Second Part

Slide 1

Welcome back.
In the first part of this lecture, we have seen how, in a QM/MM scheme, a molecular system is usually partitioned, that is a small region treated at a quantum level (computational expensive), and the rest of the system with a classical resolution (computationally less demanding).

Slide 2

Obviously, special attention has to be paid when coupling the quantum and the classical regions. In particular, in the Hamiltonian of this fictitious QM/MM system (fictitious because the real system is not separated in two parts with two different physics) the potential energy contains three types of interactions:

- interactions between particles in the quantum region, that is between electrons and nuclei;
- interactions between atoms in the classical region;
- and the interactions between quantum particles and classical atoms

According to how one takes into account these last interactions, the different QM/MM schemes can be grouped in subtractive and additive schemes.

Slide 3

In the subtractive scheme, the potential energy of the QM/MM system is obtained in three steps by performing three independent calculations:

- First, the energy of the total system is evaluated at the MM level, for example at force field level.
- Then, the energy of the isolated quantum subsystem is calculated at quantum level and added to the previous result.
- Finally, the MM energy of the QM subsystem is subtracted. This to avoid to count twice the interactions within the QM subsystem.

The main advantage of this QM/MM coupling scheme is that in the code, no communication is required between the quantum mechanical and the classical routines. This makes the implementation relatively straightforward.

However, there are also several drawbacks:

- A major disadvantage is that in the first and in the third step, a force field is required for the quantum subsystem, which may not always be available.

- In addition, this force field needs to be sufficiently flexible to describe for example the effect of chemical changes when a reaction occurs.
- A further drawback of this class of methods is the total absence of polarization of the electron density due to the classical environment. This shortcoming can be particularly problematic when modelling for example biological charge transfer processes, since these are usually mediated by the protein environment.

Therefore, for a realistic description of such cases a more consistent treatment of the interactions between the electrons and their surrounding environment is needed. This can be obtained through an additive scheme.

Slide 4

In the additive schemes, there are no calculations at different resolutions taking place separately, but the potential energy for the whole QM/MM system is a sum of three contributions:

- quantum energy terms;
- classical energy terms;
- and QM/MM coupling terms.

In contrast to the subtractive schemes, the interactions between the particles in the quantum region and the classical atoms in the MM region are treated explicitly in the last term of this expression.

This additive approach is the coupling scheme implemented in the QM/MM interface of the CP2K code that will be described in the tutorial.

- The quantum energy term usually comes from the DFT Kohn-Sham Hamiltonian.
- The classical energy term comes almost always from a classical force field and its choice can be limited by the availability of force fields implemented in the code.
- While the term describing the interactions between QM and MM regions is usually decomposed in two parts: the bonded and the non-bonded parts.

The bonded part is the part that describes the covalent interactions between “quantum” atoms, that is the atoms in the quantum region, and classical atoms in the MM region.

Slide 5

Therefore, this bonded part is present only when the boundary between the QM and MM regions cuts a covalent bond connecting a quantum atom to a classical atom, as shown in the picture if you look at the yellow fragments.

Moreover, in this case, care has to be taken when evaluating the quantum wave function, that is when you solve the Schrödinger equation for the quantum part. In fact, a straightforward cut through the QM/MM boundary would create one or more unpaired electrons in the quantum subsystem. In reality, these electrons are paired in bonding orbitals with electrons belonging to the atom on the MM side. However, now those electrons do not exist in the MM region, due to our artificial partitioning and the lower level of resolution we decided to use in this region.

In literature a number of approaches have been proposed to remedy the artefact that origins from such open valences, many of which are available in CP2K.

- For example, you can saturate the dangling valences with a monovalent capping atom, usually a hydrogen atom, at an appropriate position along the bond vector.

- Alternatively, it is possible to use the link atom pseudopotentials approach, which consists in introducing in the QM system a description of the classical atom at the border bonded to the quantum atom, through a special pseudopotential with the required valence charge. It can be shown that this method requires constraining the bond distance appropriately.

- Another approach is to place a set of hybrid orbitals on the boundary atom between the QM and MM fragments. One of these is included in the QM region for the self-consistent optimization to find the wave function of the QM region, and the others are treated as auxiliary orbitals that do not participate in the QM optimization, but they provide an effective electric field that contributes to the external potential felt by the dynamical electrons.

Slide 6

Let's come back to the QM/MM energy coupling term. Beyond the bonded part, there is always a non-bonded part that describes the interactions between the QM and MM regions not connected through a covalent bond (or a sequence up to 3 covalent bonds). This part is usually formed by a "steric" term, that is a term that takes into account the van der Waals interactions, and a pure electrostatic term.

