Tom Ouldridge
Keble College

Free Energy

1 Macroscopic Thermodynamics

You briefly met free energy last year, largely in the context of thermodynamics of macroscopic systems. For macroscopic systems in which only PV work is relevant, the quantities U, T, S, P, V, H (enthalpy), F (Helmholtz free energy), G (Gibbs free energy) and N_i (particle number of species i) are *functions of state*, in that their values are determined by the state of the system, and not by its history.

For a macrscopic equilibrium system we can specify the state using a small subset of these functions of state (for instance, the state of an ideal gas follows from knowing N, P and T – and indeed many other combinations) and all other functions of state are then specified. The values of these functions of state are assumed to be well-defined with negligible fluctuations.

2 Statistical Mechanics

If you want a more detailed discussion of the topics I skirt over here, please refer to Huang [1].

2.1 Ensembles and Partition Functions

I find it far more instructive to take the microscopic perspective of statistical mechanics. In this case, we explicitly sum over the microstates of the system. For any real system in equilibrium, the volume, number of particles and energy must be limited somehow, and there are two conceptually different ways to imagine how this might happen.

- The quantity could be absolutely fixed - for instance, a completely isolated system in a rigid box has fixed volume, energy and particle number.
- The quantity could vary, but be limited by the presence of a large reservoir with which exchange can occur. For instance, we could imagine a system in contact with a pressure reservoir (a piston?) – the volume of the system would fluctuate, but the constant external pressure stops the volume varying too much.

Different combinations of these restrictions are know as *ensembles*. Four common ones are:

- microcanonical ensemble: constant V, U and N.
- canonical ensemble: constant V, N and variable U, restrained by the presence of a reservoir of temperature T.
- isobaric-isothermal ensemble: constant N and variable U & V, restrained by temperature and pressure reservoirs.
- grand canonical ensemble: constant V and variable U & N, restrained by temperature and particle reservoirs.

One of the central assumptions of statistical mechanics is that, in the microcanonical ensemble, all microstates of a given V, U and N are equally likely. One can use this "postulate of equal a-priori probability" to derive the relative probability of states within the other ensembles.

- When in contact with a heat reservoir, the probability of being in a microstate of energy U is proportional to $\mathrm{e}^{(-U/kT)}$. This should be familiar from last year's stat mech course.
- When in contact with a pressure reservoir, the probability of being in a microstate of volume V is proportional to $\mathrm{e}^{(-PV/kT)}$.
- When in contact with a bath of particles, the probability of being in a microstate with N particles is proportional to $\mathrm{e}^{(\mu N/kT)}$, where μ is the chemical potential of the reservoir.

Note that one must make sure that identical particles are treated distinguishably. The *partition function* is constructed by summing over all possible states in the relevant ensemble, multiplied by their probability, and can be used to derive lots of useful things.

- In the canonical ensemble,
$$Z(N,V,T) = \sum_{\text{states}} \mathrm{e}^{-U/kT}. \tag{1}$$
Where the sum runs over all states of given N and V.
- In the isobaric-isothermal ensemble (note that there is a degree of subtlety in the definition of the normalizing volume V_0 – see the discussion of Corti [2]),
$$\Delta(N,P,T) = \int Z(N,V,T)\frac{\mathrm{e}^{-PV/kT}}{V_0}\mathrm{d}V. \tag{2}$$
- And in the grand canonical ensemble,
$$\Omega(\mu,P,T) = \sum_N Z(N,V,T)\mathrm{e}^{\mu N/kT}. \tag{3}$$

2.2 Fluctuations

We now think a little about how much a system in one of these ensembles will fluctuate. If a system is in the canonical ensemble, for example, V, N and T are clearly well defined. The energy, however, can vary - we can find $\langle U\rangle$ and also $\mathrm{Var}(U)$ using Z (can you see how to do this?), but, for the moment at least, we don't have a clear thermodynamic quantity U. P and μ are even harder to get a handle on.

