

Today's lecture: theory and implementation of Monte Carlo simulations in the NPT (isothermal-isobaric), μVT (Grand canonical) ensembles, and other ensembles.

General approach

One of the great benefits of Monte Carlo (MC) simulations are that we can simulate from any ensemble of interest. In this lecture, we examine some commonly-studied ensembles. Keep in mind that the general approach to deriving the methods consists of the following:

1. Determine the **microstate probability distribution**, ρ_m , for the ensemble of interest.
2. Determine a set of **Monte Carlo moves**. These moves should accomplish changes in all of the fluctuating quantities in the ensemble (e.g., energy, volume, or particle number fluctuations.)
3. Find the acceptance criterion by imposing **detailed balance**.

NPT Monte Carlo

In canonical MC, the number of particles and volume are fixed, specified as inputs to the simulation. We can then compute an estimate of the pressure using the virial relation, averaged over long simulation runs. Oftentimes, however, we are interested in specifying the pressure and determining the density (i.e., the volume of a simulation box for fixed number of particles). Such a task naturally leads one to the isothermal-isobaric ensemble.

Isothermal-isobaric ensemble (constant T, P, N)

The **isothermal-isobaric ensemble** corresponds to constant T, P, N conditions, when a system is held at constant temperature and pressure by connection to a heat and volume bath. That is, the system can exchange both energy and volume with its surroundings. The partition function is given by:

$$\begin{aligned} \Delta(T, P, N) &= \int_0^\infty e^{-\beta PV} Q(T, V, N) dV \\ &= \frac{1}{\Lambda(T)^{3N} N!} \int_0^\infty e^{-\beta PV} Z(T, V, N) dV \end{aligned}$$

The corresponding macroscopic thermodynamic potential is the **Gibbs free energy**,

$$G(T, P, N) = -k_B T \ln \Delta(T, P, N)$$

At these conditions, both system energy and volume fluctuate:

$$\wp(E, V) = \frac{\Omega(E, V, N)e^{-\beta E - \beta PV}}{\Delta(T, P, N)}$$

$$\wp(V) = \int \wp(E, V) dE$$

$$= \frac{Q(T, V, N)e^{-\beta PV}}{\Delta(T, P, N)}$$

Importantly, one microstate in the system corresponds to both a set of momenta and positions, and a value of the system volume. If we consider only the configurational coordinates,

$$\wp_m = \wp(\mathbf{r}^N, V) d\mathbf{r}^N dV$$

$$= \frac{e^{-\beta U - \beta PV}}{\Lambda(T)^{3N} N!} \times \frac{d\mathbf{r}^N dV}{\Delta(T, P, N)}$$


The differential position and volume elements must be included so that \wp_m is dimensionless.

Monte Carlo implementation

In the isothermal-isobaric ensemble, both the energy and the volume of a system fluctuate. Instead of specifying T, V, N , we specify T, P, N . That is, we specify the pressure in addition to the temperature.

In the simulation, we then accomplish energy fluctuations using displacement moves, as with canonical MC. For volume fluctuations, we can use **volume scaling moves**:

1. Pick a random change in volume ΔV in the range $[-\delta V_{\max}, \delta V_{\max}]$ and make the new volume $V \leftarrow V + \Delta V$.
2. Scale the entire simulation box uniformly along each axis.
3. Upon the volume change, scale the particle positions uniformly. If molecules are being studied, the centers of mass of the molecules are scaled uniformly, rather than the individual atoms (which would result in a severe distortion of bonds and angles.)



In scaling the positions of the particles, we keep the **dimensionless** or **scaled positions** the same:

$$\left(s_{x,1}, s_{y,1}, \dots, s_{z,N}\right) = \left(\frac{x_1}{L}, \frac{y_1}{L}, \dots, \frac{z_N}{L}\right)$$

$$\mathbf{s}^N = \frac{\mathbf{r}^N}{L}$$

We derive the acceptance criterion for volume scaling moves by imposing detailed balance:

$$\frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} = \frac{\alpha_{21} \rho_2}{\alpha_{12} \rho_1}$$

Keep in mind that here the indices 1 and 2 correspond to the sets of variables (\mathbf{r}_1^N, V_1) and (\mathbf{r}_2^N, V_2) , respectively. The proposal probability α_{12} gives the probability that we will pick (\mathbf{r}_2^N, V_2) given that we are at (\mathbf{r}_1^N, V_1) . There are two parts to this probability:

The probability that we will pick V_2 given V_1 , for $|V_2 - V_1| < \delta V_{\text{max}}$, is given by

$$\alpha(V_1 \rightarrow V_2) = \frac{1}{2\delta V_{\text{max}}}$$

Technically, this is a differential probability because volume is a continuous variable. We therefore have:

$$\int_{V-\delta V_{\text{max}}}^{V+\delta V_{\text{max}}} \alpha(V_1 \rightarrow V_2) = \int_{V-\delta V_{\text{max}}}^{V+\delta V_{\text{max}}} \frac{dV}{2\delta V_{\text{max}}} = 1$$

The second part of the proposal probability is that we will pick \mathbf{r}_2^N given \mathbf{r}_1^N . This part of our move requires some careful treatment. Recall that the positions of all of the atoms are given by

continuous variables. We might consider “one” configuration as sitting in a region $d\mathbf{r}^N$ about the coordinates \mathbf{r}^N . If we uniformly scale the positions, then the differential volume $d\mathbf{r}^N$ will change. In other words, a different amount of configurations in one volume will map to the same differential regions in configuration space in the new volume. We need to account for this change:

$$\frac{\alpha(\mathbf{r}_1^N \rightarrow \mathbf{r}_2^N)}{\alpha(\mathbf{r}_2^N \rightarrow \mathbf{r}_1^N)} = \frac{d\mathbf{r}_2^N}{d\mathbf{r}_1^N}$$

