

Biomolecular modeling

Marcus Elstner and Tomáš Kubař

Theoretical Chemical Biology, Karlsruhe Institute of Technology

(Dated: October 9, 2011)

Contents

I. Molecular mechanics: a classical description of molecules	3
A. The conceptual and chemical basis	3
B. Determination of the non-bonding parameters: Coulomb and vdW forces	8
C. Hydrogen Bonding	14
D. Determination of the bonded parameters: harmonic springs	15
E. The complete equation	18
F. Exercises	18
II. Geometry optimization and normal modes of vibration	19
A. Steepest-descent minimization (SD)	20
B. Conjugate gradient minimization	21
C. Hessian-update based methods: Newton–Raphson and quasi-Newton–Raphson	25
D. Harmonic approximation of molecular vibrations: the normal modes	26
E. Exercises	30
III. Molecular Dynamics Simulations	31
A. Integration of the equations of motion	34
B. The Euler method	35
C. The Verlet method	36
D. More advanced methods	38
E. Exercises	39
IV. Thermodynamic Ensembles	40
A. The microcanonical (NVE) ensemble: an isolated system	41

B. The canonical (NVT) ensemble: a closed system	42
C. The canonical NPT ensemble	48
D. Exercise	49
V. Non-bonded interactions	50
A. Introduction to electrostatic interaction	50
B. Periodic boundary conditions	52
C. Accelerating the calculation of non-bonded interactions – cut-off	54
D. Accelerating the calculation of electrostatic interactions – Ewald summation	56
E. Explicit solvent models – water	60
VI. Principles of statistical mechanics	62
A. Microcanonical ensemble – microstates for constant energy	62
B. Microscopic definition of entropy	67
C. Canonical ensemble – system in equilibrium with the environment	69
D. Canonical partition function – the way to the thermodynamic quantities	70
E. Exercises	70
VII. Thermodynamic properties of molecular systems	71
A. Driving forces of thermal processes. Equilibrium	71
B. Thermodynamic functions from statistical mechanics	73
C. Discrete and continuous systems	74
D. Vibrational contribution to the thermodynamic functions	77
E. Aiming at free energies	79

I. MOLECULAR MECHANICS: A CLASSICAL DESCRIPTION OF MOLECULES

A. The conceptual and chemical basis

Chemical bonding is a genuinely quantum effect, which cannot be understood on the grounds of classical physics. However, the solution of the Schrödinger equation is numerically very expensive, and only small molecules can be treated quantum mechanically (up to ≈ 100 atoms). To be able to treat larger molecules, it is necessary to find further approximations.

Two fundamental simplifications often made in quantum chemistry are the so called Born-Oppenheimer approximation (BOA) and the classical treatment of the nuclei. BOA requires the electron motion to be much faster than the motion of the nuclei (atoms), so that the electrons follow instantaneously the motion of the nuclei (being somehow ‘fixed’ to the nuclei).¹ The second approximation is the neglect of nuclear quantum effects, like tunneling effects (hydrogen) or zero-point vibrations.

With these simplifications, we have the picture of N electrons moving in the electrostatic potential of M nuclei. Then, we have to solve the Schrödinger equation for these N electrons, which can be a formidable task. Or, vice versa, we have the M nuclei ‘sitting’ within the ‘sea’ of N electrons! What is a chemical bond in such a case? What causes the attraction between two nuclei? In many cases, we do *not* have a large electron delocalization, thus there is nothing like a sea of electrons (which is the case in a metal). In organic molecules, we have two electrons in every bonding orbital, and that is how covalent bonding is explained. The main idea behind the empirical models of chemical bond is that the strength of, say, a C–H bond mediated through the bonding orbitals is similar in *all* C–H bonds. In other words, we have a localized/local phenomenon to deal with.

So, how can we model a covalent bond? Is it possible to use simple springs to approximate a bond? Consider the molecules H_2 , O_2 or N_2 : If we model the interaction between two atoms with a harmonic spring with the energy $E(x)$ given as a function of the interatomic

¹ This means, that the electrons never leave those orbitals that have been ‘assigned’ to them in the electronic ground state. But there are cases, when they ‘leave’ these ground state orbitals, e.g. in high energy collisions of nuclei. The electrons are excited then, a process not treated within the BO framework.

distance x

$$E(x) = \frac{1}{2}k(x - x_0)^2 \quad (\text{I.1})$$

then there are two parameters k and x_0 . The values of these parameters may be derived from spectroscopic measurements (bond length and vibrational frequency).²

In equilibrium, the force between the atoms

$$F(x) = -\frac{\partial E(x)}{\partial x} = k(x - x_0) \quad (\text{I.2})$$

vanishes for $x = x_0$, thus $x = x_0$ is the equilibrium geometry.

A diatomic molecule is a one-dimensional system, and thus it is equivalent to one particle with mass m connected to a spring with spring constant k .³ The force on the particle is due to Hooke's law proportional to $x - x_0$, and using Newton's second law, we have

$$F = ma = m \frac{d^2x}{dt^2} = -k(x - x_0) \quad (\text{I.3})$$

This ordinary differential equation has a solution

$$x(t) - x_0 = c_1 \cdot \sin \left[\sqrt{\frac{k}{m}} t \right] + c_2 \cdot \cos \left[\sqrt{\frac{k}{m}} t \right] \quad (\text{I.4})$$

The comparison with the known equation for harmonic motion

$$x(t) = c \cdot \sin(\omega t) \quad (\text{I.5})$$

provides the relation of the frequency with the force constant k and mass m :

$$\omega = \sqrt{\frac{k}{m}} \quad (\text{I.6})$$

From another point of view, we obtain k directly from the second derivative of the energy:

$$k = \frac{d^2 E(x)}{dx^2} \quad (\text{I.7})$$

This will hold in more complex molecules as well: the second derivatives of the energy with respect to the atomic coordinates determine the frequencies of (harmonic) motion of the atoms in a molecule.

² In addition, there is also the information about the heat of formation. However, a harmonic spring does not allow the bond to dissociate, therefore this information cannot be used until a "better" force field is used, like a Morse potential.

³ The strict derivation introduces the reduced mass of the system.

Therefore, we can parametrize a simple force field from experiment, once we know the equilibrium distances and the vibrational frequencies. Alternatively, we can obtain these values from quantum-chemical calculations.

Does this mean that we can take a protein molecule and put springs between all the atoms? If yes, does every bond need different values of k and x_0 , because every bond is in a different chemical environment? If this is the case, we would not obtain any benefit. The essence of empirical modeling of chemical bonding is that the number of necessary parameters is much smaller than the number of bonds. In other words, we will use several (not many) “types” of chemical bonds with assigned parameters to represent any bond in our molecules.

In the other extreme, can we assign (for instance) each C–H bond the same k and x_0 value? It is not quite simple, but let us see:

1. Spectroscopy

Molecules consist of units/groups that have similar properties in every molecule. For instance, every C–H has a length of $r = 1.06\text{--}1.11$ Å and a vibrational frequency $\tilde{\nu} \approx 3100$ cm⁻¹, whatever its environment in the molecule is.

2. Thermochemistry

Molecular enthalpies of formation are approximately additive, so that:



This illustrates that the C–H potential looks like the Morse potential for every C–H unit in any chemical environment. The depth of the minimum is adjusted to fit the enthalpies of formation.

3. The concept of atom type

These points indicate that bonds between atoms can be modeled by universal potentials, if one tries to identify *atom types* in similar chemical environments (groups):




FIG. 1:

- Hybridization: clearly, an sp^3 carbon with four single bonds differs from an sp^2 carbon with a double bond or even an sp carbon with a triple bond. So, we need different potentials for the C–C, C=C and C \equiv C bonds. Also, there will be various different C atom types as well, like an aromatic carbon etc. Therefore, we introduce different carbons and determine the parameters (k, x_0) using some selected molecules which are typical for this chemical environment. For carbon, we use C_2H_6 , C_2H_4 , C_2H_2 and benzene to determine the parameters (k, x_0) for these four different carbons.
- A carbon atom bonded to an oxygen is electron deficient and this directly affects its bonding to other atoms. If we have a situation like in $O=CH-C\dots$, the C–C bond will be affected and it is thus desirable to introduce an additional C type – a carbonyl carbon, which uses different force constants for the adjacent C–C bond.

Biomolecular force fields typically use about 20 different C, 10 N and 5 O and H atom types.

4. Localization of the wavefunction


The quantum mechanical basis of these findings is that the electron wave function is localized. It can be localized in a bond, or localized on a fragment of a molecule. Such a fragments may constitute of functional groups or larger pieces like amino acids, DNA bases, sugars etc. The localization is crucial in at least two respects:

- Definition of the atom types: Electron delocalization can lead to different properties of the atomic fragments, like the different carbons in the carbonyl group, benzene etc.

- Electrostatic interactions. In the force fields, atomic charges are defined — i.e. every atom is assigned a point charge and interacts with other atoms according to Coulomb’s law

$$E_{\text{QQ}} \propto \frac{Q_1 Q_2}{R_{12}} \quad (\text{I.8})$$

These charges have to be estimated carefully, and it is very helpful to use the rule of *group neutrality*: Every functional group or molecular fragment has an integer charge.



5. Bonded and non-bonded interactions

Then, there are two distinct types of interactions between atoms in an extended molecule:

- interactions mediated by, and resulting directly from the presence of a covalent bond between the atoms. We usually put springs between the atoms and have to care about bond and dihedral angles. With this treatment, we describe all the quantum-mechanical phenomena like exchange and correlation using an effective potential between two atoms (like discussed above for a diatomic molecule).
- classical Coulomb interactions and van der Waals (vdW) forces between atoms, which are both long-range. For bonded atoms, these interactions are already effectively treated through the bonded parameters. These interactions are thus excluded between atoms which are neighbors (up to the fourth neighbor).⁴

⁴ Usually, the non-bonded interaction between an atom and its third neighbor, so called 1–4 interaction, is taken into account but scaled down by a multiplicative (“fudge”) factor.

B. Determination of the non-bonding parameters: Coulomb and vdW forces

The Coulomb interaction consists of three contributions: the nucleus–nucleus repulsion

$$\frac{1}{2} \sum_{ij} \frac{Z_i \cdot Z_j}{R_{ij}} \quad (\text{I.9})$$

the nucleus–electron attraction,

$$-\sum_i \int \frac{Z_i \cdot \rho(r)}{|R_i - r|} dr \rightarrow -\sum_{ij} \frac{Z_i \cdot Q_j}{R_{ij}} \quad (\text{I.10})$$

(the electron charge density $\rho(r) \rightarrow \sum_j Q_j$ is approximated by the sum of atomic point charges Q_j),

and the classical (Hartree) electron–electron interaction term is approximated as interaction of point charges sitting on the atoms

$$\frac{1}{2} \iint \frac{\rho(r) \cdot \rho(r')}{|r - r'|} \rightarrow \frac{1}{2} \sum_{ij} \frac{Q_i \cdot Q_j}{R_{ij}} \quad (\text{I.11})$$

If we define an *effective atomic net charge* as $q_i = -Q_i + Z_i$, we can write the total Coulomb energy as:

$$E_{QQ} = \frac{1}{2} \sum_{ij} \frac{q_i \cdot q_j}{R_{ij}} \quad (\text{I.12})$$

Thus, we have to determine the effective atomic charge for every atom type (or atom).⁵ This may be conveniently done by performing quantum-chemical calculations. Making use of the localization of the wave function, we can calculate the typical molecular fragments and try to determine the charges from there. These fragments are individual amino acid residues for proteins and the bases (e.g. uracil in Fig. 2), sugars and phosphates for DNA. In the uracil example, we can see that there are three different hydrogen atom types, one nitrogen, one oxygen and two carbons. However, there are two issues associated with this procedure:

First, atomic charges are difficult to define at all; there are several schemes to calculate them and it might be difficult to judge which is the best. Nowadays, Mulliken charges are no longer used because their drawbacks have become evident. A popular strategy is to use the

⁵ The term “atom type” is used with respect to bonded and vdW interaction. Usually, the atomic charges must be determined for more specifically defined atoms.




FIG. 2:

so-called *potential-derived charges*. Here, one calculates the *electrostatic potential* (ESP) of the molecule at its surface (see Fig. 3) and then fits the atomic charges in order to reproduce ESP for a set of points on the surface. The fitting constitutes in minimizing the error R :

$$R = \sum_i (\phi_i - \phi_i^0)^2 \quad (\text{I.13})$$

with ϕ_i being ESP induced by the system of point charges at the reference point i and ϕ_i^0 being the ESP obtained previously in the quantum-chemical calculation.




FIG. 3:

Second, charges are calculated in the gas phase (i.e. for an isolated molecule), while the electrostatic situation in an aqueous solution is different — molecules are more polar. As an example, the water molecule has a dipole of about 1.8 D in the gas phase, while it is about

2.4 D in solution. To account for that, charges may be taken to be larger than the gas phase values. A widely used strategy has been to use a small basis sets in the QM calculations, as it is known that such calculations overestimate molecular dipole moments.

These force fields use static atomic charges and, therefore, neglect the effects of polarization and charge transfer between the atoms. Presently, *polarizable force fields* are becoming more and more popular. Here, *atomic polarizability* α_i is assigned to every atom i . In a linear response approach, an external electric field induces a dipole at the atom:

$$\vec{\mu}_i = \alpha_i \cdot \vec{E} \quad (\text{I.14})$$

Of course, the external electric field is generated by all other atoms present in the system. This phenomenon will be discussed later on.

A further important non-bonded contribution is the repulsion driven by Pauli exclusion principle. In contrast to the classical Coulomb interaction, it is of a purely quantum-mechanical origin. Two electrons with same spin try to avoid a large spatial overlap. A typical example is the interaction of two neutral, closed-shell systems like two He atoms.⁶ If the electron densities of both He atoms start to overlap as in Fig. 4, the Pauli repulsion sets in. This interaction is an *exchange effect* and decays exponentially with the spatial overlap. Despite its pure quantum-mechanical character, we can model this effect conveniently by an exponential term

$$E_{\text{ex}} \propto \exp[-a \cdot R_{ij}] \quad (\text{I.15})$$

As we will see later, the exponential decay is not a computationally efficient model, and most empirical force fields use a R^{-12} decay.




FIG. 4:

⁶ ...being the crudest approximation of two benzene molecules :-)

The Pauli repulsion relates to the motion of electrons with the same spin, on short intermolecular distances. In addition, we also have a correlation of electrons irrespective of their spin, which remains even on longer distances. Two electrons repel each other due to the Coulomb interaction, trying to move apart from each other as far as possible. Consider now two atomic fragments with no overlap as in Fig. 5. Due to the zero-point energy, no quantum particle is ever at rest. If an electron in the neutral fragment A fluctuates in such a way, that it creates a dipole in A, then this dipole will induce a dipole in the opposite direction in the fragment B. Thus, the fluctuations of dipoles will be correlated due to the Coulomb interaction. This correlated fluctuation of dipoles in the two fragments leads to an




FIG. 5:

effective attractive interaction of the fragments, because of the opposite orientation of the interacting dipoles.⁷

A simple model in Fig. 6 describes this qualitatively. Two systems of one negative and one positive charge each are connected with a spring at a distance r . The negative and positive charge in every pair can be separated by z_1 and z_2 , if some force acts on them. The complex system has a force constant k and (reduced) mass m with $\omega = \sqrt{k/m}$. The separation of charges in small systems brings on dipole moments of $\mu_1 = z_1 q$ and $\mu_2 = z_2 q$. The Schrödinger equation for one oscillator

$$-\frac{\hbar^2}{2m} \frac{\partial^2 \psi}{\partial z_1^2} + \frac{1}{2} k z_1^2 \psi = E \psi \quad (\text{I.16})$$

has a set of solutions given by

$$E = \left(\nu + \frac{1}{2} \right) \cdot \hbar \omega \quad (\text{I.17})$$

⁷ This effect is not described at the Hartree–Fock level of quantum mechanics, and is also missing in density-functional theory. MP2 or CI would cover it.




FIG. 6:

Now, let us couple the oscillators due to the dipole-dipole interaction:

$$V(r) = -\frac{2\mu_1\mu_2}{4\pi\epsilon_0 r^3} \quad (\text{I.18})$$

If we insert this interaction potential into the Schrödinger equation, we find

$$E(r) = -\frac{q^4\hbar\omega}{2(4\pi\epsilon_0)^2 K^2 r^6} \quad (\text{I.19})$$

Now, the force on a point charge q in an electric field of intensity E is

$$F = qE \quad (\text{I.20})$$

i.e. this force leads to a displacement z , $F = kz$:

$$qE = kz \quad (\text{I.21})$$

Therefore, the induced dipole moment of such an oscillator in the electric field (E) is

$$\mu_{ind} = qz = q^2 E/k \quad (\text{I.22})$$

At the same time, we have

$$\mu_{ind} = \alpha E \quad (\text{I.23})$$

Therefore

$$\alpha = q^2/k \quad (\text{I.24})$$

and we can write Eq. I.19 as

$$E(r) = -\frac{\alpha^2\hbar\omega}{2(4\pi\epsilon_0)^2 k^2 r^6} \quad (\text{I.25})$$

This is the well-known expression for the dispersive interaction, which is attractive with a R^{-6} dependence and dependent on the polarizabilities of the interaction partners.

