax_atoms T*

You will notice that this calculation takes 30 NGWF iterations to converge as a very tight convergence criterion for the NGWFs (*ngwf_threshold_orig 1.e-7*) has been applied in the input. This was done to ensure very accurate forces and it may be a bit extreme, but it is better to be on the safe side.

After this energy evaluation the code computes the forces and compares them with the threshold that has been defined for geometry relaxation. In this calculation you will notice that the forces are small and below the tolerance ($|F|_{max} 1.e-3 \text{ Eh/Bohr}$) that has been set for the convergence

of the geometry. As a result the code reports that the geometry relaxation has been completed with the message “*Geometry optimization completed successfully*”.

Calculation of the stress tensor

Then the calculation proceeds to evaluate the energy of the system at different distortions (strains) of the lattice vectors in order to compute the stress tensor. The beginning of this procedure is denoted by the “*STRESS: distorted cells*” message.

The stress tensor computed is summarised at the end of the first iteration of cell relaxation:

```
stress_tensor: iteration 1
```

This stress tensor is now used to change the simulation cell. Again a geometry relaxation is performed which converges at the first step. Then several single point energy calculations follow to compute a new stress tensor until we obtain the summary of the second iteration:

```
stress_tensor: iteration 2
```

Finally, we see that in the third iteration the cell has been relaxed. The relaxed cell is printed:

```
Relaxed cell:
```

```
bohr
```

```
19.07668106 0.00000000 0.00000000
```

```
0.00000000 19.07668106 0.00000000
```

```
0.00000000 0.00000000 27.83315074
```

This completes tutorial 10.

Files for this tutorial:

- T10_files.zip

TUTORIAL 11: ELECTRIFIED ELECTRODE-ELECTROLYTE INTERFACES UNDER POTENTIOSTATIC CONTROL

Author

Brad Ayers, Arihant Bhandari

Date

August 2023

9.1 Introduction

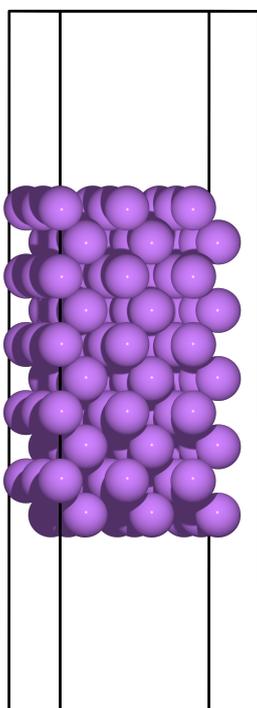
This tutorial endeavours to provide a comprehensive and illustrative example, highlighting ONETEP's ability to conduct potentiostatic calculations through the utilisation of a grand canonical DFT algorithm [Bhandari2021]. Alongside this written tutorial we have provided a Jupyter notebook that will guide you through the process of setting up and analysing a potentiostatic calculation.

Moreover, this tutorial will demonstrate the feasibility of conducting these calculations within a solvent and electrolyte medium, employing the Fisicaro soft-sphere continuum model [Fisicaro2017] for solvation and the neutralisation via the grand canonical electrolyte (NECS) model [Bhandari2020].

9.2 Setting up the calculations

We will begin this tutorial by creating lithium surfaces using a tool of your preference. For the purpose of illustration, we have employed the ASE [Larsen2017] to generate a 10-layer BCC [100] lithium surface. This can be simply visualised within the provided Jupyter notebook, and will look as follows:

Notably, in this tutorial, we've chosen a (3x3x10) supercell with a 50 Å vacuum in total, a decision that allows us to employ 9 Bohr NGWF radii, whilst ensuring ample volume for sufficient electrolyte to neutralise our surface.



9.2.1 Input files

For the purposes of this tutorial ONETEP will require only two files to work:

1. A .dat file, which contains all the information about your system (Atomic positions & Lattice vectors), as well your simulation parameters.
2. A Pseudopotential file, here we will employ the Norm-conserving on the fly generated CASTEP ones, but this is up to the users discretion.

Both of the aforementioned files, as well as a few output files can be downloaded below:

- Keywords.dat
- Li_surface.out
- Li_surface.dat
- T11_workbook.ipynb
- Li_NCP19_PBE_OTF.usp

Note that the output files required for the visualisation section of this tutorial will have to generated by the user themselves, and will need to be copied to the same directory as the Jupyter notebook.

Furthermore, Li_surface.dat is the .dat file used for your job submission, and Keywords.dat is simply the keywords that are taken in by the ASE script to generate said Li_surface.dat file.

9.2.2 The Dat file

grand canonical Parameters

Upon opening the provided Li.dat file, you will encounter the standardx parameters that ONETEP users are already familiar with: Task, Cutoff_energy, Lattice_cart, and Positions_frac.

However, after these, a new section of interest awaits, aptly named:

```
grand canonical Parameters
~~~~~