We can evaluate this expression because the scaled positions stay the same:

$$\begin{aligned} \frac{\alpha(\mathbf{r}_1^N \rightarrow \mathbf{r}_2^N)}{\alpha(\mathbf{r}_2^N \rightarrow \mathbf{r}_1^N)} &= \frac{V_2^N d\mathbf{s}_2^N}{V_1^N d\mathbf{s}_1^N} \\ &= \left(\frac{V_2}{V_1}\right)^N \end{aligned}$$

An alternate, practical way to think about this probability is to recall that computers have finite precision in storing the positions of all of the atoms, which ultimately means there are a finite number of configurations. Consider one atom in one dimension in a “volume” of length L . In principle, we could count all of the possible configurations in one volume L_1 . At a different volume $L_2 > L_1$, there would be a greater number of possible configurations available, since we would have a greater range of the particle coordinate in the new volume. By scaling the possible coordinates at L_1 to new volume L_2 , we can only generate a fraction of the possible configurations in L_2 because we always start with one of the finite possibilities at L_1 . That fraction is (L_1/L_2) . For three dimensions and N particles, the fraction is $(L_1/L_2)^{3N}$.

Putting together the two components of the move proposal probabilities,

$$\begin{aligned} \frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} &= \frac{\alpha_{21}\wp_2}{\alpha_{12}\wp_1} \\ &= \frac{(1/2\delta V_{\text{max}})}{(1/2\delta V_{\text{max}})} \left(\frac{V_2}{V_1}\right)^N \frac{\wp_2}{\wp_1} \end{aligned}$$

We are now prepared to insert the isothermal-isobaric microstate probabilities:

$$\begin{aligned} \frac{\wp_2}{\wp_1} &= \left[\frac{e^{-\beta U_2 - \beta P V_2}}{\Lambda(T)^{3N}} \times \frac{d\mathbf{r}^N dV}{\Delta(T, P, N)} \right] \left[\frac{e^{-\beta U_1 - \beta P V_1}}{\Lambda(T)^{3N}} \times \frac{d\mathbf{r}^N dV}{\Delta(T, P, N)} \right]^{-1} \\ &= e^{-\beta \Delta U - \beta P \Delta V} \end{aligned}$$

Finally, we arrive at:

$$\frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} = \left(\frac{V_2}{V_1}\right)^N e^{-\beta \Delta U - \beta P \Delta V}$$

As with canonical MC, we can satisfy this acceptance requirement using the Metropolis criterion:

$$P_{12}^{\text{acc}} = \min \left[1, \left(\frac{V_2}{V_1} \right)^N e^{-\beta\Delta U - \beta P\Delta V} \right]$$

It is often beneficial for precision reasons to rearrange this equation so that all terms are inside the exponential:

$$\begin{aligned} P_{12}^{\text{acc}} &= \min \left[1, \exp \left(N \ln \left(\frac{V_2}{V_1} \right) - \beta\Delta U - \beta P\Delta V \right) \right] \\ &= \exp \left\{ \min \left[0, N \ln \left(\frac{V_2}{V_1} \right) - \beta\Delta U - \beta P\Delta V \right] \right\} \end{aligned}$$

Thus, a typical *NPT* simulation entails the following:

1. Random particle displacement moves, with $P_{12}^{\text{acc}} = \exp\{\min[0, -\beta\Delta U]\}$. A random number r is drawn in $[0.0, 1.0)$ and the move is accepted if $P_{12}^{\text{acc}} > r$.
2. Random volume scaling moves, with $P_{12}^{\text{acc}} = \exp\{\min[0, N \ln(V_2/V_1) - \beta\Delta U - \beta P\Delta V]\}$. A random number r is drawn in $[0.0, 1.0)$ and the move is accepted if $P_{12}^{\text{acc}} > r$.

Practical considerations

Molecular systems

For molecular systems, we do not scale the positions of all of the atoms when we perform a volume scaling move, but rather the centers of mass of each molecule. Otherwise, we would greatly distort bonds in a way that would make the change in potential energy very unfavorable for acceptance of the move.

When one scales the centers of mass, the considerations for changes in the differential volume elements associated with $\alpha(\mathbf{r}_1^N \rightarrow \mathbf{r}_2^N)$ applies only to the N_{mol} degrees of freedom. Thus the N term appearing in the acceptance criterion is no longer the total number of atoms but N_{mol} :

$$P_{12}^{\text{acc}} = \min \left[1, \left(\frac{V_2}{V_1} \right)^{N_{\text{mol}}} e^{-\beta\Delta U - \beta P\Delta V} \right]$$

Log-volume scaling moves

For specified T, P that place systems near liquid-gas phase coexistence, or near the critical point, the natural volume fluctuations can be very large as the system traverses between the two phases. In these cases, the simple volume increments presented above can become inefficient because it will require many displacements ΔV to traverse between gas and liquid states.

In such cases, it becomes much more efficient to propose random moves in the logarithm of the volume rather than the volume itself. Such **log-volume scaling moves** entail the following:

1. Pick a random value $\delta (= \ln \Delta V)$ on the uniform distribution $[-\delta_{\max}, \delta_{\max}]$.
2. Let $\Delta V = \exp(\delta)$.
3. Update $V \leftarrow V + \Delta V$ and scale the particle positions as before.