The most common function which combines the Pauli repulsion and the dispersive interaction is the Lennard-Jones 12-6 potential

$$V(r) = 4\varepsilon \left(\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right) \quad (\text{I.26})$$

which is shown in Fig. 7. The minimum lies at $r_m = 2^{1/6}\sigma$ and the well depth is $-\varepsilon$, and σ and ε are treated as empirical parameters. The slight issues with the exp-6 potential

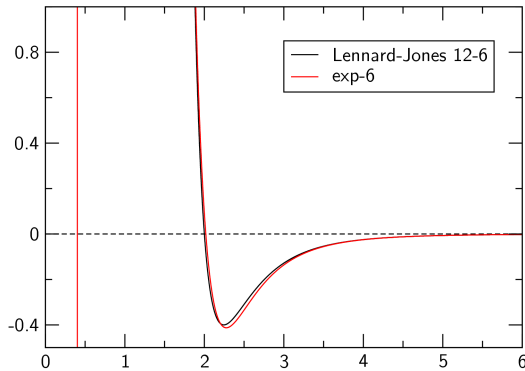


FIG. 7: The Lennard-Jones potential with $\varepsilon = 0.4$ and $\sigma = 2$, and a similar exp-6 potential.

are shown in Fig. 7 as well. It might be difficult to find appropriate parameters, which would render the vdW minimum correctly, and the function becomes negative for small r . However, the main reason for the exp-6 potential not to be used (except if vitally needed)⁸ is that it is much more computationally intensive to evaluate the exponential function than R^{-6} .

To find the vdW parameters for heteronuclear interactions, several “mixing rules” have been proposed, the simplest being:

$$\begin{aligned} \sigma_{AB} &= \frac{\sigma_{AA} + \sigma_{BB}}{2} \\ \varepsilon_{AB} &= \sqrt{\varepsilon_{AA} \cdot \varepsilon_{BB}} \end{aligned} \quad (\text{I.27})$$

To find a good set of non-bonded parameters is a challenging task, and force fields are constantly being improved, meaning that new sets of parameters are proposed. A crucial

⁸ which may be the case in the studies of phase transitions like freezing

test and calibration of a force field is the density and the heat of vaporisation of organic liquids, since these properties depend critically on the magnitude of non-bonded parameters — the charges and vdW parameters.

C. Hydrogen Bonding

Early force fields contained special potentials to describe hydrogen bonding. Hydrogen bond is the attractive interaction between a hydrogen atom and an electronegative atom like oxygen or nitrogen. The hydrogen must be covalently bonded to another electronegative atom, so that this covalent bond is polarized and the hydrogen bears a partial positive charge. A typical example is the hydrogen bond between two water molecules in Fig. 8 left. Typical binding energies are around 20 kJ/mol but may reach higher values if the binding partners are strongly polarized or even charged, or if there are several hydrogen bonds between them, like in the GC base pair in DNA, see Fig. 8 right.




FIG. 8: Water dimer (left) and the guanine:cytosine nucleobase pair (right).

Clearly, Coulomb interaction is the dominant contribution to the hydrogen bond, but vdW interaction can become important also, especially in weakly bound systems. It has been shown that they become crucial especially to describe the angular dependence of the interaction energy in complexes like $\text{H}_2\text{CO} \dots \text{H}_2\text{O}$ and similar.

Thus, force fields have everything in place to describe these phenomena (Coulomb and vdW terms), and most modern force fields do not require a special treatment for this bonding type. However, a third contribution, the *charge transfer* (Fig. 9), is not captured by force fields at all. It may be included in the other terms in an effective way.

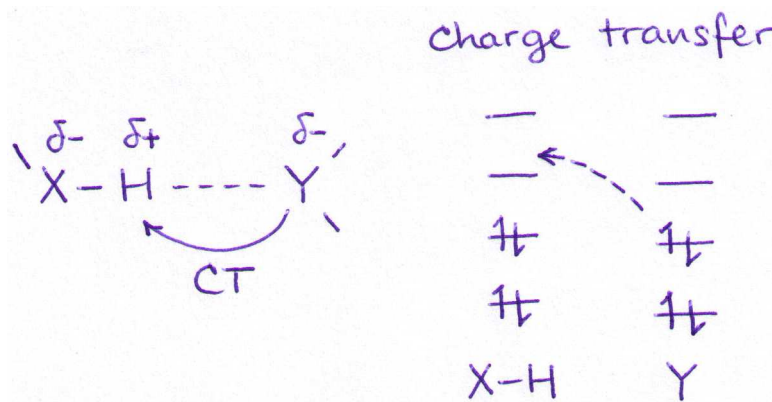


FIG. 9: Charge transfer between the charge donor (Y) and acceptor(X-H).

D. Determination of the bonded parameters: harmonic springs

1. Bonds

A bond dissociation curve typically looks like that in Fig. 1, and is well reproduced by the Morse potential with the functional form

$$E(r) = D (1 - \exp[-\alpha(r - r_0)])^2 \quad (\text{I.28})$$

In principle, the Morse potential allows for a quite good representation of the potential, in a broad range of distances. However, it is computationally quite inefficient, because of the presence of the exponential, and the Morse potential is rarely used in molecular mechanics studies.

A way to approximate virtually any function is to apply the Taylor expansion

$$E(r) = E(r_0) + \frac{dE}{dr}(r_0) \cdot (r - r_0) + \frac{1}{2} \frac{d^2E}{dr^2}(r_0) \cdot (r - r_0)^2 + \dots \quad (\text{I.29})$$

Most force fields use a harmonic approximation, i.e. the Taylor expansion cut after the second-order term. $E(r_0)$ is a constant which we set to zero, and the first derivative $\frac{dE}{dr}(r_0)$ vanishes if the function has a minimum in r_0 . Therefore, with the definition of the force constant k we have, in the second order:

$$E(r) = \frac{1}{2} k (r - r_0)^2 \quad (\text{I.30})$$

We can immediately see in Fig. 10 that the approximation of the true potential with a quadratic function is quite crude and holds only in a quite narrow interval of distances

around the minimum. Importantly, the vibrations of covalent bonds exhibit quite a small magnitude (of several tenths of ångström) and so we actually never need to describe the potential beyond this narrow interval.⁹ The application of harmonic approximation is then justified. To parametrize such a force field, we need two parameters per bond: the force constant k and the equilibrium distance r_0 .

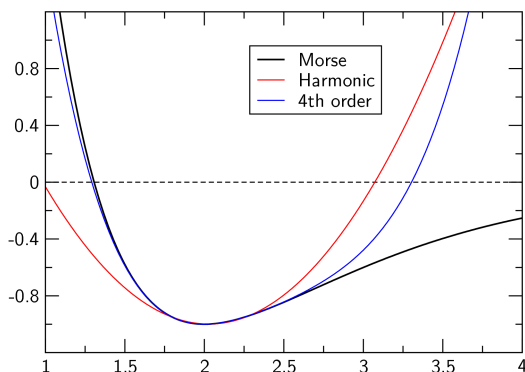


FIG. 10: Comparison of the harmonic and quartic functions with the Morse potential

To be able to compute accurate vibration frequencies, terms up to fourth order can be important to describe the curvature of the potential energy dependence; Fig. 10 compares the Morse potential with the harmonic and fourth-order functions. Also, the quartic function approximates the Morse potential a wider interval of distances.

2. Angles

As for the bonds, we apply a harmonic approximation for the angle deformation and get the potential

$$E_{\text{bend}}(\vartheta) = \frac{1}{2}k_{\vartheta}(\vartheta - \vartheta_0)^2 \quad (\text{I.31})$$

Again, we can obtain the parameters from experiment. For instance, let us consider a



⁹ of course, if we do not aim at studying chemical reactions, i.e. having covalent bonds created or broken

water molecule, which has an equilibrium bond angle of $\vartheta_0 = 104^\circ$ and a bending vibration frequency of about 1700 cm^{-1} .

3. Dihedrals

Dihedral angles describe the rotation around covalent bonds. Four atoms are needed to define this angle (ω):




Clearly, the dependence of potential energy on the dihedral angle will be described with a periodic function, and a usual way to express this is a sum of several cosine functions:

$$E(\omega) = \sum_{n=1,2,3,4,6} V_n \cos[n \cdot \omega - \gamma_n] \quad (\text{I.32})$$

with V_n being the height of the torsional barrier, n giving the periodicity ($n = 1$: 360° , $n = 2$: 180° , $n = 3$: 120°) and phase offsets γ_n .

Now, consider the C–C single and the C=C double bonds as examples. The single bond has a periodicity of 120° , i.e. we have three minima for one full rotation of 360° , and the potential energy is described as



$$E_{\text{C-C}}(\omega) = V \cdot \cos 3\omega \quad (\text{I.33})$$

The C=C double bond has a 180° -periodicity and thus only two minima for the full rotation. The energy is given by

$$E_{C=C}(\omega) = V \cdot \cos[2\omega - 90^\circ] \quad (\text{I.34})$$

E. The complete equation

Adding up all contributions, the total energy of a typical biomolecular force field reads

$$\begin{aligned} E(R^N) = & \frac{1}{2} \sum_i k_i (r_i - r_i^0)^2 + \frac{1}{2} \sum_j k_j^\vartheta (\vartheta_j - \vartheta_j^0)^2 + \frac{1}{2} \sum_n V_n \cdot \cos[n\omega - \gamma_n] \\ & + \sum_i^N \sum_{j=i+1}^N \left\{ 4\varepsilon_{ij} \left(\left(\frac{\sigma_{ij}}{r_{ij}} \right)^{12} - \left(\frac{\sigma_{ij}}{r_{ij}} \right)^6 \right) + \frac{1}{4\pi\varepsilon_0} \frac{q_i q_j}{r_{ij}} \right\} \end{aligned} \quad (\text{I.35})$$

F. Exercises

1. Show for the Lennard-Jones potential (Eq. I.26) that the minimum is at $r_m = 2^{1/6}\sigma$ and the well depth is $-\varepsilon$.
2. Calculate the first and second derivative of the general force field (Eq. I.35). Only the terms dependent on r_i and r_{ij} are required.

II. GEOMETRY OPTIMIZATION AND NORMAL MODES OF VIBRATION

Consider again a diatomic molecule. The total force-field energy

$$E(x) = \frac{1}{2}k(x - x_0)^2 \quad (\text{II.1})$$

is positive if the distance between the atoms x does not equal the reference bond length x_0 .

Consider the case, that we obtain the dimer in a geometry with the bond distance of $x_1 \neq x_0$. Now, how can we determine the optimal geometry in that case, i.e. the geometry (distance) with the lowest energy? We calculate the force F acting upon both atoms, as the negative of the gradient g of energy:

$$F(x) = -g(x) = -\frac{\partial E(x)}{\partial x} = -k(x - x_0) \quad (\text{II.2})$$

To approach towards the minimum, we calculate the force at the actual distance x_1

$$F(x_1) = -k(x_1 - x_0) \quad (\text{II.3})$$

If we move the atoms in direction of the force until the force is zero, we will reach a stationary point.

Once there, we have to check if it is indeed a minimum (and not a saddle point, for instance). We can do this by evaluating the second derivative

$$k = \frac{\partial^2 E(x)}{\partial x^2} \quad (\text{II.4})$$

If this is positive (corresponding to a real vibrational frequency $\omega = \sqrt{k/m}$), then we are in a minimum. Otherwise, the frequency is imaginary and the coordinates correspond to a saddle point on the potential energy surface (which cannot happen with just one harmonic potential). The second derivative of energy describes the curvature of the potential energy function and is directly related to the vibrational frequency.

This procedure is what we call *energy minimization* or *geometry optimization*.

Generally, our energy function depends on $3N$ atomic coordinates (x, y, z for each of N atoms). It is convenient to write all the coordinates in one vector

$$\vec{r} = (x_1, y_1, z_1, x_2, \dots, z_N) \quad (\text{II.5})$$

In three dimensions (corresponding to one atom), the gradient reads¹⁰

$$\vec{g} = \nabla E(\vec{r}) = \left(\frac{\partial E}{\partial x}, \frac{\partial E}{\partial y}, \frac{\partial E}{\partial z} \right) \quad (\text{II.6})$$

For N atoms, \vec{g} and \vec{F} are $3N$ -dimensional vectors:

$$\vec{g} = \left(\frac{\partial E}{\partial x_1}, \frac{\partial E}{\partial y_1}, \frac{\partial E}{\partial z_1}, \frac{\partial E}{\partial x_2}, \dots, \frac{\partial E}{\partial z_N} \right) \quad (\text{II.7})$$

The unit vector in the direction of the gradient is given by

$$\vec{e} = \frac{\vec{g}}{|\vec{g}|} \quad (\text{II.8})$$

A. Steepest-descent minimization (SD)

Within the method of steepest descent, the optimizer moves *iteratively* – in steps \vec{h} along the direction of the force

$$\vec{h} = \alpha \cdot \vec{e} \quad (\text{II.9})$$

The critical point here is the choice of the step size α . If the step is too long, we follow the gradient down the potential though but may miss the minimum along the gradient and go up the valley on the opposite side. If the step is too short, we may need to perform too many steps, which in turn means too many (costly) evaluations of energy and forces.




FIG. 11: Steepest descent minimization

One way to overcome this problem is to perform a *line search* along the direction \vec{e} , and find a minimum on this line. In other words, we are looking for a value of α_k such that r_{k+1}

¹⁰ using *nabla* — the formal vector of partial derivatives $\nabla \equiv \left(\frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z} \right)$




FIG. 12: Line search

is the minimum along the search direction e_k :

$$\vec{r}_{k+1} = \vec{r}_k + \alpha_k \vec{e}_k \quad (\text{II.10})$$

The interesting and important point is that two successive search directions (steps) are *orthogonal* to each other:

$$\vec{h}_{k-1} \cdot \vec{h}_k = 0 \quad (\text{II.11})$$

A problem may arise if the energy function forms a narrow valley. In such a case, the second next search direction will be similar. Therefore, an efficient strategy attempts to avoid this double work and looks for search directions that are orthogonal to *all* previous ones.




FIG. 13: Problem of SD in a narrow valley

B. Conjugate gradient minimization

Consider the Taylor series of a $3N$ -dimensional function up to the second order:

$$\begin{aligned} E(\vec{r}) &= E(\vec{o}) + \sum_i \frac{\partial E}{\partial r_i}(\vec{o}) \cdot r_i + \frac{1}{2} \sum_{ij} \frac{\partial^2 E}{\partial r_i \partial r_j}(\vec{o}) \cdot r_i r_j \\ &= c - \vec{b} \cdot \vec{r} + \frac{1}{2} \cdot \vec{r} \cdot \mathbf{A} \cdot \vec{r} \end{aligned} \quad (\text{II.12})$$

with the constants c , \vec{b} and \mathbf{A} given as

$$c = E(\vec{o}), \quad \vec{b} = -\nabla E(\vec{o}), \quad \mathbf{A}_{ij} = \frac{\partial^2 E}{\partial x_i \partial x_j}(\vec{o}) \quad (\text{II.13})$$

If we express the derivative of E in a point \vec{r} , we obtain

$$\nabla E(\vec{r}) = \mathbf{A} \cdot \vec{r} - \vec{b} \quad (\text{II.14})$$

To minimize the function E now means to find the point \vec{r}_{min} where the derivative vanishes: $\nabla E(\vec{r}_{min}) = 0$. This task corresponds to the solution of this set of linear equations:

$$\mathbf{A} \cdot \vec{r} = \vec{b} \quad (\text{II.15})$$

To do this, we could calculate the first (\vec{b}) and second (\mathbf{A} – Hessian) energy derivative and solve the set of equations.

With that, we would directly obtain the minimum \vec{r}_{min} , but

- To evaluate the second derivatives is computationally very expensive. If possible, we try to avoid such a calculation whenever we can (to calculate the gradient, we need $3N$ derivatives, while for the Hessian we need $(3N)^2$).
- Usually, the potentials of interest are not simple harmonic potentials, i.e. we only make a *local* harmonic approximation. Thus, the obtained solution \vec{r}_{min} is not the minimum we are looking for, but only a step towards the true minimum. Practically, we would have to perform such calculations iteratively, which would become very time-consuming.

Therefore, we look for methods which use information in the form of gradients *only*. As we have seen on the example of the SD method, successive gradients can have similar directions, which leads to a lot of double work. The conjugate gradient method constructs successive search directions, which are conjugate to each other in a certain way.

Practically, we perform the first step \vec{h}_1 along the gradient \vec{g} :

$$\vec{h}_1 = -\vec{g}_1 \quad (\text{II.16})$$

The second step should be ‘conjugate’ to the first step, i.e. it should not go in the same direction.

Now, how does the gradient of E change if we move along a certain direction \vec{h}_k ? Let us consider two successive steps k and $k + 1$. We have:

$$\begin{aligned}\vec{g}_k &= \mathbf{A} \cdot \vec{r}_k - \vec{b} \\ \vec{g}_{k+1} &= \mathbf{A} \cdot \vec{r}_{k+1} - \vec{b} \\ \vec{g}_{k+1} - \vec{g}_k &= \mathbf{A} \cdot (\vec{r}_{k+1} - \vec{r}_k)\end{aligned}\tag{II.17}$$

This means that going along the search direction $\vec{h}_k = -\vec{g}_k$, we obtain the change of the gradient $\Delta_k \vec{g} = \vec{g}_{k+1} - \vec{g}_k$. This is the effect of the search direction $-\vec{g}_k$. Now, when moving along the next search direction \vec{h}_{k+1} , we do not want to lose the fruit of the work done so far, i.e. we wish to keep the change of gradient that has already been made ($\Delta_k \vec{g}$). Therefore, the gradient shall remain *orthogonal*¹¹ to \vec{h}_k . This can be achieved if the change of gradient along \vec{h}_{k+1} is orthogonal to \vec{h}_k :

$$\begin{aligned}\vec{h}_k \cdot \Delta_{k+1} \vec{g} &= 0 \\ \vec{h}_k \cdot (\vec{g}_{k+2} - \vec{g}_{k+1}) &= \vec{h}_k \cdot \mathbf{A} \cdot (\vec{r}_{k+2} - \vec{r}_{k+1}) = \vec{h}_k \cdot \mathbf{A} \cdot \vec{h}_{k+1} = 0\end{aligned}\tag{II.18}$$

This condition is a generalization of the concept of orthogonality, and the two vectors \vec{h}_{k+1} and \vec{h}_k are denoted as *conjugate* (with respect to \mathbf{A}). They are orthogonal if the matrix \mathbf{A} is a unit matrix.