Slide 7

In most of the QM/MM approaches, including in CP2K, the interactions in E_QMMM_bonded and in E_QMMM_steric are handled at the force field level.

This means that for example in the bonded term you can find chemical bonds between QM and MM atoms modelled by harmonic potentials, as well as the angles defined by one QM atom and two MM atoms, but also torsions involving at most two QM atoms, which are commonly modelled by a periodic potential function.

Similarly, the steric term is usually described by a Lennard-Jones potential and the parameters come from the force field used in the MM part.

The rest of the non-bonded interactions, that is the interactions between quantum and classical atoms separated by three or more atoms in the topology, are included in the electrostatic term. This is really an interaction between classical partial charges in the MM region and the quantum charge density.

But how is this term calculated in practice? Really, several approaches have been proposed for this term.

Slide 8

In the so-called mechanical embedding, the simplest electrostatic coupling scheme, the electronic wave function is computed for an isolated QM subsystem. In other words, the classical environment cannot induce polarization on the electron density in the quantum region.
Within this approach, different ways to get E_QMMM_electrostatic have been proposed.

- The simplest and very rough approximation consists to completely neglected it. So, in this case E_QMMM_electrostatic equals to zero. Alternatively, one can assign classical charges to the atoms in the quantum region and then evaluate E_QMMM_electrostatic in the usual classical way, i.e. as a sum of pairwise Coulomb terms between these quantum atoms and the classical atoms.
- And for the assignment of classical charges to quantum atoms, people have used either a fixed set of charges for the QM atoms, for example, those given by the force field,
- or they have computed the charges somehow, by for example a least-squares fitting procedure to optimally reproduce some quantity such as the electrostatic potential at the surface of the QM subsystem, and then re-compute the charges, ideally even at each integration step.

An improvement to this mechanical embedding is to include polarization effects of the quantum region due to the presence of the classical atoms in the MM region. In fact, in this so called electrostatic embedding scheme, the electrostatic interactions between the quantum and the MM subsystem are handled during the computation of the electronic wave function, that is the E_QMMM_electrostatic, which I remind you has to be added in the quantum Hamiltonian, depends on both the classical charges and the quantum charge density.

Increasing further the level of sophistication implies to include in the model also the polarizability of the MM atoms. In this electrostatic coupling approach, called polarization embedding scheme, both regions, QM and MM, can mutually polarize each other.

Although this last embedding offers the most realistic electrostatic coupling between the quantum and the classical regions, polarizable force field for biomolecular simulations are not so effective yet. Therefore, despite progresses in the development of such force fields, QM/MM studies with polarizable MM regions are so far not so popular.

Slide 9

The CP2K QM/MM interface adopt the electrostatic embedding strategy that now I will describe in more details.

Slide 10

By construction, the QM/MM electrostatic energy term that in the electrostatic embedding one should add to the quantum Hamiltonian is the red expression in this slide. This means that the electrons (and protons) that define rho "see" the MM atoms as special nuclei with non-integer and possibly even negative charges q_I.

As it is, this expression immediately raises two problems, originating from to the peculiar short-range and long-range behaviour of this integral:

- the so-called electron spill-out
- and the very large computational cost required to evaluate the integral.

Let us start with the issue at short range.

Slide 11

A problem that may arise when using standard partial charges to describe the charge distribution in the MM subsystem is the risk of over-polarization near the boundary. Note that now there are no covalent bonds in between quantum and classical regions. However, also in this case, the point charges on the MM side may attract (or repel) the quantum electrons too strongly, which could lead to electron density "spilling out" into the MM region, as sketched in the picture.
In reality this would not be possible due to the Pauli repulsion of the electrons of the atoms in the MM regions, that however are missing in a QM/MM scheme.
This phenomenon of electronic spill-out at the boundary can become serious if a large flexible basis set or worse a nonlocal basis set is used in the QM calculations as happens for example in CP2K, because with such basis sets the electrons are fully free to delocalize.

Slide 12

This artefact can be avoided by modelling classical charges as smeared-out charges instead of the traditional point charges, where the smearing function can be a Gaussian distribution or another suitable function centered at the MM atom.

The smearing function used in CP2K to this aim is this one.

For the ones of you that are very interested to the theoretical details, this is the exact potential energy function generated by a Gaussian charge distribution.

In contrast to the point charge model, with a smearing function the Coulomb interaction between the electrons and the smeared charge distribution does not diverge if the electrons approach the MM atoms, as sketched in the picture.