For log-volume scaling moves, the proposal probability $\alpha(V_1 \rightarrow V_2)$ changes because we are now sampling non-uniformly in volume.

$$\alpha(V_1 \rightarrow V_2) \propto V_2$$

This modifies the acceptance criterion. When one proceeds through the derivation,

$$\begin{aligned} P_{12}^{\text{acc}} &= \min \left[1, \left(\frac{V_2}{V_1} \right)^{N+1} e^{-\beta \Delta U - \beta P \Delta V} \right] \\ &= \exp \left\{ \min \left[0, (N+1) \ln \left(\frac{V_2}{V_1} \right) - \beta \Delta U - \beta P \Delta V \right] \right\} \end{aligned}$$

Notice the $N + 1$ rather than N as before.

Other MC volume moves

Isotropic scaling moves are certainly not the only kinds of MC moves that can be implemented to change the volume. One can perform anisotropic moves, in which scaling is performed along only one axis. Such moves can be beneficial for systems that have fixed interfaces or walls in the other directions. In these cases, the acceptance criterion must be re-derived using a procedure similar to above.

Choice of δV_{\max}

Typically the maximum volume displacement is adjusted so that the average acceptance ration is roughly 30-50%. Too small values of this parameter result in a slow exploration of volume space. Too large will result in configurations with core overlaps upon scaling the volume down.

Computing the potential energy

After a volume scaling move, the total potential energy of the system needs to be recalculated, an operation whose expense scales as N^2 . Sometimes, it becomes possible to speed this step for specific kinds of potential energy functions. For a system of particles interacting exclusively through a pairwise energy function of the form

$$u(r_{ij}) = ar_{ij}^{-n}$$

The new energy can be computed simply from a scaling of the old energy,

$$U_2 = \left(\frac{L_2}{L_1}\right)^{-n} U_1$$

The Lennard-Jones potentials involves two terms of this form, where $n = 12$ and $n = 6$. One can maintain separate sums of each of these terms in the simulation and use the above relation to compute the new total energy for each term after volume scaling.

For systems involving bonded interactions, this kind of simplification is generally not feasible and the total energy must be re-evaluated completely at each volume scaling move.

Frequency of moves

In an *NPT* simulation, particle displacement moves typically require an expense proportional to N , while volume scaling moves scale as N^2 . It is therefore customary to attempt one volume scaling move, on average, for every N displacement moves attempted. A typical procedure is:

1. Pick a random number r on $[0.0,1.0)$.
2. If $r < (N + 1)^{-1}$, perform a volume scaling move. Otherwise, perform a displacement move.

The random component of move selection is extremely important to the correct convergence of the Markov chain. Thus, one should *not* explicitly cycle through the N particles performing displacement moves and *then* perform a volume scaling attempt. Such regularity of moves, without a random component, will ultimately bias the stationary distribution.

Check: the virial

Though the pressure is explicitly specified in an *NPT* simulation, one can still compute the average pressure through the virial relation,

$$\begin{aligned}\langle P \rangle &= \left\langle \frac{Nk_B T}{V} - \frac{W}{3V} \right\rangle \\ &= Nk_B T \left\langle \frac{1}{V} \right\rangle - \left\langle \frac{W}{3V} \right\rangle\end{aligned}$$

This average pressure is subject to statistical fluctuations in the simulation and the extent to which equilibrium can be achieved for the system of interest. However, it can be shown that in the limit of infinitely long, converged simulations, the virial pressure will equal the imposed pressure. This provides a useful check for *NPT* simulations.

Grand-canonical Monte Carlo

In a grand-canonical Monte Carlo (GCMC) simulation, we specify the temperature and chemical potential of the species of interest, for a fixed simulation volume. This leads to fluctuations in both the energy and number of particles. Like the isothermal-isobaric ensemble, this also provides a way to determine the density.

Grand canonical ensemble (constant T, V, μ)

The **grand canonical ensemble** occurs when a system is coupled to a heat bath and an infinite reservoir with which it can exchange particles. The partition function is:

$$\begin{aligned}\Xi(T, V, \mu) &= \sum_{N=0}^{\infty} Q(T, V, N) e^{\beta N \mu} \\ &= \sum_{N=0}^{\infty} \frac{Z(T, V, N)}{\Lambda(T)^{3N} N!} e^{\beta N \mu}\end{aligned}$$

The macroscopic thermodynamic potential is:

$$PV = k_B T \ln \Xi(T, V, \mu)$$

In the grand canonical ensemble, both the energy and number of particles fluctuate:

$$\begin{aligned}\wp(E, N) &= \frac{\Omega(E, V, N) e^{-\beta E + \beta \mu N}}{\Xi(T, V, \mu)} \\ \wp(N) &= \int \wp(E, N) dE \\ &= \frac{Q(T, V, N) e^{\beta \mu N}}{\Xi(T, V, \mu)}\end{aligned}$$

One microstate in the system corresponds to both a set of momenta and positions, and a value of the particle number N . If we consider only the configurational coordinates,

$$\begin{aligned}\wp_m &= \wp(\mathbf{r}^N, N) d\mathbf{r}^N \\ &= \frac{e^{-\beta U + \beta \mu N}}{\Lambda(T)^{3N} N!} \times \frac{d\mathbf{r}^N}{\Xi(T, V, \mu)}\end{aligned}$$

The differential position elements must be included so that \wp_m is dimensionless.