Now, we want to perform sequential steps as with the line search algorithm,

$$\begin{aligned}\vec{r}_{k+1} &= \vec{r}_k - \alpha_k \cdot \vec{g}_k \\ \mathbf{A} \cdot (\vec{r}_{k+1} - \vec{r}_k) &= -\alpha_k \cdot \mathbf{A} \cdot \vec{g}_k = \vec{g}_{k+1} - \vec{g}_k\end{aligned}\tag{II.19}$$

If we choose α_k such that \vec{r}_{k+1} is the minimum along the search direction \vec{g}_k , we know that the two successive gradients are orthogonal. We can multiply the last equation with \vec{g}_k to get:

$$\alpha_k = \frac{\vec{g}_k \cdot \vec{g}_k}{\vec{g}_k \cdot \mathbf{A} \cdot \vec{g}_k}\tag{II.20}$$

This equation assures the orthogonality of \vec{g}_k and \vec{g}_{k+1} , so that every search direction is orthogonal to the previous one, and determines the step size. For that, we need the gradient and the Hessian.

¹¹ or *perpendicular* in our simple 2D drawings

However, the calculation of the Hessian is often too costly, and the step should thus be determined in another way. Several algorithms to do so have been developed, and we will briefly describe the Fletcher–Reeves method here: The new search direction at \vec{r}_{k+1} is calculated as (without proof)

$$\vec{h}_{k+1} = -\vec{g}_{k+1} + \gamma_k \cdot \vec{h}_k \quad (\text{II.21})$$

with

$$\gamma_k = \frac{\vec{g}_{k+1} \cdot \vec{g}_{k+1}}{\vec{g}_k \cdot \vec{g}_k} \quad (\text{II.22})$$

assuring

$$\begin{aligned} \vec{g}_k \cdot \vec{g}_l &= 0 \\ \vec{g}_k \cdot \vec{h}_l &= 0 \\ \vec{h}_k \cdot \mathbf{A} \cdot \vec{h}_l &= 0 \end{aligned} \quad (\text{II.23})$$




FIG. 14: Conjugate gradients

- This has to be compared with the steepest descent method. In SD, the search directions $-\vec{g}_{k+1}$ and $-\vec{g}_k$ are generally not orthogonal to each other, which can lead to the case that successive steps spoil the efforts of each other. The conjugate-gradient directions \vec{h}_k are not constructed to be orthogonal, either; however, it is their property of being conjugate ($\vec{h}_k \cdot \mathbf{A} \cdot \vec{h}_l = 0$) that increases the efficiency of the minimization significantly.
- In principle, CG determines the minimum of a N -dimensional quadratic potential in N steps. Every vector in an N -dimensional space can be constructed from N basis vectors. In many cases, these are chosen to be orthogonal. It can be shown, that the

vectors \vec{h}_k form a linearly independent set of vectors, and so they represent an optimal set of search directions to find the minimum in N steps. The vector determining the path from the initial point \vec{r}_1 to the minimum \vec{r}^* can then be written as

$$\vec{r}^* - \vec{r}_1 = \sum_k \alpha_k \cdot \vec{h}_k \quad (\text{II.24})$$

C. Hessian-update based methods: Newton–Raphson and quasi-Newton–Raphson

Now, consider the vector

$$\vec{r} = \vec{r}^* - \vec{r}_1 \quad (\text{II.25})$$

where \vec{r}^* is the global minimum and \vec{r}_1 an arbitrary starting point. Since the gradient vanishes in the minimum, we find from

$$\mathbf{A} \cdot \vec{r} = \vec{b} \quad (\text{II.26})$$

that

$$\begin{aligned} \vec{r} &= \mathbf{A}^{-1} \cdot \vec{b} \\ \vec{r}^* &= \vec{r}_1 + \mathbf{A}^{-1} \cdot \vec{b} \end{aligned} \quad (\text{II.27})$$

so that we are able to find the minimum of a quadratic potential in one step, simply by inverting the Hessian \mathbf{A} . This is a very favorable property, because in CG we would have needed around N steps (N – dimensionality of the problem). The Hessian contains the information about the curvature of the function, with large curvature giving large eigenvalues. Therefore, the inverse of the Hessian leads to large steps in regions of low curvature and vice versa. This property will speed up the convergence for shallow potentials, where it would be very slow just to follow the gradient.

A trivial example is the parabola $f(x) = x^2$. For this function, the so called Newton–Raphson (NR) method as described in Eq. II.27 gives:

$$0 = r_1 - \frac{1}{2} \cdot 2 \cdot r_1 \quad (\text{II.28})$$

(with $f''(r) = 2$ and $f'(r) = 2r$). Compare this with a flat potential like $f(r) = 0.1 \cdot r^2$. The inverse of the second derivative $f''(r) = 0.2$ is $[f''(x)]^{-1} = 5$, giving a large step size.

Practically, there are some issues:

- to evaluate the Hessian and to invert it (an $\mathcal{O}(N^3)$ -operation) can be very costly, so that the CG method becomes computationally cheaper, in particular if
- the potential is not harmonic. Then, NR may even fail (converge to a saddle point, proceed in a wrong direction etc.) and it may not be the method of choice.

Therefore, one usually starts the optimization (when forces are large) with SD and CG, and NR is invoked only later – close to the minimum, where the harmonic approximation works better.

Very efficient are so-called *Hessian-update methods*, where the inverse Hessian is not calculated, but rather updated during the optimization.

$$\vec{r}_{k+1} = \vec{r}_k - \mathbf{A}_k^{-1} \cdot \vec{g}_k \quad (\text{II.29})$$

One starts with $\mathbf{A}_1^{-1} = \mathbf{I}$ (unit matrix) and the matrix is ‘updated’ in every step, so that it is identical to the inverse Hessian for large k . This means that in the beginning, one follows the gradient, and the information about the curvature is collected along the way.

All optimization techniques require a criterion to stop the procedure. Common practice is to stop if the gradient and the change of energy between two steps become small (smaller than a preset threshold).

D. Harmonic approximation of molecular vibrations: the normal modes

The displacement from equilibrium of a harmonic oscillator $r(t)$ is given by the solution of the ordinary differential equation

$$m \cdot \ddot{r} = -k \cdot r \quad (\text{II.30})$$

with k being the second derivative of the potential with respect to the displacement r . If we insert $r(t) = A \cdot \sin(\omega t)$, we immediately obtain

$$-m \cdot \omega^2 = -k \quad (\text{II.31})$$

and the vibrational frequency $\omega = \sqrt{k/m}$.

If we think of large molecules like proteins with N atoms, the problem becomes quite complicated. Every atom interacts with each other, so we have to deal with N^2 interactions,

which are usually not harmonic (Coulomb and vdW interactions, many-body interactions, although the bonded terms are harmonic).

In the *harmonic approximation*, we take the force-field energy I.35 and apply the Taylor approximation around the equilibrium positions of the atoms, with r_i being the deviations from the equilibrium (the first derivative vanishes at equilibrium):

$$E(\vec{r}) = E(\vec{r}_0) + \frac{1}{2} \sum_{i,j}^N \frac{\partial^2 E}{\partial r_i \partial r_j} \cdot r_i \cdot r_j \quad (\text{II.32})$$

With

$$k_{ij} = \frac{\partial^2 E}{\partial r_i \partial r_j} \quad (\text{II.33})$$

and setting $E(\vec{r}_0) = 0$ we have

$$E(\vec{r}) = \frac{1}{2} \sum_{i,j}^N k_{ij} \cdot r_i \cdot r_j \quad (\text{II.34})$$

Note that the k_{ij} introduced here are different from the force constants in the force field!

The force on atom i is now the (negative of the) derivative of the energy with respect to the atomic coordinate r_i :

$$F_i = -\frac{\partial E}{\partial r_i} = -\sum_j^N k_{ij} \cdot r_j \quad (\text{II.35})$$

Looking at atom i , we have

$$m_i \cdot \ddot{r}_i = F_i = -\sum_j^N k_{ij} \cdot r_j \quad (\text{II.36})$$

Since we have E , we know how to calculate the second derivatives, once we have located the minimum. The movement of atom i is now coupled to all other atoms, and the equations we have to solve become quite complicated. In case of a diagonal matrix $k_{ij} = k_{ii}\delta_{ij}$, we would have

$$m_i \cdot \ddot{r}_i = k_{ii} \cdot r_i \quad (\text{II.37})$$

for which we know the solution

$$r_i(t) = a_i \cdot \sin[\omega t] \quad (\text{II.38})$$

where we have a system of uncoupled oscillators. We can use this as an ansatz to solve the system of coupled oscillators, to obtain

$$m_i \cdot \omega^2 \cdot a_i = - \sum_j^N k_{ij} \cdot a_j \quad (\text{II.39})$$

The vector $\vec{a} = (a_1, a_2 \dots a_{3N})$ gives the amplitude of the motion, i.e. determines how much an atom is involved in the motion. We define $\lambda = \omega^2$ and the (diagonal) mass matrix $\mathbf{M}_{ii} = m_i$. With that, we can write:

$$\lambda \cdot \mathbf{M} \cdot \vec{a} = \mathbf{k} \cdot \vec{a} \quad (\text{II.40})$$

Multiplying from the left with \mathbf{M}^{-1} , we get

$$\lambda \cdot \vec{a} = \mathbf{M}^{-1} \cdot \mathbf{k} \cdot \vec{a} \quad (\text{II.41})$$

This is an eigenvalue problem for the matrix $\mathbf{M}^{-1} \cdot \mathbf{k}$, which we know how to solve. The only problem is that this matrix is not symmetric, which makes the numerical solution more difficult. Therefore, we look how to get an eigenvalue problem for a symmetric matrix.

This can be done if we define

$$\vec{b} = \mathbf{M}^{1/2} \cdot \vec{a} \quad \text{i.e.} \quad \vec{a} = \mathbf{M}^{-1/2} \cdot \vec{b} \quad (\text{II.42})$$

and get instead of II.40 this problem:

$$\lambda \cdot \mathbf{M}^{1/2} \cdot \vec{b} = \mathbf{k} \cdot \mathbf{M}^{-1/2} \cdot \vec{b} \quad (\text{II.43})$$

Now, we multiply from the left with $\mathbf{M}^{-1/2}$ and obtain

$$\begin{aligned} \mathbf{M}^{-1/2} \cdot \lambda \cdot \mathbf{M}^{1/2} \cdot \vec{b} &= \mathbf{M}^{-1/2} \cdot \mathbf{k} \cdot \mathbf{M}^{-1/2} \cdot \vec{b} \\ \lambda \cdot \vec{b} &= \tilde{\mathbf{k}} \cdot \vec{b} \end{aligned} \quad (\text{II.44})$$

with a symmetric matrix $\tilde{\mathbf{k}}$ given as


$$\tilde{\mathbf{k}}_{ij} = \frac{k_{ij}}{\sqrt{m_i} \sqrt{m_j}} \quad (\text{II.45})$$

Now, we can solve a standard eigenvalue problem

$$(\tilde{\mathbf{k}} - \lambda \cdot \mathbf{I}) \cdot \vec{b} = 0 \quad (\text{II.46})$$

to get the eigenvalues λ_ν and eigenvectors \vec{b}_ν . The $N \times N$ matrix $\tilde{\mathbf{k}}$ has N eigenvalues λ_ν and eigenvectors \vec{b}_ν , which correspond to the *normal modes* of the system. These represent the independent modes of motion of the molecule. Every eigenvector \vec{b}_ν contains the amplitudes of every atom in the system (in the directions x , y and z), contributing to the particular normal mode ν : $\vec{b}_\nu = (b_{1\nu}, b_{2\nu}, \dots, b_{N\nu})$.

For example, think of the three modes of vibration of a water molecule:



Example:

Consider a fictitious linear triatomic molecule, with two equal bond lengths and force constants, atom masses m_1 , m_2 and $m_3 = m_1$, and harmonic springs only between atoms 1–2 and 2–3. The potential energy E then reads

$$E = \frac{1}{2}k \cdot (x_1 - x_2)^2 + \frac{1}{2}k \cdot (x_2 - x_3)^2 \quad (\text{II.47})$$

with x_i being the displacements from the equilibrium positions.

We have the Hessian

$$\mathbf{k} = \begin{pmatrix} k & -k & 0 \\ -k & 2k & -k \\ 0 & -k & k \end{pmatrix} \quad (\text{II.48})$$

the mass matrix

$$\mathbf{M} = \begin{pmatrix} m_1 & 0 & 0 \\ 0 & m_2 & 0 \\ 0 & 0 & m_1 \end{pmatrix} \quad (\text{II.49})$$

and the mass-weighted Hessian

$$\mathbf{k} = \begin{pmatrix} \frac{k}{m_1} & -\frac{k}{\sqrt{m_1 m_2}} & 0 \\ -\frac{k}{\sqrt{m_1 m_2}} & 2\frac{k}{m_2} & -\frac{k}{\sqrt{m_1 m_2}} \\ 0 & -\frac{k}{\sqrt{m_1 m_2}} & \frac{k}{m_1} \end{pmatrix} \quad (\text{II.50})$$

The solution of the eigenproblem leads to three normal modes:

$$\begin{aligned}
 \lambda &= \frac{k}{\sqrt{m_1}} : & b_1 &= -b_3, \quad b_2 = 0 \\
 \lambda &= 0 : & b_1 &= b_3, \quad b_2 = \sqrt{m_2/m_1} \cdot b_1 \\
 \lambda &= k \cdot \frac{m_2+2m_1}{m_1m_2} : & b_1 &= b_3, \quad b_2 = -2\sqrt{m_1/m_2} \cdot b_1
 \end{aligned} \tag{II.51}$$

The first mode describes a motion with the central atom at rest, and m_1 and m_3 vibrating with opposite amplitudes – symmetric stretch. The second mode corresponds to pure translation of the whole molecule along the x axis. Since $a_1 = a_3 = \frac{b_1}{\sqrt{m_1}}$, we find that $a_2 = a_1$. In the third mode, the central atom moves in the opposite direction to the outer atoms – asymmetric stretch.

E. Exercises

1. Consider the function $f(x, y) = x^2 + 2y^2$. Calculate the gradient in the point (2,2) and the Hessian in (0,0). Would you reach the minimum when moving along that gradient? What are the eigenvalues of the Hessian in (0,0)? Start in the point (0,1) and move along the gradient with the step size of 1. Where do you arrive at?
2. Starting in the point (9,9), a line search in the first step finds the point (4,−1). Show that the new search direction constructed in this point with CG directly reaches (0,0).
3. Show that you arrive in the minimum in one step if you use the Newton–Raphson method.
4. Consider the function $f(x, y) = x^4 + 4x^2y^2 - 2x^2 + 2y^2$. Plot this function and determine the eigenvalues of the Hessian in the minimum (1,0) and in the saddle point (0,0). What are the eigenvectors?
5. Derive the Hessian of the artificial linear triatomic system (Eq. II.50).

III. MOLECULAR DYNAMICS SIMULATIONS

Consider a chemical reaction in solution. The molecules of solute are surrounded by many water molecules, which fluctuate and thereby assume different hydrogen bonding interactions with the solute. The different orientations of water molecules induce different dipole moments on the solute. As a result

- the energy of the system will fluctuate a lot;
- since there are a multitude of local minima close in energy, the search for a global minimum will become extremely difficult if not meaningless;
- even if a global minimum is found, the product state may have a very different global minimum than the reactant state with respect to the conformation of the water molecules. In this case, we do not know how to compare the energies at all!

The **(micro)state** of a system is completely determined, if we know the positions \vec{r}_i and the momenta \vec{p}_i of all the particles. We call the $6N$ -dimensional space spanned by the coordinates and the momenta the **phase space**, $\{\vec{r}_i, \vec{p}_i\}$.

A **trajectory** in the phase space is the sequence of points $\{\vec{r}_i(t), \vec{p}_i(t)\}$ passed by during the dynamics of the system.

Example:

Consider the simple harmonic oscillator. We can plot the time dependence of the coordinate

$$r(t) = a \cdot \cos[\omega t] \quad (\text{III.1})$$

as well as the time dependence of the velocity


$$v(t) = -a\omega \cdot \sin[\omega t] \quad (\text{III.2})$$

On the other hand, we can also plot the velocities versus the coordinates in the 2-dimensional phase space, and we see easily that the trajectory in the phase space is elliptic:

$$\left(\frac{x(t)}{a}\right)^2 + \left(\frac{v(t)}{a \cdot \omega}\right)^2 = 1 \quad (\text{III.3})$$

The total energy of the harmonic oscillator is

$$E = T + V = \frac{1}{2}m\omega^2 a^2 \quad (\text{III.4})$$



Therefore, different values of a and ω correspond to different energies. The trajectories in phase space never cross for such so-called **conservative systems** (systems conserving the total energy). If we have *friction* or some other *damping*, the oscillator is losing energy (non-conservative system) and the motion will look like spirals ending in the origin. Here, different trajectories may cross.

