Theoretical Material Science: Electronic structure theory at the computer

Exercise 11

Some rules on expected documentation from this exercise.

The computational exercises are intended as "hands-on" experience with actual, numerical electronic structure theory. Our main goal is to fill some of the basic concepts with life for real systems.

This exercise is intended as a homework exercise, to be performed at the "PC pool" during its opening hours. While the background section of the script is long, the actual exercise should be much shorter than the previous one (hence, homework).

Please still hand in your solutions to the exercise individually, as you would normally do. What is expected is a record of the basic data that we ask for (e.g., in table form), rough answers to the questions asked (answers can be short, but should be there and should indicate that you understood the meaning of your data), and plots, where required.

For required plots, please use the printer in the PC pool and append them to your exercise materials. As you may know, printed pages usually cost 5 cents per page on that printer. *However*, for the purposes of the exercise, a contingent of free pages has been agreed upon. If you are doing printouts for this class, please mention this to the administrators.

Please note that the page contingent is not gigantic – so be careful and do not print excessively many pages. A few pages at most should be sufficient for the solutions, anyway.

Please hand in all your solutions as usual at the beginning of the exercise in the following week.

Contents

1 Background

This exercise is a short "follow-up" to our previous examination of free atoms, but it conveys two important physical pieces:

- The first piece is an obvious physical property that we can calculate and that serves as one of the most important "benchmark" measurable quantity of atomic physics: The ionization potential, i.e., the energy that it costs to remove an electron from the atom entirely and take it to infinity.

- The second piece is more subtle, and concerns an unexpected pitfall of the conventional effective single-particle Ansatz on which most of our intuitive understanding of electronic structure (at least, in the solid state) is based: The fact that the electronic structure from those theories need not reflect the full symmetry of the underlying *external* potential, even in the ground state.

To be precise, it turns out that the previous exercise included a assumption that helped preserve spherically symmetric ground-state densities of free atoms.

Intriguingly, this seemingly obvious assumption not always fulfilled in our "most productive" electronic structure theories, those that are based on the approximation of a single determinant (auxiliary) wave function: Density functional theory and Hartree-Fock theory.

We will see in a moment (qualitatively) why such a simple approximation would not hold strictly, and what exactly "helped" us to preserve spherical symmetry before.

Note, already here, that this "symmetry breaking" effect is often strongest for pure Hartree-Fock theory. It appears in the same way in any density-functional theory, but for the present exercise, we will only consider Hartree-Fock.[1]

1.1 The problem

Consider (again) the electronic structure of the free Si atom as it was computed in the last exercise session (exercise number 8). For the eigenvalues in Hartree-Fock theory, you should have received a result looking somewhat like this:

```
Spin-up eigenvalues:

State    Occupation    Eigenvalue [Ha]    Eigenvalue [eV]
    1     1.00000        -68.813916         -1872.52194
    2     1.00000         -6.177093          -168.08725
    3     1.00000         -4.284405          -116.58459
    4     1.00000         -4.284405          -116.58459
    5     1.00000         -4.284405          -116.58459
    6     1.00000         -0.624976           -17.00647
    7     0.66785         -0.204245            -5.55778
    8     0.66624         -0.204243            -5.55775
    9     0.66591         -0.204243            -5.55774
   10     0.00000          0.238901             6.50083

Spin-down eigenvalues:
```

---

[1] In the local-density approximation, symmetry-breaking in the ground state happens as well, but achieving a symmetry-breaking solution without enforcing it explicitly can be a much more difficult exercise. LDA, in general, has electron-gas like behaviour at its heart, rather than just the behaviour of individual electrons under the Pauli principle. It is thus not as prone to symmetry-breaking as Hartree-Fock theory. As a result, the energy differences between symmetric and symmetry-broken states are much less severe than in Hartree-Fock theory. In that case, one also needs to "convince" the self-consistency cycle to not jump from one of several symmetry-equivalent solutions to another.

