

Finite-temperature DFT calculations using the Ensemble-DFT method in ONETEP

Álvaro Ruiz Serrano
August 2013

Extended to spin relaxation by Kevin Duff
April 2018

Contents

1	Basic principles	1
2	Free- and fixed-spin EDFT	2
3	Compilation	2
4	Commands for the inner loop	3
4.1	Basic setup	3
4.2	Tolerance thresholds	3
4.3	Advanced setup	4
5	Commands for the outer loop	4
6	Restarting an EDFT calculation	5
7	Controlling the parallel eigensolver	6
	References	7

1 Basic principles

This manual describes how to run finite-temperature calculations using ONETEP. A recent implementation uses a direct minimisation technique based on the Ensemble-DFT method [1]. The Helmholtz free energy functional is minimised in two nested loops. The inner loop performs a line-search in the space of Hamiltonian matrices (in a similar fashion as described in Ref. [2]), for a fixed set of NGWFs. Then, the outer loop optimises the NGWFs psinc expansion coefficients for a fixed density kernel. For a more detailed description and discussion of this method in ONETEP, see Ref. [3].

Using the NGWF representation, the Helmholtz free energy functional becomes:

$$A_{\mathcal{T}}[\{H_{\alpha\beta}\}, \{|\phi_{\alpha}\rangle\}] = E[\{H_{\alpha\beta}\}, \{|\phi_{\alpha}\rangle\}] - \mathcal{T}S[\{f_i\}]. \quad (1)$$

where $\{H_{\alpha\beta}\}$ is the NGWF representation of the Hamiltonian matrix, $\{|\phi_{\alpha}\rangle\}$ is the current set of NGWFs, \mathcal{T} is the electronic temperature, $E[\{H_{\alpha\beta}\}, \{|\phi_{\alpha}\rangle\}]$ is the energy functional and $S[\{f_i\}]$ is an entropy term. The occupancies of the Kohn-Sham states, $\{f_i\}$ are calculated from the energy levels, $\{\epsilon_i\}$, using the Fermi-Dirac distribution:

$$f_i(\epsilon_i) = \left(1 + \exp\left[\frac{\epsilon_i - \mu}{k_{\text{B}}\mathcal{T}}\right]\right)^{-1}. \quad (2)$$

where μ is the Fermi level. To obtain the energy eigenvalues, $\{\epsilon_i\}$, the Hamiltonian matrix is diagonalised as:

$$H_{\alpha\beta}M_i^{\beta} = S_{\alpha\beta}M_i^{\beta}\epsilon_i, \quad (3)$$

where $\{S_{\alpha\beta}\}$ are the elements of the NGWF overlap matrix, and $\{M_i^{\beta}\}$ are the expansion coefficients of the Kohn-Sham eigenstates in the NGWF basis set. At the moment, this is a cubic-scaling operation that requires dealing dense matrices, which makes it memory-demanding.

2 Free- and fixed-spin EDFT

By default in spin polarized runs, the total occupancy of each spin channel is held fixed; each spin channel has its own Fermi level determined by this constraint. Alternatively the whole system can be held at one Fermi level dictated by the conservation of the total number of electrons in the system, allowing the net spin to freely relax.

Free-spin EDFT should be appropriate for most applications unless there's a reason to hold the system fixed at a given net spin. As with any minimization with potentially many minima, the final state may depend on initial conditions. As a special case, free-spin EDFT may not be able to symmetry-break a system that wants to have any kind of spin polarization but that is initialized to have 0 net spin. The general advice for simple systems like basic ferromagnets (though this should not replace good system-specific judgment) is to slightly over-specify the expected net spin on each atom and hold the spin fixed for a few iterations before being allowed to relax. For example a cobalt cluster is expected to have a net spin per atom lower than that of an isolated atom, that decreases to bulk-like as a function of cluster size. A good initialization may be to give each atom atomic-like net spin and hold the net spin fixed for 3-5 NGWF CG iterations, then allow it to relax.

3 Compilation

By default, ONETEP is linked against the LAPACK library [4] for linear algebra. The LAPACK eigensolver DSYGVX [5], can only be executed in one CPU at a time. Therefore, EDFT calculations with LAPACK are limited to small systems (a few tens of atoms).