edft_grand_canonical T
edft_reference_potential -1.20 eV
edft_electrode_potential 0.20 V
```

Within the grand canonical formalism in ONETEP, there are two main parameters of interest:

1. Chemical Potential of the Reference Electrode

Represented as μ_e^{ref} , this value (-1.20 eV in this example) corresponds to the reference electrode's chemical potential. It serves as the benchmark against which all applied potentials are measured.

2. Applied Potential to the Working Electrode

In this case, a potential of 0.20 volts (U) is applied to the working electrode, giving us a working electrode potential of -1.40 eV (μ_e).

These parameters find application in the following equations:

- **Chemical Potential of the Working Electrode (μ_e):**

$$\mu_e = \mu_e^{\text{ref}} - e \cdot U$$

- **Number of Electrons (N_e):**

$$N_e = \sum_i f_i = \sum_i \frac{1}{1 + \exp(\beta(\epsilon_i - \mu_e))}$$

- **Charge on the System (q):**

$$q = Z_{\text{ion}} - e \cdot N_e$$

Hence, by defining the applied potential and reference electrode potential, we can establish the chemical potential of the working electrode. It is worth noting here that only with zero applied potential ($U = 0$) and the reference value set to the potential of zero charge ($\mu_e^{\text{ref}} = \mu_e^{\text{PZC}}$) will we get zero charge on the system ($q = 0$).

As a result, we gain the ability to compute the number of electrons within our system using the Fermi-Dirac distribution. Where, ϵ corresponds to the eigenvalues of the KS-equations, and μ_e is the chemical potential of our working electrode. Thereby allowing us to accurately calculate the charge present in the quantum system for a given applied potential.

For a more detailed explanation of the grand canonical formalism, please refer to [Bhandari2021].

Solvation Parameters

This section will highlight some of the more unintuitive parameters that are required to conduct a solvated calculation. They will be presented in the order in which they appear in the data file as done for the grand canonical parameters above.

This block will highlight the solvation parameters themselves, with the following block highlighting the solvent parameters

```
is_include_apolar T # Enables solvent-accessible_
↳surface-area (SASA) approximation
is_smeared_ion_rep T # Enables the smeared ion_
↳representation due to multigrid solver
is_implicit_solvent T # Enables the solvation model
is_solvation_properties T # Provides extra properties i.e.
↳ electrolyte concentration
is_dielectric_function soft_sphere # Utilising the Fisicaro soft-
↳sphere model
```

For a much more detailed overview of each key parameter and its function, please refer to the documentation here: [Solvent and Electrolyte model](#)

For the solvent itself we have the following parameters:

```
is_bulk_permittivity 90.7          # permittivity value of our
↳solvent          (ethylene carbonate in this case)
is_solvent_surf_tension 0.0506 N/m # Surface tension of the
↳selected solvent (ethylene carbonate in this case)
```

Obviously these values will vary depending on the solvent of choice, and can be found in the literature.

Electrolyte Parameters

```
is_pbe          full          # Enables the poisson-boltzmann
↳solver for electrolyte calculations
is_pbe_bc_debye_screening T    # Enables the Debye screening
↳boundary condition
is_pbe_temperature 298.15     # Sets the temperature of the
↳Boltzmann ions
is_pbe_neutralisation_scheme counterions_auto
```

The above parameters are required to conduct electrolyte calculations, and are fairly self-explanatory. However, it is worth noting that there are several neutralisation schemes available, and the one chosen here is recommended for most cases.

Additionally, the following parameters are required by DL_MG to conduct the multigrid calculations:

```
mg_max_res_ratio 1000.0
mg_max_iters_vcycle 500
mg_max_iters_newton 300
mg_vcyc_smoother_iter_pre 4
mg_vcyc_smoother_iter_post 4
```

These parameters are fairly ubiquitous, and are not expected to change between calculations. For further information, please refer to the documentation here: [DL_MG](#)

lastly, we have the following ions block, that define our electrolytes and their concentrations:

```
%BLOCK SOL_IONS
Li    +1 1.0 # all concentrations are in Molar
PF6   -1 1.0
%ENDBLOCK SOL_IONS
```

Note that additional blocks can be added here to further adapt the model such as defining the soft-sphere radii of your system.

9.2.3 Analysis

Now that we have established the parameters required to conduct a solvated calculation, we can proceed to analyse the results obtained from either your own calculations or the provided ones.

The first file of interest is the `Li.out` file, which contains the usual ONETEP output you might be familiar with, but will include additional that are of interest to us. Firstly within CG-DFT inner loop, we can find the following:

```

Iter                               Commutator      Grand pot. (L=E-TS-
↪muN)  DeltaL
# 2                                0.000001192422  -1269.
↪59513707094357  -2.05E-12

Step                               =      0.000010819110
Energy (E)                         =     -1297.517766665092
Entropy (-TS)                      =      -0.065437957621
Chemical potential (-muN)          =      27.988067551769
Grand potential (L=E-TS-muN)      =    -1269.595137070944
Est. 0K Energy 0.5*(E+L)          =    -1283.556451868018
Charge on quantum system           =      -3.995764707545
Residual Non-orthogonality        =     -0.000000000000

```

If users are familiar with the canonical inner loop printing this will look familiar, however, there are a few key differences:

- **Grand Potential** ($L = E - TS - \mu_e N$):

This is the grand potential of the system, and is the quantity that is minimised in the grand canonical formalism, rather than the Helmholtz free energy as in the canonical formalism.

- **Chemical Potential of the Working Electrode** ($-\mu_e N$):

This is the chemical potential of the working electrode, and is the quantity that is constant in the grand canonical formalism, in direct contrast to the canonical formalism, where the number of electrons is constant.

- **Charge on the System** (q):

This is the charge on the quantum system and is an important value to consider when analysing the results of a potentiostatic calculation, as surface chemistry is dependent on the charge of one's system.

Another crucial section of the `Li.out` file is the bulk concentration of the electrolyte species, which can be found in the following section:

```

+----- Chemical potential -----
↪-----+
| # | Name | Bulk conc. | mu_ideal | mu_excess | mu_
↪total |
|-----|
↪-----|

```

(continues on next page)

(continued from previous page)

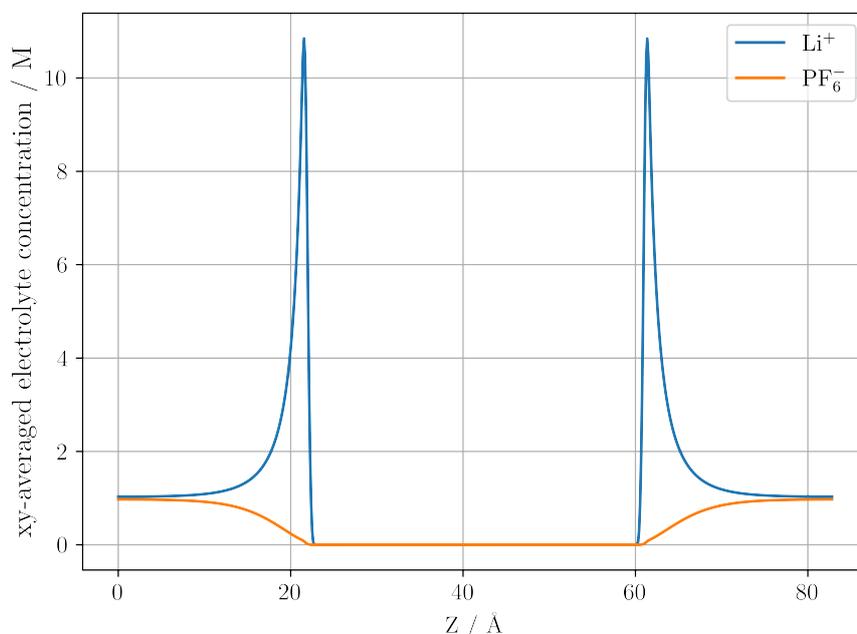
Entropic contribution	:	27.92262959414529	
Total free energy	:	-1269.59513707094357	
----- LOCAL ENERGY COMPONENTS FROM MATRIX TRACES -----			
Pseudopotential (local)	:	-443.84297910438886	
Hartree	:	-24.92214858378443	

Integrated density	:	543.99576470754505	
Integrated spin density	:	0.00000000000165	
Integrated spin density	:	0.00000001531685	
Local density approx. <S^2>	:	0.00000000765765	
Integrated density tr(KS)	:	543.99576470754528	
Integrated spin tr(KS)	:	0.00000000000153	

We won't delve into the details of each energy component here, but it is worth noting that the solvent and electrolyte contributions are present in the energy breakdown. For a more detailed explanation of each energy component, please refer to [Bhandari2021].

visualisation

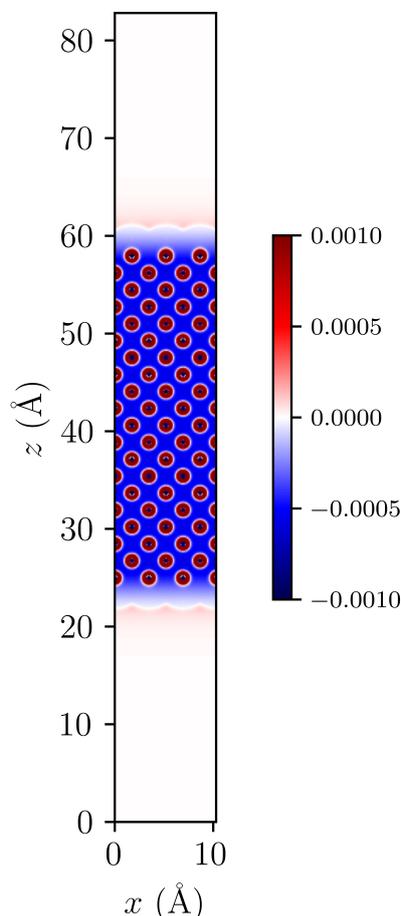
Whilst a majority of the analysis can be conducted using the `Li.out` file provided, there are a few visualisations that can be conducted using the user generated `Li_out_bion_conc_species_1.dx` and `Li_out_bion_conc_species_2.dx` files. Such as plotting the xy-averaged concentration profiles of the electrolyte species, it is worth stating that all code used to generate these plots can be found in the provided Jupyter notebook.



Looking at the figure above we can see that the concentration of the lithium electrolyte species is highest (12M) at the surface of the electrode, and decreases as we move away from the surface, forming a double layer along our surface. In contrary we can observe the PF6 electrolyte species

has a concentration approaching 0M at the surface, and increases as we move away from the surface, note that these are local concentrations and not bulk concentrations.

We can also visualise the densities of our lithium slab and the electrolyte species to illustrate this double layer effect by adding the `Li_out_bion_density_total.dx` and `Li_in_rho.dx` files together and running the provided code:



Here we can see that there is a positive charge density (our lithium ions) at the surface of the electrode, that diffuses in intensity as we move away from the surface.

9.3 Where to go next?

This tutorial has provided a brief overview of the grand canonical formalism, and how it can be employed to conduct potentiostatic calculations, as well as highlighting a few key visualisation techniques that can be conducted on the output files.

However, there are a few key areas that have not been covered in this tutorial that the reader should explore, such as:

1. **Defining their reference potentials** (μ_e^{ref}):

In this tutorial, we have defined the reference potential as -1.20 eV, however, this value is dependent on the reference electrode of choice, and can be found in the literature or by calculating one's own reference potential using the method highlighted in [Bhandari2021].

2. Cycling through potentials:

In this tutorial, we have only conducted a single potentiostatic calculation, however, it is possible to conduct a series of calculations at different potentials, and calculate properties such as capacitance etc.

9.3.1 References

TUTORIAL 12: QUANTUM EMBEDDING WITH (TIME-DEPENDENT) EMBEDDED MEAN-FIELD THEORY: HYDROGENATION AND EXCITATIONS OF PENTACENE

Author

Joseph Prentice

Date

August 2023

10.1 Introduction

10.1.1 The utility of quantum embedding

Although ONETEP [ONETEP2020] makes first principles calculations of systems containing thousands of atoms feasible, particularly when using semi-local exchange-correlation functionals, it can still be very computationally costly to treat entire systems of this size with higher level theory, such as hybrid functionals, which require the computation of exact exchange. This is particularly relevant for excited state calculations, due to the well-known underestimation of the band gap by semi-local DFT, making excitation spectra performed with semi-local DFT quantitatively, or sometimes qualitatively wrong. However, if the physics/chemistry of interest in the system is known or expected to be localised to a particular subregion – we call this the ‘active region’ – with the rest of the system acting as an environment influencing the active region, quantum embedding provides a way to achieve hybrid accuracy with a significantly reduced cost.

This is achieved by treating the active region alone with the higher level of theory, with the rest of the system (the environment) treated at a lower level of theory. In ONETEP at present, this translates to treating the active region with hybrid DFT, and the environment with semi-local DFT. This is done within a single self-consistent calculation, to ensure that the two regions are able to influence one another. As hybrid DFT is only performed on the active region, which is typically small compared to the environment, the computational cost is significantly reduced.

10.1.2 Embedded mean-field theory

There are many different quantum embedding schemes; the scheme used in ONETEP is embedded mean-field theory (EMFT), originally proposed by Fornace et al. [Fornace2015] This scheme has several advantages, including two particularly relevant to ONETEP: firstly, it partitions the system at a basis-set level, which works well with ONETEP's atom-centred NGWF basis. Secondly, it is a mean-field theory throughout, like DFT, which means that many existing methods based on DFT can be easily modified to accommodate EMFT. This includes linear-response time-dependent DFT (LR-TDDFT), which allows electronic excitations to be computed – we refer to this as TD-EMFT. For more details, please see the original EMFT paper [Fornace2015], and the papers implementing ground state EMFT [Prentice2020] and TD-EMFT [Prentice2022] in ONETEP.

EMFT is also fully compatible with the implicit solvent methods available within ONETEP, allowing for multi-level modelling of systems.

10.1.3 Pentacene

In this tutorial, we will look at both ground state EMFT and TD-EMFT, with the pentacene molecule as our test case. Ground state EMFT is demonstrated by looking at the terminal hydrogenation energy of pentacene, following Prentice et al. [Prentice2020], and TD-EMFT is demonstrated by looking at the first excitation energy of pentacene-doped p-terphenyl, following Prentice. [Prentice2022]

10.2 Ground-state EMFT: terminal hydrogenation of pentacene

10.2.1 Non-EMFT baseline calculations

The terminal hydrogenation reaction for pentacene involves two hydrogen atoms becoming bonded to the two carbon atoms at one end of the pentacene molecule.

Before using EMFT, we must first obtain the terminal hydrogenation energy without EMFT, with everything treated with first the PBE, then the B3LYP functionals. The input files required are `Pentacene.dat`, `HydrogenatedPentacene.dat`, and `H2.dat`.

If you look at these files, you will notice that we have several different labels for carbon atoms (C1, C2, etc.) and similar for hydrogen atoms – these will be used later to vary the size of our active region, by selecting different atoms to be included within it.

To begin, simply run the input files as they are to obtain the ground state energy for the three structures at the PBE level. The hydrogenation energy can then be obtained as:

$$\Delta E_{\text{hyd}} = E_{\text{Hydrogenated Pentacene}} - (E_{\text{Pentacene}} + E_{\text{H}_2}). \quad (10.1)$$

Make sure you save the `.tightbox_ngwfs` file from the `H2.dat` calculation for use later!

Now repeat this for B3LYP. In all three `.dat` files, change `xc_functional` to be B3LYP, and add the following keywords/blocks to set up Hartree-Fock exchange (deleting any species labels in the `species_swri...` block that are not present in that particular structure):

