Serial: Exciting quanta

In today's final installment, we will take a closer look at the excited states of molecules and the transitions between them.

In the previous calculations, we considered molecules only in the ground state. In such a case, we can expect to have a simple state for most molecules, where all the lowest orbitals up to some level are either occupied by two electrons, or are empty, which simplifies the calculations. Conversely, once we excite an electron to a higher orbital, this assumption no longer holds, which brings complications.

Overall, the world of excited states is much richer than if we restrict ourselves to just the ground state. As an illustration, consider a simple hydrogen molecule with only two electrons. Just as in the last episode, we can calculate the energy of the several lowest electron states for each internuclear distance. Later, we will show how we can calculate these excited states. If we plot these energies on a graph, as we did in the last episode for the dissociation curve, we get the figure 1.
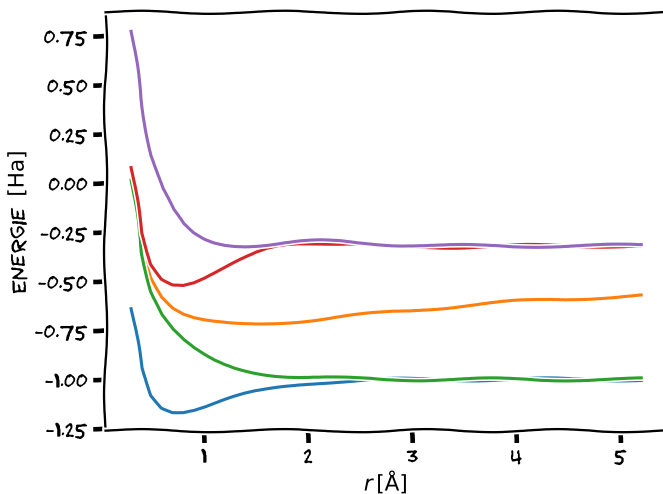

Fig. 1: Electron state energies

In it, each curve corresponds to a particular state. We see that some curves do not have the character of a bound state because their energy increases monotonically with approaching nuclei and have the lowest energy for infinitely distant nuclei. In contrast, the dissociation curves for other states have almost the same shape as the ground state curve, only with higher energy. Thus, a molecule can remain in such a state for a longer time period. At the same time, this bound excited state will have its own set of vibrational states.

Moreover, in the last episode, we discovered that for a molecule with N atoms, we must consider potential curves (surfaces) in $3N - 6$ $(3N - 5)$ dimensional space. As you can already guess, this world is extremely rich but just as complicated, and there is still intense research in this field, both at the experimental and theoretical levels.

Although the world of molecules and their excited states is complicated, we can find several universally occurring phenomena. We will demonstrate some of them in this episode.

Fluorescence

Fluorescence is one of the most well-known phenomena associated with the transition between electron states. A simplified description is that a molecule absorbs a photon of some frequency, thereby entering an excited state. It then remains there until a spontaneous emission, i.e., the photon is emitted, and the transition to the ground state occurs again.

Let us make a simple model of this process. As in the last episode, we approximate the potential curves of the ground state and the lowest excited state around the minimum by a parabola. In addition, each state will have its own set of vibrational states, as seen in figure 2.
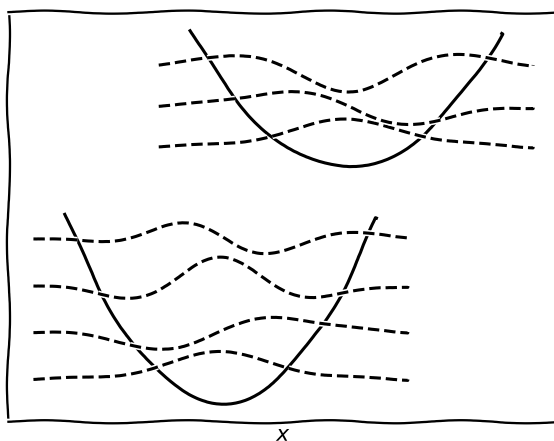

Fig. 2: Potential curve approximation

Note that the equilibrium geometry of the excited state differs from the ground state geometry. When a photon is absorbed, excitation to a higher electron state occurs. Since this is primarily a rearrangement of the electron shell, there is no change in geometry. However, as you know from the previous episode, nuclei also have a quantum character. Even though we talk about the molecule's geometry, we have to consider the wave function given by the vibrational state and the corresponding probability density. Also, we cannot directly talk about whether the ground state and the excited state have the same geometry, but only what is the probability that the ground state and the excited state will have the same geometry. Thus, the

probability that excitation occurs depends on how much the probability densities of the nuclei for the ground state and the excited state overlap. Therefore, if the equilibrium "geometry" of the excited state differs, excitation will most often end up in some higher vibrational level of the excited state, because higher vibrational levels also have a higher probability of occurring further from the equilibrium position.