Calculations on large systems are possible if, instead, ONETEP is linked against SCALAPACK library [6] during compilation time. The SCALAPACK eigensolver, PDSYGVX, can be run in parallel using many CPUs simultaneously. Moreover, SCALAPACK can distribute the storage of dense matrices across many CPUs, thus allowing to increase the total memory allocated to a given calculation in a systematic manner, simply by requesting more processors. For the compilation against SCALAPACK to take effect, the flag `-DSCALAPACK` must be specified during the compilation of ONETEP.

4 Commands for the inner loop

4.1 Basic setup

- `edft: T/F`
[Boolean, default `edft: F`].
If true, it enables Ensemble-DFT calculations.
- `edft_maxit: n`
[Integer, default `edft_maxit: 10`].
Number of EDFT iterations in the ONETEP inner loop.
- `edft_smearing_width: x units.`
[Real physical, default `edft_smearing_width: 0.1 eV`].
Sets the value of the smearing width, $k_B \mathcal{T}$, of the Fermi-Dirac distribution. It takes units of energy (eV, Hartree) or temperature. For example,
`edft_smearing_width: 1500 K`
will set $\mathcal{T} = 1500$ degree Kelvin.
- `spin: x`
[Real, default `spin: 0.0`].
For EDFT runs this value does not need to be an integer. Because we are considering an ensemble of states it can have any real value between $-n_{\text{elec}}/2$ to $n_{\text{elec}}/2$. Make sure you have enough bands to cover the more populated spin channel.
- `edft_spin_fix`
[Integer, default `edft_spin_fix: -1`].
Control for whether the net spin of the system should remain fixed at `spin`, or relax during the run. Any negative number will fix the net spin. Nonnegative numbers n will hold the net spin fixed for n iterations then let it relax for the rest of the calculation.

4.2 Tolerance thresholds

- `edft_free_energy_thres: x units.`
[Real physical, default `edft_free_energy_thres: 1.0e-6 Hartree`].
Maximum difference in the Helmholtz free energy functional between two consecutive iterations.

- `edft_energy_thres`: *x units*.
[Real physical, default `edft_energy_thres`: 1.0e-6 Hartree].
Maximum difference in the energy functional between two consecutive iterations.
- `edft_entropy_thres`: *x units*.
[Real physical, default `edft_entropy_thres`: 1.0e-6 Hartree].
Maximum difference in the entropy functional between two consecutive iterations.
- `edft_rms_gradient_thres`: *x*.
[Real, default `edft_rms_gradient_thres`: 1.0e-4].
Maximum RMS gradient $\frac{dA_{\mathcal{T}}}{df_i}$.
- `edft_commutator_thres`: *x units*.
[Real physical, default `edft_commutator_thres`: 1.0e-5 Hartree].
Maximum value of the Hamiltonian-Kernel commutator.
- `edft_fermi_thres`: *x units*.
[Real physical, default `edft_fermi_thres`: 1.0e-3 Hartree].
Maximum change in the Fermi energy between two consecutive iterations.

4.3 Advanced setup

- `edft_extra_bands`: *n*.
[Integer, default `edft_extra_bands`: -1].
Number of extra energy bands. The total number of bands is equal to the number of NGWFs plus `edft_extra_bands`. When set to a negative number, no extra bands are added.
- `edft_round_evals`: *n*.
[Integer, default `edft_round_evals`: -1].
When set to a positive integer value, the occupancies that result from the Fermi-Dirac distribution are rounded to *n* significant figures. This feature can reduce some numerical errors arising from the grid-based representation of the NGWFs.
- `edft_write_occ`: T/F.
[Boolean, default `edft_write_occ`: F].
Save fractional occupancies in a file.
- `edft_max_step`: *x*.
[Real, default `edft_max_step`: 1.0].
Maximum step during the EDFT line search.

5 Commands for the outer loop

The standard ONETEP commands for NGWF optimisation apply to the EDFT calculations as well. The only flag that is different is:

- `ngwf_cg_rotate`: T/F. [Integer, default `ngwf_cg_rotate`: T]. This flag is always true in EDFT calculations. It ensures that the eigenvectors M_i^β are rotated to the new NGWF representation once these are updated.

