

Chapter 5

Monte Carlo Simulations in Various Ensembles

In a conventional Molecular Dynamics simulation, the total energy E and the total linear momentum P are constants of motion. Hence, Molecular Dynamics simulations measure (time) averages in an ensemble that is very similar to the microcanonical (see [78]); namely, the constant-NVE- P ensemble. In contrast, a conventional Monte Carlo simulation probes the canonical (i.e., constant-NVT) ensemble. The fact that these ensembles are different leads to observable differences in the statistical averages computed in Molecular Dynamics and Monte Carlo simulations. Most of these differences disappear in the thermodynamic limit and are already relatively small for systems of a few hundred particles. However, the choice of ensemble does make a difference when computing the mean-square value of fluctuations in thermodynamic quantities. Fortunately, techniques exist to relate fluctuations in different ensembles [73]. Moreover, nowadays it is common practice to carry out Molecular Dynamics simulations in ensembles other than the microcanonical. In particular, it is possible to do Molecular Dynamics at constant pressure, at constant stress, and at constant temperature (see Chapter 6). The choice of ensembles for Monte Carlo simulations is even wider: isobaric-isothermal, constant-stress-isothermal, grand-canonical (i.e., constant- μ VT), and even microcanonical [79–84]. A more recent addition to this list is a Monte Carlo method that employs the Gibbs-ensemble technique [85], which was developed to study phase coexistence in moderately dense (multi component) fluids. The Gibbs-ensemble method is discussed in detail in Chapter 8.

As explained in section 3.1 the principal idea of importance sampling is to use a Monte Carlo procedure to generate a random walk in those regions of phase space that have an important contribution to the ensemble aver-

ages. The acceptance rules are chosen such that these configurations occur with a frequency prescribed by the desired probability distribution. In section 3.1 it is shown that such a procedure indeed yields the correct distribution of configurations. Essential in the demonstration that our Monte Carlo scheme samples the desired distribution is the condition of detailed balance. To be more precise, detailed balance, in fact, is too strong a condition, but if detailed balance is obeyed we are *guaranteed* to have a correct sampling scheme. It may very well be possible that a scheme that does not obey detailed balance still samples the correct distribution. In a Monte Carlo scheme errors are easily introduced, so one should be extremely careful. We will give some examples where we can show that detailed balance is not obeyed and the results show systematic errors. We have found that we could demonstrate that detailed balance was not obeyed in all cases where we observed strange results.

5.1 General Approach

In the following sections, we will use the following procedure to demonstrate the validity of our Monte Carlo algorithms:

1. Decide which distribution we want to sample. This distribution, denoted \mathcal{N} , will depend on the details of the ensemble.
2. Impose the condition of detailed balance,

$$K(o \rightarrow n) = K(n \rightarrow o), \quad (5.1.1)$$

where $K(o \rightarrow n)$ is the flow of configuration o to n . This flow is given by the product of the probability of being in configuration o , the probability of generating configuration n , and the probability of accepting this move,

$$K(o \rightarrow n) = \mathcal{N}(o) \times \alpha(o \rightarrow n) \times \text{acc}(o \rightarrow n). \quad (5.1.2)$$

3. Determine the probabilities of generating a particular configuration.
4. Derive the condition which needs to be fulfilled by the acceptance rules.

5.2 Canonical Ensemble

It is instructive to apply the preceding recipe to the ordinary Metropolis scheme. In the canonical ensemble, the number of particles, temperature, and volume are constant (see Figure 5.1). The partition function is

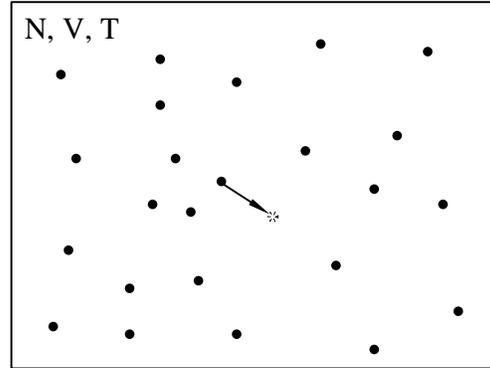


Figure 5.1: Canonical ensemble. The number of particles, volume, and temperature are constant. Shown is a Monte Carlo move in which a particle is displaced.

$$Q(N, V, T) \equiv \frac{1}{\Lambda^{3N} N!} \int d\mathbf{r}^N \exp[-\beta U(\mathbf{r}^N)], \quad (5.2.1)$$

where $\Lambda = \sqrt{h^2/(2\pi m k_B T)}$ is the thermal de Broglie wavelength. From the partition function it follows that the probability of finding configuration \mathbf{r}^N is given by distribution is

$$\mathcal{N}(\mathbf{r}^N) \propto \exp[-\beta U(\mathbf{r}^N)]. \quad (5.2.2)$$

Equations (5.2.1) and (5.2.2) are the basic equations for a simulation in the canonical ensemble.

5.2.1 Monte Carlo Simulations

In the canonical ensemble, we have to sample distribution (5.2.2). This can be done using the following scheme:

1. Select a particle at random and calculate the energy of this configuration $U(o)$.
2. Give this particle a random displacement (see Figure 5.1),

$$\mathbf{r}(o) \rightarrow \mathbf{r}(o) + \Delta(\text{Ranf} - 0.5),$$

where $\Delta/2$ is the maximum displacement. The value of Δ should be chosen such that the sampling scheme is optimal (see section 3.3). The new configuration is denoted n and its energy $U(n)$.

3. The move is accepted with a probability (see equation (3.1.17))

$$\text{acc}(o \rightarrow n) = \min(1, \exp\{-\beta[\mathcal{U}(n) - \mathcal{U}(o)]\}). \quad (5.2.3)$$

If rejected, the old configuration is kept.

An implementation of this basic Metropolis scheme is shown in Section 3.2 (Algorithms 1 and 2).

5.2.2 Justification of the Algorithm

The probability of generating a particular configuration is constant and independent of the conformation of the system

$$\alpha(o \rightarrow n) = \alpha(n \rightarrow o) = \alpha.$$

Substitution of this equation in the condition of detailed balance (5.1.1) and substitution of the desired distribution (5.2.2) gives as condition for the acceptance rules

$$\frac{\text{acc}(o \rightarrow n)}{\text{acc}(n \rightarrow o)} = \exp\{-\beta[\mathcal{U}(n) - \mathcal{U}(o)]\}. \quad (5.2.4)$$

It is straightforward to demonstrate that acceptance rule (5.2.3) obeys this condition.