After excitation, a vibrational relaxation follows, i.e., the molecule gradually moves to the lowest vibrational level of the state. That happens primarily by interacting with surrounding molecules, which receive energy in the form of heat. Vibrational relaxation is extremely fast, much faster than if the molecule was to jump to lower vibrational levels by emitting infrared photons. From there, spontaneous emission occurs, and the molecule transitions back to the electron ground state. For the same reasons as before, we cross over to a higher vibrational level again. (The molecule now has the geometry of an excited state.) Then the vibrational relaxation occurs again, and we end up in the state we started in. We can see the whole thing in the diagram 3.
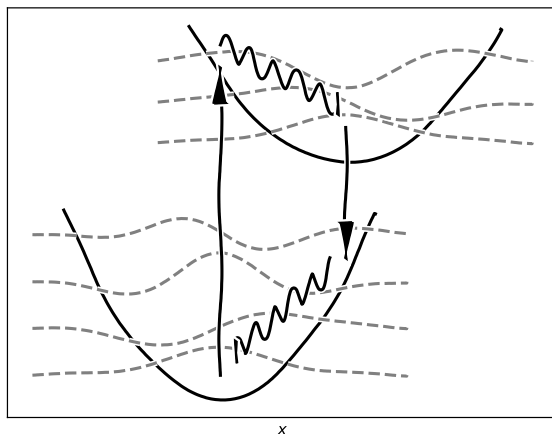

Fig. 3: Transition between states

The straight arrow indicates transitions associated with the absorption or emission of a photon, while the wavy arrow indicates nonradiative transitions. Be aware of the fact that we are plotting the potential curves of the electron states. These potentials serve only as guides, and the molecule jumps around the various vibrational states.

From the diagram, we can also see that the energy of the photons emitted by fluorescence will be lower than that of the excitation radiation. The difference between the maximum in the absorption spectrum and the maximum of fluorescence is called the Stokes shift. Moreover, the absorption and fluorescence spectra are of similar shapes (mirror images). That occurs because the magnitude of the overlap between the lowest vibrational level of the ground state with the n-th vibrational level of the excited state is similar to the extent of the overlap between the

lowest vibrational level of the excited state with the n-th vibrational level of the ground state.

Although it may seem that fluorescence takes place instantaneously (and the lifetime of excited states is indeed on the order of nanoseconds), it is a slow process on the microworld scale. For example, the mentioned vibrational relaxation is at least three orders of magnitude faster in the order of picoseconds. Because of this, we can assume that a molecule "does not manage to" emit a photon before it relaxes, and fluorescence almost always takes place from the lowest vibrational level.

In addition to fluorescence, a molecule may transition to the ground state without emitting a photon. This is called a nonradiative transition. Like vibrational relaxation, the molecule transfers energy to surrounding molecules during this transition. Consequently, this means that the molecule fluoresces fewer photons than it absorbs. The percentage of photons the molecule emits back is called the quantum yield. Depending on the molecule, the value of the quantum yield can be close to $100\,\%$ (for fluorescent dyes), and $0\,\%$ – such a molecule does not fluoresce at all. (For example, beta-carotene.)

Naturally, molecules have multiple excited states. But unlike the ground state, the differences between the energies of the different excited states are smaller. Due to this phenomenon, the vibrational spectra of the two states overlap, leading to the fact that even if the molecule is at the lowest vibrational level of some higher excited state, instead of emitting a photon, it can very quickly relax to the higher vibrational level of another state. As a result, fluorescence almost always takes place only from the lowest excited state. That is the essence of Kasha's rule, which states that the wavelength of fluorescence radiation does not depend on the wavelength at which the molecule is excited.