```
%block swri
  for_hfx 3 12 V 12 12 WE2
%endblock swri

%block species_swri-for_hfx
C
C1
C2
C3
H
H1
H2
H3
%endblock species_swri-for_hfx

hfx_use_ri for_hfx
hfx_max_l 3
hfx_max_q 12
```

For more details on the meaning of these keywords, see the HFX documentation. The most important point for our purposes here is that we can control which atoms are included in the computation of HFX through the `species_swri...` block – this will become important for our EMFT calculations.

For reasons of stability, it is best to start the B3LYP H2 calculation from the PBE-optimised NGWFs: bring back the PBE-optimised `.tightbox_ngwfs` file, and add `read_tightbox_ngwfs : T` to `H2.dat`. Make sure you still keep a copy of the PBE-optimised `.tightbox_ngwfs` file safe, as we will need it later.

Run these three calculations and compute the hydrogenation energy at the B3LYP level. ..

You should obtain a value around ??????.

10.2.2 EMFT calculations

4 carbon atoms

We can now start to use EMFT to see if we can get close to the B3LYP result without treating the entire molecule with B3LYP. Initially, our active region will just include the 4 C atoms closest to the site of the reaction, and the hydrogen atoms bonded to them. We will only need to do EMFT calculations for `Pentacene.dat` and `HydrogenatedPentacene.dat` – the hydrogen molecule will always be in the active region, so should always be treated with B3LYP, although there is one subtlety which will be introduced shortly.

To turn on EMFT, change `xc_functional` back to PBE, and add the following keywords to `Pentacene.dat` and `HydrogenatedPentacene.dat` (keep the other modifications you

already made for HFx for the moment):

```
use_emft                T
use_emft_follow         T
use_emft_lnv_only       T
block_orthogonalise     F
active_xc_functional    B3LYP
parallel_scheme         HOUSE

