

Monte Carlo Methods

Version 2.23

VESA APAJA

Spring 2005 Johannes Kepler Universität, Linz, Austria
Updated 2008-2009 University of Jyväskylä, Finland

Contents

1	History	4
2	Basic concepts	5
2.1	Sampling a distribution	9
2.2	Some terminology	16
2.3	Central Limit Theorem	18
2.4	Law of Large Numbers	22
2.5	Detailed balance condition	23
3	Markov Chain Monte Carlo	25
3.1	Metropolis–Hastings algorithm	25
4	Error estimation	31
4.1	Some definitions and clarifications	31
4.2	Biased estimators	32
4.3	Unbiased estimators	32
4.4	Derivation of $\sigma_{unbiased}$	32
4.5	Consequenses of biased sampling error	33
4.6	Block averaging	34
4.7	Resampling methods	34
4.7.1	Jackknife	34
4.7.2	Bootstrap	36
5	Variational Monte Carlo	39
6	Diffusion Monte Carlo	42
6.1	Short–time estimates of the Green’s function	43
6.2	DMC algorithms	48
6.3	Mixed and pure estimators	52
6.4	Variance optimisation	56
6.5	Fixed node and released node methods	61
6.6	Excited states	63
6.7	Second order DMC algorithms	65
6.8	DMC Example: the ground state of a He atom	69
7	Path Integral Monte Carlo	73
7.1	First ingredient: products of density matrices	74
7.2	Second ingredient: high- T density matrix	76
7.3	Connection to path integral formulation of quantum mechanics	77
7.4	Link action	79

7.5	Improving the action	79
7.6	Bose Symmetry	80
7.7	Sampling paths	82
7.8	Fermion paths: The sign problem	82
7.9	Liquid helium excitation spectrum	84
8	Thermodynamics	85
8.1	NpT Simulations	85
8.1.1	Pressure	85
8.1.2	Volume Fluctuations	88
8.2	Chemical potential	92
8.3	Thermostats	94
8.4	<i>Ab initio</i> molecular dynamics: Car–Parrinello method	97
9	Appendix	104
9.1	McMillan trial wave function for liquid ^4He	104
9.2	Cusp conditions in Coulomb systems	110

1 History

These lecture notes are mainly about quantum Monte Carlo (QMC) methods, except section 8 which discusses the thermodynamical aspects of simulations. Here are some hallmarks of QMC¹

- 1940's E. Fermi: Schrödinger equation can be written as a diffusion equation; Physical principles of *Diffusion Monte Carlo (DMC)* discovered.
- 1942 A. A. Frost: evaluation of the energy of few simple molecules by picking few representative points in the configuration space; *Variational Monte Carlo (VMC)* is born.
- 1947 N. Metropolis and S. Ulam: Diffusion Monte Carlo developed into an algorithm; "... as suggested by Fermi ...".
- 1950's R. Feynman: Path integral formulation of quantum mechanics; Physical principles of *Path Integral Monte Carlo (PIMC)* discovered
- 1962 M. H. Kalos: Solved the Schrödinger equation in the form of an integral equation: *Greens function Monte Carlo (GFMC)* is born.
- 1964 H. Conroy: Proposed, that points should be picked at random from the configuration space with probability proportional to Ψ_T^2 , the square of the trial wave function. Application to H_2^+ , H^- , HeH^{++} , He, H_2 , and Li - the first application to a fermion system!
- 1965 W. L. McMillan: Energy of liquid 4He ; Applied the Metropolis algorithm to sample points with probability proportional to Ψ_T^2 .
- 1971 R. C. Grim and R. G. Storer: Importance sampling.
- 1975 J. B. Anderson: First application of DMC to electronic systems
- 1986 D. M. Ceperley and E. L. Pollock: PIMC calculation of the properties of liquid 4He at finite temperatures.

Some more recent developments are mentioned later in these lecture notes.

¹After J. B. Anderson in *Quantum Simulations of Complex Many-Body Systems : From Theory to Algorithms*, Editors J. Grotendorst, D. Marx, and A. Muramatsu.

2 Basic concepts

Many-body theorists are interested in systems with many degrees of freedom. Often one encounters multi-dimensional integrals, like the partition function of N classical atoms,

$$Z \propto \int d\mathbf{r}_1 d\mathbf{r}_2 \dots d\mathbf{r}_N \exp \left[-\beta \sum_{i < j} v(r_{ij}) \right] , \quad (1)$$

where $\beta = 1/(k_B T)$ and $v(r_{ij})$ is the particle-particle potential. If one discretises each coordinate with only ten points, and the machine could do 10^7 points per second, would the partition function for twenty particles take 10^{53} seconds, some 10^{34} times the age of the universe (at the moment you start running the program ...).

We *must* evaluate the integrand at some points, otherwise the result wouldn't contain any information of the integrand and would be useless. So how to pick the points? And clever humanoids came up with an answer: If one cannot evaluate the integrand in any decent number of discretisation points per dimension, one might as well pick them at random from the multi-dimensional space.

The Monte Carlo way How could one calculate the area of a unit circle with random numbers? First put a box around the circle, or more economically, a box around the first quadrant. If you take random coordinates x and y uniformly distributed between 0 and 1, mark this $x \in U[0, 1]$ and $y \in U[0, 1]$. The probability that the point (x, y) is inside the circle is

$$P(\text{hit circle}) = \frac{\text{Area of a quarter unit circle}}{\text{Area of a box}} = \frac{\pi}{4} . \quad (2)$$

You know that you've hit the circle, when $x^2 + y^2 \leq 1$, so that's it: start picking up points (x, y) and count how many times you hit the circle out of N tries. This is a very ineffective way of calculating π , and the more space one has that is outside the circle the worse it gets. This is a very simple example of *importance sampling*: The better your random picks hit the relevant region in space (inside the circle), the better the efficiency. After a one night run the result was $\pi \approx 3.141592318254839$, with only 6 correct decimals. By the way, this stone-throwing calculation of π is also a simple test for the (pseudo-)random number generator.

This π -problem was inefficient application of MC mostly because the dimensionality was too low, just two.

Next consider an alternative way to do a one-dimensional integral,

$$I = \int_0^1 dx f(x) . \quad (3)$$

One can estimate this integral as the average of function $f(x)$ in the integration range,

$$I \approx I_N = \frac{1}{N} \sum_{i=1}^N f(x_i) , \quad (4)$$

where x_i are random points evenly distributed in the range $x_i \in [0, 1]$. To estimate the error for arbitrary f we can assume, that it is a random variable; a sensible assumption since we don't want to tailor this for some specific f . The *variance* or the *square of the standard deviation* is then

$$\sigma_I^2 \approx \sigma_{I_N}^2 = \frac{1}{N} \sigma_f^2 \quad (5)$$

$$= \frac{1}{N} \left[\frac{1}{N} \sum_{i=1}^N f(x_i)^2 - \left(\frac{1}{N} \sum_{i=1}^N f(x_i) \right)^2 \right] , \quad (6)$$

and it measures how much the random variable deviates from its mean, *i.e.*, from the estimate of the integral I . The inaccuracy σ_I is hence inversely proportional to the square root of the number of samples,

$$\sigma_I \sim \frac{1}{\sqrt{N}} . \quad (7)$$

Thus the error converges rather slowly, in the trapezoid rule it diminishes as $1/N^2$.

Importance sampling Eq. (6) shows that our new integration method becomes better for flatter functions, it's exact for a constant function. So let's try to make it flatter. We always start from the assumption that a method to generate random numbers $\in U[0, a]$ is already available. Define a normalised weight function w ,

$$\int_0^1 dy w(y) = 1 . \quad (8)$$

This w can be called also a probability density function (PDF). Next we write the original integral in the equivalent form

$$I = \int_0^1 dy w(y) \frac{f(y)}{w(y)} \quad (9)$$

and make a change of variables

$$x(y) = \int_0^y dy' w(y') . \quad (10)$$

Now we have

$$\frac{dx}{dy} = w(y) , \quad (11)$$

and $x(y = 0) = 0$; $x(y = 1) = 1$, so the integral reads

$$I = \int_0^1 dx \frac{f(y(x))}{w(y(x))} . \quad (12)$$

According to the above strategy this is estimated as the average of the argument evaluated at $x_i \in U[0, 1]$:

$$I \approx I_N = \frac{1}{N} \sum_{i=1}^N \frac{f(y(x_i))}{w(y(x_i))} . \quad (13)$$

If we can find a weight function w that approximated f well, then f/w is flatter than f and the variance is smaller. The applicability of this method relies on the assumptions, that

- (i) we know something useful of f and so we can look for a nice weight function w that approximates f
- (ii) we can calculate $y(x)$ for that weight function, *i.e.*, we know how to draw random numbers from the PDF w .

Since $x_i \in U[0, 1]$, points $y_i = y(x_i)$ are distributed as $dx/dy = w(y)$. Therefore more samples are taken from regions of space where w - and by construction also f - is large, and little effort is spent in regions where the weight is small. This is called *importance sampling*. Notice that (i) makes the importance-sampling method problem specific.

A simple example: If I try to hit a student with a chalk blindfolded, then I will hit more probably if I aim away from the blackboard. Or, If I was talking to walls, I could as well choose my importance sampling function to aim at the blackboard...

For one-dimensional integrals the Monte Carlo method is clearly inferior to more deterministic methods. However, if one has to evaluate the integrand N times to calculate a d -dimensional integral, then for each dimension one has $\sim N^{1/d}$ points, whose distance is $h \sim N^{-1/d}$. The error for every integration cell of volume h^d is in the trapezoid rule $\mathcal{O}(h^{d+2})$, so the total error will be N times that, $\mathcal{O}(N^{-2/d})$. So for large d this error is diminishing very slowly with increasing N . On the other hand, *the error in the Monte Carlo method goes down as $1/\sqrt{N}$ regardless of the dimension d* , so already

at $d = 4$ will the trapezoid and the Monte Carlo errors diminish equally fast with increasing N .

Critical slowing down The error in MC decreases as $\sim 1/\sqrt{N}$ for N independent measurements. This means about 100 times more work to get one more significant figure to the result. This is quite normal for the Monte Carlo method, but there is another problem specific factor that may deteriorate your sampled data points.

Samples are usually correlated, the amount of correlation is given by the *autocorrelation function*

$$\begin{aligned}\rho(t) &= \frac{\langle f(0)f(t) \rangle - \langle f(0) \rangle^2}{\langle f(0)^2 \rangle - \langle f(0) \rangle^2} \\ &\sim e^{-t/\tau_{\text{exp}}}\end{aligned}\tag{14}$$

where τ_{exp} is the (exponential) autocorrelation time. A discrete form of the autocorrelation function is ²

$$\rho_m = \frac{\sum_{i=1}^N (x_i - \bar{x})(x_{i+m} - \bar{x})}{\sum_{i=1}^N (x_i - \bar{x})^2} .\tag{15}$$

The integrated autocorrelation time τ_{int}

$$\tau_{\text{int}} = \frac{1}{2} + \sum_{t=1}^{\infty} \rho(t)\tag{16}$$

tells how correlations affect statistical errors. It is normalised so that $\tau_{\text{int}} \approx \tau_{\text{exp}}$ if $\rho \approx \exp(-t/\tau_{\text{exp}})$. It can be shown that for MC results the true error is

$$\text{error in mean} = \sqrt{\frac{2\tau_{\text{int}}}{N}} \sigma ,\tag{17}$$

where σ is the variance of N independent samples.. The bad news is that close to a critical point τ_{int} diverges. This is known as critical slowing down, and it tends to make MC calculations prohibitively slow near phase transitions. A simple example of a phase transition is that in the *Ising model*; We'll look at it more closely in Exercise 1. Near the phase transition large spin regions become correlated. The reason for critical slowing down is clear: If the basic trial step is to choose whether or not to flip a single spin, then these correlated spin regions are hardly affected at all and the phase space is not probed effectively.

²I'm often a bit sloppy and don't emphasise the slight difference between a sample average $\bar{x}$ and an expectation value $\langle x \rangle$, so don't mind the different notation.

2.1 Sampling a distribution

Without importance sampling MC is pretty useless, so we'd better find a way how to generate random numbers that sample a given distribution w either analytically or numerically.

A non-random method Let's say we have a set of points y_i that are distributed as w , and evenly distributed points x_i . Then for $i = 1, \dots, N$

$$x_i = i/N = \int_0^{y_i} dy' w(y') . \quad (18)$$

This is basically a cumulative function of $w(y)$. Hence points y_i , where i has been picked from $0, 1, \dots, N$, approximate the PDF w . They can be solved upon numerical solution of the differential equation $dx/dy = w(y)$ using a simple discretisation,

$$\frac{x_{i+1} - x_i}{y_{i+1} - y_i} = w(y_i) . \quad (19)$$

Since $x_{i+1} - x_i = 1/N$, we obtain the recurrence relation

$$y_{i+1} = y_i + \frac{1}{Nw(y_i)} , \quad (20)$$

where the initial value $y_0 = 0$ can be used.

Method based on inverse functions One can use the basic relation

$$w(y) = \frac{dx}{dy} = \frac{dy^{-1}}{dy} , \quad (21)$$

where $y^{-1} = x$ is the inverse function of y . According to Eq. (10) this is given by the cumulative function of the PDF w , so in this method we need both it and its inverse. Finding inverse functions analytically can be tricky, if not impossible. A numerical method is closer to the spirit now.

A) Discrete distributions

Let's assume we want to sample integers k , where $1 \leq k \leq N$ with probability P_k .

1. Construct the cumulative table of probabilities, $F_k = F_{k-1} + P_k$.
2. Sample $x \in U[0, 1]$, find unique k so that $F_{k-1} < x < F_k$ using, for example, the bisection method. Repeat step 2 as many times as needed.

B) Continuous distributions

This is sometimes called the *mapping method*. Now we want to sample PDF $w(x)$ where x is a continuous variable.

1. Calculate the cumulative function of $w(x)$

$$F(x) = \int_{-\infty}^x dy w(y) \quad (22)$$

either analytically or numerically.

2. Sample $z \in U[0, 1]$, solve $F(x) = z$ for x ; it samples $w(x)$.

Solving $F(x) = z$ for x can always be made quite fast by tabulating $F(x)$.

Rejection method³ Also known as the hit and miss method, or the acceptance–rejection method. This is essentially an application of the idea used earlier to compute π . Let's assume we know that the distribution is limited from above by $w(x) < M$, $x \in [a, b]$. A sequence of random numbers can be generated as follows:

1. Draw $r_1 \in U[a, b]$ and $r_2 \in U[0, M]$ (this picks a random point inside the box that is drawn around $w(x)$, see Fig. 1)
2. If $w(r_1) \geq r_2$, accept $x = r_1$ as the desired random number, else discard both r_1 and r_2 . Go to 1.

Why does this work? Take a look at the identity

$$\begin{aligned} w(x) &= \int_0^M dr_2 \theta(w(x) - r_2) \\ &= \int_a^b dr_1 \int_0^M dr_2 \delta(x - r_1) \theta(w(r_1) - r_2). \end{aligned} \quad (23)$$

The first equation states that integrating step function θ (which is 1 if $r_2 < w(x)$) to $w(x)$ gives just $w(x)$. The second form uses the Dirac delta function to put *another* random variable to the argument. So to sample $w(x)$ we need to integrate Eq. (23) using MC ideas. That is exactly what the above algorithm is doing: the second step in the algorithm just evaluates the argument of the integral (23). Notice, that thanks to the delta function we don't need to solve x from $w(x) = r_2$.

³John von Neumann in late 1950's




Figure 1: Illustration of the rejection method.

Suppose we don't know M and don't want to use a and b either. Then we could sample r_1 from some other yet unspecified PDF $g(r_1)$ and find a condition when to accept this r_1 as x . Clearly

$$\begin{aligned} g(x)h(x) &= g(x) \int_0^\infty dr_2 \theta(h(x) - r_2) \\ &= \int_0^\infty dr_1 g(r_1) \int_0^\infty dr_2 \delta(x - r_1) \theta(h(r_1) - r_2) , \end{aligned} \quad (24)$$

so if $h(x) = w(x)/g(x)$ the integral would equal $w(x)$. The step function makes the integrand vanish for $r_2 > \max[h(r_1)]$, so if we could make $h(r_1) \leq 1$ as a probability, then $r_2 \leq 1$ and we could sample $r_2 \in U[0, 1]$. This can be achieved by an extra factor c : Define $h(x) = cw(x)/g(x)$, and

$$\begin{aligned} w(x) &= \frac{1}{c} h(x) g(x) \\ &= \frac{1}{c} \int_0^\infty dr_1 g(r_1) \int_0^1 dr_2 \delta(x - r_1) \theta(h(r_1) - r_2) . \end{aligned} \quad (25)$$

We are almost done, all we need to find is some c that makes $h(x) \leq 1$. If $\max(h) < 1$, then sampling $r_2 \in U[0, 1]$ would waste some time trying to pick useless numbers between $\max(h)$ and 1. The optimal choice is clearly

$$\max(h) = 1 , \quad (26)$$

which gives $c = [\max(w/g)]^{-1} = \min(g/w)$. Now we are ready to write down the algorithm:

1. Evaluate

$$c = \min\left(\frac{g}{w}\right) \quad (27)$$

2. Draw $r_1 \in U[0, 1]$ and r_2 from PDF $g(x)$.

3. If

$$c \frac{w(r_1)}{g(r_1)} \geq r_2, \quad (28)$$

then accept $x = r_1$ as the desired random number, else discard both r_1 and r_2 . Go to 2.

A trivial case would be $g(x) = w(x)$, which would mean that we already know how to sample $w(x)$. Then $c = 1$ and all r_1 would be accepted ($1 \geq r_2$ is always valid): we say that the *acceptance ratio* is 1 in that case. If c is small, that is, in some region $g(x) \ll w(x)$, then much more trial r_1 's would be rejected and the algorithm would be ineffective. This problem is known as *under-sampling*, and it can even turn the code useless. Using layman's terms, if one is nearly overlooking some region where $w(x)$ is large, then $w(x)$ is under-sampled. The function $g(x)$ acts as an importance sampling function, and picking a bad one can be disastrous.

Standard normal distribution The standard normal distribution $\text{Normal}(x)$ is a Gaussian with mean zero and variance one,

$$\text{Normal}(x) = \frac{1}{\sqrt{2\pi}} e^{-x^2/2}. \quad (29)$$

This PDF is of special importance in MC, so we give few ways to sample random numbers x from it. A generalisation to arbitrary mean and variance is easy. If you want to sample

$$\text{Normal}(y, \mu, \sigma) = \frac{1}{\sqrt{2\pi}\sigma} e^{-(y-\mu)^2/(2\sigma^2)}, \quad (30)$$

you obviously need to spread the PDF and shift its mean to μ : $y = \sigma x + \mu$. Sampling from a multivariate Gaussian is beyond the present scope.

- Draw N random numbers $r_i \in U[0, 1]$ and calculate the sum

$$x = \sqrt{\frac{12}{N}} \sum_{i=1}^N (r_i - \frac{1}{2}) \quad (31)$$

With N large this approaches the standard normal distribution. This is a consequence of the Central Limit Theorem; a proof will be given soon. This gives a rather good $Normal(x)$ already with small N . Values $N = 6$ or $N = 12$ can be quite useful, if only crude normal distribution is sufficient.

- **Box-Müller algorithm I**⁴

1. Draw r_1 and $r_2 \in U[0, 1]$
2. Evaluate $\theta = 2\pi r_1$ and $R = \sqrt{-2 \ln r_2}$
3. $(x_1, x_2) = (R \cos(\theta), R \sin(\theta))$
4. Go to step 1.

Notice, that you get *two* random numbers (x_1 and x_2) per round, so if you have this in a subroutine it gives a “free” x for every other call. The downside is the four special functions that slow things down, so use this only if you need a truly good $Normal(x)$.

- **Box-Müller algorithm II** A modified version of the previous algorithm, uses the rejection method to avoid trig. functions.

1. Draw r_1 and $r_2 \in U[0, 1]$
2. Make $R_1 = 2r_1 - 1$ and $R_2 = 2r_2 - 1$, they sample $\in U[-1, 1]$
3. If $R = R_1^2 + R_2^2 > 1$ reject the pair and go to step 1
4. Result:

$$(x_1, x_2) = \sqrt{\frac{-2 \ln R}{R}} (R_1, R_2) \quad (32)$$

5. Go to step 1

Again, we get two random samples per round. The acceptance ratio of this algorithm is $\pi/4$, the ratio of the area of a unit circle to the area of the 2-by-2 box around it.

⁴G. Box and M. Müller, Ann. Math. Stat. 38 (1958) 610

- **Fernandéz–Rivero algorithm** A very fast algorithm for generation of Gaussian distributed random numbers⁵. Highly recommended!

1. Initialise $v_i = 1, i = 1, \dots, N$ (N is about 1024) and set $j = 0$
2. Draw K (about 16) values $\theta_k \in U[0, 2\pi]$; compute $c_k = \cos(\theta_k)$ and $s_k = \sin(\theta_k)$.
3. Set $j = (j + 1) \bmod N$, draw $i \in IU[0, N - 1]$ so that $i \neq j$
4. Draw $k \in IU[0, K - 1]$
5. Set $tmp = c_k v_i + s_k v_j$ and $v_j = -s_k v_i + c_k v_j$
6. Set $v_i = tmp$
7. If still in thermalisation go to step 3
8. Result: v_i and v_j
9. Go to 3

Here $IU[0, N]$ means random integer between 0 and N .

The Box-Müller and Fernandéz–Rivero algorithms are accurate also at the tail of $Normal(x)$, while the first algorithm clearly can get only as far from the origin as N steps allow.

Proof of the Box-Müller algorithm The algorithm is based on a common trick to integrate a Gaussian: the product of two Gaussian with Cartesian coordinates x and y can be most easily integrated in circular coordinates using

$$x^2 + y^2 = r^2. \quad (33)$$

We start the proof from the end. The Box-Müller algorithm calculates the integral

$$\begin{aligned} f(x_1, x_2) &= \int_0^1 dr_1 \int_0^1 dr_2 \delta(x_1 - \sqrt{-2 \ln r_1} \cos(2\pi r_2)) \delta(x_2 - \sqrt{-2 \ln r_1} \sin(2\pi r_2)) \\ &= \int_0^1 dr_1 \int_0^1 dr_2 \delta(r_1 - e^{-(x_1+x_2)^2/2}) \delta(r_2 - \frac{1}{2\pi} \arctan(\frac{x_2}{x_1})) \left| \frac{\partial(r_1, r_2)}{\partial(x_1, x_2)} \right| \\ &= \frac{1}{2\pi} e^{-(x_1^2+x_2^2)/2} \\ &= \frac{1}{\sqrt{2\pi}} e^{-x_1^2/2} \times \frac{1}{\sqrt{2\pi}} e^{-x_2^2/2}. \end{aligned} \quad (34)$$

⁵J. F. Fernandéz and J. Rivero, Comp. in Phys. 10 (1996) 83-88.

The key point is the Jacobian of the transformation, which is exactly the desired product of the two Gaussians (the delta-function integration is trivial after the transformation). These algorithms are always *mappings* from one or more random numbers $\in U[0,1]$ to another distribution. The last line shows that both x_1 and x_2 are sampling the standard normal distribution.

Principle of the Fernandéz-Rivero algorithm This algorithm has its roots in physics, and it is apparently a simulation. As a first step, take a look at how exponentially distributed random numbers can be generated. Consider a collection of N classical particles, that can exchange energy with each other but not with the rest of the world. From statistical physics we know, that this ensemble will equilibrate to a state, where the energies are distributed according to the Gibbs law

$$P(E) \sim e^{-E/T}, \quad (35)$$

where T is the temperature. Let the N -dimensional table E hold the energies of the particles. Initialise this table to unity, $E_i = 1$ for $1 \leq i \leq N$. Next introduce an interaction between the particles in such a way, that the total energy is conserved. Let the particles interact pairwise, pick two particles i and j at random. To conserve the total energy the sum $E_i + E_j$ must be kept constant, mark $S = E_i + E_j$. Next let the interaction to transfer a random amount of energy from particle i to particle j . The algorithm is

1. Initialise $E_i = 1$ for $1 \leq i \leq N$ and set $j = 0$
2. Set $j = (j+1) \bmod N$ and draw an integer $i \neq j$ from $i \in \text{int}[N U[0,1]]$
3. Set $S = E_i + E_j$
4. Draw $r \in U[0,1]$, set $E_i = rE_i$ and after that set $E_j = S - E_i$
5. If still in thermalisation go to step 2
6. Result: two random numbers E_i and E_j
7. Go to step 2

According to Fernandéz and Rivero, $N \sim 1000$ is enough for generations of about 10^{15} exponentially distributed random numbers. The energies have to be thermalised before they become Gibbs distributed, so one has to run the algorithm about $10 \times N$ times to let the dust settle. After that, each round produces two independent random numbers, E_i and E_j , distributed as

$w(E) = e^{-E}$. The modulo operation can be done fast with logical **and** if N is a power of 2: $N = 1024$ is just fine.

Normal distributed random numbers can be generated in a related fashion. From statistical mechanics we know, that in equilibrium velocities have a Maxwell distribution, which is nothing but a gaussian. This time we keep track of the velocities of the particles, and to conserve energy their squares must be conserved. Let the velocities of the interacting particles be v_i and v_j . Think of them as components of a two-dimensional vector, whose length squared $v_i^2 + v_j^2$ must be conserved in the interaction. Thus the interaction should rotate the vector; To speed up the rotation one can compute a finite set of rotation matrices beforehand. Although plausible, the actual proof that the algorithm does what it promises is beyond the scope of these lecture notes ⁶.

Some simple distribution The Dirac delta function can be really handy. Assume x samples $p(x)$ and y samples $q(y)$. What distribution $h(z)$ does their sum $z = x + y$ sample? Simply write

$$\begin{aligned} h(z) &= \int dx dy \delta(z - (x + y)) p(x) q(y) \\ &= \int dx p(x) q(z - x) , \end{aligned} \tag{36}$$

so the sum samples the convolution of the two distributions.

The product xy samples

$$\begin{aligned} h(z) &= \int dx dy \delta(z - xy) p(x) q(y) \\ &= \int dx dy \frac{1}{|x|} \delta(y - \frac{z}{x}) p(x) q(y) \\ &= \int dx \frac{1}{|x|} p(x) q(\frac{z}{x}) . \end{aligned} \tag{37}$$

2.2 Some terminology

Notation varies from source to source, be cautious.

Every step of the random walk is made of two parts:

- Some means to choose next point S' once the present point S is given
- Decision whether or not to accept this move from S to S'

⁶See the web page of J. F. Fernández, <http://pipe.unizar.es/~jff/papers/rn.html> for more details.

These first part is given by the transfer matrix, the second part is the acceptance probability and these two combined give the transition probability.

Transfer matrix Moving a walker from point S to S' is just a vector transformation. It is customary to define the *transfer matrix* $T(S, S')$ via

$$S' = T(S \rightarrow S')S . \quad (38)$$

This tells us how to generate a new S' from the old S . In matrix language

$$S_j = \sum_i T_{ji} S_i \quad \text{or just } S = TS , \quad (39)$$

where $S = \{S_1, S_2, \dots, S_N\}$ is a vector and T is a $N \times N$ matrix.

Transition probability (transition matrix, probability matrix) As mentioned above, the transition probability is made of two parts:

- (i) Transfer matrix $T(S \rightarrow S')$ that tells how a new state is generated from the old one, and
 - (ii) the part $A(S \rightarrow S')$ that tells how probably this new state is accepted.
- As a formula the transition probability (matrix) reads ⁷

$$P(S \rightarrow S') = T(S \rightarrow S')A(S \rightarrow S') . \quad (40)$$

In matrix language $P(S \rightarrow S')$ is the matrix element P_{ij} , so the list of all transition probabilities form a *transition matrix* P .

To illustrate the two parts of the transition probability, the probability to have a nuclear explosion depends on a governments ability to blow the bomb (T) and whether it accepts or not to use the opportunity (A). To minimise the probability the former factor is kept smaller by the Nuclear Non-proliferation Treaty (NPT) and the latter relies on the Keep-the-Nuke-Owners-Happy Program (KNOHP).

Ergodicity A Markov chain is ergodic if the walker can reach any point in configuration space in finite time, it is non-periodic, and the average return time to any state is finite. It can be very difficult to prove that the method of choosing new configurations S is ergodic, because walker(s) can get trapped in subsets of the configuration space, such as periodic loops. Some Markov chains are quasi-ergodic: There may be a high energy barrier between subspaces, so it is theoretically possible to reach all possible states, but highly

⁷Often one sees transition matrix marked with T , not with P as I've done.

improbable in a finite time simulation. One should always compare the results from multiple runs with different initial configurations.

Weight of a state S , $\pi(S)$ The weight is a positive real number that characterises the importance of S among all states. In general the $\pi(S)$ may exceed unity, so they are not probabilities. If $\pi(S_0) = 0$ the system can never be in state S_0 . If $\pi(S_1) \neq 0$ and all other $\pi(S_i) = 0$ the system stays in state S_1 .