Wave functions of excited states

As we mentioned at the beginning, the wave functions of excited states tend to be more complicated than those of the ground state. While all the levels in the ground state are occupied twice, excited states will have some of the electrons excited to a higher unoccupied orbital, as indicated in figure 4.
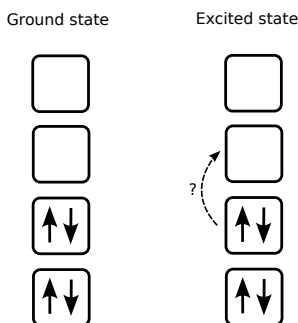

Fig. 4: Ground and excited state orbitals

We can expect to get the lowest energy state if we excite an electron from the highest occupied to the lowest unoccupied orbital. In total, we have four possibilities of how the electron spins in these two orbitals may be oriented with respect to each other.
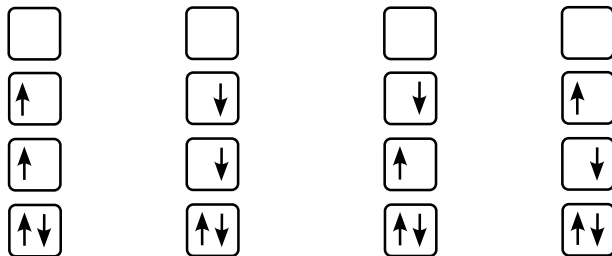
Fig. 5: Excitation options

Before we can interpret these states, we need to talk about spin. Spin is a quantum quantity interpreted as a particle's internal angular momentum, which appears as an internal magnetic moment for charged particles. Two spin numbers can characterize particles. Firstly, the total amount of spin is characteristic for a particular type of particle (for example, an electron has a spin of $1/2$). Secondly, it is the projection of the spin in a particular direction. Depending on the magnitude of the spin, the latter has only some specific values that it can take on. For example, an electron with spin $1/2$ has two possible projection values, $+1/2$ and $-1/2$. Up and down arrows indicate these two possible values in the well-known orbital diagrams.

In the same way, we can use spin to classify the states of molecules. If we have all the electrons paired, which is typical for the ground state of many molecules, we have an overall spin value of 0. We call such a state a singlet because it has a single allowed spin projection, and that is 0. Conversely, we may have two unpaired electrons whose spins are "summed", in which case a triplet state emerges. It has three possible values of projection, 1, 0, and -1.

Returning to our four possible orientations of electrons in the excited state, we see that the first two combinations from the left will be two of the potential projections of the triplet, with projections of 1 and -1. In the first case, we add the projections $1/2$ from the two electrons, while in the second case, we add the two projections $-1/2$. The other two combinations are more complicated. Both have a spin projection of zero because the contributions of the two unpaired electrons cancel out. However, if we were to address the issue in more detail, we would discover that these two combinations do not have a specific spin value; they are not pure spin states. This is similar to taking a general wave function. It is unlikely to be an intrinsic function of the Hamiltonian, and therefore we cannot assign a specific energy value to it.

Instead, we must take such combinations of spins. The result is one wave function which is a singlet, and the other, which is the third missing projection of the triplet.

Due to the interactions between the electrons, the excited triplet has a lower energy than the excited singlet. The same principles apply here as in the case of Hund's rules, which also favor electrons pointing in the same direction.

The most important implication is that since spin is conserved in nature (as is the angular momentum), radiative transitions between states of different spin are forbidden and therefore do not proceed as rapidly. That is the essence of phosphorescence, which we shall show presently.

For the sake of argument, let us add that if we start from the Schrödinger equation alone, the radiative transition should not occur at all. However, other, weaker effects may be present. In this case, relativistic effects will begin to take effect, allowing such a transition to take place, albeit more slowly.

Phosphorescence

As we have just noted, the transition between the ground state, which is usually in the singlet state, and some triplet excited state is forbidden. Therefore, we cannot use light to excite directly to the triplet state. Nevertheless, a molecule can get to this state. We excite the molecule to some singlet state, from where it can then, most often because of interaction with the angular momentum of the orbiting electrons, slip to a lower triplet state. Such a process is called an intersystem crossing.

Once the molecule has reached the triplet state, as in the case of fluorescence, it slides to the lowest of the triplet states. Once there, however, the molecule is trapped. It cannot spontaneously go to any of the other excited states because they have higher energy, which is not consistent with the law of conservation of energy. Therefore, it must wait until it finally "manages" to emit a photon and transition back to the ground state. However, because this transition is forbidden, the molecule stays in this excited state for a long time, approximately on the order of seconds, minutes, or even longer. This slow emission is then called phosphorescence.

This effect is the basis for objects glowing in the dark – while they are exposed to light, they gradually accumulate molecules that have been excited to the triplet level. If we then turn off the light, these molecules gradually radiate by phosphorescence. In this case, however, inorganic materials such as ZnS are used with various additives. In organic molecules, the phosphorescence is weaker because they have lighter atoms, and relativity is not as apparent.