There is, however, one major subtlety with this partition function and microstate distribution. The equations above assume that the particles are indistinguishable, giving rise to the $N!$ term in the denominator. In reality, in our simulation, they are **distinguishable**. Each particle always has

a definite set of coordinates and is indexed with a number in the internal routines of our simulation runs. We treat particles as distinguishable: we *pick* specific particles to displace, for example. Therefore, the true distribution function we will implement in our simulation is that for distinguishable particles, which omits the $N!$ term:

$$\rho_m = \frac{e^{-\beta U + \beta \mu N}}{\Lambda(T)^{3N}} \times \frac{d\mathbf{r}^N}{\Xi(T, V, \mu)}$$

Monte Carlo implementation

In the grand canonical ensemble, both the energy and the particle number fluctuates. Instead of specifying T, V, N , we specify T, V, μ . That is, we specify the chemical potential in addition to the temperature. In the following example, we will assume that we have a monatomic system.

In the simulation, we then accomplish energy fluctuations using displacement moves, as with canonical MC. For particle number fluctuations, we can use **particle addition and deletion moves**. For an addition move,

1. Pick a random set of coordinates in the simulation box. That is, pick x, y, z separately in the uniform distribution $[-L/2, L/2]$.
2. Insert a new particle at that position.

A random deletion move entails,

1. Randomly pick a particle of the N current particles.
2. Delete the particle.

We derive the acceptance criteria for these moves by imposing detailed balance:

$$\frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} = \frac{\alpha_{21} \rho_2}{\alpha_{12} \rho_1}$$

Keep in mind that here the indices 1 and 2 correspond to the sets of variables $(\mathbf{r}_1^{N_1}, N_1)$ and $(\mathbf{r}_2^{N_2}, N_2)$, respectively. The proposal probability α_{12} gives the probability that we will pick $(\mathbf{r}_2^{N_2}, N_2)$ given that we are at $(\mathbf{r}_1^{N_1}, N_1)$.

Particle additions

Here, we will consider a particle addition move, such that $N_2 = N_1 + 1$. For simplicity of notation we will simply use $N_1 = N$ and $N_2 = N + 1$.

The move proposal probability for insertion stems from the fact that we pick a random location in space and place the particle there,

$$\alpha_{12} = \frac{d\mathbf{r}}{V}$$

Note that, as expected, this probability is one when integrated over all possibilities, i.e., over all space,

$$\int_V \frac{d\mathbf{r}}{V} = 1$$

The proposal probability for the reverse move, however, is different. The reverse move for $2 \rightarrow 1$ involves deletion of the particle we inserted. If we randomly pick a particle to delete, there is a $1/N_2$ chance that we will pick the one that we inserted. Thus,

$$\alpha_{21} = \frac{1}{N + 1}$$

Using the detailed balance condition,

$$\begin{aligned} \frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} &= \frac{\left(\frac{1}{N+1}\right)}{\left(\frac{d\mathbf{r}}{V}\right)} \left[\frac{e^{-\beta U_2 + \beta \mu(N+1)}}{\Lambda(T)^{3(N+1)}} \times \frac{d\mathbf{r}^{N+1}}{\Xi(T, V, \mu)} \right] \left[\frac{e^{-\beta U_1 + \beta \mu N}}{\Lambda(T)^{3N}} \times \frac{d\mathbf{r}^N}{\Xi(T, V, \mu)} \right]^{-1} \\ &= \frac{V}{(N+1)\Lambda(T)^3} e^{-\beta(U_2 - U_1) + \beta \mu} \\ &= \frac{V}{(N+1)\Lambda(T)^3} e^{-\beta \Delta U + \beta \mu} \end{aligned}$$

Notice that the differential volume elements $d\mathbf{r}$ all cancelled, even though there were a different amount in each of the microstate probabilities. The differential from the move proposal probability compensated in this regard. Any move that inserts a particle should have this differential since there will always be some volume element associated with placing a particle at a location in space.

The acceptance relation includes the thermal de Broglie wavelength. In practice for very simple systems, such as the Lennard-Jones system, this constant is typically absorbed inside the chemical potential. We define a **relative chemical potential**, referenced to something that resembles the ideal gas standard chemical potential:

$$\mu' \equiv \mu - k_B T \ln \Lambda^3(T)$$

The acceptance relation then becomes,

$$\frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} = \frac{V}{N+1} e^{-\beta \Delta U + \beta \mu'}$$

This relation can be satisfied using the normal Metropolis criterion,

$$P_{12}^{\text{acc}} = \min \left[1, \frac{V}{N+1} e^{-\beta\Delta U + \beta\mu'} \right]$$

For precision reasons, this can be arranged to the following:

$$\begin{aligned} P_{12}^{\text{acc}} &= \min \left\{ 1, \exp \left[-\beta\Delta U + \beta\mu' + \ln \left(\frac{V}{N+1} \right) \right] \right\} \\ &= \exp \left\{ \min \left[0, -\beta\Delta U + \beta\mu' + \ln \left(\frac{V}{N+1} \right) \right] \right\} \end{aligned}$$

Particle deletions

For a particle deletion move, we will simply use $N_1 = N$ and $N_2 = N - 1$. A similar derivation to the one above shows that the correct criterion satisfies

$$\frac{P_{12}^{\text{acc}}}{P_{21}^{\text{acc}}} = \frac{\Lambda(T)^3 N}{V} e^{-\beta\Delta U - \beta\mu}$$

Again, we can define a relative chemical potential and absorb the thermal de Broglie wavelength inside it. Further using the Metropolis rule, we arrive at the following acceptance criterion:

$$\begin{aligned} P_{12}^{\text{acc}} &= \min \left[1, \frac{N}{V} e^{-\beta\Delta U - \beta\mu'} \right] \\ &= \exp \left\{ \min \left[0, -\beta\Delta U - \beta\mu' + \ln \left(\frac{N}{V} \right) \right] \right\} \end{aligned}$$

Note that the criteria for acceptance and deletion of moves are *not* the same. That is, this is not a symmetric MC move.