2.3 Central Limit Theorem

The Central Limit Theorem states that

Given a distribution with a mean μ and variance σ^2 , the sampling distribution of the mean approaches a normal distribution with a mean μ and a variance σ^2/N as the sample size N increases.

This is almost magic: taking the mean value of variables x with *any* probability distribution function (PDF) $f(x)$,

$$z = \frac{x_1 + x_2 + \dots + x_N}{N} \quad (41)$$

the sampling distribution $g(z)$ of the mean value approaches a normal distribution! Even more amazingly, this seems to happen already for a rather small N .

What is exactly N and a “sampling distribution”? The number of samples x_i is assumed to be infinite, and we are not going to add them all up to get the mean. So we *sample the mean* by taking a finite size specimen (of size N) and gather the mean values of them to get an idea of the true mean value.

Proof:

We wish to find $g(z)$ for large N . The PDF of N arbitrary values x_i is the product of the individual probabilities. Now we just add the constraint that their mean is z :

$$g(z) = \int dx_1 dx_2 \dots dx_N f(x_1) f(x_2) \dots f(x_N) \delta\left(z - \frac{x_1 + x_2 + \dots + x_N}{N}\right). \quad (42)$$

Writing the Dirac delta function in integral form

$$\delta(x) = \frac{1}{2\pi} \int dk e^{ikx} \quad (43)$$

we get

$$g(z) = \frac{1}{2\pi} \int dk dx_1 dx_2 \dots dx_N f(x_1) f(x_2) \dots f(x_N) \quad (44)$$

$$\begin{aligned} & \times \exp \left[ik \left(z - \frac{x_1 + x_2 + \dots + x_N}{N} \right) \right] \\ &= \frac{1}{2\pi} \int dk \left[\int dx_1 f(x_1) \exp \left[ik \left(\frac{z - x_1}{N} \right) \right] \right] \\ & \quad \times \left[\int dx_2 f(x_2) \exp \left[ik \left(\frac{z - x_2}{N} \right) \right] \right] \\ & \quad \dots \times \left[\int dx_N f(x_N) \exp \left[ik \left(\frac{z - x_N}{N} \right) \right] \right] \end{aligned} \quad (45)$$

$$= \frac{1}{2\pi} \int dk \left[\int dx f(x) \exp \left[ik \left(\frac{z - x}{N} \right) \right] \right]^N. \quad (46)$$

For large N the value of $(x_1 + x_2 + \dots + x_N)/N$ is going to be close to μ , so it's a good idea to add and subtract it:

$$g(z) = \frac{1}{2\pi} \int dk \left[\int dx f(x) \exp \left[ik \left(\frac{z - \mu + \mu - x}{N} \right) \right] \right]^N \quad (47)$$

$$= \frac{1}{2\pi} \int dk e^{ik(z-\mu)} \left[\underbrace{\int dx f(x) \exp \left[ik \left(\frac{\mu - x}{N} \right) \right]}_I \right]^N. \quad (48)$$

The integral in the brackets can now be expanded as

$$I = \int dx f(x) \left[1 + \frac{ik(\mu - x)}{N} - \frac{k^2(\mu - x)^2}{2N^2} + \dots \right] \quad (49)$$

$$= 1 + ik \frac{\mu - \int dx f(x)x}{N} - \frac{k^2}{2N^2} \int dx f(x)(\mu - x)^2 + \dots \quad (50)$$

$$= 1 - \frac{k^2 \sigma^2}{2N^2} + \dots \quad (51)$$

Hence, for large N ,

$$g(z) = \frac{1}{2\pi} \int dk e^{ik(z-\mu)} \left[1 - \frac{k^2 \sigma^2}{2N^2} \right]^N, \quad (52)$$

and since

$$\left[1 - \frac{x}{N} \right]^N \rightarrow e^{-x} \text{ as } N \rightarrow \infty \quad (53)$$

we have finally

CENTRAL LIMIT THEOREM

$$g(z) \approx \frac{1}{\sqrt{2\pi} \left[\sigma/\sqrt{N} \right]} e^{-\frac{(z-\mu)^2}{2[\sigma/\sqrt{N}]^2}} \quad (54)$$

Cauchy distribution A word of warning: Not all PDF's satisfy the Central Limit Theorem. As an example, consider the (standard) Cauchy distribution

$$w(x) = \frac{1}{\pi(1+x^2)} , \quad (55)$$

which looks a bit like the standard normal distribution (see Fig. (3)). We have

$$\int_{-\infty}^{\infty} w(x) dx = 1 , \quad (56)$$

so it's a perfectly admissible PDF. However, take a look at the mean

$$\mu = \int_{-\infty}^{\infty} xw(x) dx . \quad (57)$$

The argument is antisymmetric, so contributions from $x < 0$ and $x > 0$ cancel exactly giving $\mu = 0$, as one sees also visually, but the separate integrals for $x < 0$ and $x > 0$ diverge. If one samples N points $x_1, \dots, x_N$ from the Cauchy distribution and computes the average $N^{-1}(x_1 + \dots + x_N)$, this sum cannot converge because the points are not chosen symmetrically about $x = 0$. As a result, the Cauchy distribution has no statistical mean.

The variance is worse, the integral

$$\int_{-\infty}^{\infty} x^2 w(x) dx \quad (58)$$

is infinite also in the mathematical sense. The non-existence of these momenta pulls the carpet from under the central limit theorem. Even though you can sample $w(x)$ and - with a great deal of wishful thinking - calculate the mean over *that* set, still no matter how many data points you collect, your estimated mean is as “accurate” as with just one point!

One often uses the Cauchy distribution as a test distribution to see, how sensitive a method is to large tail contributions. By the way, the distribution of the ratio of two gaussian-distributed random variables is a Cauchy distribution.




Figure 2: The distribution of the mean value of N random numbers $\in U[-0.5, 0.5]$ for $N = 2, 3, 4, 5$ and 12 (Eq. (31)).




Figure 3: The standard Cauchy probability distribution compared with the standard normal distribution.

2.4 Law of Large Numbers

The Law of Large Numbers says that

If an event of probability p is observed repeatedly during independent repetitions the proportion of the observed frequency of that event to the number of repetitions converges towards p as the number of repetitions become large.

or equivalently

In repeated, independent trials with the same probability p of success in each trial, the chance that the percentage of successes differs from p by more than any fixed amount $\epsilon > 0$ converges to zero as the number of trials goes to infinity.

or yet another form

The sequence of averages converges to some limit number which is an expectation.

In layman's terms, if you start throwing a dice your probabilities of getting $1, \dots, 6$ will approach the true probability distribution $\{\frac{1}{6}, \dots, \frac{1}{6}\}$.

2.5 Detailed balance condition

The detailed balance condition is a sufficient, but not necessary, condition that the system reaches the equilibrium.

With *the* equilibrium I mean the correct equilibrium, it is the one where the states appear with correct weights $\pi(S)$; There are numerous ways to reach a wrong “equilibrium”! To put it another way, we want to make certain that after repeating moves according to some transition probability $P(S \rightarrow S')$ the asymptotic distribution of states S is $\pi(S)$. This is the usual starting premise: The weights of the states $\pi(S)$ are known and the task is to construct the transition probability $P(S \rightarrow S')$ that will lead to equilibrium. The detailed balance condition gives the first guideline to follow. What makes it so pleasing, is that *any* transition probability that satisfies the detailed balance condition with the weights $\pi(S)$ is acceptable!

The detailed balance condition, known also as *microscopic reversibility*, states that

$$P(S \rightarrow S')\pi(S) = P(S' \rightarrow S)\pi(S') , \quad (59)$$

where $\pi(S)$ is the weight of state S . This equality assures that there is no net traffic between the states S and S' , but it does more than that.⁸ The origin of the detailed balance condition can be traced to the so-called master equation, which states how the probability of a state (marked also as P) varies as a function of the simulation time (treated as continuous) :

$$\frac{dP(S')}{dt} = \sum_S P(S \rightarrow S')\pi(S) - \sum_S P(S' \rightarrow S)\pi(S') , \quad (60)$$

where the terms sum the “changes of getting to S' ” minus “changes of getting out of S' ”. Perhaps the easiest way to interpret the terms in the sums is to think of them as “probability of being in the initial state times the transition rate from the initial to the final state”. The master equations is like a continuity equation for the probability. It’s also characteristic to Markov processes, because the future depends only on the present state, not on the earlier history.

Now the connection between the detailed balance condition and the equilibrium of the system becomes more transparent. In equilibrium $\frac{dP(S')}{dt} = 0$,

⁸A bad simulation can get permanently stuck to state S , so there is definitely no net traffic between any states – but usually at finite temperatures this is not the *correct* equilibrium! Another form of false equilibria are *limit cycles*, which also result from irreversible changes made in a poor simulation.

so

$$\sum_S P(S \rightarrow S') \pi(S) = \sum_{S'} P(S' \rightarrow S) \pi(S') . \quad (61)$$

Obviously this equality is valid if the terms in the sums are equal, $P(S \rightarrow S') \pi(S) = P(S' \rightarrow S) \pi(S')$. But this is just the detailed balance condition (59) - the two sums are in balance term by term:

*Balance every process between a pair of states,
 $S \rightarrow S'$ in balance with $S' \rightarrow S$.*

Of course the two sums in (61) can give the same value even though the terms are not the same term by term: The detailed balance condition is not the *only* solution that is compatible with the equilibrium condition. With some imagination you may add a kind of running weight to the terms so that once the sum is done the extras cancel. In short, Monte Carlo doesn't need strict detailed balance, but it is a simple and a safe starting point.

Finally, to wrap up the discussion about detailed balance we need to show that after a repeated application of transitions with probability $P(S \rightarrow S')$ that satisfy the detailed balance condition the asymptotical distribution is $\pi(S')$. We set out to proof that

$$\pi(S') = \sum_S \pi(S) P(S \rightarrow S') . \quad (62)$$

In matrix notation this reads $\pi = \pi P$, which states that π is the (left) eigenvector of transition matrix P with eigenvalue 1.

Proof⁹:

$$\begin{aligned} \sum_S \pi(S) P(S \rightarrow S') &= \sum_S \pi(S') P(S' \rightarrow S) \\ &= \pi(S') \sum_S P(S' \rightarrow S) = \pi(S') . \end{aligned} \quad (63)$$

The first equality follows from detailed balance, the last one from the fact that P is a *stochastic matrix*. This means that the sum of row elements of P is unity and all elements are non-negative: Probability that S' is changed to some S (includes the case where S' remains S') is 1.

⁹This is not very rigorous, to give credit to the statement that this is asymptotically valid one would have to discuss the simulation time t .

Example Let the weights in a two-state system be $\pi(1) = 5$ and $\pi(2) = 3$. The detailed balance says

$$\frac{P(1 \rightarrow 2)}{P(2 \rightarrow 1)} = \frac{\pi(2)}{\pi(1)} = \frac{3}{5}. \quad (64)$$

Mark $P_{12} = P(1 \rightarrow 2)$ and $P_{21} = P(2 \rightarrow 1)$. This fixes the ratio of the two off-diagonal elements of P , and using the fact that this is a stochastic matrix we have

$$P = \begin{pmatrix} P_{11} & P_{12} \\ P_{21} & P_{22} \end{pmatrix} = \begin{pmatrix} 1 - 3x & 3x \\ 5x & 1 - 5x \end{pmatrix}. \quad (65)$$

A neat choice would be $x = 1/5$,

$$P = \begin{pmatrix} 2/5 & 3/5 \\ 1 & 0 \end{pmatrix}. \quad (66)$$

First test that the distribution π is a left eigenvector with eigenvalue 1. Indeed,

$$\pi P = (5 \quad 3) \begin{pmatrix} 2/5 & 3/5 \\ 1 & 0 \end{pmatrix} = (5 \quad 3) = \pi. \quad (67)$$

Next, successive applications of P give

$$P^{10} \approx \frac{1}{8} \begin{pmatrix} 5.01614 & 2.978186 \\ 4.96977 & 3.03023 \end{pmatrix}, \quad (68)$$

so it really seems that the rows go asymptotically to π (normalisation of π is the factor $1/8$).

3 Markov Chain Monte Carlo

3.1 Metropolis–Hastings algorithm

In year 1953 Metropolis *et al.*¹⁰ published an algorithm, which can be used to sample an almost arbitrary, multidimensional distribution $\pi(S)$. Here S can be, for example, described by a set of particle coordinates

¹⁰N. Metropolis, N. Rosenbluth, A. W. Rosenbluth, A. H. Teller and E. Teller, *Equation of state calculations by fast computing machines*, J. Chem. Phys. **21** (1953) 1087-1092. An earlier form was published by Metropolis and Ulam 1949. Allegedly the algorithm was invented at a Los Alamos dinner party.

$\mathbf{R} = \{\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N\}$ or a set of occupation numbers $\{n_1, n_2, \dots\}$ of quantum mechanical states $1, 2, \dots$. The algorithm generates the set of random numbers by means of a random walk. The sequence of points gone through is an example of a **Markov chain**, after the russian mathematician A. A. Markov (1856-1922). A Markov chain is defined as a process containing a finite number of states for which the probability of being in each state in the future depends only on the present state of the process.

It was Hastings¹¹ who generalised the algorithm to account for nonsymmetric transfer matrices, and showed how the asymmetry in proposals is compensated in the acceptance probabilities.

The original Metropolis algorithm Consider two configurations of the system, S and S' . In its original form the Metropolis algorithm samples the Boltzmann distribution. There the (unnormalised) weight of the configuration S depends on its energy $E(S)$ according to

$$\pi(S) = e^{-E(S)/T} . \quad (69)$$

Metropolis *et al.* chose the acceptance probability from S to S' to be

$$A(S \rightarrow S') = \begin{cases} e^{-\Delta E/T}, & \text{if } \Delta E = E(S') - E(S) > 0 \\ 1, & \text{else,} \end{cases} \quad (70)$$

An equivalent way of writing this is

$$A(S \rightarrow S') = \min [1, e^{-\Delta E/T}] , \quad (71)$$

where *min* means “smaller of the two arguments”. The detailed balance condition boils down to a very simple relation,

$$\frac{P(S \rightarrow S')}{P(S' \rightarrow S)} = \frac{T(S \rightarrow S')A(S \rightarrow S')}{T(S' \rightarrow S)A(S' \rightarrow S)} = \frac{\pi(S')}{\pi(S)} = e^{-\Delta E/T} . \quad (72)$$

The first two equalities go also under the name *Metropolis–Hastings* due to the explicit transfer matrix in the formula. The original Metropolis algorithm assumes a symmetric transfer matrix, $T(S \rightarrow S') = T(S' \rightarrow S)$, so the T ’s cancel out from Eq. (72) and we might as well have written

$$P(S \rightarrow S') = \min [1, e^{-\Delta E/T}] . \quad (73)$$

¹¹W.K. Hastings, *Monte Carlo sampling methods using Markov chains and their applications*, Biometrika 57, 97-109 (1970). His PhD student Peskun derived the conditions for efficient sampling using Markov chains (Peskun ordering).

If we are going uphill ($E(S') > E(S)$), then

$$\frac{P(S \rightarrow S')}{P(S' \rightarrow S)} = \frac{e^{[E(S)-E(S')]/T}}{1} = e^{-\Delta E/T}, \quad (74)$$

and if we are going downhill then

$$\frac{P(S \rightarrow S')}{P(S' \rightarrow S)} = \frac{1}{e^{[E(S')-E(S)]/T}} = e^{-\Delta E/T}, \quad (75)$$

so the detailed balance condition is satisfied and this $P(S \rightarrow S')$ leads asymptotically to the Boltzmann distribution.

Pay attention to the fact that in the detailed balance condition (72) the ratio does not involve the partition function Z , which is the normalisation factor of the Boltzmann distribution. Thus we can construct the transition matrix without ever bothering to compute Z . A similar situation occurs in quantum Monte Carlo: There is no need to normalise the wave function explicitly. Great relief!

The algorithm is

- If $e^{-\Delta E/T} > 1$ then $P(S \rightarrow S') = 1$ and the move is always accepted.
- If $e^{-\Delta E/T} < 1$ interpret it as a probability and compare it with a random number $r \in U[0, 1]$; accept it if $r < e^{-\Delta E/T}$.

This means

1. Propose a move $S \rightarrow S'$
2. If the energy decreases, accept the move: always go downhill
3. If the energy increases, accept the move with probability $e^{-\Delta E/T}$: go uphill if the temperature allows you. Go to 1.

Off-lattice Metropolis algorithm for almost any $\pi(\mathbf{R})$ Assume that the configuration of the system is $\mathbf{R}$; One says that the “random walker” is at point $\mathbf{R}$. Here is a Metropolis algorithm to sample the next point:

1. Pick randomly a trial point $\mathbf{R}'$ from the neighbourhood of $\mathbf{R}$. You can choose, for example, a random point within some distance δ of the original point.
2. calculate

$$ratio = \frac{\pi(\mathbf{R}')}{\pi(\mathbf{R})} \quad (76)$$

3. Accept $\mathbf{R}'$ as the next point if $ratio > 1$: set $\mathbf{R} = \mathbf{R}'$ (move the walker). Go to 1.
4. Draw $z \in U[0, 1]$. Accept $\mathbf{R}'$ as the next point if $ratio > z$ (in a program, set $\mathbf{R} = \mathbf{R}'$, because there is no need to keep the old point stored). Go to 1.
5. Reject $\mathbf{R}'$: The next point is the old point $\mathbf{R}$. Go to 1.

Steps 3 and 4 are often called the *Metropolis question*. A common mistake in applying the algorithm is to forget the last part of step 5. Obviously one must now and then use the same point more than once if the walker is in a region where $\pi(\mathbf{R})$ is large to describe that the walker now and then refuses to step out of high-weight regions ¹²

The free parameter δ dictates the acceptance ratio and the effectivity of sampling. By effectivity we mean that the walker is supposed to move around the phase space where $\pi(\mathbf{R})$ is nonzero within reasonable time. Hence there are two reasons why sampling can be *ineffective*:

- Walker tries to take *too long steps*; In a region with large $\pi(\mathbf{R})$ most of the trial steps would be rejected.
- Walker tries to take *too short steps* and almost all of them are accepted.

In both cases the walker is not moving around as it should. A rule of thumb says that for most effective sampling the acceptance ratio should be around 50%.

The other free parameter in the algorithm is the starting point of the walk. In principle it can be chosen freely, but it's better to start it from a point where $\pi(\mathbf{R})$ is fairly large to make the walk *thermalise* (lose memory of it's origin) faster.

The Metropolis algorithm can be used, for example, to evaluate expectation values

$$\langle f \rangle_w = \frac{\int d\mathbf{R} \pi(\mathbf{R}) f(\mathbf{R})}{\int d\mathbf{R} \pi(\mathbf{R})} . \quad (77)$$

One just takes the average value of f over the points produced by the random walk. The points in the random walk are not statistically independent, therefore the statistical error of the result is not given by Eq. (6).

¹²Random walkers are like drunkards; once the lost soul has landed in his favorite bar you may need more than one other place to suggest him before you can make him leave.

Example Look again at the two–state system with weights $\pi(1) = 5$ and $\pi(2) = 3$. Assume that the transfer matrix is symmetric; The Metropolis acceptance probability is the same as the transition probability,

$$P(S \rightarrow S') = A(S \rightarrow S') = \min \left[1, \frac{\pi(S')}{\pi(S)} \right] . \quad (78)$$

so the corresponding matrix is (don't use the above formula for $S \rightarrow S$)

$$P = \begin{pmatrix} 2/5 & 3/5 \\ 1 & 0 \end{pmatrix} . \quad (79)$$

This happens to be the matrix from chosen in the previous example, the $x = 1/5$ case of P given in Eq. (65).

Example One Markov chain generated by the stochastic matrix given above is shown in Fig. 4. I started with state 1, moved to 2 with probability $3/5$. Transition $2 \rightarrow 1$ is always accepted (downhill). I always suggested a change of state, so the transfer matrix (proposal matrix) was

$$T = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} , \quad (80)$$

which is symmetric.

Example A three–state system has stochastic matrix

$$P = \begin{pmatrix} 0.6 & 0.1 & 0.3 \\ 0.2 & 0.6 & 0.2 \\ 0.0 & 0.2 & 0.8 \end{pmatrix} \quad (81)$$

What is the underlying distribution? By random walk and counting the frequencies as in the previous example I got $\pi = (1.2309, 2.4614, 4.3076)$, some three decimals have converged.




Figure 4: Upper figure: A Markov chain for the two-state system with $\pi = (5, 3)$, lines are guides to the eye. The frequencies states 1 and 2 occurred during the depicted simulation were 5.20 and 2.82, respectively. After 10^8 steps the frequencies were 5.00010832 and 2.99989176. Lower figure: The convergence toward $\pi(2) = 5$ for three choices of x in equation (65) (three distinct transition matrices) showing that the Metropolis matrix has smallest fluctuations.

4 Error estimation

4.1 Some definitions and clarifications

Consider evaluating quantity A over a sample of size N . The mean value over the sample is

$$\overline{A} \equiv \sum_{i=1}^N A_i . \quad (82)$$

This is *not* the expectation value (ensemble average or the mean) of A , marked $\langle A \rangle$, because the sample is finite. The expectation value is the sum over every possible value of A_i ,

$$\langle A \rangle \equiv \sum_i p_i A_i , \quad (83)$$

where p_i is the probability of A_i . For a continuous distribution this is

$$\langle A \rangle \equiv \int dx p(x) A(x) . \quad (84)$$

The big difference is that

- $\overline{A}$ is a random variable that depends on the chosen sample
- $\langle A \rangle$ is a fixed number

According to the law of large numbers the mean and the expectation value coincide for an infinitely large sample,

$$\lim_{N \rightarrow \infty} \overline{A} = \langle A \rangle . \quad (85)$$

4.2 Biased estimators

The mean square deviation of the values from the mean value is called *biased estimator of the variance* σ_{biased} , with

$$\sigma_{biased}^2 = \frac{1}{N} \sum_{i=1}^N (A_i - \bar{A})^2 , \quad (86)$$

because the deviations are computed from the values $\bar{A}$ that vary from sample to sample.

4.3 Unbiased estimators

We cannot use infinite size samples, but we can still evaluate the *unbiased estimator of the variance*, meaning root mean square deviations from the true expectation value. Below I will show that

$$\sigma_{unbiased}^2 = \frac{1}{N-1} \sum_{i=1}^N (A_i - \bar{A})^2 , \quad (87)$$

so the two variances are related via

$$\sigma_{unbiased}^2 = \frac{N}{N-1} \sigma_{biased}^2 . \quad (88)$$

Statistical error estimation is based on the central limit theorem, Eq. (54). This contains *the* mean μ over all values, which is the expectation value. Thus the error estimation should utilize $\sigma_{unbiased}^2$ rather than σ_{biased}^2 ,

$$\boxed{error = \frac{\sigma_{unbiased}}{\sqrt{N}} = \sqrt{\frac{1}{N(N-1)} \sum_{i=1}^N (A_i - \bar{A})^2} , \quad (89)}$$

4.4 Derivation of $\sigma_{unbiased}$

Start from the definition of the biased estimator and subtract the constant $\langle A \rangle$:

$$\sigma_{biased}^2 = \frac{1}{N} \sum_{i=1}^N (A_i - \bar{A})^2 \quad (90)$$

$$\begin{aligned}
&= \frac{1}{N} \sum_{i=1}^N [A_i - \langle A \rangle - (\bar{A} - \langle A \rangle)]^2 \\
&= \underbrace{\frac{1}{N} \sum_{i=1}^N [A_i - \langle A \rangle]^2}_{\sigma_{unbiased}^2} - 2 \frac{1}{N} \sum_{i=1}^N (A_i - \langle A \rangle) \underbrace{(\bar{A} - \langle A \rangle)}_{no\ i} \\
&\quad + \frac{1}{N} \sum_{i=1}^N \underbrace{(\bar{A} - \langle A \rangle)^2}_{no\ i} \\
&= \sigma_{unbiased}^2 - (\bar{A} - \langle A \rangle)^2 .
\end{aligned} \tag{91}$$

The second term is the square deviation of the mean value from the expectation, but we know (or assume) that it obeys the central limit theorem:

$$(\bar{A} - \langle A \rangle)^2 = \left[\frac{\sigma_{unbiased}}{\sqrt{N}} \right]^2 = \frac{\sigma_{unbiased}^2}{N} . \tag{92}$$

Inserting this gives the relation mentioned earlier,

$$\sigma_{biased}^2 = \sigma_{unbiased}^2 \left[1 - \frac{1}{N} \right] . \tag{93}$$

4.5 Consequences of biased sampling error

- Without bias correction finite sampling underestimates variances
 $\Rightarrow$ Susceptibilities (response functions) are underestimated by factor $1 - \frac{1}{N}$

Ok, so what? Just collect more data and get smaller $1/N$. True, but – now comes the real bad news:

- $N = \#$ of statistically independent samples!
- Near criticality (like T_c in Ising model) sampling may become inefficient
 $\Leftrightarrow$ Autocorrelation times increase
 $\Leftrightarrow$ more computational time needed to increase N
- As a formula: N is related to the autocorrelation time by relation (τ given in measurement time intervals)

$$N = \frac{N_{all}}{1 + 2\tau} , \tag{94}$$

so if the autocorrelation time of your sampling increases dramatically you need a lot more simulation steps N_{all} .

4.6 Block averaging

A theorem says that the averages of sufficiently large blocks of data are statistically independent. Split the collected data, N values, to blocks each with n values. Say there are N_b blocks,

$$N = N_b \times n, \quad (95)$$

so we work with N_b mean values B_i computed for each block i :

$$\underbrace{A_1 \ A_2 \ \dots \ A_n}_{\text{mean } B_1} \quad \underbrace{A_{n+1} \ A_{n+2} \ \dots \ A_{2n}}_{\text{mean } B_2} \quad \underbrace{A_{2n+1} \ \dots \ A_N}_{\text{mean } B_{N_b}} . \quad (96)$$

Now if n is big enough¹³ the values B_i are *independent*, the error is given by the unbiased formula mentioned above,

$$\boxed{\text{error} = \sqrt{\frac{1}{N_b(N_b - 1)} \sum_{i=1}^{N_b} (B_i - \bar{B})^2}} . \quad (97)$$

If the blocks cover the whole sample of N values, then of course $\bar{B} = \bar{A}$. A prettier form is obtained by evaluating the square,

$$\boxed{\text{error} = \sqrt{\frac{1}{N_b - 1} (\overline{B^2} - \bar{B}^2)}} . \quad (98)$$

This method can be used to both primary variables (like energy; measured directly from MC simulation) and secondary variables (like specific heat, computed from energy data).

4.7 Resampling methods

The original set of data is collected by taking samples of a quantity during the simulation. That is sampling. In resampling one takes samples from the data that has been collected, so it's sampling from a sample.

4.7.1 Jackknife

Von Mises applied the jackknife resampling method in some form already in the pre-computing era (1930's), but the main work was done by Quenouille (1949) (who wanted to find a way to estimate bias) and Tukey (1958).

¹³Blocks bigger than the integrated autocorrelation time τ_{int} .

- Divide the data to N blocks, block size $> \tau_{int}$
 $\Rightarrow$ blocks are independent
- Omit the first, 2nd, 3rd, ... block from the sample of size N . Decimated set $A_{(i)}$ has i th block omitted. Case $N = 5$:

<i>Full set of blocks</i>	1	2	3	4	5		<i>average</i>	$\bar{A}$
<i>averages of blocks</i>	$\bar{a}_1$	$\bar{a}_2$	$\bar{a}_3$	$\bar{a}_4$	$\bar{a}_5$			
$A_{(1)}$		2	3	4	5		<i>average</i>	$\bar{A}_{(1)}$
$A_{(2)}$	1		3	4	5		<i>average</i>	$\bar{A}_{(2)}$
$A_{(3)}$	1	2		4	5		<i>average</i>	$\bar{A}_{(3)}$
$A_{(4)}$	1	2	3		5		<i>average</i>	$\bar{A}_{(4)}$
$A_{(5)}$	1	2	3	4			<i>average</i>	$\bar{A}_{(5)}$.

- Error of the mean can be computed using¹⁴

$$\boxed{error = \sqrt{\frac{N-1}{N} \sum_{i=1}^N (\bar{A}_{(i)} - \bar{A})^2}} . \quad (99)$$

The factor $\frac{N-1}{N}$ can be made plausible by noticing that

$$\bar{A}_{(i)} = \frac{1}{N-1} \sum_{j(\neq i)} \bar{a}_j , \quad (100)$$

so one recovers the error estimate of block data,

$$error = \sqrt{\frac{1}{N(N-1)} \sum_{i=1}^N (\bar{a}_i - \bar{A})^2} . \quad (101)$$

The resampled sets in jackknife have less data than the original set, so their variance is smaller. Thus the resampled sets are not really new measurements. Since the resampled sets deviate from the original set by only one data point, the jackknife is in this sense probing the near vicinity of the original set.