The so-called Jablonski diagram usually represents all these effects.

In it, the thick lines indicate the individual electron states and the thin lines indicate the corresponding vibrational levels for each state. In one column, singlet states are plotted on top of each other, while triplet states are to the right of them. Transitions between the different states are then indicated by arrows. A straight arrow marks a transition associated with a photon's release or emission. Conversely, a wavy arrow denotes transitions where no photon is emitted, and energy is released in another way, e.g., as heat.

How is it calculated though?

As you can guess, the Hartree-Fock method is insufficient for excited states. We have to do away with the fact that the wave function has only one particular orbital occupation. To sort it out, we take a wave function that contains several configurations with different occupied orbitals, which are summed with different weighting coefficients. Thus, in the whole problem, we optimize both the shape of the orbitals and the unknown coefficients for individual orbitals. Such a method is called a Multi-configurational self-consistent field (MCSCF).
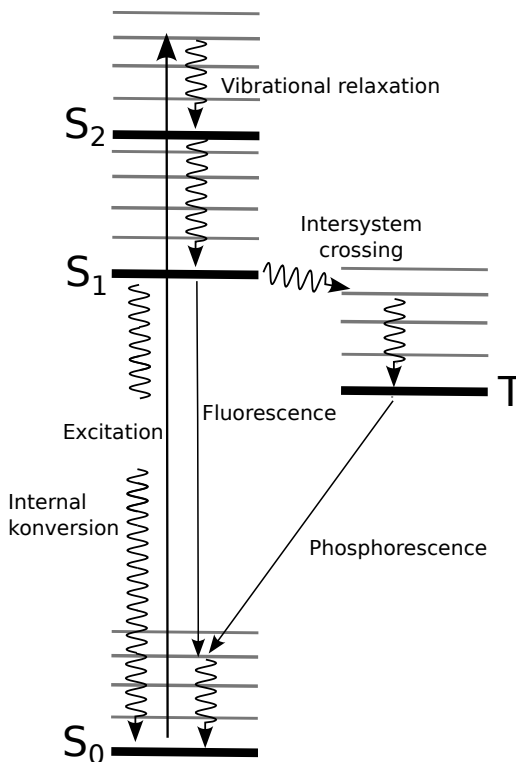
Fig. 6: Jablonski diagram

Since deciding which configurations must be taken to describe a given problem would be nontrivial, we simplify our work by selecting all orbitals relevant for describing the problem instead of specific configurations. Moreover, we consider all possible configurations that we can obtain by placing a particular number of electrons in the orbitals. A given set of orbitals is called a complete active space. The method is then called CASSCF.

Using the CASSCF method requires performing the steps we have carried out in the previous episodes first. Firstly, we optimize the geometry, then calculate the Hartree-Fock energy to obtain the initial set of orbitals. Then, we select the active space from these orbitals, and we can start the CASSCF calculation itself.

What does the calculation look like in the Psi4 program? Let's demonstrate it on a water molecule. We will calculate three states – the ground state, the first excited singlet, and the lowest triplet state. Since the states with different spins are largely independent, we can calculate them separately. We will thus have two calculations, one for the singlet, where we count two states, and the other for the triplet, where we only need one state.

As we expect the wave function of these lowest excited states to look exactly as we have just shown, we need to take two orbitals into active space. One is occupied twice in the ground state, and the other is the lowest unoccupied state. Thus, at the beginning of the input, there

will already be a known part with a base and geometry. We will input the geometry we have previously optimized to save ourselves from the complications of geometric optimization.

```
set basis cc-pvdz
molecule H20{
symmetry c1
0 1
 O~0.005700842385 -0.064489709337 0.000000000000
 H -0.791211319490 0.445805409758 0.000000000000
 H 0.700734813858 0.577693076293 0.000000000000
}
```

If you are observant, you will have noticed one extra line in the geometry, `symmetry c1`. We add this because the water molecule is symmetric, i.e., it has a double rotational axis and two planes of reflection. The Psi4 program then divides the individual orbitals according to how they behave with respect to these operations, i.e., whether they stay the same or change signs, for example. Although, in many cases, it is convenient to divide them in this way, it would unnecessarily complicate our calculation. Therefore, with this line, we have set the program to treat the molecule as if it had no symmetry.