Slide 13

Let's now move to the second issue in the calculation of E_QMMM_electrostatic, the one at long-range. In fact, due to the Coulomb long-range behaviour, the computational cost to compute the electrostatic QM/MM energy term is surprisingly large.

In fact, this term is typically evaluated by collocating on the nodes of a grid in the QM region, the contribution coming from the MM potential. Therefore, the number of operations that a direct evaluation of this quantity requires can be estimated as the number of grid points times the number of MM classical atoms. Typical grid point numbers are of order of millions in 3D (~100 times ~100 time ~100), while MM classical atoms are order of 10,000 or even more in systems of biological interest. Therefore, it is evident that in a real system a brute force computation of this term is impractical.

Many QM/MM codes face the problem by adopting strategies such as hierarchical methods or multipole techniques, which are based on the observation that the contributions to E_QMMM_electrostatic of the MM atoms further and further away from the quantum region can be replaced by less expensive expressions.

These methods are very effective, but require a fine tuning of parameters to yield optimal performance, and which lead to a loss of accuracy that makes error control sometimes difficult.

Instead CP2K solves this problem by adopting a different approach based on the representation of the electrostatic potentials of the MM atoms, that is the electrostatic charge times the smearing function, as a sum of functions with different widths, derived from the so-called Gaussian expansion of the electrostatic potential (GEEP in short) in combination with the use of a real space multi-grid technique.

Even if the details of this approach are a bit beyond what I can introduce here, let us try to provide few more elements to help you understanding this a bit better.

Slide 14

The problem of the large computational cost of evaluating E_QMMM_electrostatic corresponds to the problem to map efficiently, in both performance and accuracy) noncompact functions (the left-hand side in GEEP) on a grid in the real space.

Multigrid approach means that we use different grids, finer and coarser, to represent on a computer continuous functions in the real space, like for example the functions on the right-hand side in the GEEP expression.

Remember what we have seen in the first part of this lecture: each grid in the real space corresponds to a cutoff in the reciprocal space, but also the Gaussian functions in the reciprocal space have an intrinsic cutoff.

Now, if we map a Gaussian function on the first grid level, whose reciprocal cutoff is equal to or bigger than the intrinsic cutoff of that Gaussian function, then every Gaussian will be represented on the same number of grid points irrespective of its width. In practice, a grid with 25 points per side is usually sufficient for an optimal Gaussian representation.

Therefore, in practice, when calculating the E_QMMM_electrostatics, the GEEP decomposition allows converting the problem of mapping a non-compact smearing function on a fine grid into the mapping of N_g approximately compact Gaussian functions on grids lower or at least equal to the fine grid, plus an additional non-compact function, R_low, which however is very smooth and therefore it can be mapped on the coarsest available grid.

Why do you do that? Because the sum of the contributions of all the grids, suitably interpolated, is approximately equal to the function mapped analytically on the finest grid, i.e. the result that we want to calculate. But the total cost of decomposition, collocation on a grid and interpolation results finally a couple of order of magnitude smaller than the direct evaluation of the potential on the finest grid because many Gaussians functions contribute only partially or do not contribute at all.
In fact, as sketch in the picture, atoms whose distance from the QM box is greater than the Gaussian collocation radius do not contribute to the potential on that grid level. However, all atoms contribute to the coarsest grid level through the long-range R_low part.

The algorithm to evaluate the QM/MM electrostatic potential on the finest grid can be outlined as follows:

- Each MM atom is represented as a continuous Gaussian-like charge distribution via the smearing function.
- The electrostatic potential generating from this charge is fitted through a GEEP expansion.
- Then, every Gaussian function of the GEEP expansion is mapped on the first grid level whose reciprocal cutoff is equal to or bigger than the cutoff of that particular Gaussian function.
- This way we can limit each Gaussian to only a finite domain without loss of accuracy, so that only MM atoms embedded into the QM box, or close to it, will contribute to the finest grid, as sketch in the picture.
- Finally, in the last step of the algorithm, the contribution at each grid level is interpolated starting from the coarsest grid level up to the finest level

because the QM/MM electrostatic potential we are looking for, i.e. the one with higher accuracy, is the one on the finest grid level.

Unlike other approaches, the lack of tuning parameters makes this multigrid implementation a totally free parameter scheme, without any significant loss of accuracy. Consequently, very stable simulations can be obtained, with optimal energy conservation properties.

OK, we have arrived at the end of this brief introduction to QM/MM approaches. If you have questions or doubts on this part, I'll be waiting for you in the Q/A session in few minutes.