¹⁴An almost similar expression is valid for errors of other quantities than the mean.

Don't use jackknife to estimate the median, it will fail miserably. Median (middle value of the data) has what one calls "non-smooth statistics", a small change in the data set can make a big change in the median. This can be considered a sort of non-differentiability. Another quantity that is non-smooth is the maximum value of data.

4.7.2 Bootstrap

This is a generalization of the jackknife resampling, suggested by Efron (1979).

- Resample, *with replacement*, the existing data to get a new data set. Compute the statistics from this resampled set.

-

Original Sample : 1 2 3 4 5
1st bootstrap replica : 1 1 2 5 4
2nd bootstrap replica : 5 2 2 1 3
3rd bootstrap replica : 4 1 1 3 4
etc.

- About $N_{boot} = 1000$ replicas are often enough.
- Compute error using

$$error = \sqrt{\frac{1}{N-1} \left[\overline{A^2} - (\overline{A})^2 \right]}, \quad (102)$$

notice that one doesn't divide by $(N_{boot} - 1)$ ¹⁵

- You may block to get independent samples but it's not necessary

No theoretical calculations are needed, the method works no matter how mathematically complicated the chosen estimator (like variance) is. Just like the jackknife, the bootstrap can be used to estimate the bias and many other statistical properties as well.

The resampled sets deviate from the original one more than in the jackknife, so in this sense the bootstrap samples further away from the original set.

¹⁵The number of bootstrap replicas is unlimited, so having $1/(N_{boot} - 1)$ as a factor would lead to ever diminishing error, which cannot be.




Figure 5: The convergence of the bootstrap error estimate as a function of number of bootstrap replicas. The original data set is $\{1,2,3,4,5\}$.

Jackknife or bootstrap ?

- Jackknife requires as many resampled sets as there are original data points $\Rightarrow$ For about $N > 100$ bootstrap becomes more effective
- But for bootstrap you need about 1000 samples to get decent estimate for variance
 $\Rightarrow$ **Jackknife seems to be a better up to about 1000 data points. If you have more data use bootstrap .**
- If you don't block the data before error analysis, but choose to use statistically dependent data (=data picked more frequently than τ_{int}) you probably have lots of data $\Rightarrow$ use bootstrap.

How to debug error-estimation subroutines ?

- Program all three: blocking, jackknife, bootstrap
- Pick a simple set of five data points: 1,2,3,4,5
- Compare error estimates to $1/\sqrt{2} = 0.70717$
- If you get "NaN" in bootstrap, your $\overline{A^2} - (\overline{A})^2$ is negative. Check that you divide with correct integer counts. $\overline{A}$ should be close to 3.

- If your bootstrap error won't converge for increasing # of replicas, you haven't divided the averages $\overline{A^2}$ or $\overline{A}$ with N_{boot} .
- If your bootstrap converges to $\sqrt{2} = 1.414$, you forgot to multiply with $1/(N - 1)$ under the sqrt. In this case $N = 5$ (but don't hardwire N to 5 ;).
- After all seems to work, generate N data points using your old reliable gaussian random number generator. Analyse the data: the average should go toward 0 and the error of average should diminish as

$$\sigma_{input}/\sqrt{N} = 1/\sqrt{N}$$

(the code produces data with $\sigma_{input} = 1$).

5 Variational Monte Carlo

Eq. (77) can be readily applied to quantum many-body physics. The most straightforward application is the **variational Monte Carlo** (VMC), the simplest of **quantum Monte Carlo** (QMC) methods. If the method applies a Markov chain, then sometimes the name **Markov chain Monte Carlo** is also used.

Let $\varphi_T(\mathbf{R})$ be any trial wave function, $\mathbf{R} = (\mathbf{r}_1, \mathbf{r}_2, \dots, \mathbf{r}_N)$ is the set of $d \times N$ coordinates of the N particles in the d -dimensional system. The energy in that state is

$$\begin{aligned} E_T &= \frac{\langle \varphi_T | H | \varphi_T \rangle}{\langle \varphi_T | \varphi_T \rangle} \\ &= \frac{\int d\mathbf{R} |\varphi_T(\mathbf{R})|^2 [\varphi_T(\mathbf{R})^{-1} H \varphi_T(\mathbf{R})]}{\int d\mathbf{R} |\varphi_T(\mathbf{R})|^2} . \end{aligned} \quad (103)$$

The quantity

$$w(\mathbf{R}) = \frac{|\varphi_T(\mathbf{R})|^2}{\int d\mathbf{R} |\varphi_T(\mathbf{R})|^2} \quad (104)$$

is a perfectly good PDF and one can interpret the quantity one is averaging, namely

$$E_L(\mathbf{R}) = \varphi_T^{-1}(\mathbf{R}) H \varphi_T(\mathbf{R}) , \quad (105)$$

as the *local energy*. The energy in state φ_T is then simply

$$E_T = \int d\mathbf{R} w(\mathbf{R}) E_L(\mathbf{R}) \approx \frac{1}{M} \sum_{i=1}^M E_L(\mathbf{R}_i) , \quad (106)$$

the average of the local energy calculated over M points $\mathbf{R}_i$ that sample the distribution φ_T ; One can use the Metropolis algorithm to pick up these points as a random walk

$$\mathbf{R}_1 \rightarrow \mathbf{R}_2 \rightarrow \dots \rightarrow \mathbf{R}_M . \quad (107)$$

According to the Rayleigh-Ritz variational principle the true ground state energy is

$$E_0 \leq E_T , \quad (108)$$

so by making a good guess for the trial wave function one can easily calculate the upper bound of the ground-state energy.

Method of expected values The evaluation of averages just as sums like (106) is inefficient, because very improbable configurations occur in the random walk. The function evaluated at these configurations is far from the average, which increases the variance and slows down the MC algorithm. Instead of evaluating

$$\langle f \rangle \approx \frac{1}{M} \sum_{i=1}^M f(\mathbf{R}_i) , \quad (109)$$

one usually evaluates the average of “expected” values¹⁶. In a simulation of M -moves

$$\langle f \rangle \approx \frac{1}{M} \sum_{i=1}^M P(\mathbf{R}'_i) f(\mathbf{R}'_i) + [1 - P(\mathbf{R}_i)] f(\mathbf{R}_i) , \quad (110)$$

where $P(\mathbf{R}_i)$ is the probability that configuration $\mathbf{R}_i$ is accepted in the random walk; $\mathbf{R}_i$ and $\mathbf{R}'$ stand for the old and new configurations, respectively. If an unlikely configuration $\mathbf{R}_k$ is accepted, then $f(\mathbf{R}'_k)$ may be large, but since $P(\mathbf{R}'_k)$ is small, the product is neither large or small - this makes fluctuations smaller.

A Metropolis VMC algorithm This is for a N -particle system in volume V with a pair potential between the particles.

1. Pick up an initial configuration $\mathbf{R}_1$. For a smooth start, the wave function squared, $|\varphi_T(\mathbf{R}_1)|^2$, should not be too small. A lattice is usually a good starting configuration.
2. Choose the initial step size *step*. About $(V/N)^{1/3}$, often anything around 1 will do.
3. Set $i = 1$.
4. Pick i :th particle.
5. Draw three random numbers d_x, d_y and $d_z \in U[-\frac{1}{2}, \frac{1}{2}]$ (this is just $U[0, 1] - \frac{1}{2}$). These make the displacement vector $\mathbf{d} = (d_x, d_y, d_z)$.
6. Compute new attempted particle position $\mathbf{r}'_i = \mathbf{r}_i + \text{step } \mathbf{d}$
7. Compute

$$ratio = \left| \frac{\varphi_T(\mathbf{R}')}{\varphi_T(\mathbf{R})} \right|^2$$

¹⁶D. Ceperley, G. V. Chester and M. H. Kalos, Phys. Rev. B **16**, 3081 (1977).

8. Metropolis question: If $ratio > 1$ accept move. Else draw $t \in U[0, 1]$ and if $ratio > t$ accept move. These amount to accepting the move with probability

$$P = \min[1, ratio] .$$

If move was accepted set $\mathbf{r}_i = \mathbf{r}'_i$, else keep old $\mathbf{r}_i$.

9. Keep count N_{accept} on how many of the moves have been accepted out of N_{try} attempted moves.
10. Set $i = i + 1$. If $i < N$ go to 4.
11. Control the step size: If $100 * N_{accept}/N_{try} < 40$ set decrease $step$, if $100 * N_{accept}/N_{try} > 60$ increase $step$.
12. If still in thermalisation, go to 3.
13. Accumulate averages. Add to energy $P * E_L(\mathbf{R}_i) + (1 - P) * E_L(previous)$. Print out progress.
14. If still collecting data, go to 3

Properties of VMC

- What goes in comes out: the trial wave function φ_T dictates everything.
- Optimisation of φ_T is time consuming, but better φ_T means lower variance.
- See almost immediately if the property you added to φ_T has the desired effect.
- The perfect φ_T is the ground-state wave function and has zero variance in VMC. Every φ_T gives an upper limit to the energy (Rayleigh–Ritz variational principle).
- There is no way to know how close to the true ground state your φ_T is.
- Integrated quantities with few degrees of freedom, such as energy, are easy to evaluate. Functions with more degrees of freedom, like the pair distribution function, need more time.
- Human bias: Favors simple wave function over complicated: solid over liquid and polarised over unpolarised.

6 Diffusion Monte Carlo

The Schrödinger equation in imaginary time turns into a diffusion equation. Replacing $it/\hbar \rightarrow \tau$ one has

$$-\frac{\partial \Psi(\mathbf{R}, \tau)}{\partial \tau} = \hat{H} \Psi(\mathbf{R}, \tau) = [-D \nabla_{\mathbf{R}}^2 + V(\mathbf{R})] \Psi(\mathbf{R}, \tau) , \quad (111)$$

where $\nabla_{\mathbf{R}}^2 = \sum_i \nabla_i^2$ and is operating on $\mathbf{R}$. Here I marked $D = \hbar^2/(2m)$ (“ D ” for “diffusion constant”). Usually the potential has the form

$$V(\mathbf{R}) = \sum_{i \neq j} V(\mathbf{r}_i, \mathbf{r}_j) . \quad (112)$$

The Hamiltonian is time independent, so the formal solution reads

$$\Psi(\mathbf{R}, \tau) = e^{-\hat{H}\tau} \Psi(\mathbf{R}, 0) . \quad (113)$$

The eigenfunctions ϕ_i of $\hat{H}$ solved from the equation

$$\hat{H} \phi_i(\mathbf{R}) = E_i \phi_i(\mathbf{R}), \quad E_0 < E_1 < E_2 < \dots \quad (114)$$

form a complete basis, so we can expand

$$\Psi(\mathbf{R}, 0) = \sum_i c_i \phi_i(\mathbf{R}) , \quad (115)$$

where the sum is over all eigenstates. We have then

$$\Psi(\mathbf{R}, \tau) = \sum_i c_i e^{-E_i \tau} \phi_i(\mathbf{R}) . \quad (116)$$

Hence in the imaginary time all states with $E_i > 0$ decay exponentially, the ground state with energy $E_0 \leq E_i$ has the slowest decay rate (for bound systems highest growth rate).

Growing or decaying contributions are clumsy, so we stabilise the evolution by introducing a trial energy E_T , which we bluntly subtract from all energies (shift the energy scale)¹⁷. The energy-shifted equation

$$-\frac{\partial \Psi(\mathbf{R}, \tau)}{\partial \tau} = (\hat{H} - E_T) \Psi(\mathbf{R}, \tau) \quad (117)$$

has now the solution

$$\Psi(\mathbf{R}, \tau) = e^{-(\hat{H}-E_T)\tau} \Psi(\mathbf{R}, 0) = \sum_i c_i e^{-(E_i-E_T)\tau} \phi_i(\mathbf{R}) . \quad (118)$$

¹⁷Some sources call E_T “reference energy”.

Now all states with energy above E_T decay, and those with energy below E_T are amplified in (imaginary) time. This means that if we choose $E_T \approx E_0$, the asymptotic state that survives after the exponential decays is the true ground state $\phi_0(\mathbf{R})$. We can then start with whatever state $\Psi(\mathbf{R}, 0)$ and it will evolve into the ground state, *provided it's not orthogonal to $\phi_0(\mathbf{R})$* . In the course of the evolution E_T is updated toward E_0 , which is known more and more accurately.

Now we know in principle what the final outcome is. We still have to conceive a method to evolve $\Psi(\mathbf{R}, 0)$ in practise, since we haven't got the eigenvalues E_i . We start again from Eq. (118). The operator responsible for the time evolution is the Green's function

$$\hat{G} = e^{-(\hat{H}-E_T)\tau} , \quad (119)$$

which satisfies the operator equation

$$-\frac{\partial \hat{G}}{\partial \tau} = (\hat{H} - E_T)\hat{G} . \quad (120)$$

The coordinate space representation of $\hat{G}$ is a matrix with elements

$$G(\mathbf{R}', \mathbf{R}, \tau) = \langle \mathbf{R}' | e^{-(\hat{H}-E_T)\tau} | \mathbf{R} \rangle , \quad (121)$$

and for the many-body Hamiltonian defined in Eq. (111) the Green's function is the solution of the equation

$$-\frac{\partial G(\mathbf{R}', \mathbf{R}, \tau)}{\partial \tau} = [-D\nabla_{\mathbf{R}}^2 + V(\mathbf{R}) - E_T] G(\mathbf{R}', \mathbf{R}, \tau) . \quad (122)$$

6.1 Short-time estimates of the Green's function

At first glance Eq. (122) doesn't look any more promising than the original Schrödinger equation (111). But for *short time intervals* we may solve it approximately. A theorem by Trotter¹⁸ says, that while the kinetic and the potential parts T and V of the Hamiltonian do not commute in general ($\nabla_{\mathbf{R}}^2 V(\mathbf{R}) \neq V(\mathbf{R}) \nabla_{\mathbf{R}}^2$), we can ignore this for a very brief moment. Hence we look at the evolution from time 0 to τ , where τ is small.

If $\hat{T}$ and $\hat{V}$ would commute, we could find the Green's functions that does the "potential evolution" and the "kinetic evolution" separately. Now it's valid only up to order τ^2 :

$$G = e^{-(\hat{H}-E_T)\tau} \approx e^{-\hat{T}\tau} e^{-(\hat{V}-E_T)\tau} , \text{ for } \tau \rightarrow 0 \equiv G_T G_V . \quad (123)$$

¹⁸H. F. Trotter, Proc. Am. Math. Soc. **10** (1959) 545.

The first correction term to this approximation is

$$G - G_T G_V = \frac{1}{2}[\hat{V}, \hat{T}]\tau^2 + \mathcal{O}(\tau^3) , \quad (124)$$

and obviously there are ways to improve our G . We talk more about these refinements later.

The operator $\hat{V}$ is diagonal (leaves the state it operates on untouched). In coordinate space

$$\langle \mathbf{R}' | \hat{V} | \mathbf{R} \rangle = V(\mathbf{R}) \delta(\mathbf{R}' - \mathbf{R}) , \quad (125)$$

so

$$G_V(\mathbf{R}', \mathbf{R}, \tau) = e^{-(V(\mathbf{R}) - E_T)\tau} \delta(\mathbf{R}' - \mathbf{R}) . \quad (126)$$

This is the solution of the *rate equation*

$$\frac{\partial G_V(\mathbf{R}', \mathbf{R}, \tau)}{\partial \tau} = -(V(\mathbf{R}) - E_T) G_V(\mathbf{R}', \mathbf{R}, \tau) . \quad (127)$$

The kinetic part is not diagonal, taking $\nabla_{\mathbf{R}}^2$ definitely changes the state. The Green's function of the kinetic part is the solution of the *diffusion equation*

$$\frac{\partial G_T(\mathbf{R}', \mathbf{R}, \tau)}{\partial \tau} = -D \nabla_{\mathbf{R}}^2 G_T(\mathbf{R}', \mathbf{R}, \tau) \equiv \hat{T} G_T(\mathbf{R}', \mathbf{R}, \tau) . \quad (128)$$

This can be solved, and the matrix elements are

$$\begin{aligned} G_T(\mathbf{R}', \mathbf{R}, \tau) &= \langle \mathbf{R}' | e^{-\hat{T}\tau} | \mathbf{R} \rangle = \langle \mathbf{R}' | e^{-D\tau \nabla_{\mathbf{R}}^2} | \mathbf{R} \rangle \\ &= \int d\mathbf{K} d\mathbf{K}' \langle \mathbf{R}' | \mathbf{K}' \rangle \langle \mathbf{K}' | e^{-D\tau \nabla_{\mathbf{R}}^2} | \mathbf{K} \rangle \langle \mathbf{K} | \mathbf{R} \rangle \\ &= \int d\mathbf{K} d\mathbf{K}' e^{-i\mathbf{K}' \cdot \mathbf{R}'} \langle \mathbf{K}' | e^{-D\tau \nabla_{\mathbf{R}}^2} | \mathbf{K} \rangle e^{i\mathbf{K} \cdot \mathbf{R}} \\ &= \int d\mathbf{K} d\mathbf{K}' e^{-i\mathbf{K}' \cdot \mathbf{R}'} e^{-D\tau \mathbf{K}^2} \delta(\mathbf{K}' - \mathbf{K}) e^{i\mathbf{K} \cdot \mathbf{R}} \\ &= \int d\mathbf{K} e^{i\mathbf{K} \cdot (\mathbf{R} - \mathbf{R}')} e^{-D\tau \mathbf{K}^2} \\ &= (4\pi D\tau)^{-3N/2} e^{-\frac{(\mathbf{R} - \mathbf{R}')^2}{4D\tau}} . \end{aligned} \quad (129) \quad (130)$$

Thus G_T corresponds to a $3N$ -dimensional Gaussian spreading in time τ . To summarise, the short-time Green's function is

$$G(\mathbf{R}', \mathbf{R}, \tau) = (4\pi D\tau)^{-3N/2} e^{-\frac{(\mathbf{R}-\mathbf{R}')^2}{4D\tau}} e^{-(V(\mathbf{R})-E_T)\tau} + \mathcal{O}(\tau^2) . \quad (131)$$

In Path Integral Monte Carlo (PIMC) this form of the Green's function is usually called the **primitive approximation**. In the limit $\tau \rightarrow 0$ the Gaussian becomes narrower and approaches a Dirac delta function, thus this approximate G satisfies the boundary condition

$$G(\mathbf{R}', \mathbf{R}, 0) = \delta(\mathbf{R}' - \mathbf{R}) . \quad (132)$$

This is a natural requirement for the time evolution

$$\Psi(\mathbf{R}', \tau) = \int d\mathbf{R} G(\mathbf{R}', \mathbf{R}, \tau) \Psi(\mathbf{R}, 0) . \quad (133)$$

Ideally τ should be so small, that the approximation to the Green's function is more accurate than the statistical uncertainty. In practise one calculates the desired quantities with several values of τ in multiple MC runs, and in the end extrapolates the result to $\tau = 0$. Later Fig. 11 will demonstrate this procedure for a set of He-atom results.

In one dimension the small- τ expansion of the Green's function can be easily found with the aid of a sufficiently smooth function $f(x)$ ¹⁹. The Taylor expansion of $f(x)$ around $x = 0$ gives $(f^{(n)}(x))$ stands for the n th derivative at x)

$$\begin{aligned} \int_{-\infty}^{\infty} dx \left[\frac{1}{\sqrt{4\pi D\tau}} e^{-\frac{x^2}{4D\tau}} \right] f(x) &= \int_{-\infty}^{\infty} dx \left[\frac{1}{\sqrt{4\pi D\tau}} e^{-\frac{x^2}{4D\tau}} \right] \sum_{n=0}^{\infty} \frac{1}{n!} f^{(n)}(0) x^n \\ &= \int_{-\infty}^{\infty} dx \left[\sum_{n=0}^{\infty} \frac{1}{n!} D^n \tau^n \delta^{(2n)}(x) \right] f(x) , \end{aligned} \quad (134)$$

where the last form results after partial integrations that shift the derivatives from f to the Dirac delta function in each term. As a result,

$$G_T(\tau) = \frac{1}{\sqrt{4\pi D\tau}} e^{-\frac{(x-x')^2}{4D\tau}} = \delta(x - x') + D\tau \delta''(x - x') + \mathcal{O}(\tau^2) . \quad (135)$$

Combined with G_V , the time evolution of the wave function is

$$\Psi(x, \tau) = \Psi(x, 0) + [E_T - V(x)]\tau \Psi(x, 0) + D\tau \Psi''(x, 0) + \mathcal{O}(\tau^2) , \quad (136)$$

¹⁹R. Guardiola, *Lecture Notes in Physics: Microscopic Quantum Many-Body Theories and Their Applications*, Eds. J. Navarro and A. Polls, (Springer Verlag, Berlin 1998); Valencia lecture notes 1997.




Figure 6: The state of the system is represented by an ensemble of configurations.

so after stability has been reached, $\Psi(x, \tau) = \Psi(x, 0)$, we are left with the first-order terms that give back the Schrödinger equation

$$-D\Psi''(x, 0) + V(x)\Psi(x, 0) = E_T\Psi(x, 0) + \mathcal{O}(\tau) . \quad (137)$$

Hence in the zero- τ limit E_T is one of the eigenvalues of the Hamiltonian. Notice how a second order Green's function gives energies that are correct only to linear order

In the simulation the state $\Psi(\mathbf{R}, 0)$ is represented by an *ensemble of configurations*. In each of these N_C configurations we have N walkers located at points $\mathbf{R}$ (see Fig. 6). The Green's function (131) evolves the system via two processes:

1. Branching: According to Eq. (131) the present wave function $\Psi(\mathbf{R}, 0)$ is multiplied by the factor $e^{-(V(\mathbf{R})-E_T)\tau}$. So we calculate $V(\mathbf{R})$ for each configuration and see whether this factor is more or less than one, and create/destroy configurations accordingly. The walkers are not moved in this step: remember the $\delta(\mathbf{R}' - \mathbf{R})$!
2. Diffusion: Move from $\mathbf{R} \rightarrow \mathbf{R}'$ diffusively, in other words move each walker according to

$$r'_i = r_i + \eta , \quad (138)$$

where η is a d-dimensional Gaussian random variable with zero mean and variance $\sigma^2 = 2D\tau$. This corresponds to the kinetic part G_T of the Green's function (131).

In practise this algorithm is never used, because it has one big problem. Usually the potential is not bound from above, and walkers can easily move into positions, where the potential practically blows up. In step 1 we would need to create/destroy very many configurations and thus their total number

would fluctuate wildly. This is the DMC version of the large variance problem we encountered earlier in crude Monte Carlo. We need importance sampling to keep walkers away from the infinities of the potential, and to concentrate our efforts to regions of moderate $V(\mathbf{R})$.

Instead of $\Psi(\mathbf{R}, \tau)$, we seek to solve

$$f(\mathbf{R}, \tau) = \varphi_T(\mathbf{R})\Psi(\mathbf{R}, \tau) , \quad (139)$$

where $\varphi_T(\mathbf{R})$ is a time-independent trial wave function used for importance sampling. The imaginary-time Schrödinger equation (117) reads now

$$-\frac{\partial \Psi(\mathbf{R}, \tau)}{\partial \tau} = (\hat{H} - E_T)\Psi(\mathbf{R}, \tau) \quad (140)$$

$$\Leftrightarrow -\frac{1}{\varphi_T(\mathbf{R})} \frac{\partial f(\mathbf{R}, \tau)}{\partial \tau} = [-D\nabla_{\mathbf{R}}^2 + V(\mathbf{R}) - E_T] \frac{f(\mathbf{R}, \tau)}{\varphi_T(\mathbf{R})} \quad (141)$$

$$\Leftrightarrow -\frac{\partial f(\mathbf{R}, \tau)}{\partial \tau} = -D\varphi_T(\mathbf{R})\nabla_{\mathbf{R}}^2 \left[\frac{f(\mathbf{R}, \tau)}{\varphi_T(\mathbf{R})} \right] + [V(\mathbf{R}) - E_T] f(\mathbf{R}, \tau) \quad (142)$$

$$\Leftrightarrow -\frac{\partial f(\mathbf{R}, \tau)}{\partial \tau} = -D\nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + 2D \frac{\nabla_{\mathbf{R}} \varphi_T(\mathbf{R})}{\varphi_T(\mathbf{R})} \cdot \nabla_{\mathbf{R}} f(\mathbf{R}, \tau) - 2D \left(\frac{\nabla_{\mathbf{R}} \varphi_T(\mathbf{R})}{\varphi_T(\mathbf{R})} \right)^2 f(\mathbf{R}, \tau) + D \frac{\nabla_{\mathbf{R}}^2 \varphi_T(\mathbf{R})}{\varphi_T(\mathbf{R})} f(\mathbf{R}, \tau) + [V(\mathbf{R}) - E_T] f(\mathbf{R}, \tau) \quad (143)$$

$$\Leftrightarrow -\frac{\partial f(\mathbf{R}, \tau)}{\partial \tau} = -D\nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + D\nabla_{\mathbf{R}} \cdot [\mathbf{F}(\mathbf{R}) f(\mathbf{R}, \tau)] - D \frac{\nabla_{\mathbf{R}}^2 \varphi_T(\mathbf{R})}{\varphi_T(\mathbf{R})} f(\mathbf{R}, \tau) + [V(\mathbf{R}) - E_T] f(\mathbf{R}, \tau) , \quad (144)$$

so finally we have

$$\boxed{-\frac{\partial f(\mathbf{R}, \tau)}{\partial \tau} = -D\nabla_{\mathbf{R}}^2 f(\mathbf{R}, \tau) + D\nabla_{\mathbf{R}} \cdot [\mathbf{F}(\mathbf{R}) f(\mathbf{R}, \tau)] + [E_L(\mathbf{R}) - E_T] f(\mathbf{R}, \tau) .} \quad (145)$$

Here we have defined the *drift force* (or *drift velocity, quantum force*)

$$\mathbf{F}(\mathbf{R}) \equiv \nabla_{\mathbf{R}} \ln |\varphi_T(\mathbf{R})|^2 = 2 \frac{\nabla_{\mathbf{R}} \varphi_T(\mathbf{R})}{\varphi_T(\mathbf{R})} , \quad (146)$$

and the *local energy* (this was used also in VMC)

$$E_L(\mathbf{R}) \equiv \varphi_T(\mathbf{R})^{-1} \hat{H} \varphi_T(\mathbf{R}) . \quad (147)$$

Notice that these are calculated from the *trial state* $\varphi_T(\mathbf{R})$, so the expressions for $\mathbf{F}(\mathbf{R})$ and $E_L(\mathbf{R})$ are usually derived in analytic form before starting the DMC time evolution. In the appendix we show how to evaluate the local energy and the drift force for the McMillan trial wave function designed for two- or three-dimensional liquid ^4He . The first term in Eq. (146) is the usual diffusion term. The drift force guides the diffusion process to regions where $\varphi_T(\mathbf{R})$ is large (direction of gradient).

The Green's function corresponding to Eq. (146) is a modification of Eq. (131). One possible solution is

$$G(\mathbf{R}', \mathbf{R}, \tau) = (4\pi D\tau)^{-3N/2} e^{-\frac{(\mathbf{R}-\mathbf{R}'-\mathbf{F}(\mathbf{R})D\tau)^2}{4D\tau}} e^{\tau(E_L(\mathbf{R})-E_T)} + \mathcal{O}(\tau^2) . \quad (148)$$

This is not unique: Should we drift the walkers before or after we've diffused them ($\mathbf{F}(\mathbf{R})$ or $\mathbf{F}(\mathbf{R}')$)? This is once again a commutation problem, drift and diffusion don't commute. Often one uses $\frac{1}{2}[\mathbf{F}(\mathbf{R}) + \mathbf{F}(\mathbf{R}')]$ and similarly in the local energy $\frac{1}{2}[E_L(\mathbf{R}) + E_L(\mathbf{R}')]$. We'll talk about modifications later.

Notice that (148) has no explicit dependence on $V(\mathbf{R})$; It's up to our trial wave function $\varphi_T(\mathbf{R})$ to lead to a drift force which guides the walkers to configurations where $E_L(\mathbf{R})$ has low values and where $V(\mathbf{R})$ is not causing problems.