%block species_ngwf_regions
C1 H1
C C2 C3 H H2 H3
%endblock species_ngwf_regions
```

A brief explanation of each keyword:

- `use_emft`: this turns on EMFT, so that the active region and environment are treated with different levels of theory.
- `use_emft_follow`: this toggles whether a non-EMFT calculation at the lower level of theory (in this case, PBE) is done first, and then uses that as a starting point for the EMFT calculation. For this to happen, the value should be `T`.
- `use_emft_lnv_only`: this toggles whether EMFT is used to optimise both the NGWFs and the density kernel (value is `F`), or just the density kernel (value is `T`). Typically, EMFT is only used to optimise the density kernel (`T`), as NGWF optimisation is poorly behaved under EMFT – NGWFs can unphysically optimise towards the region described by the level of theory that predicts the lowest energy, and the block orthogonalisation procedure designed to counteract this (discussed shortly) makes the NGWF optimisation stall. The NGWFs are optimised at the lower level of theory, and then fixed – the error this introduces is typically <1% of the difference between the high and low levels of theory. For more details, see Prentice et al. [Prentice2020]
- `block_orthogonalise`: this toggles whether a block orthogonalisation procedure is applied to the NGWFs before using EMFT. This transforms the NGWFs of the environment so that they are orthogonal to the NGWFs of the active region, so the off-diagonal blocks of the overlap matrix are 0. This prevents the emergence of unphysical solutions that can occur in some systems. For this to happen, the value should be `T`.
- `active_xc_functional`: this selects the functional that will be used in the active region, whilst `xc_functional` selects the functional used in the environment.
- `parallel_scheme`: this decides how MPI processes should be split between the regions. `HOUSE` means that the processes are distributed proportionally to the number of atoms within each region; `SENATE` means that the processes are distributed equally between all regions; and `NONE` means that each region will use all the processes in turn. `HOUSE` is strongly preferred.
- `block species_ngwf_regions`: this assigns species to regions, with one region per line. The first line is the active region by default.

The first four keywords should be turned to `T` in the order they are listed in. The first three keywords should almost always be `T` for an EMFT calculation, with `block_orthogonalise`

turned on if the calculation proves unstable without it. Here, we leave it off.

We also need to modify the HFX set-up to match the fact we only want HFX in the active region. To do this, simply delete any species in the `species_swri...` block that are *not* in the active region. Remember that each species in the active region should be on its own line in the `species_swri...` block, whereas all the species in the active region should be on the first line of the `species_ngwf_regions` block.

Once these additions/modifications have been made, run the calculations for pentacene and hydrogenated pentacene.

Before we can compute the hydrogenation energy from these results, there is one more subtlety. As we optimised the active region NGWFs at the lower level of theory, but the active region density kernel with the higher level of theory, we need to do the same in our hydrogen molecule for consistency. To do this, bring back the `H2.tightbox_ngwfs` file you saved from the PBE calculation earlier, and re-run your B3LYP `H2.dat` calculation with the following modifications/additions:

```
read_tightbox_ngwfs : T
maxit_ngwf_cg : 0
```

You can now use these three results to compute the hydrogenation energy with an active region of this size. ..

You should get a result of around ???????

Larger active regions

Next, expand the active region to include the 8 carbon atoms closest to the reaction site. To do this, add the C2 and H2 species to the active region (remember to remove them from the environment region!), and modify the HFX set-up to match. Rerun the pentacene and hydrogenated pentacene calculations, and compute the hydrogenation energy (you don't need to rerun the hydrogen molecule calculation, as you can just reuse the result obtained using PBE-optimised NGWFs).

Finally, expand the active region further to include the 12 carbon atoms closest to the reaction site, by adding C3 and H3 to the active region. Re-calculate the hydrogenation energy.

If you plot the hydrogenation energy vs. the size of the active region, you should see the hydrogenation energy approach the full B3LYP result. ..

,getting within ?????? with only around half of the molecule included in the active region.

This demonstrates the ability of EMFT to obtain high level results at a reduced cost, even when the boundary between regions cuts through covalent bonds.

This also demonstrates the importance of selecting the appropriate active region. In systems made up of weakly bonded parts (e.g. molecular crystals, solvated systems), the appropriate active region will often be obvious – it will be the molecule or molecules of interest (examples of multiple-molecule active regions could include a nearest-neighbour dimer or a solute along with some nearest-neighbour solvent atoms). In extended covalent or ionic systems, the choice of active region may be more difficult, and should be carefully converged, in a similar way to that shown in this tutorial.

Further investigations

To further investigate the use of EMFT in ONETEP, you could look at the effects of:

- changing the active region further – perhaps including more C atoms, excluding H atoms, etc.,
- using block orthogonalisation,
- using other functionals for either the high or low level of theory – semi-local functionals can be used for the higher level, although this is of course not expected to produce a significant advantage,

and many other possibilities.

10.3 TD-EMFT: excitations of pentacene-doped p-terphenyl

10.3.1 Non-EMFT benchmark

Here, we will be looking at the S_0 to S_1 transition in pentacene, which is the lowest excited state observed in TDDFT. This is significantly affected by the environment. In particular, we are interested in pentacene-doped p-terphenyl, which is of interest for room-temperature maser applications, and how the p-terphenyl environment affects the excitation energy.

To give us a reference for isolated pentacene, we first need to perform a high-level TDDFT calculation for pentacene. We will again use B3LYP as our high-level theory. The input file is `Pentacene_isolated.dat` – as this tutorial assumes you are already familiar with running TDDFT calculations with ONETEP, we will not go into any detail, and this calculation can just be run as it is. ..

You should find the lowest energy excitation to be ?????.

10.3.2 TD-EMFT calculation

We now perform a TD-EMFT calculation for a pentacene molecule surrounded by 6 p-terphenyl molecules, as extracted from the pentacene-doped p-terphenyl molecular crystal. The input file is `Pentacene_in_p-terphenyl.dat`. This can be run as it is, but one point regarding EMFT should be noted first. The general set-up of the EMFT calculation is precisely the same as for ground state EMFT, with one addition: the `species_tddft_kernel` block. By using this block, we can specify which species we will restrict our excitations to be localised on. Given that in a TD-EMFT calculation we expect the excitations of interest to be localised within the active region, the species contained within the `species_tddft_kernel` block should be a subset of those in the active region. Typically, the two will be identical, i.e., the contents of the `species_tddft_kernel` block should be the same as the first line of the `species_ngwf_regions` block. For more details, see the LR-TDDFT documentation.

Run this calculation – this may take some time. ..

You should find that the lowest energy excitation is ????

If you compare this to the results in Prentice [Prentice2022], you can see that the result is very close to the experimental value of 2.09 eV.

You can plot the resulting excitation as a `.cube` file, and visualise it using e.g. VESTA.

10.4 Files for this tutorial

- `Pentacene.dat`
- `HydrogenatedPentacene.dat`
- `H2.dat`
- `Pentacene_isolated.dat`
- `Pentacene_in_p-terphenyl.dat`
- `C_NCP19_PBE_OTF.usp`
- `H_NCP19_PBE_OTF.usp`

10.4.1 References

TUTORIAL 13: ASE ONETEP INTERFACE

Author

Tom Demeyere

Date

August 2023

Preamble & Download links

Here is the tutorial for the ONETEP interface in ASE. It is a work in progress, but if you wish to give it a try, please find the two files:

- `learn_ase_in_y_minutes.py`
- `onetep_interface.py`

The first file is a general tutorial for people who are not used to ASE. The second file is a small file that explains how to use the ONETEP interface with ASE.

Below is the documentation for the ONETEP interface in ASE, as plan to be included in the ASE documentation.

11.1 Introduction

ONETEP is a linear-scaling density functional theory code which exploit the near-sightness of the electronic density. It uses a set of atom-centered local orbitals (denoted NGWFs) which are optimised in situ to enable calculations with a minimal number of orbitals.

This interface makes it possible to use ONETEP as a calculator in ASE. You need to have a copy of the ONETEP code (and an appropriate license) to use this interface.

11.2 Environment variables

The environment variable `ASE_ONETEP_COMMAND` must hold the command to invoke the ONETEP calculation. The variable must be a string with a link to the ONETEP binary, and any other specific settings required for your environment (srun, mpirun, ...)

You can setup this environment variable in your shell configuration file:

```
$ export ASE_ONETEP_COMMAND="export OMP_NUM_THREADS=4; mpirun -n 6 ~
↳ /onetep/bin/onetep.arch"
```

Or within python itself:

```
>>> environ["ASE_ONETEP_COMMAND"]="export OMP_NUM_THREADS=4; mpirun.
↳ -n 6 ~/onetep/bin/onetep.arch"
```

ASE will automatically redirect stdout and stderr to the appropriate files, namely “\$LABEL.out” and “\$LABEL.err” where label is the name used for your ONETEP calculations

11.3 Pseudopotentials

ONETEP accepts PAW datasets in the abinit format, and NCP pseudopotentials with formats USP and recpot. Support has recently been added for the upf format, for both PAW and NCPP potentials. Pseudopotentials are passed directly to the Onetep calculator as a dictionary definition. If no pseudopotentials are passed ASE will try to guess the files based on the element used and the `pseudo_path` variable.

```
# Explicitly providing each path
calc = Onetep(pseudopotentials = {'H': '/path/to/pseudos/H.usp', 'O
↳ ': '/path/to/pseudos/O.usp'})
# Using pseudo_path
calc = Onetep(pseudo_path = '/path/to/pseudos', pseudopotentials = {
↳ 'H': 'H.usp', 'O': 'O.usp'})
# ASE will try to guess them
calc = Onetep(pseudo_path = '/path/to/pseudos')
```

For ASE to correctly guess the pseudopotentials, it is best to use a `pseudo_path` that contains only one pseudopotential file for each element.

11.4 ONETEP Calculator

Simple calculations can be setup calling the Onetep calculator without any parameters, in this case ONETEP's default parameters will be used. For more complex cases using the keywords parameters is necessary. The 'keywords' parameters is a dictionary, in which each of the keys is a string that should be a ONETEP keyword, and the corresponding value is what you want to set that keyword to in the input.

11.5 Examples

Here is an example python script which sets up a calculation on a water molecule:

```
from ase.build import molecule
from ase.calculators.onetep import Onetep
from os import environ

# water molecule from ASE database, centered in a ~ 24 Å box
wat = molecule('H2O')
wat.center(12)
environ["ASE_ONETEP_COMMAND"]="export OMP_NUM_THREADS=8; mpirun -n
↳2 ~/onetep/bin/onetep.arch"
# Output will be in "water.out"
calc = Onetep(label = 'water', xc = 'PBE', paw = True, pseudo_path_
↳= '/path/to/pseudos')
wat.calc = calc
wat.get_potential_energy()
```

Here is a more complex example, this time setting up a Pt13 cluster and running a geometry optimisation on 64 cores:

```
from os import environ

import numpy as np

from ase.build import molecule
from ase.calculators.onetep import Onetep
from ase.cluster import Octahedron
from ase.optimize.sciopt import SciPyFminBFGS
# Pt13 from ase.cluster
nano = Octahedron('Pt', 3, 1)
nano.set_cell(np.eye(3)*12)
nano.center()

label = 'pt13'

environ["ASE_ONETEP_COMMAND"]="export OMP_NUM_THREADS=8; mpirun -n
↳8 ~/onetep/bin/onetep.arch"
```

(continues on next page)

(continued from previous page)

```

# ONETEP default are atomic units, one can specify 'cutoff_energy' ↵
↵: '600 eV' if needed.
keywords = {
    'xc' : 'rpbe',
    'do_properties' : True,
    'cutoff_energy' : 35,
    'output_detail': 'verbose',
    'elec_energy_tol': 1.0e-5/len(atoms),
}

# Output will be in "pt13.out",
# append = True will not overwrite file at each step
calc = Onetep(
    label = label,
    edft = True,
    append = True,
    pseudo_path = '/path/to/pseudos',
    keywords = keywords)

nanoparticle.calc = calc

opt = SciPyFminBFGS(atoms = nano, trajectory = label + ".traj", ↵
↵logfile = label + ".log")
opt.run(fmax=0.01)

```

Here is an example of setting up an EELS and LDOS calculation on an N-substituted graphene sheet, demonstrating several more advanced functionalities (eg tags, species groups, and overrides to pseudopotentials and atomic solver strings):