State	Occupation	Eigenvalue [Ha]	Eigenvalue [eV]
1	1.00000	-68.805356	-1872.28899
2	1.00000	-6.167874	-167.83640
3	1.00000	-4.265994	-116.08361
4	1.00000	-4.265994	-116.08361
5	1.00000	-4.265994	-116.08361
6	1.00000	-0.478362	-13.01690
7	0.00000	0.038724	1.05374
8	0.00000	0.038724	1.05374
9	0.00000	0.038724	1.05374
10	0.00000	0.281893	7.67069

From atomic physics, we know that there should be four valence electrons: Two $3s$ electrons, and two $3p$ electrons. Indeed, this is so. There are two eigenvalues for $3s$ at -17 eV (spin up) and -13 eV (spin down), respectively. In addition, each spin channel contains three degenerate $3p$ levels.

The oddity, which some of you noticed already, is here. If the $3p$ levels are supposed to be degenerate, they should be filled with equal numbers of electrons. Since there are only two p electrons to be distributed, that makes $2/3$ electrons per $3p_x$, $3p_y$, and $3p_z$ level, respectively. This value is, in fact, reflected in the occupation numbers above.

On second thought, however, a $2/3$-electron is not exactly a textbook particle. To a physicist used to metallic systems, fractional occupation averages for a single-electron level may somehow seem tolerable. A chemist (loosely speaking), on the other hand, might be inclined to have you placed in jail.

How did we get there?

In a (numerically) perfect world, we start our s.c.f. cycle with a spherical free atom. From there, we solve the first, not yet self-consistent step of the problem, in which all $3p$ orbitals must still be numerically exactly the same (they came from a spherical potential). As a result, even if we "occupy" these levels with an exact step function, we can only assign $2/3$ of an electron to each one.

Alternatively, we could decide to pick two specific $3p$ orbitals and make them "special" by filling them with two whole electrons, leaving all the others empty. In that case, however, we have already broken the symmetry of our system. The self-consistent field cycle may or may not converge from this point forward, but the key point is that our atom will have lost its symmetry.

1.2 The meaning of our problem

How could we have prevented this problem from happening? From a purely quantum-mechanical point of view, the reasoning is now rather simple.

Remember that the effective single-particle levels above only arose as factors in a single many-body wave function. For Hartree-Fock theory, this many-body wave function is defined as a *single* Slater determinant of interacting electrons. In Kohn-Sham density-functional theory, it is a *single* Slater determinant of non-interacting auxiliary orbitals (the *physical* results of Kohn-Sham theory are the density and the total energy).

In both cases, to restore a fully symmetric quantum mechanical wave function, we would thus have to symmetrize the many-electron wave function itself, by an appropriate superposition of *several* Slater determinants. More than one Slater determinant, however, are explicitly not foreseen in a single-determinant theory.

On the other hand, it is not entirely clear that symmetry-breaking electronic solutions should be forbidden in practice. The reality is that symmetry-breaking states of matter exist and are important in physics. Antiferromagnetic solids are an important conceptual example. From a puristic point of view, the correct solution might be a superposition of two or more equivalent

antisymmetric states, but in fact, a solid is usually in one state or the other, not in both (the time it takes to tunnel between one macroscopic antiferromagnetic state and its symmetry equivalent would be prohibitive).

Finally, since we know fractional occupation numbers in metallic systems, it is important to remember where they came from (see the "Jellium" part of the lecture). Here, fractional level occupation arose in the form of the Fermi function, as a consequence of *thermodynamics*, not of quantum mechanics.

In thermodynamics, however, the Fermi function arose as a *probability* to find a many-body system in a specific state, based on the idea that there exists an *ensemble* of equivalent systems that we are probing.

Even for a free atom, a formal argument can thus be made that fractional occupation numbers reflect not the properties of a single quantum-mechanical many-electron state, but rather the properties of *an ensemble average of individual systems in different quantum-mechanical states*. Such ensemble averages can be defined consistently through their *density matrices*. Whether such an interpretation (as thermodynamic averages) makes sense for a single atom calculation and for a given theory (e.g., Hartree-Fock) is quite another matter, but we will not pursue this question here.

The qualitative considerations above are summarized as background only and will not be pursued further in this script, nor will it be needed for the exercise below. If you want to know more, however, the aspects mentioned above—ensemble averages vs. pure quantum mechanical states— have a long history in the field.