5.3 Microcanonical Monte Carlo

Most experimental observations are performed at constant N, P, T ; sometimes at constant μ, V, T ; and occasionally at constant N, V, T . Experiments at constant N, V, E are very rare, to say the least. Under what circumstances, then, would anyone wish to perform Monte Carlo simulations at constant N, V , and E ? We suppose that, if you are interested in the simulation of dense liquids or solids, the answer would be “hardly ever”. Still there are situations where a microcanonical Monte Carlo method, first suggested by Creutz [84], may be of use. In particular, you might be worried that a poor random-number generator may introduce a bias in the sampling of the Boltzmann distribution or in the unlikely case that the exponentiation of the Boltzmann factor $\exp\{-\beta[\mathcal{U}(n) - \mathcal{U}(o)]\}$ may account for a nonnegligible fraction of the computing time.

The microcanonical Monte Carlo method uses no random numbers to determine the acceptance of a move. Rather, it uses the following procedure. We start with the system in a configuration \mathbf{q}^N . Denote the potential energy for this state by $\mathcal{U}(\mathbf{q}^N)$. We now fix the total energy of the system at a value $E > \mathcal{U}$. To this end, we introduce an additional degree of freedom that carries the remainder of the energy of the system: $E_D = E - \mathcal{U}$. E_D must always be nonnegative. Now we start our Monte Carlo run.

1. After each trial move we compute the change in potential energy of the system,

$$\Delta u = u(\mathbf{q}'^N) - u(\mathbf{q}^N).$$

2. If $\Delta u < 0$, we accept the move and increase the energy carried by the demon by $|\Delta u|$. If $\Delta u > 0$, we test if the demon carries enough energy to make up the difference. Otherwise, we reject the trial move.

Note that no random numbers were used in this decision. Using elementary statistical mechanics it is easy to see that, after equilibration, the probability density to find the demon with an energy E_D is given by the Boltzmann distribution:

$$\mathcal{N}(E_D) = (k_B T)^{-1} \exp(-E_D/k_B T).$$

Hence, the demon acts as a thermometer. Note that this method does not really simulate the microcanonical ensemble. What is kept (almost) constant is the total potential energy. We can, however, mimic the real N, V, E ensemble by introducing a demon for every quadratic term in the kinetic energy. We then apply the same rules as before, randomly selecting a demon to pay or accept the potential energy change for every trial move.

Microcanonical Monte Carlo is rarely, if ever, used to simulate molecular systems.

5.4 Isobaric-Isothermal Ensemble

The isobaric-isothermal (constant-NPT) ensemble is widely used in Monte Carlo simulations. This is not surprising because most real experiments are also carried out under conditions of controlled pressure and temperature. Moreover, constant-NPT simulations can be used to measure the equation of state of a model system even if the virial expression for the pressure cannot be readily evaluated. This may be the case, for instance, for certain models of nonspherical hard-core molecules, but also for the increasingly important class of models where the (nonpairwise additive) potential energy function is computed numerically for each new configuration. Finally, it is often convenient to use constant-NPT Monte Carlo to simulate systems in the vicinity of a first-order phase transition, because at constant pressure the system is free (given enough time, of course) to transform completely into the state of lowest (Gibbs) free energy, whereas in a constant-NVT simulation the system may be kept at a density where it would like to phase separate into two bulk phases of different density but is prevented from doing so by finite-size effects.

Monte Carlo simulations at constant pressure were first described by Wood [79] in the context of a simulation study of two-dimensional hard

disks. Although the method introduced by Wood is very elegant, it is not readily applicable to systems with arbitrary continuous potentials. McDonald [80] was the first to apply constant-NPT simulations to a system with continuous intermolecular forces (a Lennard-Jones mixture), and the constant-pressure method of McDonald is now being used almost universally and that is discussed next.

5.4.1 Statistical Mechanical Basis

We will derive the basic equations of constant-pressure Monte Carlo in a way that may appear unnecessarily complicated. However, this derivation has the advantage that the same framework can be used to introduce some of the other non-NVT Monte Carlo methods to be discussed later. For the sake of convenience we shall initially assume that we are dealing with a system of N identical atoms. The partition function for this system is given by

$$Q(N, V, T) = \frac{1}{\Lambda^{3N} N!} \int_0^L \cdots \int_0^L \mathbf{dr}^N \exp[-\beta \mathcal{U}(\mathbf{r}^N)]. \quad (5.4.1)$$

It is convenient to rewrite equation (5.4.1) in a slightly different way. We have assumed that the system is contained in a cubic box with diameter $L = V^{1/3}$. We now define scaled coordinates \mathbf{s}^N by

$$\mathbf{r}_i = L \mathbf{s}_i \quad \text{for } i = 1, 2, \dots, N. \quad (5.4.2)$$

If we now insert these scaled coordinates in equation (5.4.1), we obtain

$$Q(N, V, T) = \frac{V^N}{\Lambda^{3N} N!} \int_0^1 \cdots \int_0^1 \mathbf{ds}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]. \quad (5.4.3)$$

In equation (5.4.3), we have written $\mathcal{U}(\mathbf{s}^N; L)$ to indicate that \mathcal{U} depends on the real rather than the scaled distances between the particles. The expression for the Helmholtz free energy of the system is

$$\begin{aligned} F(N, V, T) &= -k_B T \ln Q \\ &= -k_B T \ln \left(\frac{V^N}{\Lambda^{3N} N!} \right) - k_B T \ln \int \mathbf{ds}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)] \\ &= F^{\text{id}}(N, V, T) + F^{\text{ex}}(N, V, T). \end{aligned} \quad (5.4.4)$$

In the last line of this equation we have identified the two contributions to the Helmholtz free energy on the previous line as the ideal gas expression plus an excess part. Let us now assume that the system is separated by a piston¹ from an ideal gas reservoir (see Figure 5.2). The total volume of the

¹Actually, there is no need to assume a real piston. The systems with volume V and $V_0 - V$ may both be isolated systems subject to their individual (periodic) boundary conditions. The only constraint that we impose is that the sum of the volumes of the two systems equals V_0 .