6.2 DMC algorithms

The Diffusion Monte Carlo (DMC) algorithm corresponding to the Green's function (148) has now **drift, diffusion and branching**, Here is one possible procedure²⁰:

1. Initialise N_C configurations (typically 100-500), where walkers are distributed according to $|\varphi_T(\mathbf{R})|^2$ (not necessarily normalised to be a PDF). We don't want them all destroyed in the branching process, so we need to generate as diverse a set as possible. A typical way is to use VMC to sample from $\varphi_T(\mathbf{R})$ and pick from the VMC run configurations now and then to our **zeroth generation of configurations**. The total number of configurations in a generation is its **population**.
2. Each configuration of the present generation is taken in turn. Move the i^{th} walker of a configuration according to

$$\mathbf{r}_i^{\text{new}} = \mathbf{r}_i + D\tau\mathbf{F}(\mathbf{R}) + \eta , \quad (149)$$

²⁰D. Ceperley, Reynolds ...

where η is a d -dimensional Gaussian random variable with mean zero and variance $2D\tau$. If you have a function `gauss` that gives Gaussian-distributed random numbers with mean zero and variance 1, you have to multiply it with $\sqrt{2D\tau}$. If $\mathbf{r}(\mathbf{k}, i, j)$ is the k :th component of the position vector of the i :th particle in j :th configuration and $\mathbf{F}(\mathbf{k}, i, j)$ is the corresponding component of the drift force, then the move (149) is

$$\mathbf{r}(\mathbf{k}, i, j) = \mathbf{r}(\mathbf{k}, i, j) + D * \tau * \mathbf{F}(\mathbf{k}, i, j) + \text{sqrt}(2. * D * \tau) * \text{gauss}()$$

Don't re-evaluate $\mathbf{F}$ until all coordinates in the configuration has been updated!

3. Calculate a weight for the move,

$$ratio = \frac{|\varphi_T(\mathbf{R}')|^2 G(\mathbf{R}', \mathbf{R}, \tau)}{|\varphi_T(\mathbf{R})|^2 G(\mathbf{R}, \mathbf{R}', \tau)} , \quad (150)$$

and accept the move $\mathbf{R} \rightarrow \mathbf{R}'$ with probability $\min(1, ratio)$; this is just a Metropolis question similar to the one used in VMC²¹. This step guarantees, that the algorithm satisfies the detailed balance condition. Moreover, it is the only way to make sure that we are sampling the correct distribution at finite τ . This accept-reject test is easy to convert back to the test appropriate for VMC - a very useful property.

In DMC application that are second order or higher (to be discussed later) this rejection step is sometimes left out. The reason is that it may lead to a new problem of "persistent configurations", configurations that get stuck. These lead to erraneous energy vs. τ results and to wrong extrapolations to zero τ .

4. Take each configuration in turn and make N_{copies} copies of it, where

$$N_{\text{copies}} = \text{int} \left[e^{-\tau(E_L(\mathbf{R}) - E_T)} + \xi \right] , \quad (151)$$

where $\xi \in U[0, 1]$ (int= "integer part of"). We want to have *on average* $e^{-\tau(E_L(\mathbf{R}) - E_T)}$ copies of the configuration $\mathbf{R}$. If this is not an integer, we treat the remainder as a probability: if it's bigger than ξ round the number upward. In math,

$$\langle e^{-\tau(E_L(\mathbf{R}) - E_T)} \rangle = \langle \text{int} \left[e^{-\tau(E_L(\mathbf{R}) - E_T)} + \xi \right] \rangle . \quad (152)$$

²¹The Greens function in the *ratio* is the transfer matrix in the Metropolis-Hastings detailed balance formula, Eq. (72)

In case $N_{\text{copies}} = 0$ that configuration is killed; This is one reason why we start with large N_C , the other one is that small number of configurations introduces *bias* to the result.

5. Get an estimate the ground-state energy E_0 by taking the average of the local energies of *current* generation.
6. Adjust the trial energy E_T so that the population is roughly stable. You can use, for example, a feedback mechanism

$$E_T = E_0 + \kappa \ln \frac{N_{C\text{target}}}{N_C}, \quad (153)$$

where E_0 is the best guess for the ground-state energy we have so far, $N_{C\text{target}}$ is the population we'd like to have, and N_C is the one we have now. Parameter κ should have the smallest positive value that does the job; too large κ will introduce bias to the result.

Bias means that something is causing the assumption of randomness to fail. The source of bias may be statistical (e.g. bad pseudo-random number generator) or systematic (e.g. non-random choices made somewhere).

Some dislike the idea of replicating configurations as in step 3 in the above algorithm. Indeed, there is a way to avoid this, known as **reweighting**. Instead of actually making copies or killing configurations, we attach to each one of them a weight which describes essentially the relative number of descendants that configuration would have had in the replicating algorithm up to that moment. We give that weight to all quantities measured from that configuration. In a reweighting DMC algorithm one would replace the replication step with the calculation of the average local energy over the configuration by adding up each local energy with the weight $\exp[-(E_L(\mathbf{R}) - E_T)\tau]$. For i :th configuration one computes

```
...
weight(i) = weight(i)*exp(-(EL-ET)*tau)    ! update weight
E          = E + weight(i)*EL                ! weight contributions
...
```

Here EL is the local energy of the i :th configuration, $E_L(\mathbf{R}_i)$, and ET is the trial or reference energy that estimates the ground state energy or the energy of some other desired eigenstate. This method will be applied in section 6.8 to a ^4He atom. In general, reweighting is known to lead to slightly larger variances than replicating.

The problem with reweighting is that the cumulative weights tend to either decrease toward zero or to increase without limit. A configuration that

has a tiny weight is useless to the computation and one with extremely high weight will dominate over all other configurations leading to poor statistics. Hence in a longer DMC run one has to **renormalise the weights** in some fashion, which is not causing bias in itself. You cannot simply split half configurations whose weight exceeds a fixed value, say 2. A rule of thumb is that unless every configuration has some chance to be split your algorithm is biased. What you can do is you can choose the configurations to be split by favouring the high-weight ones. This can be done, for example, by a wheel of fortune depicted in Fig. 7. The change that you hit a higher-weight configuration is larger. Next, pick up at random any other configuration and replace it with the chosen configuration and let both configurations now carry half of the weight the chosen configuration had before splitting.




Figure 7: A way to choose a configuration, so that every one has a change but higher-weighted ones are more probable. The lengths of the angular segments are the weights of the configurations. The angle of the pointer is chosen at random.

Running a DMC algorithm in VMC mode The time step must be kept small and the acceptance ratio of a good DMC run is always $> 99.99\%$. The Metropolis step 3 in the algorithm may seem unnecessary, and often very good results can be obtained without it. A serious algorithm should anyhow contain it. Another nice feature by this Metropolis test is that the algorithm can now be run as VMC: If we leave out the Green's function in test (150) (set it to unity), drift in step 2 as well as the branching step 3, we are left with a VMC algorithm. And a very effective one, it has N_C independent configurations.

Why would one want to run VMC if one has a DMC code? Because

1. Unless you have a sophisticated DMC code that gives you pure estimators (see section 6.3 below), you need the expectation values $\langle \varphi_T | O | \varphi_T \rangle$ of all quantities O that you wish to evaluate but don't commute with the Hamiltonian, $[O, H] \neq 0$.
2. The DMC code needs as input a set of *uncorrelated* configurations. I don't say "random", because if you start with a completely random set, you may soon realise that none of them had any descendant and they were all killed. This may happen, for example, if you simulate N hard-core particles and your "random" configurations all have overlapping hard-core regions: infinite local energy and no offspring. VMC gives the trial wave function a change to set things almost correct.

Commonly one starts with just one configuration, runs the VMC part for a while and picks up the N_C *uncorrelated* configurations. These are then run for a while using only VMC to get a good estimate for the energy of the trial wave function,

$$E_{\text{VMC}} = \langle \varphi_T | H | \varphi_T \rangle . \quad (154)$$

VMC steps are much faster to do than DMC steps, because the Gaussian step η can be much larger: τ can be much larger, because VMC doesn't have to worry about any Green's function: no short time estimates are used.

The DMC evolution needs to be run for a while to let the excited states die out. The estimate E_0 will be fluctuating around a fixed value and your subsequent generations will (hopefully) be sampling the distribution $f(\mathbf{R}, \tau) = \varphi_T(\mathbf{R})\Psi(\mathbf{R}, \tau)$. It's imaginary time to start collecting data for ground-state averages.

6.3 Mixed and pure estimators

Let's take a closer look at how averages are calculated using an importance sampling DMC algorithm. We need to be cautious, because the algorithm is *not* sampling the ground state directly, but the product state $f(\tau) = \varphi_T \Psi(\tau)$ that approaches $\varphi_T \phi_0$ as $\tau \rightarrow \infty$.

At time τ the average of the local energy $E_L(\tau)$ is (I'll drop the $\mathbf{R}$'s in the arguments to simplify notation)

$$\begin{aligned} E_{\text{mixed}} &= \frac{\int d\mathbf{R} f(\tau) E_L}{\int d\mathbf{R} f(\tau)} = \frac{\int d\mathbf{R} f(\tau) [\varphi_T^{-1} H \varphi_T]}{\int d\mathbf{R} f(\tau)} \\ &\rightarrow \frac{\int d\mathbf{R} \phi_0 H \varphi_T}{\int d\mathbf{R} \phi_0 \varphi_T} . \end{aligned} \quad (155)$$

This called a **mixed estimator**, because it is *not* the expectation value $\langle \phi_0 | H | \phi_0 \rangle$, but the mixed matrix element $\langle \varphi_T | H | \phi_0 \rangle$. We still get the correct ground-state energy directly out of importance sampling DMC, because

$$E_{\text{mixed}} = \frac{\int d\mathbf{R} \phi_0 \sum_i a_i E_i \phi_i}{\int d\mathbf{R} \phi_0 \sum_i a_i \phi_i} = E_0 ,$$

where the last step follows from orthogonality of the energy eigenstates. So the mean value of E_L is the ground-state energy also if importance sampling is used. What about the variance? We have

$$\begin{aligned} \sigma_E^2 &= \langle E_L^2 \rangle - E_0^2 = \frac{\int d\mathbf{R} f(\tau) E_L^2}{\int d\mathbf{R} f(\tau)} - E_0^2 \\ &\rightarrow \frac{\int d\mathbf{R} \varphi_T \phi_0 [\varphi_T^{-1} \hat{H} \varphi_T] [\varphi_T^{-1} \hat{H} \varphi_T]}{\int d\mathbf{R} \varphi_T \phi_0} - E_0^2 \\ &= \frac{\int d\mathbf{R} \phi_0 \left(\hat{H} \varphi_T \right) \varphi_T^{-1} \left(\hat{H} \varphi_T \right)}{\int d\mathbf{R} \varphi_T \phi_0} - E_0^2 \\ &= \frac{\int d\mathbf{R} \phi_0 \left(\sum_i a_i E_i \phi_i \right)^2 \left(\sum_i a_i \phi_i \right)^{-1}}{a_0} - E_0^2 . \end{aligned} \tag{156}$$

If we have the perfect trial wave function ($\varphi_T = \phi_0$; $a_i = 0$ for $i \neq 0$), then $\sigma_E^2 = 0$. Same thing happens if you are evaluating the energy of an excited state and use exactly the correct wave function as φ_T . This is the famous **zero variance property**:

The variance in the DMC energy will vanish for a trial wave function, which coincides with any eigenstate of the Hamiltonian.

This is valid also for VMC, as one can easily prove. For a trial wave function that is an exact eigenstate this is trivial: The local energy will be just a constant, so no matter how you move your walkers you will – of course – always get that constant as the energy.

The better the trial wave function, the smaller the error bars .²²

Let's see what happens if you try to evaluate a quantity $\hat{A}$ that does *not* commute with the Hamiltonian. Then you get also

$$\begin{aligned} A_{\text{mixed}} &= \frac{\int d\mathbf{R} f(\tau) \left[\varphi_T^{-1} \hat{A} \varphi_T \right]}{\int d\mathbf{R} f(\tau)} \\ &= \frac{\int d\mathbf{R} \phi_0 \hat{A} \varphi_T}{\int d\mathbf{R} \phi_0 \varphi_T} = \frac{\int d\mathbf{R} \phi_0 \hat{A} \sum_i a_i \phi_i}{a_0} , \end{aligned} \tag{157}$$

²²Old saying in the Monte Carlo Jungle.

but this is where the road ends, because ϕ_i 's are not eigenstates of $\hat{A}$. In this situation one commonly uses the so-called **extrapolated estimator**

$$\langle \phi_0 | A | \phi_0 \rangle \approx 2\langle \varphi_T | \hat{A} | \phi_0 \rangle - \langle \varphi_T | A | \varphi_T \rangle , \quad (158)$$

which assumes that the mixed estimator is “halfway” between the true ground-state and the VMC expectation values. Then one does a kind of linear extrapolation. More rigorously, one makes a (functional) Taylor expansion of the DMC and VMC estimators in terms of $\Delta = \phi_0 - \varphi_T$,

$$\begin{aligned} \langle \phi_0 | \hat{A} | \varphi_T \rangle &= \langle \phi_0 | \hat{A} | \phi_0 \rangle + \int d\mathbf{R} \psi_0 \left[\langle \phi_0 | \hat{A} | \phi_0 \rangle - \hat{A} \right] \Delta + \mathcal{O}(\Delta^2) \\ \langle \varphi_T | \hat{A} | \varphi_T \rangle &= \langle \phi_0 | \hat{A} | \phi_0 \rangle + 2 \int d\mathbf{R} \psi_0 \left[\langle \phi_0 | \hat{A} | \phi_0 \rangle - \hat{A} \right] \Delta + \mathcal{O}(\Delta^2) , \end{aligned} \quad (159)$$

so a combination of these gives

$$\langle \phi_0 | \hat{A} | \phi_0 \rangle = 2\langle \phi_0 | \hat{A} | \varphi_T \rangle - \langle \varphi_T | \hat{A} | \varphi_T \rangle + \mathcal{O}(\Delta^2) . \quad (160)$$

Extrapolated estimators are needed very frequently, you cannot even determine the kinetic and potential energies directly: $\hat{H}$ commutes neither with $\hat{T}$ nor with $\hat{V}$.

There is also a way to get more accurate ground-state expectation values for quantities that don't commute with $\hat{H}$, known as **pure estimators**. A pure estimator means any estimator that is unbiased, *i.e.*, not affected by the choice of the trial wave function. The total energy in DMC is independent of φ_T , so it *is* already a pure estimator. But quantities that don't commute with the Hamiltonian need special treatment to get their pure estimator. I'm not going to describe the method in any detail, just mention that one needs to keep account of how long the configurations in the DMC run survive and how many descendants they have²³.

²³See, for example, Jordi Boronat's work in the 90's, published in Phys. Rev. **B**.

Growth estimator In DMC, the ground state energy can be obtained at least from two sources: (i) as the average of the local energies $E_L(\mathbf{R})$ as in VMC or (ii) from the so-called growth estimator. The latter is based on the notion, that after the excited states have died out, the population (number of configurations) evolves during time τ according to the equation

$$N(\tau) = e^{(E_0 - E_T)\tau} N(0) . \quad (161)$$

Solving for E_0 one finds

$$E_0 = E_T + \frac{1}{\tau} \ln \frac{N(\tau)}{N(0)} . \quad (162)$$

This looks like the population control equation (153), just inverted to give E_0 . Indeed, one has to refrain from adjusting the population manually during the evaluation period in Eq. (162). In practical application of the growth estimator it is profitable to have populations $N(\tau)$ and $N(0)$ evaluated at as distinct times as possible. The relation to the local energy is clear, since the population changes via branching, which is in controlled by the local energy. The growth estimator is unpopular due to the large fluctuations in the population during a simulation, resulting in a large variance.

The significance of branching in the DMC algorithm In the importance sampling algorithm the total energy is the average over local energies $E_L(\mathbf{R}) = \varphi_T^{-1} H \varphi_T$. Hence the functional form of E_L is completely determined by the trial wave function - which could be very poor. The reason why averaging over E_L in DMC does give the true ground-state energy is that the configurations $\mathbf{R}$ where one evaluates $E_L(\mathbf{R})$ are biased due to branching²⁴.

The situation is even more dramatic if one doesn't use importance sampling at all. Remember that we introduced it to DMC only to improve efficiency. Then the total energy is the average of potential energies $V(\mathbf{R})$ at the chosen configurations $\mathbf{R}$!. Now the branching has a huge job in biasing the configurations so, that the average of these potential energies gives the total ground-state energy. We get the correct ground-state energy by computing the value of a *wrong* function at well-chosen points, $\langle V(\mathbf{R}) \rangle_{\mathbf{R}} = E_0$!

²⁴Of course, reweighting can be used to achieve the same effect.

6.4 Variance optimisation

Of course we could find the lowest energy trial state, and that would, in some sense, be the best trial state for the ground state. The trouble with this most naïve optimisation scheme is that we need to evaluate $\langle E \rangle$ for many sets of parameters, and each one may take a long time to evaluate, depending on how long the VMC calculation takes. Still after a lengthy evaluation, there will be some error in each energy, and from the energy vs. parameter plot it will be nearly impossible to determine the location of the minimum accurately. For recent articles on wave function optimisation see Assaraf and Caffarel²⁵ and Umrigar and Filippi²⁶

Correlated sampling Statistical fluctuations can render a direct minimisation unstable. One can hope, that there is some “systematics” behind the fluctuations and that the trial wave function with different parameters α give rise to “similar” fluctuations. If this is the case, then the cause of our trouble with fluctuations is that using *different* configurations $\mathbf{R}$ for different α mixes this systematics. The idea of using the *same set of configurations* to evaluate the expectation values of a several version of the same quantity is known as correlated sampling. If the fluctuations are somehow systematic, then

The differences of the correlated averages are much more accurate than the averages themselves.

One is interested in small energy differences, for example, in dealing with ³He impurities in liquid ⁴He or in pinpointing the Wigner crystallisation or paramagnetic-to-ferromagnetic transitions in electron gas. In these cases all interesting data is hidden in minute energy differences. Separate computation of electron gas and Wigner crystal energies as a function of density easily lead to ridiculously large uncertainties in the Wigner crystallisation density.

To elucidate correlated sampling, consider the expectation value of the local energy with different parameters α . It can be written in the form

$$\begin{aligned} \langle E(\alpha) \rangle &\equiv \frac{\int d\mathbf{R} |\varphi_T(\mathbf{R}, \alpha)|^2 E_L(\mathbf{R}; \alpha)}{\int d\mathbf{R} |\varphi_T(\mathbf{R}, \alpha)|^2} \\ &= \frac{\int d\mathbf{R} |\varphi_T(\mathbf{R}, \alpha_0)|^2 \left[\frac{|\varphi_T(\mathbf{R}, \alpha)|^2}{|\varphi_T(\mathbf{R}, \alpha_0)|^2} E_L(\mathbf{R}; \alpha) \right]}{\int d\mathbf{R} |\varphi_T(\mathbf{R}, \alpha_0)|^2 \frac{|\varphi_T(\mathbf{R}, \alpha)|^2}{|\varphi_T(\mathbf{R}, \alpha_0)|^2}}. \end{aligned} \quad (163)$$

²⁵R. Assaraf and M. Caffarel, Phys. Rev. Lett. 83, 4682 (1999).

²⁶C. J. Umrigar and C. Filippi, Phys. Rev. Lett. 94, 150210 (2005).

Using the weights defined in Eq. (167) and approximating the integrals as sums over M configurations (as we always do in MC) we get

$$\langle E(\alpha) \rangle \approx \frac{\sum_{i=1}^M |\varphi_T(\mathbf{R}_i; \alpha_0)|^2 [w(\mathbf{R}_i) E_L(\mathbf{R}_i; \alpha)]}{\sum_{i=1}^M |\varphi_T(\mathbf{R}_i; \alpha_0)|^2 w(\mathbf{R}_i)} . \quad (164)$$

For every value of α , this can be evaluated by a random walk $\mathbf{R}_i$ ²⁷ that samples $\varphi_T(\mathbf{R}_i; \alpha_0)$, *i.e.*, a wave function at *fixed* α_0 . Moreover, in correlated sampling you use the *same sequence of random numbers*, so the random walks $\mathbf{R}_i$ are the same for every α . Because of the same walk, the results $\langle E(\alpha) \rangle$ will be correlated. Their variances are larger than for uncorrelated sampling, but if the fluctuations “cancel in the same direction”, the differences

$$\langle E(\alpha_1) \rangle - \langle E(\alpha_2) \rangle , \quad (165)$$

are more accurate. In most cases we just want to know if we are going upwards or downwards.

One general notion is very useful, namely that the fluctuations of $\langle ab \rangle$ are usually much larger than those of $\langle ab \rangle - \langle a \rangle \langle b \rangle$ (“covariance”). This is especially true if a and b are uncorrelated, in which case the covariance is zero. Umrigar and Filippi (see a footnote) use this idea and *insert* terms like $\langle a \rangle \langle b \rangle$ that will vanish for an infinite sample to get expressions that contain only covariances.

Objective functions The zero variance property can be used as an optimisation criteria in finding a better trial wave function. It is often a more stable process to optimise the variance of the energy, not the energy itself. One possible quantity to minimise – the **objective function** – is²⁸

$$\sigma_{\text{opt}}^2 = \frac{\sum_{i=1}^{N_{\text{opt}}} [E_L(\mathbf{R}_i) - E_{\text{guess}}]^2 w(\mathbf{R}_i)}{\sum_{i=1}^{N_{\text{opt}}} w(\mathbf{R}_i)} , \quad (166)$$

where the weight function is

$$w(\mathbf{R}_i) = \left| \frac{\varphi_T(\mathbf{R}_i)}{\varphi_{T,\text{fixed}}(\mathbf{R}_i)} \right|^2 . \quad (167)$$

²⁷A set of configurations $\mathbf{R}_i$ is usually generated by performing a random walk, so $\mathbf{R}_i$ is a “set of configurations” or a “random walk” or a “random path” depending on how one visualises the thing. Of course, a “set of configurations” is slightly more general and may even be non-random.

²⁸For other objective functions and further discussion, see P. R. C. Kent, R. J. Needs, and G. Rajagopal, Phys. Rev. **B59** (1999) 12344.

Instead of taking a new set of randomly picked configurations one is using the *same set of configurations* $\mathbf{R}_i = \{\mathbf{R}_1, \mathbf{R}_2, \dots, \mathbf{R}_{N_{\text{opt}}}\}$, which sample a *fixed* $\varphi_{T,\text{fixed}}$. This is again correlated sampling.

About E_{guess} , Umrigar *et al.*²⁹ say (their E_g is our E_{guess}):

The initial value of E_g in Eq. (1) is taken to be the expectation value of the energy for the initial trial wave function (HF-J or MCSFC-J). Once an improved trial wave function is found, another MC run is performed with this improved function and the expectations value of the energy of this MC run is used as the new E_g . Typically only two iterations are required to obtain a self-consistent value of E_g .

Let us test the process for the hydrogen atom, where we can compute σ_{opt}^2 exactly for the trial wave function in the limit $N_{\text{opt}} \rightarrow \infty$ in (166). We use the trial wave function

$$\varphi_T(\mathbf{R}) = e^{-\alpha r} , \quad (168)$$

where r is the distance of the electron from the proton (at origin). The local energy in atomic units is

$$E_L(\mathbf{R}) = -\frac{1}{2}\alpha + \frac{\alpha - 1}{r} , \quad (169)$$

so the exact ground-state energy is $E_0 = -0.5$ (Hartrees) at $\alpha = 1$. Simple integration gives

$$\sigma_{\text{opt}}^2 = \frac{5}{4}\alpha^4 - 3\alpha^3 + \alpha^2(2 - E_{\text{guess}}) + 2\alpha E_{\text{guess}} + E_{\text{guess}}^2 . \quad (170)$$

Figure 8 shows σ_{opt}^2 as function of α for few guessed energies E_{guess} . Notice that if $E_{\text{guess}} > E_0$ there is another, lower minimum at low α , and in this case the optimisation process is unstable and even infinite number of configurations won't help!

In short, the process of finding the minimum of σ_{opt}^2 can be *unstable* if $E_{\text{guess}} > E_0$, and the proposal by Umrigar *et al.* cited above should be applied with care. Even so, the variance σ_{opt}^2 does have an absolute minimum at the exact ground state, because then $E_L(\mathbf{R}) = E_0$ for all $\mathbf{R}$ and

$$\begin{aligned} \sigma_{\text{opt}}^2 &\geq \frac{\sum_{i=1}^{N_{\text{opt}}} [E_0 - E_{\text{guess}}]^2 w(\mathbf{R}_i)}{\sum_{i=1}^{N_{\text{opt}}} w(\mathbf{R}_i)} \\ &= (E_0 - E_{\text{guess}})^2 \geq 0 . \end{aligned} \quad (171)$$

²⁹C. J. Umrigar, K. G. Wilson, and J. W. Wilkins, Phys. Rev. Lett. **60** (1988) 1719.




Figure 8: The *exact* objective function (166) in the limit $N_{\text{opt}} \rightarrow \infty$ as a function of α for a hydrogen atom with trial wave function (168). The guessed energies are $E_{\text{guess}} = 0.1, 0.0, -0.1, \dots, -0.8$, the curves for $E_{\text{guess}} > E_0$ are those with another minimum at small α . The bold line shows the case $E_{\text{guess}} = E_0 = -0.5$.




Figure 9: The figure illustrates the serious misconception that energy optimisation and variance optimisation are equivalent ways to find the best parameters. The first 25000 steps are VMC data from energy optimised Padé–Jastrow wave function of He atom with just one parameter ($a = 0.1437$): The energy is $E = -2.878$ and the variance $\sigma^2 = 0.11$. The last 25000 step are variance optimised data ($a = 0.3334$): $E = -2.869$ and $\sigma^2 = 0.08$. So the higher energy result has lower variance and vice versa!

Hence the limiting value $\sigma_{\text{opt}}^2 = 0$ requires also $E_{\text{guess}} = E_0$. This can be seen also in the case of hydrogen in Fig. 8, where the zero value at $\alpha = 1$ corresponds to exactly $E_{\text{guess}} = E_0 = -0.5$.

One point worth remembering is, that if your trial wave function has exactly the correct form with n unknown parameters, then you need just n independent configurations to find out their values! It's simple fitting or – even better – solving a set of linear equations. With imperfect trial wave functions there is no guarantee, that the parameters giving smallest variance of some form, like σ_{opt}^2 , correspond to the best trial wave function. This can be seen by considering an optimisation process, where one of the configurations happens to give a much larger contribution than the rest. In that case the variance would be close to zero, but the parameters would be very poor indeed! In some cases – one is the helium atom problem with two 1S orbitals – your optimisation process may fail to converge to the “best” value. In that case you might try another objective function. If the potential energy diverges for some configurations, then you'd better take care of the cusp conditions (see appendix). See section 6.8 for more details about the He-atom problem.

Figure 9 shows how energy optimisation and variance optimisation may disagree about what state is closest to the true ground state. The fact that a lower energy state may have higher variance does not contradict the zero-variance property, it merely indicates that the search of the “best” wave function in the limited Hilbert space spanned by the possible trial wave functions has local minima in either energy or variance or both.

To conclude, variance optimisation is not a black box that always produces the optimal parameters. But if it does, then you have the best possible trial wave function and your DMC code will be much more efficient.

6.5 Fixed node and released node methods

The ground state of a many-body Bose system can always be chosen to be real and positive, so the quantity $f(\mathbf{R}) = \varphi_T(\mathbf{R})\phi_0(\mathbf{R})$ can be safely interpreted as density in the diffusion process. Excited states and Fermi systems are more problematic, because the exact wave function has nodes (zeros)³⁰. The negative areas themselves are not bothering us, because one can define

$$\phi(\mathbf{R}) = \phi^+(\mathbf{R}) - \phi^-(\mathbf{R}) , \quad (172)$$

where both ϕ^+ and ϕ^- are positive and use them as densities. The Schrödinger equation can be written separately for both parts. Similar decomposition is done for the trial wave function, too. Eq. (172) is trivially satisfied by, for example,

$$\phi^\pm(\mathbf{R}) = \frac{1}{2} [|\phi(\mathbf{R})| \pm \phi(\mathbf{R})] . \quad (173)$$

Now that I mentioned Fermi systems, it would seem appropriate to give an example of what the trial wave function might look like. The Slater-Jastrow wave function is given by

$$\varphi_T(\mathbf{R}) = \det(A_\uparrow) \det(A_\downarrow) e^{U(\mathbf{R})} , \quad (174)$$

where $A_\uparrow$ and $A_\downarrow$ are Slater matrices made of spin up and spin down single-particle orbitals; the Slater matrix is

$$A = \begin{pmatrix} \varphi_1(\mathbf{r}_1) & \varphi_1(\mathbf{r}_2) & \dots & \varphi_1(\mathbf{r}_N) \\ \varphi_2(\mathbf{r}_1) & \varphi_2(\mathbf{r}_2) & \dots & \varphi_2(\mathbf{r}_N) \\ \vdots & \vdots & \ddots & \vdots \\ \varphi_N(\mathbf{r}_1) & \varphi_N(\mathbf{r}_2) & \dots & \varphi_N(\mathbf{r}_N) \end{pmatrix} . \quad (175)$$

Orbitals φ_i can be, *e.g.*, plane waves or molecular orbitals. The term $e^{U(\mathbf{R})}$ is the Jastrow factor (see the appendix for details).