Practical considerations

Molecular systems

For rigid molecules, upon insertion we must choose both a center of mass position and an orientation. Due to the inclusion of orientational variables, the acceptance criterion will then entail a “volume” of orientational space due to one molecule. For diatomic molecules, this volume is 4π . For larger molecules, the volume is $8\pi^2$. The difference comes from the fact that diatomic molecules only require two angles to describe orientation, while higher-order polyatomic molecules require three. In either case, this constant can be absorbed into the relative chemical potential. The acceptance criterion then depends on the number of center-of-mass positions:

$$P_{12}^{\text{acc}} = \min \left[1, \frac{V}{N_{\text{mol}} + 1} e^{-\beta\Delta U + \beta\mu'} \right] \quad \text{for insertions}$$

$$P_{12}^{\text{acc}} = \min \left[1, \frac{N_{\text{mol}}}{V} e^{-\beta \Delta U - \beta \mu'} \right] \quad \text{for deletions}$$

For nonrigid molecules, one must select a conformation of the molecule before inserting it. The selection of a conformation will affect the move proposal probabilities and thus modify the acceptance criterion. One approach may be to take these conformations from a **reservoir** of structures extracted from a single-molecule canonical simulation at the same temperature. In this case, the part of the move proposal probability corresponding to selection of the molecule's internal degrees of freedom will be proportional to $\exp(-\beta U)$.

Computing the potential energy

After a particle addition move, the change in potential energy of the system needs to be recalculated. Since only a single particle is added, one only needs to compute the N pairwise interaction energies with the newly-added particle. Similarly for a particle deletion, only the N pairwise energies of the deleted molecule need to be computed (prior to deletion). Thus the expense of these moves scales as N rather than N^2 . This makes grand canonical simulation particularly efficient for exploring variations in density.

Frequency of moves

In an μVT simulation, both particle displacement and insertion/deletion moves require an expense proportional to N for pairwise potentials. Thus typical fractions of attempted moves dedicated to additions/deletions are 30-50%, with displacements the remaining balance. Keep in mind that the type of move should be chosen randomly, by drawing from a random distribution.

Efficiency of insertion and deletion moves

Particle insertion and deletion moves can offer an important advantage relative to other kinds of moves: they are often thought to enhance ergodicity and the exploration of configuration space because they enable particles to suddenly appear and disappear in the simulation box. That is, they often help speed the equilibration portion of the simulation because they impart such drastic changes on the system configuration.

On the other hand, dense systems like liquids can render insertion/deletion moves very inefficient. In these cases, an insertion frequently causes a core overlap with an existing particle; the energy becomes very high, and the move is rejected. Unlike displacement and volume scaling moves, there is no parameter to tune to achieve a target acceptance ratio. It is not infrequent to find reported acceptance ratios for GCMC simulations that in the range 0.1%-1.0% for liquids. For crystalline solids, these moves are completely impractical as presented here.

Multicomponent systems

For multicomponent systems, a GCMC simulation entails a separate chemical potential for each molecular species. In this case, acceptance and deletion moves must be performed for each kind of molecule using the appropriate chemical potential and component number of molecules. Alternatively, it is possible to perform a partial GCMC simulation in which only a subset of the chemical potentials is specified and the number of molecules of the other species is fixed.

Comparison of isothermal-isobaric and GCMC simulations

Thermodynamic state and fluctuations

Both NPT and μVT simulations enable one to specify two thermodynamic field variables: the temperature and either the pressure P or chemical potential μ . In this respect, there is only one extensive parameter that defines the size of the system (N or V). These simulations are convenient because they enable one to uniquely specify the **intensive thermodynamic state** of the system at the outset. Keep in mind that, for a single component system, the Gibbs phase rule dictates that only two such parameters are required.

In addition, both of these ensembles capture both **energy and density fluctuations**. As we will see in later lectures, such fluctuations are important to understanding phase equilibria, notably liquid-gas phase transitions. Ultimately, we are able to relate the observed fluctuations to the underlying density of states. As a preview, consider the isothermal-isobaric ensemble. For specified T and P , the observed joint distribution of energy and volume fluctuations obeys the relation:

$$\wp(E, V) \propto \Omega(E, V, N) e^{-\beta E - \beta P V}$$

If we measure the distribution on the LHS—for example, via a histogram in an NPT simulation—we can back out the density of states to within an unknown, multiplicative constant. This enables us to begin computing entropies and free energies in our simulations. Moreover, there are important techniques for using this perspective to predict conditions of phase equilibrium.

Advantages and disadvantages

Both of the methods described explore both density and energy fluctuations. The choice of the isothermal-isobaric or the grand canonical ensemble for a simulation can be motivated by several factors:

- For molecular systems of nonrigid molecules, NPT simulations are far easier to implement.

- For atomic systems or molecular systems of rigid molecules, μVT simulations can be much more efficient because the simulation expense for a density change scales as N rather than N^2 for pairwise potentials in the general case. However, speedups may be possible for particular potentials in the NPT ensemble.
- μVT simulations can speed equilibration due to the fast exploration of configuration space that insertion and deletion moves promote.
- For very dense systems, grand canonical insertions and deletions are rarely accepted, and the NPT ensemble is more efficient.
- For structured phases, like crystalline solids, the NPT ensemble is much preferred.