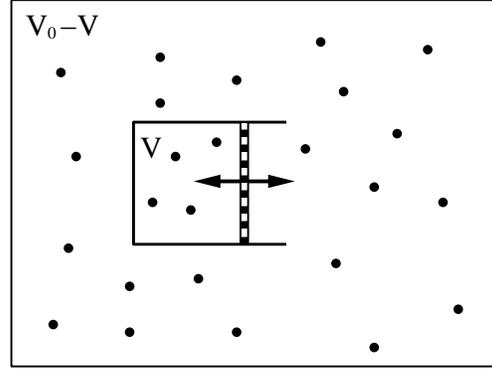


Figure 5.2: Ideal gas (m particles, volume $V_0 - V$) can exchange volume with an N -particle system (volume V).

system plus reservoir is fixed at a value V_0 . The total number of particles is M . Hence, the volume accessible to the $M - N$ ideal gas molecules is $V_0 - V$. The partition function of the total system is simply the product of the partition functions of the constituent subsystems:

$$Q(N, M, V, V_0, T) = \frac{V^N (V_0 - V)^{M-N}}{\Lambda^3 M N! (M - N)!} \int \mathbf{ds}^{M-N} \int \mathbf{ds}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]. \quad (5.4.5)$$

Note that the integral over the \mathbf{s}^{M-N} scaled coordinates of the ideal gas yields simply 1. For the sake of compactness, we have assumed that the thermal wavelength of the ideal gas molecules is also equal to Λ . The total free energy of this combined system is $F^{\text{tot}} = -k_B T \ln Q(N, M, V, V_0, T)$. Now let us assume that the piston between the two subsystems is free to move, so that the volume V of the N -particle subsystem can fluctuate. Of course, the most probable value of V will be the one that minimizes the free energy of the combined system. The probability density $\mathcal{N}(V)$ that the N -particle subsystem has a volume V is given by

$$\mathcal{N}(V) = \frac{V^N (V_0 - V)^{M-N} \int \mathbf{ds}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]}{\int_0^{V_0} dV' V'^N (V_0 - V')^{M-N} \int \mathbf{ds}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L')]} \quad (5.4.6)$$

We now consider the limit that the size of the reservoir tends to infinity ($V_0 \rightarrow \infty$, $M \rightarrow \infty$, $(M - N)/V_0 \rightarrow \rho$). In that limit, a small volume change of the small system does not change the pressure P of the large system. In other words, the large system works as a manostat for the small system. In that case, we can simplify equations (5.4.5) and (5.4.6). Note that in the limit $V/V_0 \rightarrow 0$, we can write

$$(V_0 - V)^{M-N} = V_0^{M-N} [1 - (V/V_0)]^{M-N} \rightarrow V_0^{M-N} \exp(-mV/V_0).$$

Note that for $M - N \rightarrow \infty$, $\exp(-(M - N)V/V_0) \rightarrow \exp(-\rho V)$. But, as the reservoir contains an ideal gas, ρ can be written as βP . With these substitutions, the combined partition function (5.4.5) can be written as

$$Q(N, P, T) \equiv \frac{\beta P}{\Lambda^{3N} N!} \int dV V^N \exp(-\beta P V) \int d\mathbf{s}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)], \quad (5.4.7)$$

where we have included a factor βP to make $Q(N, P, T)$ dimensionless. This gives, for equation (5.4.6);

$$\mathcal{N}_{N, P, T}(V) = \frac{V^N \exp(-\beta P V) \int d\mathbf{s}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]}{\int_0^{V_0} dV' V'^N \exp(-\beta P V') \int d\mathbf{s}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]}. \quad (5.4.8)$$

In the same limit, the difference in free energy between the combined system and the ideal gas system in the absence of the N -particle subsystem is the well-known Gibbs free energy G :

$$G(N, P, T) = -k_B T \ln Q(N, P, T). \quad (5.4.9)$$

Equation (5.4.8) is the starting point for constant-NPT Monte Carlo simulations. The idea is that the probability density to find the small system in a particular configuration of the N atoms (as specified by \mathbf{s}^N) at a given volume V is given by:

$$\begin{aligned} \mathcal{N}(V; \mathbf{s}^N) &\propto V^N \exp(-\beta P V) \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)] \\ &= \exp\{-\beta[\mathcal{U}(\mathbf{s}^N, V) + P V - N \beta^{-1} \ln V]\}. \end{aligned} \quad (5.4.10)$$

We can now carry out Metropolis sampling on the reduced coordinates \mathbf{s}^N and the volume V .

In the constant-NPT Monte Carlo method, V is simply treated as an additional coordinate, and trial moves in V must satisfy the same rules as trial moves in \mathbf{s} (in particular, we should maintain the symmetry of the underlying Markov chain). Let us assume that our trial moves consist of an attempted change of the volume from V to $V' = V + \Delta V$, where ΔV is a random number uniformly distributed between over the interval $[-\Delta V_{\max}, +\Delta V_{\max}]$. In the Metropolis scheme such a random, volume changing move will be accepted with the probability

$$\begin{aligned} \text{acc}(o \rightarrow n) &= \min(1, \exp\{-\beta[\mathcal{U}(\mathbf{s}^N, V') - \mathcal{U}(\mathbf{s}^N, V) \\ &\quad + P(V' - V) - N \beta^{-1} \ln(V'/V)]\}). \end{aligned} \quad (5.4.11)$$

Instead of attempting random changes in the volume itself, one might construct trial moves in the box length L [80] or in the logarithm of the volume [86]. Such trial moves are equally legitimate, as long as the symmetry of the underlying Markov chain is maintained. However, such alternative

schemes result in a slightly different form for equation (5.4.11). The partition function (5.4.7) can be rewritten as

$$Q(N, P, T) = \frac{\beta P}{\Lambda^{3N} N!} \int d(\ln V) V^{N+1} \exp(-\beta P V) \int d\mathbf{s}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]. \quad (5.4.12)$$

This equation shows that, if we perform a random walk in $\ln V$, the probability of finding volume V is given by

$$\mathcal{N}(V; \mathbf{s}^N) \propto V^{N+1} \exp(-\beta P V) \exp[-\beta \mathcal{U}(\mathbf{s}^N; L)]. \quad (5.4.13)$$

This distribution can be sampled with the following acceptance rule:

$$\begin{aligned} \text{acc}(o \rightarrow n) = & \min(1, \exp[-\beta[\mathcal{U}(\mathbf{s}^N, V') - \mathcal{U}(\mathbf{s}^N, V) \\ & + P(V' - V) - (N+1)\beta^{-1} \ln(V'/V)])]. \end{aligned} \quad (5.4.14)$$

5.4.2 Monte Carlo Simulations

The frequency with which trial moves in the volume should be attempted is dependent on the efficiency with which volume space is sampled. If, as before, we use as our criterion of efficiency

$$\frac{\text{sum of squares of accepted volume changes}}{t_{\text{CPU}}},$$

then it is obvious that the frequency with which we attempt moves depends on their cost. In general, a volume trial move will require that we recompute all intermolecular interactions. It therefore is comparable in cost to carrying out N trial moves on the molecular positions. In such cases it is common practice to perform one volume trial move for every cycle of positional trial moves. Note that, to guarantee the symmetry of the underlying Markov chain, volume moves should not be attempted periodically after a fixed number (say N) positional trial moves. Rather, at every step there should be a probability $1/N$ to attempt a volume move instead of a particle move. The criteria determining the optimal acceptance of volume moves are no different than those for particle moves.