Now, let us return to the molecule in solution. It does not make much sense to look at a single structure; rather, we may wish to average the property of interest over all available structures. For example, we follow the trajectory in phase space for some time and evaluate the energy for many **snapshots** $\{\vec{r}_i(t_k), \vec{p}_i(t_k)\}$ along the trajectory. So, we obtain values of energy E_k in time instants t_k ($k = 1, \dots, M$) and calculate the average:

$$\langle E \rangle = \frac{1}{M} \sum_{k=1}^M E_k \quad (\text{III.5})$$

We do this for the product and for the reactant, obtain average energies for both states, and thus the energy difference. In principle, this is a very good idea, but there are a few issues:

- How can we assure that we have enough points (snapshots) to represent all possible (or at least relevant) conformations of the molecular system?
- How do we obtain the trajectory? \rightarrow perform MD! Then, how to consider the temperature?
- Suppose we know the structure of the reactant, then how do we get the structure of the product and, if desired, the whole reaction path? This is a difficult problem and we will have to spent some time on that.
- What about entropy? Does the average of energy provide sufficient information? This is another difficulty, so we will have to look deeper into thermodynamics and statistics.

Here, we are at the core of biomolecular simulations: It is easy to understand the total energy of the force field, and how geometry optimization works. The main issue is to make proper use of the energy function to get the *thermodynamic properties* right. Biomolecular simulations are all about thermodynamics, in possible contrast to quantum chemistry, where people look at small molecules with quite simple structures and with a focus just on total energies. At the moment, we will concentrate on the former two points. The latter two points will be discussed in next chapters.

If we have a trajectory, we can calculate so-called **time averages** for the properties of interest, like structural properties or velocities. For any property $A(t)$, this works the same way as for energy, so that we may generalize the Eq. III.5:

$$\langle A(t) \rangle_t = \frac{1}{t_1 - t_0} \int_{t_0}^{t_1} A(t) dt \quad (\text{III.6})$$

On the other hand, an experimental sample contains a huge number of molecules. This number is so large that we can assume all relevant conformations of molecule and solvent to be present. Think of a large number of harmonic oscillators, then for every point in the phase space \vec{r}_i, \vec{p}_i , there will be for every point (on the ellipse) an oscillator, which is at that point. Such a collection of systems is called an **ensemble**.

Now, let us count how many molecules (oscillators) in the ensemble are found in a phase-space point \vec{r}_i, \vec{p}_i as well as in all the other points. Doing this, we obtain the **phase space density**, i.e. the number of realizations of each point in the phase space (per a unit of volume):

$$\rho(\vec{r}, \vec{p}) \quad (\text{III.7})$$

If we integrate ρ over the entire phase space, we obtain the total number of systems (oscillators or molecules):

$$Z = \int \rho(\vec{r}, \vec{p}) d\vec{r} d\vec{p} \quad (\text{III.8})$$

We can use this phase space density to calculate the **ensemble average**:

$$\langle A \rangle_e = \frac{1}{Z} \int A \cdot \rho(\vec{r}, \vec{p}) d\vec{r} d\vec{p} \quad (\text{III.9})$$

Experimentally, we always measure ensemble averages; however, in a simulation, we always have a single system for which we calculate time averages. To be able to compare

both, we always assume that the systems we study are **ergodic**. This means that the system passes through all points of the phase space constituting the real ensemble during the MD, provided the simulation is long enough. If the system is indeed ergodic, then these averages are equal:

$$\langle A \rangle_e = \langle A \rangle_t \quad (\text{III.10})$$

We leave the statistics for a moment and have a look, how we can get the trajectories in computer simulations.

A. Integration of the equations of motion

The total energy $E = T + V$ is often termed *Hamilton function* H , or Hamiltonian in quantum mechanics. The Hamilton formalism is a practical way to derive equations of motion, once we know the total energy:

$$\dot{r}_i = \frac{\partial H}{\partial p_i} \quad \dot{p}_i = -\frac{\partial H}{\partial r_i} \quad (\text{III.11})$$

We cannot prove these equations here, but give an example (momentum $p = m \cdot v$):

$$\begin{aligned} H &= \frac{1}{2} \frac{p^2}{m} + \frac{1}{2} k r^2 \\ \dot{r} &= \frac{\partial H}{\partial p} = \frac{p}{m} \\ \dot{p} &= -\frac{\partial H}{\partial r} = -k \cdot r \end{aligned} \quad (\text{III.12})$$

which, if put together ($\dot{p} = m \cdot \ddot{r}$ from the eqn for \dot{r}), give the well-known equation of motion of the harmonic oscillator. These are **ordinary differential equations** (ODE) of 1st and 2nd order.

A 1st-order differential equation may look like e.g.

$$\dot{x} = -k \cdot x \quad (\text{III.13})$$

which has the solution $x(t) = A \cdot \exp[-k \cdot t]$. Examples of phenomena in nature following such an exponential law are radioactive decay or the dynamics of populations.

More generally, we can write these ODEs as

$$\dot{x} = f(x, t) \quad (\text{III.14})$$

A 2nd-order ODE then follows as

$$\ddot{x} = f(x, \dot{x}, t) \quad (\text{III.15})$$

An example is the equation of motion of a harmonic oscillator $\ddot{x} = -\frac{k}{m} \cdot x$, or, when we introduce damping

$$\ddot{x} = -\zeta \cdot \dot{x} - \frac{k}{m} \cdot x \quad (\text{III.16})$$

The damping is assumed to be proportional to velocity, which is very often the case.

We can reduce the 2nd-order ODE to two 1st-order ODEs by introducing v :

$$\begin{aligned} \dot{x} &= v \\ \dot{v} &= -\zeta \cdot v - \frac{k}{m} \cdot x \end{aligned} \quad (\text{III.17})$$

Now, we have to solve these ODEs numerically. There are several methods available, which we will briefly discuss.

B. The Euler method

To solve a 1st-order ODE

$$\dot{r} = f(r, t) \quad (\text{III.18})$$

we proceed by a Taylor expansion ($\Delta t = t - t_0$):

$$r(t) = r(t_0) + \dot{r}(t_0) \cdot \Delta t + \frac{1}{2} \ddot{r}(t_0) \cdot \Delta t^2 + \dots \quad (\text{III.19})$$

The **Euler approximation** is a first-order approximation, so that we neglect the second and higher orders:

$$r(t) \approx r(t_0) + \dot{r}(t_0) \cdot \Delta t \quad (\text{III.20})$$

We start the numerical integration at time t_0 , take a time step Δt and get the value of r and v in the time $t = t_0 + \Delta t$. For our MD simulation, we have

$$\begin{aligned} r(t_0 + \Delta t) &= r(t_0) + v(t_0) \cdot \Delta t \\ v(t_0 + \Delta t) &= v(t_0) + a(t_0) \cdot \Delta t \\ a(t_0) &= -\frac{1}{m} \frac{\partial H}{\partial r} = -\frac{F}{m} \end{aligned} \quad (\text{III.21})$$

where H is the Hamilton function. Instead, we can use the potential energy V only, and the derivative of V is the force upon the atom.

To start the MD, we have to specify the *initial conditions* – the positions r_0 and the velocities v_0 , and calculate the force at the positions x_0 to get the accelerations a_0 . Then, we can use the Euler equations to get the positions and velocities at the time $t_0 + \Delta t$.

This means, to obtain a trajectory over a time interval T , we have to evaluate the forces on all atoms $M = T/\Delta t$ times, i.e. to perform M steps. Therefore, the computational cost of the calculation of forces determines how many steps we can afford to make.

The length of the time step is a crucial parameter in the simulation. There is a numerical and a chemical constraint on the step size:

- Numerical: Since we neglect the contributions in Δt^2 and higher orders, we introduce an error in the order of Δt^2 ($\mathcal{O}(\Delta t^2)$). A possibility is to make the step very short, so that the error is small. But then, we need to perform too many steps to simulate a certain time T . on the other hand, if we make the step too long, we introduce an error in the integration. We will notice this error in the total energy: $E = T + V$ should be conserved, and if we monitor the total energy during the MD we can see, how good our integration is.
- Chemical: The fastest motion in biological systems is the movement of hydrogen atoms, with the period of around 10 fs. A rule of thumb recommends the step size to be at least an order of magnitude smaller than the period of the fastest motion. Practically, this leads to a step size of 1 fs.

The step size (of typically 1 or several fs) is the fundamental time unit in MD. This means, we need 1 million calculations of energy and forces to get a trajectory in a total length of 1 ns. For large systems, multi-nanoseconds simulations can be routinely done, and even microsecond regime may be reached for smaller ones.

C. The Verlet method

The numerical error of the Euler method is too large to allow for time steps of 1 fs. Therefore, methods with errors in $\mathcal{O}(\Delta t^3)$ have to be used. A popular one is the Verlet method. Here, we make a virtual step in positive time and in ‘negative’ time, and expand

up to second order:

$$\begin{aligned} r(t + \Delta t) &= r(t) + \dot{r}(t) \cdot \Delta t + \frac{1}{2} \ddot{r}(t) \cdot \Delta t^2 \\ r(t - \Delta t) &= r(t) - \dot{r}(t) \cdot \Delta t + \frac{1}{2} \ddot{r}(t) \cdot \Delta t^2 \end{aligned} \quad (\text{III.22})$$

We may add both equations to eliminate the velocity \dot{r} :

$$\begin{aligned} r(t + \Delta t) &= 2 \cdot r(t) - r(t - \Delta t) + \ddot{r}(t) \Delta t^2 \\ \ddot{r}(t) &= a(t) = -\frac{1}{m} \frac{\partial V}{\partial r}(t) = -\frac{F(t)}{m} \end{aligned} \quad (\text{III.23})$$

This equation seems to be a little strange, as it requires not only the positions $r(t)$ and the accelerations $a(t)$, but in addition the positions one step back $r(t - \Delta t)$! This may look like a problem, because at the start of a MD, we know only $r(t_0)$, $v(t_0)$ (and $a(t_0)$ via forces) but not $r(t_0 - \Delta t)$. This would mean that Verlet could not be started. We have to use the initial velocities to calculate the positions an imaginary step back $r(t_0 - \Delta t)$ to start Verlet:

$$r(t_0 - \Delta t) = r(t_0) - v(t_0) \cdot \Delta t \quad (\text{III.24})$$

The algorithm does not contain velocities explicitly; if these are of interest, they may be obtained as

$$\dot{r}(t) = v(t) = \frac{r(t + \Delta t) - r(t - \Delta t)}{2 \cdot \Delta t} \quad (\text{III.25})$$

This form of the algorithm is called the Verlet normal form.

Several equivalent variations of the algorithm have been developed, like the **velocity Verlet**.¹² Here are new positions calculated first

$$r(t + \Delta t) = r(t) + v(t) \cdot \Delta t + \frac{1}{2} a(t) \cdot \Delta t^2 \quad (\text{III.26})$$

followed by the evaluation of forces (and thus accelerations) in the new positions. New velocities are then obtained as

$$v(t + \Delta t) = v(t) + \frac{1}{2} (a(t) + a(t + \Delta t)) \cdot \Delta t \quad (\text{III.27})$$

which are used in the next calculation of positions r .

¹² Another often used variant is the leap-frog.

Here, MD can be started with the knowledge of x_0 and v_0 ; however, in every next step, new positions $r(t + \Delta t)$ must be calculated first so that accelerations may be updated in order to obtain $v(t + \Delta t)$. Although both mentioned forms of the Verlet algorithm are equivalent, the velocity Verlet exhibits better numerical precision.¹³

D. More advanced methods

The Verlet method (with the error of $\mathcal{O}(\Delta t^3)$) is a very approximative one though, but it is routinely used in MD simulations for its favorable performance due to the necessity to evaluate only second derivatives of positions (accelerations). Plenty of more accurate methods to solve ODEs have been developed, and we will briefly mention a few of them.

A straightforward way to improve the quality of integrator would be to involve several extra terms from the Taylor expansion, i.e. to cut the expansion at some further point.

The classical Runge–Kutta method evaluates the first derivative in several points (selected with a special algorithm, k_i) in every step, and the function r is then integrated using a weighted average of these values:

$$r_{n+1} = r_n + \frac{1}{6}\Delta t \cdot (\dot{r}(k_1) + 2 \cdot \dot{r}(k_2) + 2 \cdot \dot{r}(k_3) + \dot{r}(k_4)) \quad (\text{III.28})$$

This reduces the error to $\mathcal{O}(\Delta t^5)$ (fourth-order method).

The Gear’s predictor–corrector family of methods provides solutions correct to an order of choice. Here, new position etc. are calculated (‘predicted’) from the Taylor expansion as usually. Then, the forces are calculated in the new positions to give accelerations. These are then used to ‘correct’ the new positions, and new positions are then calculated finally. This method is rarely used due to the considerable computational requirements, although it may provide quite accurate solutions of equations of motion as well as make a longer time step possible.

¹³ The numerical problem of normal Verlet is that it adds a small but important term $\ddot{r}(t_0)\Delta t^2$ to a large term calculated as a difference $2r(t) - r(t - \Delta t)$.

E. Exercises

1. Derive the equations of motion for the linear triatomic system with Hamiltonian

$$H = \frac{1}{2} \sum_{i=1}^3 m_i v_i^2 + \frac{1}{2} k (x_1 - x_2)^2 + \frac{1}{2} k (x_2 - x_3)^2 \quad (\text{III.29})$$

2. Derive the equations of motion for the particle i (r_i and p_i) using the Hamilton formalism for

$$H = \frac{1}{2} \sum_i m_i v_i^2 + \frac{1}{2} \sum_{ij} V_{ij} \quad (\text{III.30})$$

V_{ij} is given by the Lennard-Jones potential discussed in Chapter 1. These are the equations we have to solve when doing real MD for a system of uncharged particles without covalent bonds, like argon atoms.

IV. THERMODYNAMIC ENSEMBLES

A system of classical particles interacting with a potential V is a **deterministic system**. Knowing the **initial conditions**, which consist of the positions and velocities of all particles at the initial time, we can calculate the **trajectory** of the system, i.e. the positions $r(t)$ and velocities $v(t)$ at all future times t . We can have **analytic** solutions, as in the case of the harmonic oscillator, where we have the trajectories in a closed form:

$$x(t) = x_0 \cdot \cos[\omega t] \quad v(t) = -v_0 \cdot \sin[\omega t] \quad (\text{IV.1})$$

or, if the system becomes too complicated, we have to compute the trajectories **numerically**, for instance using the Verlet method.

Note that so-called **chaotic** systems are also strictly deterministic. For completely specified initial conditions, the system is perfectly predictable. The point here is that two trajectories which are initially close in phase space, may depart exponentially from each other, while they would stay close in non-chaotic systems (e.g. harmonic oscillator). We say that the solution of the underlying differential equations is unstable. This may happen already for seemingly simple systems like a double pendulum!




FIG. 15: For a ‘chaotic’ system, two initially close trajectories depart exponentially. Therefore, small differences in the initial conditions may lead to very different behavior of the system, although being completely deterministic.

A motion becomes **stochastic**, if we do not have the information about all degrees of freedom, like for a dust particle moving in an erratic fashion. If we could keep track about the movements of all air molecules, the particle motion would be completely deterministic for us. We have to describe systems using **statistical mechanics**, if we cannot keep track of all degrees of freedom.

We are now interested to find simulation techniques, where we can control the temperature. The temperature is a crucial parameter, which determines if a certain part of the phase space is to be reached during the MD. The phase space will be sampled differently at high temperatures than at low ones, and different ensembles will be generated. In particular, we have to find a way to model the system, that we simulate the right phase space density.




FIG. 16: A high energy $E = E_{\text{kin}} + E_{\text{pot}}$ allows to sample more different parts of the phase space. The difference $E - E_{\text{pot}}$ corresponds to the kinetic energy E_{kin} .

A. The microcanonical (NVE) ensemble: an isolated system

A system is called **isolated** if it exchanges neither energy (in the form of heat or work) nor matter (particles) with the environment. In this case, the total energy of the system is constant and given by the sum of the kinetic energy E_{kin} and the potential energy E_{pot}

$$E = E_{\text{kin}} + E_{\text{pot}} \quad (\text{IV.2})$$

Like for the harmonic oscillator, the kinetic and potential energies fluctuate in time as they are being transformed into each other all the time, keeping the total energy constant. This is what we describe when using the plain Verlet method for a large molecule after specifying the initial conditions. Looking into different regions of the molecule, the kinetic energy can be very different locally.

It is known from kinetic theory of gases that the kinetic energy is related to the temperature T

$$\langle E_{\text{kin}} \rangle = \frac{3}{2} N k T \quad (\text{IV.3})$$

with $\langle E_{\text{kin}} \rangle = \frac{1}{2} \sum_i m_i \langle v_i^2 \rangle$. So, if we stuck a thermometer into various regions of the entire system (with N particles), we would measure different temperatures, fluctuating in time.

This is not quite the situation we normally have in the experimental setup (a test tube with a sample). Usually, our systems are in **thermodynamic equilibrium** with the environment, so that they have the same temperature (and optionally pressure) as the environment.




FIG. 17: Left: Isolated system. Right: Closed system; exchange of heat in the canonical ensemble.