The Gibbs Ensemble

The Gibbs ensemble is, in some sense, a combination of both isothermal-isobaric and GCMC simulation methods. It was devised by Panagiotopoulos in 1987 as a method for studying phase coexistence. While it falls in the general class of phase equilibria methods that we will discuss in more detail in later lectures, it is beneficial for us to discuss it here given its connection to the approaches just described.

A distinguishing feature of the Gibbs ensemble is that it entails *two* simulation boxes (“subsystems”) that are coupled to each other. The coupling between them is done in such a way as to achieve coexistence between two phases.

Statistical-mechanical formulation

Consider equilibrium between two phases A and B. Phases exist spontaneously on their own without the use of partitions or membranes or other interventions, and since they are in equilibrium with each other, they can exchange energy, volume, and particles. When such is the case, the conditions of equilibrium are given by maximization of the entropy subject to the constraints of constant total energy, volume, and particles. For two phases, this becomes:

$$\max[S_A(E_A, V_A, N_A) + S_B(E_B, V_B, N_B)]$$

In general, at a maximum, we have:

$$dS_A + dS_B = 0$$

$$\left(\frac{dS_A}{dE_A}\right) dE_A + \left(\frac{dS_A}{dV_A}\right) dV_A + \left(\frac{dS_A}{dN_A}\right) dN_A + \left(\frac{dS_B}{dE_B}\right) dE_B + \left(\frac{dS_B}{dV_B}\right) dV_B + \left(\frac{dS_B}{dN_B}\right) dN_B = 0$$

$$\frac{1}{T_A} dE_A + \frac{P_A}{T_A} dV_A - \frac{\mu_A}{T_A} dN_A + \frac{1}{T_B} dE_B + \frac{P_B}{T_B} dV_B - \frac{\mu_B}{T_B} dN_B = 0$$

But since $dE_A + dE_B = 0$, $dV_A + dV_B = 0$, and $dN_A + dN_B = 0$, by the constancy of the total energy, volume, and number of particles:

$$\left(\frac{1}{T_A} - \frac{1}{T_B}\right) dE_A + \left(\frac{P_A}{T_A} - \frac{P_B}{T_B}\right) dV_A - \left(\frac{\mu_A}{T_A} - \frac{\mu_B}{T_B}\right) dN_A = 0$$

Recognizing that dE_A , dV_A , and dN_A can all be varied independently in this equation, the conditions for equilibrium between two phases are given by:

$$T_A = T_B \quad P_A = P_B \quad \mu_A = \mu_B$$

We can also consider the case when the two systems are also in contact with a very large heat bath that maintains them at constant temperature. When such is the case, the partition function for the aggregate system is the product of the partition functions for the two subsystems. If the two systems are not allowed to exchange volume or particles, the total partition function is:

$$\begin{aligned} Q(T, V_A, V_B, N_A, N_B) &= Q_A(T, V_A, N_A) Q_B(T, V_B, N_B) \\ &= \left[\frac{1}{N_A! \Lambda(T)^{3N_A}} \int e^{-\beta U(\mathbf{r}_A^{N_A})} d\mathbf{r}_A^{N_A} \right] \left[\frac{1}{N_B! \Lambda(T)^{3N_B}} \int e^{-\beta U(\mathbf{r}_B^{N_B})} d\mathbf{r}_B^{N_B} \right] \end{aligned}$$

On the other hand, if the two phases are constrained to *exchange* particles and volume, then we must sum the total system partition function over all possible allocations of volume and particles between the two phases for fixed $V = V_A + V_B$ and $N = N_A + N_B$:

$$\begin{aligned} Q(T, V, N) &= \sum_{N_A=0}^N \int_0^V Q_A(T, V_A, N_A) Q_B(T, V - V_A, N - N_A) dV_A \\ &= \sum_{N_A=0}^N \int_0^V \left[\frac{1}{N_A! \Lambda(T)^{3N_A}} \int_{V_A} e^{-\beta U(\mathbf{r}_A^{N_A})} d\mathbf{r}_A^{N_A} \right] \left[\frac{1}{(N - N_A)! \Lambda(T)^{3(N - N_A)}} \int_{V - V_A} e^{-\beta U(\mathbf{r}_B^{N_B})} d\mathbf{r}_B^{N_B} \right] dV_A \end{aligned}$$

Combining the integrals and simplifying,

$$Q(T, V, N) = \frac{1}{\Lambda(T)^{3N}} \sum_{N_A=0}^N \int_0^V \frac{1}{N_A! (N - N_A)!} \left[\int_{V_A} e^{-\beta U(\mathbf{r}_A^{N_A})} d\mathbf{r}_A^{N_A} \right] \left[\int_{V - V_A} e^{-\beta U(\mathbf{r}_B^{N_B})} d\mathbf{r}_B^{N_B} \right] dV_A$$

In the total system, one microstate corresponds to the specification of the positions of all of the atoms in each system ($\mathbf{r}_A^N, \mathbf{r}_B^N$) and the values of N_A and V_A :

$$\wp(\mathbf{r}_A^N, \mathbf{r}_B^N, N_A, V_A) \propto \frac{1}{N_A! (N - N_A)!} e^{-\beta U(\mathbf{r}_A^{N_A})} e^{-\beta U(\mathbf{r}_B^{N_B})}$$

Monte Carlo implementation

The Gibbs ensemble simulates two phases that are in equilibrium. Each phase corresponds to its own, separate simulation box in the computer. The phases are both connected to a heat bath to maintain constant temperature conditions. Phase equilibrium is achieved by allowing the boxes to exchange volume and particles. The phases are non-interacting other than this fact; i.e., there are no potential energies between particles in different simulation boxes. There are then three types of MC moves in a Gibbs ensemble simulation:

Particle displacement moves are performed in a manner similar to the ensembles we have just discussed. A random particle is picked out of the two simulation boxes. This move creates energy fluctuations in the subsystems.