The problem is that we usually have very little or no idea where the nodes of ϕ are. The trial wave function is supposed to approximate the true solution, so we can choose one which has nodes. But this choice is forced upon the final solution:

$$\varphi_T(\mathbf{R}_0) = 0 \Rightarrow f(\mathbf{R}_0, \tau) = \varphi_T(\mathbf{R}_0)\psi(\mathbf{R}_0, \tau) = 0 , \quad (176)$$

³⁰We really need to impose those nodes somehow, without them the Bose case will converge to the ground state and the Fermi case will not obey the Pauli principle and cease to be a Fermi case.

and no amount of imaginary time evolution will make the zeroes to change a bit. If, as usual, the nodes of φ_T and ϕ do not coincide, then near the nodes we would be badly undersampling ϕ ; the result will certainly be biased and unreliable. Anyhow, this is exactly what is done in the **fixed node method**³¹.

One can see how the walkers avoid the nodes by looking at the drift force $\mathbf{F}(\mathbf{R})$, see Fig. 10. The drift is away from the nodes (nodal surfaces); At the nodes $\mathbf{F}(\mathbf{R})$ is infinite and walkers shouldn't be able to cross nodal surfaces. Due to the finite time step τ some walkers do cross, and usually one just deletes such spurious configurations from the set. This is not the best way, and in literature it's argued that it would be better just to reject moves that would cross a nodal surface. Of course, if the nodes of φ_T happened




Figure 10: The drift force near the nodes of the wave function (trial or exact) and the decomposition proposed in Eq.(173).

to coincide with the exact ones, then ϕ would also be correctly sampled in the fixed node DMC. Until recently it was believed, that with the fixed-node DMC one always gets a variational upper bound for the exact eigenvalues. However, Foulkes, Hood and Needs³² showed that this is true only if the state in question is non-degenerate or has only accidental degeneracy.

Many of the shortcomings of the fixed node method has been known ever since it was first applied, and something had to be done to somehow let the nodes of φ_T move towards the true nodes of ϕ as the computation progresses. In the **released node method** one lets the walkers to cross nodal surfaces. One assigns a “colour” (red/blue) or a sign (± 1) to each walker, which keeps track of whether the walker has crossed the nodal surface odd or even times. Each walker crossing a nodal surface “pulls” the surface with it to some extent. Only to some extent, because a single walker has only statistical

³¹This section relies heavily on Jordi Boronat’s section in Valencia lecture notes (1997).

³²W. M. C. Foulkes, R. Q. Hood and R. J. Needs, Phys. Rev. B **60**, 4558 (1999).

significance, and a red walker in the midsts of blue ones does not mean the true nodal surface makes a wild loop to engulf that solitary walker. Thus we need to restrict the effect of a single walker, and usually it is done by letting the walker carry it's original colour only for a finite "lifetime" τ_r . Just like a hot chicken cooling in a box of icy chickens.

Next one must let the walkers to cross nodal surfaces, not only spuriously but in honest. One way is to use a slightly modified trial function, such as

$$\varphi'_T(\mathbf{R}) = \sqrt{|\varphi_T(\mathbf{R})|^2 + a^2} . \quad (177)$$

With proper value of a this approaches φ_T far away from the nodes. Parameter a controls the amount of flow across nodal surfaces, and for larger a we take shorter lifetime τ_r and vice versa.

The fact that we have tinkered with our trial wave function can be compensated by attaching to each walker a weight

$$W(\mathbf{R}) = \sigma(\mathbf{R}) \frac{|\varphi_T(\mathbf{R})|}{\varphi'_T(\mathbf{R})} , \quad (178)$$

where the sign carried by the walker, $\sigma(\mathbf{R})$, is $+1$ (-1) for even (odd) number of crossings. It can be shown that using $\varphi'_T(\mathbf{R})$ affects only the variance, and not the mean of the total energy. The released node energy has also the lifetime τ_r as a parameter, and it's customary to define

$$E_{\text{RN}}(\tau_r) = \frac{\sum_{\tau \leq \tau_r} E_L(\mathbf{R}) W(\mathbf{R})}{\sum_{\tau \leq \tau_r} W(\mathbf{R})} , \quad (179)$$

where the sums are over walkers that have survived and have lifetime less than τ_r . We are interested in the asymptotic region of large τ_r , and one usually plots the sequence of energies for $\tau_r^1 < \tau_r^2 < \dots < \tau_r^{\text{max}}$, which reveals the systematic error in the released-node energy.

6.6 Excited states

Here we take a closer look at how excited-state eigenvalues could be obtained from a DMC run. The list is by no means complete, the question has no satisfactory solution and is under constant research.

Using orthogonality $\langle \varphi_T | \phi_0 \rangle = 0$ If the trial wave function is orthogonal to the ground state, the DMC process stabilises to an excited state with proper choice of the trial energy E_T . In the large- τ limit all high-energy

states have died out and we are left with the two lowest ones. According to (172) we have

$$\Psi(\tau \rightarrow \infty) = \phi = \phi^+ - \phi^- , \quad (180)$$

where the two components are

$$\phi^\pm = c_0 \phi_0 e^{(E_1 - E_0)\tau} \pm c_1 \phi_1 . \quad (181)$$

This form reveals where the problem lies. The difference $(\phi^+ - \phi^-)$ depends only on the first excited state, so the expectation value of the local energy will be stabilising to E_1 . The variance, however, depends on $(\Psi H)^2$, so the orthogonality $\langle \varphi_T | \phi_0 \rangle = 0$ does not eliminate the ground state contribution from the variance. And that contribution is increasing exponentially, as (181) shows. The strategy is then to achieve a quick decay to the first excited state before the ground-state “noise” overwhelms the signal. This is happening only if our trial state is very close to the true first excited state.

Using decay curves In principle, we could extract excited-state information from the decay curves. By fitting

$$E(\tau) = \frac{\sum c_i a_i e^{-(E_i - E_T)\tau} E_i}{\sum c_i a_i e^{-(E_i - E_T)\tau}} \quad (182)$$

to the DMC energy we could get E_i . However, this has turned out to be impractical, because $E(\tau)$ depends on two exponentially decaying functions and the data is dominated by statistical noise. It would be more easy to fit, instead, the expectation value of the Green’s function, which has only one exponent function:

$$I(\tau) = \langle \varphi_T | e^{-(H - E_T)\tau} | \varphi_T \rangle . \quad (183)$$

This function can be sampled during the random walk by evaluating the **cumulative branching weight**

$$I(\tau) = \langle \varphi_T | \prod_{k=0}^N e^{-(E_L - E_T)\tau} | \varphi_T \rangle . \quad (184)$$

The time dependence is simply

$$I(\tau) = \sum a_i^2 e^{-(E_i - E_T)\tau} , \quad (185)$$

which is much easier to fit. But again, the data is governed by statistical noise³³. Fitting to exponentially decaying curves is an inverse Laplace transform - a perfectly stable task in the absence of statistical noise. If the data is noisy, the Laplace problem can be regularised by the use of a statistical **maximum entropy method**, in which prior knowledge is used to bias the inversion³⁴. One faces a similar situation in path integral Monte Carlo (PIMC) calculations of the dynamic structure function, where the numerical inaccuracy shows up as rounded peaks in $S(k, \omega)$ ³⁵

Using the zero variance property One way to search for the excited states owes to the fact that they are eigenstates of the Hamiltonian, which means that they have lower (zero) variance than other energies. In theory, one could take a sweep through the energies above the ground-state energy and look for minima in the variance. Unfortunately, statistical noise deems this method prohibitively expensive³⁶.

Matrix Monte Carlo method by Ceperley and Bernu³⁷ Take m states, which hopefully have large overlap with the first m excited states of the system. Then McDonald's theorem states, that if you simultaneously diagonalise both the overlap matrix and the Hamiltonian, then the n th eigenvalue is an upper limit to the exact energy of the n th state. One can then use the DMC states $f(\tau)$ to construct these matrices, and it can be shown that the eigenvalues approach the exact ones exponentially fast. Unfortunately, also the statistical noise increases exponentially fast. The method has been applied to small molecules and two-dimensional electron gas. There is a nice, short introduction to the method in the Valencia lecture notes by J. Boronat.

6.7 Second order DMC algorithms

A DMC algorithm is said to be first order, if the timestep error in the total energy is linear, $E(\tau) = E_0 + a\tau$. It would be desirable to find a second order algorithm, $E(\tau) = E_0 + b\tau^2$, so that larger timesteps could be used, possibly without any $\tau \rightarrow 0$ extrapolations. This section will follow closely the discussion given by Chin³⁸.

In section 6.1 we made a short-timestep approximation to the Green's

³³D. Ceperley and B. Bernu. J. Chem. Phys. , 89 (1988) 6316.

³⁴M. Caffarel and D. Ceperley, J. Chem. Phys. 97 (1992) 8415.

³⁵M. Boninsegni and D. Ceperley, J. Low Temp. Phys. 104 (1996) 339.

³⁶A. J. Williamson, PhD thesis, Cavendish Laboratory, September 1996.

³⁷D. M. Ceperley and B. J. Bernu, J. Chem. Phys. **89** (1988) 6316.

³⁸S. A. Chin, Phys. Rev. **A42** (1990) 6991.

function. To begin with, the Campbell-Baker-Hausdorff formula states that

$$e^A e^B = e^{A+B+\frac{1}{2}[A,B]+\frac{1}{12}[A-B,[A,B]]+\dots} \quad (186)$$

This shows that the exact Green's function³⁹

$$G = e^{-(H-E_T)\tau} = e^{-(T+V-E_T)\tau} \quad (187)$$

can be approximated as

$$G \approx e^{-(H'-E_T)\tau} , \quad (188)$$

where

$$H' = H - \frac{1}{2}\tau[H, V] + \frac{1}{12}\tau^2[H - 2V, [H, V]] + \dots \quad (189)$$

But this is just the exact Hamiltonian H plus a small perturbation, and one can formally compute the result of the DMC calculation. Keeping only the first two terms in (189) leads to the short-time approximation to the Green's function derived earlier in these lecture notes. The ground state of that H' is

$$|\phi'_0\rangle = |\phi_0\rangle + \frac{1}{2}\tau V|\phi_0\rangle + \dots \quad (190)$$

The first order term in (189) vanishes between eigenstates,

$$\langle\phi_i|[H, V]|\phi_i\rangle = 0 , \quad (191)$$

so although the eigenstates are first order in τ , the eigenvalues are quadratic in τ :

$$E'_0 = E_0 - \frac{1}{24}\tau^2\langle\phi_0|[V, [H, V]]|\phi_0\rangle + \dots \quad (192)$$

This means that *without* importance sampling the primitive approximation to the Green's function (as it will be called in the PIMC section 7) is already quadratic. Still, the energy (192) is not showing the τ -dependence of the energy that really comes out from an importance sampling DMC code. If this surprises you, remember that if one introduces a trial wave function one is not sampling $|\phi_0|^2$, but $\phi_0\varphi_T$. Hence the energy that is the average local energy over the distribution $\phi_0\varphi_T$ behaves as

$$\begin{aligned} E_{DMC} &\equiv \frac{\langle\phi_0|H|\varphi_T\rangle}{\langle\phi_0|\varphi_T\rangle} \\ &= E_0 + \frac{1}{2}\tau \frac{\langle\phi_0|[H, V]|\varphi_T\rangle}{\langle\phi_0|\varphi_T\rangle} + \dots , \end{aligned} \quad (193)$$

³⁹Chin thinks of the Hamiltonian as a matrix, hence G is a *transfer matrix*.

which is just first order.

Quadratic E_{DMC} can be rather easily found by looking at the Campbell-Baker-Hausdorff formula. The undesired first order term will cancel if one takes a product of operators at half the timestep and another set with reversed order. One such approximate Green's function would be

$$G \approx e^{-\frac{1}{2}\tau V} e^{-\frac{1}{2}\tau T} e^{-\frac{1}{2}\tau T} e^{-\frac{1}{2}\tau V} , \quad (194)$$

which approximates the Hamiltonian

$$H' = H + \frac{1}{24}\tau^2[(2H - V), [H, V]] + \dots , \quad (195)$$

and the corresponding importance sampled DMC energy is

$$E_{DMC} = E_0 + \frac{1}{24}\tau^2 \left[\langle \phi_0 | \Delta H | \phi_0 \rangle - \frac{\langle \phi_0 | \Delta H | \phi_0 \rangle}{\langle \phi_0 | \phi_0 \rangle} \right] , \quad (196)$$

where

$$\Delta H = [(2H - V), [H, V]] . \quad (197)$$

Now we have at hand an approximate Green's function that leads to a quadratic DMC energy. However, the implementation of the Langevin algorithm that is used to solve the underlying Fokker-Planck equation can still turn the code first order. I will not discuss the subject at length, for more details see Chin's 1990 article. I just remind you what was noted below Eq. (148), that the way one combines the *deterministic* drift force to the *stochastic* Gaussian move is not unique. It is easy to imagine that a non-symmetrical way of doing this will give rise to a first order DMC algorithm. The recipe is similar as in the Green's function: take intermediate steps, very much like in solving differential equations using Runge-Kutta and the kind.

At first glance it would seem that developing an even higher order DMC algorithm is an easy task. But Suzuki showed 1991⁴⁰, that in the expansion

$$e^{(T+V)\tau} = \prod_{i=1}^N e^{a_i T \tau} e^{b_i V \tau} \quad (198)$$

one can go only up to second order without having some of the coefficients a_i and b_i negative. But a negative a_i would mean diffusion backwards in time, which is impossible to simulate!

⁴⁰M. Suzuki, J. Math. Phys. 32, 400 (1991).

Suzuki’s “no-go theorem” means only that higher than second order algorithms must be based on a different kind of operator splitting than the one in Eq. (198).

Expansion of the Green’s function is not just a Monte Carlo task. Recently Auer, Krotscheck and Chin published an application of fourth order algorithm (error $\mathcal{O}(\tau^4)$) to solve non-trivial, three-dimensional Schrödinger equations⁴¹

A second order algorithm The following five stage algorithm will give a second order energy convergence. In the following $\mathbf{r}(\mathbf{k}, \mathbf{i})$ is the k th coordinate of i th walker in configuration $\mathbf{j}$ (not marked):

1. Take a half-size timestep with a Gaussian of variance $D\tau$:

$$\mathbf{r}(\mathbf{k}, \mathbf{i}) = \mathbf{r}(\mathbf{k}, \mathbf{i}) + \text{sqrt}(D*\tau)*\text{gauss}()$$

2. Compute the drift force using the new coordinates, $\mathbf{F}(\mathbf{R})$.
3. Move walkers according to that drift for half a timestep; keep the old configuration $\mathbf{R}$ in memory for later use:

$$\mathbf{rp}(\mathbf{k}, \mathbf{i}) = \mathbf{r}(\mathbf{k}, \mathbf{i}) + 0.5d0*D*\tau*\mathbf{F}(\mathbf{k}, \mathbf{i})$$

This gives you the configuration $\mathbf{R}'$.

4. Compute the drift force $\mathbf{F}(\mathbf{R}')$; I call it $\mathbf{FRP}$.
5. In configuration $\mathbf{R}$ (not in $\mathbf{R}'$!), take a half-size Gaussian step and move walkers full step with drift force $\mathbf{F}(\mathbf{R}')$:

$$\mathbf{rp}(\mathbf{k}, \mathbf{i}) = \mathbf{r}(\mathbf{k}, \mathbf{i}) + D*\tau*\mathbf{FRP}(\mathbf{k}, \mathbf{i}) + \text{sqrt}(D*\tau)*\text{gauss}()$$

This gives you a new configuration $\mathbf{R}'$, which not the same as in step 3! That old $\mathbf{R}'$ is not needed any more, and I make use of the memory space.

This algorithm is Eq. (25) in Chin’s article, where it was given in the form

$$y_i = x_i + \sqrt{\Delta t/2} \xi_i \quad (199)$$

$$x'_i = y_i + \Delta t G_i((\mathbf{y} + \frac{1}{2}\Delta t \mathbf{G}(\mathbf{y})) + \sqrt{\Delta t/2} \xi_i), \quad (200)$$

⁴¹J. Auer, E. Krotscheck, and S. A. Chin, J. Chem. Phys. 115 (2001) 6841.

where ξ is Gaussian random number with unit variance (as our η before). In the algorithm x_i and y_i are both $\mathbf{R}$, because x_i is not needed later and it's unnecessary to keep it in store. Also G is *half* our drift force and $D = 1/2$. After these changes one recovers the algorithm given above.

After these steps you have a new trial configuration $\mathbf{R}'$ that you can accept or discard (use the modified Metropolis question with the Green's functions). These steps are done for each coordinate, each walker and in the outmost loop, for each configuration in the present generation. Only after all the steps are done you have the next generation ready.

6.8 DMC Example: the ground state of a He atom

The He atom can be used to show how DMC works in practise. The He ion is in the origin (with infinite mass in the present approximation) and the coordinates of the two electrons are $\mathbf{r}_1$ and $\mathbf{r}_2$. Their spins are assumed to be antiparallel. If ∇_i operates on the i th particle, the Schrödinger equation for bound states of helium in atomic units reads⁴²

$$\left[-\frac{1}{2}\nabla_1^2 - \frac{1}{2}\nabla_2^2 - \frac{2}{r_1} - \frac{2}{r_2} + \frac{1}{r_{12}} \right] \varphi(\mathbf{r}_1, \mathbf{r}_2) = E\varphi(\mathbf{r}_1, \mathbf{r}_2) , \quad (201)$$

where $r_{12} = |\mathbf{r}_1 - \mathbf{r}_2|$ is the distance between the electrons.

Next you need to write down the trial wave function and find the expression for the local energy. A simple choice would be the product of two 1S hydrogen orbitals,

$$\varphi_T(\mathbf{r}_1, \mathbf{r}_2) = e^{-\alpha r_1} e^{-\alpha r_2} , \quad (202)$$

which gives the local energy

$$\begin{aligned} E_L &= \varphi_T(\mathbf{r}_1, \mathbf{r}_2)^{-1} H \varphi_T(\mathbf{r}_1, \mathbf{r}_2) \\ &= e^{\alpha r_1} e^{\alpha r_2} \left[-\frac{1}{2}\nabla_1^2 - \frac{1}{2}\nabla_2^2 - \frac{2}{r_1} - \frac{2}{r_2} + \frac{1}{r_{12}} \right] (e^{-\alpha r_1} e^{-\alpha r_2}) \\ &= \underbrace{-\alpha^2 + \frac{\alpha}{r_1} + \frac{\alpha}{r_2}}_{\text{kinetic part}} \underbrace{-\frac{2}{r_1} - \frac{2}{r_2} + \frac{1}{r_{12}}}_{\text{potential part}} . \end{aligned} \quad (203)$$

Notice that the “kinetic part” is *not* the kinetic energy of the trial state: The “kinetic part” can be negative for some choices of r_1 and r_2 ! It has only a statistical meaning, the averages give the kinetic and the potential energy, but one has to remember that since $[H, T] \neq 0$ and $[H, V] \neq 0$, one gets by

⁴²Atomic units: energy in Hartrees (1 Ha = 2 Ry), distance in Bohr radii.

straightforward averaging only their mixed estimates. The “kinetic part” is an approximate energy contribution brought into the local energy by the use of a trial wave function, and its sole purpose is to try to cancel potential energy fluctuations when r_1 , r_2 and r_{12} are changing. In real world large negative potential energy is compensated by large positive kinetic energy.

Variance optimisation is tricky, because this 1S-orbital trial state is so far from the true ground state of He. For example, this trial wave function won’t keep the electrons apart. Minimum energy will be reached with $\alpha = 1.6875$, but the variance is at minimum at a slightly higher value. This value depends on the (fixed) set of configurations you happen to be using for optimisation: If the set has many electrons near the origin, the variance optimisation will try to make the Coulomb terms cancel as well as it can and at the same time try to minimise the energy, hence α would be between 1.6875 and 2.0. On the other hand, if your configurations happen to be such that the potential energies are moderate, then the Coulomb cancellation will seem less acute and the optimal value of α would be closer to 1.6875.

I applied the reweighting algorithm with the wheel-of-fortune type of weight renormalisation. All attempted moves were accepted, so the algorithm did not rigorously satisfy the detailed balance condition. As is evident from Fig. 11, characteristic to the used algorithm is that the **timestep error** is linear for small τ , but for larger time steps this linearity breaks down. Not shown in the figure is that a long computation with $\alpha = 1.6875$ has large jumps in the variance due to electrons coming too close to the nucleus. Although this parameter gives the lowest trial-state energy, it is not the most stable when it comes to energy variance. Fig. 11 also shows the results for $\alpha = 2$, with some very large timesteps. The energy fluctuations are now in better control, as Fig. 12 shows. The negative spikes are characteristic to the $\alpha = 1.6875$ -calculation, and even larger ones occur now and then causing sudden jumps to the energy variance. Smaller positive spikes are seen in the $\alpha = 2.0$ data; they originate from the $1/r_{12}$ electron repulsion, but they are less catastrophic to the variance.

The kinetic and potential energies of the trial state with $\alpha = 2.0$ are $\langle \varphi_T | T | \varphi_T \rangle = 4.0007 \pm 0.0009$ and $\langle \varphi_T | V | \varphi_T \rangle = -6.7507 \pm 0.0008$, respectively. The DMC calculation gave, after extrapolation to zero timestep, $\langle \phi_0 | T | \varphi_T \rangle = 3.375$ and $\langle \phi_0 | V | \varphi_T \rangle = -6.280$ (no error analysis was made). Hence the extrapolated estimators are $E_{kin} = 2.750$ and $E_{pot} = -5.809$ (unknown error). The total energy computed from these is -3.059 with much larger error bars than in the total energy computed directly (-2.90465 in the present simulation vs. the exact value -2.90474). Deviations between the trial state energies and the DMC values give some measure to how close the trial state is to the true ground state: quite far in the present case. For




Figure 11: The dependence of the calculated total energy of a He atom of the time step in DMC. The error bars are obtained via **block averaging**. The upper panel show the results for $\alpha = 1.6875$, the lower for $\alpha = 2.0$. The extrapolated zero time step result from the upper panel data is -2.90374 Ha (five correct desimals).

the record, the total energy of the trial state was $E_{\text{trial}} = -2.7500 \pm 0.0002$ from the same simulation. As always, the total energy comes out more accurately than the kinetic or the potential energy. Part of the potential energy fluctuations has been cancelled by the kinetic energy fluctuations.

Fig. 13 shows the DMC energy as a function of the time step obtained using the quadratic DMC algorithm given earlier. The error bars show the statistical error estimated using block averaging. The computation was continued until the σ of all data was 0.0005, therefore the small time-step results have larger error bars. The least square fit shown in the figure is $f(\tau) = -2.90286 - 2.1263 \tau^2$, the horizontal line is the exact energy -2.90374. The shown results are obtained after a set of short runs, and the error bars




Figure 12: Plots of the local energy averages taken over 2000 configurations during over a set of generations for $\alpha = 1.6875$ (upper panel) and $\alpha = 2.0$ (lower panel). The time step was 0.001. The plot shows only a tiny fraction of the data collected for the results given in Fig. 11.

are too large for serious use. However, the overall quadratic behaviour of the time-step error is evident.




Figure 13: Time-step error in the quadratic DMC code, with a trial state that has 2S orbital states.

7 Path Integral Monte Carlo

Path Integral Monte Carlo (PIMC) is a Monte Carlo application of Feynman's path integral formulation of quantum mechanics. It is very powerful and leads to in principle exact results at *finite temperature*. I shall be talking about bosons in equilibrium, but some remarks on fermion will be made in the end. The main reference to ideas presented in this section are from the review article by David Ceperley⁴³

Our task is to sample the coordinate space according to $|\Psi(\mathbf{R})|^2$ so that the energy eigenstates are visited according to the Maxwell-Boltzmann probability $e^{-\beta E_N}$. The basic quantity at finite T is the **thermal density matrix**, which in operator form reads

$$\hat{\rho}(T) \equiv e^{-\beta \hat{H}}. \quad (204)$$

In coordinate space the density matrix reads (I shall drop “thermal” for brevity) $|\mathbf{R}\rangle = \sum_n \phi_n(\mathbf{R})|\phi_n\rangle$;

$$\rho(\mathbf{R}, \mathbf{R}'; T) \equiv \langle \mathbf{R} | e^{-\beta H} | \mathbf{R}' \rangle = \sum_n \phi_n(\mathbf{R})^* \phi_n(\mathbf{R}') e^{-\beta E_n}, \quad (205)$$

⁴³D. Ceperley, *Path integrals in the theory of condensed helium*, Rev. Mod. Phys. **67** (1995) 279.

where the sum is over all states n and ϕ_n are the eigenstates of the Hamiltonian with energy E_n . Since these eigenstates are unknown for N -body systems, we must find a more practical way to compute the density matrix.

All static properties of the system can be written in terms of the density matrix. The expectation value of an arbitrary operator A at temperature T is

$$\langle A \rangle \equiv Z^{-1} \sum_n \langle \phi_n | A | \phi_n \rangle e^{-\beta E_n} \quad (206)$$

$$= Z^{-1} \int d\mathbf{R} d\mathbf{R}' \rho(\mathbf{R}, \mathbf{R}'; T) \langle \mathbf{R} | A | \mathbf{R}' \rangle, \quad (207)$$

where

$$Z = \sum_n e^{-\beta E_n} = \sum_n \langle n | e^{-\beta H} | n \rangle \quad (208)$$

$$\begin{aligned} &= \int d\mathbf{R} \sum_n \langle n | \mathbf{R} \rangle \langle \mathbf{R} | e^{-\beta H} | n \rangle \\ &= \int d\mathbf{R} \sum_n \langle \mathbf{R} | e^{-\beta H} | n \rangle \langle n | \mathbf{R} \rangle \\ &= \int d\mathbf{R} \langle \mathbf{R} | e^{-\beta H} | \mathbf{R} \rangle = \int d\mathbf{R} \rho(\mathbf{R}, \mathbf{R}; T) \end{aligned} \quad (209)$$

is the partition function. Notice, that Z depends on the *diagonal* part of the density matrix.

7.1 First ingredient: products of density matrices

From the PIMC point of view the single most important property of the density matrix is that the product of two density matrices is a density matrix; in operator form

$$\hat{\rho}(\beta_1) \rho(\beta_2) = e^{-\beta_1 \hat{H}} e^{-\beta_2 \hat{H}} = e^{-(\beta_1 + \beta_2) \hat{H}} = \hat{\rho}(\beta_1 + \beta_2), \quad (210)$$

or with explicit temperatures

$$\hat{\rho}(T_1) \hat{\rho}(T_2) = \hat{\rho}\left(\frac{T_1 T_2}{T_1 + T_2}\right). \quad (211)$$

If we take, for example, M temperatures $T_1 = T_2 = \dots = T_M = T$, then (210) states that

$$\hat{\rho}(T)\hat{\rho}(T)\dots\hat{\rho}(T) = \hat{\rho}(T)^M = \hat{\rho}\left(\frac{T}{M}\right), \quad (212)$$

so the product property makes it possible to express a low temperature density matrix (rhs) as a product of higher temperature density matrices. In coordinate space (210) is a convolution integral

$$\rho(\mathbf{R}_1, \mathbf{R}_3, \beta_1 + \beta_2) = \int d\mathbf{R}_2 \rho(\mathbf{R}_1, \mathbf{R}_2, \beta_1) \rho(\mathbf{R}_2, \mathbf{R}_3, \beta_2). \quad (213)$$

Let us define the **time step** τ ,

$$\tau \equiv \frac{\beta}{M}. \quad (214)$$

We call this “time”, because the time evolution of a state is given by the operator $e^{-iHt/\hbar}$, while we are interested in $e^{-\beta H}$. Hence we are actually doing **imaginary time** quantum mechanics like we did in DMC, except that now

$$t \sim -i\hbar\beta. \quad (215)$$

This gives the idea to call β/M a time step ⁴⁴.

The density matrix is now divided into a M -segment, discrete **path**:

$$\begin{aligned} \rho(\mathbf{R}_0, \mathbf{R}_M; \beta) &= \int d\mathbf{R}_1 d\mathbf{R}_2 \dots d\mathbf{R}_{M-1} \rho(\mathbf{R}_0, \mathbf{R}_1; \tau) \rho(\mathbf{R}_1, \mathbf{R}_2; \tau) \\ &\dots \rho(\mathbf{R}_{M-2}, \mathbf{R}_{M-1}; \tau) \rho(\mathbf{R}_{M-1}, \mathbf{R}_M; \tau). \end{aligned} \quad (216)$$

Thus the density matrix requires a computation of a multidimensional convolution integral; exactly what Monte Carlo integration is good at.