The canonical ensembles allow for the exchange of heat and work, but not matter (particles).¹⁴ If we keep the volume, temperature and number of particles constant, we call the system an **NVT ensemble**. If we keep pressure, temperature and the number of particles constant, allowing the volume to change, it is called an **NPT ensemble**.

B. The canonical (NVT) ensemble: a closed system

In the **canonical ensemble**, the thermal contact of the system with the environment leads to an exchange of energy in the form of heat, until the same temperature T as in the environment is reached. Strictly spoken, temperature is only defined if there is such a thermal contact, and therefore not applicable to the closed systems as discussed above.

For a classical system in thermal equilibrium is the velocity of a particle/atom (its component in any direction) and the speed¹⁵ given by the **Maxwell–Boltzmann distribution**

¹⁴ There is another thermodynamic ensemble with constant volume, temperature and chemical potential, allowing the number of particles to change = exchange of matter with the environment (grand canonical ensemble).

¹⁵ *Velocity* is the vector \vec{v} and its magnitude is denoted as the scalar *speed* (v).

$$p(v_{x,i}) = \sqrt{\frac{m_i}{2\pi kT}} \cdot \exp\left[-\frac{m_i v_{x,i}^2}{2kT}\right] \quad (\text{IV.4})$$

$$p(v_i) = 4\pi \left(\frac{m_i}{2\pi kT}\right)^{3/2} \cdot v_i^2 \cdot \exp\left[-\frac{m_i v_i^2}{2kT}\right] \quad (\text{IV.5})$$




FIG. 18: The distribution of velocity in the x direction (left) and speed (right) of a N_2 molecule (ideal gas approximation).

With this, the **equipartition theorem** (Gleichverteilungssatz) can be derived, which ascribes every degree of freedom the same amount of kinetic energy of

$$\left\langle \frac{1}{2} m_i v_{x,i}^2 \right\rangle = \frac{1}{2} kT \quad (\text{IV.6})$$

Since every atom i has three degrees of freedom x_i , y_i and z_i (and $v_i^2 = v_{x,i}^2 + v_{y,i}^2 + v_{z,i}^2$), we find for the total kinetic energy of the system

$$\langle E_{\text{kin}} \rangle = \left\langle \sum_i \frac{1}{2} m_i v_i^2 \right\rangle = \left\langle \sum_i \frac{1}{2} m_i v_{x,i}^2 + \frac{1}{2} m_i v_{y,i}^2 + \frac{1}{2} m_i v_{z,i}^2 \right\rangle = \frac{3}{2} N kT \quad (\text{IV.7})$$

The point is that this special velocity distribution is a property of systems in contact with a heat reservoir, which is not the case for the microcanonical ensemble.

We can use this formula to control the temperature in our system in a very simple way. An algorithm that allows to control the temperature during an MD simulation is called a **thermostat**.

1. Scaling of velocities

We start the simulation using one of the variants of the Verlet integrator, having specified the initial positions and velocities. Usually, we know the positions from (experimental struc-

ture etc.), and velocities can be drawn randomly from the Maxwell–Boltzmann distribution corresponding to a desired temperature (using a random number generator).

In the course of the simulation, the actual temperature will deviate from the target temperature, which may be understood as the temperature of the heat bath in equilibrium with the simulated molecular system:

$$T(t) = \frac{2}{3} \frac{E_{\text{kin}}(t)}{Nk} \neq T_{\text{ref}} \quad (\text{IV.8})$$

Note a discrepancy with the previous definition of thermodynamic temperature (Eq. IV.3) — temperature should correspond to the *average* of kinetic energy over time. However, Eq. IV.8 presents some sort of ‘instantaneous temperature,’ i.e. ‘temperature’ in a certain time t . This quantity will fluctuate in the course of simulation, whereas the true thermodynamic temperature is an average over this simulation.

The instantaneous temperature is just another name for the mean kinetic energy (Eq. IV.3), which is in turn determined by the velocities of the atoms. Now then, a simple idea how to bring the instantaneous temperature to the target is to *scale* the velocities of all atoms by a certain multiplicative factor λ . The required magnitude of this factor may be estimated by casting the ‘new’ velocities $\lambda \cdot v_i$ into the expression for the temperature and making this equal to T_{ref} :

$$\begin{aligned} T_{\text{ref}} &= \frac{1}{\frac{3}{2}Nk} \cdot \frac{1}{2} \sum_i m_i (\lambda \cdot v_i)^2 = \\ &= \lambda^2 \cdot \frac{1}{\frac{3}{2}Nk} \cdot \frac{1}{2} \sum_i m_i v_i^2 = \\ &= \lambda^2 \cdot T \end{aligned} \quad (\text{IV.9})$$

From this, the value of λ follows simply as

$$\lambda = \sqrt{\frac{T_{\text{ref}}}{T}} \quad (\text{IV.10})$$

So, we have to scale the velocities of all atoms by this factor, in order to obtain the target temperature T_{ref} exactly.

This is a very crude way of controlling the temperature. From time to time, we knock the system by rescaling the velocities, affecting the ‘natural’ way of evolution of the system. This is quite drastic and can lead to artifacts. More importantly, the system does not represent any canonical ensemble, i.e. its **phase space density** ρ is not that of a canonical

ensemble. As for the velocities, we never make sure that their distribution is correct (the Maxwell–Boltzmann distribution). This is very important because we calculate all properties of interest as averages:

$$\langle A \rangle = \frac{1}{Z} \int \rho \cdot A \, d\vec{r} d\vec{p} \quad (\text{IV.11})$$

With the velocity scaling method, it could happen that our system would **sample** the phase space incorrectly, leading to wrong averages.

2. Berendsen thermostat

A way to avoid the severe turns in the dynamics is to adjust the velocities more smoothly, resigning on the target temperature to be recovered in every step (like it is done with simple scaling).

With the Berendsen coupling scheme, the simulated system is considered as coupled to an infinite thermal bath with temperature T_{ref} . We ask the temperature to change between two time steps according to:

$$\frac{dT}{dt} = \frac{1}{\tau} (T_{\text{ref}} - T) \quad (\text{IV.12})$$

so that the rate of change of temperature (due to the change of velocities) is proportional to the deviation of the actual temperature from the target. The constant of proportionality is the relaxation time τ , and this 1st-order differential equation corresponds to an exponential decay of temperature towards the target. Thus, we want to change the temperature by

$$\Delta T = \frac{\Delta t}{\tau} (T_{\text{ref}} - T) \quad (\text{IV.13})$$

which will be achieved by scaling the velocities by a factor λ as above:

$$\Delta T = T_{\text{ref}} - T = (\lambda^2 - 1) \cdot T \quad (\text{IV.14})$$

$$\lambda = \sqrt{1 + \frac{\Delta t}{\tau} \left(\frac{T_{\text{ref}}}{T} - 1 \right)} \quad (\text{IV.15})$$

For $\tau = \Delta t$, we get the simple velocity scaling; for the usually applied larger values ($\tau = 0.1 - 1$ ps), the temperature of system is fluctuating around the target temperature (of the imaginary external bath).

This **fluctuation** of temperature is a desired property. For the canonical ensemble, we can calculate the *variance* of ‘instantaneous temperature’ T :

$$\sigma_T^2 = \langle (T - \langle T \rangle)^2 \rangle = \langle T^2 \rangle - \langle T \rangle^2 \quad (\text{IV.16})$$

and obtain the relative variance

$$\frac{\sigma_T^2}{\langle T \rangle^2} = \frac{2}{3N} \quad (\text{IV.17})$$

For large number of particles (atoms) N , the magnitude of fluctuations approaches zero. However, for finite-sized systems, we observe a visible fluctuation of temperature, which is indeed a feature of the canonic ensemble. So, if we kept the kinetic energy (and thus the instantaneous temperature) constant by means of the simple velocity scaling, we would not obtain the correct fluctuations.

There are several drawback with the velocity rescaling algorithms:

- It does not generate a rigorous canonical ensemble.
- Various parts of the system (different individual molecules, or solute \times solvent) may exhibit different temperatures, while the temperature of the entire system is correct. This discrepancy may be maintained for extended periods of time.
- It gradually moves the energy from the fastest modes of motion to the slowest/weakest ones, violating the equipartition theorem. The bond stretching and angle bending modes are usually the fastest modes, and the loss of energy of these modes manifests itself as a ‘freezing’ of the molecules. On the other hand, the three translations (and, for aperiodic systems, the three rotations) of the entire system are the slowest modes, and thus those that actually gain the energy. This problem is known as the ‘flying (or spinning) ice cube.’

3. Andersen thermostat

The Andersen thermostat has an entirely different working principle. From time to time, some particles (atoms) are selected to undergo a ‘collision’ with the particles of a heat bath, which changes their velocities suddenly. Such an algorithm exhibits a certain stochastic character.

This thermostat works as follows:

- Start the MD with a standard integrator (Verlet...)
- Select randomly the particles (atoms) that shall undergo a collision with the heat bath.
- For these particles, assign new velocities by a draw from the correct Maxwell-Boltzmann distribution. All other particles remain unaffected.

The advantage of the Andersen thermostat is that if correctly implemented, it really generates a canonical ensemble. However, for this to come true, the rate of collisions with imaginary particles must be neither too low (insufficient) nor too high (in that case, the collisions would dominate the dynamics of the system over the equations of motion).

4. Nosé-Hoover thermostat

The Nosé-Hoover thermostat rigorously represents a canonical ensemble and is thus usually the method of choice. However, it is conceptionally and mathematically slightly difficult to understand. We present the fundamental ideas here.

The heat reservoir is treated as a part of the system and is assigned an extra degree of freedom s , with which a fictitious mass Q is associated. The equations of motion are then derived for this *extended system* with $3N + 1$ degrees of freedom.

The equations of motion for the atoms are modified to

$$\ddot{\mathbf{r}}_i = \frac{\mathbf{F}_i}{m_i} - s \cdot \dot{\mathbf{r}}_i \quad (\text{IV.18})$$

with the second term formally corresponding to a kind of ‘friction’ due to the bath. Further, there is another equation of motion for the bath parameter s :

$$\dot{s} = \frac{1}{Q} (T - T_{\text{ref}}) \quad (\text{IV.19})$$

The strength of coupling of the real system to the bath is determined by the ‘mass’ parameter Q . This may be made more intuitive by introducing a time parameter τ :

$$Q = \frac{\tau^2 \cdot T_{\text{ref}}}{4\pi^2} \quad (\text{IV.20})$$

This parameter has the meaning of a *period of oscillation* of the kinetic energy between the real system and the bath. Here, it is important to note the difference between the τ in the Berendsen and in the Nosé-Hoover couplings: The Berendsen method consists in

the exponential damping of temperature difference, with coefficient τ . In the Nosé–Hoover framework, an oscillatory relaxation of temperature is achieved, with period τ .

This way, the prerequisites of the thermostat are incorporated in the equations of motion, and the thermostat thus constitutes an integral part of the integrator, rather than an *a posteriori* correction, as is the case of all previously mentioned ones. The Nosé–Hoover thermostat is used frequently and it can be shown that the ensemble average taken with the Nosé–Hoover ensemble is identical to that of the canonical ensemble.

C. The canonical NPT ensemble

Rather than under the condition of constant volume, chemical experiments are usually performed under constant pressure, often the atmospheric pressure. Therefore, it is of importance to be able to implement such conditions in MD simulations as well.

Similar to temperature, it is possible to calculate the pressure in the simulation. This is usually done by using the quantity called *the virial of force*:

$$\Xi = -\frac{1}{2} \sum_{i < j} \vec{r}_{ij} \cdot \vec{F}_{ij} \quad (\text{IV.21})$$

where \vec{r}_{ij} is the vector connecting atoms i and j , and \vec{F}_{ij} is the force acting between these atoms. The pressure then follows as¹⁶

$$P = \frac{2}{3V} \cdot (E_{\text{kin}} - \Xi) = \frac{2}{3V} \cdot \left(\frac{1}{2} \sum_i m_i \cdot |\vec{v}_i|^2 + \frac{1}{2} \sum_{i < j} \vec{r}_{ij} \cdot \vec{F}_{ij} \right) \quad (\text{IV.22})$$

A surprising feature of pressure calculated in this way is that it fluctuates immensely along a simulation and can even assume negative values!

Once a means to *calculate* the pressure is available, it is possible to develop an algorithm to *maintain* it at a constant value. These *barostats* are equivalents of the Berendsen or Nosé–Hoover algorithms to maintain constant temperature. Instead of velocities, it is the volume of the simulation box that is varied in order to adjust pressure; this is done either directly (Berendsen) or again via an additional degree of freedom (Parrinello–Rahman).

¹⁶ Note the difference between the kinetic theory of gases and this definition — the idea of the particles colliding with the wall of the container is not present here.




FIG. 19: Temperature and pressure in an NPT simulation of DNA oligomer in water ($T_{\text{ref}} = 300$ K, $P_{\text{ref}} = 1.0$ bar).

D. Exercise

Prove the equipartition theorem (Eq. IV.6) using the Maxwell–Boltzmann distribution (Eq. IV.5) as the phase-space density to evaluate the ensemble average $\langle v_{x,j}^2 \rangle_{\text{ens}}$ (Eq. III.9). To do this, you have to solve integrals of the forms

$$\int_0^\infty x^2 \exp[-ax^2] dx$$

and

$$\int_0^\infty \exp[-ax^2] dx$$

which may be found at http://en.wikipedia.org/wiki/Lists_of_integrals.

V. NON-BONDED INTERACTIONS

There are several reasons to take particular care of the non-bonded interactions:

- They are a key to understand the structure, function and in particular the efficiency of action of many proteins. It is the electrostatic and/or van der Waals interaction of the protein with the ligand that is responsible for the efficiency of a reaction, color of the chromophore etc. The solvent surrounding is co-responsible for the particular structure of nucleic acids, polypeptides and proteins (hydrophobic-hydrophilic residues).
- The non-bonded interactions are treated in MM by two-body potentials, and the computational effort of $\mathcal{O}(N^2)$ dominates the overall requirements for large molecular systems. Therefore, the non-bonded (above all, the long-range electrostatic) interactions represent a good target for optimizations.
- Solvation plays a crucial role in determining the structure and function of biomolecules. However, the necessary amount of water is often extremely large, becoming the main source of computational cost.¹⁷ Therefore, there is a need to efficiently describe the solvation effects, which are of a predominantly electrostatic character (due to the large dipole moment of the water molecule).

A. Introduction to electrostatic interaction

The electrostatic interaction energy of two point charges q and Q separated by a distance r is given by Coulomb's law

$$E^{\text{el}} = \frac{1}{4\pi\epsilon_0} \cdot \frac{q \cdot Q}{r} \quad (\text{V.1})$$

Of importance is the concept of *electrostatic potential* (ESP), induced at the point \vec{r} by a point charge Q located at \vec{r}_1 :

$$\Phi(\vec{r}) = \frac{1}{4\pi\epsilon_0} \cdot \frac{Q}{|\vec{r} - \vec{r}_1|} \quad (\text{V.2})$$

¹⁷ Typically, the simulated molecular system consists from more than 90 % of water, so that more than 80 % of the computational time is spent by calculating the forces among the (uninteresting) water molecules around the (interesting) solute.

If we know the electrostatic potential at a point \vec{r} in space, then we can obtain the total electrostatic energy of a charge q at this point:

$$E^{\text{el}}(\vec{r}) = \Phi(\vec{r}) \cdot q \quad (\text{V.3})$$

In this way, we can have an ‘electrostatic potential energy surface’ in analogy to mechanics. There, if we know the topography of the Alps, then we immediately know the potential energy of a person with a weight of 70 kg, at any point. In a similar way, if we know the electrostatic potential induced by the atoms of a protein, then we can readily obtain for instance the binding energy of a point charge (like a metal cation) at any place. The electrostatic potential induced by a number of point charges Q_i follows simply as a sum

$$\Phi(\vec{r}) = \frac{1}{4\pi\epsilon_0} \sum_i \frac{Q_i}{|\vec{r} - \vec{r}_i|} \quad (\text{V.4})$$

with the energy of a point charge q at \vec{r} given by Eq. V.3.

In case of a continuous charge distribution, we have to consider the charge density $\rho = Q/V$, with $\rho(\vec{r})$ being the charge density at the point \vec{r} . Then, $Q_i = \rho(\vec{r}_i) \cdot V_i = \rho(\vec{r}_i) \cdot \Delta V$ is the charge in the i -th volume element V_i . Summing over all elements, one obtains

$$\Phi(\vec{r}) = \frac{1}{4\pi\epsilon_0} \sum_i \frac{\rho(\vec{r}_i) \cdot \Delta V}{|\vec{r} - \vec{r}_i|} \quad (\text{V.5})$$

If we make the volume elements infinitesimally small, this changes to (with $d^3\vec{r} = dV$)

$$\Phi(\vec{r}) = \frac{1}{4\pi\epsilon_0} \int \frac{\rho(\vec{r}_1)}{|\vec{r} - \vec{r}_1|} d^3\vec{r}_1 \quad (\text{V.6})$$

Finally, the electrostatic energy of a charge distribution $\rho(\vec{r})$ follows as

$$E = \frac{1}{2} \int \Phi(\vec{r}) \cdot \rho(\vec{r}) dV = \frac{1}{8\pi\epsilon_0} \iint \frac{\rho(\vec{r}_1) \cdot \rho(\vec{r})}{|\vec{r} - \vec{r}_1|} d^3\vec{r} d^3\vec{r}_1 \quad (\text{V.7})$$

The main task is to get the electrostatic potential from a charge distribution. For that, one has to solve **Poisson’s equation**

$$\nabla^2 \Phi(\vec{r}) = -\frac{\rho(\vec{r})}{\epsilon} \quad (\text{V.8})$$

(differential equation for Φ as a function of \vec{r}), or, if the permittivity ϵ is not constant,

$$\nabla (\epsilon \nabla \Phi(\vec{r})) = -\rho(\vec{r}) \quad (\text{V.9})$$

As an example let us have a look at the ESP of a gaussian charge distribution. This distribution centered around the origin of coordinate system with a width σ is given as

$$\rho(r) = Q \cdot \frac{1}{\sigma^3 \sqrt{2\pi}^3} \cdot \exp \left[-\frac{r^2}{2\sigma^2} \right] \quad (\text{V.10})$$

The corresponding solution of the Poisson equation is

$$\Phi(r) = \frac{1}{4\pi\epsilon} \cdot \frac{Q}{r} \cdot \text{erf} \left[\frac{r}{\sqrt{2}\sigma} \right] \quad (\text{V.11})$$

with erf being the error function. Here, if we move far enough from the center of the charge distribution (r is large), the error function converges to unity and the ESP is very near to that of a point charge placed in the origin (Eq. V.2). This is in accord with experience – a point charge and a well-localized charge distribution interact with distant charges in the same way. Actually, we need not go so far in order to see that – the error function assumes a value of 0.999 already at the distance of 2.4σ .