In Fig. 1 I consider the two factors that contribute to the probability of being in a microstate of energy U in the canonical ensemble of a typical system. As you can see, they tend to be in opposition: the number of states available grows with U, but the probability of being in any single one of those states is suppressed by the $\mathrm{e}^{-U/kT}$. The results in a peaked distribution of the energy. As the system gets larger, the width of the peak (standard deviation of U) gets smaller relative to $\langle U\rangle$. This follows from a similar sort of argument to the fact that if you roll a dice 100 times, your average score will be closer to 3.5 than if you roll it 5 times.

So, for a thermodynamically large system, the fluctuations become *negligible*. In other words, $\langle U\rangle = U$ can be treated as a thermodynamic property of the system. Similarly, in the isobaric-isothermal and grand canonical ensembles, fluctuations about $\langle V\rangle$ and $\langle N\rangle$ become negligible for thermodynamically large systems.

The net result of this is that whichever ensemble a system is in, if it is thermodynamically large the quantities U, T, S, P, V and μ are all well defined. So why bother with different ensembles at all? They are actually at their most useful when we consider splitting up the partition function, which I will discuss in section 3. Another point to bear in mind is that for systems that are not thermodynamically large, fluctuations remain important and this equivalence does not hold.

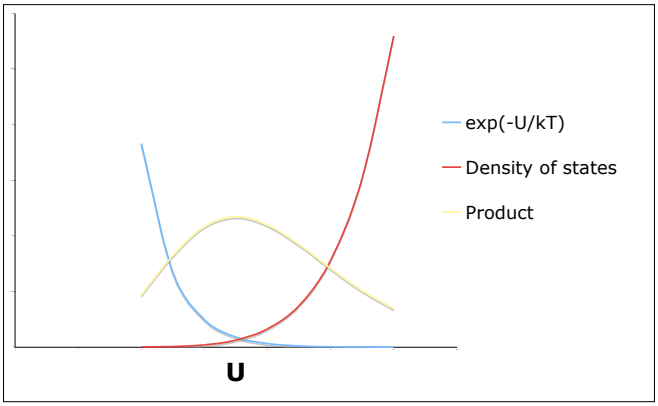

Figure 1: Schematic representation of the factors contributing to the distribution of energies in the canonical ensemble. The Boltzmann factor (blue) decreases with U, but the number of states available at a given energy (red) increases with U. The distribution of energies therefore contains a peak.

2.3 Free Energies

Last year you defined Helmholtz (F) an Gibbs (G) free energies for a thermodynamic system.

$$\begin{aligned} F &= U - TS \\ G &= U - TS + PV \end{aligned} \tag{4}$$

A second postulate of statistical mechanics is that entropy is a measure of the number of microstates occupied by the system, quantified by:

$$S = k \sum_{\text{states}} p_{\text{state}} \ln p_{\text{state}} \tag{5}$$

You should be able to show that, given this definition, we can infer $F = -kT \ln Z$, and $G = -kT \ln \Delta$. This is how I always think of free energies – logs of partition functions, which themselves are weighted sums over all states available to the system. In this way they include information about the energy and entropy associated with the system.

2.4 Chemical Potential

Thermodynamically, the chemical potential of a species is defined as the overall energy change of the system when a particle is added with the entropy and volume kept constant, i.e.,

$$\mu = \left(\frac{\partial U}{\partial N} \right)_{S,V}. \tag{6}$$

This isn't actually all that helpful. But, from the relationship between U, F and G, it is easy to show that:

$$\mu = \left(\frac{\partial F}{\partial N} \right)_{T,V} = \left(\frac{\partial G}{\partial N} \right)_{T,P}. \tag{7}$$

So μ is the change in F if a particle is added at constant T, V and the change in G if it is added at constant T, P. This will be helpful later.

In biophysics, there tend to be two important contributions to the chemical potential for species in dilute solution:

$$\mu \approx \mu_0(T) + kT \ln(C/C_0) + qV. \tag{8}$$

μ_0 depends on external conditions, and in most cases simply cancels out of any calculation. C is the concentration, measured relative to an arbitrary reference concentration C_0. So a system with a high concentration will have a high chemical potential. The final term takes electrostatic energy into account: q is the charge per particle and V the Voltage experienced.