Next, we need to set which orbitals will be in active space. We do this by manually setting the number of orbitals occupied twice in all configurations and the number of orbitals in active space. Water has a total of 10 electrons (8 from the oxygen and one from each of the hydrogens), so in the ground state, there are five orbitals occupied twice, of which we want one in active space. Therefore, the number of orbitals occupied in all configurations is 4. On the other hand, we want two active orbitals. We set this using the following lines:

```
set restricted_docc [4]
set active [2]
```

We still need to set other things, though. First, we want the program to check if we have the correct multiplicity. We can do this with

```
set calc_s_squared true
set S 0
```

where the second line sets the multiplicity value directly. Note that the total spin value is given directly, i.e., for singlet 0, triplet 1, etc.

Ultimately, we have to say that we want two states. But since each state has its "optimal" orbitals, we have to tell the program not to try to find optimal orbitals for either state but rather to create an "averaged" orbital suitable for both states equally.

```
set num_roots 2
set avg_states [0,1]
```

The second line sets from which states (numbered from 0) the program calculates the average. You can see that if you leave only [0] in the square brackets, the program will decrease the energy of the lowest ground state, and on the other hand, the energy of the excited state will increase significantly. Conversely, [1] will reduce the energy of the excited state at the expense of the ground state.

Now you need to calculate the energy using `energy('casscf')`. The whole input will look like this:

```
set basis cc-pvdz
molecule H20{
symmetry c1
0 1
 O 0.005700842385 -0.064489709337 0.000000000000
 H -0.791211319490 0.445805409758 0.000000000000
 H 0.700734813858 0.577693076293 0.000000000000
}

set restricted_docc [4]
set active [2]

set calc_s_squared true
set S 0

set num_roots 2
set avg_states [0,1]

energy('casscf')
```

In the output, we find a piece of text before the end:

```
 ==> Energetics <==

 SCF energy = -76.027032779942914
 Total MCSCF energy = -75.868406011147982

 ==> MCSCF root 0 information <==

 MCSCF Root 0 energy = -76.012806678099437

 Active Space Natural occupation numbers:

 A 1.999576 A 0.000424

 The 4 most important determinants:

 * 1 0.999894 ( 0, 0) 5AX
 * 2 -0.014555 ( 1, 1) 6AX
 * 3 0.000000 ( 0, 1) 5AA 6AB
 * 4 0.000000 ( 1, 0) 5AB 6AA


 ==> MCSCF root 1 information <==

 MCSCF Root 1 energy = -75.724005344196783
```

```
Active Space Natural occupation numbers:

A 1.000000 A 1.000000

The 4 most important determinants:

* 1 0.707107 ( 0, 1) 5AA 6AB
* 2 0.707107 ( 1, 0) 5AB 6AA
* 3 0.000000 ( 1, 1) 6AX
* 4 0.000000 ( 0, 0) 5AX
```

Beware, in this case, we are not interested in the `Total MCSCF energy` line because this energy is also "averaged", but the energy of the individual states (here called `root`), which are below. Thus, we see that the energy of the lowest state is $-76.012\,806\,678\,099\,437$ Ha, while the first excited state has an energy of $-75.724\,005\,344\,196\,783$ Ha. Now we modify the input for the triplet – change the line `set S 0` to `set S 1` and delete the lines where we wanted two states.

We get input like this,

```
set basis cc-pvdz
molecule H20{
symmetry c1
0 1
 O 0.005700842385 -0.064489709337 0.000000000000
 H -0.791211319490 0.445805409758 0.000000000000
 H 0.700734813858 0.577693076293 0.000000000000
}

set restricted_docc [4]
set active [2]

set calc_s_squared true
set S 1

energy('casscf')
```

and after execution, the energy of the triplet state $-75.770\,492\,145\,528\,493$ Ha.

Subtracting the energies of the two singlet states gives an excitation energy of 0.28 Ha, which is 7.62 eV. That is certainly not a bad result since the experimental value is about 7.4 eV.

In conclusion, it is worth mentioning that the CASSCF method is one of the strongest hammers in the quantum chemist's arsenal, and is used for calculations even with the most entangled molecules. Thus, it is suitable, for example, for calculations of transition metal complexes, excited states, or transition structures in chemical reactions.

A few words to conclude

This concludes our series on quantum chemistry. I hope you enjoyed it and that you learned something new. And perhaps it is already clear to you that the statement "chemistry is just applied physics" is true, but that in practice, the application of physics to molecules is far from trivial.