B. Periodic boundary conditions

The most frequent objective of MD simulations is to describe a molecular system in aqueous solution. The problem that we readily encounter is that we have to make the system as small as possible, in order to reduce the computational cost. The most straightforward way to do so is to consider only a single molecule of the solute (e.g. a protein or DNA species) with a smallest possible number of solvent (water) molecules. A typical size of such a system with several thousand water molecules is in the range of units of nanometer. Here, a serious issue occurs: while we are trying to describe the behavior of a molecule in *bulk solvent*, every point in such a small system is actually very close to the *surface*. The surface layer of a system has always properties very different from those of the bulk phase, and with such a setup, we would simulate something else than what we aim at.

An elegant way to avoid this problem is to implement the **periodic boundary conditions** (PBC). Here, the molecular system is placed in a box with a regular geometrical shape (the possibilities are listed below). This polyhedron is virtually replicated in all spatial directions, with *identical* positions (and velocities) of all particles, as shown in Fig. 20. This way, the system is made *infinite* – there is no surface in the system at all. The atoms in the vicinity of the wall of the *simulation cell* (like the black circle in Fig. 20) interact with the atoms in the neighboring replica.




FIG. 20: Replication of the unit cell (grey) using periodic boundary conditions. Interactions of an atom (black) with the nearest images of all other atoms (red).

This method is not quite perfect as it introduces artificial periodicity in the system – all the replicas of the simulation cell look the same, making the thermodynamics of the system incorrect in principle.¹⁸ However, this treatment is much better than simulating a too small system with artificial boundary with vacuum.

Practically, the implementation has the following features:

- Only the coordinates of the unit cell are recorded.
- If a particle leaves the box, then it enters the box from the other side.
- Carefull accounting of the interaction of the particles is necessary. The simplest approach is to make an atom interact only with the $N - 1$ particles within the closest periodic image, i.e. with the nearest copy of every other particle (**minimum image convention**). This is to avoid the interaction of an atom with two different images of another atom, or even with another image of itself. If the box is cubic with boxsize L , then each atom can interact only with all atoms closer than $L/2$. Evidently, PBC have to be synchronized with the applied cut-offs, see below.

The unit cell may have a simple shape – cubic or orthorhombic, parallelepiped (specially, rhomboeder), or hexagonal prism; but also a more complicated like truncated octahedral

¹⁸ For instance, the *entropy* of the entire system is obviously too small, because of its (wrong) translational symmetry. As a general rule, this is rarely a problem.

or rhombic dodecahedral. In the latter two cases, the corresponding PBC equations are quite complicated; the advantage of such shapes for the simulation of spherical objects (like globular proteins in solvent) is that there are no voluminous distant corners which increase the amount of solvent and thus the computational complexity (like in the case of cubic /orthorhombic box). Two-dimensional objects like phase interfaces are usually treated in a *slab* geometry.




FIG. 21: Truncated octahedron (left) and rhombic dodecahedron (right).

C. Accelerating the calculation of non-bonded interactions – cut-off

As mentioned above, the evaluation of non-bonded terms becomes a bottleneck for large molecular systems, and in order to make simulations of extended systems possible, it is necessary to limit their computational cost.

The simplest and crudest approach is to neglect the interaction of atoms that are further apart than a certain distance r_c . This so-called *cut-off* is commonly used with the rapidly decaying ($1/r^6$) Lennard-Jones interaction, which indeed nearly vanish already for moderate distances r_c , commonly around 10 Å. However, with the slowly decaying electrostatic interaction ($1/r$), this would lead to a sudden jump (discontinuity) in the potential energy; even worse, this would be a disaster for the forces, which would become infinite at that point.

A better idea would be to *shift* the whole function by $V(r_c)$, so that there is no jump at r_c anymore. We would have

$$V^{\text{sh}}(r) = \begin{cases} V(r) - V(r_c), & \text{for } r \leq r_c, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{V.12})$$

However, the gradients (forces) are at r_c still not continuous! To eliminate this force jump,

it is possible to apply a *shift-force* potential ($V' \equiv dV/dr$):

$$V^{\text{sf}}(r) = \begin{cases} V(r) - V(r_c) - V'(r_c) \cdot (r - r_c), & \text{for } r \leq r_c, \\ 0, & \text{otherwise.} \end{cases} \quad (\text{V.13})$$

The obvious drawback of this method is that the Coulomb energy is changed!




FIG. 22: Electrostatic interaction energy of two unit positive charges, evaluated using Coulomb's law and the various modifications.

A further alternative is the *switch potential*. Here, an additional potential is applied starting from a certain distance r_1 , which brings the Coulomb interaction gradually to zero, as shown in Fig. 22. The drawback of this method is, that the forces are altered in the cut-off region.

Both methods can be applied to either the energy or the forces: when applied to the energy, the forces follow through differentiation, and vice versa, when applied to forces, the energy follows through integration.

Generally, the cut-off schemes can be based either only on atomic distances, or on functional groups. Usually, the latter is employed to assure charge conservation in the Coulomb interaction.

D. Accelerating the calculation of electrostatic interactions – Ewald summation

In many cases, above all if highly charged molecular systems (like DNA or some proteins) are simulated, the use of cut-offs is a bad approximation. For instance, the artificial forces if using a switching function may lead to the accumulation of ions in the regions of solution in the cut-off distance (measured from DNA). Therefore, it is desirable to abandon the minimum image convention and the cut-offs, and rather sum up the long-range Coulomb interaction between *all* the replicas of the simulation cell

Let us introduce a vector \vec{n} , which runs over all the replicas of the cell, denoting them uniquely:

- For $|\vec{n}| = 0$, we have $\vec{n} = (0, 0, 0)$ – the central unit cell.
- For $|\vec{n}| = L$, we have $\vec{n} = (0, 0, \pm L)$, $\vec{n} = (0, \pm L, 0)$, $\vec{n} = (\pm L, 0, 0)$ – the six neighboring unit cells.
- Further, we continue with $|\vec{n}| = \sqrt{2} \cdot L$ and the 12 cells neighboring over an edge, etc.

With this vector, we can write the sum of Coulomb interactions over all replicas as

$$E^{\text{Coul}} = \frac{1}{2} \sum_{i,j} \sum_{\text{images } \vec{n}} \frac{q_i \cdot q_j}{|\vec{r}_{ij} + \vec{n}|} \quad (\text{V.14})$$

for indices i and j running over all atoms in the unit cell (r_{ij} is then their distance). This expression is an infinite sum which has special convergence problems. Such a sum decays like $1/|\vec{n}|$ and is a *conditionally convergent* series, meaning that it converges ($\sum_{i=1}^{\infty} a_i < \infty$) but does not converge absolutely ($\sum_{i=1}^{\infty} |a_i|$ cannot be summed up). The problem is that the convergence of such a sum is slow and, even worse, dependent on the order of summation. So, a conditionally convergent series may add up to *any* (!) value, as shown in this example:

$$\begin{aligned} \text{I: } S &= 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \frac{1}{7} - \frac{1}{8} + \dots \\ \text{II: } \frac{1}{2}S &= \frac{1}{2} - \frac{1}{4} + \frac{1}{6} - \frac{1}{8} + \dots \\ \text{I + II: } \frac{3}{2}S &= 1 + \frac{1}{3} - \frac{1}{2} + \frac{1}{5} + \frac{1}{7} - \frac{1}{4} + \frac{1}{9} + \frac{1}{11} - \frac{1}{6} + \dots \\ &= 1 - \frac{1}{2} + \frac{1}{3} - \frac{1}{4} + \frac{1}{5} - \frac{1}{6} + \dots = S \quad (\text{sic!}) \end{aligned} \quad (\text{V.15})$$

Therefore, we need a clever way to evaluate the potential resulting from the interaction of all images of all charges

$$\Phi(\vec{r}_i) = \sum_j \sum_{\text{images } |\vec{n}|} \frac{q_j}{|\vec{r}_{ij} + \vec{n}|} \quad (\text{V.16})$$

in order to evaluate the Coulomb energy of the charges q_i in the unit cell.

$$E^{\text{Coul}} = \frac{1}{2} \sum_i q_i \cdot \Phi(\vec{r}_i) \quad (\text{V.17})$$

The idea of the Ewald methods is to convert the difficult, slowly convergent series to the sum of two series, which both converge much more rapidly, like

$$\sum \frac{1}{r} = \sum \frac{f(r)}{r} + \sum \frac{1 - f(r)}{r} \quad (\text{V.18})$$

where $\sum 1/r$ represents the difficult series that we have to deal with. Whereas the terms on the right-hand side look more complicated, they actually exhibit a much more rapid convergence than $\sum 1/r$ in our case, and such an awkwardly looking ‘decomposition’ is the preferred way to go.

Since the summing over point charges leads to convergence problems with conditionally convergent sums, the Ewald method uses rather *normal distributions* of charge of the same magnitude:

$$q_j \rightarrow q_j \cdot \left(\frac{\alpha}{\sqrt{\pi}} \right)^3 \exp[-\alpha^2 \cdot |\vec{r}_j|^2] \quad (\text{V.19})$$

To get the electrostatic potential induced by this smeared charge distribution, Poisson’s equation (Eq. V.8) has to be solved. This leads to the potential being represented by the so-called *error function*:¹⁹

$$\Phi(\vec{r}) = q_j \cdot \frac{\text{erf}[\alpha \cdot r]}{r} \quad (\text{V.20})$$

¹⁹ The error function is defined as the definite integral of the normal distribution

$$\text{erf}[x] = \frac{2}{\sqrt{\pi}} \int_0^x \exp[-t^2] dt$$

and the complementary error function as

$$\text{erfc}[x] = 1 - \text{erf}[x]$$

and, in the special case of $\vec{r} = \vec{o}$:

$$\Phi(\vec{o}) = q_j \cdot \frac{2\alpha}{\sqrt{\pi}} \quad (\text{V.21})$$

If we sum up the potentials given by Eq. V.20 for all charges, we obtain

$$\Phi(\vec{r}_i) = \sum_j \sum_{\text{images } |\vec{n}|} q_j \cdot \frac{\text{erf}[\alpha \cdot |\vec{r}_{ij} + \vec{n}|]}{|\vec{r}_{ij} + \vec{n}|} \quad (\text{V.22})$$

This has to be compared with the potential induced by the point charges (Eq. V.16). The difference between Eq. V.16 and Eq. V.22 is given by the complementary error function.

The genuine potential induced by the point charges can then be expressed as

$$\Phi(\vec{r}_i) = \sum_j \sum_{\text{images } |\vec{n}|} q_j \cdot \frac{\text{erfc}[\alpha \cdot |\vec{r}_{ij} + \vec{n}|]}{|\vec{r}_{ij} + \vec{n}|} \quad (\text{V.23})$$

$$+ \sum_j \sum_{\text{images } |\vec{n}|} q_j \cdot \frac{\text{erf}[\alpha \cdot |\vec{r}_{ij} + \vec{n}|]}{|\vec{r}_{ij} + \vec{n}|} \quad (\text{V.24})$$

The first term V.23 called the *real-space contribution* is shown graphically in Fig. 23 (top). From a certain (quite small) distance, the point charges and the gaussian charge distributions (with opposite signs) cancel each other. This distance depends on the gaussian width α – a small gaussian width would lead to a rapid convergence.




FIG. 23: Top: Real-space contribution to the Ewald sum consists of the original point charges (red) and gaussian charge distributions (blue) of the same magnitude but opposite sign. Bottom: Reciprocal-space contribution.

On the other hand, the second term V.24 called the *reciprocal-space contribution* is best evaluated in the form (\vec{k} – the reciprocal lattice vector of periodic images)

$$E^{\text{rec}} = \frac{1}{2V\varepsilon_0} \cdot \sum_{\vec{k} \neq \vec{o}} \frac{1}{k^2} \cdot \exp\left[-\frac{|\vec{k}|^2}{4\alpha^2}\right] \cdot \left| \sum_j q_j \cdot \exp[-i \cdot \vec{k} \cdot \vec{r}_j] \right|^2 \quad (\text{V.25})$$

The usually applied *Fourier transform* techniques²⁰ need a large gaussian width α for fast convergence, therefore the value of α is a necessary compromise between the requirements for the real- and reciprocal-space calculations. All in all, both mentioned contributions exhibit quite favorable convergence behavior, making the evaluation of the electrostatic potential due to all periodic images feasible.

After calculating these two terms (Fig. 23), yet another one has to be taken into account: Since we have broadened charge distributions, they do interact with themselves, as shown in Fig. 24. This interaction has been brought about by the substitution of point charges by gaussian charge distributions, and thus it must be subtracted from the final result.




FIG. 24: Interaction of the charge with the gaussian distribution

The potential of a broadened gaussian is given by Eq. V.21, which leads to Coulomb energy of

$$E^{\text{self}} = \sum_j q_j \cdot \Phi(\vec{o}) = \sum_j q_j \cdot q_j \cdot \frac{2\alpha}{\sqrt{\pi}} \quad (\text{V.26})$$

At the end, we have three energy contributions: one from the ‘real-space’ evaluation of Φ^{real} in Eq. V.23, which gives

$$E^{\text{real}} = \frac{1}{2} \sum_j q_j \cdot \Phi^{\text{real}}(\vec{r}_j) \quad (\text{V.27})$$

one from the ‘reciprocal-space’ evaluation of Φ^{rec} in Eq. V.25 and the ‘self-energy’ in Eq. V.26, so that

$$E^{\text{Ewald}} = E^{\text{real}} + E^{\text{rec}} - E^{\text{self}} \quad (\text{V.28})$$

²⁰ A popular implementation is the FFTW (Fastest Fourier Transform in the West), with a favorable computational cost scaling as $\mathcal{O}(N \cdot \ln N)$.

E. Explicit solvent models – water

The most simulations of biomolecules are performed with water as the solvent, to mimic the physiological or *in vitro* experimental conditions. If a not too concentrated solution is to be simulated, then the necessary amount of the solvent is quite large, often many thousand molecules.

For instance, in a typical simulation of a DNA oligomer with ten base pairs (see Fig. 25), the dimensions of the PBC box are $3.9 \times 4.1 \times 5.6$ nm, and there are 630 atoms in the DNA molecules, 8346 atoms of water and 18 sodium counterions. The macroscopic concentration of DNA in this ‘sample’ reaches an astonishingly large value of 18 mmol/L!²¹ At the same time, 86 % of all *pair interactions* are those where each of the partner atoms belongs to a water molecule,²² and the most remaining interactions involve one atom of water. This is a huge portion, and the smallest possible at the same time, as we have the minimal number of water molecules.





FIG. 25: Typical setup of the simulation of a DNA oligomer.

We can see that the most interactions involve water, and that is why it is necessary to turn our attention to the description of water in the simulations. The model of water must be made simple enough in order to reduce the computational complexity, but at the same time it is necessary not to compromise the accuracy of the description.

²¹ Due to the commonly accepted criteria, such a box is the smallest possible. Thus, the amount of water is also the smallest possible, and the concentration the highest possible.

²² There are 8346 water atoms, that is roughly 8346^2 interactions water–water, and 8994 atoms altogether, corresponding to 8994^2 pair interactions. The ratio of these figures is 0.86.

Many simple water models have been developed so far. They are usually *rigid*, so that the bond lengths as well as angles remain constant during the simulation. A molecule is composed of at least three sites (corresponding to atoms in this case), but possibly also as many as six sites – three atoms and optional dummy particles corresponding to a ‘center’ of electron density, or to the lone electron pairs on the oxygen atom.



The most frequently used atomic model of water is the TIP3P (very similar is the SPC). A TIP3P molecule consists of three atoms connected by three rigid bonds. A charge is placed on every atom (-0.834 on the O and $+0.417$ on the Hs), while *only* the oxygen atom possesses non-zero Lennard-Jones parameters.²³

If the negative charge is placed on a dummy atom M rather than on the oxygen, then the electric field around the molecule is described better. This idea is implemented e.g. in the TIP4P model.

A further improvement may be achieved if two dummy particles L bearing negative charge are placed near the oxygen atom, to mimic the lone electron pairs. Consequently, such a five-site model (like TIP5P) describes the directionality of hydrogen bonding and derived effects (radial distribution function, temperature of highest density) better than less sophisticated models.