Decomposing the partition function (208) into M path segments gives

$$Z = \int d\mathbf{R} \langle \mathbf{R} | e^{-\beta H} | \mathbf{R} \rangle \quad (217)$$

$$\begin{aligned} &= \int d\mathbf{R} d\mathbf{R}_1 \dots d\mathbf{R}_M \langle \mathbf{R} | e^{-\beta\tau} | \mathbf{R}_1 \rangle \langle \mathbf{R}_1 | e^{-\beta\tau} | \mathbf{R}_2 \rangle \\ &\dots \langle \mathbf{R}_M - 1 | e^{-\beta\tau} | \mathbf{R}_M \rangle \langle \mathbf{R}_M | e^{-\beta\tau} | \mathbf{R} \rangle. \end{aligned} \quad (218)$$

So the partition function is a *path integral over all closed loops*. Often $\hat{A}$ is diagonal and the thermal average is computed as a trace

$$\langle A \rangle = Z^{-1} \int d\mathbf{R} d\mathbf{R}_1 \dots d\mathbf{R}_M \rho(\mathbf{R}, \mathbf{R}_1; \tau) \rho(\mathbf{R}_1, \mathbf{R}_2; \tau) \dots \rho(\mathbf{R}_M, \mathbf{R}; \tau) \langle \mathbf{R} | \hat{A} | \mathbf{R} \rangle$$

⁴⁴Later we will encounter another example where temperature and time step are related: the Nosé–Hoover thermostat, Eq. (316).

There is a clear analogy to classical polymers, each particle being one polymer. The above trace corresponds to a **ring polymer**. The lower the temperature, the longer the polymer.

7.2 Second ingredient: high- T density matrix

We wouldn't be any much wiser if we couldn't evaluate the individual segments of the paths. It turns out, however, that the density matrix at high temperature can be found, at least approximately. With $\hat{H} = \hat{T} + \hat{V}$ we have

$$\begin{aligned} e^{-\tau\hat{H}} &= e^{-\tau(\hat{T}+\hat{V})} = e^{-\tau\hat{T}} e^{-\tau\hat{V}} e^{\frac{\tau^2}{2}[\hat{V},\hat{T}]} \\ &\approx e^{-\tau\hat{T}} e^{-\tau\hat{V}} \text{ for small } \tau . \end{aligned} \quad (220)$$

The last line is called the **primitive approximation**. In coordinate space it is

$$\rho(\mathbf{R}_0, \mathbf{R}_2; \tau) \approx \int d\mathbf{R}_1 \langle \mathbf{R}_0 | e^{-\tau\hat{T}} | \mathbf{R}_1 \rangle \langle \mathbf{R}_1 | e^{-\tau\hat{T}} | \mathbf{R}_2 \rangle . \quad (221)$$

Now you should hear a bell ringing: *but we've done this already!* Indeed, in section 6.1 we evaluated exactly the matrix elements of the kinetic and potential operators and called them Green's functions. I just copy the results here:

$$\langle \mathbf{R}_1 | e^{-\tau\hat{T}} | \mathbf{R}_2 \rangle = (4\pi\lambda\tau)^{-3N/2} e^{-\frac{(\mathbf{R}_1 - \mathbf{R}_2)^2}{4\lambda\tau}} , \quad (222)$$

and

$$\langle \mathbf{R}_1 | e^{-\tau\hat{V}} | \mathbf{R}_2 \rangle = e^{-\tau V(\mathbf{R}_1)} \delta(\mathbf{R}_1 - \mathbf{R}_2) . \quad (223)$$

Here λ is the **thermal (de Broglie) wavelength** – it replaces the diffusion constant D :

$$\lambda = \sqrt{\frac{2\pi\hbar^2}{mk_B T}} . \quad (224)$$

If you take a look at the derivations of these formulae, you notice that something could go wrong because of the finite size of the simulations system. The k -space is not actually continuous, and in (130) we approximated the k -space sums with integrals. This is accurate only if

$$\lambda\tau \ll L^2 \quad (225)$$

in a simulation box with sides L . If this condition is not met, the Gaussian (222) will reach across the simulation box, leading to spurious effects.

Using the primitive approximation Eq. (216) can be put into following form:

$$\rho(\mathbf{R}_0, \mathbf{R}_M; \beta) \approx \int d\mathbf{R}_1 d\mathbf{R}_2 \dots d\mathbf{R}_{M-1} (4\pi\lambda\tau)^{-3NM/2} \times \exp \left[- \sum_{m=1}^M \left(\frac{(\mathbf{R}_{m-1} - \mathbf{R}_m)^2}{4\lambda\tau} + \tau V(\mathbf{R}_m) \right) \right] . \quad (226)$$

7.3 Connection to path integral formulation of quantum mechanics

From Eq. (226) we can easily get to Feynman's path integral formulation of quantum mechanics. First, shift from the thermal Green's function to a dynamic Green's function used in DMC by replacing $\lambda \rightarrow D = \hbar^2/(2m)$. Next, shift back to real time and take the limit $M \rightarrow \infty$ in (226). Then

$$\begin{aligned} \frac{(\mathbf{R}_{m-1} - \mathbf{R}_m)^2}{4\lambda\tau} &\sim \frac{(\mathbf{R}_{m-1} - \mathbf{R}_m)^2}{4 \frac{\hbar^2}{2m} \frac{i\Delta t}{\hbar M}} \\ &\rightarrow -\frac{i}{\hbar} \frac{1}{2} m \left(\frac{d\mathbf{R}}{dt} \right)^2 dt, \text{ as } M \rightarrow \infty, \end{aligned} \quad (227)$$

so the exponent term in (226) corresponds to

$$\exp \left\{ \frac{i}{\hbar} \int dt \left[\frac{1}{2} m \left(\frac{d\mathbf{R}}{dt} \right)^2 - V(\mathbf{R}) \right] \right\} = \exp \left\{ \frac{i}{\hbar} S(\mathbf{R}, t) \right\}, \quad (228)$$

where $S = \int dt L$ is the **classical action**, time integral of the Lagrangian L . We recover the famous Feynman-Kac formula of quantum mechanics,

$$\langle x', t' | x, t \rangle = \int \mathcal{D} e^{\frac{i}{\hbar} S[x(t)]}, \quad (229)$$

which states that the probability amplitude of a point particle to get from point (x, t) to point (x', t') is obtained by integrating over all possible paths between the points, attaching to each path a weight $e^{\frac{i}{\hbar} S}$. The measure $\mathcal{D}$ is just a commonly used shorthand for "all paths and proper normalisation". The normalisation, by the way, is in many cases the tricky part!




Figure 14: Two possible paths connecting fixed end points. The time scale is divided into discrete time slices. The quantum amplitude $\langle x', t' | x, t \rangle$ is calculated by adding all possible paths with *classical* weight $\exp \{iS[x(t)]/\hbar\}$.

Fig. 14 shows two possible paths from point (x, t) to point (x', t') . Time is discretised to **time slices** of finite width, like in a simulation.

Eq. (229) is not only the most straightforward connection between classical and quantum mechanics, it is also an excellent starting point of the whole quantum mechanics of point-like particles (quantum field theory gives the rest). The classical limit (Ehrenfest's theorem) is easy to derive. For $\hbar \rightarrow 0$ the integrand oscillates very rapidly and only paths with stationary S will contribute: we get $\delta S = 0$, the classical least action principle. Schrödinger equation follows, too, by considering a case where one makes a small variation to the path keeping the starting point (x, t) fixed. The path-integral formulation of quantum mechanics has been reviewed by its inventor, Richard Feynman⁴⁵

Visualisation of paths Paths can be visualised by drawing world-line plots that resemble Minkowski diagrams used in special relativity (see Ceperley's review for examples). One must remember, however, that the imaginary time axis is now a thermal quantity and not related to the true dynamics of the particles. Coordinate space plots that cover some fixed amount of time steps are also useful. There one must remember, that the discretisation to time slices is entirely artificial, and a plot obtained by simply connecting all the

⁴⁵R. P. Feynman, *Space-time approach to non-relativistic quantum mechanics*, Rev. Mod. Phys. **20** (1948) 367.

beads (points visited at discrete times) does not give full credit to the real situation. A more realistic picture can be obtained by taking into account, that a particle may visit all points within the thermal wavelength of any calculated bead.

7.4 Link action

A **link** is the pair of adjacent time slices, $\{\mathbf{R}_{m-1}, \mathbf{R}_m\}$, in Eq. (226). The **link action** is minus the logarithm of the *exact* density matrix:

$$S_m \equiv -\ln [\rho(\mathbf{R}_{m-1}, \mathbf{R}_m; \tau)] . \quad (230)$$

Thus the exact density matrix reads

$$\rho(\mathbf{R}_0, \mathbf{R}_M; \beta) = \int d\mathbf{R}_1 \dots \mathbf{R}_{M-1} \exp \left[- \sum_{m=1, M} S_m \right] . \quad (231)$$

It is customary to separate the actions coming from the kinetic and potential parts of the Hamiltonian, which are

$$K_m = \frac{3M}{2} \ln(4\pi\lambda\tau) + \frac{(\mathbf{R}_{m-1} - \mathbf{R}_m)^2}{4\lambda\tau} \quad (232)$$

and the “inter-action” $U_m = S_m - K_m$, which is often used in the symmetrised form. Put together, the primitive approximation for the link action is

$$S_m^{prim.} = \frac{3M}{2} \ln(4\pi\lambda\tau) + \underbrace{\frac{(\mathbf{R}_{m-1} - \mathbf{R}_m)^2}{4\lambda\tau}}_{\text{“Spring term”}} + \underbrace{\frac{\tau}{2} [V(\mathbf{R}_{m-1}) + V(\mathbf{R}_m)]}_{\text{“Potential term”}} . \quad (233)$$

In polymer analogy the spring term holds the polymer together and the potential term gives the inter-polymer interaction. For many quantities thermal averages translate to ring polymers.

7.5 Improving the action

The superfluid temperature is only 2.172 K for ^4He liquid, and the number of time slices to reach the superfluid region turns out to be $M \sim 1000$ if one uses the primitive approximation. A better action is needed, which allows to take longer time steps. The path is a **fractal**, so finding better actions is not the same thing as finding a better integrator of a differential equation (*e.g.* leap-frog in MD). The strategy is to use exact properties of the action as a guidance.

The Feynman–Kac formula gives the inter-action as the average over all free-particle paths,

$$e^{-U_m} \equiv e^{-U(\mathbf{R}_{m-1}, \mathbf{R}_m; \tau)} = \left\langle \exp \left[- \int_0^\tau dt V(\mathbf{R}(t)) \right] \right\rangle_{RW} , \quad (234)$$

where $\langle \dots \rangle_{RW}$ denotes average over all random walk paths from $\mathbf{R}_{m-1}$ to $\mathbf{R}_m$. Often we have (or know) only the pair potential between particles, so we have

$$e^{-U(\mathbf{R}_{m-1}, \mathbf{R}_m; \tau)} = \left\langle \prod_{i < j} \exp \left[- \int_0^\tau dt V(r_{ij}(t)) \right] \right\rangle_{RW} , \quad (235)$$

where $r_{ij}(t)$ is the distance between particles i and j . If one neglects three-body and higher correlations between particles, the product can be taken out of the average,

$$e^{-U(\mathbf{R}_{m-1}, \mathbf{R}_m; \tau)} \approx \prod_{i < j} \left\langle \exp \left[- \int_0^\tau dt V(r_{ij}(t)) \right] \right\rangle_{RW} . \quad (236)$$

Now the average is the interaction of a pair of particles - a quantity that can be computed beforehand. This is the core idea of the **pair-product action**, and the improvement upon the primitive action is significant. Furthermore, it applies also to Coulomb and hard-core potentials, while many other improved action methods fail for nonintegrable potentials. A short list of possible improved actions:

- Pair-product action
- Cumulant action
- Harmonic action
- WKB action (semiclassical)

7.6 Bose Symmetry

In the beginning of the PIMC section I stated without hesitation that the energy eigenstates are visited according to the Maxwell–Boltzmann probability $e^{-\beta E_N}$. This is correct only for distinguishable particles. Bose-Einstein statistics requires that we symmetrise the density matrix (subindex B refers to bosons),

$$\rho_B(\mathbf{R}_0, \mathbf{R}_1; \beta) = \frac{1}{N!} \sum_P \rho(\mathbf{R}_0, P\mathbf{R}_1; \beta) , \quad (237)$$

where the sum is over all particle-label permutations P . The permutation of coordinates can be done on either the first or the second argument of the density matrix, or to both; it makes no difference. This involves the summation of $N!$ terms and is a computational burden. Fortunately all terms in the sum are positive and we can evaluate it by *sampling*: the bosonic simulation is a random walk through both the coordinate space *and* the permutation space. This cannot be done for fermions, because even and odd permutations have large cancellation and a MC evaluation of the sum is out of the question. I'll return to this problem later.

The partition function is a sum over all closed paths; Bose-symmetrised partition function is obtained by letting the paths close to any permutation of the starting points.

Superfluidity Partition function is obtained by taking all closed paths. Adding the Bose symmetrisation, this gives rise to exchange paths, where, say, particle 1 takes the place of particle 2 and vice versa. A special situation arises, when the exchange involves the periodic image of the other particle, because then the extended plot (all periodic boxes side by side) shows that the looping path goes through the entire system. The finite size of the simulation box causes finite size effects, but the idea of **winding paths**, paths that span across the entire system, remains valid. Thus the superfluid transition can be seen visually in PIMC. The best theoretical value of the lambda transition ⁴⁶ in ⁴He is 2.19 ± 0.02 K at SVP (saturated vapor pressure), very close to the experimental value 2.172 K. Via his path integral formulation of quantum mechanics, Feynman⁴⁷ has answered two important questions:

- Does the strong inter-particle potential invalidate the idea of Bose-Einstein condensation in helium? (No, it is still ok).
- Why should there be a transition at low temperature? (Feynman divides exchange paths into polygons and shows that at some low temperature large polygons that involve of the order of N particles start to contribute in the partition function).

⁴⁶E. L Pollock and K. J. Runge, Phys. **B46** (1992) 3535.

⁴⁷R. P. Feynman, Phys. Rev. **90** (1953) 1116; Phys. Rev. **91** (1953) 1291; Phys. Rev. **91** (1953) 1301.

7.7 Sampling paths

As Eq. (231) shows, we want to sample paths from the distribution

$$w(\mathbf{R}_0, \dots, \mathbf{R}_{M-1}) = \frac{1}{Z} \exp \left[- \sum_{m=1, M} S_m \right], \quad (238)$$

where the normalisation is again the partition function.

Consider first how to sample a single point in the path, the elementary operation in PIMC. We wish to move a bead $\mathbf{R}$ at time τ , which is connected to two fixed beads $\mathbf{R}_0$ and $\mathbf{R}_2$, at times 0 and 2τ , respectively. The simplest choice would be the **classic sampling**, where the bead is displaced uniformly inside a cube, whose size is adjusted to get near 50 % acceptance. In the **free-particle sampling** the new position is sampled from

$$w(\mathbf{R}) = \exp \left[- \frac{(\mathbf{R} - \mathbf{R}_{mid})^2}{2\lambda\tau} - U(\mathbf{R}, \mathbf{R}_1) - U(\mathbf{R}, \mathbf{R}_2) \right], \quad (239)$$

where $\mathbf{R}_{mid} = \frac{1}{2}(\mathbf{R}_1 + \mathbf{R}_2)$, so this is a Gaussian centered at $\mathbf{R}_{mid}$ with width $\sqrt{\lambda\tau}$. If there is a repulsive potential between the particles, the moves suggested by the free-particle sampling will be rejected whenever the beads become too close; Further improvement would be for example a correlated Gaussian.

Unfortunately the computer time needed for a free-particle path to diffuse a fixed distance will scale as M^3 . Hence large values of M will be needed and this fact combined with a non-efficient action (primitive action) can easily ruin the whole PIMC routine.

Moving the whole path at once is likely to be inefficient, as most of the attempted moves will be rejected. It is better to move path segments. One speaks of **multi-level Metropolis algorithms that uses bisection**. Their further description is beyond the scope of these lecture notes.

7.8 Fermion paths: The sign problem

For fermions indistinguishability is imposed by summing again over all permutations, but now terms which come from odd number of permutations have negative sign (compare with (237)):

$$\rho_F(\mathbf{R}_0, \mathbf{R}_1; \beta) = \frac{1}{N!} \sum_P (-1)^{\zeta_P} \rho(\mathbf{R}_0, P\mathbf{R}_1; \beta), \quad (240)$$

where ζ_P is the parity of permutation P . For bosons the factor is simply $(+1)^{\zeta_P} = 1$.

Now we can prove for both bosons and fermions the statement made earlier, namely that the permutation of coordinates can be done on either the first or the second argument of the density matrix, or to both. (Anti)symmetrising with respect to both arguments gives

$$\rho_{B/F}(\mathbf{R}_0, \mathbf{R}_1; \beta) \equiv \frac{1}{(N!)^2} \sum_P (\pm 1)^{\zeta_P} \sum_{P'} (\pm 1)^{\zeta_{P'}} \rho(P\mathbf{R}_0, P'\mathbf{R}_1; \beta) . \quad (241)$$

Taking a particular pair of permutations P and P' we can always find a third one such that $P'' = PP'$, and $\zeta_{P''} = \zeta_P + \zeta_{P'}$. Thus we can write the above equations in the form

$$\rho_{B/F}(\mathbf{R}_0, \mathbf{R}_1; \beta) = \frac{1}{(N!)^2} \sum_P (\pm 1)^{\zeta_P} \sum_{PP''} (\pm 1)^{\zeta_P + \zeta_{P''}} \rho(P\mathbf{R}_0, PP''\mathbf{R}_1; \beta) . \quad (242)$$

Sum over all permutations PP'' while P is fixed is the same as summing over all permutations P'' , so we get

$$\rho_{B/F}(\mathbf{R}_0, \mathbf{R}_1; \beta) = \frac{1}{(N!)^2} \sum_P \sum_{P''} (\pm 1)^{\zeta_{P''}} \rho(P\mathbf{R}_0, PP''\mathbf{R}_1; \beta) . \quad (243)$$

Now P is applied to both coordinates, so it merely relabels dummy variables. Hence all $N!$ terms give the same contribution and we can reduce this to

$$\rho_{B/F}(\mathbf{R}_0, \mathbf{R}_1; \beta) = \frac{1}{N!} \sum_{P''} (\pm 1)^{\zeta_{P''}} \rho(\mathbf{R}_0, P''\mathbf{R}_1; \beta) . \quad (244)$$

This completes our poof that permuting both arguments is equivalent with permuting one argument.

The expectation value of operator O in a Fermi system is

$$\langle O \rangle = \text{Tr}[\rho_F O] = \frac{1}{N!} \sum_P (-1)^{\zeta_P} \int d\mathbf{R} d\mathbf{R}_1 \dots d\mathbf{R}_{M-1} \quad (245)$$

$$\rho(\mathbf{R}, \mathbf{R}_{M-1}, \tau) \dots \rho(\mathbf{R}_1, \mathbf{R}', \tau) O(\mathbf{R}', \mathbf{R})|_{\mathbf{R}'=P\mathbf{R}} . \quad (246)$$

The Fermi antisymmetrisation is a sum of alternating series of large terms of slowly decreasing magnitude. We have no precise formula available for any of the terms, and already Feynman himself found impractical any attempts to analytically rearrange the terms in such way that they are positive. In Bose case one could rather easily sample the symmetrisation sum, but now one faces the problem of how to sample an alternating sum. Ceperley has given arguments, that the signal-to-noise ratio of such samplings vanishes

exponentially. In particular, the efficiency of a **direct-fermion method** (straightforward sampling as in the Bose case) scales as e^{-N} , where N is the number of fermions. Clearly the method is useless for all physically interesting systems. The Restricted Path Integral Monte Carlo (RPIMC) is one of the attempts to organise the terms in the sum is related to the fixed node DMC. Here one imposes an approximate nodal structure, with the aid of a trial density matrix ρ_T , to the density matrix ρ_F , and if this is correct, then the solution is exact. The theory behind RPIMC is beyond the scope of these lecture notes, I just mention that the free energy obtained using RPIMC satisfies the variational principle⁴⁸

7.9 Liquid helium excitation spectrum

Neutron scattering measures directly the dynamic structure function $S(k, \omega)$, which is related to the density-density correlation function $\mathcal{X}(k, \omega)$ via

$$S(k, \omega) = -\frac{1}{\pi} \Im m \mathcal{X}(k, \omega) , \quad (247)$$

where $\mathcal{X}(k, \omega)$ is the Fourier transform of the real space the density-density correlation function, $\mathcal{X}(\mathbf{r}_1, \mathbf{r}_2, t)$. The liquid is homogeneous and isotropic (same density everywhere and no direction dependences), therefore only one wave vector k appears. Explicitly, $S(k, \omega)$ is given by

$$S(k, \omega) = \frac{1}{2\pi N} \int_{-\infty}^{\infty} dt e^{i\omega t} \langle \hat{\rho}_k(t) \rho_{-k}(t) \rangle \quad (248)$$

$$= \frac{1}{NZ} \sum_{m,n} \delta(\omega - E_m - E_n) e^{-\beta E_n} |\langle m | \hat{\rho}_k | n \rangle|^2 , \quad (249)$$

where the density operator $\hat{\rho}_k$ ⁴⁹,

$$\hat{\rho}_k = \sum_{i=1}^N e^{i\mathbf{k} \cdot \mathbf{r}_i} \quad (250)$$

computed at the particle coordinates $\mathbf{r}_i$. The quantity that comes out of PIMC is written in imaginary time and it is the Laplace transform of the dynamic structure function,

$$F(k, t) = \int_{-\infty}^{\infty} d\omega e^{-\omega t} S(k, \omega) \quad (251)$$

$$= \frac{1}{NZ} \text{Tr} \{ \hat{\rho}_{-k} e^{-Ht} \hat{\rho}_k e^{-H(\beta-t)} \} . \quad (252)$$

⁴⁸The RPIMC free energy is equal or larger than the exact free energy.

⁴⁹Just a Fourier transform of a sum of Dirac delta functions.

Mathematically there is no harm done, because $F(k, t)$ and $S(k, \omega)$ contain exactly the same information and a Laplace inversion is well defined. The problem arises because the simulated $F(k, t)$ always has some noise in it, and perfect inversion is impossible. The time-domain structure function $F(k, t)$ is an exponentially decaying curve plus tiny features on top. And it is these small feature that give all $S(k, \omega)$, so very little noise can cover the useful part of the information. Even sharp features in $S(k, \omega)$, such as delta-function peaks, will show only as unimpressive exponentially decaying curves in $F(k, t)$.

8 Thermodynamics

8.1 NpT Simulations

If one wishes to study, for example, the liquid-solid phase transition or transitions from one solid phase to another, the best choice would be a NpT , not a NVT system. Such phase transitions are accompanied by a change in the volume, so in a NVT simulation one may have chosen the volume V which is characteristic to neither phase. In this case the system phase separates, and the simulation is no longer showing the phase transition we were interested in. After the simulated system has phase separated it is very difficult to distinguish which phase is giving which signal.

To simulate NpT systems one obviously has to find out how one should change the volume during the simulation. This was first done by McDonald in 1972⁵⁰, and an elegant derivation from statistical mechanics is due to Frenkel and Smit⁵¹

Let us first take a look at how the pressure is measured from a simulation, and see after that how to put up a NpT simulation.

8.1.1 Pressure

Pressure can be calculated from

$$p = k_B T \left(\frac{\partial \ln Z_N}{\partial V} \right)_{N,T}, \quad (253)$$

⁵⁰I. R. McDonald, Mol. Phys. **23** (1972) 41.

⁵¹D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press, New York, 1996).

The **canonical partition function** of classical 3D systems ($\beta = 1/(k_B T)$) is given by

$$Z_N = \frac{1}{\lambda^{3N} N!} \int_V d\mathbf{R} e^{-\beta V(\mathbf{R})} , \quad (254)$$

where the integrals run over volume V . Here $V(\mathbf{R})$ is the potential energy⁵² and λ is the deBroglie wavelength. The kinetic part factors out in classical systems, and won't affect the results. I shall use the symbol Z_N , although in the classical context the normalised version of integral (254) is often called **configuration integral** and marked Q_N .

The volume is now fluctuating (in a way to be specified later), so we *scale the coordinates* to get V out of the integration limits. In a cubic geometry $V = L^3$ and we scale

$$\begin{aligned} x &\rightarrow xL \\ y &\rightarrow yL \\ z &\rightarrow zL . \end{aligned} \quad (255)$$

Instead of $\mathbf{R}$, I shall use $\mathbf{S}$ for scaled coordinates. The scaled canonical partition function reads

$$Z_N \propto V^N \int d\mathbf{S} e^{-\beta V(\mathbf{S}, L)} , \quad (256)$$

where integration is over unit volume, viz.

$$\int d\mathbf{S} \equiv \int_0^1 dx_1 \int_0^1 dy_1 \int_0^1 dz_1 \dots \int_0^1 dx_N \int_0^1 dy_N \int_0^1 dz_N , \quad (257)$$

and the extra argument of the potential, L in $V(\mathbf{S}, L)$, indicates the scaling length of the coordinates.

The volume derivative is

$$\begin{aligned} \left(\frac{\partial Z_N}{\partial V} \right)_{N,T} &\propto NV^{N-1} \int d\mathbf{S} e^{-\beta V(\mathbf{S}, L)} \\ &- \beta V^N \int d\mathbf{S} e^{-\beta V(\mathbf{S}, L)} \left(\frac{\partial V(\mathbf{S}, L)}{\partial V} \right)_{N,T} . \end{aligned} \quad (258)$$

The volume derivative of the potential is

$$\left(\frac{\partial V(\mathbf{S}, L)}{\partial V} \right)_{N,T} = \sum_{i=1}^N \sum_{k=1}^3 \left(\frac{\partial V(\mathbf{S}, L)}{\partial (\mathbf{r}_i)_k} \right)_{N,T} \left(\frac{\partial (\mathbf{r}_i)_k}{\partial V} \right)_{N,T} \quad (259)$$

⁵²I shall use the same letter for potential and volume; volume has no arguments.

$$(\text{scale}) = \sum_{i=1}^N \sum_{k=1}^3 (\mathbf{F}_i)_k \left(\frac{\partial (V^{1/3} \mathbf{s}_i)_k}{\partial V} \right)_{N,T} \quad (260)$$

$$= \sum_{i=1}^N \sum_{k=1}^3 (\mathbf{F}_i)_k \frac{1}{3} V^{-2/3} (\mathbf{s}_i)_k \quad (261)$$

$$= \frac{1}{3} V^{-2/3} \sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{s}_i \quad (262)$$

$$(\text{descale}) = \frac{1}{3V} \sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i . \quad (263)$$

Here $\mathbf{F}_i$ is the force acting on i th particle. Finally,

$$\left(\frac{\partial Z_N}{\partial V} \right)_{N,T} \propto \frac{N}{V} \int d\mathbf{S} e^{-\beta V(\mathbf{S},L)} - \frac{\beta}{V} \int d\mathbf{S} e^{-\beta V(\mathbf{S},L)} \frac{1}{3} \sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i \quad (264)$$

$$= \frac{N}{V} Z_N - Z_N \frac{\beta}{V} \frac{1}{Z_N} \int d\mathbf{S} e^{-\beta V(\mathbf{S},L)} \frac{1}{3} \sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i \quad (265)$$

$$= \frac{N}{V} Z_N - Z_N \frac{\beta}{V} \left\langle \frac{1}{3} \sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i \right\rangle , \quad (266)$$

where the last form comes from the definition of a classical statistical average (canonical average). From this form we get

$$pV = Nk_B T - \frac{1}{3} \underbrace{\left\langle \sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i \right\rangle}_{\text{"virial"}} , \quad (267)$$

known as the **virial theorem**. In the case of a pair potential we can elaborate the virial term to get

$$\sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i = \sum_{i=1}^N \left(\sum_{j(\neq i)}^N \mathbf{F}_{ij} \right) \cdot \mathbf{r}_i = \sum_{i \neq j} \mathbf{F}_{ij} \cdot \mathbf{r}_i , \quad (268)$$

where $\mathbf{F}_{ij}$ is the force acting on the i th particle by the j th particle. Since $\mathbf{F}_{ij} = -\mathbf{F}_{ji}$, we have

$$\sum_{i=1}^N \mathbf{F}_i \cdot \mathbf{r}_i = \frac{1}{2} \sum_{i \neq j} (\mathbf{F}_{ij} \cdot \mathbf{r}_i + \mathbf{F}_{ji} \cdot \mathbf{r}_j) \quad (269)$$

$$= \sum_{i < j}^N (\mathbf{F}_{ij} \cdot \mathbf{r}_i - \mathbf{F}_{ij} \cdot \mathbf{r}_j) \quad (270)$$

$$= \sum_{i < j}^N \mathbf{F}_{ij} r_{ij} \quad (271)$$

$$= - \sum_{i < j}^N \frac{\partial V(r_{ij})}{\partial r_{ij}} r_{ij} , \quad (272)$$

so the virial theorem reads in this case

$$pV = Nk_B T + \frac{1}{3} \left\langle \sum_{i < j}^N \frac{\partial V(r_{ij})}{\partial r_{ij}} r_{ij} \right\rangle . \quad (273)$$

Evidently one can write the pressure in terms of the radial distribution function $g_2(r)$, namely (in 3D)

$$p = \rho k_B T - \frac{2}{3} \pi \rho^2 \int_0^\infty dr r^3 \frac{\partial V(r)}{\partial r} g_2(r) , \quad (274)$$

but using (267) or (273) pressure can be measured independently.