Volume exchange moves are performed in which one phase (simulation box) decreases in volume by an amount ΔV and the other phase increases in volume by exactly this amount. Thus, volume is exchanged. In the thermodynamic limit, as we have seen before, this results in equality of the average pressures between the two phases. Upon exchanging volume, we can scale the particle positions in each phase as before in the NPT ensemble. The acceptance criterion for this kind of move follows from the probability distribution above and considerations similar to those discussed for the isothermal-isobaric ensemble:

$$P_{12}^{\text{acc}} = \min \left[1, \frac{(V_A + \Delta V)^{N_A} (V_B - \Delta V)^{N_B}}{V_A^{N_A} V_B^{N_B}} e^{-\beta \Delta U_A - \beta \Delta U_B} \right]$$

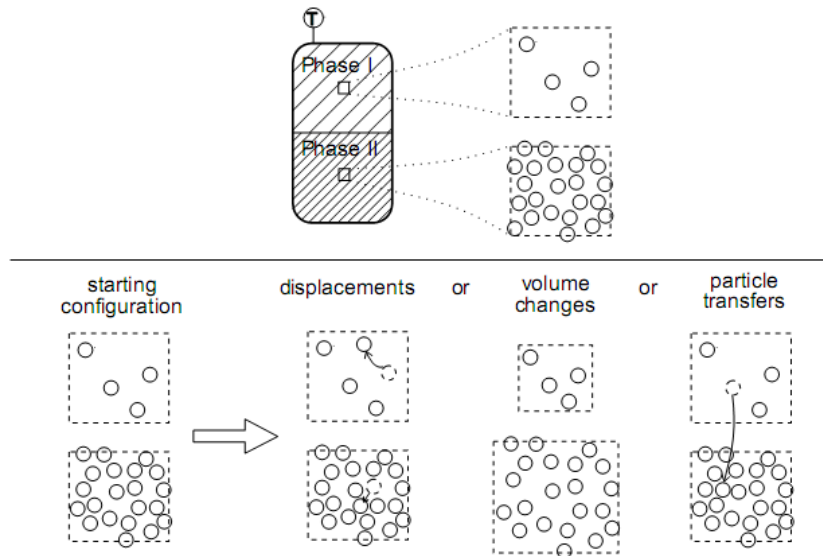
Here, the values of V_A, V_B, N_A, N_B are those prior to proposing the move. ΔV is the change in volume proposed for phase 1, chosen randomly in the range $[-\delta V_{\text{max}}, \delta V_{\text{max}}]$. The terms ΔU_A and ΔU_B are the changes in potential energy due to volume scaling in phases A and B, respectively.

Particle exchange moves are performed in which one particle in one phase is removed and added to a random location in the other. In the thermodynamic limit, this exchange of mass results in equality of the chemical potential between the two phases. In practice, a particle exchange move resembles a simultaneous deletion and insertion move. The acceptance criterion follows the microstate distribution above and is derived using ideas similar to those discussed for the GCMC ensemble:

$$P_{12}^{\text{acc}} = \min \left[1, \frac{N_A V_B}{(N_B + 1) V_A} e^{-\beta \Delta U_A - \beta \Delta U_B} \right] \quad \text{particle transferred } A \rightarrow B$$

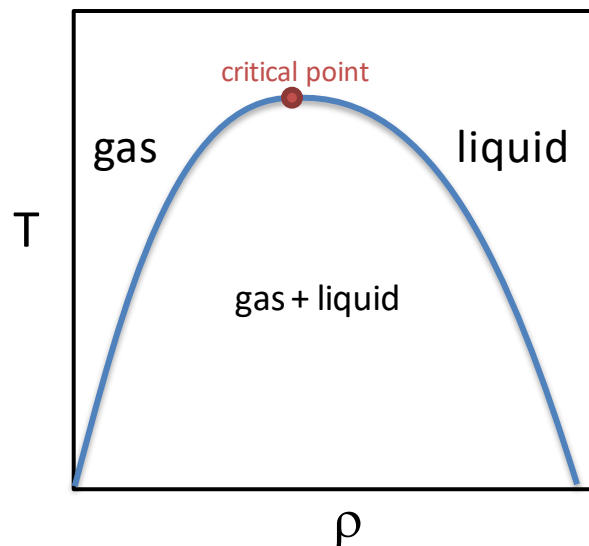
$$P_{12}^{\text{acc}} = \min \left[1, \frac{N_B V_A}{(N_A + 1) V_B} e^{-\beta \Delta U_A - \beta \Delta U_B} \right] \quad \text{particle transferred } B \rightarrow A$$

The three kinds of moves are depicted in the following diagram, taken from [Panagiotopoulos, J. Phys.: Condens. Matter 12, R25 (2000)]:



Running a Gibbs ensemble simulation

Consider liquid-vapor equilibrium. The phase diagram in the $T - \rho$ plane might look something like:



If the temperature and total system density [$\rho = (N_A + N_B)/(V_A + V_B)$] is within the two-phase region, a Gibbs ensemble simulation will have the following behavior: the volume and particle numbers in each simulation box will adjust so that one of the boxes will correspond to the liquid phase and one to the gas phase. That is, the two subsystems will, in a sense, “phase-separate” into gas and liquid phases, although there is no explicit interface between them. The reason for

this behavior is that the phase-separation ultimately minimizes the free energy of the overall system.