²³ This makes it possible to additionally optimize the algorithm for the calculation of energy and forces.

VI. PRINCIPLES OF STATISTICAL MECHANICS

In the last chapters, we often used the notion that a particular ensemble generated by a simulation does not represent the canonical ensemble. This means that the phase space is sampled differently in the simulation than what would correspond to the canonical ensemble. But what does the **canonical probability distribution function** look like? Let us first have a look to the microcanonical ensemble.

Interestingly, the statistics is simpler for quantum mechanics (QM) than for classical mechanics (CM). The reason is that in QM, we can talk about discrete states, while everything in CM is continuous. In CM, have to select small volumes of phase space and calculate the density there. On the other hand, in QM, we can use for instance the discrete energy states of molecules or quantum oscillators as simple examples, as we will do in the following.

A. Microcanonical ensemble – microstates for constant energy

If the energy of the system is the same for every *microstate*, then we assume that every microstate of the system occurs with the same probability. With microstate we understand the particular distribution of energy among the particles in the system. Another important concept is that of the *configuration* or macrostate, which is defined by the occupations of energy level by the individual indistinguishable particles — see Fig. 26 for an explanation.




FIG. 26: A system of three particles with two available energy levels. In one of the possible *configurations* (left), two indistinguishable particles (grey) occupy the level E1 while the remaining one sits on the level E2. This configuration is composed of three *microstates* (right) which differ in the exact distribution of particles (with various symbols) among the energy levels.

The important point is that since all microstates have *equal* probabilities, the probability of a configuration is given by the *number of microstates* that correspond to the configuration. Thus, if we wish to know the probability of a configuration, we have to count the microstates

that compose the configuration.

1. *Example – counting the microstates*

Consider a system of three particles that possess three identical energy quanta altogether. These three energy quanta can be distributed as shown in Fig. 27. We find ten **microstates** and three classes of systems, which we shall call **configurations** A, B and C:

Now, we wish to count, in *how many ways* we can distribute the energy quanta, thus *how many microstates* we have. In principle, we will do it in the following way: One particle obtains l quanta, the next m , and the last one n quanta of energy, where $l + m + n = 3$.




FIG. 27: Possible microstates with three energy quanta in the system.

There are three possibilities to divide three energy quanta between three particles: $(3,0,0)$, $(2,1,0)$ and $(1,1,1)$ — these are the configurations A, B and C, respectively. For every configuration, we have now to assign the quanta specifically to the individual particles. This runs as follows:

- We start by pick a first particle and giving it the largest number of energy quanta in the configuration. We have *3 choices*.
- Then, we assign the second-largest number of quanta to another particle. There are *2 choices*.
- There is only *1 choice* particle left to accomodate the smallest number of quanta.

- This makes $3 \cdot 2 \cdot 1 = 3! = 6$ choices in total, for the given configuration.

However, there are some *degeneracies*, so that we do not always obtain six different microstates:

Config A: Let us give particle 1 three quanta. Then, it does not matter if we choose particle 2 first to obtain zero quanta and particle 3 next, or particle 3 obtains zero quanta first and particle 2 follows. The result is the same as seen in Fig. 27. There are two particles that obtain the same number of quanta, and two assignments are thus indistinguishable, leading to an identical result. However, we have counted both ways, and to obtain the right number of microstates we have to divide by $2 \cdot 1$ (two times one redundant assignment). The number of microstates is thus

$$\frac{3!}{2!} = 3$$

Config B: All three particles are given a different number of quanta, so that no degeneracy actually occurs. The number of possibilities to distribute the quanta is thus truly

$$3! = 6$$

Config C: Trivially, there is only one microstate for this configuration. In detail, we have to assign the particles three identical numbers of energy quanta, and so we obtain the number of microstates by dividing by $3!$:

$$\frac{3!}{3!} = 1$$

Thus, there are ten microstates for our system; these have been already presented in Fig. 27.

2. Number of microstates for a configuration; the dominating configuration

Generally, if there are N particles, then there are N ways to pick the first, $N - 1$ ways to pick the second etc. Thus, we have $N!$ ways to build up the system. But then, if n_a particles have to accomodate the same number of energy quanta, then the real number of different microstates is obtained by dividing by $n_a!$. Then, we find the number of microstates W for a system of N particles in a certain configuration as

$$W = \frac{N!}{n_a! \cdot n_b! \cdot \dots} \quad (\text{VI.1})$$

with n_i being the numbers of particles accommodating the same number of energy quanta, which we shall also call *occupation numbers* of the energy levels.

For a large number of particles N , W follows as an extensively large number, and thus it is practical to consider the logarithm of it:

$$\ln W = \ln \frac{N!}{n_a! \cdot n_b! \cdot \dots} = \ln[N!] - \ln[n_a!] - \ln[n_b!] - \dots = \ln[N!] - \sum_i \ln[n_i!] \quad (\text{VI.2})$$

and using Stirling's approximation $\ln a! = a \cdot \ln a - a$ we find

$$\ln W = N \cdot \ln N - \sum_i n_i \cdot \ln n_i \quad (\text{VI.3})$$

We may introduce the *fraction* of particles in state i , or the probability of a particle to be found in state i , as

$$p_i = \frac{n_i}{N} \quad (\text{VI.4})$$

Then, it follows that

$$\ln W = \sum_i n_i \cdot \ln N - \sum_i n_i \cdot \ln n_i = - \sum_i n_i \cdot \ln \frac{n_i}{N} = -N \cdot \sum_i p_i \cdot \ln p_i \quad (\text{VI.5})$$

Now, we can calculate the number of microstates corresponding to a given configuration. The crucial point is now to find out which configuration has the largest *weight*, in other words, is the most likely. This will be the configuration with the largest number of corresponding microstates — the task is thus to find the maximum of W as the function of occupation numbers n_i .

Method of Lagrange multipliers:

We want to find the maximum of a function $f(\vec{x})$ under the condition that certain constraints $y_k(\vec{x}) = 0$ are fulfilled. Then, we can search for the maximum by taking these constraints into account in the following way:

$$\begin{aligned} \frac{\partial}{\partial x_i} \left(f(\vec{x}) - \sum_k \lambda_k \cdot y_k(\vec{x}) \right) &= 0 \\ \frac{\partial}{\partial \lambda_i} \left(f(\vec{x}) - \sum_k \lambda_k \cdot y_k(\vec{x}) \right) &= 0 \end{aligned} \quad (\text{VI.6})$$

with x_i running over the components of \vec{x} . So, we add the constraints to the function and set all the derivatives to zero, as usually.

This may be applied to the microcanonical ensemble: We are looking for the configuration with maximum weight of a system of N particles distributed among energy levels ε_i , with the total energy E . Subject to maximization is then the weight

$$\ln W = N \cdot \ln N - \sum_i n_i \cdot \ln n_i \quad (\text{VI.7})$$

under the normalization conditions

$$\sum_i n_i - N = 0 \quad (\text{VI.8})$$

$$\sum_i n_i \cdot \varepsilon_i - E = 0 \quad (\text{VI.9})$$

which express the constant number of particles and the constant energy.

The application of Lagrange's method leads to the equations

$$\frac{\partial}{\partial n_i} \left[\ln W + \alpha \cdot \left(\sum_j n_j - N \right) - \beta \cdot \left(\sum_j n_j \cdot \varepsilon_j - E \right) \right] = 0 \quad (\text{VI.10})$$

$$\frac{\partial \ln W}{\partial n_i} + \alpha - \beta \cdot \varepsilon_i = 0 \quad (\text{VI.11})$$

once we used $-\alpha$ and β for the Lagrange multipliers. This leads to the solution

$$\frac{n_i}{N} = \exp[\alpha - \beta \cdot \varepsilon_i] \quad (\text{VI.12})$$

where the parameter α may be obtained from the condition VI.8 as

$$\exp \alpha = \frac{1}{\sum_j \exp[-\beta \cdot \varepsilon_j]} \quad (\text{VI.13})$$

so that

$$\frac{n_i}{N} = \frac{\exp[-\beta \cdot \varepsilon_i]}{\sum_j \exp[-\beta \cdot \varepsilon_j]} \quad (\text{VI.14})$$

In the microcanonical ensemble, the parameter β might be estimated from the condition VI.9 in the form

$$\frac{E}{N} = \frac{\sum_j \varepsilon_j \cdot \exp[-\beta \cdot \varepsilon_j]}{\sum_j \exp[-\beta \cdot \varepsilon_j]} \quad (\text{VI.15})$$

An important observation is that if the number of particles N is huge, there is always one configuration with weight much larger than that of all other configurations. This *dominating configuration* then actually determines the properties of the system. The occupation numbers n_i obtained above correspond to such a dominating configuration.

B. Microscopic definition of entropy

The configurations $p_i = \frac{1}{N}, \frac{1}{N}, \frac{1}{N}, \dots$ and $p_i = 1, 0, 0, \dots$ are the extreme cases. In the former case, $\ln W$ is maximal (for large N)

$$\ln W = N \cdot \ln N$$

whereas in the latter case it is minimal

$$\ln W = \ln 1 = 0$$

We define the **microscopic entropy** as

$$S = -k_B \cdot \ln W \tag{VI.16}$$

where the universal constant k_B is the so-called Boltzmann constant. This property tells us something about the travel of the system through the configuration space. If entropy is small, few states are occupied; if it is large, then many states are occupied. We may see the requirement of the maximal number of microstates corresponding to the optimal configuration in the microcanonical ensemble as the requirement of *maximal entropy*.

Entropy can be related to the **order** in the system. If entropy is small, only a small part of the configuration space is accessible, and we consider the system to be ordered. On the other hand, if entropy is large, an extended part of the configuration space is covered. Think of the books on your desk. When they are all stapled on one place on your desk, you consider this state to be ordered, while when they are freely distributed over the entire room, you would probably call this state disordered.²⁴

There is also another route to the entropy, via the so-called **information entropy**. If entropy is minimal, we have perfect knowledge about a system — it is in the state i because $p_i = 1$. If entropy reaches its maximum, we have no information at all, as every possibility is equally likely. Think again of the books in your room: if $S = 0$ then $p_1 = 1$ and you know that a particular book is on your desk. If $S = k_B \cdot \ln N$, you know that the book is in one of the N possible places in your room. In this case, you have no idea where to look first!

²⁴ Jan Černý (Charles University in Prague, Department of Cellular Biology) once coined the term *anthropy* for “entropy of human origin” as is the case of the books on your desk.

Second law of thermodynamics

The second law claims that entropy increases in every irreversible process:

$$\frac{\partial}{\partial t} S \geq 0.$$

For a closed system, the microscopic entropy does not have this property. This can be most easily seen in a quantum mechanical example:

Consider particles in a box with a wavefunction Ψ , for which we solve the Schrödinger equation

$$-i\hbar\dot{\Psi} = \hat{H}\Psi$$

We can expand this in the set of eigenfunctions of the box ϕ_i

$$\Psi = \sum_i c_i \phi_i$$

to get

$$\begin{aligned} -i\hbar\dot{\Psi} &= \hat{H}\Psi \\ -i\hbar \cdot c_i \dot{\phi}_i &= \sum_j c_j H_{ij} \phi_j \end{aligned}$$

We see that the probability to find the particle in state ϕ_i

$$p_i = |c_i|^2$$

does not change during the dynamics:

$$\dot{p}_i = 0$$

This leads to the entropy change of

$$\frac{\partial}{\partial t} S = k_B \cdot \frac{\partial}{\partial t} \left(-k_B \sum_i p_i \cdot \ln p_i \right) = 0$$

The microscopic equations do not lead to any change in entropy, in contradiction with the second law. Yet, we can understand that this has to be so: We know the initial conditions of a system p_i . Then, since our equations of motion are perfectly deterministic, we know how every trajectory evolves — we know p_i at every instant! This way, the change of information entropy is zero!

In the example of your books on the desk. Initially, they are all on the desk. Since you know all the equations of motion of every book, you know exactly where every book is at any later time T . So, you have perfect knowledge of

C. Canonical ensemble – system in equilibrium with the environment

For the microcanonical ensemble, we optimized $\ln W$, and so looked for the configuration of probabilities p_i such that gave the largest number of states, which maximized the entropy. In our example, W was the number of microstates for a certain configuration of particles over the energy states ε_i in the system $(\varepsilon_1, \varepsilon_2, \dots)$. We saw that for the occupation numbers (2,1,0) in the conformation B we got the largest number of states. This was the dominating configuration in the sense that the most microstates belonged to it. Therefore, if we wanted to know the configuration with the largest number of microstates, we had to vary the n_i in order to maximize $\ln W$. And this was the most probable configuration.

Now, we shall consider the system to be in contact with the environment. Under such circumstances, it is the temperature rather than energy that remains constant. The famous **Boltzmann distribution** of p_i

$$p_i = \frac{\exp[-\beta \cdot \varepsilon_i]}{\sum_j \exp[-\beta \cdot \varepsilon_j]} \quad (\text{VI.17})$$

is valid, with the denominator

$$Q = \sum_j \exp[-\beta \cdot \varepsilon_j] \quad (\text{VI.18})$$

being designated as the *canonical partition function* (Zustandssumme).

To derive what β means, we have to remember some basic thermodynamics: The energy as the basic thermodynamic potential depends on the extensive variables entropy, volume and number of particles:

$$E = E(S, V, N) = TS - pV - \mu N \quad (\text{VI.19})$$

Therefore, the *thermodynamic temperature* comes into play by means of

$$\frac{\partial E}{\partial S} = T \quad \text{or} \quad \frac{\partial S}{\partial E} = \frac{1}{T} \quad (\text{VI.20})$$

Now, we can use the microscopic definition of S to calculate

$$\frac{1}{T} = \frac{\partial S}{\partial E} = \frac{\partial S}{\partial \beta} \cdot \frac{\partial \beta}{\partial E} = -k_B \cdot \sum_i \frac{\partial p_i}{\partial \beta} (-\beta \varepsilon_i - \ln Q) \cdot \left(\sum_i \frac{\partial p_i}{\partial \beta} \varepsilon_i \right)^{-1} = k_B \cdot \beta \quad (\text{VI.21})$$

when using

$$\sum_i \frac{\partial p_i}{\partial \beta} = \frac{\partial}{\partial \beta} \sum_i p_i = \frac{\partial}{\partial \beta} 1 = 0$$

Finally, we have

$$\beta = \frac{1}{k_B T} \quad (\text{VI.22})$$

For continuous systems, the spacing between the energy levels becomes infinitesimally small and we can write

$$\rho(\vec{r}, \vec{p}) = \frac{1}{Q} \cdot \exp \left[-\frac{E(\vec{r}, \vec{p})}{k_B T} \right] \quad (\text{VI.23})$$

with the canonical partition function

$$Q = \int \exp \left[-\frac{E(\vec{r}, \vec{p})}{k_B T} \right] d\vec{r} d\vec{p} \quad (\text{VI.24})$$

D. Canonical partition function – the way to the thermodynamic quantities

The partition function Q seems to be a purely abstract quantity, but the very opposite is true! In order to characterize the thermodynamics of a system, we need just to evaluate Q and the desired thermodynamics observable follow as functions of Q , in principle. This illustrates the purpose for which the statistical thermodynamics has been developed: it makes us able to derive the (macroscopic) thermodynamic properties of a system from the knowledge of (microscopic) properties of the molecules that compose the system, with the partition functions connecting the microscopic and the macroscopic.

As an example, we may calculate the mean total energy of a system from the dependence of canonical partition function Q on the parameter β or the temperature T :

$$\langle E \rangle = -\frac{\partial \ln Q}{\partial \beta} = k_B T^2 \cdot \frac{\partial \ln Q}{\partial T} \quad (\text{VI.25})$$

E. Exercises

1. You have 20 (identical) cookies and six boxes (for six of your friends, which makes the boxes distinguishable). How many possibilities do you have to distribute them as follows: $\{1,0,3,5,10,1\}$ (this is the distribution of cookies in the 6 boxes). Use Eq. VI.1.
2. Prove Eq. VI.21.
3. Calculate the average energy of the classical harmonic oscillator using the Boltzmann distribution.

VII. THERMODYNAMIC PROPERTIES OF MOLECULAR SYSTEMS

We got to know the principles of statistical mechanics and its significance as the way from the properties of particles to the thermodynamic properties of ensembles. The role of the mediator of information is played by the *partition function*.

In this chapter, we will see

- how the thermodynamic equilibrium is characterized, which quantities are of interest and how these may be derived from the partition function,
- how the partition function is connected to the phase-space density,
- how the ensemble partition function may be derived from the partition function of a single molecule,
- that MD simulation provides an alternative way to thermodynamic quantities,
- that it is difficult to obtain free energies from normal simulations.

A. Driving forces of thermal processes. Equilibrium

Classical thermodynamics introduces the concepts of *thermodynamic equilibrium* and *spontaneous process*, and identifies the quantities that are maximized/minimized in the equilibrium and show a definite change in the course of a spontaneous process.

So if the system does not exchange energy with the surroundings, as is the case in the microcanonical ensemble, the equilibrium is reached if *entropy* S is maximized. Consequently, a process proceeds spontaneously if it introduces an increase of entropy ($\Delta S > 0$).

The case of microcanonical ensemble is simple as there is no exchange of energy with the surroundings, but the situation is more complex if we pass to the canonical ensemble. Here, it is necessary to consider a super-system composed of the system of interest *together* with the surroundings, once we wish to identify the equilibrium and the spontaneity of processes. Of course, it is near-to-impossible to estimate the entropy of such a super-system, as we cannot handle the whole universe in a calculation. Thus, an alternative criterion has to be sought.