```

import numpy as np

from ase.build import graphene_nanoribbon
from ase.calculators.onetep import Onetep
from ase.io import write
from numpy.linalg import norm
from numpy.random import choice

sheet = graphene_nanoribbon(10, 10, type='zigzag', vacuum = 10)

# Get all distances to center of mass
com = sheet.get_center_of_mass()
distances_to_com = norm(sheet.positions - com, axis = 1)

# Find atoms close to com and change one randomly to N
p, = np.where(distances_to_com < 5)
to_nitro = choice(p)
sheet[to_nitro].symbol = 'N'

```

(continues on next page)

(continued from previous page)

```

shell_rad = np.array([1.5, 2.5, 3.0, 4.0, 4.5])

tags = np.zeros(len(sheet), dtype=np.int32)

# We want to tag atoms that are close to the introduced nitrogen
for idx, rad in enumerate(reversed(shell_rad)):
    # All distances N-C
    dist = norm(sheet[to_nitro].position - sheet.get_positions(),
    ↪axis = 1)
    # Which ones are closest to rad?
    p, = np.where(dist < rad)
    # Cannot be the nitrogen itself
    p = p[p != to_nitro]
    # Tags them
    tags[p] = len(shell_rad) - idx

sheet.set_tags(tags)

tags = ['' if i == 0 else i for i in tags]

species = np.unique(np.char.add(sheet.get_chemical_symbols(), tags))

keywords = {
    'species_core_wf' : ['N /path/to/pseudo/corehole.abinit'],
    'species_solver' : ['N SOLVE conf=1s1 2p4'],
    'pseudo_path': '/Users/tomdm/PseudoPotentials/SSSP_1.2.1',
    'xc' : 'PBE',
    'paw': True,
    'do_properties': True,
    'cutoff_energy' : '500 eV',
    'species_ldos_groups': species,
    'task' : 'GeometryOptimization'
}

calc = Onetep(
    label = 'N_doped_graphene_001',
    keywords = keywords
)

# Checking the input before running the calculation
write('to_check.dat', sheet, format='onetep-in', keywords =
    ↪keywords)

sheet.calc = calc
# Will actually run the geometry optimisation
# using ONETEP internal BFGS
sheet.get_potential_energy()

```

Quickly restart with solvation effect using the soft sphere model

```
from ase.io import read
from ase.io.onetep import get_onetep_keywords

# Read from the previous run...
optimized_sheet = read("N_doped_graphene_001.out")

# Function to retrieve keywords dict from input file...
keywords = get_onetep_keywords('N_doped_graphene_001.dat')

# We add solvation keywords
keywords.update(
    {
        'is_implicit_solvent': True,
        'is_include_apolar': True,
        'is_smear_ion_rep': True,
        'is_dielectric_model': 'fix_cavity',
        'is_dielectric_function' : 'soft_sphere'
    }
)

optimized_sheet.calc = Onetep(...)