Without proof, I quote that with the aid of the $g_2(r)$ one can easily compute the internal energy and compressibility, which are given by

$$\langle U \rangle / N = 2\pi\rho \int_0^\infty dr r^2 V(r) g_2(r) , \quad (275)$$

and

$$\frac{\kappa}{\kappa_{\text{ideal gas}}} = 1 + 4\pi\rho \int_0^\infty dr r^2 [g_2(r) - 1] , \quad (276)$$

respectively.

8.1.2 Volume Fluctuations

Let us now consider a situation, where the system described so far is actually a tiny part of a much bigger *ideal gas* system. The surrounding system acts as a **heat bath**⁵³, see Fig. 15.

The partition function of the total system is the product of the partitions functions of the subsystem and the surrounding system. If the total system

⁵³Energy flows but no particles move between the systems.




Figure 15: A small, non-ideal system whose volume can change (hence the piston) is part of a *huge* ideal gas system.

has N_0 particles in volume V_0 , then ⁵⁴

$$\begin{aligned} Z(N, N_0, V, V_0, T) & \\ = \frac{V^N (V_0 - V)^{N_0 - N}}{N! (N_0 - N)!} \frac{1}{\lambda^{3N_0}} \underbrace{\int d\mathbf{S}'^{(N_0 - N)} d\mathbf{S}^N e^{-\beta V(\mathbf{S}, L)}}_{=1} . \end{aligned} \quad (277)$$

From this expression we may read the probability $P(V)$, that the subsystem has volume V :

$$P(V) = \frac{V^N (V_0 - V)^{N_0 - N} \int d\mathbf{S} e^{-\beta V(\mathbf{S}, L)}}{\int_0^{V_0} dV' V'^N (V_0 - V')^{N_0 - N} \int d\mathbf{S} e^{-\beta V(\mathbf{S}, L')}} . \quad (278)$$

Next we need to introduce the pressure p to this probability. Generally that's a tough one, but if we go to the limit $N_0 \rightarrow \infty$, $V_0 \rightarrow \infty$, so that the average density $\rho = (N_0 - N)/V_0$ is fixed, the total system is mostly ideal gas and small changes in the pressure of the subsystem has no effect on the pressure of the total system. Hence we may approximate

$$(V_0 - V)^{N_0 - N} = V_0^{N_0 - N} \left[1 - \frac{V}{V_0} \right]^{N_0 - N} \quad (279)$$

⁵⁴The ideal gas part has, of course, $e^{-\beta V_{\text{ideal}}(\mathbf{S}', L)} = e^0 = 1$.

$$\rightarrow V_0^{N_0-N} e^{-(N_0-N)V/V_0} \quad (280)$$

$$= V_0^{N_0-N} e^{-\rho V} \quad (281)$$

$$\approx V_0^{N_0-N} e^{-\beta p V}, \quad (282)$$

where the last line follows if we use the ideal gas law $p = \rho k_B T \equiv \rho/\beta$. Now we can integrate (277) over V , factor ρ is introduced to keep the outcome dimensionless:

$$\int dV \rho Z(N, N_0, V, V_0, T) \quad (283)$$

$$= \frac{V_0^{N_0-N}}{N!(N_0-N)!} \frac{\beta p}{\lambda^{3N_0}} \int dV V^N e^{-\beta p V} \int d\mathbf{S}^N e^{-\beta V(\mathbf{S}, L)} \quad (284)$$

$$= \underbrace{\frac{V_0^{N_0-N}}{(N_0-N)! \lambda^{3(N_0-N)}}}_{\text{heat bath}} \quad (285)$$

$$\times \underbrace{\frac{\beta p}{N! \lambda^{3N}} \int dV V^N e^{-\beta p V} \int d\mathbf{S}^N e^{-\beta V(\mathbf{S}, L)}}_{Z(N, p, T)}, \quad (286)$$

where we have identified the partition functions of the heat bath and the desired partition function of the NpT ensemble, $Z(N, p, T)$. The probability $P(V)$ is now

$$P(V) = \frac{V^N e^{-\beta p V} \int d\mathbf{S}^N e^{-\beta V(\mathbf{S}, L)}}{\int_0^{V_0} dV' V'^N e^{-\beta p V'} \int d\mathbf{S}^N e^{-\beta V(\mathbf{S}, L')}}. \quad (287)$$

This is the starting point of NpT Monte Carlo simulations. The denominator is just a normalisation factor, which, as we have learnt, is not actually needed in MC!⁵⁵ The probability to find the subsystem in a configuration $\mathbf{S}$ and volume V is

$$\begin{aligned} P(\mathbf{S}, V) &\propto V^N e^{-\beta p V} e^{-\beta V(\mathbf{S}, L)} \\ &= e^{-\beta[V(\mathbf{S}, L) + pV - Nk_B T \ln V]}. \end{aligned} \quad (288)$$

The quantity in brackets can be interpreted as a generalised Hamiltonian, with an extra variable V . Also, you might have noticed that the volume-dependent quantities in the brackets is actually the **enthalpy** of the system. Eq. (288) tries to tell us that we can simulate NpT systems by making *trial volume changes* similar to the changes in particle coordinates. Let

⁵⁵See, e.g., Eq. (76); only the *ratio* of probabilities is needed.

us assume that we try to change the volume $V \rightarrow V' = V + \Delta V$, where $\Delta V \in U[-\Delta V_{max}, \Delta V_{max}]$. This can be put into the Metropolis algorithm by accepting the volume change with probability

$$\begin{aligned} P(V \rightarrow V') &= \min \left(1, \frac{P(\mathbf{S}, V')}{P(\mathbf{S}, V)} \right) \\ &= \min \left(1, e^{-\beta[V(\mathbf{S}, L') - V(\mathbf{S}, L) + p(V' - V) - Nk_B T \ln(V'/V)]} \right) . \end{aligned} \quad (289)$$

This just the usual Metropolis question: Accept if the ratio is larger than unity or a random number $\in U[0, 1]$.

After a volume change has been made one has to recompute the potential energy. Except for few special cases where the effect of V on $V(\mathbf{S}, L)$ is easy to compute, this update is computationally expensive, and takes about the same time as N particle moves. Power-law potentials such as the Lennard-Jones potential are exceptional; one can write the n th power as

$$U_n = \sum_{i < j} \varepsilon \left(\frac{\sigma}{r_{ij}} \right)^n = L^{-n} \sum_{i < j} \varepsilon \left(\frac{\sigma}{s_{ij}} \right)^n , \quad (290)$$

which gives simply $U_n(L') = (L/L')^n U_n(L)$. And yet, this is true only for *untruncated* potential; if you truncated potential after a cutoff distance r_c the scaling is no longer trivial.

Usually one is forced to make volume changes less frequently than particle moves, but

Never do anything at non – random intervals.

Even though we cannot let the volume change happen at *every* step of the simulation, it is imperative to let it happen at *any* step of the simulation. If you want to try a volume change on average every n th step, then do it like this:

```

!
! basic MC loop
! -----
do istep = 1, Nstep
  i = int(rand()*n) + 1
  if(i < n) then
    call move_particles
  else
    call change_volume
  end if
end do

```

Here $\text{rand}() \in U[0, 1]$ and N is the number of particles. This way you won't introduce bias to your results that would reflect your choice of a fixed volume-change interval.

The working of a NpT Monte Carlo simulation should be always checked by comparing the externally set pressure against the pressure actually measured *from* the simulation using the virial theorem (267). This will reveal programming errors and tell if the equilibration has been sufficient.

8.2 Chemical potential

Let's return to the NVT system and evaluate the chemical potential μ . Roughly speaking, μ is the energy gained/loosed upon insertion of one particle to the system. It is therefore natural to ask what happens if one indeed *does* insert a "test" particle in the simulation.

For simplicity I shall describe only the case when all particles are of the same species. First, the chemical potential is defined via Gibbs or Helmholtz free energies or via entropy,

$$\mu = \left(\frac{\partial G(N, P, T)}{\partial N} \right)_{TP} = \left(\frac{\partial F(N, V, T)}{\partial N} \right)_{TV} = -T \left(\frac{\partial S}{\partial N} \right)_{EV} . \quad (291)$$

The Helmholtz free energy can be calculated from the canonical partition function using (in 3D)

$$\begin{aligned}
 F(N, V, T) &= -k_B T \ln Z_N \\
 &= -k_B T \ln \left(\frac{1}{\lambda^{3N} N!} \int d\mathbf{S} e^{-\beta V(\mathbf{S}, L)} \right) \\
 &= \underbrace{-k_B T \ln \left(\frac{1}{\lambda^{3N} N!} \right)}_{\text{ideal gas}} - k_B T \ln \left(\int d\mathbf{S} e^{-\beta V(\mathbf{S}, L)} \right) .
 \end{aligned} \quad (292)$$

It is worth mentioning, that F is not in a form of a canonical average over phase space, and cannot be directly measured either in a simulation or in real experiments. Same is true also for G and S ; quantities like that are referred to as **thermal quantities**. A reference system must be used; usually ideal gas (used here) or a harmonic crystal.

Sometimes energy differences are all we need. Obviously one has, for sufficiently large N ,

$$\begin{aligned}\mu &= F(N+1, V, T) - F(N, V, T) = F(N, V, T) - F(N-1, V, T) \\ &= -k_B T \ln \frac{Z_{N+1}}{Z_N} = -k_B T \ln \frac{Z_N}{Z_{N-1}} .\end{aligned}\quad (293)$$

If $\mathbf{S}^{(N)}$ refers to a configuration of N particles, the first form gives

$$\mu = \mu_{\text{ideal}} - k_B T \ln \left(\frac{\int d\mathbf{S}^{(N+1)} e^{-\beta V(\mathbf{S}^{(N+1)})}}{\int d\mathbf{S}^{(N)} e^{-\beta V(\mathbf{S}^{(N)})}} \right) . \quad (294)$$

Our strategy is to split the interaction of the $(N+1)$ -particle system in to the interactions of the N -particle system and to the interaction of the $(N+1)$ st particle with the rest of the particles. The non-ideal, “excess” part μ_{ex} can be rewritten as

$$\mu_{\text{ex}} = -k_B T \ln \left(\frac{\int d\mathbf{S}^{(N+1)} e^{-\beta V(\mathbf{S}^{(N)})} e^{-\beta [V(\mathbf{S}^{(N+1)}) - V(\mathbf{S}^{(N)})]} }{\int d\mathbf{S}^{(N)} e^{-\beta V(\mathbf{S}^{(N)})}} \right) \quad (295)$$

$$= -k_B T \ln \left(\frac{\int d\mathbf{r}_{N+1} \int d\mathbf{S}^{(N)} e^{-\beta [V(\mathbf{S}^{(N+1)}) - V(\mathbf{S}^{(N)})]} }{\int d\mathbf{S}^{(N)} e^{-\beta V(\mathbf{S}^{(N)})}} \right) . \quad (296)$$

Again, this contains a canonical average over the N -particle configurations and marking with $\Delta V(\mathbf{r}_{N+1})$ the potential energy of the $(N+1)$ th particle with the rest of the particles, we have simply

$$\mu = \mu_{\text{ideal}} - k_B T \ln \int d\mathbf{r}_{N+1} \langle e^{-\beta \Delta V(\mathbf{r}_{N+1})} \rangle_N , \quad (297)$$

or equivalently

$$e^{-\beta \mu} = Z_N^{\text{ideal}} \left\langle e^{-\beta e^{-\beta \Delta V(\mathbf{r}_{N+1})}} \right\rangle_N . \quad (298)$$

These formulae describe the **Widom test particle method**, known also as the **particle insertion method**:

1. Perform NVT Monte Carlo simulation for a system of N particles
2. Measure the average of $e^{-\beta\Delta V(\mathbf{r}_{N+1})}$ by inserting a fictitious test particle at a random position, and calculating its interaction energy with the rest of the system as if it were real. Repeat the insertion many times to get good statistics.

This technique works fine for fluids at moderate densities around the critical point, but runs into trouble at high densities around liquid–solid coexistence, because the probability to insert the test particle at a place where $e^{-\beta\Delta V(\mathbf{r}_{N+1})} \neq 0$ is very small.

As a final remark, few modifications are needed in a NpT system, in 3D

$$\mu(p) = \underbrace{-k_B T \ln \left(\frac{k_B T}{\beta p \lambda^3} \right)}_{\text{ideal gas}} \quad (299)$$

$$- k_B T \ln \left\langle \frac{\beta p V}{(N+1)} \int d\mathbf{r}_{N+1} e^{-\beta\Delta V(\mathbf{r}_{N+1})} \right\rangle . \quad (300)$$

8.3 Thermostats

We’ve talked about NVT and NpT systems and not really paid any attention to how exactly are classical fixed- T simulations done. What I describe here is mostly applicable to molecular dynamics (MD), but there is a certain random element involved and parts of the code are more Monte Carlo than MD.

Classical particles obey Boltzmann statistics, and inserting the kinetic energy we see, that the velocities have familiar distribution (in 3D)

$$P(v) = \left(\frac{1}{2\pi m k_B T} \right)^{3/2} e^{-mv^2/(2k_B T)} , \quad (301)$$

and the (translational) kinetic energy per particle is (equipartition of energy)

$$m \langle v_i^2 \rangle = \frac{3}{2} k_B T . \quad (302)$$

In part I of these computational physics lectures we applied this relation to do simulations at a given T by using a very simple –and crude– method of *velocity scaling*. Regardless of the initial velocities the MD or MC, after equilibrium has been reached velocities have the Boltzmann distribution,

but T is fluctuating around wrong value. Then we just scaled the velocities to spread the velocity distribution to the T we want. After a while this procedure has to be repeated, since T is drifting away from the given value.

Relation (302) measures the temperature in a micro-canonical ensemble and one can show that the condition “constant T ” is *not* equivalent to “kinetic energy per particle is constant”. This might sound surprising, but once you realise that the latter condition implies that there are no kinetic energy fluctuations you start to wonder how can this be true in a canonical fixed- T system where energies *are supposed to fluctuate*. Specific heat can be evaluated using total energy fluctuations, so why shouldn’t the kinetic energy fluctuate? It should, and we can actually compute the fluctuation in a NVT system. Let T_{kin} be the **kinetic temperature** defined by (302). Then⁵⁶

$$\begin{aligned}\sigma_{T_{\text{kin}}}^2 &\equiv \langle T_{\text{kin}}^2 \rangle - \langle T_{\text{kin}} \rangle^2 \\ &= \frac{1}{N} \frac{\langle p^4 \rangle - \langle p^2 \rangle^2}{\langle p^2 \rangle^2} \langle T_{\text{kin}} \rangle^2 \\ &= \frac{2}{3N},\end{aligned}\tag{303}$$

which is definitely non-zero. The conclusion is, that in a canonical ensemble the kinetic temperature fluctuates and in fact the crude method of velocity scaling does not produce *any* known ensemble.

Better methods to reach fixed T are based on the idea of a heat bath. The question is then, how does one bring the heat bath into contact with the simulated system; the mechanism is called a **thermostat**. If properly executed, MD and MC simulations within the canonical ensemble will give same results.

Andersen Thermostat In MD, one can pick a random particle and give it a new velocity sampled from the Gaussian velocity distribution (301). This coupling mimics the collision between the chosen particle and the heat bath and is characterised by the frequency ν one makes these abrupt velocity changes. The probability that a particle undergoes a collision in a time step Δt is $\nu \Delta t$. The final results should not depend on ν . Few remarks about this method:

- Produces canonical ensemble (indirect evidence, produces the known properties of a canonical ensemble)
- Static properties (potential energy, equation of state etc.) are well produced

⁵⁶Frenkel & Smit, p. 126-127.

- Dynamic properties are *not* well produced (can you think of why?)

Nosé-Hoover Thermostat Nosé introduced a more gentle way to establish the contact with the heat bath; the present implementation is by Hoover. The theory is rather lengthy, but at the same time most interesting, because it falls to a broader category of **how to replace a computationally expensive iterative optimisation with a cheap extended dynamical scheme**. The latter involves an extended Lagrangian, where one has additional coordinates and their canonical momenta. This results algorithms, that do the optimisation “on the fly”.

Nosé wrote the Lagrangian in the form

$$\mathcal{L}_{\text{Nose}} = \sum_{i=1}^N \frac{1}{2} m_i s^2 \dot{\mathbf{r}}_i^2 - V(\mathbf{R}) + \frac{1}{2} Q \dot{s}^2 - g k_B T \ln s , \quad (304)$$

where s is an auxiliary coordinate and Q is the effective mass associated with s , and g is a parameter. The conjugate momenta are

$$\mathbf{p}_i \equiv \frac{\partial \mathcal{L}}{\partial \dot{\mathbf{r}}_i} = m_i s^2 \dot{\mathbf{r}}_i \quad (305)$$

$$p_s = \frac{\partial \mathcal{L}}{\partial \dot{s}} = Q \dot{s} , \quad (306)$$

so the Hamiltonian reads

$$H_{\text{Nose}} = \sum_{i=1}^N \frac{\mathbf{p}_i^2}{2m_i s^2} + V(\mathbf{R}) + \frac{p_s^2}{2Q} + g k_B T \ln s . \quad (307)$$

From this we can derive the equations of motion,

$$\dot{\mathbf{r}}_i = \frac{\partial H_{\text{Nose}}}{\partial \mathbf{p}_i} = \frac{\mathbf{p}_i}{m_i s^2} \quad (308)$$

$$\dot{\mathbf{p}}_i = -\frac{\partial H_{\text{Nose}}}{\partial \mathbf{r}_i} = -\frac{\partial V(\mathbf{R})}{\partial \mathbf{r}_i} \quad (309)$$

$$\dot{s} = \frac{\partial H_{\text{Nose}}}{\partial p_s} = \frac{p_s}{Q} \quad (310)$$

$$\dot{p}_s = -\frac{\partial H_{\text{Nose}}}{\partial s} = \frac{1}{s} \left(\sum_i \frac{p_i^2}{m_i s^2} - g k_B T \right) . \quad (311)$$

With N atom, this describes a micro-canonical system with $6N + 2$ degrees of freedom. In this ensemble one can show that the time average of A is given by

$$\overline{A} = \langle A(\mathbf{r}, \mathbf{p}/s) \rangle_{\text{Nose}} = \langle A(\mathbf{r}, \mathbf{p}/s) \rangle_{NVT} , \quad (312)$$

where the last equality follows if one chooses $g = 3N + 1$. Hence the momentum $\mathbf{p}$ is not the momentum that shows physically, and therefore we call it **virtual momentum**. Few other variables turn out to be virtual as well, and they are all related to the real, physical variables by scaling relations. If **primed variables are real and unprimed virtual**, then

$$r' = r \quad (313)$$

$$p' = p/s \quad (314)$$

$$s' = s \quad (315)$$

$$\Delta t' = \Delta t/s. \quad (316)$$

This implies that s appears as a scaling factor to time. It is very annoying to have a variable time step in a simulation, therefore one usually prefers real variables, which obey the following equations of motion:

$$\dot{\mathbf{r}}'_i = \frac{\mathbf{p}'_i}{m_i} \quad (317)$$

$$\dot{\mathbf{p}}'_i = -\frac{\partial V(\mathbf{R}')}{\partial \mathbf{r}'_i} - \left(\frac{s'p'_s}{Q}\right) \mathbf{p}'_i \equiv -\frac{\partial V(\mathbf{R}')}{\partial \mathbf{r}'_i} - \xi \mathbf{p}'_i \quad (318)$$

$$\frac{1}{s} \dot{s}' = \frac{s'p'_s}{Q} \equiv \xi \quad (319)$$

$$\frac{d\left(\frac{s'p'_s}{Q}\right)}{dt'} = \frac{1}{Q} \left(\sum_i \frac{p'_i}{m_i} - g k_B T \right) \equiv \dot{\xi}, \quad (320)$$

where we have followed Hoover and defined the **thermodynamical friction coefficient** $\xi \equiv s'p'_s/Q$. Furthermore, Hoover was able to prove, that no other set of equations of the same form can lead to a canonical distribution.

8.4 *Ab initio* molecular dynamics: Car–Parrinello method

This subsection is here mostly because the method is based on an extended Lagrangian just like the Nosé–Hoover thermostat, and also because everybody should have at least heard of it. One indication of the popularity of *ab initio* molecular dynamics and the Car–Parrinello method is that they have their own PACS number (71.15Pd), and about 2000 citations on year 2000. As mentioned earlier, the Car–Parrinello method can be viewed as “how to replace a computationally expensive iterative optimisation with a cheap extended dynamical scheme”. In a more narrow perspective, one can

also describe it as a “way to exploit the quantum mechanical adiabatic time–scale separation of fast electronic and slow nuclear motion by transforming that into an adiabatic energy–scale separation in the framework of dynamical systems theory of classical mechanics”.

In some cases one can assume that the atoms as a whole form a meaningful unit and that one can approximate the atom–atom interaction with a fixed pair potential, such as the Lennard–Jones 6-12 potential or the Aziz potential in He–He interaction. How interesting it may be, this approach is teaching us more about statistical mechanics and thermodynamics than it does about real materials, and chemical reactions are out of question. The pair potential view runs into trouble if the atoms can spare electrons to form chemical bonds. Table 1 shows how poorly real metals obey Lennard–Jones predictions⁵⁷. Also real semiconductors deviate clearly from pair–potential results.

- $E_c/(k_B T_m)$: Cohesive energy per melting temperature. Clearly LJ crystals are more strongly bound than real metallic crystals.
- E_v/E_c : Vacancy formation energy per cohesion energy. If a vacancy is created in a solid by removing an atom with coordination number Z , the change of coordination number from Z to $Z - 1$ for Z atoms costs energy E_v . The cohesive energy E_c is the energy needed to change the coordination number of a single atom from Z to zero. Apart from corrections due to lattice relaxation, the two energies would be the same for a pair potential.
- C_{12}/C_{44} : Ratio of two elastic constants of a cubic crystal. According to the so–called Cauchy relation this ratio is exactly 1 for pair potentials in cubic crystals.

Property	Cu	Ag	Pt	Au	LJ
$E_c/(k_B T_m)$	30	28	33	33	13
E_v/E_c	0.33	0.36	0.26	0.25	~ 1
C_{12}/C_{44}	1.5	1.9	3.3	3.7	1

Table 1: Some properties of real (FCC) metals and the Lennard–Jones prediction. The comparison shows the inadequacy of a pair potential.

Electrons have a life of their own, and ideally the calculation should take as an input only the number and types of atoms in the system. Supplemented with a possible external potential, the Coulomb interaction would be

⁵⁷After F. Ercolessi, *A Molecular Dynamics Primer*, Trieste (1997).

sufficient and chemical bonds or effective atom-atom potentials should come out naturally from quantum mechanics; methods capable for this are called first-principles or *ab initio* methods. Solving the full N -body Schrödinger equation is not attempted, instead one usually assumes that the interaction of an electron with all other electrons can be described by a mean-field effective potential. Apart from strongly-correlated systems, this is a very good approximation.

Think about a simulation, where the ions are originally not in their equilibrium positions. To optimise the structure we change the ion positions, and electrons will settle down to a new equilibrium. The electron response to the shifted ionic positions is fast and – to a very good approximation – adiabatic. From the practical side this means that one searches for a self-consistent solution for a set of coupled partial differential equations; it involves either matrix diagonalisation or minimisation of an energy functional.

Up to mid 80's it was believed, that one would have to equilibrate the electronic structure very carefully after every ionic move. This is known as **Born–Oppenheimer dynamics** and it is very time consuming unless one is ready to reduce some accuracy in the energy conservation. Car and Parrinello⁵⁸ devised a new sort of dynamics, which doesn't attempt to optimise the electronic structure exactly after every step, but still ensures that there is *no net force exerted on the ions*.

Application to a classical system: polarisable molecules The total energy of a polarisable molecule fluid is given by⁵⁹

$$E = E_0 - \sum_i \mathbf{E}_i \cdot \boldsymbol{\mu}_i + \frac{1}{2\alpha} \sum_i \mu_i^2, \quad (321)$$

where $\mathbf{E}_i$ is the local electric field acting on molecule i , α is the polarisability of the molecules and $\boldsymbol{\mu}_i$ is the dipole moment induced on it. For dipolar molecules

$$\mathbf{E}_i = \sum_j T_{ij} \cdot \boldsymbol{\mu}_j^{\text{tot}}, \quad (322)$$

where T_{ij} is the dipole–dipole tensor and $\boldsymbol{\mu}_j^{\text{tot}}$ is the induced plus the permanent dipole moment of molecule j .

Now assume we do MD on such a system. When ions move, the electron cloud follows approximately adiabatically the movement, and at every

⁵⁸R. Car and M. Parrinello, Phys. Rev. Lett. **55**, 2471 (1985).

⁵⁹This part follows closely D. Frenkel and B. Smit, *Understanding Molecular Simulation: From Algorithms to Applications* (Academic Press 1996).

instance the polarisation energy is at minimum:

$$\boldsymbol{\mu}_i = \alpha \mathbf{E}_i . \quad (323)$$

This means that for N molecules we would have to solve $3N$ equations per each time step. If it's solved iterative, a partially converged solution would easily leave the electron cloud dragging behind the molecular motion and the local electric field would have a systematic force in the induced dipoles and the total energy of the system would drift. This is one example of the so-called *lag effect* of *ab initio* molecular dynamics.

In a Car–Parrinello approach one treats the induced dipole moments as dynamical variables and uses the Lagrangian

$$\mathcal{L}(\mathbf{R}, \boldsymbol{\mu}^{(N)}) = \frac{1}{2} \sum_{i=1}^N (m \dot{\mathbf{r}}_i^2 + \mathcal{M} \dot{\boldsymbol{\mu}}_i^2) - E , \quad (324)$$

where $\mathcal{M}$ is the fictitious mass associated with dipole motion. The equations of motion are

$$\mathcal{M} \ddot{\boldsymbol{\mu}}_i = -\frac{\boldsymbol{\mu}_i}{\alpha} + \mathbf{E}_i . \quad (325)$$

The rhs is a generalised force acting on dipoles. Notice, how the solutions (323) are recovered if one puts this force to zero (the mass $\mathcal{M}$ to zero). This would mean a return to the iterative scheme and the computational overhead would be lost. In the Car–Parrinello approach the electron clouds would fluctuate around the minimum energy configuration symmetrically and no drift is experienced in the total energy. In order to maintain the dipoles close to their ground–state configuration and still have them adiabatically respond to molecular motion one has to have $\mathcal{M} \ll m$ and the kinetic temperature associated with the dipole motion low (see Eq. (302)). This is a non–trivial problem, it turns out that there is heat exchange between the translational motion and the dipole motion and thus the dipole temperature cannot be simply fixed to a constant.

***Ab initio* molecular dynamics in electron structure calculations** The basic idea behind the Car–Parrinello method is to treat the expansion coefficients of the electronic wave function as dynamical variables and to do MD on *them*. With a chosen basis set $\{\phi\}$ one expands⁶⁰

$$\psi = \sum_n c_n \phi_n , \quad (326)$$

⁶⁰For an extensive review see M. C. Payne, M. P. Teter, D. C. Allan, T. A. Arias and J. D. Joannopoulos, Rev. Mod. Phys. 64 (1992) 1045.

so now we think as if $c_n = c_n(t)$, where t is the simulation time. In MD one would write the Lagrangian in terms of the atomic positions $\mathbf{R}$ and the unit cell dimensions $\mathbf{B}$. As a constraint, the electronic wave functions must be kept orthogonal,

$$\langle \psi_i | \psi_j \rangle = \delta_{ij} , \quad (327)$$

which can be incorporated to the Lagrangian using the method of Lagrange multipliers. The full Car–Parrinello Lagrangian reads

$$\mathcal{L} = \sum_i \mu \langle \dot{\psi}_i | \dot{\psi}_i \rangle - E_{KS}[\{\psi_i\}, \mathbf{R}_I, \mathbf{B}] + \sum_{i,j} A_{ij} (\langle \psi_i | \psi_j \rangle - \delta_{ij}) , \quad (328)$$

where I refers to ions and subscript i to i th electron. Here μ is the fictitious effective mass associated with the expansion coefficients of the Kohn–Sham one–electron wave functions and $E_{KS}[\dots]$ is the Kohn–Sham energy functional. Notice, how the fictitious dynamics of the electronic wave functions has taken the place of the kinetic energy, and how the Kohn–Sham energy has replaced the potential energy in a conventional Lagrangian. The Lagrange multipliers A_{ij} act like additional forces that maintain the orthogonality of the electronic states during the MD simulation. The Lagrange equations of motion

$$\frac{d}{dt} \left(\frac{\partial \mathcal{L}}{\partial \dot{\psi}_i^*} \right) - \frac{\partial \mathcal{L}}{\partial \psi_i^*} = 0 \quad (329)$$

give

$$\mu \ddot{\psi}_i = -H_{KS} \psi_i + \sum_{i,j} A_{ij} \psi_j , \quad (330)$$

where H_{KS} is the Kohn–Sham Hamiltonian. Car and Parrinello solved these equations of motions using the Verlet algorithm. An auxiliary damping of the coefficients $c_n(t)$ is often necessary to reduce overshoots and speed up convergence. Finally, it should be emphasised, that only the ground state c_n are physical, because the Kohn–Sham states are have no clear physical interpretations off the ground state.