Gibbs ensemble simulations are extremely convenient because they enable one to map out very quickly the vapor-liquid phase diagram:

1. Choose a temperature of interest T .
2. Pick an overall density ρ estimated to be somewhere deep within the two-phase region. Pick a system size, either V or N , and then compute the value of the other variable using the density.
3. Initialize the two simulation boxes with the same volume and number of particles.
4. Equilibrate the simulations by performing a long run with particle displacements, volume exchanges, and particle exchanges.
5. At the end of equilibration, one box should be the liquid phase and one should be the gas phase.
6. Use a production period to compute the average properties of each phase, such as $\langle V_i \rangle$, $\langle N_i \rangle$, $\langle \rho_i \rangle$, and $\langle U_i \rangle$, where $i = G$ or L .

Notice that this approach also enables us to compute latent heats, since:

$$\begin{aligned} h_i &\equiv \frac{H_i}{N_i} \\ &= \frac{E_i + P_i V_i}{N_i} \\ &= \frac{\frac{3}{2} N_i k_B T + U_i + P_i V_i}{N_i} \end{aligned}$$

All of the terms on the RHS of this equation can be computed from simulation averages.

While the Gibbs ensemble is extremely attractive for computing phase equilibria, it has some notable limitations. It tends to perform poorly for:

- low temperatures – In such cases, the liquid phase is very dense and particle exchange moves that insert a molecule into the liquid phase tend not to be accepted.
- near the critical point – Here, density fluctuations become very large and it becomes difficult to distinguish which simulation box corresponds to the gas versus the liquid phase. That is, both simulation boxes begin to explore both phases.

- for structured phases – Crystalline solids are difficult to treat using this approach, due to the particle insertion moves.

In later lectures, we will examine a much more robust and statistically optimal method for computing phase equilibrium curves based on histograms. While the Gibbs ensemble is still widely used to study phase equilibria, histograms are now the state-of-the-art approach.

Microcanonical Monte Carlo

In the microcanonical ensemble, one specifies all extensive variables explicitly, E, V, N . The microcanonical partition function is simply the density of states,

$$\Omega(E, V, N) = \frac{1}{h^{3N} N!} \int \delta[H(\mathbf{p}^N, \mathbf{r}^N) - E] d\mathbf{r}^N d\mathbf{p}^N$$

Here, the delta indicates a Dirac delta function. Similar to the canonical ensemble, one can perform the integration over momenta space analytically. This requires the mathematics of delta functions, but the result is fairly straightforward. For a monatomic system,

$$\begin{aligned} \Omega(E, V, N) &= \frac{1}{h^{3N} N!} \int \delta[K(\mathbf{p}^N) + U(\mathbf{r}^N) - E] d\mathbf{r}^N d\mathbf{p}^N \\ &= \frac{1}{h^{3N} N!} \int \delta \left[\frac{1}{2m} \sum_i \mathbf{p}_i^2 + U(\mathbf{r}^N) - E \right] d\mathbf{r}^N d\mathbf{p}^N \end{aligned}$$

To evaluate the integral over momenta, one makes the change of variables $P^2 = \sum_i \mathbf{p}_i^2$ and switches momentum space to $3N$ hyperspherical coordinates. We will not examine the mathematics of this transformation in detail. The result is simply given:

$$\Omega(E, V, N) = \frac{1}{\Gamma(3N/2) N!} \left(\frac{\sqrt{2\pi m}}{h} \right)^{3N} \int [E - U(\mathbf{r}^N)]^{\frac{3N-2}{2}} \theta[E - U(\mathbf{r}^N)] d\mathbf{r}^N$$

Here, $\theta(x)$ gives the Heavyside step function, $\theta(x \geq 0) = 1$ and $\theta(x < 0) = 0$. This function prevents imaginary numbers from appearing in the integral due to the term to its left, which entails a square-root operation. The function $\Gamma(3N/2)$ is the gamma function, and it returns a prefactor.

This expression for the partition function shows that the configurational probability distribution for the microcanonical ensemble is given by

$$\wp(\mathbf{r}^N) \propto [E - U(\mathbf{r}^N)]^{\frac{3N-2}{2}} \theta[E - U(\mathbf{r}^N)]$$

Using this fact, we can devise a Monte Carlo simulation that performs simulations in the micro-canonical ensemble:

1. Specify E, V, N .
2. Start with an initial configuration for which $U(\mathbf{r}^N) < E$. If the potential energy of the initial configuration is greater than E , it has zero probability in the ensemble because there is no possible selection of momenta such that $U + K = E$ and $K \geq 0$ always.
3. Perform single-particle displacement moves. The acceptance criterion is follows from above:

$$P_{12}^{\text{acc}} = \min \left[1, \left(\frac{E - U_2}{E - U_1} \right)^{\frac{3N-2}{2}} \theta(E - U_2) \right]$$

In other words, moves that would take the system to $U_2 > E$ are outright rejected.

MC simulations in the microcanonical ensemble are not frequently performed, but there may be particular cases where this approach is desirable (for example, to compute reaction pathway free energies in the microcanonical ensemble).