1.3   Now what happened in practice?

The fact that we *could* get formal fractional occupation numbers for degenerate single-electron levels is now clear. If degenerate levels at the Fermi level do arise in an actual calculation, however, one must usually still enforce them against some numerical difficulties.

The standard way to do this is to impose a "broadening function," which assigns fractional occupation numbers to any orbital close to the Fermi level. Once could, for instance, use the Fermi function. From a formal point of view, this choice would reflect thermodynamics correctly. For numerical practice, though, other choices exist and may be more convenient. One popular option is a Gaussian-derived broadening around the Fermi level (with the shape of an error function). To ensure stability of the self-consistency cycle in FHI-aims, such a broadening is enforced by default (with a width of 0.01 eV), and this is what stabilized the spherically symmetric level occupation of the last exercise.

If one simply reduces this broadening to a very small value (10^{-6} eV), residual numerical inaccuracies are sometimes enough to push a system into a stable symmetry-broken solution. We will see the consequences below.

2   Reminder: FHI-aims at the "PC pool"

As a quick reminder, here is (again) an overview of the most important pieces needed for FHI-aims at the "PC pool".

- To run the code, create the necessary input files (`control.in` and `geometry.in`) in a working directory of your choice.

- The calling sequence for FHI-aims is:

```
mpirun -np 1 aims.VERSION.mpi.x | tee calculation.out
```

This should be enough to produce the output needed, based on the input files already generated during exercise 8.

As before, although you should not need it, there is a complete manual (pdf) for the FHI-aims code located in
/media/public/TFKP_2012/FHI-aims.pdf .

3 Symmetry breaking: Exercise 11a)

- In exercise 8, we addressed free atoms with an implicit assumption that their ground state is spherically symmetric. In fact, in Hartree-Fock theory as well as DFT, this is usually not the case. Follow the detailed instructions in the exercise to break the symmetry of the Si atom in Hartree-Fock theory. What does the density look like?

To begin, repeat a single calculation for the Si atom in Hartree-Fock theory, as in the previous exercise 8. You should recover the symmetric, fractionally occupied result quoted in the present script, above.

Next, modify the `control.in` file by adding the following line:

```
occupation_type gaussian 0.000001
```

This choice enforces a much narrower broadening about the Fermi level, so that small numerical differences between formally equivalent eigenlevels may drive the system into a symmetry-broken self-consistent solution.

What is the result? What are the eigenvalues, occupation numbers, and what is the total energy? How different are they from the symmetric solution? Which state is the ground state in Hartree-Fock theory—the present (non-symmetric) one, or the symmetric state from the last time?

To plot the three-dimensional density, remember the following construct:

```
output cube total_density
cube filename total_density_uncharged.cube
```

in your `control.in` file. Does the density look as you expect?

4 Ionization potential: Exercise 11 b)

- A classic measurable property of atoms is the ionization potential, tabulated, e.g., at
http://en.wikipedia.org/wiki/Ionization_energy .
For the Si atom, use the total energy difference between the neutral and +1 charged atoms to compute the ionization potential in Hartree-Fock theory. Compare your results to the experimental value, and to the highest occupied single-particle eigenvalues of the neutral Si atom. What do you find? How good is Koopman's theorem? If you had used the symmetric atoms of the previous exercise, how good would the result have been?

Look up the experimental ionization energy as stated above.

Then, compute the ionization energy from the definition:

$$IP = E(N) - E(N-1) \quad , \tag{1}$$

where N is the number of electrons.

We already have the total energy for the symmetry-broken ground state of the Si atom. To compute $E(N-1)$, repeat the calculation, but set the charge to "+1" in `control.in`. Make sure that the ground state does not show any fractional occupation numbers. How good is the result for the IP?

If you had, instead, applied Koopman's theorem to the *neutral* Hartree-Fock result, how good would the result have been?

Finally, compare the same set of results, but for the spherically symmetric solutions for the Si atom and ion, i.e., with fractional occupation numbers (default broadening). How good would the results have been?