If the equations of motion are solved in the absence of the orthonormal constraints the time evolution of the one electron wave functions are found to be either oscillatory or convergent to a single degenerate ground state.

If the ions were fixed, the electronic total energy would be conserved by these equations of motion. Also the one–electron wave functions would remain orthonormal at all times. However, this is true only if one does zero–timestep MD, and in practise one has to do a separate orthonormalisation

step after each MD step, e.g., by the Gram–Schmidt orthogonalisation or – like Car and Parrinello – an iterative method:

$$\psi'_i = \psi_i - \frac{1}{2} \sum_j \langle \psi_j | \psi_i \rangle \psi_j . \quad (331)$$

This algorithm does not conserve the norms, and an extra normalisation step is needed after each iteration.

Once the electronic wave functions are self-consistent, the forces exerted on ions are computed from the Hellman-Feynman theorem. Finite temperature can be incorporated by simply giving the ions random initial velocities and by controlling them afterwards by a Nosé–Hoover thermostat. A further problem is that in metals the one-electron states are partially filled. In that case one introduces a Fermi occupation function, which for numerical reason is smeared by introduction of “electronic temperature”. Each state near the estimated Fermi surface is weighted by the occupied portion of the reciprocal space volume it represents. The Fermi surface itself is found by sampling the k -space within the first Brillouin zone; However, for a large system the number of k -points needed is usually small (in many cases one can restrict to the Γ -point only).

In their original article, Car and Parrinello minimised the energy using simulated annealing. Nowadays simulated annealing has been replaced by direct second order search of the energy minimum, known as the conjugate-gradient method⁶¹. For extended systems the natural basis is plane waves, but a for long time the major hindrance was, that the rapidly oscillating charge density in the core region of an atom requires very many plane waves, which slows down computations. And yet, for most practical purposes the core electron are inert. This fact led to the idea of so-called pseudopotentials, which are artificial potentials that contain all the salient features of a valence electron moving through the solid. Pseudopotentials are smoothly varying, thus the number of basis states can be reduced dramatically and larger simulation systems become feasible than in the more fundamental all-electron calculations. Also relativistic effects can be incorporated to pseudopotentials.

Just for the record, the Kohn–Sham energy functional for doubly occupied electron states is

$$E[\{\psi_i\}, \mathbf{R}_I, \mathbf{B}] = 2 \sum_i \int d\mathbf{r} \psi_i \left(-\frac{\hbar^2}{2m} \nabla^2 \right) \psi_i + \underbrace{\int d\mathbf{r} V_{\text{ion}}(\mathbf{r}) n(\mathbf{r})}_{\text{ion-electron Coulomb}} \quad (332)$$

⁶¹Also the numerical integration of the equations of motion such as the Verlet algorithm is several times slower than evaluation after analytic integration; Look for names Williams and Soler if you are interested in improvements.

$$+\frac{e^2}{2}\int d\mathbf{r}d\mathbf{r}'\frac{n(\mathbf{r})n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}+\underbrace{E_{xc}[n(\mathbf{r})]}_{\text{exchange--correlation}}+\underbrace{E_{\text{ion}}(\{\mathbf{R}_I\})}_{\text{ion Coulomb}}$$

and the Kohn–Sham equations read

$$\left[-\frac{\hbar^2}{2m}\nabla^2+V_{\text{ion}}(\mathbf{r})+V_H(\mathbf{r})+V_{xc}(\mathbf{r})\right]\psi_i(\mathbf{r})=\varepsilon_i\psi_i(\mathbf{r})\;, \quad (333)$$

where

$$V_H(\mathbf{r})=e^2\int d\mathbf{r}'\frac{n(\mathbf{r}')}{|\mathbf{r}-\mathbf{r}'|}\;, \quad (334)$$

and

$$V_{xc}(\mathbf{r})\;=\;\frac{\delta E_{xc}[n(\mathbf{r})]}{\delta n(\mathbf{r})}\;. \quad (335)$$

9 Appendix

9.1 McMillan trial wave function for liquid ^4He

In his PhD thesis William L. McMillan (1936-1984) made an early application of Monte Carlo to study the ground state of liquid and solid helium⁶². Back then the He-He potential was taken to be the 6-12 Lennard-Jones potential,

$$V(r) = 4\varepsilon \left[\left(\frac{\sigma}{r} \right)^{12} - \left(\frac{\sigma}{r} \right)^6 \right] , \quad (336)$$

where $\varepsilon = 10.22$ K and the hard core radius of He atoms is $\sigma = 2.556$ Å, and it's customary to express distances in units σ . The trial wave function should keep the atoms apart, and McMillan chose the form

$$\varphi_T(\mathbf{R}) = \prod_{i<j=1}^N f(r_{ij}) , \quad r_{ij} = |\mathbf{r}_i - \mathbf{r}_j| , \quad (337)$$

with two-body correlations

$$f(r) = e^{-\frac{1}{2}(\beta/r)^\mu} . \quad (338)$$

McMillan found that the best values are $\beta \sim 2.6 - 2.7$ and $\mu = 5$. This wave function is easy and fast to evaluate numerically, and the exponential form is quite convenient, since one needs mostly the expression

$$\frac{|\varphi_T(\mathbf{R}')|^2}{|\varphi_T(\mathbf{R})|^2} . \quad (339)$$

Local energy The Hamiltonian is

$$H = -\frac{\hbar^2}{2m} \sum_{i=1}^N \nabla_i^2 + \sum_{i<j} V(r_{ij}), \quad (340)$$

so the local energy on reads

$$E_L(\mathbf{R}) = -\frac{\hbar^2}{2m} \sum_{i=1}^N \frac{\nabla_i^2 \varphi_T(\mathbf{R})}{\varphi_T(\mathbf{R})} + \sum_{i<j} V(r_{ij}) . \quad (341)$$

This form gives some hint of what is *not* meant by “local energy”: The potential part of it is just the “classical” potential energy of particles located

⁶²W. L. McMillan, Phys. Rev. 138 (1965) A442-A451.

at $\mathbf{R}$, hence the first term should describe the kinetic energy of that system. Right? No, because thinking that way the “kinetic energy of i :th walker” can be negative! So can the whole sum be for some configuration $\mathbf{R}$! This is just a warning that one shouldn’t consider the terms too seriously as the kinetic and the potential energy.

Let’s do the job for a slightly more general case. Use a **one-body correlation function** $u(\mathbf{r}_i)$ for inhomogeneous cases and a **two-body correlation function** to keep particles apart:

$$\varphi_T(\mathbf{R}) = e^{\frac{1}{2}[\sum_i u(\mathbf{r}_i) + \sum_{i < j} u(\mathbf{r}_i, \mathbf{r}_j)]} . \quad (342)$$

where McMillan had $u(\mathbf{r}_i, \mathbf{r}_j) = -(\beta/r_{ij})^\mu$. If only the two-body correlation part is present, then this is called a Jastrow trial wave function. It is much used in any strongly interacting many-body problems, not just liquid helium. To simplify calculations we assume that $u(\mathbf{r}_i, \mathbf{r}_j)$ depends only on the length r_{ij} and not on the spatial direction, hence $u = u(r_{ij})$.

The following formulas come handy; I’ll write intermediate steps in 2D, final forms are valid in any dimension. First of all, we introduce the shorthand notation $u_{ij} = u(r_{ij})$. Let ∇_k denote the gradient with respect to the position vector of particle k . We have

$$\begin{aligned} \nabla_k r_{ij} &\equiv \left(\frac{\partial}{\partial x_k}, \frac{\partial}{\partial y_k} \right) \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2} \\ &= r_{ij}^{-1} (x_i - x_j, y_i - y_j) (\delta_{ik} - \delta_{jk}) \\ &= \hat{\mathbf{r}}_{ij} (\delta_{ik} - \delta_{jk}) , \end{aligned} \quad (343)$$

where the hat signifies a unit vector. Similarly we have

$$\begin{aligned} \nabla_k \cdot \mathbf{r}_{ij} &= \left(\frac{\partial}{\partial x_k}, \frac{\partial}{\partial y_k} \right) \cdot (x_i - x_j, y_i - y_j) \\ &= D(\delta_{ik} - \delta_{jk}) . \end{aligned} \quad (344)$$

From these we get

$$\begin{aligned} \nabla_k^2 r_{ij} &= \nabla_k \cdot \hat{\mathbf{r}}_{ij} (\delta_{ik} - \delta_{jk}) \\ &= \nabla_k \cdot [r_{ij}^{-1} \mathbf{r}_{ij}] (\delta_{ik} - \delta_{jk}) \\ &= -r_{ij}^{-2} \nabla_k r_{ij} \cdot \mathbf{r}_{ij} (\delta_{ik} - \delta_{jk}) + D r_{ij}^{-1} (\delta_{ik} - \delta_{jk})^2 \\ &= (D - 1) r_{ij}^{-1} (\delta_{ik} - \delta_{jk})^2 \\ &= (D - 1) r_{ij}^{-1} (\delta_{ik} + \delta_{jk}) . \end{aligned} \quad (345)$$

Then

$$\nabla_k u_i = u'_i \nabla_k r_i = u'_k \hat{\mathbf{r}}_k \quad (346)$$

$$\nabla_k u_{ij} = u'_{ij} \nabla_k r_{ij} = u'_{kj} \hat{\mathbf{r}}_{kj} - u'_{ik} \hat{\mathbf{r}}_{ik} \quad (347)$$

$$(348)$$

and

$$\nabla_k^2 u_i = u''_k \quad (349)$$

$$\begin{aligned} \nabla_k^2 u_{ij} &= u''_{ij} (\nabla_k r_{ij})^2 + u'_{ij} \nabla_k^2 r_{ij} \\ &= [u''_{ij} + u'_{ij} (D-1) r_{ij}^{-1}] (\delta_{ik} + \delta_{jk}) \\ &= u''_{ik} + u'_{ik} (D-1) r_{ik}^{-1} + (i \leftrightarrow j) . \end{aligned} \quad (350)$$

Here $u'(r_{ij})$ stands for the derivative du_{ij}/dr_{ij} and D is the dimension.

Back to the main task. Defining

$$\varphi_T(\mathbf{R}) \equiv e^{U(\mathbf{R})} , \quad U(\mathbf{R}) \equiv \frac{1}{2} \left[\sum_i u_i + \sum_{i < j} u_{ij} \right] \quad (351)$$

we get

$$\nabla_k \varphi_T = \varphi_T \nabla_k U(\mathbf{R}) , \quad (352)$$

and so

$$\begin{aligned} \varphi_T^{-1} \nabla_k^2 \varphi_T &= \varphi_T^{-1} \nabla_k \cdot [\varphi_T \nabla_k U(\mathbf{R})] \\ &= [\nabla_k U(\mathbf{R})]^2 + \nabla_k^2 U(\mathbf{R}) . \end{aligned} \quad (353)$$

The gradient is now

$$\nabla_k U(\mathbf{R}) = \frac{1}{2} \left[u'_k \hat{\mathbf{r}}_k + \sum_{j(>k)} u'_{kj} \hat{\mathbf{r}}_{kj} - \sum_{i(<k)} u'_{ik} \hat{\mathbf{r}}_{ik} \right] \quad (354)$$

$$= \frac{1}{2} \left[u'_k \hat{\mathbf{r}}_k + \sum_{i(\neq k)} u'_{ki} \hat{\mathbf{r}}_{ki} \right] , \quad (355)$$

and the other one is

$$\nabla_k^2 U(\mathbf{R}) = \frac{1}{2} \left[u''_k + \sum_{i(\neq k)} u''_{ik} + (D-1) u'_{ik} r_{ik}^{-1} \right] . \quad (356)$$

From these we get the local energy

$$E_L(\mathbf{R}) = -\frac{\hbar^2}{2m} \sum_k [[\nabla_k U(\mathbf{R})]^2 + \nabla_k^2 U(\mathbf{R})] + \sum_{i<j} V(r_{ij}) , \quad (357)$$

and the drift (“force”) defined in Eq. (146),

$$\mathbf{F}_k(\mathbf{R}) = 2\nabla_k U(\mathbf{R}) . \quad (358)$$

The local energy can be now computed from the drift,

$$E_L(\mathbf{R}) = -\frac{\hbar^2}{2m} \sum_k \left[\frac{1}{4} \mathbf{F}_k^2 + \frac{1}{2} \nabla_k \cdot \mathbf{F}_k \right] + \sum_{i<j} V(r_{ij}) , \quad (359)$$

If the trial wave function is the exact eigenstate, then this simplifies to

$$E_L(\mathbf{R})_{\text{exact wf}} = \frac{\hbar^2}{8m} \sum_k \mathbf{F}_k^2 + \sum_{i<j} V(r_{ij}) \quad (360)$$

For a good trial wave function the local energy should be as close to a constant (ground state energy) as possible. Inserting a general two-body correlation to the local energy expression one gets

$$\begin{aligned} & -\frac{\hbar^2}{2m} \sum_k \left[\frac{1}{4} \left(\sum_{i(\neq k)} u'_{ki} \hat{\mathbf{r}}_{ki} \right)^2 + \frac{1}{2} \sum_{i(\neq k)} u''_{ik} + (D-1) u'_{ik} r_{ik}^{-1} \right] + \sum_{i \neq j} \frac{1}{2} V(r_{ij}) \quad (361) \\ & = -\frac{\hbar^2}{2m} \frac{1}{4} \sum_j \left(\sum_{i(\neq j)} u'_{ji} \hat{\mathbf{r}}_{ji} \right)^2 + \sum_{i \neq j} \left[-\frac{\hbar^2}{2m} \left(\frac{1}{2} u''_{ij} + \frac{1}{2} (D-1) u'_{ij} r_{ij}^{-1} \right) + \frac{1}{2} V(r_{ij}) \right] . \end{aligned}$$

The first term is a square of a sum, and cannot be written like the rest of the terms. It reads

$$\begin{aligned} & \sum_k \left(\sum_{i(\neq k)} u'_{ki} \hat{\mathbf{r}}_{ki} \right)^2 = \sum_k \left(\sum_{i(\neq k)} u'_{ki} \hat{\mathbf{r}}_{ki} \right) \left(\sum_{j(\neq k)} u'_{kj} \hat{\mathbf{r}}_{kj} \right) \\ & = \sum_{k \neq i} (u'_{ki})^2 + \sum_{i \neq j \neq k} u'_{ki} u'_{kj} \hat{\mathbf{r}}_{ki} \cdot \hat{\mathbf{r}}_{kj} \\ & = \sum_{i \neq j} \left[(u'_{ij})^2 + \sum_{k(\neq i,j)} u'_{ki} u'_{kj} \hat{\mathbf{r}}_{ki} \cdot \hat{\mathbf{r}}_{kj} \right] \quad (362) \end{aligned}$$

$$= \sum_{i \neq j} \left[(u'_{ij})^2 + \sum_{k(\neq i,j)} \nabla_k u_{ki} \cdot \nabla_k u_{kj} \right] . \quad (363)$$

This term is a source of anisotropy in the liquid, especially showing at high density. The density dependence is present already in the 2-body correlation part, hidden in the coordinate i and j occurring in the real system.




Figure 16: The three-body term generated by two-body correlations in the $[\nabla_k U(\mathbf{R})]^2$ local energy term.

For a repulsive potential $u'(r)$ is large at small distances, so the three-body term gets largest contributions from nearest neighbours, which are arranged to a suitable stacking structures at elevated densities.

Forgetting this term one would get

$$\approx \sum_{i \neq j} \left[-\frac{\hbar^2}{2m} \left(\frac{1}{4}(u'_{ij})^2 + \frac{1}{2}u''_{ij} + \frac{1}{2}(D-1)u'_{ij}r_{ij}^{-1} \right) + \frac{1}{2}V(r_{ij}) \right], \quad (364)$$

so obviously choosing $u(r)$ that solves the equation

$$-\frac{\hbar^2}{2m} \left(\frac{1}{4}(u'(r))^2 + \frac{1}{2}u''(r) + \frac{1}{2}(D-1)u'(r)r^{-1} \right) + \frac{1}{2}V(r) = C \quad (365)$$

one gets *nearly* a constant local energy for every $\mathbf{R}$,

$$E_L(\mathbf{R}) \approx \sum_{i \neq j} C = N(N-1)C = N(E_0/N). \quad (366)$$

where the last equality target the assumed ground state energy per particle E_0/N . Practise shows that this two-body correlation keeps atoms too far apart, and the three-body term that gains importance at increasing density. This determines an approximate $u(r)$ up to a constant, which doesn't matter since it enters the normalization of the wave function.

Specifically, the local energy of the McMillan trial wave function is

$$\begin{aligned} u'_{ij} &= \frac{\partial}{\partial r_{ij}} \left[-\left(\frac{\beta}{r_{ij}} \right)^\mu \right] = \mu \beta^\mu r_{ij}^{-\mu-1} \\ u''_{ij} &= -\mu(\mu+1)\beta^\mu r_{ij}^{-\mu-2}, \end{aligned} \quad (367)$$

so

$$\begin{aligned}
E_L^{\text{McMillan}}(\mathbf{R}) &= 2 \underbrace{\frac{\hbar^2}{8m} \left[\mu(\mu+2-d) \beta^\mu \sum_{i \neq j} r_{ij}^{-\mu-2} \right]}_{T_1} \\
&\quad - \underbrace{\frac{\hbar^2}{8m} (\mu\beta^\mu)^2 \sum_i \left[\sum_{j(\neq i)} r_{ij}^{-\mu-2} \mathbf{r}_{ij} \right]^2}_{T_2} + \sum_{i < j} V(r_{ij}) \\
&= 2T_1(\mathbf{R}) - T_2(\mathbf{R}) + \sum_{i < j} V(r_{ij}) .
\end{aligned} \tag{368}$$

The kinetic part of the local energy has two terms, which correspond to evaluating the expectation values $\langle \varphi_T | \nabla_{\mathbf{R}}^2 | \varphi_T \rangle$ and $\langle \nabla_{\mathbf{R}} \varphi_T | \nabla_{\mathbf{R}} \varphi_T \rangle$. Although mathematically equivalent, these two forms give different variance in Monte Carlo. Sometimes mathematical manipulations *may* lead to increase of variance. A rule of thumb says, that one should evaluate expressions in their most basic form. The mathematical equivalence of the two forms of the kinetic energy ensures that their *mean* values should be the same. This is a very useful test: Evaluate expectation values of both forms and see if they approach the same value. If not, then either expression has programming error or the derivatives are not correct.

For completeness we spell out also the drift force corresponding to the McMillan trial wave function :

$$[\mathbf{F}(\mathbf{R})]_i = \mu\beta^\mu \sum_{j(\neq i)} r_{ij}^{-\mu-2} \mathbf{r}_{ij} . \tag{369}$$

For a repulsive potential the quantity in the exponent has to be negative to keep walkers apart. The choice $\mu = 5$ pulls them strongly apart: the He-He potential is very repulsive in short range.

Figs. 17 and 18 show some selected results for the McMillan correlation factor. For comparison we show the fully optimised u_{ij} , as given by the HyperNetted Chain (HNC) theory. Clearly the McMillan two-body correlation function is able to keep atoms apart, but it doesn't describe the structure near 5 Å, and hence the radial distribution oscillates with smaller amplitude than the experimental one. The DMC radial distribution by Boronat and Casulleras⁶³ is practically identical to the experimental one.

The local energy of the Slater-Jastrow trial wave function (174) is also

⁶³J. Boronat and J. Casulleras, Phys. Rev. **B49** (1994) 8920.




Figure 17: The McMillan correlation factor for few parameters, compared with the two-body correlations optimised fully using the HNC theory and the Aziz (1979) He-He potential at saturation density $\rho = 0.02185\text{\AA}^{-3}$.

fairly simple to evaluate, if we define

$$L(\mathbf{R}) = \ln [\det(A_{\uparrow}) \det(A_{\downarrow})] U(\mathbf{R}) , \quad (370)$$

so that again

$$\nabla_{\mathbf{R}}^2 \varphi_T(\mathbf{R}) = [\nabla_{\mathbf{R}}^2 L(\mathbf{R}) + (\nabla_{\mathbf{R}} L(\mathbf{R}))^2] \varphi_T(\mathbf{R}) . \quad (371)$$

The gradient of the Slater determinant is obtained using the identity

$$\frac{\partial}{\partial x_k} \ln \det(A) = \text{Tr} \left[A^{-1} \frac{\partial}{\partial x_k} A \right] . \quad (372)$$

9.2 Cusp conditions in Coulomb systems

We must make sure that the trial wave function satisfies the so-called **cusp condition**, so that the infinities of the potential energy are cancelled by those of the kinetic energy⁶⁶. Here we deal with the conditions for an N -electron system⁶⁷

⁶⁶T. Kato, Comm. Pure Appl. Math. 10, 151 (1957).

⁶⁷This appendix follows closely the one given by M. L. Stedman in his PhD thesis.




Figure 18: The radial distribution function corresponding to the McMillan wave function using N walkers with periodic boundary conditions of 3D liquid helium at saturation density. For comparison we show the experimental results by Svensson *et al.*⁶⁵

The Hamiltonian of N electrons is

$$H = -\frac{1}{2} \sum_{i=1}^N \nabla_i^2 + V(\mathbf{R}) . \quad (373)$$

Move electrons 1 and 2 so that other electrons and the total center of mass $\mathbf{R}_{CM}$ are fixed. The N -electron wave function is

$$\Psi(\mathbf{R}) \equiv \Psi(\mathbf{r}_1, \dots, \mathbf{r}_N) = \Psi'(\mathbf{r}, \mathbf{R}_{CM}, \mathbf{r}_3, \dots, \mathbf{r}_N) , \quad (374)$$

where $\mathbf{r} \equiv \mathbf{r}_1 - \mathbf{r}_2$. Assume next that the trial wave function is of the Jastrow form,

$$\Psi'(\mathbf{r}, \mathbf{R}_{CM}, \mathbf{r}_3, \dots, \mathbf{r}_N) \propto e^{-u(\mathbf{r})} f(\mathbf{r}) , \quad (375)$$

where f obeys the following rules:

1. $f(0) \neq 0$, if the spins at $\mathbf{r}_1$ and at $\mathbf{r}_2$ are anti-parallel
2. $f(0) = 0$, if the spins at $\mathbf{r}_1$ and at $\mathbf{r}_2$ are parallel
3. $f(-\mathbf{r}) = -f(\mathbf{r})$, if the spins at $\mathbf{r}_1$ and at $\mathbf{r}_2$ are parallel.

First we need to rewrite the Hamiltonian using $\mathbf{r}$ and $\mathbf{R}_{CM}$. One has

$$\nabla_1^2 + \nabla_2^2 = 2\nabla_{\mathbf{r}}^2 + \frac{1}{2}\nabla_{\mathbf{R}_{cm}}^2, \quad (376)$$

where ∇_i on the left hand side operates on i th particle.

Proof:

We have

$$\begin{aligned} \mathbf{X} &= \frac{\mathbf{r}_1 + \mathbf{r}_2}{2} \\ \mathbf{x} &= \mathbf{r}_1 - \mathbf{r}_2, \end{aligned} \quad (377)$$

so

$$\begin{aligned} \frac{\partial}{\partial(\mathbf{r}_1)_k} &= \frac{\partial \mathbf{X}_k}{\partial(\mathbf{r}_1)_k} \frac{\partial}{\partial \mathbf{X}_k} + \frac{\partial \mathbf{x}_k}{\partial(\mathbf{r}_1)_k} \frac{\partial}{\partial \mathbf{x}_k} \\ &= \frac{1}{2} \frac{\partial}{\partial \mathbf{X}_k} + \frac{\partial}{\partial \mathbf{x}_k} \\ \frac{\partial}{\partial(\mathbf{r}_2)_k} &= \frac{1}{2} \frac{\partial}{\partial \mathbf{X}_k} - \frac{\partial}{\partial \mathbf{x}_k}, \end{aligned} \quad (378)$$

and further

$$\nabla_1^2 = \sum_{k=1}^3 \frac{\partial^2}{\partial(x_1)_k^2} \quad (379)$$

$$= \sum_{k=1}^3 \left(\frac{1}{2} \frac{\partial}{\partial \mathbf{X}_k} + \frac{\partial}{\partial \mathbf{x}_k} \right)^2 \quad (380)$$

$$= \sum_{k=1}^3 \frac{1}{4} \frac{\partial^2}{\partial \mathbf{X}_k^2} + \frac{\partial^2}{\partial \mathbf{X}_k \partial \mathbf{x}_k} + \frac{\partial^2}{\partial \mathbf{x}_k^2}.$$

Similarly we get

$$\nabla_2^2 = \sum_{k=1}^3 \frac{1}{4} \frac{\partial^2}{\partial \mathbf{X}_k^2} - \frac{\partial^2}{\partial \mathbf{X}_k \partial \mathbf{x}_k} + \frac{\partial^2}{\partial \mathbf{x}_k^2}, \quad (381)$$

hence

$$\nabla_1^2 + \nabla_2^2 = \sum_k \frac{1}{2} \frac{\partial^2}{\partial \mathbf{X}_k^2} + 2 \frac{\partial^2}{\partial \mathbf{x}_k^2} = \frac{1}{2} \nabla_{\mathbf{R}_{CM}}^2 + 2 \nabla_{\mathbf{r}}^2, \quad (382)$$

what we set out to prove.

The Hamiltonian reads now

$$H = - \left(\frac{1}{4} \nabla_{\mathbf{R}_{CM}}^2 + \nabla_{\mathbf{r}}^2 \right) - \frac{1}{2} \sum_{i=3}^N \nabla_i^2 + V(\mathbf{r}, \mathbf{R}_{CM}, \mathbf{r}_3, \dots, \mathbf{r}_N) . \quad (383)$$

Our goal is to keep the local energy $E_L = \varphi_T^{-1} H \varphi_T$ finite as $\mathbf{r} \rightarrow 0$. The potential part diverges as $1/r$.

$$\lim_{\mathbf{r} \rightarrow 0} \frac{1}{e^{-u(\mathbf{r})} f(\mathbf{r})} \left(-\nabla_{\mathbf{r}}^2 + \frac{1}{r} \right) e^{-u(\mathbf{r})} f(\mathbf{r}) < \infty , \quad (384)$$

and the expanding the operator gives

$$\begin{aligned} \left(-\nabla_{\mathbf{r}}^2 + \frac{1}{r} \right) e^{-u(\mathbf{r})} f(\mathbf{r}) &= \left(\frac{2u'}{r} + u'' - u'^2 \right) e^{-u(\mathbf{r})} f(\mathbf{r}) + \frac{e^{-u(\mathbf{r})}}{r} f(\mathbf{r}) \\ &\quad + 2u' e^{-u(\mathbf{r})} \hat{\mathbf{r}} \cdot \nabla_{\mathbf{r}} f(\mathbf{r}) - e^{-u(\mathbf{r})} \nabla_{\mathbf{r}}^2 f(\mathbf{r}) , \end{aligned} \quad (385)$$

where $u' \equiv du/dr$.

Electrons with anti-parallel spins Since now $f(0) \neq 0$ there is only one singular terms in (386) and the cusp condition reduces to

$$\lim_{r \rightarrow 0} \frac{2u' + 1}{r} < \infty , \quad (386)$$

so we must have

$$\boxed{\left. \frac{du}{dr} \right|_{r=0} = -\frac{1}{2} .} \quad (387)$$

Electrons with parallel spins Taking into account the antisymmetry of $f(\mathbf{r})$, the Taylor expansion of $f(\mathbf{r})$ around zero is

$$f(\mathbf{r}) = \mathbf{a} \cdot \mathbf{r} + \mathcal{O}(r^3) , \quad (388)$$

so the derivatives are

$$\begin{aligned} \nabla_{\mathbf{r}} f(\mathbf{r}) \Big|_{r=0} &= \mathbf{a} \\ \nabla_{\mathbf{r}}^2 f(\mathbf{r}) \Big|_{r=0} &= 0 . \end{aligned} \quad (389)$$

The operator in the local energy (384) is

$$\left(-\nabla_{\mathbf{r}}^2 + \frac{1}{r}\right) e^{-u(\mathbf{r})} f(\mathbf{r}) = \left(\frac{4u'}{r} + u'' - u'^2 + \frac{1}{r}\right) e^{-u(\mathbf{r})} f(\mathbf{r}) , \quad (390)$$

so the only singular terms in the cusp condition give

$$\lim_{r \rightarrow 0} \left(\frac{4u'}{r} + \frac{1}{r}\right) < \infty , \quad (391)$$

which gives the cusp condition for two electrons with parallel spins:

$$\boxed{\left.\frac{du}{dr}\right|_{r=0} = -\frac{1}{4} .} \quad (392)$$