In one class of potential energy functions, volume trial moves are very cheap; namely, those for which the total interaction energy can be written as a sum of powers of the interatomic distances:

$$\begin{aligned} u_n &= \sum_{i < j} \epsilon(\sigma/r_{ij})^n \\ &= \sum_{i < j} \epsilon[\sigma/(Ls_{ij})]^n, \end{aligned} \quad (5.4.15)$$

or, possibly, a linear combination of such sums (the famous Lennard-Jones potential is an example of the latter category). Note that U_n in equation (5.4.15) changes in a trivial way if the volume is modified such that the linear dimensions of the system change for L to L' :

$$U_n(L') = \left(\frac{L}{L'}\right)^n U_n(L). \quad (5.4.16)$$

Clearly, in this case, computing the probability of acceptance of a volume-changing trial move is extremely cheap. Hence such trial moves may be attempted with high frequency, for example, as frequent as particle moves. One should be very careful when using the scaling property (5.4.16) if at the same time one uses a cutoff (say r_c) on the range of the potential. Use of equation (5.4.16) implicitly assumes that the cutoff radius r_c scales with L , such that $r'_c = r_c(L'/L)$. The corresponding tail correction to the potential (and the virial) should also be recomputed to take into account both the different cutoff radius and the different density of the system.

Algorithms 2, 10 and 11 show the basic structure of a simulation in the NPT-ensemble.

Finally, it is always useful to compute the virial pressure during a constant pressure simulation. On average, the virial pressure should always be equal to the applied pressure. This is easy to prove as follows. First of all, note that the virial pressure $P_v(V)$ of an N -particle system at volume V is equal to

$$P_v(V) = - \left(\frac{\partial F}{\partial V} \right)_{NT}. \quad (5.4.17)$$

In an isothermal-isobaric ensemble, the probability-density $\mathcal{P}(V)$ of finding the system with volume V is equal to $(\exp[-\beta(F(V) + PV)]) / Q(\text{NPT})$, where

$$Q(\text{NPT}) \equiv \beta P \int dV \exp[-\beta(F(V) + PV)].$$

Let us now compute the average value of the virial pressure:

$$\begin{aligned} \langle P_v \rangle &= \frac{\beta P}{Q(\text{NPT})} \int dV (\partial F(V) / \partial V) \exp[-\beta(F(V) + PV)] \\ &= \frac{\beta P}{Q(\text{NPT})} \int dV \beta^{-1} (\partial \exp[-\beta F(V)] / \partial V) \exp(-\beta P V) \\ &= \frac{\beta P}{Q(\text{NPT})} \int dV P \exp[-\beta(F(V) + PV)] \\ &= P. \end{aligned} \quad (5.4.18)$$

The third line in this equation follows from partial integration.

Algorithm 10 (Basic NPT-Ensemble Simulation)

PROGRAM mc_npt	basic NPT-ensemble simulation
do icycl=1,ncycl	perform ncycl MC cycles
ran=ranf()*(npart+1)+1	
if (ran.le.npart) then	
call mcmove	perform particle displacement
else	
call mcvol	perform volume change
endif	
if (mod(icycl,nsamp).eq.0)	
+ call sample	sample averages
enddo	
end	

Comments to this algorithm:

1. *This algorithm ensures that, after each MC step, detailed balance is obeyed and that per cycle we perform (on average) n_{part} attempts to displace particles and one attempt to change the volume.*
2. *Subroutine mcmove attempts to displace a randomly selected particle (Algorithm 2), and subroutine mcvol attempts to change the volume (Algorithm 11), and subroutine sample updates ensemble averages every n_{sampl} cycle.*

Thus far we have limited our discussion of Monte Carlo at constant pressure to pure, atomic systems. Extension of the technique to mixtures is straightforward. The method is also easily applicable to molecular systems. However, in the latter case, it is crucial to note that only the center-of-mass positions of the molecules should be scaled in a volume move, never the relative positions of the constituent atoms in the molecule. This has one practical consequence; namely, that the simple scaling relation (5.4.16) can never be used in molecular systems with site-site interactions. The reason is that, even if the center-of-mass separations between molecules scale as a simple power of the system size, the site-site separations do not.

5.4.3 Applications

Case Study 7 (Equation of State of the Lennard-Jones Fluid)

Simulations at constant pressure can be used also to determine the equa-

Algorithm 11 (Attempt to Change the Volume)

SUBROUTINE mcvol	attempts to change the volume
call toterg(box,eno)	total energy old conf.
vo=box**3	determine old volume
lnvn=log(vo)+(ranf()-0.5)*vmax	perform random walk in $\ln V$
vn=exp(lnvn)	
boxn=vn**(1/3)	new box length
do i=1,npart	
x(i)=x(i)*boxn/box	rescale center of mass
enddo	
call toterg(boxn,enn)	total energy new conf.
arg=-beta*((enn-eno)+p*(vn-vo)	
+ -(npart+1)*log(vn/vo)/beta)	appropriate weight function!
if (ranf().gt.exp(arg)) then	acceptance rule (5.2.3)
do i=1,npart	REJECTED
x(i)=x(i)*box/boxn	restore the old positions
enddo	
endif	
return	
end	

Comments to this algorithm:

1. A random walk in $\ln V$ is performed using acceptance rule (5.4.14).
2. The subroutine toterg calculates the total energy. Usually the energy of the old configuration is known, therefore this subroutine is called only once.

tion of state of a pure component. In such as simulation the density is determined as a function of the applied pressure and temperature. Figure 5.3 shows that, for the Lennard-Jones fluid, the results of an N,P,T simulation compare very well with those obtained in Case Study 1. In simulations of models of real molecules one would like to know whether at atmospheric conditions the model fluid has the same density as the real fluid. One would need to perform several N,V,T simulations to determine the density at which the pressure is approximately 1 atm. In an N,P,T simulation one would obtain this result in a single simulation. Furthermore, 1 atm. is a relatively low pressure, and one would need long simulations to determine the pressure from an N,V,T simulation, whereas the density in general is determined accurately from an N,P,T simulation.

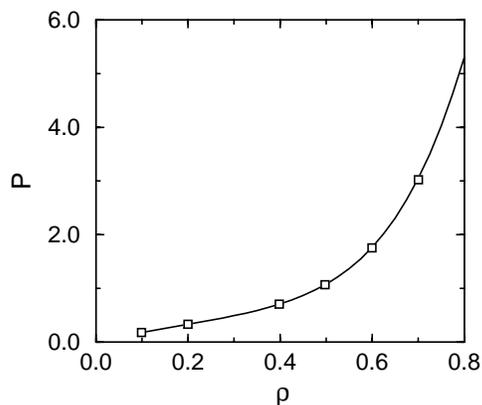


Figure 5.3: Equation of state of the Lennard-Jones fluid as obtained from N,P,T simulations; isotherms at $T = 2.0$. The solid line is the equation of state of Johnson *et al.* [55] and the squares are the results from the simulations ($N = 108$).