...
```

PDF VERSION OF ALL TUTORIALS

Note: A pdf version of the tutorials can be downloaded [here](#)

INDICES AND TABLES

- genindex
- modindex
- search

BIBLIOGRAPHY

- [Bowler2012] D. R. Bowler, and T. Miyazaki, *O(N) methods in electronic structure calculations*, Reports on Progress in Physics, 75 (2012).
- [Dziedzic2011] J. Dziedzic, H. H. Helal, C.-K. Skylaris, A. A. Mostofi, and M. C. Payne, M. C., *Minimal parameter implicit solvent model for ab initio electronic-structure calculations*, EPL, 95 (2011).
- [Fox2014] S. J. Fox, J. Dziedzic, T. Fox, C. S. Tautermann, and C.-K. Skylaris, *Density functional theory calculations on entire proteins for free energies of binding: Application to a model polar binding site*, Proteins: Structure, Function and Bioinformatics, 82 (2014).
- [Gundelach2021] L. Gundelach, T. Fox, C. S. Tautermann, and C.-K. Skylaris, *Protein–ligand free energies of binding from full-protein DFT calculations: convergence and choice of exchange–correlation functional*, Physical Chemistry Chemical Physics, 23 (2021).
- [Manz2012] T. A. Manz, and D. S. Sholl, *Improved Atoms-in-Molecule Charge Partitioning Functional for Simultaneously Reproducing the Electrostatic Potential and Chemical States in Periodic and Nonperiodic Materials*, Journal of Chemical Theory and Computation, 8 (2012).
- [Mobley2017] D. L. Mobley, and M. K. Gilson, Michael K., *Predicting Binding Free Energies: Frontiers and Benchmarks*, Annual Review of Biophysics, 46 (2017).
- [Prentice2020_T8] J. C. A Prentice, J. Aarons, J. C. Womack, A. E. A. Allen, L. Andrinopoulos, L. Anton, R. A. Bell, A. Bhandari, G. A. Bramley, R. J. Charlton, R. J. Clements, D. J. Cole, G. Constantinescu, F. Corsetti, S. M. M. Dubois, K. K. B. Duff, J. M. Escartin, A. Greco, Q. Hill, L. P. Lee, E. Linscott, D. D. O’Regan, M. J. S. Phipps, L. E. Ratcliff, A. Ruiz Serrano, E. W. Tait, G. Teobaldi, V. Vitale, N. Yeung, T. J. Zuehlsdorff, J. Dziedzic, P. D. Haynes, N. D. M. Hine, A. A. Mostofi, M. C. Payne, and C.-K. Skylaris, *The ONETEP linear-scaling density functional theory program*, Journal of Chemical Physics, 152 (2020).
- [Wilkinson2014] K. A. Wilkinson, N. D. M. Hine, and C.-K. Skylaris, *Hybrid MPI-OpenMP parallelism in the ONETEP linear-scaling electronic structure code: Application to the delamination of cellulose nanofibrils*, Journal of Chemical Theory and Computation, 10 (2014).

- [Womack2018] J. C. Womack, L. Anton, J. Dziedzic, P. J. Hasnip, M. I. J. Probert, and C.-K. Skylaris, *DL-MG: A Parallel Multigrid Poisson and Poisson-Boltzmann Solver for Electronic Structure Calculations in Vacuum and Solution*, *Journal of Chemical Theory and Computation*, 14 (2018).
- [Genheden2010] S. Genheden, J. Kongsted, P. Soderhjelm, and U. Ryde, *Nonpolar solvation free energies of protein-ligand complexes*, *Journal of Chemical Theory and Computation*, 11 (2010).
- [Cornell2003] R.M.Cornell et al, in *The Iron Oxides*, John Wiley & Sons, Ltd, 2003, pp. 9-38.
- [Parkinson2016] G.S.Parkinson, *Surface Science Reports*, vol. 71, no. 1, pp. 272–365, 1 Mar. 1, 2016.
- [Naveas2023] Naveas M. et al, *iScience* 26, 106033, February 17, 2023.
- [Huang2016] Huang X. et al, *J.Phys.Chem C* 2016, 120, 4919-4930.
- [Si2020] Si et al, *J. Chem. Phys.* 152, 024706 (2020).
- [O-Regan2012] D.D.O'Regan, N. D. M. Hine, M. C. Payne and A. A. Mostofi, *Phys. Rev. B* 85, 085107 (2012).
- [Cococcioni2005] M.Cococcioni and S. de Gironcoli, *Phys. Rev. B* 71, 035105 (2005).
- [O-Regan2010] D.D.O'Regan, N. D. M. Hine, M. C. Payne and A. A. Mostofi, *Phys. Rev. B* 82, 081102 (2010).
- [Anisimov1991] J.Z.V.I. Anisimov and O. K. Andersen, *Phys. Rev. B* 44, 943 (1991).
- [Anisimov1997] V.I. Anisimov, F. Aryasetiawan, and A. I. Liechtenstein, *J. Phys.: Condens. Matter* 9, 767 (1997).
- [Dudarev1998] S.L. Dudarev, *Phys. Rev. B* 57, 3 (1998).
- [Bhandari2020] Bhandari, A.; Anton, L.; Dziedzic, J.; Peng, C.; Kramer, D.; Skylaris, C.-K. *Electronic Structure Calculations in Electrolyte Solutions: Methods for Neutralization of Extended Charged Interfaces*. *The Journal of Chemical Physics* 2020, 153 (12), 124101. <https://doi.org/10.1063/5.0021210>.
- [Fisicaro2017] Fisicaro, G.; Genovese, L.; Andreussi, O.; Mandal, S.; Nair, N. N.; Marzari, N.; Goedecker, S. *Soft-Sphere Continuum Solvation in Electronic-Structure Calculations*. *Journal of Chemical Theory and Computation* 2017, 13 (8), 3829–3845. <https://doi.org/10.1021/acs.jctc.7b00375>.
- [Larsen2017] Hjorth Larsen, A.; Jørgen Mortensen, J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dułak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; Hermes, E. D.; Jennings, P. C.; Bjerre Jensen, P.; Kermode, J.; Kitchin, J. R.; Leonhard Kolsbjerg, E.; Kubal, J.; Kaasbjerg, K.; Lysgaard, S.; Bergmann Maronsson, J. *The Atomic Simulation Environment—a Python Library for Working with Atoms*. *Journal of Physics: Condensed Matter* 2017, 29 (27), 273002. <https://doi.org/10.1088/1361-648x/aa680e>.
- [Bhandari2021] Bhandari, A.; Peng, C.; Dziedzic, J.; Anton, L.; Owen, J. R.; Kramer, D.; Chris-Kriton Skylaris. *Electrochemistry from First-Principles in the grand canoni-*

cal Ensemble. *Journal of Chemical Physics* 2021, 155 (2). <https://doi.org/10.1063/5.0056514>.

- [ONETEP2020] J. C. A. Prentice, J. Aarons, J. C. Womack, A. E. A. Allen, L. Andrinopoulos, L. Anton, R. A. Bell, A. Bhandari, G. A. Bramley, R. J. Charlton, R. J. Clements, D. J. Cole, G. Constantinescu, F. Corsetti, S. M. M. Dubois, K. K. B. Duff, J. M. Escartin, A. Greco, Q. Hill, L. P. Lee, E. Linscott, D. D. O'Regan, M. J. S. Phipps, L. E. Ratcliff, A. Ruiz Serrano, E. W. Tait, G. Teobaldi, V. Vitale, N. Yeung, T. J. Zuehlsdorff, J. Dziedzic, P. D. Haynes, N. D. M. Hine, A. A. Mostofi, M. C. Payne, and C.-K. Skylaris, *The ONETEP linear-scaling density functional theory program*, *J. Chem. Phys.* 152, 174111 (2020).
- [Prentice2020] J. C. A. Prentice, R. J. Charlton, A. A. Mostofi, and P. D. Haynes, *Combining Embedded Mean-Field Theory with Linear-Scaling Density-Functional Theory*, *J. Chem. Theory Comput.* 16, 354 (2020).
- [Prentice2022] J. C. A. Prentice, *Efficiently Computing Excitations of Complex Systems: Linear-Scaling Time-Dependent Embedded Mean-Field Theory in Implicit Solvent*, *J. Chem. Theory Comput.* 18, 1542 (2020).
- [Fornace2015] M. E. Fornace, J. Lee, K. Miyamoto, F. R. Manby, and T. F. Miller, *Embedded Mean-Field Theory*, *J. Chem. Theory Comput.* 11, 568 (2015).

INDEX

A

ASE_ONETEP_COMMAND, 80

E

environment variable

ASE_ONETEP_COMMAND, 80