3 Splitting-up the Partition Function

3.1 Thermodynamic Systems

Up until now, it has been implicitly assumed that when calculating anything, we included all possible states of the system in our ensemble. However, it is often useful to consider splitting up the partition function in a physically useful way, and considering a sum over a limited set of microstates.

Consider a system in the canonical ensemble containing of two species of particles, A and B. A and B can associate to form the dimer AB. Assuming that, other than during dimerization, the species are approximately ideal, let us calculate the partition function for the system *if we imagine it is restricted to having a fixed number of dimers* N_{AB}. In other words, we sum over states in which there are N_{AB} dimers, and for now ignore other possibilities. Let us call this reduced partition function $Z(N_{AB})$. We could sum up these reduced partition functions to give the total partition function if we wanted.

$$Z = \sum_{N_{AB}=0}^{N} Z(N_{AB}). \tag{9}$$

The point of the partition function is that it represents the sum of 'statistical weights' (relative probabilities) of microstates. If we sum over all possible microstates, the quantity is not actually that helpful directly. However, if we split it up using some 'order parameter' – such as N_{AB} – we can compare $Z(N_{AB})$ for different possible values of N_{AB} *to see which value is most likely.* The relative probabilities are simply given by the ratio of the partial partition functions.

We can also define a free energy as a function of N_{AB}:

$$F(N_{AB}) = -kT \ln Z(N_{AB}). \tag{10}$$

Therefore the value of N_{AB} that minimizes $F(N_{AB})$ is the most probable. In the thermodynamic limit, fluctuations are negligible relative to N_{AB} that minimizes $F(N_{AB})$, meaning that the value of N_{AB} that minimizes $F(N_{AB})$ *is the value we see.*

So to find the value of N_{AB}, we can differentiate $F(N_{AB})$. As the separate species are approximately ideal, and assuming that the total number of A and B is N, we find:

$$F(N_{AB}) = F_A(N - N_{AB}) + F_B(N - N_{AB}) + F_{AB}(N_{AB}), \tag{11}$$

where F_x is the contribution to the free energy from species x. For an ideal mixture, there is no interaction between different species and hence we can factorize the partition function, meaning that the free energy splits into additive terms for each species. Looking for the minimum, we find:

$$0 = \frac{\partial F_A(N - N_{AB})}{\partial N_{AB}} + \frac{\partial F_B(N - N_{AB})}{\partial N_{AB}} + \frac{\partial F_{AB}(N_{AB})}{\partial N_{AB}} = -\mu_A - \mu_B + \mu_{AB} \tag{12}$$

So if we know enough about μ of each species, we can calculate the degree of dimerization. This is an extremely widely used result, and something I myself use in my own research.

Hopefully now you are seeing the use of calculating the free energy as a function of an internal order parameter (N_{AB} in our case). What about if we had been in the isobaric-isothermal ensemble? We would have obtained the same result in the thermodynamic limit, $\mu_{AB} = \mu_A + \mu_B$. But there is a subtlety – the fact that the system is constrained differently (constant P rather than constant V) means that the state that the system finds this balance in will be different. Think about this for a while.

3.2 Other Examples

We can split up the partition function in any way we like - later this year, you will find the magnetization of a simple model for a magnet. One way to do this is to consider $F(M)$, the free energy as a function of total magnetization, and minimize it. Another instructive example is to consider two boxes containing the same type of particle which can exchange particles. For variety's sake, let's assume they are held at constant T, P. If the total number of particles is N, a sensible reaction coordinate would be N_1, the number of particles in one of the boxes. As an exercise, you should be able to show that the equilibrium is given by $\mu_1 = \mu_2$ by minimizing the total free energy as a function of N_1: $G(N_1) = G_1(N_1) + G_2(N_2)$.