Case Study 8 (Phase Equilibria from Constant Pressure Simulations)

In Case Studies 1 and 7 N,V,T or N,P,T simulations are used to determine the equation of state of a pure component. If these equation of state data are fitted to an analytical equation of state (for example, the van der Waals equation of state or more sophisticated forms of this equation), the vapor-liquid coexistence curve can be determined from Maxwell's equal area construction. Although this way of determining a coexistence curve is guaranteed to work for all systems, it requires many simulations and, therefore, is a rather cumbersome route. Alternative routes have been developed to determine vapor-liquid coexistence from a single simulation. In this case study we investigate one of them: zero pressure simulation.

A zero pressure simulation provides a quick (and dirty) way to obtain an estimate of the liquid coexistence density. If we perform a simulation at zero pressure and start with a density greater than the liquid density, the average density obtained from a simulation that is not too long will be close to the coexistence density. Such a simulation should not be too long because the probability exists that the system will undergo a large fluctuation in density. If this fluctuation is towards a lower density the system size can become infinitely large, since the equilibrium density that corresponds to zero pressure is exactly zero. Figure 5.4 shows the results for the Lennard-Jones fluid. Not too close to the critical temperature, a reasonable estimate of the liquid density can be obtained via a single simulation. Important to note is that this estimate deviates systematically from the true coexistence densities and

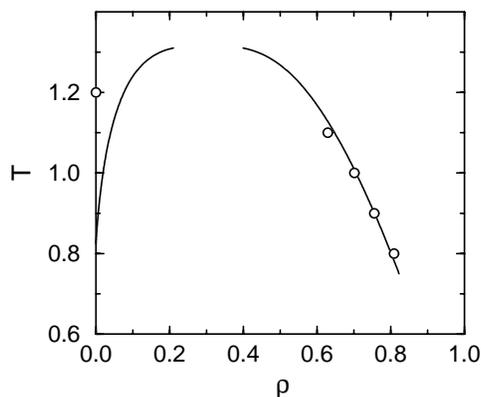


Figure 5.4: Vapor-liquid coexistence of the Lennard-Jones fluid; for each temperature the solid lines give at a given temperature the coexisting gas density (left curve) and the coexisting liquid density (right curve). The circles are the average densities obtained from N,P,T simulations at zero pressure. Important to note is that for $T = 1.2$ the zero pressure method fails to predict coexistence.

this technique should not be used to determine the coexistence curve. This technique is very useful to obtain a first estimate of the coexistence curve.

In contrast to the zero pressure simulations, the histogram technique does lead to a correct estimate of the coexistence density. This technique is often used in lattice models because it allows for a systematic investigation of the finite-size effects (see, for example, [87–90]).

5.5 Isotension-Isothermal Ensemble

The NPT-MC method is perfectly adequate for homogeneous fluids. However, for inhomogeneous systems, in particular crystalline solids, it may not be sufficient that the simulation box can change size. Often we are interested in the transformation of a crystal from one structure to another or even in the change of the shape of the crystalline unit cell with temperature or with applied stress. In such cases it is essential that the shape of the simulation box has enough freedom to allow for such changes in crystal structure without creating grain boundaries or other highly stressed configurations. This problem was first tackled by Parrinello and Rahman [91, 92], who developed an extension of the constant-pressure Molecular Dynamics technique introduced by Andersen [93]. The extension of the Parrinello-Rahman method to

Monte Carlo simulations is straightforward (actually, the method is quite a bit simpler in Monte Carlo than in Molecular Dynamics).

To our knowledge, the first published account of constant-stress Monte Carlo is a paper by Najafabadi and Yip [81]. At the core of the constant-stress Monte Carlo method lies the transformation from the scaled coordinates \mathbf{s} to the real coordinates \mathbf{q} . If the simulation box is not cubic and not orthorhombic, the transformation between \mathbf{s} and \mathbf{r} is given by a matrix h : $r_\alpha = h_{\alpha\beta} s_\beta$. The volume of the simulation box V is equal to $|\det h|$. Without loss of generality we can choose h to be a symmetric matrix. In the constant-stress Monte Carlo procedure certain moves consist of an attempted change of one or more of the elements of h . Actually, it would be equally realistic (but not completely equivalent) to sample the elements of the metric tensor $G = hh$, where h is the transpose of h . If only hydrostatic external pressure is applied, the constant-stress Monte Carlo method is almost equivalent to constant-pressure Monte Carlo.² Under nonhydrostatic pressure (e.g., uniaxial stress), there is again some freedom of choice in deciding how to apply such deforming stresses. Probably the most elegant method (and the method that reflects most closely the statistical thermodynamics of deformed solids) is to express all external deforming stresses in terms of the so-called thermodynamic tension (see, e.g., [94]).

5.6 Grand-Canonical Ensemble

The ensembles we have discussed so far have the total number of particles imposed. For some systems, however, one would like to obtain information on the average number of particles in a system as a function of the external conditions. For example, in adsorption studies one would like to know the amount of material adsorbed as a function of the pressure and temperature of the reservoir with which the material is in contact. A naive but theoretically valid approach would be to use the Molecular Dynamics technique (microcanonical ensemble) and simulate the experimental situation; an adsorbent in contact with a gas (see Figure 5.5). Such a simulation is possible for only very simple systems. In real experiments, equilibration may take minutes or even several hours depending on the type of gas molecules. These equilibration times would be reflected in a Molecular Dynamics simulation, the difference being that a minute of experimental time takes of the order of 10^9 seconds on a computer. Furthermore, in most cases, we are not interested in the properties of the gas phase, yet a significant amount of

²Except that one should never use the constant-stress method for uniform fluids, because the latter offer no resistance to the deformation of the unit box and very strange (flat, elongated, etc.) box shapes may result. This may have serious consequences because simulations on systems that have shrunk considerably in any one dimension tend to exhibit appreciable finite-size effects.