A way to concentrate on our molecular system and to be able to omit the surroundings, is to introduce a new thermodynamic function. In case of the canonical ensemble, the key function is the *Helmholtz free energy*

$$F = U - TS \quad (\text{VII.1})$$

The equilibrium state in canonical ensemble exhibits a minimum of Helmholtz energy, and F decreases in the course of a spontaneous process.

This is the fundamental property we are interested in, because

- $F = F(T, V)$: F depends on the variables T and V , which are experimentally controllable, while $U = U(S, V)$ depends on S and V . We have no means to control entropy in experiments. In particular, F is the energetic property which is measured when T and V are constant, a situation we often model in our simulations.
- $\Delta F = F_f - F_i = W^{\max}$ is the *maximum amount of work* that the system can perform between an initial (i) and final (f) state. In the first law $dU = \delta Q + \delta W$, we can cast in a formulation of the second law $TdS \geq \delta Q$ to obtain²⁵

$$dU \leq TdS + \delta W$$

and for the amount of work

$$\delta W \geq W^{\max} = dU - TdS = dF$$

Therefore, the work is always greater or equal the change of free energy.²⁶ In other words, a certain amount of internal energy dU can *never* be converted to work, because a part of it is always lost as an increase of entropy.

- In principle, energy minimization as we have discussed before, does not make any sense. The energy is simply conserved, but within the entire universe (the universe thus samples a microcanonical ensemble). Once we have defined a quantity like *free energy*, the idea of a minimization procedure is restored: Systems will evolve in order to minimize the free energy; however, this is nothing else than the maximization of entropy of the super-system (universe).²⁷

²⁵ This inequality becomes an equality for a *reversible* process.

²⁶ Note that if the system *performs* work, then the value of δW is *negative*. So, the value of W^{\max} is the *most negative possible*, and its magnitude represents truly the maximum possible work.

²⁷ In an NPT ensemble, the same role is played by the *Gibbs free energy* $G = H - TS = U + pV - TS$.

B. Thermodynamic functions from statistical mechanics

Now, we want to apply the formalism to calculate enthalpies and free energies of molecules. The main trick we learned is, that we only have to calculate the **partition function** Q , since from there we get all thermodynamic functions:

$$U = \langle E \rangle = k_B T^2 \frac{\partial \ln Q}{\partial T} \quad (\text{VII.2})$$

$$S = k_B T \cdot \frac{\partial \ln Q}{\partial T} + k_B \cdot \ln Q \quad (\text{VII.3})$$

$$F = -k_B T \cdot \ln Q \quad (\text{VII.4})$$

$$P = k_B T \cdot \left(\frac{\partial \ln Q}{\partial V} \right)_T \quad (\text{equation of state}) \quad (\text{VII.5})$$

$$H = U + pV \quad (\text{VII.6})$$

$$G = F + pV = H - TS \quad (\text{VII.7})$$

Therefore, the computational problem reduces to the need to determine Q .

The problem of evaluation of the partition function may be greatly simplified for two special cases. If the system is composed of n *identical distinguishable particles*, as is the ideal crystal, then the ensemble partition function Q is obtained from the *molecular partition function* q as

$$Q = q^n \quad (\text{VII.8})$$

In the other case, if the n particles composing the system are *identical and indistinguishable*, as in the gas phase, then another relation is valid:

$$Q = \frac{q^n}{n!} \quad (\text{VII.9})$$

Note the immensely reduced need of effort that we have to make in these cases: It is only necessary to evaluate the *molecular partition function* q for *one* molecule of the substance in the studied system, and we obtain directly the thermodynamic quantities of interest, via the ensemble partition function Q .

C. Discrete and continuous systems

Discrete systems

In the chapter about statistical mechanics, we considered a (typically, but not necessarily) quantum-mechanical system, i.e. a system with *discrete energy levels* E_i . For that, we found a partition function

$$Q = \sum_i \exp[-\beta E_i] \quad (\text{VII.10})$$

and a (discrete) Boltzmann distribution function

$$p_i = \frac{1}{Q} \exp[-\beta E_i] \quad (\text{VII.11})$$

which is the probability to find the system in state i with energy E_i .

A prominent example is the harmonic oscillator with energy levels

$$E_i = \left(i + \frac{1}{2}\right) \cdot \hbar\omega \quad (\text{VII.12})$$




FIG. 28: States of a quantum-mechanical harmonic oscillator. The thermal distribution of energy over the vibrational states leads to the number of microstates and a value for the entropy/free energy.

Continuous systems

On the other hand, we discussed the dynamics of molecules, where temperature allows to sample a certain part of the conformational space of a protein. This means, that at a certain temperature only certain values of the coordinates (x_i) and also only certain values of momenta (p_i , cf. the Maxwell distribution) are reached, and thus only a **part of the phase space is sampled**. The phase space density $\rho(\vec{r}, \vec{p})$ gives the probability to find a system at positions \vec{r} and momenta \vec{p} . Now, since we force the dynamics to generate a canonical ensemble, we know the following:




FIG. 29: Complex energy landscape $E(x, p)$. Note that the p -axis is not shown; for $p = 0$, the system ‘lies’ on the potential energy curve (black line); for $p > 0$, various trajectories (sampling various phase-space volumes) are accessible indicated by the blue and red curves.

- Every point in the phase space has potential energy depending on the coordinates \vec{x} ($V = V(\vec{x})$), coming from the force field, and kinetic energy related to \vec{p} . This way, every point (\vec{r}, \vec{p}) in the phase space is assigned a certain energy value.
- From the canonical distribution we know that the probability to find the system in a state with energy E is

$$p(\vec{r}, \vec{p}) = \rho(\vec{r}, \vec{p}) = \frac{1}{Q} \cdot \exp \left[-\frac{E(\vec{r}, \vec{p})}{k_B T} \right] \quad (\text{VII.13})$$

- To obtain the partition function Q , we now have to substitute the summation (in the discrete case) by an integral over the entire phase space:

$$Q = \int \exp \left[-\frac{E(\vec{r}, \vec{p})}{k_B T} \right] d\vec{x} d\vec{p} \quad (\text{VII.14})$$

The canonical distribution function gives the probability of finding the system at a point (\vec{r}, \vec{p}) in the phase space. Typically, the system will be only sampling a part of the phase space, with non-zero probabilities, as shown schematically in Fig. 30.

At the start of an MD simulation, we only know the form of the potential (force field) and kinetic energy. **The fundamental aim of MD simulations is to render the correct phase-space density.** Then, we obtain the thermodynamic potentials U , H etc. as time averages, if the simulation is ergodic. However, this holds only if the phase-space density is truly that of the canonical ensemble (if the simulation sampled the canonical ensemble correctly); otherwise, the derivations in the last chapter would not be valid!




FIG. 30: Part of the phase space, sampled during the dynamics. Right: example of a classical harmonic oscillator.

Thus, we have the following agenda:

- Perform a MD to get a trajectory in phase space. We have to set up the simulation in order to sample the canonical ensemble; for instance, we have to choose the right thermostat! Only then will the distribution of points of the trajectory in the phase space be the same as that given by the canonical distribution.
- Now we can use the ergodic theorem: The time series we get can be used to evaluate time averages of thermodynamic quantities, which will be equal to the ensemble averages of those.

In quantum chemistry, we commonly use both the discrete and the continuous situation, to compute thermal corrections to the total energy E in order to get H , S , F or G . Here, we consider the two situations as follows:

- If we have simple molecules (like benzene) with only one local minimum of energy, we can use the harmonic approximation. The energy is given by the ground state energy of the molecule plus the translational, rotational and vibrational contributions to the energy. From the vibrational Hamiltonian we compute the Boltzmann distribution, from that the partition function and then in turn the thermodynamic potentials.
- For molecules where a multitude of local minima can be reached, a single minimum does not make sense any more. Here, we perform MD simulations to sample the phase space appropriately and evaluate the time averages of the thermodynamic quantities of interest.²⁸

²⁸ These will be equal to the ensemble averages due to the ergodic theorem.

D. Vibrational contribution to the thermodynamic functions

Consider a simple molecule with one or only a small number of well characterizable minima. For one of the minima, we can calculate the electronic energy (here: the force field energy), the translational, rotational and vibrational energy, i.e. we have:

$$E = E^{\text{el}} + E^{\text{trans}} + E^{\text{rot}} + E^{\text{vib}} \quad (\text{VII.15})$$

The partition function follows as

$$\begin{aligned} Q &= \exp \left[-\beta (E^{\text{el}} + E^{\text{trans}} + E^{\text{rot}} + E^{\text{vib}}) \right] = \\ &= \exp[-\beta E^{\text{el}}] \cdot \exp[-\beta E^{\text{trans}}] \cdot \exp[-\beta E^{\text{rot}}] \cdot \exp[-\beta E^{\text{vib}}] = \\ &= Q^{\text{el}} \cdot Q^{\text{trans}} \cdot Q^{\text{rot}} \cdot Q^{\text{vib}} \end{aligned} \quad (\text{VII.16})$$

Since we often need $\ln Q$, we have:

$$\ln Q = \ln Q^{\text{el}} + \ln Q^{\text{trans}} + \ln Q^{\text{rot}} + \ln Q^{\text{vib}} \quad (\text{VII.17})$$

1. Electronic states

We usually consider the molecule to occupy the ground electronic state only, as the electronic excitation energy is quite high. Then, we set the ground state energy arbitrarily to zero ($E^{\text{el}}(0) = 0$) and the electronic partition function is equal to unity:

$$Q^{\text{el}} = \exp[-\beta E^{\text{el}}(0)] + \exp[-\beta E^{\text{el}}(1)] + \dots \approx 1 + 0 + \dots = 1 \quad (\text{VII.18})$$

This unity may be neglected in the product in Eq. VII.17.

2. Translational and rotational contribution

These contributions will not be discussed in detail. They are calculated for the quantum-mechanical particle in a box and for a quantum-mechanical rotator. One obtains

$$U^{\text{trans}} = \frac{3}{2} k_{\text{B}} T \quad (\text{VII.19})$$

$$U^{\text{rot}} = \frac{3}{2} k_{\text{B}} T \quad (\text{VII.20})$$

For both kinds of motion, also entropic contributions can be calculated. With those, we can estimate free energies.

3. Vibrational contribution

To get this contribution, we have to optimize the geometry of a molecule and calculate its vibrational frequencies ω_k . Each normal mode represents one harmonic oscillator k . For example, a water molecule has three modes, $k = 1, 2, 3$. Every harmonic oscillator has energy levels

$$E_k^m = \left(m + \frac{1}{2}\right) \cdot \hbar\omega_k \quad (\text{VII.21})$$

where $E_k^0 = \frac{1}{2}\hbar\omega_k$ is called the *zero point vibrational energy*. Let us evaluate the partition function of one harmonic oscillator:²⁹

$$\begin{aligned} Q_k &= \sum_{m=0}^{\infty} \exp \left[-\beta \left(m + \frac{1}{2} \right) \hbar\omega_k \right] = \exp \left[-\frac{1}{2}\beta\hbar\omega_k \right] \cdot \sum_{m=0}^{\infty} \exp [-\beta m\hbar\omega_k] = \\ &= \frac{\exp \left[-\frac{1}{2}\beta\hbar\omega_k \right]}{1 - \exp [-\beta\hbar\omega_k]} \end{aligned} \quad (\text{VII.22})$$

We can then derive the internal energy U_k

$$\ln Q_k = -\frac{1}{2}\beta\hbar\omega_k - \ln [1 - \exp[-\beta\hbar\omega_k]] \quad (\text{VII.23})$$

$$U_k = -\frac{\partial \ln Q_k}{\partial \beta} = \frac{1}{2}\hbar\omega_k - \left(\frac{\hbar\omega_k \exp[-\beta\hbar\omega_k]}{1 - \exp[-\beta\hbar\omega_k]} \right) = \hbar\omega_k \left(\frac{1}{2} + \frac{1}{\exp[\beta\hbar\omega_k] - 1} \right) \quad (\text{VII.24})$$

We have to do this for the $N - 6$ vibrational degrees of freedom of the molecule — each molecule thus consists of $N - 6$ harmonic oscillators. We then get for the internal energy

$$U = \sum_{k=1}^{N-6} U_k = \sum_{k=1}^{N-6} \hbar\omega_k \left(\frac{1}{2} + \frac{1}{\exp[\beta\hbar\omega_k] - 1} \right) \quad (\text{VII.25})$$

The zero-point vibrational energy contributions are not temperature dependent, and they are usually added to the electronic part of the energy (see above), so that the vibrational part of the internal energy is defined merely as

$$U^{\text{vib}} = \sum_{k=1}^{N-6} \left(\frac{\hbar\omega_k}{\exp[\beta\hbar\omega_k] - 1} \right) \quad (\text{VII.26})$$

and for the corresponding entropy contribution and the Helmholtz free energy we find

$$S^{\text{vib}} = k_B \ln Q^{\text{vib}} + \frac{U^{\text{vib}}}{T} = \sum_{k=1}^{N-6} \left(\frac{\hbar\omega_k}{\exp[\beta\hbar\omega_k] - 1} - \ln [1 - \exp[-\beta\hbar\omega_k]] \right) \quad (\text{VII.27})$$

$$F^{\text{vib}} = -k_B T \ln Q = \sum_{k=1}^{N-6} k_B T \ln [1 - \exp[-\beta\hbar\omega_k]] \quad (\text{VII.28})$$

²⁹ In the last step, we use the formula for the sum of geometric series ($\sum_{k=0}^{\infty} x^k = \frac{1}{1-x}$).

4. The PV term

To obtain the *enthalpy* H , we have to add the PV term to the internal energy. This term would be difficult to estimate unless we did not adopt the approximation of ideal gas:

$$PV = Nk_{\text{B}}T \quad (\text{VII.29})$$

Then, we obtain the enthalpy and the Gibbs free energy:

$$H = U + pV = U + Nk_{\text{B}}T \quad (\text{VII.30})$$

$$G = F + pV = F + Nk_{\text{B}}T \quad (\text{VII.31})$$

Many quantum chemical as well as molecular mechanics programs (like Gaussian, respectively CHARMM) perform a calculation of the thermal contributions by default, whenever vibrational analysis is requested. (This is because the calculation of vibrational frequencies is the most time consuming step and the evaluation of thermodynamic functions is done virtually ‘for free.’)

E. Aiming at free energies

For a molecule with many conformations, we have to approach the phase-space density with MD simulations ($\vec{r} = \{r_1, \dots, r_{3N}\}$, $\vec{p} = \{p_1, \dots, p_{3N}\}$):

$$\rho(\vec{r}, \vec{p}) = \frac{\exp[-\beta E(\vec{r}, \vec{p})]}{Q} \quad (\text{VII.32})$$

which is the (canonical) probability of finding the system at the point (\vec{r}, \vec{p}) .

The central point now is, how long an MD simulation we can perform. If we integrate the equations of motion for 1 ps, we will have 1,000 points in the trajectory; if we extend it to 1 ns, we already have a million points etc. Here, a nanosecond is close to the limit of what can be done!

Imagine we have simulated for 1 ps: Then, we will barely have sampled the points (\vec{r}, \vec{p}) for which $\rho(\vec{r}, \vec{p}) \leq \frac{1}{1000}$, meaning that any points with high energy will hardly be reached, while the low-energy region may already have been sampled very well.

We get a problem if we are willing to calculate the averages:

$$U = \langle E \rangle = \sum_{i=1}^M E(\vec{r}_i, \vec{p}_i) \cdot \rho(\vec{r}_i, \vec{p}_i) \quad (\text{VII.33})$$




FIG. 31: High-energy points B and C, which are badly sampled during the MD

Here, $\rho(\vec{x}, \vec{p})$ tells us the relative probability of ‘visit’ of the particular point (\vec{r}, \vec{p}) , i.e. how often the MD trajectory passes through this point. If we do not sample long enough, ρ will vanish for points with large energies $E(\vec{r}, \vec{p})$, and we will miss these large energies in the average.

One could think, that since we miss large values, the average of any quantity will be systematically wrong. However, for the internal energy this is not very serious, because the canonical probability distribution

$$\rho(\vec{r}, \vec{p}) = \frac{\exp[-\beta E(\vec{r}, \vec{p})]}{Q} \quad (\text{VII.34})$$

is very small if the energy is high, and the neglect of such points brings on no error.

For free energies, however, the situation is much worse, as we can write

$$\begin{aligned} F &= -k_B T \ln Q = k_B T \ln \frac{1}{Q} = k_B T \ln \left[\frac{c^{-1} \cdot \iint \exp[\beta E(\vec{r}, \vec{p})] \cdot \exp[-\beta E(\vec{r}, \vec{p})] d\vec{r} d\vec{p}}{Q} \right] = \\ &= k_B T \ln \left[\iint \exp[\beta E(\vec{r}, \vec{p})] \cdot \rho(\vec{r}, \vec{p}) d\vec{r} d\vec{p} \right] - \ln c \end{aligned} \quad (\text{VII.35})$$

(where the complicated integral in the numerator on the first line is just a resolution of the integral of unity, which corresponds to $c^{-1} = (8\pi^2 V)^N$; the canonical probability distribution ρ was cast).

Now, we have a real problem, since the large energy values enter an exponential term in the calculation of the free energy; the high-energy regions thus may contribute significantly. So, if we have too few points in these high-energy regions, we may find large errors in the calculated averages. Therefore, a really good idea is needed here to improve the sampling.