6 Restarting an EDFT calculation

- `write_hamiltonian`: T/F.
[Boolean, default `write_hamiltonian`: F].
Save the last Hamiltonian matrix on a file.
- `read_hamiltonian`: T/F.
[Boolean, default `read_hamiltonian`: F].
Read the Hamiltonian matrix from a file, and continue the calculation from this point.
- `write_tightbox_ngwfs`: T/F.
[Boolean, default `write_tightbox_ngwfs`: T].
Save the last NGWFs on a file.
- `read_tightbox_ngwfs`: T/F.
[Boolean, default `read_tightbox_ngwfs`: F].
Read the NGWFs from a file and continue the calculation from this point.

NOTE: if a calculation is intended to be restarted at some point in the future, then run the calculation with

```
write_tightbox_ngwfs: T
write_hamiltonian: T
```

to save the Hamiltonian and the NGWFs on disk. Two new files will be created, with extensions `.ham` and `.tightbox_ngwfs`, respectively. Then, to restart the calculation, set

```
read_tightbox_ngwfs: T
read_hamiltonian: T
```

to tell ONETEP to read the files that were previously saved on disk. Remember to keep a backup of the output of the first run before restarting the calculation.

NOTE: the density kernel is not necessary to restart an EDFT calculation. However, it is necessary to calculate the electronic properties of the system, once the energy minimisation has completed. To save the density kernel on a file, set:

```
write_denskernel: T
```

to generate a `.dkn` file containing the density kernel. To read in the density kernel, set

```
read_denskernel: T
```

7 Controlling the parallel eigensolver

Currently, only the SCALAPACK PDSYGVX parallel eigensolver is available. A complete manual to this routine can be found by following the link in Ref. [7]. If ONETEP is interfaced to SCALAPACK, the following directives can be used:

- `eigensolver_orfac`: x.
[Real, default `eigensolver_orfac`: 1.0e-4].
Precision to which the eigensolver will orthogonalise degenerate Hamiltonian eigenvectors. Set to a negative number to avoid reorthogonalisation with the SCALAPACK eigensolver.
- `eigensolver_abstol`: x.
[Real, default `eigensolver_abstol`: 1.0e-9].
Precision to which the parallel eigensolver will calculate the eigenvalues. Set to a negative number to use SCALAPACK defaults.

The abovementioned directives are useful in calculations where the SCALAPACK eigensolver fails to orthonormalise the eigenvectors. In such cases, the following error will be printed in the input file:

```
(P)DSYGVX in subroutine dense_eigensolve returned info= 2.
```

Many times (although not always) this error might cause the calculation to fail. If this situation occurs, set

```
eigensolver_orfac: -1  
eigensolver_abstol: -1
```

in the input file and restart the calculation. SCALAPACK will not reorthonormalise the eigenvectors. Instead, an external Löwdin orthonormalisation process [8] will be triggered. This is usually more efficient for larger systems.

References

- [1] N. Marzari, D. Vanderbilt, and M. C. Payne. Ensemble density-functional theory for ab initio molecular dynamics of metals and finite-temperature insulators. *Phys. Rev. Lett.*, 79(7):1337–1340, 1997.
- [2] C. Freysoldt, S. Boeck, and J. Neugebauer. Direct minimization technique for metals in density functional theory. *Phys. Rev. B*, 79(24):241103, 2009.
- [3] Á. Ruiz-Serrano and C.-K. Skylaris. A variational method for density functional theory calculations on metallic systems with thousands of atoms. *J. Chem. Phys.*, 139(5):054107, 2013.
- [4] Lapack. <http://www.netlib.org/lapack/>.
- [5] Lapack DSYGVX eigensolver. <http://netlib.org/lapack/double/dsygvx.f>.
- [6] ScaLapack. <http://www.netlib.org/scalapack/>.
- [7] ScaLapack PDSYGVX eigensolver. <http://www.netlib.org/scalapack/double/pdsygvx.f>.
- [8] Per-Olov Lowdin. On the non-orthogonality problem connected with the use of atomic wave functions in the theory of molecules and crystals. *J. Chem. Phys.*, 18(3):365–375, 1950.