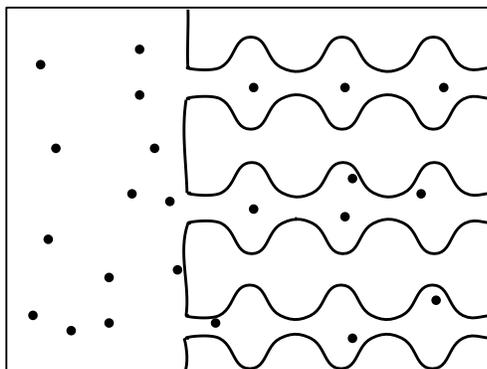


Figure 5.5: Adsorbent (for example a zeolite) in direct contact with a gas

CPU time will be spent on the simulation of this phase, and finally, in such a simulation, there is an interface between the gas phase and the zeolite. In the interfacial region the properties of the system are different from the bulk properties in which we are interested. Since in a simulation the system is relatively small, we have to simulate a very large system to minimize the influence of this interfacial region.³

Most of these problems can be solved by a careful choice of ensembles. For adsorption studies, a natural ensemble to use is the grand-canonical ensemble (or μ, V, T ensemble). In this ensemble, the temperature, volume, and chemical potential are fixed. In the experimental setup, the adsorbed gas is in equilibrium with the gas in the reservoir. The equilibrium conditions are that the temperature and chemical potential of the gas inside and outside the adsorbent must be equal.⁴ The gas that is in contact with the adsorbent can be considered as a reservoir that imposes a temperature and chemical potential on the adsorbed gas (see Figure 5.6). We therefore have to know only the temperature and chemical potential of this reservoir to determine the equilibrium concentration inside the adsorbent. This is exactly what is mimicked in the grand-canonical ensemble: the temperature and chemical potential are imposed and the number of particles is allowed to fluctuate during the simulation. This makes these simulations different from the conventional ensembles, where the number of molecules is fixed.

³Such a simulation, of course, would be appropriate if the interest is in just this region.

⁴Note that the pressure is *not* defined inside the zeolite; therefore, the pressure cannot be an equilibrium quantity. However, the pressure is related to the chemical potential via an equation of state, and it is always possible to calculate the pressure of the gas that corresponds to a given chemical potential and vice versa.

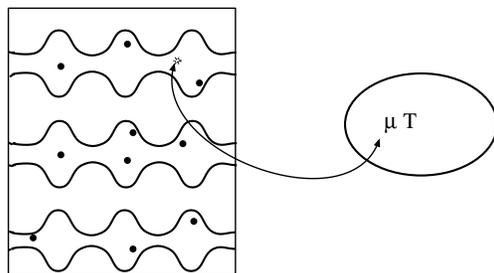


Figure 5.6: Adsorbent in contact with a reservoir that imposes constant chemical potential and temperature by exchanging particles and energy

5.6.1 Statistical Mechanical Basis

In section 3.1.2, we introduced the Metropolis sampling scheme as a method to compute thermal averages of functions $A(\mathbf{r}^N)$ that depend explicitly on the coordinates of the molecules in the N -body system under study. Examples of such mechanical properties are the potential energy or the virial contribution to the pressure. However, the Metropolis method could not be used to determine the integral $\int d\mathbf{r}^N \exp[-\beta U(\mathbf{r}^N)]$ itself. The latter quantity measures the effective volume in configuration space that is accessible to the system. Hence, the original Metropolis scheme could not be used to determine those thermodynamic properties of a system that depend explicitly on the configurational integral. Examples of such thermal properties are the Helmholtz free energy F , the entropy S , and the Gibbs free energy G . However, although the Metropolis method cannot be used to measure, for instance, free energies directly, it can be used to measure the difference in free energy between two possible states of an N -body system. This fact is exploited in the grand-canonical Monte Carlo method first implemented for classical fluids by Norman and Filinov [82], and later extended and improved by a number of other groups [83,95–102]. The basic idea of the grand-canonical Monte Carlo method is explained next.

To understand the statistical mechanical basis for the grand-canonical Monte Carlo technique, let us return to equation (5.4.5) of section 5.4. This equation gives the partition function of a combined system of N interacting particles in volume V and $M - N$ ideal gas molecules in volume $V_0 - V$:

$$Q(N, M, V, V_0, T) = \frac{V^N (V_0 - V)^{M-N}}{\Lambda^{3M} N! (M - N)!} \int d\mathbf{s}^{M-N} \int d\mathbf{s}^N \exp[-\beta U(\mathbf{s}^N)].$$

Now, instead of allowing the two systems to exchange volume, let us see what happens if the systems can also exchange particles (see Figure 5.7). To be more precise, we assume that the molecules in the two subvolumes are

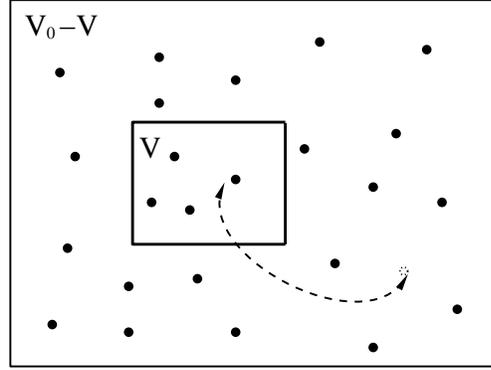


Figure 5.7: Ideal gas ($M - N$ particles, volume $V_0 - V$) can exchange particles with a N -particle system (volume V).

actually identical particles. The only difference is that when they find themselves in volume V , they interact and, when they are in volume $V_0 - V$, they do not. If we transfer a molecule i from a reduced coordinate \mathbf{s}_i in the volume $V_0 - V$ to the same reduced coordinate in volume V , then the potential energy function \mathcal{U} changes from $\mathcal{U}(\mathbf{s}^N)$ to $\mathcal{U}(\mathbf{s}^{N+1})$. The expression for the total partition function of the system, including all possible distributions of the M particles over the two subvolumes is

$$Q(M, V, V_0, T) = \sum_{N=0}^M \frac{V^N (V_0 - V)^{M-N}}{\Lambda^{3M} N! (M - N)!} \int d\mathbf{s}^{M-N} \int d\mathbf{s}^N \exp[-\beta \mathcal{U}(\mathbf{s}^N)]. \quad (5.6.1)$$

Following the approach of section 5.4, we now write the probability density to find a system with $M - N$ particles at reduced coordinates \mathbf{s}^{M-N} in volume $V' \equiv V_0 - V$ and N particles at reduced coordinates \mathbf{s}^N in volume V :

$$\mathcal{N}(\mathbf{s}^M; N) = \frac{V^N V'^{M-N}}{Q(M, V, V', T) \Lambda^{3M} N! (M - N)!} \exp[-\beta \mathcal{U}(\mathbf{s}^N)]. \quad (5.6.2)$$

Let us now consider a trial move in which a particle is transferred from V' to the same scaled coordinate in V . First we should make sure that we construct an underlying Markov that is symmetric. Symmetry, in this case, implies that the a priori probability to move a particle from V' to V should be equal to the a priori probability of the reverse move. The probability of acceptance of a trial move in which we move a particle to or from volume V is determined by the ratio of the corresponding probability densities (5.6.2):

$$\alpha(N \rightarrow N + 1) = \frac{V(M - N)}{V'(N + 1)} \exp(-\beta[\mathcal{U}(\mathbf{s}^{N+1}) - \mathcal{U}(\mathbf{s}^N)]) \quad (5.6.3)$$

$$\alpha(N+1 \rightarrow N) = \frac{V'(N+1)}{V(M-N)} \exp(-\beta[\mathcal{U}(\mathbf{s}^N) - \mathcal{U}(\mathbf{s}^{N+1})]). \quad (5.6.4)$$