This result has an important corollary. Instead of two separate boxes, we could think of our argument applying to any two regions of a system between which particles are able to move. Our argument implies that the equilibrium state is one in which μ in these two regions is equal - so the equilibrium state of a system through which particles are able to move has the same chemical potential everywhere. This is very important at biological membranes – if we apply a concentration difference across a membrane, a small flux of particles will set up a voltage difference, allowing μ to be the same on either side without the concentrations being equal (see Eqn. 8).

3.3 Small Systems

In small systems, we can again split up the partition function using internal reaction coordinates, and define free energies as a function of these coordinates. The only difference is that fluctuations are not negligible, and hence we cannot say that the system will spend all of its time at the point that minimizes the free energy. Instead, if we have two states A and B in m canonical ensemble, the relative occupation is given by:

$$\frac{P_A}{P_B} = \exp\left(-(F_A - F_B)/kT\right). \tag{13}$$

You should satisfy yourself that this is true by plugging in the partial partition functions. What can we say about the relative rates $k_{A\to B}$ and $k_{B\to A}$? You should prove that they are given by:

$$\frac{k_{B\to A}}{k_{A\to B}} = \exp\left(-(F_A - F_B)/kT\right). \tag{14}$$

Note that here $N_A k_{A\to B}$ gives the flux from A to B, so $k_{A\to B}$ is a rate per system in state A. The argument is essentially the same one used to derive the relationship between Einstein coefficients in lasers.

3.3.1 Myosin

When we analyzed myosin, we actually split up the partition function in a very complicated way. First, we had a discrete reaction coordinate which had six states given by differences in what was bound to myosin.

1. ATP bound to myosin, myosin bound to actin.
2. ATP bound to myosin, myosin not bound to actin.
3. ADP + P bound to myosin, myosin not bound to actin.
4. ADP + P bound to myosin, myosin bound to actin.
5. ADP bound to myosin, myosin bound to actin.
6. no ATP or ADP bound to myosin, myosin bound to actin.

We then divided these states up using a continuous reaction coordinate, which is the relative z-coordinate of the myosin linker and the actin filament. So we had a two dimensional reaction coordinate, and a free energy that looks like $G_i(z)$. Obviously, states 2 and 3 don't care about the relative position, so $G_{2/3}(z)$ look flat and the others have a minimum about some preferred binding orientation. By plotting the free energy like we did, it was possible to analyze how the motor was applying a force, and elucidate the essential features of the process.

Small technical point – the system we are actually considering is not very well defined. The only way I can make it rigorous is to consider a single myosin head in a large bath of solution. Within this solution, ADP and ATP are not equilibriated – so the free energy of the system before step 1 is obtained by only considering states with a fixed number of ATP and ADP. We then analyze the changes as the system takes one step towards equilibriation as a single ATP is converted into ADP + P.

4 Biophysics Technicalities

In biophysics, the distinction between Gibbs and Helmholtz free energies is often fudged. Strictly speaking, systems are generally at constant pressure, and so the isobaric-isothermal ensemble is probably appropriate. However, when trying to derive anything, we often treat water implicitly, so we can focus on the interesting solute. But the majority of pressure in most biological systems is due to the water, and the volume is relatively insensitive to what the solute does. So in a sense, the solute molecules see an approximately constant volume ensemble. As a result, the distinction between the two free energies is not that important and often they are used interchangeably.

The other important thing to remember about biological systems is that they are generally kept out of equilibrium – that is what being alive is. So all of this equilibrium stuff can be useful in telling us what is going on, but it often isn't the whole story. For instance, ions are frequently pumped across membranes, which means things are out of equilibrium – but our equilibrium calculations can tell us something about how hard you have to pump to set up a certain non-equilibrium steady-state.

References

[1] K. Huang. *Statistical Mechanics, Second Edition.* John Wiley & Sons, Inc., New York, 1987.

[2] David S. Corti. Isothermal-isobaric ensemble for small systems. *Phys. Rev. E*, 64(1):016128, 2001.