Now let us consider the limit that the ideal gas system is very much larger than the interacting system: $M \rightarrow \infty, V' \rightarrow \infty, (M/V') \rightarrow \rho$. Note that for an ideal gas the chemical potential μ is related to the particle density ρ by

$$\mu = k_B T \ln \Lambda^3 \rho.$$

Therefore, in the limit $(M/N) \rightarrow \infty$, the partition function (5.6.1) becomes

$$Q(\mu, V, T) \equiv \sum_{N=0}^{\infty} \frac{\exp(\beta\mu N) V^N}{\Lambda^{3N} N!} \int d\mathbf{s}^N \exp[-\beta\mathcal{U}(\mathbf{s}^N)], \quad (5.6.5)$$

and the corresponding probability density

$$\mathcal{N}_{\mu VT}(\mathbf{s}^N; N) \propto \frac{\exp(\beta\mu N) V^N}{\Lambda^{3N} N!} \exp[-\beta\mathcal{U}(\mathbf{s}^N)]. \quad (5.6.6)$$

Equations (5.6.5) and (5.6.6) are the basic equations for Monte Carlo simulations in the grand-canonical ensemble. Note that, in these equations, all explicit reference to the ideal gas system has disappeared.

5.6.2 Monte Carlo Simulations

In a grand-canonical simulation, we have to sample the distribution (5.6.6). Acceptable trial moves are

1. *Displacement of particles.* A particle is selected at random and given a new conformation (for example in the case of atoms a random displacement). This move is accepted with a probability

$$\text{acc}(s \rightarrow s') = \min \left(1, \exp\{-\beta[\mathcal{U}(\mathbf{s}'^N) - \mathcal{U}(\mathbf{s}^N)]\} \right). \quad (5.6.7)$$

2. *Insertion and removal of particles.* A particle is inserted at a random position or a randomly selected particle is removed. The creation of a particle is accepted with a probability

$$\text{acc}(N \rightarrow N+1) = \min \left[1, \frac{V}{\Lambda^3(N+1)} \exp\{\beta[\mu - \mathcal{U}(N+1) + \mathcal{U}(N)]\} \right] \quad (5.6.8)$$

and the removal of a particle is accepted with a probability

$$\text{acc}(N \rightarrow N-1) = \min \left[1, \frac{\Lambda^3 N}{V} \exp\{-\beta[\mu + \mathcal{U}(N-1) - \mathcal{U}(N)]\} \right]. \quad (5.6.9)$$

Appendix F demonstrates how the chemical potential of the reservoir can be related to the pressure of the reservoir. Algorithm 12 shows the basic structure of a simulation in the grand-canonical ensemble.

Algorithm 12 (Basic Grand-Canonical Ensemble Simulation)

PROGRAM mc_gc	basic μ VT-ensemble simulation
do icycl=1,ncycl	perform ncycl MC cycles
ran=int(ranf()*(npart+nexc))+1	
if (ran.le.npart) then	
call mcmove	displace a particle
else	
call mcexc	exchange a particle with the reservoir
endif	
if (mod(icycl,nsamp).eq.0)	
+ call sample	sample averages
enddo	
end	

Comments to this algorithm:

1. This algorithm ensures that, after each MC step, detailed balance is obeyed. Per cycle we perform on average n_{part} attempts to displace particles and n_{exc} attempts to exchange particles with the reservoir.
2. Subroutine `mcmove` attempts to displace a particle (Algorithm 2), subroutine `mcexc` attempts to exchange a particle with a reservoir (Algorithm 13), and subroutine `sample` samples quantities every n_{samp} cycle.

5.6.3 Justification of the Algorithm

It is instructive to demonstrate that the acceptance rules (5.6.7)–(5.6.9) indeed lead to a sampling of distribution (5.6.6). Consider a move in which we start with a configuration with N particles and move to a configuration with $N + 1$ particles by inserting a particle in the system. Recall that we have to demonstrate that detailed balance is obeyed:

$$K(N \rightarrow N + 1) = K(N + 1 \rightarrow N),$$

with

$$K(N \rightarrow N + 1) = \mathcal{N}(N) \times \alpha(N \rightarrow N + 1) \times \text{acc}(N \rightarrow N + 1).$$

In Algorithm 12 at each Monte Carlo step the probability that an attempt is made to remove a particle is equal to the probability of attempting to add one:

$$\alpha_{\text{gen}}(N \rightarrow N + 1) = \alpha_{\text{gen}}(N + 1 \rightarrow N),$$

Algorithm 13 (Attempt to Exchange a Particle with a Reservoir)

<pre> SUBROUTINE mcexc if (ranf().lt.0.5) then if (npart.eq.0) return o=int(npart*ranf()+1) call ener(x(o),eno) arg=npart*exp(beta*eno) + /(zz*vol) if (ranf().lt.arg) then x(o)=x(npart) npart=npart-1 endif else xn=ranf()*box call ener(xn,enn) arg=zz*vol*exp(-beta*enn) + /(npart+1) if (ranf().lt.arg) then x(npart+1)=xn npart=npart+1 endif endif return end </pre>	<p>attempt to exchange a particle with a reservoir</p> <p>decide to remove or add a particle</p> <p>test whether there is a particle</p> <p>select a particle to be removed</p> <p>energy particle o</p> <p>acceptance rule (5.6.9)</p> <p>accepted: remove particle o</p> <p>new particle at a random position</p> <p>energy new particle</p> <p>acceptance rule (5.6.8)</p> <p>accepted: add new particle</p>
---	---

Comment to this algorithm:

1. We have defined: $zz = \exp(\beta\mu)/\Lambda$. The subroutine `ener` calculates the energy of a particle at a given position.

where the subscript gen refers to the fact that α measures the probability to generate this trial move. Substitution of this equation together with equation (5.6.6) into the condition of detailed balance gives:

$$\begin{aligned}
 \frac{\text{acc}(N \rightarrow N+1)}{\text{acc}(N+1 \rightarrow N)} &= \frac{\exp[\beta\mu(N+1)]V^{N+1} \exp[-\beta\mathcal{U}(\mathbf{s}^{N+1})]}{\Lambda^{3(N+1)}(N+1)!} \\
 &\quad \times \frac{\Lambda^{3N}N! \exp[\beta\mathcal{U}(\mathbf{s}^N)]}{\exp(\beta\mu N)V^N} \\
 &= \frac{\exp(\beta\mu)V}{\Lambda^3(N+1)} \exp\{-\beta[\mathcal{U}(\mathbf{s}^{N+1}) - \mathcal{U}(\mathbf{s}^N)]\}.
 \end{aligned}$$

It is straightforward to show that acceptance rules (5.6.8) and (5.6.9) obey this condition.

The most salient feature of the grand-canonical Monte Carlo technique is that in such simulations the chemical potential μ is imposed, while the number of particles N is a fluctuating quantity. During the simulation we may measure other thermodynamic quantities, such as the pressure P , the average density $\langle\rho\rangle$, or the internal energy $\langle U\rangle$. As we know μ , we can derive all other thermal properties, such as the Helmholtz free energy or the entropy. This may seem surprising. After all, in section 3.1 we stated that Metropolis sampling cannot be used to sample absolute free energies and related quantities. Yet, with grand-canonical Monte Carlo we seem to be doing precisely that. The answer is that, in fact, we do not. What we measure is not an absolute but a relative free energy. In grand-canonical Monte Carlo, we are equating the chemical potential of a molecule in an ideal gas at density ρ (for the ideal gas case we know how to compute μ) and the chemical potential of the same species in an interacting system at density ρ' .

Grand-canonical Monte Carlo works best if the acceptance of trial moves by which particles are added or removed is not too low. For atomic fluids, this condition effectively limits the maximum density at which the method can be used to about twice the critical density. Special tricks are needed to extend the grand-canonical Monte Carlo method to somewhat higher densities [100]. Grand-canonical Monte Carlo is easily implemented for mixtures and inhomogeneous systems, such as fluids near interfaces. In fact, some of the most useful applications of the grand-canonical Monte Carlo method are precisely in these areas of research. Although the grand canonical Monte Carlo technique can be applied to simple models of nonspherical molecules, special techniques are required since the method converges very poorly for all but the smallest polyatomic molecules. In section 13.5 some of these techniques are discussed.

5.6.4 Applications

Case Study 9 (Equation of State of the Lennard-Jones Fluid)

In Case Studies 1 and 7, we used N,V,T simulations and N,P,T simulations to determine the equation of state of the Lennard-Jones fluid. A third way to determine the equation of state is to impose the temperature and chemical potential and calculate the density and pressure. An example of such a calculation is shown in Figure 5.8. This is not a very convenient method since both the pressure and density are fluctuating quantities, hence we will have an error in both quantities, while in the N,P,T ensemble either the pressure or density is imposed and therefore known without a statistical error. Of course, a grand-canonical simulation is useful if we want to have information on the chemical potential of our system.

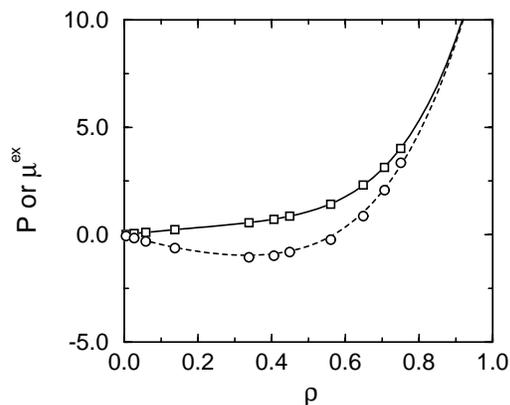


Figure 5.8: Equation of state of the Lennard-Jones fluid; isotherm at $T = 2.0$. The solid line is the equation of state of Johnson *et al.* [55] and the squares are the results from grand-canonical simulations (with volume $V = 250.047$). The dotted line is the excess chemical potential as calculated from the equation of state of Johnson *et al.* and the circles are the results of the simulations.

Example 2 (Adsorption Isotherms of Zeolites)

Zeolites are crystalline inorganic polymers which form a three-dimensional network of micropores (see Figure 5.9). These pores are accessible to various guest molecules. The large internal surface, the thermal stability, and the presence of thousands of acid sites make zeolites an important class of catalytic materials for petrochemical applications. For a rational use of zeolites, it is essential to have a detailed knowledge of the behavior of the adsorbed molecules inside the pores of the zeolites. Since this type of information is very difficult to obtain experimentally, simulations are an attractive alternative. One of the first attempts to study the thermodynamic properties of a molecule adsorbed in a zeolite was made by Stroud *et al.* [103]. Reviews of the various applications of computer simulations of zeolites can be found in [104].

For small adsorbents such as methane or the noble gases, grand-canonical Monte Carlo simulations can be applied to calculate the adsorption isotherms in the various zeolites [105–111]. An example of an adsorption isotherm of methane in the zeolite silicalite is shown in Figure 5.10. These calculations are based on the model of Goodbody *et al.* [107]. The agreement with the experimental data is very good, which shows that for these well characterized systems simulations can give data that are comparable with experiments.

For long-chain alkanes (butane and longer) it is very difficult to perform a

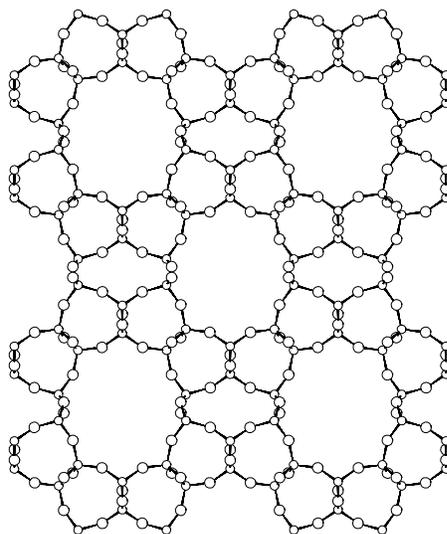


Figure 5.9: Example of a zeolite structure (Theta-1), the pore size is ca. $4.4 \times 5.5 \text{ \AA}$. The Si atoms have four bonds and the O atoms two.

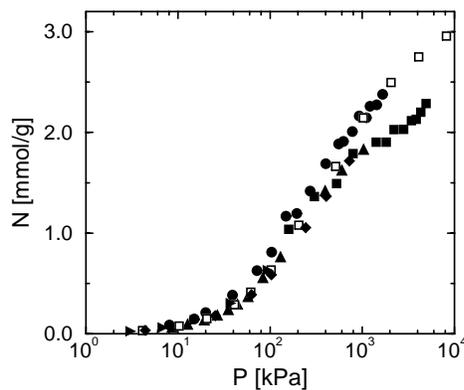


Figure 5.10: Adsorption isotherms of methane in silicalite, showing the amount of methane adsorbed as a function of the external pressure. The black symbols are experimental data (see [112] for details). The open squares are the results of grand-canonical simulations using the model of [107].

successful insertion, in almost all attempts one of the atoms of the molecule will overlap with one of the atoms of the zeolite. As a consequence the number of attempts has to be astronomically large to have a reasonable number of successful exchanges with the reservoir. In Chapter 13 we show how this can problem be solved.